The Actual Cost of Retail Crypto Trading*

Dominik Rösch† Thomas Shohfi‡ Jenny Stanco§ Michael Walz ¶

First Draft: March 2026
This Draft: September 2026

Abstract

We execute several hundred crypto trades to compare costs at major crypto exchanges with costs of simultaneous trades across three benchmarks: trades executed at Robinhood, trades reported by Kaiko, and trades on token-tracking ETFs reported by TAQ data. Effective spreads at major crypto exchanges are substantially higher than these benchmarks, ranging from 253 to 835 basis points. Overall, the results show that retail crypto trading costs are economically large, might be underestimated using common benchmarks, and that substantially cheaper trading alternatives appear to exist. Regulation, such as enforcing best execution, establishing a crypto NBBO, and creating a consolidated tape, could help bring these costs in line with those in other markets and make benchmarks more reliable to accurately reflect retail execution costs.

Keywords: benchmarks, crypto, crypto assets, cryptocurrency, digital assets, effective spreads, payment for order flow (PFOF), retail trading, Bitcoin, Ethereum, Dogecoin, Solana, Coinbase, Kraken, Crypto.com, Robinhood
JEL Codes: C93, D43, D53, G12, G18

*We thank James Cook for excellent research assistance. We thank Cuyler Strong, Tom Ruchti, and seminar participants at the University at Buffalo, the U.S. Securities and Exchange Commission's Crypto Working Group, and Wayne State University for helpful comments and suggestions. We thank OneMarket Data for the use of their OneTick software. This work is supported in part by NSF ACI-1541215. Shohfi acknowledges financial support from the Wayne State Undergraduate Research Opportunities Program (UROP).

All errors or omissions are the responsibility of the authors.

†School of Management, University at Buffalo, drosch@buffalo.edu
‡Mike Ilitch School of Business, Wayne State University, tshohfi@wayne.edu
§Department of Finance, Towson University, jstanco@towson.edu
¶Mike Ilitch School of Business, Wayne State University, mwalz@wayne.edu

1 Introduction

Retail participation has and continues to be an important part of cryptocurrency markets over the past decade (Cong et al., 2023; Kogan et al., 2024). Platforms such as Coinbase, Kraken, Crypto.com, and Robinhood have made spot crypto trading accessible to millions of individual investors through simplified interfaces and often commission-free marketing. Whereas early crypto markets were dominated by technologically sophisticated participants, today's environment increasingly resembles mainstream retail brokerage markets in scale and accessibility. Despite this rapid growth, few studies have examined retail trading in crypto asset markets (e.g., Kogan et al., 2024; Pursiainen and Toczynski, 2023; Divakaruni and Zimmerman, 2024) and comparatively little is known about the execution quality retail crypto traders receive.

This paper provides a transaction-level analysis of retail crypto execution costs across major retail platforms. Using 456 trades across four widely traded tokens—Bitcoin (BTC), Ethereum (ETH), Solana (SOL), and Dogecoin (DOGE)—we compute effective spreads, realized spreads, and price impact using both displayed and all-in prices, i.e., displayed prices inclusive of any explicit fees. Effective spreads follow the conventional definition of twice the signed distance between the execution price and the midpoint. They therefore express costs on a round-trip-equivalent basis, rather than the cost of an individual purchase or sale. We benchmark execution quality relative to both (local) exchange-level midpoints and a synthetic three-exchange Best Bid and Offer (3BBO) midpoint constructed from the best quotes across our three venues.

Retail trading costs are large and differ substantially across platforms. Using the 3BBO midpoint, average all-in effective spreads across platform–token combinations range from approximately 253 to 835 basis points on Coinbase, Kraken, and Crypto.com. Accounting for explicit fees changes the ranking of platforms. Coinbase's displayed execution prices are close to the midpoint, but its separately charged fees make it the most expensive platform in our sample. Kraken's displayed prices already imply substantial costs, which increase further after adding its separately charged fees. Crypto.com's charges are embedded in its displayed prices. These differences show why comparisons based only on displayed prices or separately itemized fees provide an incomplete account of retail trading costs.

We compare execution quality of our crypto trades with three benchmarks. First, we compare trades executed on crypto exchanges with simultaneous trades executed on Robinhood, a payment for order flow broker. Second, we compare execution quality of our trades to the execution quality as reported by Kaiko at the same time and on the same exchange. Third, we compare execution quality of our trades to the execution quality of the relevant token-tracking ETFs at the same time.

We find that effective spreads from trades executed on Robinhood and estimated from Kaiko data and from ETFs, are much smaller. All-in effective spreads on Robinhood range

from 133 to 158 basis points.

Average local effective spreads calculated from exchange-reported trades (Kaiko) are below three basis points for every platform–token combination. These estimates exclude explicit fees and are far below the all-in costs we incur. For Kraken and Crypto.com, a substantial difference remains when we use displayed execution prices without adding explicit fees. The gap therefore extends beyond fees added to an otherwise comparable order-book execution price.

ETF trades provide a complementary benchmark for obtaining exposure to cryptocurrency prices. Bitcoin and Ethereum ETF effective spreads are substantially below the all-in spreads on our corresponding crypto trades. The available Solana ETF observations have much wider spreads, underscoring that the comparison varies across tokens. ETFs also differ from direct token ownership in trading hours, transferability, and holding costs. We therefore interpret their secondary-market spreads as a benchmark for alternative price exposure, rather than a complete measure of the cost of owning cryptocurrency.

Overall, these results indicate that common benchmarks (Kaiko data) substantially understate the transaction costs faced by retail customers and that substantially cheaper trading alternatives (Robinhood and ETFs) appear to exist.

We then decompose spreads into realized spreads, a common way to estimate liquidity providers profits. Whether we use effective or realized spreads, results only marginally change, indicating that order processing and platform margins dominate adverse selection as drivers of transaction cost (Huang and Stoll, 1997).

Research on equities and options shows that conventional market data do not always capture the execution costs investors incur. Using simultaneous equity trades, Schwarz et al. (2025) document substantial differences in execution prices across retail brokers, even for otherwise identical orders. In options markets, Muravyev and Pearson (2020) show that traders who strategically time their executions incur lower costs than conventional spread estimates suggest. These findings underscore the importance of matching transaction-cost measures to the investors and trading practices under study.

The institutional setting also matters for interpreting execution costs. U.S. equity and options markets have consolidated quotation systems, and FINRA Rule 5310 requires member firms to exercise reasonable diligence in seeking favorable execution terms for customer orders. Within this setting, O'Hara and Ye (2011) find that fragmentation across trading venues need not impair market quality. Spot crypto trading differs: venues maintain separate order books without a comparable market-wide NBBO, and securities-market execution and disclosure requirements do not apply uniformly across crypto assets and platforms. Retail interfaces may also combine embedded price components with separately charged fees, making their costs difficult to compare using exchange quotes alone. These differences motivate direct measurement of retail execution costs, although our comparisons do not identify the causal effects of regulatory arrangements.

Classical microstructure models explain spreads through adverse selection, inventory risk, and order processing costs (Glosten and Milgrom, 1985; Stoll, 1989, 2000). Applying these explanations to retail crypto transactions requires distinguishing the cost of liquidity in an exchange order book from the additional charges associated with accessing it through a retail interface. Existing crypto research examines related but distinct frictions. Makarov and Schoar (2020) document persistent cross-exchange price differences, particularly across countries, and the role of limits to arbitrage. Easley et al. (2019) study transaction fees on the Bitcoin blockchain. We complement this work by directly measuring execution prices and fees through basic retail interfaces and assessing whether exchange-based benchmarks represent the costs faced by their customers.

This paper makes four primary contributions. First, we extend the "actual retail price" framework developed for equities by Schwarz et al. (2025) to crypto asset markets. By applying consistent effective spread measurement techniques across asset classes, we provide a direct cross-market comparison of retail execution quality. We also provide evidence that payment for order flow-based transactions (i.e. Robinhood), the current standard in equity markets, generally offer better execution quality than fee-based brokers.

Second, we contribute to the broader market microstructure literature by analyzing spread formation in a fragmented and lightly regulated trading environment. Our results suggest that consolidated quoting and best execution obligations play an economically meaningful role in constraining retail spreads. Third, we add to the emerging literature on crypto market design by focusing explicitly on retail transaction prices rather than solely on cross-exchange arbitrage or volatility transmission. In doing so, we provide new evidence on the economic incidence of embedded spread pricing. In particular, execution-cost estimates derived from Kaiko data or token-tracking ETF trades substantially understate the transaction costs we observe through basic retail crypto interfaces.

Finally, our findings contribute to the literature on retail trading more broadly. As retail participation expands across asset classes, understanding the full cost of trading—including embedded spreads—is essential for evaluating welfare, competition, and market design. Crypto markets serve as a natural experiment in how transaction costs evolve when regulatory architecture differs markedly from that of traditional securities markets.

In sum, this paper positions retail crypto execution within the canon of market microstructure research. By measuring effective and realized spreads in an environment without consolidated quoting or formal best execution mandates, we provide new evidence on how institutional design shapes retail trading costs.

2 Data

We use data from three different sources throughout our analysis, "Actual Trade Data", "Kaiko", and "TAQ". We conduct trades ourselves on three crypto exchanges and benchmark

these to trades that we conduct on Robinhood. We call this data the "Actual Trade Data". We further benchmark our trades against trades as reported by "Kaiko" and against trades of token-tracking ETFs as reported by "TAQ".

2.1 Actual Trade Data

We conduct actual trades across four U.S.-based retail brokers: Robinhood, Coinbase, Kraken, and Crypto.com. Details for each broker are available in the Appendix Table A.1. Robinhood and Coinbase are publicly traded companies while Kraken and Crypto.com are privately owned.

We selected these brokers because they allowed account sign-ups in Michigan, from where we accessed the trading platforms and because they are some of the largest brokers in the United States offering retail trading in US dollar terms. According to CoinMarketCap, as of March 2026, Coinbase was the largest exchange or broker offering U.S.-dollar crypto trading services; Crypto.com ranked third and Kraken fifth.[1] Robinhood is not included in the CoinMarketCap ranking. However, public records indicate that Robinhood executes $635.6 million in crypto volume per day, which would make it one of the largest crypto trading platforms in the United States.[2]

Robinhood is the only broker that uses payment for order flow (PFOF), while the other three brokers are fee-based. Robinhood also offers a smart router trading option that does not use PFOF. While our primary empirical results focus on standard Robinhood transactions, we also conduct robustness tests using trades executed through Robinhood's Smart Order Router. Across benchmarks, our conclusions are unchanged: we find no significant differences when comparing our trades with standard Robinhood executions versus Smart Order Router executions.

To emulate the simplest possible retail trading operations, we use the most basic broker functionality offered (e.g., standard Coinbase, Kraken, etc. as opposed to Coinbase Advanced or Kraken Pro, both of which require minimum volumes for reducing transaction fees).

We execute retail trades in synchronized pairs between Robinhood (PFOF) and the other three exchanges. Trades began in April 17, 2025, ending in September 3, 2025, and ranged randomly from $100 to $400 in dollar value. Table 1 provides the distribution (of trades in Panel A and dollar volume in Panel B) of the 456 total trades that we executed.

[Place Table 1 about here]

Our sample is designed to provide broad and balanced coverage across brokers, tokens,

[1] See `https://coinmarketcap.com/rankings/exchanges/`.

[2] See `https://robinhood.com/us/en/newsroom/this-year-in-crypto-2025/` for details on Robinhood's crypto trading volume in 2025.

and trade directions. We execute both buy and sell orders for each token on each broker, yielding 456 trades in total, with similar numbers of buys and sells, 236 and 220, respectively. Our trades are executed contemporaneously and paired directly against Robinhood: every buy or sell order on Coinbase, Kraken, or Crypto.com (e.g., a BTC purchase on Coinbase or an ETH sale on Kraken) is matched with a simultaneous order of the identical dollar amount on Robinhood. Trading activity is also comparable across tokens, ranging from 108 trades in DOGE to 124 trades in BTC (Section 2.2 explains how we chose which tokens we trade). Dollar volume is likewise distributed across assets and brokers, with somewhat larger volume for BTC and Robinhood, but no single token or trade direction dominating the sample. In total, our sample represents more than $73,000 of notional executed trade value.

Data collection for each trade includes the dollar trade size, the number of tokens traded, the displayed trade price, and explicit fees. Note that we do not observe prevailing bid and ask prices. Crypto exchanges often do not provide bid and ask prices, as indicated by their trading interface, see Appendix Figure A.2.

Displayed prices are the prices provided by the broker at the time of execution. For some brokers, such as Coinbase and Kraken, the displayed price excludes the explicit fee, while for Crypto.com, the displayed price already incorporates the explicit fee. To ensure comparability across brokers, we calculate an all-in price as the dollar trade size divided by the number of tokens traded, as shown in Equation 1:

$$\text{All-in Price} = \frac{\text{Dollar Trade Size}}{\text{Number of Tokens Traded}} \tag{1}$$

For example, consider a BTC purchase on Coinbase where the displayed price is $84,737.775, the dollar trade size is $100.17, and the number of BTC purchased is 0.0011266. The all-in price is calculated as $100.17 divided by 0.0011266, which equals $88,913.545. For this transaction, Coinbase reports an explicit fee of $4.704422685. The total dollar amount paid by the investor (Dollar Trade Size) is equal to the value of the Bitcoin purchased at the displayed price plus the explicit fee, calculated as $84,737.775 multiplied by 0.0011266 plus $4.704422685, which totals $100.17. Thus, the all-in per token price captures the total cost paid by the investor, including both the displayed execution price and any per token explicit fees ($\frac{\text{Dollar Explicit Fee}}{\text{Number of Tokens Traded}}$), as defined in Equation 2:

$$\text{All-in Price} = \text{Displayed Price} + \text{Trade Direction} \times \text{Explicit Fee} \tag{2}$$

Note that Equation 2 does not apply to Crypto.com, where the displayed price already incorporates the explicit fee.

2.2 Kaiko Data

We use tick-by-tick data from Kaiko, a private firm that provides trade and top-of-book quotes timestamped to the millisecond and recorded in UTC. This data source is commonly used in other studies (Makarov and Schoar, 2020; Li et al., 2021; Dyhrberg et al., 2023; Foley et al., 2023; Divakaruni and Zimmerman, 2024). We obtain data from the three crypto exchanges we used for our "actual" trades (Coinbase, Kraken, and Crypto.com). The data includes the buy/sell indicator, the transaction price in USD, and the number of tokens traded. The dollar trade size for each trade is calculated as the product of the transaction price and the number of tokens traded. The best bid and ask data per exchange are an event-driven record of top-of-book quotes where a new observation is recorded for every update.[3] From these, we construct a "3BBO" using the best bid and ask across the three exchanges. While this follows the logic of the NBBO in equity markets, we term it the 3BBO because it is constructed from only the three exchanges in the Kaiko data.

We focus on four crypto tokens for our experiment. The first are BTC and ETH, the two largest crypto asset tokens by market capitalization as of March 2026.[4] We also include SOL and DOGE, both of which are in the top ten crypto assets by market capitalization and are relatively popular amongst retail investors due in part to their low single unit prices. We provide the daily total trading volume (in USD) over our sample period for each crypto asset in Figure 1.

[Place Figure 1 about here]

Trading volume is highly skewed across both assets and exchanges. Coinbase and Crypto.com generally dominate trading volume across all tokens, except Dogecoin. BTC and ETH account for the largest dollar volumes, with frequent days above $1 billion (on both Coinbase and Crypto.com), while SOL and DOGE trade at lower levels, at, respectively, around $200 and $50 million per day on Coinbase. While crypto trading volume is difficult to measure because of concerns about wash trading (see Cong et al., 2023), we restrict our sample to regulated exchanges, where such concerns should be less pronounced.

To compare our transactions with Kaiko's trade records, we examine a one-minute window of Kaiko data immediately preceding and following each of our trades.

[3] Kraken does not have best bid and ask for Dogecoin available from Kaiko.

[4] See `https://coinmarketcap.com`.

2.3 TAQ (ETF) Data

We obtain tick-by-tick transaction and quote records from the Trade and Quote (TAQ) database, accessed through Wharton Research Data Services (WRDS). Our sample consists of Exchange-Traded Funds (ETFs) tracking BTC, ETH, and SOL.[5] The underlying crypto ETF tickers in our sample are as follows: (i) ETFs tracking BTC: IBIT, FBTC, GBTC, BITB, HODL, BTC, ARKB, BRRR, BTCO, BTCW, and EZBC; (ii) ETFs tracking ETH: ETHA, ETHE, ETH, FETH, ETHW, ETHV, QETH, and EZET; and (iii) ETFs tracking SOL : SSK.[6] Transaction prices are extracted from the TAQ trades files, while prevailing bid and ask quotes are obtained from the TAQ complete NBBO files.

To compare our transactions with the ETF trade records, we examine a one-minute window of TAQ data immediately preceding and following each of our trades.

2.4 Variable construction and summary statistics

For both our actual trades and the benchmark transactions reported by Kaiko and TAQ, we estimate the quoted spread (QSP), effective spread (ESP), realized spread (RSP), price improvement, and price impact as given below.

We calculate the quoted spread (QSP) for trade t on exchange i as shown in Equation 3:

$$\text{QSP}_{i,t} = \frac{\text{Ask}_{i,t} - \text{Bid}_{i,t}}{\text{Mid}_{i,t}} \tag{3}$$

where $Ask_{i,t}$ and $Bid_{i,t}$ are the best ask and bid prevailing trade t on exchange i, and $Mid_{i,t}$ is the midpoint between the best bid and ask quote at the time of trade t for exchange i. We do not construct quoted spreads directly for our actual trades because we do not observe the prevailing bid and ask prices, as discussed in Section 2.1. However, for each actual trade, we match the trade to the prevailing bid and ask prices reported by Kaiko for the same exchange and use these quotes to calculate $\text{Mid}_{i,t}$.

We measure the effective spread for trade t on exchange i as shown in Equation 4:

$$\text{ESP}_{i,t} = 2 \times \text{BuySell}_{i,t} \times \frac{\text{Price}_{i,t} - \text{Mid}_{i,t}}{\text{Mid}_{i,t}} \tag{4}$$

where $Mid_{i,t}$ is defined as before, $BuySell_{i,t}$ is the buy/sell indicator, taking +1 (-1) for a buyer (seller) initiated order at the time of trade t for exchange i. For our actual trade data and for trades from Kaiko we know the true trade initiator, but for trades from TAQ we follow Lee and Ready (1991) to determine whether a trade is buyer or seller initiated. $Price_{i,t}$ is the trade price.

[5]The first DOGE ETF is introduced after the end of our sample period, so no data are available.

[6]Crypto.com ETH does not have corresponding ETF transactions because these trades occur between 8:00 PM and 4:00 AM the following day, when no TAQ data are available. We obtain only 13 observations for SOL ETFs because the product was only recently introduced and remains relatively illiquid.

For our actual trade data, we compute two effective spread measures: one based on the displayed transaction price and another based on the all-in price. Because the all-in price incorporates both the displayed price and explicit fees, the effective spread calculated using the all-in price is given in Equation 5:

$$\text{ESP}_{i,t}^{\text{All-in}} = \text{ESP}_{i,t}^{\text{Displayed}} + 2 \times \frac{\text{Explicit Fee}_{i,t}}{\text{Mid}_{i,t}} \tag{5}$$

Note that Equation 5 does not apply to Crypto.com, where the displayed price already incorporates the explicit fee. In this case,

$$\text{ESP}_{i,t}^{\text{All-in}} = \text{ESP}_{i,t}^{\text{Displayed}} \tag{6}$$

We measure the realized spread for trade t on exchange i as defined in Equation 7:

$$\text{RSP}_{i,t} = 2 \times \text{BuySell}_{i,t} \times \frac{\text{Price}_{i,t} - \text{Mid}_{i,t+1\,\text{minute}}}{\text{Mid}_{i,t}} \tag{7}$$

We measure the round-trip price improvement for trade t on exchange i as defined in Equation 8:

$$\text{PriceImprovement}_{i,t} = \text{QSP}_{i,t} - \text{ESP}_{i,t} \tag{8}$$

We measure the price impact for trade t on exchange i as defined in Equation 9:

$$\text{PriceImpact}_{i,t} = 2 \times \text{BuySell}_{i,t} \times \frac{\text{Mid}_{i,t+1\,\text{minute}} - \text{Mid}_{i,t}}{\text{Mid}_{i,t}} \tag{9}$$

Panels A and B of Table 2 provides summary statistics for our Kaiko and TAQ data: the number of trades, the average dollar trade size, the total dollar trading volume, and the average quoted and effective spreads over all trades during the one-minute before and after each of our actual trades. Average values are calculated as follows: we first compute the equally weighted average within the one-minute window surrounding each of our actual trades. We then take the equally weighted mean within each subset (e.g., Coinbase BTC). Trade counts and dollar trading volume are reported as totals across the corresponding windows.

[Place Table 2 about here]

Most trading, as reported by Kaiko, occurs on Coinbase and Crypto.com and on BTC and ETH, consistent with Figure 1. For example, during the one minute before and after each of our 18 BTC trades on Coinbase, we find in total 16,913 BTC trades on Coinbase with a total trading volume of almost $30 million.

The average dollar trade size for ETH, SOL, and DOGE on Coinbase ranges from $254 to $311, which is reasonably comparable to the size of our own trades. In contrast to the average trade size on Coinbase, the average dollar trade size is generally larger on Kraken and Crypto.com, reaching up to $2,755 and $3,204 for BTC, respectively.

We calculate quoted and effective spreads using both local and 3BBO midquotes. For example, the local effective spread for Coinbase is calculated as shown in Equation 10:

$$\text{Local ESP}_{\text{Coinbase},t} = 2 \times \text{BuySell}_{\text{Coinbase},t} \times \frac{\text{Price}_{\text{Coinbase},t} - \text{Mid}_{\text{Coinbase},t}}{\text{Mid}_{\text{Coinbase},t}} \tag{10}$$

The 3BBO effective spread for Coinbase is calculated as shown in Equation 11:

$$\text{3BBO ESP}_{\text{Coinbase},t} = 2 \times \text{BuySell}_{\text{Coinbase},t} \times \frac{\text{Price}_{\text{Coinbase},t} - \text{Mid}_{\text{3BBO},t}}{\text{Mid}_{\text{3BBO},t}} \tag{11}$$

We estimate very small average quoted and effective spreads. The average local quoted (effective) spread of the 16,913 BTC trades on Coinbase is less than 0.1 (0.5) bps. Across all exchanges and tokens, local quoted and effective spreads consistently remain below 3 bps. In fact, local quoted spreads are below one basis point in all but two cases (Coinbase DOGE at 1.05 bps and Kraken SOL at 1.49 bps). Furthermore, the 3BBO quoted spreads are frequently negative, reflecting potential arbitrage opportunities across these three exchanges. The 3BBO effective spreads on Coinbase are positive, while most of the 3BBO effective spreads for Kraken and Crypto.com are negative.

Negative effective spreads arise if the trade price is above (below) the midpoint for a seller (buyer) and may indicate extreme price improvements, potentially because of hidden liquidity. For example, if there is an iceberg hidden buy order that is not fully visible in the order book, a market sell order may execute at a price above the quoted midpoint, resulting in a negative effective spread.

In equity markets, buy and sell indicators are often inferred using the Lee and Ready algorithm, under which negative effective spreads are mechanically impossible. A trade price above (below) the midpoint price is automatically considered a buyer (seller) initiated trade. In contrast, our dataset provides the actual trade direction, allowing negative effective spreads to arise.

Negative spreads could also indicate stale quoted prices. While the average and the median time between trades and quotes in our data is below one second for both Coinbase and Crypto.com, the average delay is approximately ten seconds for Kraken.

Panel B reports the same summary statistics for token-tracking ETFs using TAQ data. We separate the results by crypto exchange because the ETF data are sampled only around our actual trade times, which differ across exchanges. Thus, in Panel B, the exchange distinction reflects differences in sample timing. Local midquotes are therefore undefined and QSP and ESP values are reported only using the NBBO midquote.

Continuing the previous example, we find 12,128 trades on BTC tracking ETFs around our 18 BTC trades on Coinbase. We find that average quoted and effective spreads are much larger than in Panel A. The average quoted (effective) spread of the 12,128 trades on BTC tracking ETFs is around 4.8 (3.5) bps. Panel B also shows much larger variations in spreads across ETFs. Spreads for ETFs tracking BTC are in generally small, below 7 bps. Spreads for ETFs tracking ETH and especially SOL are, in general, larger, ranging from around 3 to 34 bps for ETH and 300 to 400 bps for SOL.

3 Results

We now provide several results indicating that transaction costs on crypto exchanges are large, might be underestimated when relying on publicly available data (Kaiko), and that cheaper alternatives seem to exist (token-tracking ETFs).

3.1 Univariate Results

We calculate effective spreads for our actual trades, using displayed prices as reported by our brokers and matched with quotes as reported from Kaiko prevailing our trades. This allows us to determine the prevailing best bid and ask on each exchange at the exact moment of trade, providing a benchmark to evaluate execution quality. As in Table 2, we calculate local and 3BBO effective spreads.

[Place Table 3 about here]

Table 3 summarizes the transaction cost estimates for our actual trades across trading platforms. The first column shows the explicit fee ratio ($\frac{\text{Explicit Fee}}{\text{All in price}}$), expressed in basis points. Coinbase, Kraken, and Crypto.com all include spreads in displayed prices and the explicit fee ratios starting at approximately 0.8%.[7] Coinbase, Kraken, and Crypto.com operate and, to the best of our knowledge, execute retail trades on their own exchanges in addition to their brokerage platforms.

In contrast, Robinhood does not charge explicit fees for standard PFOF transactions; instead, those transactions include a 0.85% spread embedded in the displayed price. While Robinhood acquired crypto exchange Bitstamp in June 2025, their current execution methods allow execution across multiple exchanges. Appendix Table A.2 provides additional details for each broker, including screenshots of account trading interfaces and disclosed spreads and fees.

[7]Traders can bypass these explicit fees to a limited extent via subscriptions. For example, a "Coinbase One" subscription costs $4.99 per month and waives explicit fees for up to $500 of monthly trading volume.

Columns 2 and 3 of Table 3 report the effective spreads, comparing the trade price to the local midquote. This approach limits the sample to Coinbase, Kraken, and Crypto.com. For Robinhood trades, which do not have a local exchange benchmark, we instead use the quote midpoint from the paired exchange in the transaction. For example, in simultaneous Coinbase–Robinhood trade pairs, we use Coinbase's midpoint as the benchmark for Robinhood trades.

The effective spreads in Column 2 using displayed prices span a wide range from -1.5 to 481 bps. Coinbase exhibits small negative effective spreads across all tokens except ETH, for which the effective spread is positive but minimal at 0.3 bp. Robinhood offers the second lowest transaction costs, with effective spreads ranging from 133 to 147 bps. Kraken and Crypto.com have the largest effective spreads, ranging from 253 to 481 bps.

Column 3 reports effective spreads calculated using all in prices. Once explicit fees are incorporated, Coinbase becomes the most expensive broker, with effective spreads ranging from 702 to 834 bps. Thus, although Coinbase appears to offer exceptionally narrow effective spreads based on displayed prices in Column 2, its substantial explicit fees result in the highest overall trading costs. Kraken is the second most expensive platform, with effective spreads ranging from 493 to 591 bps, followed by Crypto.com, with spreads ranging from 253 to 481 bps. Robinhood becomes the least expensive platform, with effective spreads ranging from 133 to 147 bps. These results suggest that Coinbase's favorable displayed prices obscure the high cost of its explicit fees relative to competitors.

In Columns 4 and 5, which utilize the 3BBO mid price, the results for Coinbase, Kraken, and Crypto.com are mostly similar to those in Columns 2 and 3, confirming that the results are insensitive to the use of either midpoint benchmark.[8]

It is worth noting that these large fees in the crypto market make it impractical for us to construct a sample of crypto trades as large as Schwarz et al. (2025) in the equity market.

Table 4 compares our transaction costs with several alternative benchmarks: simultaneous trades on Robinhood (Columns 1 and 4), trades from the Kaiko exchange order book within a one minute window before and after each of our actual trades (Columns 2 and 5), and ETF trades within a one minute window before and after each of our actual trades (Columns 3 and 6). Panel A calculates effective spreads using displayed prices provided by the broker at the time of execution, while Panel B calculates effective spreads using all-in prices that include explicit fees. Columns 1 to 3 use the midpoint quote from the Kaiko exchange that corresponds to the retail broker's local exchange, whereas Columns 4 to 6 use the Kaiko 3BBO midpoint quote.

To ensure that our benchmarks (trades from Kaiko and TAQ) are not driven by any

[8]Under the 3BBO framework, the difference in effective spreads across platforms is proportional to their nominal price difference because the shared 3BBO midquote benchmark cancels out in the numerator. We retain the ESP terminology here to maintain consistency with the subsequent tables.

potential extreme observation, we do not simply use the nearest trade in time, but instead compute the equally weighted average across all trades executed during the pre- and post-trade one-minute intervals.

[Place Table 4 about here]

Column 1 reports the mean difference in effective spread, measured in basis points, between trades executed on Coinbase, Kraken, or Crypto.com and simultaneous trades on Robinhood (e.g., $ESP_{Coinbase,t}^{Our\ trades} - ESP_{Robinhood,t}^{Our\ trades}$). For BTC and ETH, Kraken exhibits the highest effective spreads, ranging from 174 to 249 bps above simultaneous Robinhood trades. Crypto.com is generally less expensive than Kraken but remain significantly more costly than Robinhood. Crypto.com spreads range from 84 to 109 bps higher. For SOL, Crypto.com exhibits the highest effective spread at 314 bps above Robinhood, followed by Kraken at 270 bps. Coinbase is significantly cheaper than Robinhood.

Column 2 compares our transaction costs with trades from the exchange order book within a one minute window before and after each of our actual trades (e.g.,$ESP_{Coinbase,t}^{Our\ trades} - ESP_{Coinbase,t}^{Kaiko\ trades}$). Similar to the results in Column 1, Kraken exhibits the highest effective spreads for BTC and ETH, followed by Crypto.com.

Column 3 compares our transaction costs with ETF trades executed within a one minute window before and after each of our actual trades (e.g., $ESP_{Coinbase,t}^{Our\ trades}$ - $ESP(NBBO)_{Coinbase,t}^{ETF\ trades}$).[9] Kraken and Crypto.com exhibit higher effective spreads than ETFs, whereas effective spreads for Coinbase are lower than those for ETFs. For Robinhood, effective spreads are higher relative to ETFs for BTC and ETH, but lower for SOL. The lower spreads for SOL are likely due to the small number of observations obtained for SOL ETFs as a consequence of their recent introduction and relatively low liquidity during the sample period. Columns 4 to 6 use the Kaiko 3BBO midpoint quotes to calculate effective spreads for our trades (e.g., Column 4: $ESP(3BBO)_{Coinbase,t}^{Our\ trades} - ESP(3BBO)_{Robinhood,t}^{Our\ trades}$;
Column 5: $ESP(3BBO)_{Coinbase,t}^{Our\ trades} - ESP(3BBO)_{Coinbase,t}^{Kaiko\ trades}$;
Column 6: $ESP(3BBO)_{Coinbase,t}^{Our\ trades} - ESP(NBBO)_{Coinbase,t}^{ETF\ trades}$). The results for Coinbase, Kraken, and Crypto.com remain largely similar to those reported in Columns 1 to 3. Relative to ETF trades in Column 6, Robinhood exhibits transaction costs that are approximately 129 bps higher for BTC and 100 bps higher for ETH, while appearing less expensive than SOL ETFs.

Panel B calculates effective spreads using all-in prices that include explicit fees, and the results remain broadly consistent with those reported in Panel A using displayed prices. Except for Coinbase, which now is identified as the most expensive crypto exchange,

[9] As in Section 2.4, we use the $ESP(NBBO)$ because there is no $ESP(Local)$ for ETF trades. Note that $ESP(NBBO)_{Coinbase,t}^{ETF\ trades}$ represents the $ESP(NBBO)$ for ETF trades executed within a one minute window before and after each of our Coinbase trades.

consistent with results of Table 3. Overall, the all-in, round-trip-equivalent transaction costs (effective spreads) we incurred through the basic retail interfaces of Coinbase, Kraken, and Crypto.com is substantially higher, often several percentage points higher, than the corresponding costs implied by each of our three benchmarks.

To make all-in trading costs fully comparable between our actual crypto trades and ETF trades, however, the ETF effective spreads would also need to incorporate explicit transaction costs. These costs are generally small for U.S. investors and are unlikely to materially affect our conclusions. For example, Interactive Brokers charges no commission for U.S.-listed stocks and ETFs under IBKR Lite, while IBKR Pro Fixed charges $0.005 per share, subject to a $1 minimum and a maximum of 1% of trade value.[10]

It is important to note that all three exchanges exhibit significantly higher all-in transaction costs through their basic retail interfaces than those implied by their order-book data. This difference is not driven solely by separately charged fees: even before such fees are added, the costs implied by the displayed execution prices on Kraken and Crypto.com exceed their corresponding order-book estimates. For Coinbase and Kraken, incorporating separately itemized fees widens the gap further; for Crypto.com, the relevant charge is already embedded in the displayed execution price.

This divergence likely reflects differences in access channels. Our analysis examines trades executed through basic retail interfaces, whereas the more favorable order-book prices may reflect executions through advanced trading platforms, such as Kraken Pro. Nevertheless, the results raise an important question about representativeness. Although exchange-reported order-book data may accurately characterize executions in the professional market segment, they may not capture the prices available to customers trading through basic retail interfaces.

3.2 Multivariate Results

Next, we examine the execution quality measures in a multivariate framework. In addition to quoted spreads and effective spreads, we further calculate realized spread and price impact using the one-minute future midpoint, as well as price improvement.[11] Robinhood is the omitted reference category, so all coefficients are interpreted relative to this baseline group. Panel A of Table 5 includes results without controls while Panel B includes the natural logarithm of dollar trade size as a control variable and token fixed effects. We use heteroskedasticity-robust standard errors.

[Place Table 5 about here]

[10] See https://www.interactivebrokers.com/en/pricing/commissions-stocks.php?region=americas, accessed on 2026-09-25.

[11] Using 5-minute and 10-minute future midpoints to calculate realized spreads and price impact produce similar results.

Using the local exchange midpoint benchmark, Coinbase exhibits lower local effective spreads than the baseline group, Robinhood, for displayed prices, though this difference is statistically insignificant (-142.26 bps). In contrast, Kraken and Crypto.com exhibit significantly higher displayed price effective spreads than Robinhood, by 222.08 bps and 200.62 bps, respectively. However, when evaluating all-in prices, Coinbase becomes the most expensive trading platform, with an all-in effective spread that is 606.62 bps higher than Robinhood's. Kraken is the second most expensive venue, exceeding Robinhood by 421.30 bps, while Crypto.com remains significantly more costly by 200.62 bps.

The regression results for realized spreads closely mirror those of the effective spreads. Because realized spreads capture trade profitability before revenue-sharing or inventory adjustments, these highly significant, positive coefficients indicate that market makers or liquidity providers on Kraken and Crypto.com retain substantial execution revenues. On Coinbase, this source of revenue is entirely driven by explicit fees, as evidenced by the swing from a negative displayed realized spread (-141.03 bps) to a large, statistically significant all-in realized spread (607.86 bps).

The results for price improvement are the opposite of the spread metrics. Platforms with higher effective and realized spreads show significantly negative price improvement relative to the baseline. For example, using displayed prices, Kraken and Crypto.com reduce price improvement by 221.87 bps and 200.79 bps, respectively, compared to Robinhood.

The results remain robust in Panel B after controlling for trade size and adding token fixed effects. This implies that the execution cost differences across these platforms are stable and not driven by variations in trade volume or underlying asset composition.

[Place Table 6 about here]

Table 6 reports the regression results based on the 3BBO midquote benchmark. These are our preferred specifications because they account for cross-venue pricing opportunities, rendering the performance differences across trading platforms sharper and more economically meaningful. Overall, the results remain structurally consistent with the local midpoint specifications in Table 5.

These findings are consistent with the equity payment-for-order-flow (PFOF) literature (e.g., Battalio and Holden, 2001; Schwarz et al., 2025). In our setting, the PFOF-based platform (Robinhood) provides lower all-in transaction costs than fee-based retail platforms, whose explicit fees substantially increase total trading costs.

[Place Table 7 about here]

[Place Table 8 about here]

Tables 7 and 8 report the regression results for our execution quality measures: quoted spread, effective spread, realized spread, price impact, and price improvement. These measures are based on trades and quotes in the exchange order book at the time of our actual trades using Kaiko data. Because the Kaiko data cover only Coinbase, Kraken, and Crypto.com, these analyses are restricted to those three exchanges, with Coinbase serving as the reference group. Table 7 reports the results based on the local mid price, while Table 8 is based on the 3BBO mid price.

Table 7, Column 1 reports the regression results for local quoted spreads. Relative to the baseline group, Coinbase, there is no statistically significant difference for Kraken, while Crypto.com exhibits slightly smaller local quoted spreads. In Panel B, after controlling for trade size and token fixed effects, Kraken's local quoted spread is slightly higher at the 10% significance level, while Crypto.com remains slightly lower. In Columns 2, 3 and 4, there are no statistically significant differences across the exchanges for local effective spreads, realized spreads, and price improvement, respectively. In Column 5, the price impact component is generally small and statistically indistinguishable across brokers, suggesting that execution venue rather than adverse selection drives the cost differences. This is consistent with the notion that retail crypto trades are relatively uninformed, with brokers competing primarily on the spread and fee dimensions rather than managing adverse selection risk. After controlling for trade size and token fixed effects in Panel B, the results for Columns 2 through 5 remain largely consistent.

Table 8 reports the results based on the 3BBO mid price. Relative to the Coinbase baseline, there is no significant difference in 3BBO quoted spreads for Kraken or Crypto.com. The remaining variables—effective spread, realized spread, price improvement, and price impact—remain statistically insignificant across trading platforms. Panel B specifications, which control for trade size and token fixed effects, remain entirely consistent. Taken together, these results indicate that execution quality measures do not vary dramatically among the main trading platforms.

Overall, the results in this section naturally raise the question of why our trading costs are so large, even though the effective spreads reported by crypto exchanges (and captured by Kaiko) are small, and alternative ETFs also offer better execution.

One possible explanation is that exchanges obscure trading costs by displaying only a single displayed price rather than separate bid and ask prices. In addition, not all exchanges clearly indicate how much spread and fees are charged, which makes cross platform comparisons difficult. Even so, the magnitude of the costs remains strikingly high.

4 Policy Implications

Our findings document large differences in retail crypto trading costs across platforms and substantial gaps between those costs and exchange-based benchmarks. These comparisons do not identify the effects of regulation, but they point to three policy priorities: clearer retail cost disclosures, more comparable market-wide benchmarks, and better information about alternative ways to obtain crypto exposure.

4.1 Retail Cost Disclosure and Best Execution

Regulators should prioritize standardized disclosure of all-in trading costs across retail platforms and access channels. Disclosures should distinguish displayed execution prices, separately charged fees, and any disclosed charges embedded in prices, and report their combined effect on transaction costs. Aggregate execution statistics should identify the trading interface, order size, and reference benchmark. These distinctions matter because low order-book spreads need not imply low costs through a platform's basic retail interface, as we find, for example, for Coinbase which uses variable fees per transaction.

Best-execution standards for retail crypto transactions warrant consideration alongside these disclosures. Assessments should distinguish execution quality from separately charged fees while considering the total cost to the customer and the availability of alternative execution arrangements. Related evidence illustrates the potential value of these comparisons. Ernst et al. (2025b) document price improvement from off-exchange crypto liquidity relative to an exchange-based benchmark. In equities, Ernst et al. (2025a) show that brokers allocate more orders to wholesalers with better past execution performance, giving wholesalers incentives to compete for future order flow.

The comparison across business models also argues for evaluating outcomes rather than fee labels alone. Robinhood's PFOF-based executions have lower all-in costs than the basic retail interfaces of the three exchanges in our sample. This finding does not identify the effect of PFOF, but it shows that the absence of PFOF is insufficient to ensure lower customer costs. Execution assessments should therefore consider all-in costs together with routing incentives and potential conflicts of interest.

4.2 Consolidated Reporting and Cross-Exchange Benchmarks

Policymakers should consider standardized reporting of crypto quotes and transactions as a basis for cross-exchange execution benchmarks. Such reporting would make prices easier to compare and provide a common reference for evaluating execution quality. A credible benchmark requires clear rules for venue inclusion, synchronized timestamps, quote freshness, executable depth, and the treatment of fees. Negative quoted spreads in our synthetic 3BBO illustrate the importance of these choices: crossed recorded quotes

do not by themselves establish profitable arbitrage after accounting for timing, access constraints, and trading costs.

Consolidating order-book data alone would not resolve the measurement problem documented here. A benchmark could accurately summarize exchange quotes while omitting charges incurred through basic retail interfaces. Consolidated market reporting should therefore be complemented by disclosures of actual retail execution prices and fees. This combination would allow investors and regulators to assess both underlying market conditions and the additional costs associated with particular access channels, without presuming that a consolidated benchmark would itself eliminate those costs.

4.3 ETFs as an Alternative Access Channel

An additional policy lever that has received increasing attention is the approval and expansion of cryptocurrency exchange-traded funds (ETFs). While ETFs are often discussed primarily in the context of investor access, a substantial body of literature suggests that they also have important implications for market quality, liquidity, and price efficiency. For example, ETF arbitrage may affect underlying-market liquidity, consistent with the role of arbitrage as liquidity provision in segmented markets (Rösch, 2021). ETFs could also place competitive pressure on retail crypto platforms by offering an alternative access channel.

However, our findings suggest that improving institutional spot liquidity is insufficient to directly resolve the substantial effective spreads faced by retail investors. The large all-in effective spreads documented in our study seem primarily driven by opaque dealer markups within fragmented retail platforms, rather than a lack of underlying global liquidity.

Instead of mechanically fixing retail execution costs, the primary policy benefit of the spot crypto ETF is that it acts as a transparent, highly regulated substitute. By providing a traditional access point with clearly defined fees and tightly constrained secondary market spreads, the ETF establishes a public benchmark for execution costs. Retail investors can effectively bypass the larger spreads of crypto brokers by routing their demand through traditional brokerage accounts. Over time, this substitution effect may introduce competitive pressure. To prevent a potential exodus of market share to crypto ETF issuers, retail crypto platforms may be forced to compress their hidden markups and transition to more retail-friendly fee structures.

5 Conclusion

This paper provides the first systematic measurement of all-in retail trading costs in U.S. cryptocurrency markets, using a controlled experiment across four major brokers and four

tokens over several months. Our key finding is that all-in effective spreads in crypto market are extraordinarily large, dwarfing the trading costs in equity and options markets. The magnitude of these costs implies that retail crypto investors pay hundreds or thousands of dollars per year in avoidable trading costs, costs that are largely invisible because they are at least in part bundled into prices rather than disclosed as explicit commissions. More importantly, benchmarks based on Kaiko data or token-tracking ETF trades substantially understate the transaction costs we observe through basic retail crypto interfaces.

This finding has important implications for studies that evaluate cryptocurrency trading strategies using exchange-based transaction-cost estimates. For example, Liu et al. (2022) find that several strategies remain profitable after accounting for estimated fees and bid–ask spreads of less than 1% (p. 1167). While our estimates based on Kaiko data are similarly small, the costs implied by our actual trades are substantially larger, with all-in transaction costs ranging from approximately 2.5% to more than 8%. These markedly higher costs suggest that strategies that appear profitable using exchange-based cost estimates may generate negative net returns for retail investors.

We document substantial variation in all-in costs across brokers. Robinhood's PFOF-based executions have the lowest costs, while trading through the basic retail interfaces of Coinbase, Kraken, and Crypto.com is substantially more expensive. Separately charged fees are particularly important at Coinbase, where they turn displayed prices close to the midpoint into the highest all-in costs in our sample. These findings show that a fee-based business model does not necessarily deliver lower transaction costs.

Our multivariate analysis confirms these patterns after controlling for token and dollar trade size. 3BBO-based regressions reveal that Coinbase is the most expensive option, running about 600 bps higher than Robinhood. Kraken and Crypto.com are also pricier, though by smaller margins of about 480 and 185 bps, respectively. Price impact is economically small and statistically indistinguishable across brokers, indicating that execution venue selection—not adverse selection—drives the cross-broker cost differences. Smaller retail trades face higher per-unit costs, with implications for the equitable distribution of market-access costs across the retail investor population.

Negative 3BBO quoted spreads for BTC, ETH, and SOL indicate crossed quotes across exchanges, although establishing executable arbitrage requires accounting for quote timing, market access, and trading costs. Our findings support standardized, Rule 605-style disclosures of all-in retail execution costs as an immediate policy priority. Policymakers should also consider consolidated quote and transaction reporting, a crypto NBBO with clear venue eligibility criteria, and best-execution standards that account for accessible liquidity and total customer costs. Consolidated order-book data alone might remain insufficient if they omit charges incurred through variable explicit fees, as used by Coinbase.

These measures could facilitate comparison and competition and potentially reduce trading costs. In equities, Ernst et al. (2025a) show that brokers allocate more orders to

wholesalers with better execution performance, creating incentives to compete on execution quality. Our evidence motivates greater transparency in crypto markets by documenting substantial cost differences across retail platforms that conventional exchange-based benchmarks do not capture.

References

Battalio, R., Holden, C. W., 2001. A simple model of payment for order flow, internalization, and total trading cost. Journal of Financial Markets 4, 33–71.

Battalio, R., Schultz, P., 2011. Regulatory uncertainty and market liquidity: The 2008 short sale ban's impact on options markets. Journal of Finance 66, 2013–2053.

Ben-David, I., Franzoni, F., Moussawi, R., 2018. Do etfs increase volatility? Journal of Finance 73, 2471–2535.

Boulton, T. J., Shohfi, T., Walz, M., Miller, B., 2025. How does payment for order flow influence markets? evidence from robinhood crypto token introductions. Working paper, Miami University.

Chowdhury, K. P., Shohfi, T., 2026. The global and local impact of the introduction of bitcoin and ethereum exchange traded products (etps) on bitcoin and ethereum liquidity and returns. Available at SSRN 6905798 .

Cong, L. W., Li, X., Tang, K., Yang, Y., 2023. Crypto wash trading. Management Science 69, 6427–6454.

Divakaruni, A., Zimmerman, P., 2024. Uncovering retail trading in bitcoin: The impact of covid-19 stimulus checks. Management Science 70, 2066–2085.

Dyhrberg, A. H., Foley, S., Svec, J., 2023. When bigger is better: The impact of a tiny tick size on undercutting behavior. Journal of Financial and Quantitative Analysis 58, 2387–2416.

Easley, D., O'Hara, M., Basu, S., 2019. From mining to markets: The evolution of bitcoin transaction fees. Journal of Financial Economics 134, 91–109.

Ernst, T., Malenko, A., Spatt, C. S., Sun, J., 2025a. What does best execution look like? Working paper, University of Maryland.

Ernst, T., Spatt, C. S., Routledge, B., Zetlin-Jones, A., 2025b. Dark crypto. Available at SSRN 5186551 .

Foley, S., Krekel, W., Mollica, V. B., Svec, J., 2023. Not so fast: Identifying and remediating slow and imprecise cryptocurrency exchange data. Finance Research Letters 51, 103401.

Glosten, L., Milgrom, P., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. Journal of Financial Economics 14, 71–100.

Hadad, E., Malhotra, D., Nippani, S., 2024. Trading commodity etfs: Price behavior, investment insights, and performance analysis. Journal of Futures Markets 44, 1257–1276.

Hendershott, T., Khan, S. A., Riordan, R., 2025. Option auctions. Review of Financial Studies 39, 783–834.

Huang, R., Stoll, H., 1997. The components of the bid–ask spread: A general approach. Review of Financial Studies 10, 995–1034.

Huang, X., Jorion, P., Schwarz, C., 2025. Some anonymous options trades are more equal than others. Cornell SC Johnson College of Business Research Paper .

Kogan, S., Makarov, I., Niessner, M., Schoar, A., 2024. Are cryptos different? evidence from retail trading. Journal of Financial Economics 159, 103897.

Lee, C. M. C., Ready, M. J., 1991. Inferring trade direction from intraday data. Journal of Finance 46, 733–746.

Lettau, M., Madhavan, A., 2018. Exchange-traded funds 101 for economists. Journal of Economic Perspectives 32, 135–154.

Li, T., Shin, D., Wang, B., 2021. Cryptocurrency pump-and-dump schemes. Available at SSRN 3267041 .

Liu, X., Zhang, T. T., Zhang, Y., 2026. A hidden cost of etf investing: Retail demand shocks and limits to arbitrage. Journal of Banking & Finance p. 107621.

Liu, Y., Tsyvinski, A., Wu, X., 2022. Common risk factors in cryptocurrency. Journal of Finance 77, 1133–1177.

Makarov, I., Schoar, A., 2020. Trading and arbitrage in cryptocurrency markets. Journal of Financial Economics 135, 293–319.

Muravyev, D., Pearson, N., 2020. Options trading costs are lower than you think. Review of Financial Studies 33, 4973–5014.

O'Hara, M., Ye, M., 2011. Is market fragmentation harming market quality? Journal of Financial Economics 100, 459–474.

Pan, J., Poteshman, A. M., 2006. The information in option volume for future stock prices. Review of Financial Studies 19, 871–908.

Pursiainen, V., Toczynski, J., 2023. Retail investors' cryptocurrency investments. Research Paper 23-51, Swiss Finance Institute.

Rösch, D., 2021. The impact of arbitrage on market liquidity. Journal of Financial Economics 142, 195–213.

Schwarz, C., Barber, B., Huang, X., Jorion, P., Odean, T., 2025. The "actual retail price" of equity trades. Journal of Finance 80, 2507–2541.

Stoll, H., 1989. Inferring the components of the bid–ask spread: Theory and empirical tests. Journal of Finance 44, 115–134.

Stoll, H. R., 2000. Friction. Journal of Finance 55, 1479–1514.

Tables & Figures

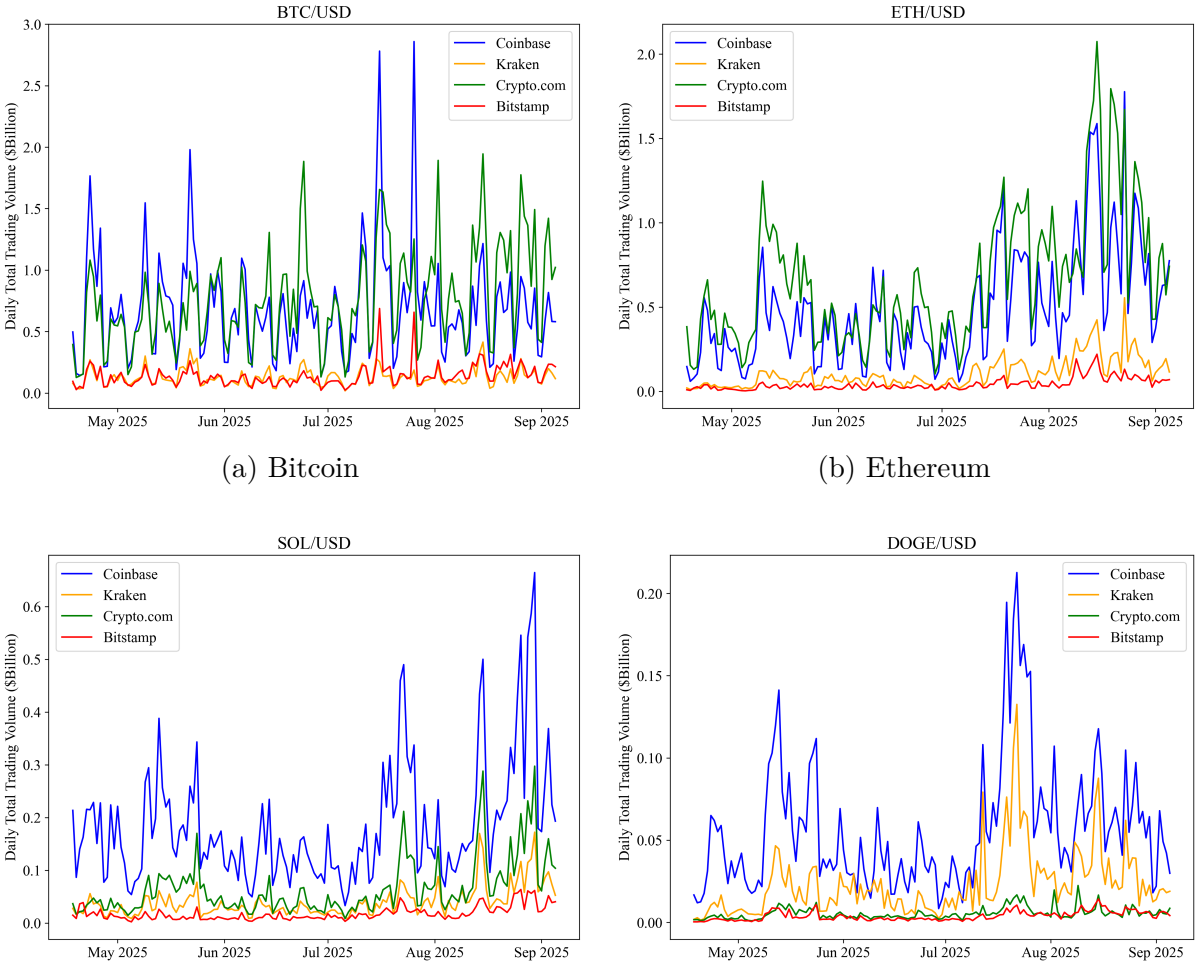

(a) Bitcoin

(b) Ethereum

(c) Solana

(d) Dogecoin

Figure 1: Daily Total Trading Volume across Exchange for Each Token.
This figure displays daily total trading volume (in $billion) for each of the four tokens (Bitcoin, Ethereum, Solana, and Dogecoin) across four exchanges (Coinbase, Kraken, Crypto.com, and Bitstamp) using Kaiko data.

Table 1: Transaction and Dollar Volume Summary for Actual Trades
This table presents a breakdown of our trading activity from April 17, 2025 to September 3, 2025 across exchanges for each token. Panel A details the transaction counts categorized by buy and sell orders, including aggregate totals per exchange and per token. Panel B reports our total dollar trading volume for the same categories.

Panel A – Transaction count by exchange and symbol

Exchange		BTC	ETH	SOL	DOGE	Total Buy/Sell per exchange
Coinbase	Buy	10	10	10	10	40
	Sell	8	9	10	9	36
Kraken	Buy	7	10	10	10	37
	Sell	7	8	10	10	35
Crypto.com	Buy	15	10	8	8	41
	Sell	15	10	7	7	39
Robinhood	Buy	32	30	28	28	118
	Sell	30	27	27	26	110
Total trades per token		124	114	110	108	
					Total Buy	236
					Total Sell	220

Panel B – Total dollar trading volume by exchange and symbol

Exchange		BTC	ETH	SOL	DOGE
Coinbase	Buy	$1973.21	$1267.04	$1751.49	$1700.00
	Sell	$1908.36	$1181.09	$1850.16	$1335.00
Kraken	Buy	$1550.59	$2000.00	$2000.19	$1520.00
	Sell	$1389.96	$1695.00	$1700.52	$1400.00
Crypto.com	Buy	$2300.00	$1500.00	$900.00	$900.00
	Sell	$2080.00	$1300.00	$700.00	$700.00
Robinhood	Buy	$5823.80	$4767.05	$4651.68	$4120.00
	Sell	$5378.32	$4176.09	$4250.68	$3435.00
Total value per token		$22,404.24	$17,886.27	$17,804.72	$15,110.00

Table 2: Summary Statistics for Kaiko and ETF Benchmarks
This table reports summary statistics computed within a one-minute window before and after each of our actual trades. Panel A covers Kaiko exchange order book data, and Panel B covers ETF trade data. Reported metrics include the total number of trades ($\#Trades$), average dollar trade size ($\$Size$), total dollar trading volume ($\$Volume$), average quoted spread in bps (QSP), and average effective spread in bps (ESP). Average values are calculated as follows: we first compute the equally weighted average within the one-minute window surrounding each of our actual trades. We then take the equally weighted mean within each subset (e.g., Coinbase BTC).

Panel A – Kaiko

					Local		3BBO	
		$\#Trades$	$\$Size$	$\$Volume$	QSP	ESP	QSP	ESP
Coinbase	BTC	16,913	1,640	29,517,197	0.0951	0.4317	-3.0446	0.1616
	ETH	9,078	308	3,039,663	0.1647	0.5888	-3.0854	0.8577
	SOL	17,992	311	6,439,643	0.6178	1.1704	-3.0618	1.4126
	DOGE	15,632	254	4,189,105	1.0469	2.9786	-0.3498	2.9439
Kraken	BTC	929	2,755	3,689,072	0.3759	1.4100	-3.0572	-0.9234
	ETH	505	1,049	512,886	0.3632	0.0563	-3.3594	-2.7378
	SOL	1,017	1,959	3,267,647	1.4888	1.2551	-1.7898	-0.6530
	DOGE	147	146	21,352			-1.0345	0.7170
Crypto.com	BTC	11,589	3,204	45,782,574	0.0011	0.5322	-3.2795	-0.3229
	ETH	22,774	2,149	48,198,376	0.0307	1.1829	-2.9357	-0.9226
	SOL	10,257	888	8,958,106	0.4958	1.6904	-2.9460	1.2995
	DOGE	1,904	146	355,509	0.5040	2.7589	-0.7340	-0.6396

Panel B – ETF

					Local		NBBO	
		$\#Trades$	$\$Size$	$\$Volume$	QSP	ESP	QSP	ESP
Coinbase	BTC	12,128	172	3,167,407			4.836	3.494
	ETH	16	247	3,514			33.974	28.399
Kraken	BTC	12,621	187	2,349,770			1.795	1.212
	ETH	675	335	224,012			4.722	3.147
Crypto.com	BTC	361	131	41,892			6.864	4.905
	SOL	13	118	1,281			409.0	302.8
Robinhood	BTC	25,111	165	5,559,158			4.412	3.144
	ETH	691	269	227,526			26.661	22.086
	SOL	13	118	1,281			409.0	302.8

Table 3: Transaction Costs for Our Actual Trades
This table reports the explicit fee ratio and the effective spreads (both in basis point) for our trades across exchanges for each token. The explicit fee ratio is defined as the per-token explicit fee relative to the all-in price, according to the following formula:

$$Explicit\ Fee\ Ratio_{i,t} = \frac{Explicit\ Fee_{i,t}}{All\ in\ Price_{i,t}}$$

Effective spreads relative to each exchange's midpoint are calculated using the following formula:

$$ESP(Local)_{i,t} = 2 \times BuySell_{i,t} \times \frac{Price_{i,t} - Local\ Mid_{i,t}}{Local\ Mid_{i,t}}$$

Effective spreads relative to the 3BBO midpoint are calculated using the following formula:

$$ESP(3BBO)_{i,t} = 2 \times BuySell_{i,t} \times \frac{Price_{i,t} - 3BBO\ Mid_{i,t}}{3BBO\ Mid_{i,t}}$$

$BuySell_{i,t}$ is the buy/sell indicator and takes a value of +1 (-1) for a buyer-initiated (seller-initiated) order at the time of trade t for exchange i. Two prices are used in the calculation: the displayed price (Displayed) and the all-in price (All-in), defined as total dollar trade size divided by the total number of tokens traded. Mid prices are from Kaiko, BuySell and trade prices are from our actual trade data. Note that, in contrast to Coinbase and Kraken, the displayed price for Crypto.com already incorporates the explicit fee.

		(1)	(2)	(3)	(4)	(5)
		Explicit Fee Ratio	*ESP(Local)*		*ESP(3BBO)*	
			Price Used			
			Displayed	All-in	Displayed	All-in
Coinbase	BTC	350.8	-1.5023	702.3	-1.9167	702.0
	ETH	415.8	0.2972	834.3	0.4774	834.5
	SOL	375.5	-0.4042	752.0	-1.1750	751.2
	DOGE	376.1	-0.1997	702.5	-0.0869	702.7
Kraken	BTC	99.43	294.4	493.3	293.3	492.1
	ETH	99.47	387.4	587.0	386.7	586.3
	SOL	99.52	391.4	590.5	390.9	590.0
	DOGE	99.62			567.8	767.0
Crypto.com	BTC	80	253.4	253.4	253.3	253.3
	ETH	80	277.9	277.9	278.1	278.1
	SOL	80	480.7	480.7	480.3	480.3
	DOGE	80	468.4	468.4	468.2	468.2
Robinhood	BTC	0	142.9	142.9	142.9	142.9
	ETH	0	147.1	147.1	147.1	147.1
	SOL	0	133.1	133.1	132.5	132.5
	DOGE	0	145.1	145.1	157.7	157.7

Table 4: Comparison of Transaction Costs Between Our Actual Trades and Benchmarks This table reports the mean differences in effective spread (ESP), measured in basis points, between our trades and several alternative benchmarks. The alternatives include simultaneous trades on Robinhood, trades from the Kaiko exchange order book within a one minute window before and after each of our actual trades, and ETF trades from TAQ within a one minute window before and after each of our actual trades. Panel A computes ESP using displayed prices, while Panel B uses all-in transaction prices. t-values are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A – Displayed price

		(1)	(2)	(3)	(4)	(5)	(6)
		$ESP(Local)$			$ESP(3BBO)$		
		vs Robinhood	vs Kaiko	vs ETF	vs Robinhood	vs Kaiko	vs ETF
Coinbase	BTC	-117.13***	-1.93	-3.84***	-117.13***	-2.08	-4.22***
		(-103.41)	(-1.65)	(-6.11)	(-103.44)	(-1.37)	(-4.55)
	ETH	-131.26***	-0.29	-28.59***	-131.26***	-0.38	-28.46***
		(-60.58)	(-0.97)	(-4.36)	(-60.57)	(-0.71)	(-4.71)
	SOL	-119.10***	-1.57***		-119.09***	-2.59**	
		(-211.40)	(-4.47)		(-211.57)	(-2.34)	
	DOGE	-125.37***	-3.18		-125.38***	-3.03	
		(-116.17)	(-0.01)		(-116.08)	(-0.01)	
Kraken	BTC	172.65***	293.02***	293.22***	172.64***	294.20***	292.06***
		(66.01)	(191.39)	(206.96)	(65.92)	(132.83)	(148.01)
	ETH	248.38***	387.31***	338.32***	248.35***	389.42***	330.99**
		(65.88)	(91.17)	(66.40)	(65.80)	(92.87)	(47.32)
	SOL	269.80***	390.17***		269.81***	391.58***	
		(180.20)	(334.22)		(180.73)	(200.76)	
	DOGE				397.48***	567.10***	
					(27.60)	(39.88)	
Crypto.com	BTC	84.27***	252.84***	321.40***	84.27***	253.64***	321.26***
		(7.50)	(22.40)	(131.79)	(7.50)	(22.32)	(143.97)
	ETH	108.57***	276.70***		108.57***	278.98***	
		(38.29)	(97.00)		(38.30)	(99.07)	
	SOL	313.23***	479.03***	179.77**	313.24***	479.00***	176.59**
		(53.27)	(91.43)	(4.01)	(53.19)	(84.49)	(3.97)
	DOGE	298.03***	465.62***		298.02***	468.84***	
		(23.59)	(38.02)		(23.59)	(37.90)	
Robinhood	BTC			129.86***			129.31***
				(38.08)			(37.68)
	ETH			101.22***			99.51***
				(12.56)			(14.58)
	SOL			-131.11**			-134.30**
				(-3.32)			(-3.43)

Panel B – All-in price

		(1)	(2)	(3)	(4)	(5)	(6)
		$ESP(Local)$			$ESP(3BBO)$		
		vs Robinhood	vs Kaiko	vs ETF	vs Robinhood	vs Kaiko	vs ETF
Coinbase	BTC	586.70***	701.89***	707.04***	586.74***	701.80***	706.73***
		(13.76)	(16.44)	(15.91)	(13.76)	(16.44)	(15.89)
	ETH	702.73***	833.70***	875.23***	702.73***	833.61***	875.37***
		(15.98)	(19.41)	(9.71)	(15.98)	(19.38)	(9.63)
	SOL	633.34***	750.87***		633.27***	749.77***	
		(14.29)	(16.90)		(14.29)	(16.91)	
	DOGE	577.37***	699.56*		577.39***	699.74*	
		(13.65)	(1.92)		(13.65)	(1.92)	
Kraken	BTC	371.53***	491.91***	492.11***	371.52***	493.07***	490.94***
		(121.04)	(230.67)	(238.64)	(120.87)	(177.32)	(206.59)
	ETH	448.00***	586.92***	541.74***	447.93***	589.01***	534.33***
		(122.69)	(133.14)	(104.88)	(122.39)	(140.50)	(75.43)
	SOL	468.88***	589.26***		468.90***	590.67***	
		(205.90)	(278.86)		(206.46)	(218.52)	
	DOGE				596.68***	766.30***	
					(41.96)	(54.56)	
Crypto.com	BTC	84.27***	252.84***	321.40***	84.27***	253.64***	321.26***
		(7.50)	(22.40)	(131.79)	(7.50)	(22.32)	(143.97)
	ETH	108.57***	276.70***		108.57***	278.98***	
		(38.29)	(97.00)		(38.30)	(99.07)	
	SOL	313.23***	479.03***	179.77**	313.24***	479.00***	176.59**
		(53.27)	(91.43)	(4.01)	(53.19)	(84.49)	(3.97)
	DOGE	298.03***	465.62***		298.02***	468.84***	
		(23.59)	(38.02)		(23.59)	(37.90)	
Robinhood	BTC			129.86***			129.31***
				(38.08)			(37.68)
	ETH			101.22***			99.51***
				(12.56)			(14.58)
	SOL			-131.11**			-134.30**
				(-3.32)			(-3.43)
	DOGE						

Table 5: Regression Results for Liquidity Measures for Our Trades using local mid
This table reports regression results for various measures of execution quality of our actual trades across trading platforms. Robinhood is the omitted reference category, so all coefficients are interpreted relative to this baseline group. The dependent variables include effective spreads relative to local mid (ESP ($Local$)), realized spreads relative to local mid (RSP ($Local$)), and price improvement relative to local mid ($PIMPROVE$ ($Local$)). For each measure, results are presented across two prices: the displayed price from our transaction (Displayed), and the all-in price from our transaction (All-in) calculated as total dollar trade size divided by the total number of tokens traded. Panel A presents results without controls, while Panel B includes the natural logarithm of dollar trade size as a control variable and token fixed effects. We use heteroskedasticity-robust standard errors. t-values are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A – Without controls

	(1)	(2)	(3)	(4)	(5)	(6)
	ESP ($Local$)		RSP ($Local$)		$PIMPROVE$ ($Local$)	
price used	Displayed	All-in	Displayed	All-in	Displayed	All-in
Coinbase	-142.26	606.62***	-141.03	607.86***	142.29	-606.60***
	(-1.50)	(6.26)	(-1.49)	(6.27)	(1.50)	(-6.26)
Kraken	222.08***	421.30***	221.59***	420.80***	-221.87***	-421.08***
	(6.71)	(12.72)	(6.68)	(12.68)	(-6.70)	(-12.71)
Crypto.com	200.62***	200.62***	199.77***	199.77***	-200.79***	-200.79***
	(5.75)	(5.75)	(5.72)	(5.72)	(-5.76)	(-5.76)
Constant	141.82***	141.82***	141.33***	141.33***	-141.40***	-141.40***
	(4.36)	(4.36)	(4.33)	(4.33)	(-4.34)	(-4.34)
Num.Obs.	416	416	416	416	416	416
R^2	0.068	0.199	0.067	0.199	0.068	0.199
Token FE	NO	NO	NO	NO	NO	NO

Panel B – With controls

	(1)	(2)	(3)	(4)	(5)	(6)
	ESP ($Local$)		RSP ($Local$)		$PIMPROVE$ ($Local$)	
price used	Displayed	All-in	Displayed	All-in	Displayed	All-in
Coinbase	-145.84	608.34***	-144.60	609.58***	145.79	-608.39***
	(-1.59)	(6.55)	(-1.58)	(6.56)	(1.59)	(-6.55)
Kraken	227.84***	440.25***	227.06***	439.46***	-227.57***	-439.98***
	(8.87)	(16.49)	(8.82)	(16.47)	(-8.86)	(-16.47)
Crypto.com	200.04***	185.38***	199.39***	184.72***	-200.17***	-185.50***
	(5.58)	(5.11)	(5.54)	(5.08)	(-5.58)	(-5.12)
log($\$Size$)	-18.34	-101.90*	-17.49	-101.05*	18.37	101.93*
	(-0.36)	(-1.93)	(-0.34)	(-1.92)	(0.36)	(1.93)
Constant	206.07	625.71**	201.28	620.92**	-206.18	-625.82**
	(0.79)	(2.34)	(0.77)	(2.32)	(-0.79)	(-2.34)
Num.Obs.	416	416	416	416	416	416
R^2	0.070	0.207	0.069	0.207	0.070	0.207
Token FE	YES	YES	YES	YES	YES	YES

Table 6: Regression Results for Liquidity Measures for Our Trades Using 3BBO Mid
This table reports regression results for various measures of execution quality of our actual trades across trading platforms. Robinhood is the omitted reference category, so all coefficients are interpreted relative to this baseline group. The dependent variables include effective spreads relative to 3BBO mid (ESP (3BBO)), realized spreads relative to 3BBO mid (RSP (3BBO)), and price improvement relative to 3BBO mid ($PIMPROVE$ (3BBO)). For each measure, results are presented across two prices: the displayed price from our transaction (Displayed), and the all-in price from our transaction calculated as total dollar trade size divided by the total number of tokens traded (All-in). Panel A presents results without controls, while Panel B includes the natural logarithm of dollar trade size as a control variable and token fixed effects. We use heteroskedasticity-robust standard errors. t-values are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A – Without controls

	(1)	(2)	(3)	(4)	(5)	(6)
	ESP (3BBO)		RSP (3BBO)		$PIMPROVE$ (3BBO)	
price used	Displayed	All-in	Displayed	All-in	Displayed	All-in
Coinbase	-145.60	603.29***	-144.35	604.54***	145.62	-603.27***
	(-1.64)	(6.62)	(-1.62)	(6.63)	(1.64)	(-6.62)
Kraken	275.08***	474.28***	275.27***	474.47***	-274.89***	-474.10***
	(22.09)	(38.08)	(22.03)	(37.95)	(-22.23)	(-38.33)
Crypto.com	197.42***	197.42***	196.89***	196.89***	-197.67***	-197.67***
	(15.65)	(15.65)	(15.66)	(15.66)	(-15.75)	(-15.75)
Constant	144.93***	144.93***	144.32***	144.32***	-147.34***	-147.34***
	(94.59)	(94.59)	(90.82)	(90.82)	(-96.25)	(-96.25)
Num.Obs.	456	456	456	456	456	456
R^2	0.158	0.352	0.157	0.352	0.158	0.352
Token FE	NO	NO	NO	NO	NO	NO

Panel B – With controls

	(1)	(2)	(3)	(4)	(5)	(6)
	ESP (3BBO)		RSP (3BBO)		$PIMPROVE$ (3BBO)	
price used	Displayed	All-in	Displayed	All-in	Displayed	All-in
Coinbase	-146.91*	605.35***	-145.71	606.55***	146.88*	-605.37***
	(-1.66)	(6.75)	(-1.65)	(6.76)	(1.66)	(-6.75)
Kraken	272.81***	484.23***	272.84***	484.25***	-272.76***	-484.17***
	(24.05)	(37.88)	(23.89)	(37.75)	(-23.99)	(-37.79)
Crypto.com	200.47***	185.27***	200.14***	184.93***	-200.54***	-185.34***
	(15.77)	(13.58)	(15.71)	(13.51)	(-15.76)	(-13.57)
log($\$Size$)	-19.40	-98.35***	-18.17	-97.12***	19.50	98.45***
	(-0.57)	(-2.73)	(-0.53)	(-2.70)	(0.57)	(2.73)
Constant	206.24	602.94***	199.67	596.38***	-209.89	-606.59***
	(1.16)	(3.23)	(1.12)	(3.19)	(-1.18)	(-3.24)
Num.Obs.	456	456	456	456	456	456
R^2	0.167	0.367	0.166	0.367	0.167	0.367
Token FE	YES	YES	YES	YES	YES	YES

Table 7: Regression Results for Liquidity Measures for Kaiko Trades Using Local Mid
This table reports regression results for various measures of execution quality of Kaiko trades across trading platforms. Coinbase is the omitted reference category, so all coefficients are interpreted relative to this baseline group. The dependent variables include Displayed spreads using local mid (QSP ($Local$)), effective spreads relative to local mid (ESP ($Local$)), realized spreads relative to local mid (RSP ($Local$)), price improvement relative to local mid ($PIMPROVE$ ($Local$)), and price impact using local mid ($PIMPACT$ ($Local$)). The price used for calculating ESP ($Local$), RSP ($Local$), and $PIMPROVE$ ($Local$) is the last execution trade price in each exchange's order book at the time of our trades. Panel A presents results without controls, while Panel B includes the natural logarithm of dollar trade size as a control variable and token fixed effects. We use heteroskedasticity-robust standard errors. t-values are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A – Without controls

	(1) QSP ($Local$)	(2) ESP ($Local$)	(3) RSP ($Local$)	(4) $PIMPROVE$ ($Local$)	(5) $PIMPACT$ ($Local$)
Kraken	0.20	0.14	-1.60	0.06	1.74
	(0.99)	(0.41)	(-1.01)	(0.13)	(1.09)
Crypto.com	-0.19**	-0.09	-2.17	-0.10	2.09
	(-2.48)	(-0.16)	(-1.22)	(-0.18)	(1.23)
Constant	0.45***	-0.27	0.47	0.72**	-0.74
	(8.01)	(-1.01)	(0.39)	(2.50)	(-0.60)
Num.Obs.	208	208	208	208	208
R^2	0.035	0.001	0.009	0.000	0.009
Token FE	NO	NO	NO	NO	NO

Panel B – With controls

	(1) QSP ($Local$)	(2) ESP ($Local$)	(3) RSP ($Local$)	(4) $PIMPROVE$ ($Local$)	(5) $PIMPACT$ ($Local$)
Kraken	0.34*	-0.10	-2.16	0.44	2.06
	(1.78)	(-0.29)	(-1.39)	(0.97)	(1.34)
Crypto.com	-0.09*	-0.28	-2.22	0.18	1.94
	(-1.66)	(-0.41)	(-1.17)	(0.26)	(1.10)
log($Size)	-0.04	-0.35	0.37	0.32	-0.72
	(-0.33)	(-0.59)	(0.22)	(0.51)	(-0.43)
Constant	0.22	2.02	-0.82	-1.80	2.84
	(0.36)	(0.61)	(-0.09)	(-0.52)	(0.32)
Num.Obs.	208	208	208	208	208
R^2	0.213	0.025	0.014	0.054	0.011
Token FE	YES	YES	YES	YES	YES

Table 8: Regression Results for Liquidity Measures for Kaiko Trades Using 3BBO Mid
This table reports regression results for various measures of execution quality of Kaiko trades across trading platforms. Coinbase is the omitted reference category, so all coefficients are interpreted relative to this baseline group. The dependent variables include quoted spreads using 3BBO mid (QSP ($3BBO$)), effective spreads relative to 3BBO mid (ESP ($3BBO$)), realized spreads relative to 3BBO mid (RSP ($3BBO$)), price improvement relative to 3BBO mid ($PIMPROVE$ ($3BBO$)), and price impact using 3BBO mid ($PIMPACT$ ($3BBO$)). The price used for calculating ESP ($3BBO$), RSP ($3BBO$), and $PIMPROVE$ ($3BBO$) is the last execution trade price in each exchange's order book at the time of our trades. Panel A presents results without controls, while Panel B includes the natural logarithm of dollar trade size as a control variable and token fixed effects. We use heteroskedasticity-robust standard errors. t-values are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A – Without controls

	(1) QSP ($3BBO$)	(2) ESP ($3BBO$)	(3) RSP ($3BBO$)	(4) $PIMPROVE$ ($3BBO$)	(5) $PIMPACT$ ($3BBO$)
Kraken	0.17	-0.56	-1.63	0.73	1.06
	(0.48)	(-0.67)	(-1.00)	(0.84)	(0.72)
Crypto.com	-0.27	0.04	-1.74	-0.31	1.78
	(-0.84)	(0.06)	(-1.03)	(-0.40)	(1.12)
Constant	-2.39***	-0.49	0.15	-1.89***	-0.64
	(-10.16)	(-1.26)	(0.13)	(-4.07)	(-0.55)
Num.Obs.	228	228	228	228	228
R^2	0.008	0.003	0.006	0.007	0.006
Token FE	NO	NO	NO	NO	NO

Panel B – With controls

	(1) QSP ($3BBO$)	(2) ESP ($3BBO$)	(3) RSP ($3BBO$)	(4) $PIMPROVE$ ($3BBO$)	(5) $PIMPACT$ ($3BBO$)
Kraken	0.08	-0.44	-1.67	0.51	1.23
	(0.26)	(-0.51)	(-0.99)	(0.59)	(0.80)
Crypto.com	-0.05	-0.25	-1.66	0.21	1.41
	(-0.16)	(-0.32)	(-0.93)	(0.24)	(0.86)
log($\$Size$)	0.10	-0.84	0.79	0.94	-1.62
	(0.33)	(-0.97)	(0.49)	(1.10)	(-1.11)
Constant	-3.69**	4.13	-3.39	-7.82*	7.52
	(-2.16)	(0.89)	(-0.40)	(-1.70)	(0.98)
Num.Obs.	228	228	228	228	228
R^2	0.235	0.015	0.010	0.080	0.012
Token FE	YES	YES	YES	YES	YES

Appendix

Table A.1: Brokers Used for Experiment
For the brokers used to trade cryptocurrencies in our experiment, the table summarizes whether they accept payment for order flow (PFOF) for cryptocurrency trading, their corporate structure, and their regulatory environment.

Exchange	**PFOF**	**Public/Private**	**Regulatory Environment**
Robinhood	Yes	Public: Ticker: HOOD (IPO July 2021). Acquired Bitstamp June 2, 2025.	SEC-registered; U.S. regulated.
Robinhood Smart Router	No	Internal Tool	Operates under the Robinhood regulatory umbrella but functions as a non-PFOF execution path.
Coinbase	No	Public: Ticker: COIN (IPO April 2021).	SEC-registered; FINRA and SIPC member. U.S. regulated.
Kraken	No	Private	Wyoming-chartered SPDI (Kraken Financial); MSB with FinCEN. Partially U.S. regulated for custody.
Crypto.com	No	Private	U.S. registered MSB; operates globally with localized licenses across multiple jurisdictions.

Table A.2: Fee Information by Brokerage
This table provides details about each brokerage including screenshots of accounts and disclosed spreads/fees.

Exchange	Spreads	Fees	Screenshots
Robinhood Market Maker Routing	0.85% spread included in displayed price.	No fees.	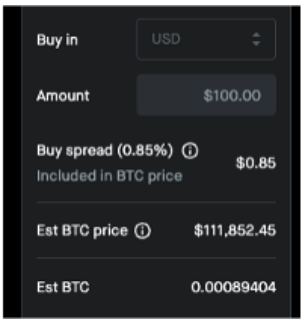
Robinhood Smart Router		0.85% flat fee for 30-day trading volume under $50K, decreases for higher volumes according to table. 0.02% goes to the executing partner exchange(s) and is included in fee. $0.01 minimum fee.	
Coinbase	1% spread included in displayed price.	Trading fee is "'calculated at the time you place your order and can be influenced by factors such as your chosen payment method, order size, market conditions, jurisdictional location, asset, and other costs we incur to facilitate your transaction.'"	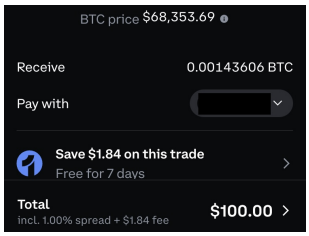
Kraken	Spread included in displayed price.	Fixed 1% trading fee.	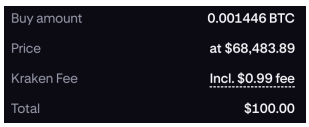

Exchange	Spreads	Fees	Screenshots
Crypto.com	Spread included in displayed price.	Estimates from actual trading find that the fee is approximately 0.7937%.	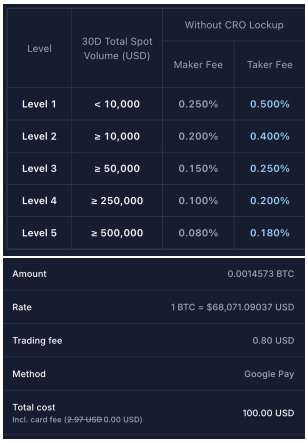

Level	30D Total Spot Volume (USD)	Without CRO Lockup	
		Maker Fee	Taker Fee
Level 1	< 10,000	0.250%	0.500%
Level 2	≥ 10,000	0.200%	0.400%
Level 3	≥ 50,000	0.150%	0.250%
Level 4	≥ 250,000	0.100%	0.200%
Level 5	≥ 500,000	0.080%	0.180%

Amount	0.0014573 BTC
Rate	1 BTC = $68,071.09037 USD
Trading fee	0.80 USD
Method	Google Pay
Total cost Incl. card fee (~~2.97 USD~~ 0.00 USD)	100.00 USD

Elastic in Cash, Inelastic in Repo: Hedge Funds in the Treasury and Repo Markets *

Andrea Poinelli[†], Loriana Pelizzon[‡], Davide Tomio[§], Benoit Nguyen[¶], and Tobias Linzert[‖].

This draft: September 2, 2026 - Click here for the latest version

Abstract

Sovereign bond markets are a cornerstone of the financial system, and their functioning is tightly linked to repo markets, where investors finance long positions and source bonds for short sales. We show, theoretically and empirically, that repo prices are set in the cash bond market: when demand for cash bonds exceeds available supply, arbitrageurs accommodate the excess by shorting bonds and borrowing them in the repo market, opening a wedge between the policy and repo rates, i.e., generating specialness. An elastic supply of collateral, in turn, limits how much of the excess demand is capitalized into bond prices. Using regulatory data covering the universe of repos backed by German sovereign bonds, we identify the final borrowers and lenders of securities and estimate the first demand and supply elasticities for a repo market. Supply, dominated by the public sector, accounts for 87% of the aggregate elasticity. Demand, driven by hedge funds, is strongly inelastic: a 10% increase in borrowing costs reduces their borrowing by only 1%. Despite being among the most price-elastic investors in cash bond markets, hedge funds are inelastic in repo, as their borrowing sustains relative-value positions whose size is pinned by the preferred-habitat demand they intermediate: their elasticity is inherited from their cash-market counterparties rather than being a primitive. Specialness thus emerges as the equilibrium price of cash-market demand pressure—of which collateral scarcity from central bank purchases is a special case—tying safe-asset pricing and the transmission of monetary policy to the same imbalances in the cash bond market.

JEL Classification Codes: E51, E52, E58, G21.

Keywords: Monetary Policy, Repo Market, Bond Market, Safe Assets.

---

*The views expressed here are our own and do not necessarily reflect those of the Deutsche Bundesbank, the European Central Bank (ECB), or the eurosystem. We thank Nina Boyarchenko, Manav Chaudhary, Stefania DAmico, Ruggero Japelli, Gabriele La Spada, Matteo Leombroni, Angelo Ranaldo and the participants in the University of Chicago Booth Treasury Markets Conference, 8th Short-Term Funding Markets Conference, ChaMP 3rd Workshop in Banca d'Italia, ECB RCC1 Workshop, 12th University of York-BoE-BdF Asset Pricing Workshop, and seminar participants at Texas A&M University, the Università della Svizzera Italiana, the Bank of England, ECB's Directorate General Market Operations. We are grateful to ECB's Directorate General Market Operations's staff for their help with the data, and Imène Rahmouni-Rousseau for her continued support.

[†]European Central Bank; Deutsche Bundesbank; Goethe University Frankfurt; andrea.poinelli@bundesbank.de

[‡]Leibniz Institute for Financial Research SAFE; Goethe University Frankfurt; Ca' Foscari University of Venice; CEPR; pelizzon@safe-frankfurt.de

[§]Texas A&M University, Mays Business School; CEPR; tomio@tamu.edu

[¶]European Central Bank; benoit.nguyen@ecb.europa.eu

[‖]European Central Bank; tobias.linzert@ecb.europa.eu

Non-technical summary

Repo and sovereign bond markets are deeply interconnected. Investors use repo markets both to finance bond holdings and to borrow securities needed for short sales. Yet most existing explanations for repo "specialness"—the premium paid to borrow a particular bond—focus on collateral shortages or dealer balance-sheet stress. This paper proposes a different explanation: repo specialness originates from imbalances in the cash bond market.

The mechanism is straightforward. Many investors, such as pension funds, insurance companies, and mutual funds, have strong incentives or regulatory requirements to hold physical government bonds. When demand for these bonds exceeds available supply, hedge funds step in to satisfy the imbalance. They short scarce bonds and hedge the position using futures contracts, profiting from the gap between cash and futures prices. Because short positions require borrowing bonds in repo markets, excess demand in the cash market mechanically generates repo borrowing demand.

A key implication follows. The quantity hedge funds need to borrow is determined primarily by the imbalance in the cash bond market, not by repo market conditions. As a result, hedge fund demand for borrowing securities is relatively insensitive to borrowing costs, even when those costs become very high. In other words, hedge funds behave as highly price-sensitive traders in cash bond markets, but as highly price-insensitive borrowers in repo markets. By contrast, public-sector institutions such as central banks and government debt managers behave in the opposite way: relatively inelastic in the cash market, but highly elastic suppliers of collateral in repo markets.

We test this mechanism using the Securities Financing Transactions Data Store (SFTDS), a regulatory dataset covering nearly all repo transactions involving euro-area counterparties. Because most European repo trades are centrally cleared, the raw data contain multiple reports of the same transaction and often obscure the identity of the ultimate borrower and lender. This is the first paper to clean and exploit SFTDS at scale, developing a novel deduplication algorithm that reconstructs ultimate borrower-lender pairs by looking through central counterparties—which clear 80–90% of repo volume—and consolidating duplicate reports of the same economic transaction. The resulting dataset provides the most comprehensive "who-to-whom" view of the German repo market to date.

Our findings reveal a highly polarized market structure. Public-sector institutions account for about 60% of net securities lending, while offshore hedge funds account for 50%–80% of net securities borrowing. Traditional long-term investors such as pension funds and insurance companies are largely absent from repo borrowing activity. The 2022 specialness episode illustrates the mechanism: cash-futures arbitrage opportunities widened, hedge funds expanded their short positions, and repo specialness increased. When arbitrage opportunities normalized in 2023, repo specialness declined.

Our results yield three implications for policy. First, repo specialness serves as a timely signal of cash-market demand pressure, providing central banks with an observable indicator for monitoring monetary policy transmission. Second, the design and pricing of securities lending facilities have first-order effects on specialness: our sectoral decomposition indicates that the public sector's facilities reduced specialness by approximately 15 basis points at the peak of the 2022 episode. Third, because repo demand inelasticity originates in the cash market, specialness exists as long as profitable basis-trade opportunities persist—unless bond supply fully satisfies preferred-habitat demand or securities lending takes place at the policy rate. Supply-side interventions in the cash market must therefore be accompanied by monitoring of repo-market conditions.

1 Introduction

Sovereign bonds are a cornerstone of the financial system, and their prices are shaped by the interaction of cash and repo markets: intermediaries and arbitrageurs finance bond holdings and source bonds for short sales in the repo market, so sovereign bond and repo markets clear jointly. The price of borrowing a specific bond—repo specialness, the wedge between the policy rate and the bond's repo rate—is capitalized into the bond's price and yield curve (Duffie, 1996; DAmico and Pancost, 2022; Jappelli, Subrahmanyam, and Pelizzon, 2023), and the divergence of repo rates from the policy rate impairs the transmission of monetary policy to short-term rates (Duffie and Krishnamurthy, 2016; Nguyen, Tomio, and Vari, 2023; Ballensiefen, Ranaldo, and Winterberg, 2023).

The nexus between the two markets has moved to the center of policy debates as hedge funds have become dominant borrowers in government-bond repo, funding the basis trades whose rapid unwinding is a leading concern for Treasury market stability (Barth and Kahn, 2025). Yet the bond and repo markets are typically studied in isolation, and remarkably little is known about who sets the price in the repo market or what determines its level. We provide the first answer with a model in which repo prices are set by demand pressure originating from the cash bond market, testing it with the first estimation of demand and supply elasticities for the repo market.

The existing literature treats specialness as a crisis phenomenon, owing to dealer balance-sheet constraints (Andersen, Duffie, and Song, 2019; Brunnermeier and Pedersen, 2009); as a byproduct of unconventional monetary policy, driven by the collateral scarcity that central bank purchases induce (Arrata, Nguyen, Rahmouni-Rousseau, and Vari, 2020; Corradin and Maddaloni, 2020; Pelizzon, Subrahmanyam, and Tomio, 2025); or as a result of trading frictions, with search costs and dealer market power generating dispersion in funding rates (Infante, 2020; Eisenschmidt, Ma, and Zhang, 2020). Specialness, however, prevails even when balance sheets are unconstrained and collateral supply is stable. We show that specialness arises in a repo market that clears competitively at a single price, with unconstrained arbitrageurs, as the equilibrium response to demand pressure from the cash bond market—of which central bank purchases are a special case.

Our analysis is organized around the model's central prediction: hedge funds—among the most price-elastic investors in cash bond markets (Chaudhary, Fu, and Zhou, 2024; Jansen, Li, and Schmid, 2024)—are price-inelastic in the repo market. Unlike the decision to enter the basis trade, repo borrowing is not discretionary (Duffie, 1996): it is an input required to sustain relative-value positions whose size is set elsewhere, namely, by the preferred-habitat

investors' demand that hedge funds satisfy by short selling the bond in the cash bond market.

The economic mechanism is simple. Preferred-habitat investors—pension funds, insurance companies, bond mutual funds—value physical bonds for regulatory and institutional reasons that make synthetic substitutes imperfect. When their demand exceeds available supply, the cash bond becomes more expensive than the synthetic substitute: bases, such as the cash-futures basis or the asset-swap spread, open up.[1] Hedge funds exploit the wedge by shorting the bond against an offsetting synthetic position, borrowing the bond in the repo market to deliver it, effectively expanding the supply of bonds available to preferred-habitat investors. The cash-market imbalance thus translates one-for-one into repo borrowing demand. On the other side of the repo market, dealers and the public sector lend securities, with supply that is upward-sloping in specialness due to intermediation costs and institutional design. Specialness is the price at which the repo market supports hedge funds' intermediation of cash-market demand.

This mechanism delivers a theoretical insight that applies beyond the repo market. In demand-system asset pricing, estimated elasticities are commonly interpreted as preference or technology parameters of the estimated sector. Our model shows that for an intermediary whose position in one market is mechanically tied to its position in another, the elasticity is instead an equilibrium object: the arbitrageur's elasticity in one market is determined, through market clearing, by the elasticity of its counterparties in the *other* market.

In our setting, the linkage works as follows. Because hedge funds arbitrage the spread between the basis and specialness, their willingness to adjust repo borrowing is not limited by their own price sensitivity, but by how much quantity their cash-market counterparties release when prices move. The slope of hedge funds' demand for borrowing bonds on the repo market: the price sensitivity of preferred-habitat investors.[2] Symmetrically, the slope of their cash-market supply schedule is governed by the repo-market friction. The repo market is the cleanest laboratory for this cross-market determination: hedge funds trade there only to fund relative-value positions, never as an independent strategy, so their demand in the repo market is just as inelastic as preferred-habitat investors' demand is in the cash market.

[1] The cash-futures basis compares the price of the cash bond with the price implied by the corresponding futures contract; the asset-swap spread compares the bond's yield with the fixed rate of a maturity-matched interest-rate swap.

[2] Hedge funds' own frictions—transaction costs, price impact, balance-sheet charges—would enter as well, but Appendix A.4 shows that the equilibrium elasticity never exceeds the cash-market depth, and approaches it as those frictions become small relative to the rigidity of preferred-habitat demand.

Equilibrium prices follow from this intermediation structure. Specialness reflects the intersection of hedge funds' inelastic derived demand with imperfectly elastic securities-lending supply. Any shift in cash-market conditions—in preferred-habitat demand or bond supply—moves specialness through hedge funds' derived demand, and thereby moves bond prices through the specialness premium embedded in yields. The model predicts that time variation in market-wide specialness is driven by the size of cash-market imbalances, not by changes in price sensitivities.

Testing these predictions requires observing who ultimately borrows and lends each bond for every transaction—a level of detail that, to our knowledge, no available dataset provides for the U.S. Treasury market. We obtain such visibility for the European repo market from the Securities Financing Transactions Data Store (SFTDS), a regulatory dataset that records who-to-whom transactions for the universe of European repos. Ours is the first paper to use it to study repo pricing and its link to investor positions, and to estimate a demand system for a repo market, which we do using the Optimal Granular Instrumental Variables (O-GIV) framework of Chaudhary et al. (2024). Identification exploits the granularity of the market: the idiosyncratic changes in the positions of large sectors are, under the model's factor structure, uncorrelated with aggregate demand shifts and thus serve as instruments for price changes.

We focus on the repo market for German government bonds: the euro area's safe asset and the world's second-largest safe-asset repo market, with approximately EUR 200 billion in outstanding transactions. Its structure—preferred-habitat demand, hedge fund basis trades, and repo-financed positions—mirrors that of the U.S. Treasury market.[3] We combine SFTDS with confidential data on the ECB's securities lending operations and identify ultimate borrowers and lenders, looking through intermediaries and central counterparties.

The market topology matches the model. Repo participation is effectively decoupled from bond ownership: insurance companies and pension funds hold over EUR 120 billion in German government bonds yet are virtually absent from the repo market, while hedge funds—predominantly domiciled in the Cayman Islands and Bermuda, the same jurisdictions that dominate U.S. Treasury basis-trade activity (Barth and Kahn, 2025)—hold few bonds yet absorb 50% of net securities borrowing directly, and 80% once their positions through foreign dealers are included. The public sector, through the ECB's securities lending program and the German debt management office's retained holdings, provides approximately 60%

[3] In recent years, U.S. basis trades have typically involved buying the cash bond and selling futures, whereas Bund richness implies the opposite legs. The mechanism—repo-financed arbitrage linking the cash and derivative markets, with the repo leg pinned by the cash position—is analogous across geographies.

of net securities lending, while euro-area banks and dealers run matched books. Preferred-habitat investors generate the demand pressure in the cash market but do not appear in the repo market; the hedge funds that accommodate their demand in the cash bond market do so by borrowing the securities in the repo market.

The elasticity estimates confirm the model's central prediction. Hedge funds exhibit a repo elasticity of 0.11: a 10% increase in borrowing costs reduces their demand by only 1%. The estimate is economically small but statistically different from zero—precisely what the model predicts when preferred-habitat investors retain some, albeit limited, price sensitivity—and consistent with the cash-market evidence in Koijen, Koulischer, Nguyen, and Yogo (2021). The market as a whole shows an aggregate repo-market elasticity of 1.2, only 13% of which originates on the demand side and 87% on the supply side. Among suppliers, the public sector is elastic, with an institutional nuance the model helps interpret: the German debt management office supplies bonds broadly, while the ECB's facility is more responsive, consistent with its role as a backstop.

Comparing our repo-market estimates with cash-market elasticities for the same sectors shows the cross-market inversion at the heart of the framework. Hedge funds are among the most elastic investors in the cash bond market (Jansen et al., 2024; Chaudhary et al., 2024) but the least elastic in the repo market; the public sector displays the opposite pattern, purchasing bonds according to policy mandates regardless of price while lending them elastically in repo. The inversion is what the mechanism predicts for an intermediary that exercises discretion in one market and covers positions in the other. The determination, moreover, runs in both directions: just as hedge funds' repo elasticity reflects preferred-habitat investors' price sensitivity, their high cash-market elasticity rests on the elastic securities-lending supply of dealers and the public sector in the repo market. Neither elasticity is a hedge-fund primitive.

The time series supports the model's remaining prediction that specialness moves with quantities, not as a result of changing price sensitivities.[4] During the ECB's rapid tightening in 2022, every measure of cash-market imbalance (cash-futures bases, asset-swap spreads, and pricing errors from fitted yield curves) widened, consistent with intensifying preferred-habitat demand. Hedge funds expanded their short positions, bond borrowing on the repo market surged, and specialness reached historically unprecedented levels. As imbalances receded in late 2022 and 2023—with increased issuance, moderating excess demand, and expanded securities lending by the German debt management office—borrowing and specialness normalized. Sub-sample estimates show hedge fund elasticities virtually unchanged

[4] Quarter-end spikes in specialness are the exception: they reflect dealers' balance-sheet constraints around reporting dates rather than cash-market demand pressure.

across the quantitative easing and tightening regimes. The dynamics ran through positions rather than price sensitivity, as the model implies when repo elasticity is anchored by the structural rigidity of preferred-habitat demand.

Our results yield three policy implications. First, aggregate repo specialness is a timely, observable signal of cash-market demand pressure, and it co-moves with the bases and pricing errors that central banks monitor for transmission. Second, the design and pricing of securities lending facilities have first-order effects on both repo pricing—our sectoral decomposition indicates that public-sector facilities reduced specialness by approximately 15 basis points at the peak of the 2022 episode—and cash bond pricing. Third, because repo demand inelasticity originates in the cash market, specialness exists as long as profitable basis-trade opportunities persist. It can be contained only by satisfying preferred-habitat demand in the cash market or by lending securities at the policy rate.

The topology and the elasticities together carry a final financial-stability message that extends well beyond our setting: the marginal price-setters in the repo market for a major safe asset are offshore, non-bank entities outside the domestic regulatory perimeter. The transmission of monetary policy through repo rates thus increasingly depends on actors whom domestic regulators cannot directly supervise—and through whom shocks originating far from euro-area fundamentals can propagate into the market that anchors the euro-area risk-free rate. The same hedge funds, running the same basis trades from the same jurisdictions, are the dominant (cash) borrowers in U.S. Treasury repo. Our estimates provide the first evidence on how price-sensitive that borrowing is—evidence not yet available for the U.S. market. To the extent that the same derived-demand logic operates there, hedge funds' financing demand is similarly inelastic to repo rates, and shifts in their basis-trade positions would transmit directly into the rates from which SOFR—the benchmark that anchors U.S. dollar pricing—is constructed. As dispersion across money market interest rates is a primary indicator of the level of the pass-through inefficiency of monetary policy (Duffie and Krishnamurthy, 2016), hedge funds' participation in repo market deserve regulatory attention.

The remainder of the paper is organized as follows. Section 2 develops the theoretical framework. Section 3 describes the data and documents the structure of the repo market. Section 4 presents the empirical methodology, and Section 5 reports the elasticity estimates. Section 6 analyzes hedge fund behavior and the link between cash and repo markets. Section 7 concludes.

Literature Review

Our paper contributes to several strands of the literature. The first is the literature on repo pricing and its linkage to bond pricing, pioneered by Duffie (1996), who models repo demand as a byproduct of bond short-selling and shows that specialness is reflected in the price of the underlying bond, and by Krishnamurthy (2002) who shows that the on-the-run Treasury's premium over the off-the-run equals the expected cumulative cost of shorting it in the repo market, and that this premium is pinned down by investors who hold the new bond for liquidity reasons and lend out only part of their holdings. In both papers, specialness is an attribute of a single security relative to a close substitute, and the friction that generates it is tied to the auction cycle, so both the specialness and the associated price premium dissipate by the next auction. Subsequent work on repo pricing emphasized dealer balance-sheet constraints (Andersen et al., 2019; D'Avernas, Petersen, and Vandeweyer, 2025), liquidity spirals (Brunnermeier and Pedersen, 2009), collateral scarcity from central bank asset purchases (Arrata et al., 2020; Corradin and Maddaloni, 2020; Jappelli et al., 2023; Pelizzon et al., 2025), and its implication on monetary policy effectiveness (Duffie and Krishnamurthy, 2016; Nguyen et al., 2023). We complement this work by showing that specialness arises as a normal-times equilibrium outcome of cash-market demand pressures—of which collateral scarcity from asset purchases is a special case—and that the demand pressure we identify is neither security-specific nor tied to the auction cycle, so that specialness is a persistent, market-wide price. Our paper thereby provides the first market-wide empirical test of the mechanism proposed at the security level by Duffie (1996), using comprehensive transaction data to estimate a demand system for the repo market.

We contribute to the demand-system approach to asset pricing, pioneered by Koijen and Yogo (2019) and extended to sovereign bond markets by Koijen et al. (2021), Jansen et al. (2024), and Chaudhary et al. (2024). Most relevant to our work, Jansen et al. (2024) estimate sector-level demand functions for the U.S. Treasury market and document that hedge funds and primary dealers accommodate demand imbalances from preferred-habitat investors by holding short positions, a setting shared by Jappelli et al. (2023). Chaudhary et al. (2024) similarly find that hedge funds are the most elastic investors in the U.S. Treasury cash market. We show that the hedge funds' cash-market role has an exact repo-market counterpart, and that neither elasticity is a primitive: hedge funds' cash-market elasticity is governed by the frictions of the repo market where they source the bonds they short, while their repo-market elasticity is inherited from the preferred-habitat investors they accommodate in the cash market.

On the empirical side, we provide new evidence on the structure of the repo market and

hedge funds' prominent role in it. Our work complements Mancini, Ranaldo, and Wrampelmeyer (2015), who provide the first in-depth analysis of the euro-area repo market, Hermes, Schmeling, and Schrimpf (2025), who use SFTDS to study U.S.-dollar repo positions of European banks, and Ampudia, Formoso da Silva, Hesse, Putz, Schobert, von Landesberger, and Wohlert (2025), who study the impact of interventions in the cash and repo markets for German collateral. A nascent literature highlights the importance of hedge funds in the repo market. Barth and Kahn (2025) document that basis trades account for the large majority of hedge fund positions in the U.S. Treasury repo market, while Li, Ma, Mendicino, and Supera (2026) show that bank lending to non-bank financial institutions such as hedge funds affects credit supply to the real economy. We are the first to use transaction-level regulatory data to map the full topology of a safe-asset repo market, link investor positions to prices, and estimate sector-level elasticities.

2 Theoretical Framework

This section develops a simple one-period model designed to capture the mechanism underlying the elasticities estimated in our empirical analysis. The model features three markets: cash bonds, futures, and repo. The central agents are hedge funds, which link the cash bond and repo markets: they allow the former to clear by borrowing bonds in the latter, thereby increasing supply. Hedge funds are motivated by profiting from the wedge in price between the physical bond and a synthetically equivalent futures position. The wedge originates from preferred habitat investors' preference for holding the cash bond, which makes it "rich" relative to the futures.[5] Because shorting the bond requires borrowing it in the repo market, cash-market imbalances translate into hedge funds' repo bond borrowing demand. This demand, combined with the securities-lending supply from dealers and the central bank, determines the level of specialness at which the repo market clears.

The model is characterized by three traded assets: a cash bond, a futures contract on the same underlying bond, and a risk-free asset available in perfectly elastic supply at rate r_s, which can be interpreted as the policy rate. We derive demand and supply functions for the three markets from the agents' objective functions. The net supply of futures is zero, while the physical supply of the bond is exogenous. There are four types of agents: preferred habitat investors (PH), hedge funds (H), dealers (D), and the central bank (CB). Preferred habitat investors and hedge funds interact in the cash bond market. Hedge funds,

[5]While we refer to this synthetic asset as a futures contract, it may equally be thought of as an asset or interest-rate swap, and the richness can be in line with the definition of a bond's convenience yield.

dealers, and the central bank interact in the repo market. We present the objective functions of the four agents, then solve for cash- and repo-market equilibria.

2.1 The Hedge Fund

Hedge funds have mean-variance preferences and implement a range of strategies, including relative-value arbitrage. They trade in both the cash bond and futures markets and obtain bonds in the repo market to support short positions in the cash market. The relative-value strategy exploits the cash-futures wedge, or basis, defined as $x = r_F - r_T > 0$, where r_F is the futures bond- and r_T is the cash bond-yield. The basis represents the per-unit return advantage of a long futures position relative to a long cash-bond position.

Hedge funds implement this trade by shorting the bond and taking an offsetting long futures position, earning gross arbitrage profits equal to x before repo and margin costs. A positive basis, therefore, reflects a situation in which the cash bond is rich relative to futures: obtaining exposure through the physical bond is more expensive (lower yield) than replicating it synthetically, consistent with a premium paid by preferred habitat investors for holding the underlying security.

To short the bond, the hedge fund must borrow it in the repo market. Let $\chi > 0$ denote the specialness spread, i.e., the spread between the policy rate and the repo rate received from lending cash against the borrowed bond. Specialness is the per-unit cost of borrowing the scarce bond in the repo market to short-sell it in the cash market.

The representative hedge fund chooses a short position in the cash bond, $D_H < 0$, and a futures position with two components. The first is a hedging leg, denoted $D_F^H > 0$, which supports the basis trade and satisfies

$$D_F^H = -D_H. \tag{1}$$

The second is an additional directional futures position, N_F^H, which can be long or short.

The cash bond and the futures contract reference the same underlying bond, so their returns are perfectly correlated: $\text{corr}(r_T, r_F) = 1$. Since the basis-trade legs satisfy $D_F^H = -D_H$, the variance of the basis-trade component is zero. The directional trade exposes the fund to the common variance in bond and futures returns, σ^2, which, for simplicity, we normalize to 1.

Both the basis-trade futures leg and the directional futures position require margin. Let $m > 0$ denote the margin requirement per unit of futures exposure, and let $k > 0$ denote the opportunity cost of one unit of posted margin. Total margin usage must not exceed the fund's equity capital $E_H > 0$:

$$m\left(|D_F^H| + |N_F^H|\right) \leq E_H. \tag{2}$$

A hedge fund that shorts the bond and borrows it on the repo market pays χ per unit of the bond demanded, $|D_H|$. Once specialness and funding costs are accounted for, the basis trade earns per-unit profits of $x - \chi - mk$, and the directional futures position earns $N_F^H(r_F - r_s) - m|N_F^H|$. The basis trade is profitable if $x > \chi + mk$. We set up the model assuming that hedge funds are price-takers in both the cash and repo markets, and relax this assumption in Section 2.7.

The hedge fund's objective function is therefore:

$$\max_{D_F^H,\, N_F^H} \quad D_F^H(x - \chi - mk) + N_F^H(r_F - r_s) - km|N_F^H| - \frac{1}{2B}(N_F^H)^2 \tag{3}$$

subject to (1) and (2), where $B > 0$ is hedge fund risk tolerance. The directional futures position is standard and is characterized by a no-trade region induced by the margin friction.

Lemma 1 (Optimal basis trade)**.** *Two cases arise:*

(i) Basis trade profitable *($x > \chi + mk$): the hedge fund optimally sets $D_F^H = -D_H$, earning a riskless profit of $D_F^H(x - \chi - mk)$. Since this profit is linear and increasing in D_F^H, the hedge fund is willing to take as large a short bond position as possible.*

(ii) Basis trade unprofitable *($x \leq \chi + mk$): the hedge fund does not engage in the basis trade. $D_F^H = 0$ and $D_H = 0$.*

Lemma 2 (HF directional positions)**.** *The optimal directional futures position is*

$$N_F^{H*} = \begin{cases} 0, & \text{if } |r_F - r_s| \leq km, \\ B(r_F - r_s - km), & \text{if } r_F - r_s > km, \\ B(r_F - r_s + km), & \text{if } r_F - r_s < -km. \end{cases}$$

2.2 The Preferred Habitat Investors

Preferred habitat investors—pension funds, insurance companies, bond mutual funds—hold the physical bond both for structural reasons (regulatory mandates, liability matching,

benchmark tracking, Holm-Hadulla, Leombroni, Mota, Papoutsi, and Corell, 2024) and in response to the term premium (Vayanos and Vila, 2021). The representative preferred habitat investor chooses $D_{PH} \geq 0$ to maximize:

$$\max_{D_{PH} \geq 0} \quad D_{PH}\big(r_T - r_s\big) - \frac{1}{2\eta}\big(D_{PH} - \bar{D}\big)^2, \tag{4}$$

where $\bar{D} > 0$ is the target allocation, capturing mandate-driven demand, and $\eta > 0$ is the price sensitivity, governing willingness to deviate from the target in response to the term premium. The quadratic penalty $\frac{1}{2\eta}(D_{PH} - \bar{D})^2$ captures the costs of deviating from the mandate. We define structural excess demand as $\Delta_S \equiv \bar{D} - S_T$, the gap between target allocation and physical bond supply.

The first-order condition yields:

$$D^*_{PH} = \bar{D} + \eta\big(r_T - r_s\big). \tag{5}$$

Demand has two components: a structural component $\bar{D}$, independent of returns, and a price-sensitive component $\eta(r_T - r_s)$. A small η corresponds to a setting in which preferred habitat investors are unwilling to deviate from their target allocation. A zero $\bar{D}$, instead, recovers a standard mean-variance investor with no target and risk tolerance η.

2.3 The Repo Dealer

The representative dealer is active only in the repo market, where it chooses the quantity $D^D_R \geq 0$ of bonds to lend. For each unit lent, the dealer earns the specialness spread χ, but incurs a convex cost of sourcing securities. This reduced-form cost captures balance-sheet charges, search frictions, and operational costs associated with locating and delivering scarce collateral.[6] The dealer solves

$$\max_{D^D_R \geq 0} \quad D^D_R \chi - \frac{\gamma}{2}(D^D_R)^2,$$

where $\gamma > 0$ governs the marginal cost of supplying bonds. The first-order condition yields the dealer's optimal supply, which is upward-sloping in specialness:

$$D^{D*}_R = \frac{\chi}{\gamma}. \tag{6}$$

[6]These dealers can be thought of as custodian banks that lend the holdings of preferred habitat investors, retaining any profit from securities lending.

2.4 The Central Bank

The central bank lends securities through a facility associated with a legacy portfolio of asset purchases $\bar{Q} > 0$. We model central-bank lending as a reduced-form tradeoff between alleviating repo-market distortions and limiting balance-sheet usage to avoid crowding out private intermediation, consistent with the ECB's stance that "the aim of securities lending is to support bond and repo market liquidity without unduly curtailing normal repo market activity."[7]

Specifically, the central bank chooses D_R^{CB} to solve

$$\min_{D_R^{CB} \in [0, \bar{Q}]} \frac{\mu}{2}\chi^2 + \frac{\lambda}{2}(D_R^{CB})^2, \tag{7}$$

where μ is the specialness-aversion parameter, governing how much the central bank dislikes repo rates diverging from the policy rate, and $\lambda > 0$ is the footprint-cost parameter, capturing operational, institutional, or political costs of lending and crowding out private securities lenders. We solve this problem explicitly after deriving the equilibrium condition.

While the model features the central bank as the sole public-sector lender of securities, empirically we estimate separate elasticities for the two entities that constitute the public sector: the central bank and debt management office (DMO). Like the central bank, DMOs operate securities lending facilities—to manage liquidity, finance short-term budget needs, and mitigate scarcity in their bonds (Ampudia et al., 2025). The two entities may weigh the costs—crowding out private lenders—and the benefit—limiting specialness—differently, and the DMOs may additionally value the financing costs saved by collecting lending fees. Any such objective, however, that combines a quadratic footprint cost with specialness aversion or fee revenue implies lending supply that is linear in specialness. Aggregate public sector lending is therefore equivalent to that of a single lender with composite parameters (μ, λ), and we model the public sector as a single entity.

2.5 The Cash Market Equilibrium

We next characterize cash-market clearing. We first build intuition by assuming that r_T is exogenous, then present the equilibrium with endogenous r_T in Section 2.7. Let S_T denote the exogenous physical supply of the bond available to the private sector, net of any central-bank holdings. The cash market clears when preferred-habitat investor holdings

[7] See https://www.ecb.europa.eu/mopo/implement/app/lending/html/index.en.html.

plus hedge-fund holdings equal net supply. When the basis trade is profitable, hedge funds accommodate the excess demand of preferred habitat investors by shorting the bond. In that case, their equilibrium position is pinned down by cash-market clearing: $D_H^* = S_T - D_{PH}^*$.

In the empirically relevant case in which preferred habitat investors demand more bonds than are physically available, $D_{PH}^* > S_T$ and hence $D_H^* < 0$: hedge funds supply bonds to the cash market through short sales. By contrast, if the basis trade is not profitable, hedge funds are absent from the cash market and preferred habitat investors hold the available supply.

Lemma 3 (Cash-market equilibrium)**.** *Suppose the hedge fund's equity constraint does not bind.*

(i) Profitable basis trade: *If $x > \chi + mk$, the basis trade is profitable and the hedge fund's supply of bonds to the cash market is perfectly elastic. The hedge fund earns the same per-unit riskless profit $x - \chi - mk$ on every bond, so its supply curve in the cash market is horizontal. The hedge fund's equilibrium bond supply is determined entirely by preferred habitat investors' excess demand relative to the physical supply S_T:*

$$|D_H^*| = D_{PH}^* - S_T = \Delta_S + \eta(r_T - r_s) > 0.$$

(ii) Unprofitable basis trade: *If $x \leq \chi + mk$, the hedge fund does not supply bonds to the cash market: $D_H = 0$, regardless of the prevailing price. Preferred habitat investors can only hold the physical supply: $D_{PH} = S_T$.*

Lemma 3 highlights the first step of the narrative the model is designed to capture. Conditional on a profitable basis trade and no capital limits, the hedge fund's cash-market position is pinned down by the imbalance created by preferred habitat investors' demand for the physical bond, which the hedge fund accommodates in full.

2.6 The Repo Market Equilibrium

On the repo market, hedge funds borrow bonds to support their short cash positions. Dealers and the central bank supply those bonds through securities lending. Repo-market clearing requires that hedge fund bond borrowing equal total securities-lending supply:

$$|D_H^*| = D_R^D + D_R^{CB}. \tag{8}$$

Substituting (6) into (8) and assuming that $|D_H^*| > D_R^{D*}$ yields equilibrium specialness as a function of hedge fund borrowing, central-bank lending and repo dealer's marginal costs:

$$\chi = \gamma \left(|D_H^*| - D_R^{CB} \right). \tag{9}$$

Specialness equals the dealer's marginal cost evaluated at the residual demand that the central bank does not absorb. The central bank internalizes this relationship when choosing how much to lend. Substituting (9) into (7), the interior solution is

$$D_R^{CB*} = \frac{\mu\gamma^2}{\mu\gamma^2 + \lambda} |D_H^*|. \tag{10}$$

The central bank absorbs a share $\mu\gamma^2/(\mu\gamma^2 + \lambda)$ of hedge fund borrowing. This share is increasing in μ (CB cares more about specialness), increasing in γ (dealers' intermediation is costly, so the CB compensates), and decreasing in λ (securities lending affects CB footprint). The interior solution is valid whenever $D_R^{CB*} \leq \bar{Q}$; otherwise the portfolio constraint binds and $D_R^{CB} = \bar{Q}$.

Substituting (10) into (9), specialness can be expressed as a function of hedge fund demand:

$$\chi = \gamma |D_H| \left(1 - \frac{\mu\gamma^2}{\mu\gamma^2 + \lambda} \right) = \frac{\gamma\lambda}{\mu\gamma^2 + \lambda} |D_H| = \phi |D_H|, \tag{11}$$

where

$$\phi \equiv \frac{\gamma\lambda}{\mu\gamma^2 + \lambda} \tag{12}$$

is the repo-market friction parameter. It summarizes the marginal specialness cost per unit of hedge fund borrowing, given the equilibrium responses of the dealer and central bank.

We can now state the central implication of the model. When the basis trade is profitable, hedge fund repo bond borrowing is determined by equilibrium in the cash market, as the hedge fund must borrow exactly the quantity of bonds that it shorts. Hedge fund bond borrowing on the repo market is therefore pinned down by the cash-market imbalance and is insensitive to repo specialness—that is, inelastic. Hedge fund demand elasticity in the repo market is thus bimodal, as the following proposition describes.

Proposition 1 (Repo borrowing demand)**.** *Suppose the hedge fund's equity constraint does not bind.*

(i) Profitable basis trade: *If $x > \chi + mk$, the short position D_H is determined entirely by cash-market clearing. Hedge fund repo borrowing, $|D_H|$, is therefore inelastic to specialness.*

(ii) Unprofitable basis trade: *If $x \leq \chi + mk$, the hedge fund does not trade the basis or borrow bonds in the repo market: $|D_H| = 0$.*

Proposition 1 characterizes the fund's demand conditional on cash-market clearing; it does not say that the fund is insensitive to prices. On the contrary: holding the cash price fixed, the fund's demand is perfectly elastic with respect to the arbitrage spread—its partial price elasticities are unbounded. The zero elasticity of Proposition 1 is what remains of that unbounded willingness to trade once the quantity is pinned by cash-market clearing. Section 2.7 develops this distinction between partial and total elasticities with endogenous bond prices, and Appendix A.4 formalizes it.

Finally, we characterize repo supply. Combining the dealer and central-bank problems yields simple upward-sloping supply schedules.

Proposition 2 (Repo-market supply)**.** *In the repo market, securities-lending supply is upward-sloping in specialness.*

(i) Dealer supply is $D_R^D = \frac{1}{\gamma}\chi$. A lower dealer cost parameter γ makes private supply more elastic.

(ii) Central-bank supply is $D_R^{CB} = \frac{\mu\gamma}{\lambda}\chi$ at the interior solution. Higher specialness aversion μ, lower footprint cost λ, or higher dealer cost γ makes public supply more elastic.

(iii) Aggregate repo supply is

$$D_R^D + D_R^{CB} = \left(\frac{1}{\gamma} + \frac{\mu\gamma}{\lambda}\right)\chi = \frac{\mu\gamma^2 + \lambda}{\gamma\lambda}\,\chi. \tag{13}$$

Aggregate supply is upward-sloping in specialness. Its slope is decreasing in λ (costlier central-bank lending reduces supply elasticity) and increasing in μ (a more specialness-averse central bank lends more at every level of χ).

Equilibrium specialness is determined by the intersection of hedge fund demand (Proposition 1) at $|D_H|$ with the aggregate upward-sloping supply schedule (13).

Taken together, Propositions 1 and 2 deliver the mechanism that underlies our empirical analysis. Preferred habitat investors' demand makes the cash bond rich relative to futures. Hedge funds respond by shorting the bond and going long futures whenever the basis exceeds repo and margin costs, and their supply in the cash bond market is perfectly elastic. Because maintaining the arbitrage position requires borrowing the bond in repo, the cash-market imbalance translates directly into repo borrowing demand that is independent

of specialness—i.e., inelastic. Since the supply of securities lending from dealers and the central bank is not perfectly elastic, equilibrium specialness is positive.

In this sense, specialness does not reflect collateral scarcity in a mechanical sense; it arises from the interaction between cash-market demand for the physical bond and the elasticity of repo-market supply. Any shift in cash-bond demand—through preferred habitat investors' price sensitivity η, target allocation $\bar{D}$, or the term premium $r_T - r_s$—or in bond supply S_T translates into changes in specialness. Naturally, specialness also depends on frictions arising from dealers' balance-sheet charges, search frictions, and operational costs, as well as on the design of the central bank's securities lending facilities.

2.7 Endogenous Basis

The benchmark results are derived under the assumptions that the basis x and bond return r_T are exogenous and that the hedge fund does not internalize its own price impact on the bond and repo markets. We relax these assumptions, and show that the inelasticity of hedge funds in the repo market arises as a limiting case of the baseline model when preferred-habitat demand is sufficiently rigid. In this section, we model hedge funds to be perfectly competitive—thus behaving as price takers, yet having a price impact as a group.

The cash-market clearing condition $D_H = S_T - D_{PH}$ allows us to solve for the endogenous return, $r_T(|D_H|) = r_s + \frac{|D_H| - \Delta_S}{\eta}$, or equivalently for the endogenous basis:

$$x\big(|D_H|\big) = r_F - r_T\big(|D_H|\big) = \underbrace{r_F - r_s + \frac{\Delta_S}{\eta}}_{\equiv\ \bar{x}} - \frac{|D_H|}{\eta}. \tag{14}$$

We denote by $\bar{x}$ the natural basis, i.e., the basis that would prevail if hedge funds were absent:

$$\bar{x} \equiv r_F - r_s + \frac{\Delta_S}{\eta}. \tag{15}$$

It decomposes into the futures premium $r_F - r_s$ and the demand premium Δ_S/η. The latter captures the price distortion from preferred habitat demand: it is large when structural excess demand is high (Δ_S large) and preferred habitat investors are inflexible (η small).

Because the hedge fund's cash-market short position and repo borrowing are the same quantity $|D_H|$, its behavior in one market depends on conditions in the other. To understand how the fund's behavior in one market depends on the other, we derive conditional schedules: the fund's position as a function of one market's price, taking the other market's clearing

condition as given. Substituting either the endogenous basis (14) or specialness (11) into the zero profit condition $x = \chi + mk$ allows us to characterize these conditional schedules.

Substituting the cash-market clearing condition (14) into the zero-profit function yields the fund's repo-market schedule as a function of χ:

$$|D_H|^{Repo}(\chi) = \eta\big(\bar{x} - \chi - mk\big).$$

Repeating this procedure with the repo-market condition (11) yields the cash-market schedule as a function of r_T:

$$|D_H|^{Cash}(x) = \frac{x - mk}{\phi} = \frac{r_F - r_T - mk}{\phi}.$$

The slopes of the two schedules are

$$\text{for the repo market: } \frac{\partial |D_H|^{Repo}}{\partial \chi} = -\eta, \qquad \text{for the cash market: } \frac{\partial |D_H|^{Cash}}{\partial x} = \frac{1}{\phi}$$

and represent total elasticities, the result of collapsing a multi-asset system into a single-asset system.

A total elasticity embeds the general-equilibrium feedback from cross-asset substitution, measuring the equilibrium effect of a change in repo prices, after all other prices have adjusted (Chaudhary et al., 2024). The distinction matters here because the two objects behave very differently: with zero profits, the fund's partial elasticities—its responses to either price holding the other fixed—are unbounded, yet the total elasticities in each market are finite, bounded by the friction of the other market. An increase in specialness is capitalized into the bond price, the spread re-widens, and what survives of the fund's unbounded willingness to trade is exactly the quantity its cash-market counterparties release: η. The total elasticity is what our empirical framework recovers. Section 5 uses the mapping between the estimated elasticities and the model's conditional schedules to interpret hedge funds' behavior across the two markets. Appendix A.4 provides a formal derivation, embedding the model in a multi-asset demand system.

These schedules reveal a cross-market pattern that follows from the model's structure. Because a single quantity $|D_H|$ links two markets, the fund's conditional schedule slope in each market is determined by the friction it internalizes from the *other* market. The intuition is that when the funds expand in one market, the cost of doing so is governed by how much the other market tightens. Higher η (more flexible preferred habitat investors) makes funds more responsive to repo conditions, because the cash-market price impact is smaller and

the fund can adjust more freely. A less frictional repo market (lower ϕ) makes the fund more responsive to cash-market price changes, because the cost of adjusting repo borrowing is smaller.

Equilibrium quantities follow from the zero profit condition, where $\beta \equiv \frac{1}{\eta} + \phi$ captures the two equilibrium price responses that pin the competitive quantity: $\frac{1}{\eta}$ on the cash market and ϕ on the repo market. The following propositions summarize the model when the basis is endogenous.

Proposition 3 (Equilibrium quantities)**.** *The basis trade is profitable if and only if $\bar{x} > mk$, where $\bar{x} = r_F - r_s + \Delta_S/\eta$. Conditional on profitability and under perfect competition, $|D_H| = (\bar{x} - mk)/\beta$, the basis is $x = \bar{x} - \frac{\bar{x}-mk}{\eta\beta}$, specialness is $\chi = \phi \cdot \frac{\bar{x}-mk}{\beta}$, and profits are zero: $x - \chi - mk = 0$.*

Proposition 4 (Cash-market conditional schedule)**.** *Conditional on the repo market clearing, the hedge fund's intermediation is increasing in the basis x with slope $1/\phi$, or equivalently decreasing in the bond return r_T with slope $-1/\phi$, where $\phi = \gamma\lambda/(\mu\gamma^2 + \lambda)$ is the repo-market friction parameter. The slope is governed by the repo-market parameters: a less frictional repo market (lower ϕ) makes the hedge fund more responsive to the basis.*

Proposition 5 (Repo-market conditional schedule)**.** *Conditional on the cash market clearing, the hedge fund's repo borrowing is decreasing in specialness with slope $-\eta$, where η is the cash-market depth parameter. The slope is governed by the cash-market friction: more flexible preferred habitat investors (higher η) make the hedge fund more responsive to repo conditions. Since the hedge fund's cash-market short position and repo borrowing are the same quantity, the cross-market structure follows: in each market, the hedge fund's conditional schedule slope is determined by the friction in the other market.*

The baseline result of hedge funds' inelasticity on the repo market is effectively recovered as $\eta \to 0$, when preferred habitat investors' demand becomes perfectly inelastic at $\bar{D}$. The price impact $1/\eta \to \infty$ dominates β and $|D_H| \to \Delta_S$. The cost of obtaining a bond on the repo market becomes irrelevant, as the quantity $|D_H|$ is pinned on the cash side. The hedge fund's conditional schedule slope on the repo market $\eta \to 0$: the fund becomes perfectly inelastic, borrowing a fixed quantity of bonds regardless of χ. On the cash market, the fund's conditional schedule slope $1/\phi$ is unchanged—it is governed by repo friction, not by η. However, the equilibrium quantity is now pinned down by the preferred habitat investors' excess demand, so repo-market frictions are irrelevant for the cash-market clearing.

3 The Repo Market: Data and Topology

We analyze the repo market for German central government collateral from March 2021 to June 2024. This section presents the main datasets used in our analysis—with a specific focus on the most novel source, the Securities Financing Transactions Data Store (SFTDS)—and a first look at the data's descriptive statistics.

3.1 Data

The Securities Financing Transactions Regulation (SFTR), introduced in 2016 following Financial Stability Board recommendations, mandates reporting of securities financing transactions—including repos, buy–sell backs, securities lending, and margin lending—to authorized trade repositories.[8] As a result, SFTDS provides a near-exhaustive view of repo activity involving euro-area counterparties, substantially improving upon earlier datasets such as MMSR, which covered only selected dealers and did not allow for CCP look-through.

Because each transaction is reported multiple times by different counterparties, deduplication is required. A single CCP-cleared transaction typically generates four reports when all parties are euro-area domiciled. We describe our deduplication algorithm in Appendix A.1. In brief, it links ultimate securities lenders and borrowers, tracks any CCP intermediation, and consolidates all reports corresponding to the same economic transaction. Central banks and government debt management offices are exempt from reporting, although their counterparties must still report transactions conducted with them.

From the cleaned transactions, we construct daily stock positions that account for repo maturities. The methodology is detailed in Appendix A.1. Counterparties are identified by Legal Entity Identifier (LEI), country of domiciliation, and ESA 2010 sector classification. Because this classification does not distinguish dealers from other deposit-taking monetary (MFIs) or other financial institutions (OFIs), we supplement it with Bloomberg data to identify dealers explicitly. We refer to non-dealer banks as MFIs throughout.

We thus group counterparties into eight sectors: banks (MFIs excluding central banks), hedge funds (investment funds excluding MMFs), government (primarily euro-area debt management offices), other financial institutions (OFIs), dealers, central banks, insurance corporations and pension funds (ICPFs), and money market funds (MMFs). Non-financial

[8] See Financial Stability Board, *Policy Framework for Addressing Shadow Banking Risks in Securities Lending and Repo*, and EU Regulation 2015/2365.

corporations and residual sectors are excluded, as their presence on this market is minimal. Details are provided in Appendix A.2.

To address missing reporting by central banks and government entities, we supplement SFTDS with two additional sources. For the eurosystem, we use confidential data on securities lending operations from the ECBs Market Operations Database (MOPDB). For government activity, we rely on public institutional information. In the case of the Deutsche Finanzagentur, securities lending is primarily conducted via Eurex on a limited set of retained-issuance bonds. Because neither Eurex nor the Deutsche Finanzagentur reports these transactions, but final borrowers report trades with Eurex, we attribute such lending activity to the Deutsche Finanzagentur.[9]

SFTDS reporting is restricted to transactions involving at least one EU-domiciled counterparty. This limits visibility into repo positions of non-euro area dealers. As illustrated in Figure A-2, when a euro-area dealer intermediates between a euro-area bank and a non-euro area hedge fund, both legs of the transaction are observed. By contrast, when the dealer itself is non-euro area domiciled, only the initial sourcing of the bond is reported. This asymmetry mechanically inflates the net borrowing positions of non-euro area dealers, which largely mirror the positions of their hedge fund counterparties. Consequently, non-euro area dealers appear as net borrowers rather than operating matched books.

We include all repo transactions collateralized by German Bunds. We complement the repo data with bond-level characteristics from the Centralised Securities Database (CSDB) and sectoral bond holdings from the eurosystems Securities Holdings Statistics (SHS). We compute asset swap (ASW) spreads for each bond as the difference between its yield (from MTS) and the interpolated Overnight Index Swap (OIS) rate at the same remaining maturity, with STR as the floating leg. ASWs capture cash-market expensiveness beyond interest rate and default risk and are closely related to repo-market expensiveness.

Overall, through transaction-level cleaning, harmonization, and reconstruction of stock positions, this paper provides the first comprehensive and granular view of investor participation in the repo market for German collateral.

[9]While the unreported sector may include limited CCP proprietary lending, market participants indicate that these volumes are negligible. Figure A-1 shows that the unreported sector is consistently a net lender and lends more when specialness is high.

3.2 The Topology of the Repo Market for German Collateral

Figure 1 shows daily flows and stocks of repos collateralized by German government bonds. Between 2021 and 2024, outstanding repo stocks ranged from EUR 200 to EUR 500 billion per day, equivalent to roughly one-fifth of total Bund supply. These quantities hide a significant heterogeneous participation rate to the repo market. Figure 2 decomposes repo stocks into borrowing and lending positions by sector, while Table 1 reports yearly average repo positions alongside cash bond holdings. On the borrowing side, dealers, banks, and hedge funds dominate. On the lending side, the primary providers are public sector entities—central banks and debt management offices—followed by dealers and banks.

Insurance corporations and pension funds (ICPFs) exhibit almost no direct repo activity, despite holding EUR 121–144 billion in German government bonds (around 8% of total outstanding supply). This reflects the fact that ICPFs typically delegate securities lending to agent intermediaries, whose activity is recorded under banks and dealers repo positions. In contrast to the U.S. Treasury repo market, money market funds play only a negligible role for German collateral. Banks, OFIs, and dealers hold relatively small bond inventories—EUR 3–11 billion—but repo out more than EUR 100 billion per day, consistent with delegated and proprietary lending. Overall, repo market activity is largely decoupled from bond ownership across sectors.

To assess net positions, we compute net repo stocks separately for euro-area (EA) and non-euro area (non-EA) entities, reflecting reporting constraints. Figure 3 shows that the public sector, the eurosystem and the Deutsche Finanzagentur, is the largest net lender of securities, providing around 60% of all securities, while non-euro area hedge funds are the largest net borrowers, taking in nearly 50% of securities. These hedge funds are predominantly domiciled in the Cayman Islands and Bermuda (Figure A-3).

Banks and dealers primarily act as intermediaries. Euro-area banks are marginal net lenders, consistent with agent lending, while euro-area dealers operate matched books. Non-euro area dealers appear as net borrowers (37% of total demand), reflecting reporting limitations rather than true balance sheet exposure. Their net positions closely track those of non-euro area hedge funds, indicating that hedge funds are the ultimate borrowers. With full reporting coverage, we suspect that non-euro area dealers would likely also appear as matched-book intermediaries.

Finally, Figure 4 illustrates securities flows in 2022. The euro-area public sector emerges as the de facto securities lender, with foreign dealers and hedge funds as the ultimate

borrowers. Banks and dealers mainly intermediate flows from euro-area entities to non-euro area investors, while euro-area banks lend more securities than they borrow, underscoring their role as agent lenders.

4 Equilibrium Model of the Repo Market: Econometric Framework

This paper estimates an equilibrium model of repo specialness that jointly incorporates investor demand across sectors and allows for heterogeneous responses to repo specialness and macroeconomic conditions. We conduct the empirical analysis using the methodological contribution of Chaudhary et al. (2024), the Optimal Granular Instrumental Variable (O-GIV). O-GIV extends the standard GIV framework introduced by Gabaix and Koijen (2024). The GIV framework constructs instruments by extracting idiosyncratic demand shocks from large investors and uses these shocks to estimate supply and demand elasticities.

GIV, like any other instrumental variable setup, relies on two assumptions, relevance and exclusion restriction. Relevance is satisfied if the idiosyncratic shocks are economically meaningful and originate from a small number of large institutions: When investors are sufficiently large, their idiosyncratic shocks to quantities can move aggregate prices. The exclusion restriction is satisfied if the shocks used as instruments are truly idiosyncratic. In particular, the variation captured by the instrument must be orthogonal to shocks that affect all investors, i.e., common factors. Under these two conditions, Chaudhary et al. (2024) provide a framework for estimating *sector-specific* elasticities. While their application focuses on the cash Treasury market, we apply O-GIV to the repo market.

Their setting is as follows. Assume sector i's demand for borrowing securities, $\Delta q_{i,t}$, can be expressed as a log-linear function of repo specialness, Δp_t—i.e., the repo price, defined as the eurosystem deposit facility rate minus the average repo rate for German government collateral—and other demand shifters $\upsilon_{i,t}$:

$$\Delta q_{i,t} = \overline{\Delta q}_i - \zeta_i \Delta p_t + \upsilon_{i,t}. \tag{16}$$

Here, $\overline{\Delta q}_i$ denotes sector fixed effects. Equation 16 represents sector i's net demand in the repo market, and ζ_i captures the response of sector i's net securities borrowing to changes in repo prices. A large ζ_i indicates a very elastic sector, which aggressively adjusts its supplied or demanded quantity in response to price changes. By contrast, a near-zero ζ_i characterizes

a sector composed of price-inelastic investors, whose security-borrowing decisions are largely insensitive to repo pricing.

We decompose demand shifters into common factors η_t and sector-specific idiosyncratic shocks $u_{i,t}$. For example, an increase in uncertainty may induce hedge funds to expand cash-market positions in an attempt to profit from price discrepancies. To increase the profitability of these positions, hedge funds rely on leverage obtained by borrowing cash or securities in the repo market. Banks may also respond to heightened uncertainty by relying more on short-term funding, potentially taking positions that mirror those of hedge funds with varying intensity, or even adopting opposite strategies. In this setting, uncertainty is a common factor: it affects repo market activity across sectors while eliciting heterogeneous sectoral responses, captured by the sector-specific coefficient λ_i. Beyond uncertainty and other common factors η_t, sector-specific "bets" are captured by the idiosyncratic component $u_{i,t}$. This yields:

$$\Delta q_{i,t} = \overline{\Delta q}_i - \zeta_i \Delta p_t + \lambda_i \eta_t + u_{i,t}. \tag{17}$$

We relate prices and quantities through market clearing: in equilibrium, net securities borrowed must equal net securities lent. Equivalently, the market-share-weighted sum of sector flows satisfies $\sum_i S_i \Delta q_{i,t} = 0$. Imposing market clearing in Equation 17 and solving for equilibrium repo specialness gives, with S denoting size-weighted aggregation across sectors,

$$\Delta p_t = \frac{1}{\zeta_S} \left(\overline{\Delta q}_S + \lambda_S \eta_t + u_{S,t} \right), \tag{18}$$

where

$$\zeta_S = \sum_i S_i \zeta_i \tag{19}$$

is the aggregate elasticity. Equation 18 links changes in repo specialness to changes in aggregate net demand for securities in the repo market, scaled by the inverse of the aggregate elasticity.

Because prices and quantities are jointly determined in equilibrium, we use an instrumental variable approach to address simultaneity. Under the relevance and exclusion restriction conditions, the (O-)GIV framework constructs instruments by extracting investors' idiosyncratic demand shocks and using them as instruments for repo specialness faced by other sectors.

The endogenous variable is sector i's net repo position, $Q_{i,t}$. Net positions are daily outstanding stocks, computed as securities borrowed minus securities lent (EUR bn). We detrend net positions using a three-month moving average $\bar{Q}_{i,t}$, following Kubitza, Sigaux,

and Vandeweyer (2025). The resulting variable, $\Delta q_{i,t}$, is expressed as a percentage deviation from the moving average:

$$\Delta q_{i,t} = \frac{Q_{i,t} - \bar{Q}_{i,t}}{|\bar{Q}_{i,t}|}, \tag{20}$$

a transformation we apply to all other quantities, including price, as well. The sample is partitioned into euro-area and non–euro area sectors, for a total of 15 sectors.[10]

We compute sector i's daily market share as its three-month moving average net position relative to the corresponding aggregate across sectors, and use the time-series average market share S_i over the sample:

$$S_i = \frac{1}{T} \sum_{t=1}^{T} \frac{|\bar{Q}_{i,t}|}{\sum_j |\bar{Q}_{j,t}|}. \tag{21}$$

Figure 5 compares value weights to equal weights by sector. Equal weights (the inverse of the number of sectors) provide a benchmark in which all sectors participate with equal footprint, while value weights S_i capture actual market shares in terms of net borrowing positions. The figure shows that the repo market for German collateral is highly concentrated in a small number of large sectors. This granularity strengthens the relevance of O-GIV, as idiosyncratic shocks from large sectors can meaningfully affect repo prices.

Common factors are central to identification because they allow us to separate idiosyncratic shocks from aggregate shocks shared across sectors, supporting the exclusion restriction. We group controls into four macro-categories: Monetary policy quantities, repo demand quantities, repo supply quantities, and bond supply quantities. We detail which measures fall in each of these categories in the next section. We apply the same transformation as in Equation 20 to all common factors.

We extract idiosyncratic shocks to sector positions by residualizing each sector's net position $\Delta q_{i,t}$ with respect to sector fixed effects and common factors. This step is central to identification: conditioning on common factors isolates sector-specific shocks that can be used as instruments. The first-stage regression is:

$$\Delta q_{i,t} = \alpha_i + \lambda_i \eta_t + u_{i,t}, \tag{22}$$

where α_i denotes sector fixed effects, η_t is the vector of common factors (expressed as percentage deviations from three-month moving averages), and $u_{i,t}$ captures idiosyncratic shocks. We use Newey–West standard errors to account for serial correlation.

[10] Owing to their limited (direct) participation to the repo market, non-financial corporations, households, and non–euro area money market funds are excluded.

Because prices and quantities are jointly determined, we use an instrumental variable approach to identify the effect of quantities on repo specialness and address simultaneity. Under the relevance and exclusion restriction conditions, O-GIV constructs instruments by using estimated idiosyncratic shocks from other sectors as instruments for changes in repo prices.

Formally, Chaudhary et al. (2024) introduce a generalized method of moments (GMM) estimator—the Optimal-GIV estimator—to recover sector-specific elasticities $\hat{\zeta}_i$. Identification is based on the exclusion restriction and on weighting moment conditions by sector market shares to increase instrument relevance. The exclusion restriction implies:

$$\mathbb{E}[u_{i,t}u_{-i,t}] = 0 \qquad \text{for any } i \neq -i, \tag{23}$$

i.e., idiosyncratic shocks from sector i are orthogonal to idiosyncratic shocks from any other sector $-i$. The Optimal-GIV estimator $\hat{\zeta}$ solves the moment condition:

$$\widehat{\mathbb{E}}\Big[\hat{u}_{i,t}(\hat{\zeta}_i)\,\hat{u}_{S(-i),t}(\hat{\zeta})\Big] = 0, \tag{24}$$

where $\hat{u}_{S(-i),t}(\hat{\zeta}) = \sum_{j\neq i} S_j \hat{u}_{j,t}(\hat{\zeta}_j)$ denotes the size-weighted idiosyncratic shocks of all sectors other than i. The weights S_j give more influence to sectors with larger market shares, reflecting that their shocks move prices more strongly. Accordingly, for each sector i, the estimator selects $\hat{\zeta}_i$ so that the sample covariance between $\hat{u}_{i,t}(\hat{\zeta}_i)$ and $\hat{u}_{S(-i),t}(\hat{\zeta})$ is zero.

Substituting $\hat{u}_{i,t}(\hat{\zeta}_i) = \Delta q_{i,t} + \hat{\zeta}_i \Delta p_t$ into Equation 24 yields:[11]

$$\hat{\zeta}_i = -\frac{\widehat{\mathbb{E}}\Big[\,\Delta q_{i,t}\,\hat{u}_{S(-i),t}(\hat{\zeta})\,\Big]}{\widehat{\mathbb{E}}\Big[\,\Delta p_t\,\hat{u}_{S(-i),t}(\hat{\zeta})\,\Big]}, \tag{25}$$

which can be resolved recursively (Appendix A.3). This expression corresponds to the inverse of the canonical IV estimator: the numerator is the covariance between the endogenous variable and the instrument (constructed from first-stage residuals), and the denominator is the covariance between the price change and the same instrument.

[11] For simplicity, we abstract from common factors in this derivation.

5 The Repo Market Elasticity

In this section, we present our main results from the O-GIV estimation detailed in the previous section. We include several robustness checks and conclude with the quantification of how much each sector and factor contribute to the evolution of specialness.

5.1 Loadings on Common Factors Across Sectors

We begin by estimating the relationship between sector-level net securities borrowing positions and a broad set of common factors, controlling for sector fixed effects, as specified in Equation 22. The dependent variable is the percentage deviation of sector is net securities borrowing from its three-month moving average. The same transformation is applied to the common factors, which we group into four macro-categories, monetary policy quantities, repo demand quantities, repo supply quantities, and bond supply quantities:

- *Monetary policy quantities:* These variables capture interest rate expectations, surprises, and duration extraction. They include: The 1-month OIS 1M; the 10Y-1Y,1Y-3M, and 3M–1M OIS spread; and the 1M OIS–deposit facility rate spread.
- *Repo supply quantities:* These varaibles proxy for capital provision by repo-dealers. We include the three-month cross-currency basis swap, which links repo supply to FX markets, and Duffie (2023)'s measure of liquidity-versus-volatility deviations for the euro-area bond market (see Besugo, Cinquin, Nguyen, Paoli, and Szydlo, 2025) as a proxy for intermediation constraints.
- *Repo demand quantities:* These measures capture repo-demand driven by arbitrage opportunities. We include the skewness of Euribor futures options-implied density; the Merrill Lynch Swaption Option Volatility Estimate (MOVE); the 3M Bund volatility premium; deviations of cash Bund prices from a Nelson-Siegel curve; and generic 10-year Bund futures open interest. We also include the Bund–KfW spread, a measure of European bank's illiquidity, and the Bund–French spread.
- *Bond supply quantities:* the eurosystem balance sheet, Deutsche Finanzagentur holdings, and the free float of German government bonds, computed as issued amounts minus holdings by the Deutsche Finanzagentur and the eurosystem.[12]

[12]We obtain data for the common factors from Bloomberg, MTS, LSEG, and the ECB Statistical Data Warehouse. We compute a proxy for the Deutsche Finanzagentur's holdings from public information by projecting retained-at-issuance amounts to maturity, implicitly assuming no secondary-market sales.

Results for the largest sectors are reported in Table 2. These sectors—euro-area government, banks, and central bank; non-euro area hedge funds, dealers, and other financial institutions (OFIs)—each account for more than 5% of average outstanding repo positions. Together, they represent 83% (91%) of aggregate repo supply (demand). Results for the remaining sectors are reported in Table A-2.

We classify a sector as a supplier (demander) of repo market liquidity if it is, on average over the sample period, a net lender (borrower) of securities. Several patterns emerge. Non-euro area hedge funds load positively on all proxies for investment opportunities and pricing dislocations, including MOVE, the Bund volatility premium, and deviations from the Nelson-Siegel yield curve. Monetary policy variables, such as the yield curve slope, also strongly explain hedge fund borrowing behavior: hedge funds increase net securities borrowing when these measures rise. Bond supply variables are likewise important. Hedge fund borrowing increases with expansions in the public sector balance sheet and declines when free float rises.

Non-euro area dealers borrowing positions are more strongly correlated with monetary policy variables such as OIS 1M. They lend more when collateral availability—proxied by free float—is higher. The public sector, comprising the euro-area central bank and government, supplies more securities when repo market conditions deteriorate, as indicated by larger pricing deviations from the Nelson-Siegel model or elevated interest rate volatility. Higher bond holdings by these institutions are associated with greater net lending in the repo market, generating a negative correlation between their bond holdings and repo positions. Free float loads positively for both the government and the central bank, indicating that public sector lending declines when collateral becomes more readily available. Hedge funds, by contrast, display the opposite pattern.

Euro-area banks load significantly on the cross-currency basis swap and the Kreditanstalt für Wiederaufbau (KfW)-Bund yield spread, which proxy for overall euro-area banking liquidity. Overall, common factors explain a substantial share of the variation in sectoral net borrowing positions, with peak R^2 values reaching 45% for hedge funds and 42% for government.

Figure A-4 plots the daily sum of residuals across sectors from the first-stage regressions. These residuals represent idiosyncratic demand and supply shocks, which form the basis for the estimation of sector-specific elasticities in the next subsection. At this stage, net securities borrowing positions have been purged of common influences and sector fixed effects, yielding the inputs for the second-stage elasticity estimation.

5.2 Specialness Elasticity

We now turn to estimates of repo market specialness elasticities. We first present the elasticity of aggregate repo demand and then examine sector-specific elasticities. The bottow row of Table 3 reports the aggregate Bund repo market elasticity, ζ_S. We estimate an aggregate elasticity of 1.2, implying that a 1% demand shock increases repo specialness by approximately 0.83%, the inverse of the aggregate elasticity. In our sample, this corresponds to a 0.21 basis point increase in specialness for a EUR 110 million increase in net repo borrowing. This estimate is comparable in magnitude to elasticity estimates for the U.S. Treasury market (close to 1; see Kaminska and Zinna, 2020; Jansen et al., 2024; Chaudhary et al., 2024) but substantially smaller than those found for the U.S. equity market (around 5; Gabaix and Koijen, 2021), as illustrated in Figure 6. These comparisons are indicative rather than exact: our elasticity is defined over a rate spread and over positions that net to zero across the market, so its units differ from those of holdings-based estimates for cash markets.[13] Importantly, the sector and factor decompositions of Sections 5.3–5.4 are invariant to this scaling, so the paper's central results do not depend on it.

The remainder of Table 3 reports sector-level elasticities ζ_i and corresponding market multipliers $M = 1/\zeta_i$, along with average market shares S_i. As before, sectors are classified as suppliers or demanders based on average net lending or borrowing over the sample period. The table shows the stark contrast between demand and supply elasticity.

Non-euro area hedge funds and dealers (the demand side) are by far the most price-inelastic participants, with elasticities of 0.11 and 0.05, respectively. As discussed earlier, dealers plausibly act on behalf of hedge funds. A 1% increase in net securities borrowing by non-euro area hedge funds (dealers) increases repo specialness by approximately 9% (20%). In economic terms, an additional EUR 380 million (EUR 330 million) in repo borrowing by hedge funds (dealers) raises specialness by about 2.34 (5.2) basis points.[14] Alternatively, a 1% rise in specialness cuts hedge funds' and dealers' repo borrowing by only a tenth or a twentieth of a percent.

The estimated elasticity of 0.11 is economically small but statistically different from zero, consistent with the model with endogenous basis in Section 2.7 rather than with the baseline's knife-edge prediction of exactly zero elasticity. When the basis is determined in equi-

[13] Existing estimates for Treasury, corporate-bond, and equity markets relate quantities to percentage changes in asset prices and scale holdings by assets in positive net supply. Repo positions net to zero, so no direct analogue of shares outstanding exists, and the unavailability of investor-level bond holdings prevents scaling repo positions by holdings of the underlying collateral.

[14] In our sample, average specialness is 26 basis points, while average net positions are EUR 38 billion for non-euro area hedge funds and EUR 33 billion for non-euro area dealers.

librium, the hedge fund's repo-market elasticity is proportional to the cash-market depth parameter η; through the lens of the model, a small but positive repo elasticity thus indicates that preferred-habitat investors retain some, albeit limited, price sensitivity—in line with the cash-market demand-system literature (Koijen et al., 2021).

The estimate of 0.11 is the hedge funds' *total* elasticity with respect to the repo price: it measures the equilibrium response of their repo quantities to a change in specialness, after all other prices have adjusted (Chaudhary et al., 2024). A low total elasticity, however, does not imply that hedge funds' repo positions are insensitive to cash prices. In the model, in fact, the opposite holds, as Appendix A.4 formalizes. Because hedge funds arbitrage the spread between the basis and specialness, their repo borrowing responds to the repo price and to the cash bond price with partial elasticities of equal magnitude: the two prices enter the fund's problem only through the arbitrage spread $x - \chi - mk$ (equation (3) and Lemma 1), so, holding one price fixed, a change in the other moves the spread, and thus repo borrowing, one-for-one.

As Section 2.7 anticipates, an increase in specialness triggers two large, opposing responses in hedge funds' repo borrowing. Holding the bond price fixed, the arbitrage spread compresses and funds cut their borrowing sharply—the direct, own-price response. In equilibrium, however, the rise in specialness is fully reflected in an increase in the bond price through the zero-profit condition (Proposition 3; cf. also Duffie, 1996), and the spread, re-widened by the price increase, pulls borrowing back up through the equally large cross-price response—the indirect response (for this decomposition of total elasticities into partial responses and price passthrough, see Chaudhary, Fu, and Li, 2023). The direct and indirect responses offset almost fully: on net, a rise in specialness lowers hedge funds' repo borrowing only at rate η, the cash-market depth parameter of Proposition 5. The residual is exactly η because repo borrowing must equal preferred-habitat investors' excess demand (Lemma 3): when the bond price rises one-for-one with specialness, preferred-habitat investors release bonds at rate η per unit of price increase, and the short position—and with it repo borrowing—contracts by exactly η. Hedge funds' measured total repo elasticity is thus not their own price sensitivity, which is effectively unbounded, but that of their counterparties in the cash market. This resolves the apparent tension: precisely because hedge funds are perfectly elastic in the cash market—absorbing whatever imbalance preferred-habitat investors leave—their demand in the repo market is just as inelastic as preferred-habitat investors' demand is in the cash market.

On the supply side, the public sector—the German debt management office and the Eurosystem—provides highly elastic securities lending, but the two institutions do so in markedly different

ways, by institutional design. The debt management office, with an elasticity of 0.64, lends at every level of specialness: its facility grants broad access, requiring only clearing membership, and is pried at terms competitive with the market, so investors tap it routinely and its supply responds more moderately to price. The Eurosystem, with an elasticity of 1.28, is the more responsive lender—by design rather than by appetite. Its facility charges the higher of 20 basis points and the prevailing market rate, so at low specialness the facility is dominated by the market: at a specialness of 10 basis points, borrowing from the Eurosystem costs twice the market rate. The facility becomes competitive only once specialness exceeds the 20 basis point threshold, and its lending is therefore concentrated precisely where specialness is high. The backstop pricing generates the high measured elasticity: the Eurosystem effectively supplies securities when specialness rises, and only then.

Intermediaries such as euro-area banks and dealers are highly elastic, implying limited price impact. In aggregate, the public sector and euro-area intermediaries account for 88% of repo supply, while non-euro area hedge funds and dealers account for 82% of repo demand. These six sectors alone contribute more than 85% of the aggregate elasticity. Of the overall elasticity of 1.2, approximately 13% originates from demanders and 87% from suppliers, indicating that repo demand is highly inelastic, while supply is highly elastic.

Figure 7 visualizes these findings by comparing market multipliers across sectors. Demand-side participants—non-euro area hedge funds and dealers—exhibit the largest multipliers, whereas supply-side participants—the public sector and euro-area banks—exhibit small multipliers.

Table 4 compares sector-level elasticities in the repo and cash bond markets. Cash market elasticities are taken from Koijen et al. (2021), which provides the most recent sector-level estimates for the euro area.[15] We uncover a clear negative relationship between elasticities in the two markets, illustrated in Figure 13.

In the cash market, the euro-area central bank exhibits zero price elasticity, reflecting bond purchases that are independent of prices and driven solely by monetary policy objectives. In contrast, the central bank is elastic in the repo market, limiting its ability to mitigate specialness during stress episodes. While securities lending facilities are designed not to interfere with the monetary policy stance, they could nevertheless play a stabilizing role by easing access to scarce collateral during systemic specialness episodes. The government is

[15]The sectoral mapping is imperfect because Koijen et al. (2021) relies on ESA 2010 classifications and treats the foreign sector as a residual. Our sectoral breakdown is more granular and focuses explicitly on German government bonds. Sample periods also differ, although both cover periods of eurosystem balance sheet expansion.

inelastic in the cash market and only weakly elastic in the repo market, reflecting interventions that are largely insensitive to prices and aimed at minimizing funding costs.

Finally, the foreign sector—dominated by non-euro area hedge funds and their dealer intermediaries—displays the most striking asymmetry: hedge funds are highly elastic in the cash market but extremely inelastic in the repo market. This behavior is economically rational. Profits from cash market strategies dominate repo funding costs, leading hedge funds to maintain repo borrowing even when specialness rises sharply. Repo borrowing thus serves primarily to finance cash market strategies, reinforcing and amplifying specialness episodes.

This asymmetry is a direct prediction of the model in Section 2.7. Proposition 5 shows that the hedge fund's conditional repo-market schedule slope is $-\eta$, governed by the cash-market depth parameter, while Proposition 4 shows that the cash-market schedule slope is ϕ, governed by the repo-market friction parameter. When hedge funds are intermediating, they face an elastic supply on the repo market and inelastic demand on the cash market, and they, in turn, provide elastic supply on the cash market and inelastic demand on the repo market. The empirical cross-market elasticity pattern in Figure 13 aligns with our proposed economic mechanism.

With these elasticity estimates in hand, Section 5.3 assesses the contribution of different sectors to repo specialness between 2021 and 2024. Before doing so, we present robustness checks for our baseline estimates.

5.3 Sector Decomposition

Using sector-specific elasticities, we decompose deviations in repo specialness into contributions from individual sectors. Formally, the contribution of sector i to specialness is given by:

$$\Delta\hat{p}_{i,t} = \frac{1}{\hat{\zeta}_S} S_i \left(\widehat{\overline{\Delta q}}_i + \hat{\lambda}_i \eta_t + \hat{u}_{i,t} \right). \tag{26}$$

We compute average daily contributions at the quarterly frequency and group sectors into four categories: (i) non-euro area hedge funds and dealers, (ii) the public sector (euro-area governments and central banks), (iii) intermediaries (euro-area banks and dealers), and (iv) remaining sectors. Results are shown in Panel A of Figure 8.

Natural securities borrowers—non-euro area hedge funds and dealers—are the primary drivers of repo specialness dynamics. In the first part of the sample, their net borrowing contributes almost exclusively to positive deviations in specialness. In the second part,

their contribution turns negative as they sharply reduce net borrowing and move toward balanced positions (Figure 3). By contrast, the public sector adopts the opposite stance, acting as a net lender over most of the sample. Through its securities lending facilities, it mitigates the impact of hedge fund borrowing on specialness, particularly in the first half of the sample. This offsetting role is illustrated in Panel B of Figure 8, which cumulates sectoral contributions to specialness over time, expressed in basis points.[16]

At the peak of specialness in 2022Q3, the government (central bank) reduced specialness by approximately 5 (10) basis points relative to its three-month average. Absent public sector lending facilities, specialness would have been roughly 15 basis points higher. At the same time, non-euro area hedge funds and their dealer intermediaries cumulatively increased specialness by more than 30 basis points relative to its three-month average. While this exercise does not constitute a formal counterfactual,[17] it provides intuition for how sectoral borrowing and lending translate into cumulative changes in the *level*, rather than the *change*, of specialness.

Overall, market multipliers indicate that changes in hedge fund participation exert the largest price impact. Consistent with this, the sector decomposition shows that hedge funds were the primary drivers of the 2022 specialness episode. This naturally raises the question of whether the subsequent easing reflects changes in hedge fund elasticities or instead changes in their repo positions.

The sample period spans two distinct monetary policy regimes. From January 2021 to November 2022, the euro area experienced prolonged balance sheet expansion, with elevated excess liquidity reflecting pandemic-era interventions such as the PEPP, PSPP, and large-scale TLTROs. This regime shifted in mid-2022, when the ECB initiated a tightening cycle, raised policy rates in July, and returned to positive rates in September. In October 2022, the Governing Council recalibrated TLTRO pricing to better align central bank operations with tighter monetary conditions. This triggered substantial early repayments, with nearly EUR 300 billion—about 14% of the outstanding EUR 2 trillion—repaid in November 2022, marking the effective start of balance sheet contraction.

While our baseline specification controls for common factors such as the central bank balance sheet and the monetary policy stance, elasticities are estimated as averages over the full

[16] We back out basis point contributions from the three-month percentage deviation transformation in Equation 20, applied to prices. Specifically, we scale $\Delta\hat{p}_{i,t}$ by the observed three-month average $\overline{p_t}$, divide by the number of sectors, and cumulate over time. Figure A-5 shows the strong positive correlation between predicted specialness $\Delta\hat{p}_t$ and observed deviations Δp_t, indicating good model fit.

[17] A full counterfactual would require recursively recomputing three-month average specialness in the absence of specific sectors.

sample and evaluated at period-average sector sizes. Consequently, declining specialness in the second half of the sample could reflect either changes in hedge fund elasticities or changes in their net repo positions. To distinguish between these mechanisms, we re-estimate the O-GIV separately for the balance sheet expansion period (January 2021–November 2022) and the contraction period (December 2022–June 2024).

Results, shown in Figure 10, indicate that hedge fund elasticities remain virtually unchanged across regimes. This robust finding implies that the decline in specialness is driven primarily by reduced hedge fund borrowing—or increased lending—rather than by changes in price sensitivity. With the exception of the central bank, repo supply becomes slightly more elastic during quantitative tightening, while demand becomes more inelastic. The market-wide elasticity increases modestly from 1.2 to 1.4. The broad stability in estimates is consistent with the model. Because the hedge fund's repo-market elasticity is governed by the structural parameter η—the price sensitivity of preferred-habitat investors—rather than by the level of specialness or the monetary policy regime, it should remain approximately constant across subsamples, as we observe.

5.4 Factor Decomposition

Equation 26 also allows us to disentangle the contributions of common factors and sector-specific shocks to repo specialness. The contribution of the k-th common factor is:

$$\Delta\hat{p}_{\eta(k),t} = \frac{1}{\hat{\zeta}_S} S_i \hat{\lambda}_i(k) \eta_t(k). \tag{27}$$

Importantly, this decomposition does not measure the direct effect of common factors on specialness. Instead, it captures how common factors affect sector-level repo borrowing and how these quantity responses transmit to prices. Contributions are computed as the product of sector-level first-stage coefficients $\hat{\lambda}_i(k)$ (Tables 2 and A-2), factor realizations $\eta_t(k)$, and sector market shares.

Figure 9 presents the factor decomposition aggregated to the quarterly frequency, grouping factors into five categories: bond supply quantities, window-dressing dummies, repo supply quantities, monetary policy variables, and repo demand variables. Figure A-6 reports contributions of individual factors.

Bond supply quantities. Bond supply is decomposed into eurosystem holdings, Deutsche Finanzagentur holdings, and free float. Up to June 2022, these factors jointly contributed to

a decline in specialness (Figure 9), driven primarily by central bank holdings (Figure A-6). While balance sheet expansion is often associated with higher specialness, our decomposition captures an alternative channel: larger central bank holdings are associated with increased securities lending through central bank facilities, mitigating specialness. As balance sheet expansion ended in mid-2022 and contraction began, this mechanism reversed, and bond supply factors contributed positively to specialness. Overall, bond supply effects tended to reduce specialness over the sample, consistent with increased collateral circulation via public sector facilities.

Window-dressing. Month-end and quarter-end effects are captured using countdown dummies for the final 18 days of each month. Aggregated to the quarterly frequency, window-dressing contributes positively to specialness (Figure 9). Figure 11 examines sector-specific contributions during the three weeks preceding quarter-end, distinguishing periods before and after March 2022. Prior to March 2022, euro-area intermediaries benefited from leverage ratio exemptions that reduced balance sheet costs. After the exemption expired, intermediaries contributions to specialness became more pronounced, a pattern confirmed in Figure 11. Figure A-6 further shows that euro-area banks and dealers are the primary drivers of specialness during post-exemption quarter-ends. Overall, these results highlight the role of balance sheet constraints and demonstrate that quarter-end effects reflect both common and sector-specific dynamics.

Repo supply quantities. Among repo supply variables, the one-year EUR/USD cross-currency basis swap accounts for the largest contribution (Figure A-6). The basis turned from negative to positive in late 2023, possibly reflecting the end of the euro-area tightening cycle. This shift indicates strong demand for euro funding and translated into a large negative contribution to specialness between 2023Q4 and 2024Q1. As shown in Table 2, these effects are driven by increased securities lending (or reduced borrowing) by non-euro area hedge funds and dealers, as well as by the euro-area government.

Monetary policy quantities. Contributions in this category are dominated by the spread between the one-year and three-month OIS rates, which proxies short-term interest rate expectations. The spread widened early in the sample, remained elevated through early 2023, and then narrowed as expectations shifted toward the end of the tightening cycle. In late 2023, the spread turned negative, consistent with anticipated rate cuts in 2024. Table A-2 shows that this factor loads significantly on non-euro area banks lending activity. Correspondingly, Figure A-6 shows negative contributions to specialness in the first half of the sample and positive contributions in the second half.

Repo demand quantities. This group proxies broader financial conditions affecting repo borrowing across sectors, including volatility (MOVE), hedging activity (Bund futures open interest), and flight-to-quality and flight-to-safety episodes. While common demand factors explain only a limited share of specialness at the quarterly frequency (Figure 9), this reflects offsetting effects across individual factors. The French-German government bond spread, capturing flight to safety, contributes positively to specialness and loads strongly on hedge fund borrowing. By contrast, flight-to-quality effects load on central bank lending, contributing negatively to specialness. Volatility measures load negatively on euro-area public sector and intermediaries, as these sectors increase lending during volatile periods. Finally, pricing errors from the Nelson-Siegel model correlate negatively with securities borrowing by both the public sector and non-euro area hedge funds, implying that higher pricing errors are associated with increased lending and lower specialness.

5.5 Robustness

To assess the robustness of our price elasticity estimates, we report in this section how our main results vary as we alter sample definitions, first-stage controls, and the transformation of the main variables of interest. Our results are mostly unchanged, pointing to a stable ranking of elastic and inelastic investors.

To account for unobserved common drivers of investors' quantities, we augment the first stage with latent common factors extracted from the residuals of the baseline regression. The O-GIV identification requires idiosyncratic shocks to be orthogonal across sectors, so any remaining common variation would violate the exclusion restriction. We proceed in three steps: we estimate the baseline first-stage regression in Equation 22, including sector fixed effects and observed common factors; we extract the first three principal components of the sector-level residuals; and we re-estimate the first stage including the three principal components alongside the original controls and sector fixed effects. First-stage results for the largest sectors are reported in Table 5. Relative to the baseline estimates in Table 2, the principal components increase explanatory power on both the supply and demand sides, especially for euro-area banks. We report the second-stage results in Table 6 and Figure 14, together with the baseline elasticities. Sector-level elasticity estimates are broadly unchanged, with the only notable difference being euro-area banks, whose elasticity falls closer to that of the other suppliers. The resulting aggregate elasticity declines from 1.20 to 0.97.

For our estimates to qualify as total—rather than partial—elasticities for the repo market, the explanatory variables in the first stage should account for the general-equilibrium clear-

ing of all other markets, including that for cash bonds (Chaudhary et al., 2024). While our baseline includes a wide set of interest-rate controls, it omits measures of the German yield curve. We repeat the analysis including changes in the 10-year point of the German yield curve in the first stage. The resulting elasticity estimates, reported in Figure A-8 of the Appendix, are nearly identical to the baseline.[18]

Our elasticity estimates should be interpreted as applying to the aggregate market for repo positions, even though repo contracts are collateralized by heterogeneous Bunds. We estimate the model separately for sample splits based on a bond's time to maturity and scarcity. In both cases, we rank each day's transactions by the relevant bond characteristic, form the top and bottom halves by volume, compute average quantities and prices within each half, and apply our baseline transformations. We report the results in Figure A-9 of the Appendix. Hedge funds' elasticity for longer time-to-maturity bonds is significantly lower than for shorter ones, consistent with hedge funds' activity concentrating in futures–bond basis trades on the 10-year contract—by far the most liquid European bond future—which lock in their repo demand for the underlying long-term bonds. Repos backed by short-term bonds, by contrast, are rarely tied to basis positions, leaving hedge funds freer to respond to repo costs. As for scarcity, central-bank securities lending is more elastic for high-scarcity bonds, consistent with the central bank stepping in more aggressively to tamp down high specialness.

The variables in our main specifications are transformed into deviations from their three-month moving average, following Kubitza et al. (2025). This transformation represents a middle ground between the two extremes of estimating the model in daily changes—deviations from a one-day moving average—and in levels—deviations from an infinitely long moving average. The choice trades off statistical well-behavedness against noise. To confirm that our results are robust to this choice, we replicate the analysis using deviations from one- and six-month moving averages, reporting the results in Figure A-10 of the Appendix. The spread between demand and supply elasticities widens with the length of the look-back window; even with a one-month window, however, demand is significantly more inelastic than supply. All estimates remain close to the baseline.

Finally, scaling by realized trading positions may not be fully exogenous to prices—unlike scaling by an asset's amount outstanding—since investors can adjust positions more readily. We repeat the analysis with the scaling variable lagged by four months and obtain elasticities generally in line with the baseline, reported in Figure A-11 of the Appendix.

[18] Bund yields are nearly spanned by the explanatory variables we included in the baseline's first-stage: The correlation between actual and predicted yields is 64%.

6 The Role of Hedge Funds

Previous sections estimate an equilibrium model of demand and supply in the repo market. We show that the repo market for German government collateral features an elastic supply—largely provided by the public sector—and a highly inelastic demand. This demand inelasticity is driven by non-euro area hedge funds. In the sectoral decomposition in Section 5.3, we show that these investors were responsible for the sharp increase in repo specialness in 2022. During this period, hedge funds became the largest net borrowers of securities and thus the marginal investors shaping specialness. This role of hedge funds in 2022 is consistent with evidence based on confidential Bund auction group data presented by Ampudia et al. (2025) suggesting that hedge funds were net sellers during this period in the cash Bund market.

To understand why hedge funds borrowed so aggressively, Panel A of Figure 12 documents a concurrent increase in several measures of mispricing across markets, including the cash–futures bond basis, asset swap spreads, and pricing errors from the Nelson-Siegel model.[19] The widening of these measures indicates the emergence of profitable arbitrage opportunities.

In what follows, we use the cash–futures bond basis as an illustrative example.[20] A widening of the positive cash–futures basis—shown in Panel B of Figure 12—implies that the cash price of a deliverable bond exceeds the implied bond price embedded in the corresponding futures contract. A hedge fund seeking to exploit this mispricing would sell the bond in the cash market and buy the corresponding futures contract, a strategy commonly referred to as the *basis trade*. The short cash position is financed by borrowing the bond in the repo market. Large short positions in the cash market are therefore consistent with large net borrowing positions in the repo market.

Crucially, hedge funds access the repo market as price-inelastic investors, sustaining elevated borrowing even as funding costs rise. As long as arbitrage profits exceed repo specialness—i.e., the cost of funding the trade—this behavior is economically rational. Basis-trade

[19] The observed mispricing likely reflects the interaction of several concomitant factors. The rapid pace of monetary policy tightening following the inflation surge induced abrupt repricing across the yield curve and elevated interest rate volatility, which may have limited arbitrageurs ability to enforce price parity between cash and derivatives markets. At the same time, pressures on dealers balance sheets were increasing during this period (Besugo, Nguyen, Poinelli, and Scheicher, 2025), potentially impairing their intermediation capacity, which is central to the implementation of arbitrage trades.

[20] This relative-value trade is only illustrative. We acknowledge that other trading strategies that require securities borrowing may also have contributed to hedge funds elevated repo demand.

positions are maintained as long as cash-market profits dominate repo-market costs.[21] As discussed in Section 5.2, while hedge funds are inelastic in the repo market, the existing literature finds them to be elastic in the cash market: when arbitrage opportunities narrow, both cash and repo positions unwind.

Consistent with this mechanism, by the end of 2022 mispricing measures had declined, hedge fund borrowing in the repo market had fallen, and repo specialness had receded. More generally, elevated repo borrowing by hedge funds reflects equilibrium demand for cash-market positions that require securities borrowing, i.e., hedge funds short-selling bonds to accommodate strong cash-market demand. Repo specialness therefore emerges as a by-product of securities borrowing associated with cash-market clearing and should be interpreted as a macroeconomic outcome rather than an idiosyncratic feature of individual securities.

7 Conclusion

This paper provides a unified theoretical and empirical framework for understanding repo specialness as an equilibrium outcome of cross-market positioning, rather than as a purely mechanical consequence of collateral scarcity. Our model shows that the preferred habitat investors' preference for physical bonds generates excess cash-market demand, which hedge funds accommodate by shorting bonds and hedging with futures. Because shorting requires borrowing in the repo market, the cash-market imbalance pins down hedge fund repo demand independently of specialness—making them inelastic borrowers despite being the most elastic participants in the cash market. Allowing for endogenous bond yields shows that hedge funds' repo-market elasticity is inherited from preferred-habitat investors: it equals the price sensitivity of the investors whose demand they accommodate in the cash market.

We take these predictions to the data using the first demand-system estimation for a repo market, applying the O-GIV framework to novel SFTDS transaction-level data for German government collateral. The estimated aggregate elasticity of 1.2 masks a stark asymmetry: non-euro area hedge funds exhibit an elasticity of just 0.11, while public sector supply—provided nearly monopolistically by the German debt management office and the Eurosystem—is highly elastic. Hedge funds are the primary drivers of specialness, with their elevated borrowing closely tracking the widening of the cash-futures basis and other

[21] When repo specialness increases, the funding leg of the trade becomes more expensive, limiting arbitrage capacity and allowing the basis to persist; see, e.g., Pelizzon et al. (2025).

mispricing measures. Their retrenchment from the market, rather than a change in price sensitivity, explains the normalization of specialness.

The joint analysis of the topology of the repo market and sector-level elasticities reveals a previously unappreciated feature of the collateralized funding market: the marginal price-setters in the European sovereign repo market are offshore, non-bank entities outside the euro-area regulatory perimeter. Financial stability in the European sovereign bond market and the effectiveness of monetary policy transmission through repo rates increasingly depend on actors whom euro-area regulators cannot directly supervise.

These findings carry implications for monetary policy. As sovereign bond markets have taken center stage in unconventional monetary and fiscal interventions, our results show that the impact of such operations depends critically on how bond investors—hedge funds in particular—accommodate purchases and sales through the repo market. Repo specialness thus serves as an observable signal of cash-market demand pressure: widening bases coincide with specialness episodes, giving central banks an indicator for monitoring monetary policy transmission. Securities lending, in turn, becomes a policy lever: the design and pricing of lending facilities have first-order effects on specialness and, through it, on cash bond prices. The lever has limits, however: because repo demand inelasticity originates in the cash market, specialness cannot be eliminated as long as profitable basis-trade opportunities persist, unless bond supply satisfies preferred-habitat demand or securities are lent at the policy rate. Cash-market interventions should therefore be accompanied by monitoring of repo-market conditions: in the repo market, as our framework shows, prices are set by demand pressures that originate elsewhere.

References

Ampudia, M., P. Formoso da Silva, S. Hesse, A. Putz, F. Schobert, J. von Landesberger, and A. Wohlert (2025). Supply-side interventions in the Bund repo market. Technical report, ECB Working Paper, forthcoming.

Andersen, L., D. Duffie, and Y. Song (2019). Funding value adjustments. *The Journal of Finance 74*(1), 145–192.

Arrata, W., B. Nguyen, I. Rahmouni-Rousseau, and M. Vari (2020). The scarcity effect of QE on repo rates: Evidence from the euro area. *Journal of Financial Economics 137*(3), 837–856.

Ballensiefen, B., A. Ranaldo, and H. Winterberg (2023). Money Market Disconnect. *The Review of Financial Studies*.

Baltzer, M., K. Schlepper, and C. Speck (2022). The Eurosystem's asset purchase programmes, securities lending and Bund specialness. *Deutsche Bundesbank Discussion Paper No 39/2022*.

Barth, D. and R. J. Kahn (2025). Hedge funds and the treasury cash-futures basis trade. *Journal of Monetary Economics*, 103823.

Bassi, C., M. Grill, F. Hermes, H. Mirza, C. O'Donnell, and M. Wedow (2024). Enhancing repo market transparency: the EU Securities Financing Transactions Regulation.

Bassi, C., M. Grill, F. Hermes, H. Mirza, C. ODonnell, and M. Wedow (2025). Enhancing Repo Market Transparency: The EU Securities Financing Transactions Regulation. *Journal of Financial Regulation 11*(1), 98–118.

Besugo, R., J.-B. Cinquin, B. Nguyen, L. Paoli, and J. Szydlo (2025). Plumbing matters: a cross-market functioning index. *ECB Working Papers Series, forthcoming*.

Besugo, R., B. Nguyen, A. Poinelli, and M. Scheicher (2025). Dealers costs of intermediation in fixed income markets: Empirical results for the euro area. Technical report, SUERF Policy Brief No 1226.

Brunnermeier, M. K. and L. H. Pedersen (2009). Market liquidity and funding liquidity. *The review of financial studies 22*(6), 2201–2238.

Carrera de Souza, T. and T. Hudepohl (2024). Frictions in scaling up central bank balance sheet policies: How Eurosystem asset purchases impact the repo market. *Journal of Banking & Finance 158*, 107037.

Chaudhary, M., Z. Fu, and J. Li (2023). Corporate bond multipliers: Substitutes matter. *The Review of Financial Studies*.

Chaudhary, M., Z. Fu, and H. Zhou (2024). Anatomy of the Treasury Market: Who Moves Yields?

Corradin, S. and A. Maddaloni (2020). The importance of being special: Repo markets during the crisis. *Journal of Financial Economics 137*(2), 392–429.

d'Avernas, A., D. Petersen, and Q. Vandeweyer (2025). The central bank's balance sheet and treasury market disruptions.

D'Avernas, A., D. Petersen, and Q. Vandeweyer (2025). The Central Bank's Balance Sheet and Treasury Market Disruptions. *ECB Working Paper No. 2025/3066*.

Duffie, D. (1996). Special Repo Rates. *Journal of Finance 51*(2), 493–526.

Duffie, D. (2023). Resilience redux in the us treasury market. In *Jackson Hole Symposium, Federal Reserve Bank of Kansas City, September 2023*.

Duffie, D., P. Dworczak, and H. Zhu (2016). Benchmarks in search markets. *SSRN Electronic Journal*.

Duffie, D. and A. Krishnamurthy (2016). Passthrough efficiency in the feds new monetary policy setting. In *Designing Resilient Monetary Policy Frameworks for the Future. Federal Reserve Bank of Kansas City, Jackson Hole Symposium*, Volume 1.

DAmico, S., R. Fan, and Y. Kitsul (2018). The Scarcity Value of Treasury Collateral: Repo-Market Effects of Security-Specific Supply and Demand Factors. *Journal of Financial and Quantitative Analysis 53*(5), 2103–2129.

DAmico, S. and N. A. Pancost (2022). Special Repo Rates and the Cross-Section of Bond Prices: The Role of the Special Collateral Risk Premium. *Review of Finance 26*(1), 117–162.

ECB (2024). Euro money market study 2024.

Eisenschmidt, J., Y. Ma, and A. L. Zhang (2020). Monetary policy transmission in segmented markets. *SSRN Electronic Journal*.

ESMA (2024). Securities Financing Transactions markets 2024. Technical report, ESMA.

Gabaix, X. and R. S. J. Koijen (2021). In search of the origins of financial fluctuations: The inelastic markets hypothesis. *NBER Working Paper No. w28967*.

Gabaix, X. and R. S. J. Koijen (2024). Granular Instrumental Variables. *Journal of Political Economy 132*(7), 2274–2303.

Greppmair, S. and S. Jank (2023). Collateral scarcity and market functioning: Insights from the Eurosystem securities lending facilities. *Deutsche Bundesbank Discussion Paper No 31/2023*.

Hermes, F., M. Schmeling, and A. Schrimpf (2025). The International Dimension of Repo: Five New Facts. *SSRN Working Paper*.

Holm-Hadulla, F., M. Leombroni, L. Mota, M. Papoutsi, and F. Corell (2024). Drivers of convenience yields. *Available at SSRN 4656963*.

Huh, Y. and S. Infante (2021). Bond market intermediation and the Role of Repo. *Journal of Banking & Finance 122*, 105999.

Infante, S. (2020). Private money creation with safe assets and term premia. *Journal of Financial Economics 136*(3), 828–856.

Jank, S. and E. Moench (2018). The impact of eurosystem bond purchases on the repo market. *Bundesbank Research Brief 21.*.

Jansen, K. A., W. Li, and L. Schmid (2024). Granular treasury demand with arbitrageurs. Technical report, USC Marshall School of Business.

Jappelli, R., M. G. Subrahmanyam, and L. Pelizzon (2023). Quantitative Easing, the Repo Market, and the Term Structure of Interest Rates.

Jordan, B. D. and S. D. Jordan (1997). Special repo rates: An empirical analysis. *The Journal of Finance 52*(5), 2051–2072.

Kaminska, I. and G. Zinna (2020). Official demand for us debt: Implications for us real rates. *Journal of Money, Credit and Banking 52*(2-3), 323–364.

Koijen, R. S., F. Koulischer, B. Nguyen, and M. Yogo (2021). Inspecting the mechanism of quantitative easing in the euro area. *Journal of Financial Economics 140*(1), 1–20.

Koijen, R. S. J. and M. Yogo (2019). A Demand System Approach to Asset Pricing. *Journal of Political Economy 127*(4), 1475–1515.

Krishnamurthy, A. (2002). The bond/old-bond spread. *Journal of financial Economics 66*(2-3), 463–506.

Kubitza, C., J.-D. Sigaux, and Q. Vandeweyer (2024). Cross-Currency Basis Risk and International Capital Flows. *Fama-Miller Working Paper, Chicago Booth Research Paper 24-18*.

Kubitza, C., J.-D. Sigaux, and Q. Vandeweyer (2025). The implications of cip deviations for international capital flows. Technical report, ECB Working Paper.

Li, J., Y. Ma, C. Mendicino, and D. Supera (2026). Bank to non-bank lending and the reallocation of credit.

Mancini, L., A. Ranaldo, and J. Wrampelmeyer (2015). The Euro Interbank Repo Market. *Review of financial studies forthcomin*(7), 1747–1779.

Nguyen, B., D. Tomio, and M. Vari (2023). Safe asset scarcity and monetary policy transmission. *Banque de France Working Paper Series no. 934*.

Pelizzon, L., M. G. Subrahmanyam, and D. Tomio (2025). Central BankDriven Mispricing. *Journal of Financial Economics 166*, 104004.

Schaffner, P., A. Ranaldo, and K. Tsatsaronis (2019). Euro Repo Market Functioning: Collateral is King. *Journal of Financial Markets 5*(1).

Vayanos, D. and J. Vila (2021). A preferredhabitat model of the term structure of interest rates. *Econometrica 89*(1), 77–112.

Figures

Figure 1. Flows and stock quantities in the repo market

In this figure, we show the size of the repo market for German government collateral, using two different metrics, a flow measure and a stock measure. In Panel A we show the flow of repo transactions, which quantifies the size of new trades happening on a given day, regardless of the term of the cotract. In Panel B, we report the stock measure, which captures the quantity of securities currently on loan on any given day. For example, a three-day transaction will be counted three times in the measure shown in Panel B, but only once, on the day the trade was initiated, in Panel A. Repo volumes date are from SFTDS, and eurosystem repo lending data are from MOPDB. Data are limited to repo transactions that use German central government bonds as collateral. The sample period is from January 2021 to June 2024.

A: Flows

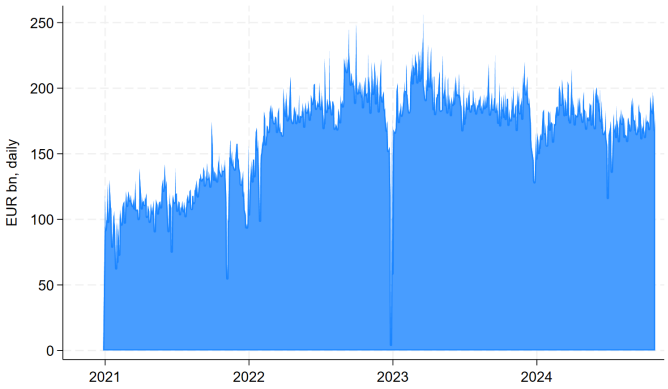

B: Stocks

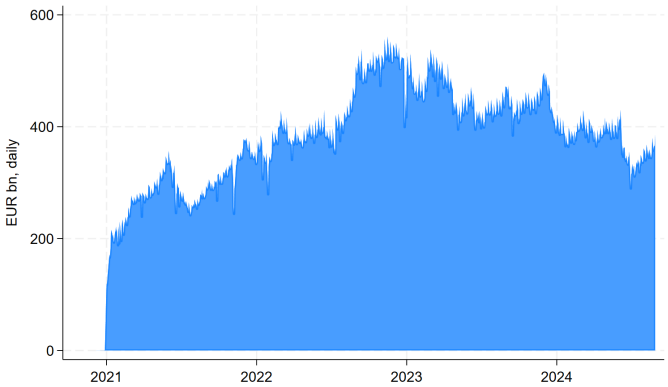

Figure 2. Gross outstanding amounts of securities lent and borrowed by sector

In this figure, we show the repo gross positions of different sectors on the repo market, measured as the stock of securites borrowed (Panel A) or lent (Panel B), in billions of euro. We net out intra-sector transactions, so that we can interpret the quantities as measuring the obligation of a sector toward the rest of the world. We do not report volumes of money-market funds, as they are minimal. Repo transactions data are from SFTDS, while data for the position by the euro area's central banks are from MOPDB and capture the eurosystem securities lending facility's outstanding volume. For the latter, we only consider securities lent against cash, to better approximate repo the economic exposure of a repo transaction. We focus on transactions collaterlized by German government bonds, and the subsample covers the period from March 2021 to June 2024.

A: Borrowed stocks of securities

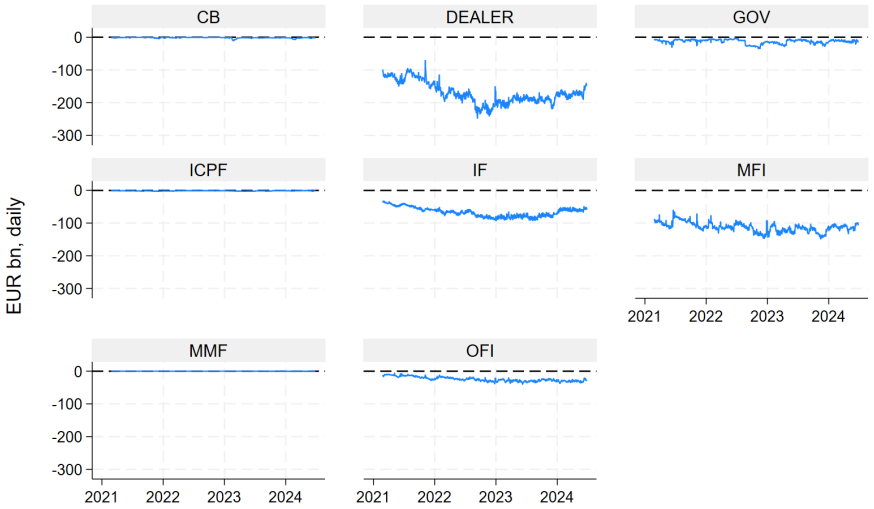

B: Lent stocks of securities

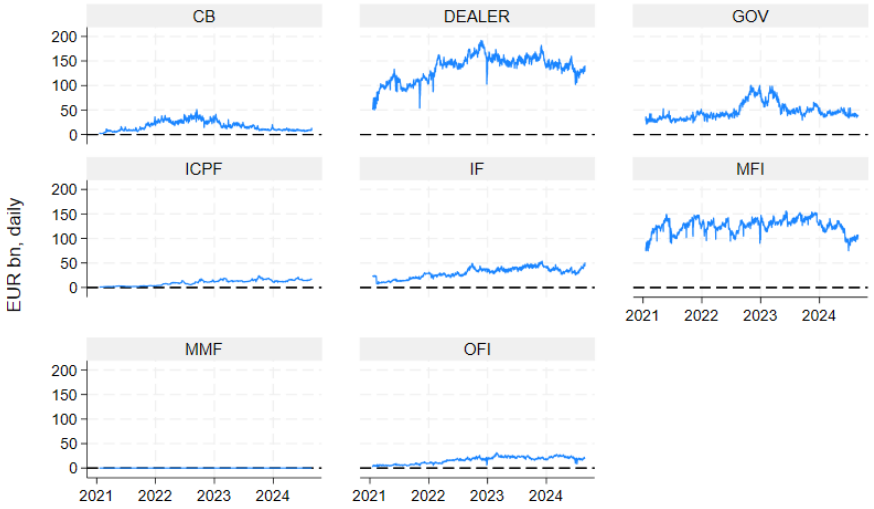

Figure 3. Securities lending positions by sector and geography

In this figure, we show the net position on the repo market, by sector, calculated as the amount of securities lent minus securities borrowed. A positive (negative) amount means a sector is a net-security lender (borrower) on the repo market. The measure captures the stock positions on the repo market, i.e., the outstanding amount of repo transactions taking into account the term of each repo contract. We show the quantity for euro-area and non-euro Area sectors, separately. We exclude positions of money market funds, insurance companies, and pension funds, as they are minimal. Repo data are sourced from SFTDS. Data for the securities lending by the central bank are obtained from MOPDB and are limited to cash-collateralized securities-lending transactions. All transactions in the sample are collateralized with German government bonds. The sample covers March 2021 to June 2024.

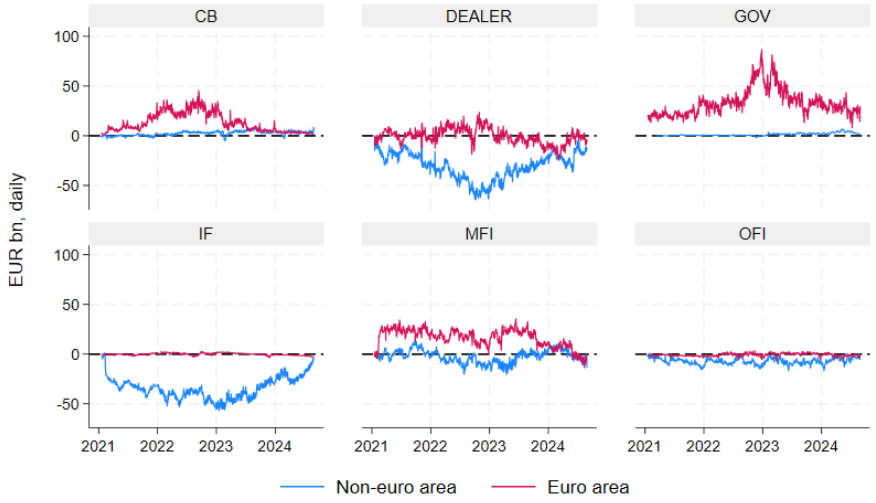

Figure 4. The structure of the repo market

In this figure, we show the structure of the repo market. For each couple of investors, we show the average end-of-day oustanding amount of securities on loan, in billions of euro, for 2022. Sectors on the left are net securities lenders, sectors on the right are net securiites borrowers. The "Other" includes the smaller players, i.e., money market funds, non-euro area banks, euro-area hedge funds, insurance companies and pension funds, and foreing government and central banks. We consider repo transactions that have German government bond as collateral, and source repo data from SFTDS, while the eurosystem's cash-against-collateral securities lending data are from MOPDB.

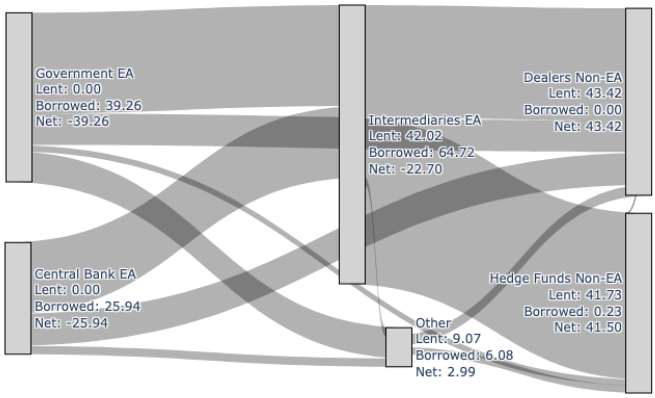

Figure 5. The concentration of the repo market for German government collateral

In this figure, we compare the equal weights versus the value weights - the repo market participation shares - by sector. Equal weights are computed as the inverse of the number of sectors and represent a theoretical benchmark where all sectors participate in the repo market for German collateral with equal footprint. The value weights are the actual participation shares of the different sectors in terms of net securities borrowing positions, see Equation 21. We consider repo transactions that have German government bond as collateral, and source repo data from SFTDS, while the eurosystem's cash-against-collateral securities lending data are from MOPDB.

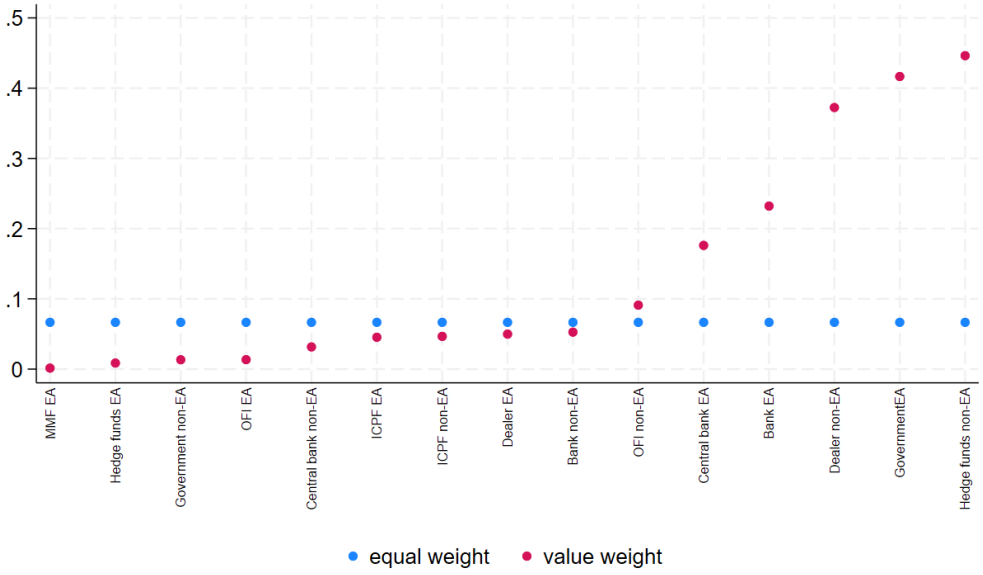

Figure 6. Market multipliers in the literature

In this figure, we compare our estimated market multiplier—i.e., the inverse of the price-elasticity ζ from Eq. 19—for the repo market to existing estimates in the literature for several asset classes. Jansen et al. (2024), Chaudhary et al. (2024), and Kaminska and Zinna (2020) estimate the elasticity for the US Treasury market, Kubitza et al. (2025) for the FX derivatives market and Gabaix and Koijen (2021) for the U.S. stock market.

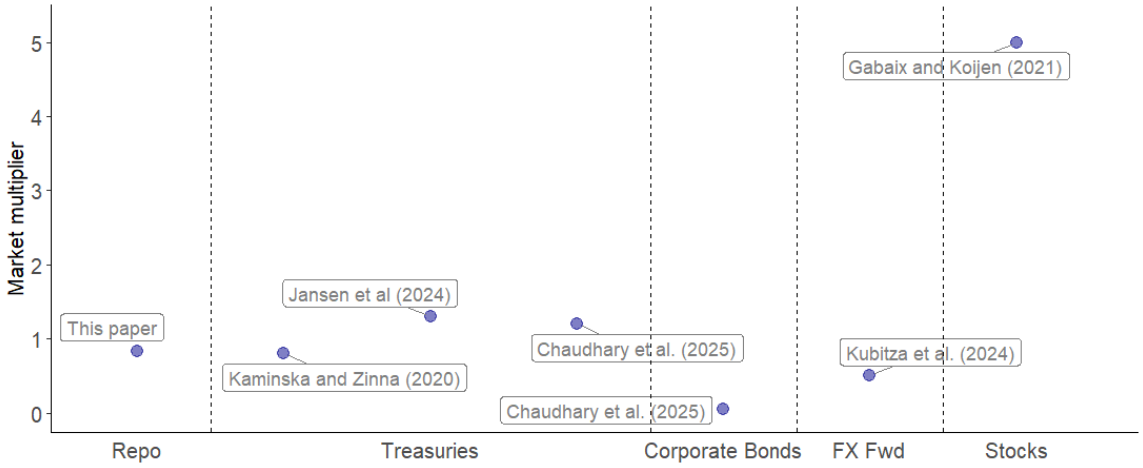

Figure 7. Market multiplier by sector

In this figure we compare the market multipliers, the inverse of the price-elasticity ζ_i, for the largest participants to the repo market. We exclude sectors whose outstanding positions make up less than 5% of the total. These sectors cumulatively represent 83% (91%) of the supply (demand) in the market. We allocate a sector to supply (demand) if the sector is a net lender (borrower) of securities, on average, over our sample period. The market multipliers are estimates of $1/\zeta_i$ from Eq. 16. Our data includes repo transactions that use German government bonds as collateral for the March 2021–June 2024 period. Repo data are from SFTDS, and eurosystem repo lending data are from MOPDB

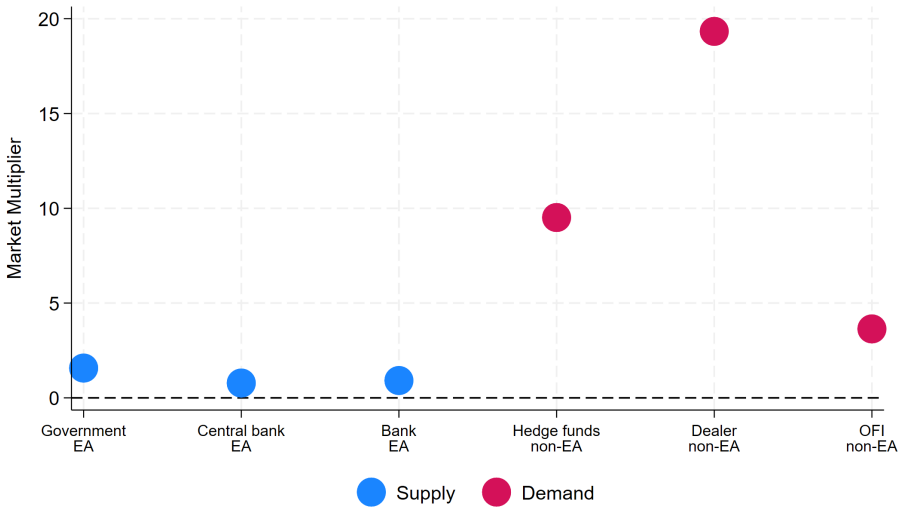

Figure 8. Contribution to repo specialness changes by sector

This figure shows the decomposition of the changes in the specialness spread among the contributions of different sectors. We define changes in specialness as the relative deviation from the three-month moving average, and changes in quantities accordingly. In Panel A, we average the daily contributions at the quarter level, and classify the sectors into the public sector, the euro-area government and central bank; intermediaries, euro-area banks and dealers; and non-euro Area hedge funds and their dealers. The "Other" sector aggregates the contributions of non-euro Area central bank, governments, and banks, and insuarance companies, pension funds, and money market funds. The sector decomposition is calculated as $\Delta\hat{p}_{i,t} = \frac{1}{\hat{\zeta}_S} S_i(\hat{\Delta q}_i + \hat{\lambda}_i \eta_t + \hat{u}_{i,t})$, for sector-$i$ and day-t, where $\hat{\zeta}_S$ is reported in Table 3. In Panel B, we report the cumulated contribution of each sector. To obtain this, we scale $\Delta\hat{p}_{i,t}$ by the observed three-month average $\overline{p_t}$, divide it by the number of sectors, and cumulate it across time. In the same panel, we also show the 3-month moving average of specialness. We estimate the demand-system using outstanding positions in repo contracts collaterlized by German government bonds, for the period betwee from March 2021 to June 2024. Repo data are from SFTDS, and eurosystem repo lending data are from MOPDB.

A: Specialness change by quarter

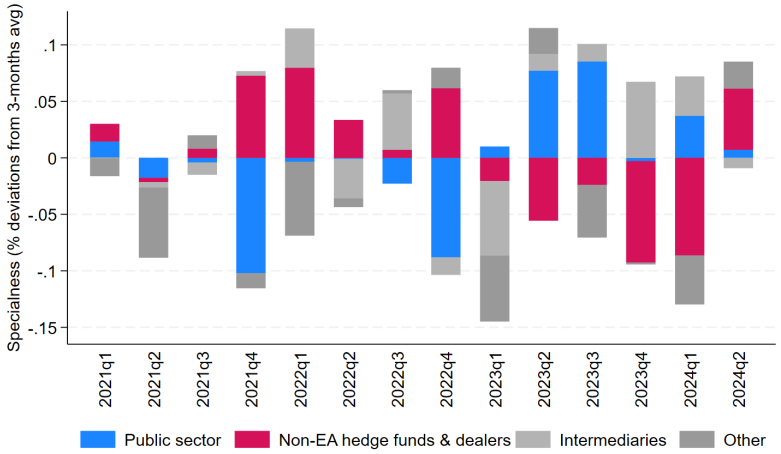

B: Cumulative specialness innovations

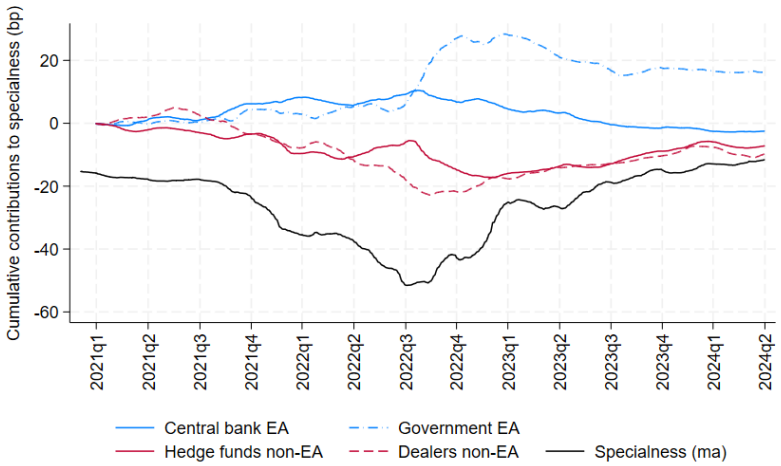

Figure 9. Decomposition of repo specialness contributions by common factors

This figure shows the decomposition of the changes in the specialness spread among common factors. We define changes in specialness as the relative deviation from the three-month moving average, and changes in quantities accordingly. We average the daily contributions at the quarter level, and classify the factors into monetary policy, repo demand, repo supply, and bond supply factors. The factor decomposition for factor $\eta(k)$ is calculated as $\Delta\hat{p}_{\eta(k),t} = \frac{1}{\hat{\zeta}_S} S_i \hat{\lambda}_i(k)\eta_t(k)$, for and day-$t$, where $\hat{\zeta}_S$ is reported in Table 3. The sample period goes from March 2021 to June 2024 and the sample consists of repos with underlying collateral a German government bond. Repo data are from SFTDS, and eurosystem repo lending data are from MOPDB.

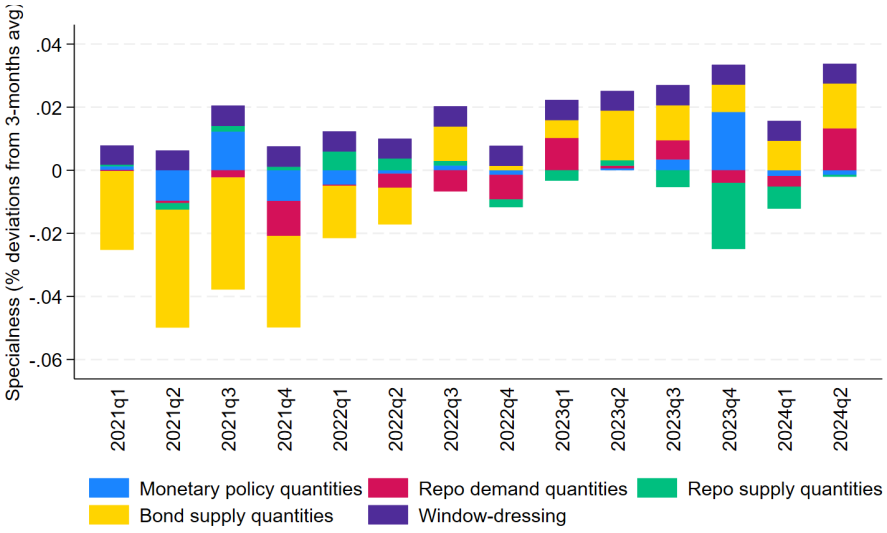

Figure 10. Repo market elasticities during QE and QT

In this figure, we compare the price-elasticity ζ_i of the largest sectors in the repo market across periods of varying monetary policy. The first sub-sample covers the eurosystem' balance sheet expansion, which terminated with the large repayment of targeted long-term refinancing operations, March 2021–November 2022, while in the second period the ECB's balance sheet is reducing, December 2022–June 2024. We include sectors whose positions account for at least 5% of the total repo outstanding, which cumulatively represent 83% (91%) of the supply (demand). The price elasticities ζ_i in Eq. 16 are estimated using the optimal GIV method shown in Section 4. The estimation weights S_i are sample-specific. Repo data are from SFTDS, and eurosystem repo lending data are from MOPDB. We estimate the model on repo transactions collateralized by German government bonds.

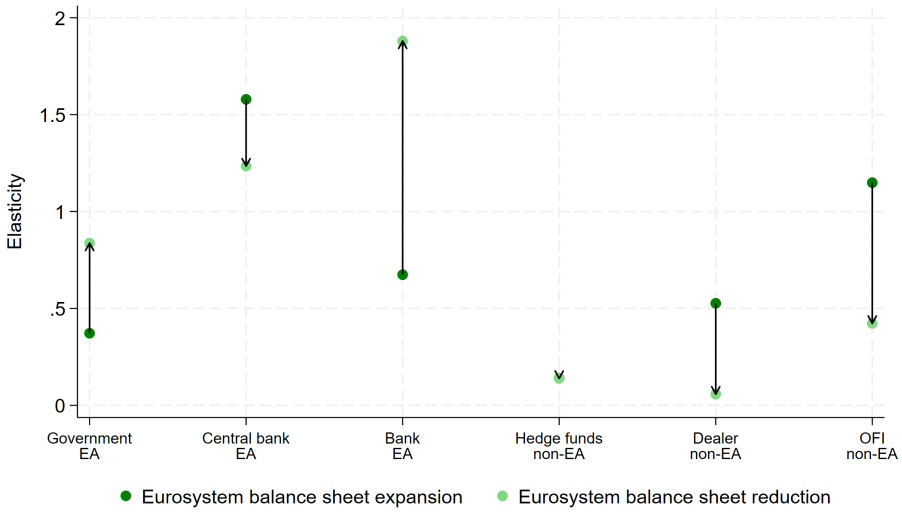

Figure 11. Decomposition of repo specialness contributions of EA intermediaries at quarter-end

In this figure, we show the cyclicality of sectoral contribution to specialness within a quarter. That is, we calculate the contribution $\Delta\hat{p}_{i,t} = \frac{1}{\hat{\zeta}_S} S_i(\hat{\tilde{\Delta}} q_i + \hat{\lambda}_i \eta_t + \hat{u}_{i,t})$, for sector-$i$ and day-t, where $\hat{\zeta}_S$ is reported in Table 3, aggregate it by sector type, and average it by the number of days to the end of the quarter. We group investors in euro-Area intermediaries (banks and dealers) and the rest. We show the results separately for two sub-samples, before and after March 2022. In the earlier part of the sample, euro-area intermediaries benefitted from a leverage ratio exemption, which allowed them to exclude sizeable central bank exposures, contributing to lower balance sheet costs. The sample consists of repos with underlying collateral a German government bond. Repo data are from SFTDS, and eurosystem repo lending data are from MOPDB.

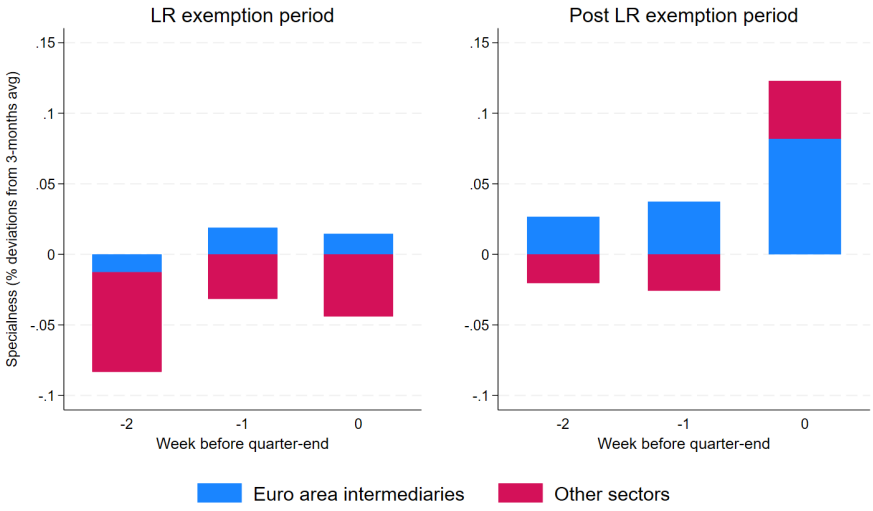

Figure 12. Mispricing measures and hedge funds positions

In Panel A, we show three measures of asset mispricing against the net aggregate securities borrowing positions by non-euro Area hedge funds. The three measures consist of the cash-futures bond basis, the average government-OIS spread (asset swap), and the pricing errors of a Nelson-Siegel curve. The future-bond basis is calculated for the next-next-delivery 10-year contract, while the other two measures are based on the universe of traded German government bonds. The mispricing measures are expressed as percentiles over the sample period, to allow for a comparison across measures. The hedge fund net-borrowing positions are in billion of euros of German governmnet bonds borrowed on the repo market. In Panel B, we show the development of repo specialness and the cash-futures bond basis. The former is measured in basis points, while the latter is annualized and shown as percentage difference between futures and cash bond price. Repo volumes and prices are from SFTDS, pricing errors and asset swaps are calculated using MTS pricing data, and the Bund cash-futures basis is sourced from Bloomberg and uses the euro short-term rate (ESTR) as funding rate.

A: Mispricing measures against hedge fund repo positions

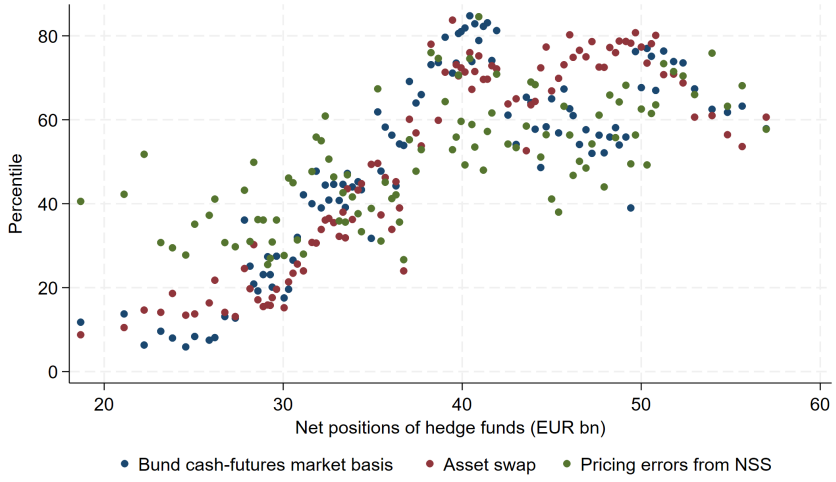

B: Cash-future bond basis and specialness

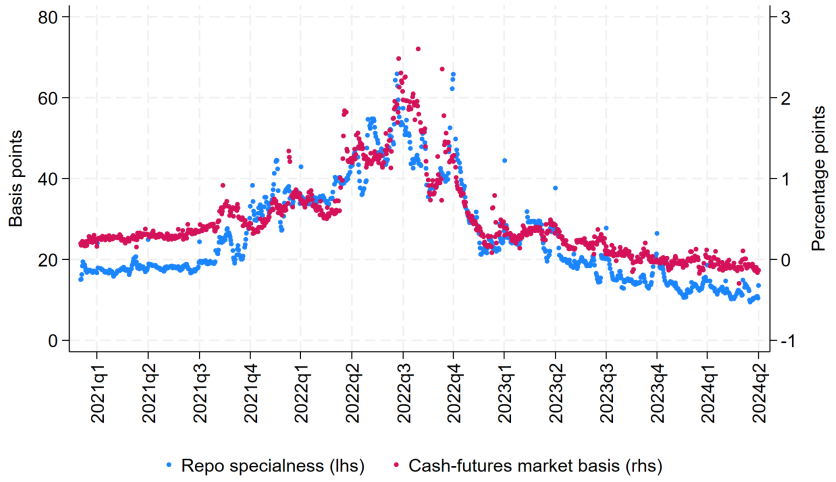

Figure 13. The inverse relation between cash- and repo-market elasticity

In this figure, we plot a sector's price elasticity in the cash market from Koijen et al. (2021) against its price elasticity of the repo market from Table 3. We impute an elasticity of zero for the eurosystem in the cash market and for a sector when its elasticity in either the cash or repo market is negative. The foreign sector' elasticity in the repo marker is the average of the repo market elasticity of non-EA dealers and hedge funds.

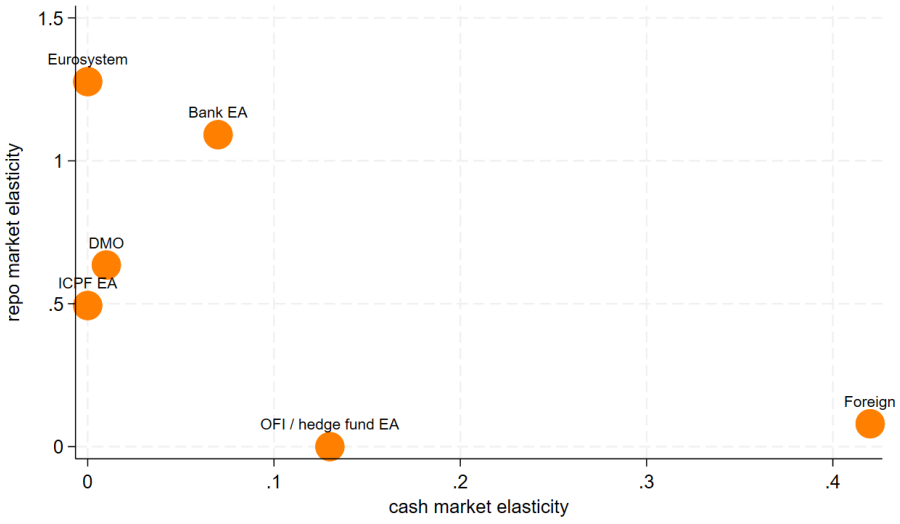

Figure 14. Elasticity by sector: estimation with PCA

In this figure we compare the price-elasticity ζ_i for the largest participants to the repo market when three principal components are added to the common factors. The pricipal components are calculated on the residuals from the first stage, see Eq. 22. We first estimate the first stage equation with common factors and sector fixed effects, we extract sector's residuals, we apply the PCA on these residuals and we re-estimate the first stage using common factors, sector fixed effects and the three principal components. We exclude sectors whose outstanding positions make up less than 5% of the total. These sectors cumulatively represent 83% (91%) of the supply (demand) in the market. We allocate a sector to supply (demand) if the sector is a net lender (borrower) of securities, on average, over our sample period. The elasticities are computed from Eq. 25. Our data includes repo transactions that use German government bonds as collateral for the March 2021–June 2024 period. Repo data are from SFTDS, and eurosystem repo lending data are from MOPDB

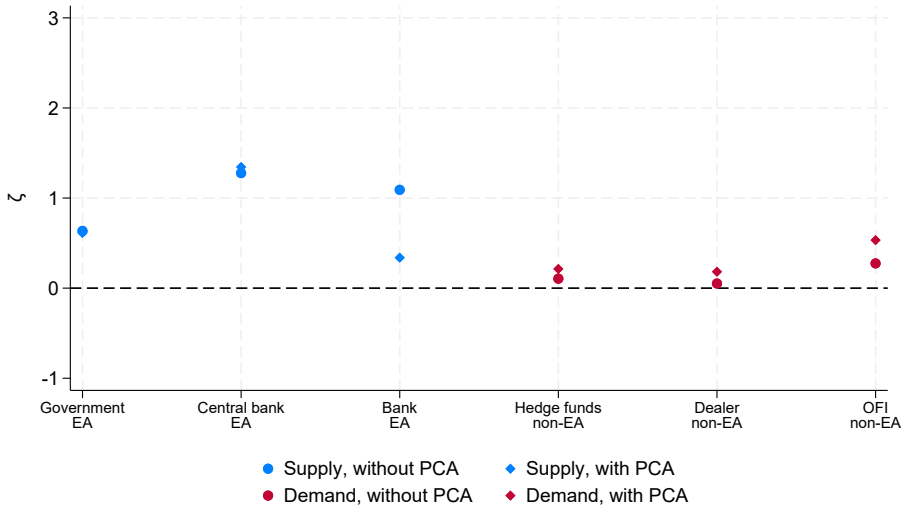

Tables

Table 1
Securities holdings and amounts borrowed and lent on the repo market, by sector

In this table, we show the amount of securities held, and the amounts of securities lent and borrowed on the repo market. We report the average by sector at the yearly level, and restrict the sample to German government bonds and repo transactions that use them as collateral. The amounts borrowed and lent on the repo market represents averages of end-of-day stocks against all counterparties. All quantities are expresse in billion of euros. eurosystem borrowing and lending figures are limited to cash-collateralized transactions. Holdings of dealers, banks (monetary financial institutions, MFI), and other financial institutions (OFI) are reported as one, as the holding dataset does not distinguish between them. Non-euro area holdings are computed as a residual of the nomial value oustanding of German bonds. We exclude money market funds, households and non-financial corporations from the euro area, as their participation to the repo market is negligible. Our sample extends from January 2021 to June 2024. Holdings data are from SHSS, repo volumes from SFTDS, and eurosystem repo lending data are from MOPDB.

Year	Area	Sector	Holdings	Borrowed	Lent
2021	**euro Area**	ICPF	121	0.2	0.2
		Bank	11	61.5	83.0
		OFI		3.5	2.1
		Dealer		65.7	65.1
		Hedge funds	117	1.9	1.6
		Government	172	10.0	32.0
		Central Bank	554	0.0	9.1
	Non-euro Area		554	142.0	92.0
2022	**euro Area**	ICPF	121	0.1	3.9
		Bank	3	67.0	83.9
		OFI		6.3	5.7
		Dealer		108.1	113.9
		Hedge funds	148	1.8	2.3
		Government	169	13.6	53.0
		Central Bank	655	0.0	25.9
	Non-euro Area		534	215.9	124.3
2023	**euro Area**	ICPF	141	0.1	7.1
		Bank	3	80.0	99.4
		OFI		6.9	7.2
		Dealer		122.4	119.2
		Hedge funds	178	2.2	2.7
		Government	203	14.2	56.2
		Central Bank	621	0.0	10.6
	Non-euro Area		616	215.4	139.3
2024	**Euro Area**	ICPF	144	0.3	6.5
		Bank	3	81.8	88.3
		OFI		7.1	6.0
		Dealer		111.9	103.0
		Hedge funds	178	2.2	1.2
		Government	208	11.0	40.3
		Central Bank	577	0.0	3.7
	Non-euro Area		700	172.4	137.8

Table 2
First stage of Optimal-GIV - large sectors

In this table, we show the coefficients resulting from the first stage of the O-GIV where sector i net securities borrowing volumes in the repo market are residualised against a broad set of common factors and sector fixed effects, see equation 22. The dependent variable is the percentage deviation of sector i ' net securities borrowing volumes from its 3-months moving average. The same transformation is applied to common factos, which are listed in Section 3. We exclude sectors whose outstanding positions make up less than 5% of the total. These sectors cumulatively represent 83% (91%) of the supply (demand) in the market. We allocate a sector to supply (demand) if the sector is a net lender (borrower) of securities, on average, over our sample period. *, **, and *** indicate that parameters are significantly different from zero at the 10%, 5%, or 1% level, respectively. Significance is based on autoregression and heteroskedasticity-robust Newey-West standard errors.

	Government EA	Central Bank EA	Banks EA	Dealers non-EA	OFI non-EA	Hedge funds non-EA
OIS 1M	-0.021^{*}	-0.054^{***}	-0.018	0.039^{***}	0.022	-0.012
OIS1M DFR spread	-0.001	-0.008	0.000	0.003	0.002	-0.002
3M1M OIS spread	-0.019^{**}	0.001	0.004	0.008	0.005	-0.018^{**}
1Y3M OIS spread	-0.002	-0.021^{***}	-0.005	0.011^{**}	0.009	0.011^{***}
Slope YC	-0.085^{***}	0.004	-0.053	-0.006	0.120^{*}	0.056^{***}
Cross-ccy basis	-0.047^{**}	-0.002	0.077^{**}	-0.043^{*}	-0.066	-0.019^{*}
Duffie	0.002	0.003	0.002	0.004	-0.011	0.000
KFW vs DE spread	0.023	-0.207^{***}	0.164^{**}	-0.069^{*}	-0.172^{*}	-0.020
OID skew	-0.019	0.113^{*}	0.005	-0.009	0.197^{**}	-0.067^{**}
SMOVE	-0.168	-0.038	-0.345	0.492^{**}	-0.042	0.233^{**}
Bund vol prem	0.030	-0.094^{**}	0.053	-0.043^{*}	0.026	-0.052^{***}
NSS dev	-0.103	-0.414^{**}	0.088	0.088	0.607^{***}	-0.237^{***}
Bund OpenInt	-0.034	0.552^{*}	-0.344	-0.120	1.421^{**}	0.243
DE vs FR spread	0.052	0.130	0.685	-0.337	-0.075	0.187^{*}
Free float	3.700^{***}	3.080^{**}	-1.913	-2.155^{**}	-3.363	-1.773^{***}
ECB holdings	1.508^{**}	-2.506^{*}	-3.714^{*}	-1.655	-1.912	1.135^{*}
DFA holdings	-3.127^{***}	2.324^{**}	2.901^{**}	0.921^{*}	1.312	2.018^{***}
Constant	-0.005	0.036	0.007	0.063	0.107	0.004
R-squared	0.43	0.42	0.14	0.30	0.17	0.48
Observations	823	823	823	823	823	823
Day-of-the-month FE	Yes	Yes	Yes	Yes	Yes	Yes
Weights	42%	18%	23%	37%	9%	45%
Side	Supply	Supply	Supply	Demand	Demand	Demand

Table 3
Sector Elasticities ζ_i

In this Table, we show the sector-level elasticities ζ_i. The table reports the market share of the sector calculated as in Eq. 21, the sector price-elasticity ζ_i calculated from Eq. 25, the market multiplier or the inverse of the price-elasticity and the contribution to the aggregate elasticity. We allocate a sector to supply (demand) if the sector is a net lender (borrower) of securities, on average, over our sample period.

Sector	**Share**	ζ	**M**	$\zeta\times$**Share**
Supply				
Euro area				
Government	42%	0.64	1.56	21%
Central bank	18%	1.28	0.78	18%
Bank	23%	1.09	0.91	20%
Dealer	5%	5.38	0.18	21%
ICPF	5%	0.49	2.04	2%
Hedge funds	1%	−0.85	−1.17	−1%
Non-euro area				
Government	1%	−1.41	−0.70	−1%
Central bank	3%	2.50	0.40	6%
ICPF	5%	0.48	2.08	2%
Demand				
Euro area				
OFI	1%	−0.89	−1.12	−1%
MMF	0%	0.11	9.10	0%
Non-euro area				
Bank	5%	1.60	0.62	7%
Dealer	37%	0.05	20.00	1%
OFI	9%	0.27	3.70	2%
Hedge funds	45%	0.11	9.10	4%
Aggregate	200%	1.20	0.83	100%

Table 4
Repo vs. Cash Market Elasticities

In this table, we compare a sector's price elasticity in the cash market from Koijen et al. (2021) against its price elasticity of the repo market from Table 3.

	Linzert et al. (2025) *Repo market elasticity*	**Koijen et al. (2021)** *Cash market elasticity*	
Non-euro area			
Government	−1.41		0.42
Central bank	2.50		
Bank	1.60		
Dealer	0.05		
OFI	0.27		
Hedge funds	0.11		
ICPF	0.48		
Euro area			
Government	0.64	0.01	
Central bank	1.28	–	
Bank	1.09	0.07	
Dealer	5.38	–	
OFI	−0.89		0.13
Hedge funds	−0.85		
ICPF	0.49	−0.34	
MMF	0.11	–	
Households	–	0.01	

Table 5
First stage of Optimal-GIV with PCs - large sectors

In this table, we show the coefficients resulting from the first stage of the O-GIV where sector i net securities borrowing volumes in the repo market are residualised against a broad set of common factors and sector fixed effects, see equation 22. The dependent variable is the percentage deviation of sector i ' net securities borrowing volumes from its 3-months moving average. The same transformation is applied to common factos, which are listed in Section 3. We exclude sectors whose outstanding positions make up less than 5% of the total. These sectors cumulatively represent 83% (91%) of the supply (demand) in the market. We allocate a sector to supply (demand) if the sector is a net lender (borrower) of securities, on average, over our sample period. *, **, and *** indicate that parameters are significantly different from zero at the 10%, 5%, or 1% level, respectively. Significance is based on autoregression and heteroskedasticity-robust Newey-West standard errors.

	Government EA	Central Bank EA	Banks EA	Dealers non-EA	OFI non-EA	Hedge funds non-EA
OIS1M DFR spread	−0.001	−0.008*	0.000	0.003	0.002	−0.002
3M1M OIS spread	−0.019**	0.001	0.004	0.008	0.005	−0.018***
1Y3M OIS spread	−0.002	−0.021***	−0.005	0.011***	0.009*	0.011***
KFW vs DE spread	0.023	−0.207***	0.164***	−0.069**	−0.172**	−0.020
DE vs FR spread	0.052	0.130	0.685**	−0.337***	−0.075	0.187***
OIS 1M	−0.021**	−0.054***	−0.018	0.039***	0.022	−0.012
Slope YC	−0.085***	0.004	−0.053	−0.006	0.120***	0.056***
OID skew	−0.019	0.113*	0.005	−0.009	0.197***	−0.067***
SMOVE	−0.168	−0.038	−0.345	0.492***	−0.042	0.233***
Bund vol prem	0.030*	−0.094***	0.053*	−0.043**	0.026	−0.052***
NSS dev	−0.103*	−0.414***	0.088	0.088	0.607***	−0.237***
Bund OpenInt	−0.034	0.552*	−0.344	−0.120	1.421***	0.243**
Cross-ccy basis	−0.047**	−0.002	0.077***	−0.043**	−0.066*	−0.019**
Duffie	0.002	0.003	0.002	0.004	−0.011	0.000
Free float	3.700***	3.080***	−1.913	−2.155***	−3.363*	−1.773***
ECB holdings	1.508**	−2.506**	−3.714**	−1.655*	−1.912	1.135***
DFA holdings	−3.127***	2.324***	2.901***	0.921***	1.312	2.018***
PC 1	−0.003	0.047***	0.121***	−0.081***	0.008	−0.049***
PC 2	−0.007	0.075***	0.016	−0.012	−0.046***	0.012***
PC 3	−0.048***	0.013	0.108***	0.017*	0.191***	0.003
Constant	−0.005	0.036	0.007	0.063**	0.107	0.004
R-squared	0.56	0.59	0.56	0.63	0.55	0.79
Observations	823	823	823	823	823	823
Day-of-the-month FE	Yes	Yes	Yes	Yes	Yes	Yes
Weights	42%	18%	23%	37%	9%	45%
Side	Supply	Supply	Supply	Demand	Demand	Demand

Table 6
Sector Elasticities ζ_i with PCs

In this Table, we show the sector-level elasticities ζ_i for the baseline estimation of the first stage (Table 2) and for the estimation of the first stage with three pricipal components (Table 5). The table reports the market share of the sector calculated as in Eq. 21, the sector price-elasticity ζ_i calculated from Eq. 25, and an alternative estimate of ζ_i using principal component analysis (PCA). We allocate a sector to supply (demand) if the sector is a net lender (borrower) of securities, on average, over our sample period.

Sector	**Share**	ζ	ζ with PCA
Supply			
Euro area			
Government	42%	0.64	0.61
Central bank	18%	1.28	1.34
Bank	23%	1.09	0.34
Dealer	5%	5.38	1.17
ICPF	5%	0.49	−0.32
Hedge funds	1%	−0.85	−1.99
Non-euro area			
Government	1%	−1.41	−0.47
Central bank	3%	2.50	1.65
ICPF	5%	0.48	−0.12
Demand			
Euro area			
OFI	1%	−0.89	0.02
MMF	0%	0.11	−0.67
Non-euro area			
Bank	5%	1.60	2.44
Dealer	37%	0.05	0.18
OFI	9%	0.27	0.53
Hedge funds	45%	0.11	0.21
Aggregate	200%	1.20	0.97

Appendix

A.1 SFTDS De-duplication Algorithm

This appendix is intended to provide a detailed overview of the de-duplication algorithm employed for the analysis of SFTDS transaction data and report the external validity cross-checks conducted to ensure the data quality of this deduplication.

A.1.1 Regrouping observations belonging to the same economical transaction

The first step consists of a light cleaning of the raw SFTDS data, namely the exclusion of transactions above EUR 50 bn and below EUR 1000 and the exclusion of trades which do not report at least volumes, counterparties, collateral and rates. We focus on new repo transactions (`report_type = 'New'` and `trade_type='RpTrad'`). SFTDS identifies transactions based on Unique Transaction Identifiers ('UTI') that are usually not directly exploitable. For instance, UTIs keep information about the direction of trade ('B' or 'S'), or use sometimes incremental numbering depending on the reporting counterparty. We first clean the UTI by identifying these patterns, namely removing leading and trailing '0','X','B','S','BO' and 'SO' from the UTI strings, and removing the numbering in the last 5th and 6th digits for the UTI strings registered by LCH. To further ensure that we indeed regroup observations related to the same transaction, we check the exact matching of the trade characteristics, using a MD5 hash function (from the R package `openssl`) that recreates a unique identifier from the following variables:

Variable Name	**Description**
`exctndt`	Execution date
`valdt`	Value date
`term_fxd_mtrtydt`	Maturity date
`trunc_prncplamt_valdtamt_eur`	Truncated principal amount
`round_intrst_rate_derived_sdw`	Repo rate, rounded at the second decimal
`assttp_scty_id`	ISIN of the collateral

The creation of the hash allows us to detect double-counting that even cleaned UTI cannot flag. This is the case, for instance, for interoperability, which trades. Interoperability refers to arrangements between two CCPs to intermediate a single repo transaction. They allow clearing members of one CCP to centrally clear trades carried out with members of another CCP, without needing to be a member of the second CCP. LCH and CC&G, both members

of LSE Group have an interoperability agreement for Italian government bonds collateral. In SFTDS, an interoperability transaction involving LCH and CC&G and two final non-CCP counterparties in EEA would feature 6 observations, but with different UTIs. Observable patterns in the UTI would not allow to match the observations for a given transaction. We then proceed by creating the hash as explained above, check whether both LCH and CC&G appear as reporting agents or counterparties, override the UTI with the hash when it is the case.

A.1.2 Building repo positions from transaction flows

After cleaning all transactions, we compute the outstanding amounts of repos for each investor. For each unique transaction identifier (UTI), we calculate the number of days between the settlement date and the maturity date, and expand the dataset accordingly by adding one row for each business day during which the repo remains outstanding. Since accurate information on the maturity dates of open repos is unavailable, we exclude these transactions from the analysis.

A.2 Bloomberg Sector Classification

Our sector classification uses a mix of Bloomberg sector classifications (Sector 4 and Fund type) and ESA 2010 classification. Investors (LEIs) are allocated to hedge funds when fund type is "Hedge Fund", "FCP", "OEIC", "SICAV","Open-End Fund", "Closed-End Fund", "ETF", "Fund of Funds", "Private Equity Fund", "Unit Trust", "Variable Annuity", "Investment Trust" or sector level 4 is "Private Equity".

Investors are allocated to ICPF using ESA classification "S129" and "S128" we also add investors with Bloomberg sector 2 classification "Insurance" or fund type "Open-End Pension".

We use ESA sector S123 for MMFs. For NFC we use Bloomberg sector 2 "Utilities", "Consumer Discretionary Products", "Consumer Discretionary Services", "Health Care", "Industrial Services", "Materials", "Media", "Oil & Gas", "Real Estate", "Retail & Wholesale - Staples", "Retail & Whsle - Discretionary", "Software & Tech Services", "Telecommunications", "Consumer Staple Products".

For Central Bank we use the Bloomberg industry subgroup "Central Bank". For Government we use sector 2 "National", "Regional & Local" and ESA sector classification S1311 and

S1312. We also allocate to Government all the transactions where the counterparty is unreported, see Section 3.

For OFI, we use Bloomberg sector 4 "Consumer Finance", "Instl Trust, Fiduciary & Custody", "Commercial Finance", "Diversified Banks", "Institutional Brokerage", "Investment Companies", "Investment Management", "Other Financial Services", "Unknown", "Wealth Management", "Mortgage Finance", "Security & Cmdty Exchanges".

MFI are investors belonging to Bloomberg sector 4 "Banks". The remaining non-allocated investors are hedge funds. Some LEIs of international institutions and banks are re-allocated manually, we refer to our online code for details.

A.3 Estimation of the demand elasticities

The elasticities ζ_i are estimated by granular IV (OGIV). For each sector i, the estimator takes the canonical IV form

$$\hat{\zeta}_i \;=\; -\,\frac{\hat{\mathbb{E}}\Big[\Delta q_{i,t}\,\hat{u}_{S(-i),t}(\hat{\zeta})\Big]}{\hat{\mathbb{E}}\Big[\Delta p_t\,\hat{u}_{S(-i),t}(\hat{\zeta})\Big]}, \qquad \hat{u}_{S(-i),t} = \sum_{j\neq i} S_j\,\hat{u}_{j,t}, \tag{A-28}$$

where the instrument $\hat{u}_{S(-i),t}$ is the size-weighted aggregate of the *other* sectors' idiosyncratic demand shocks. Because these shocks are themselves functions of the elasticities, the estimator is solved by iteration. Starting from an initial guess for $(\hat{\zeta}_1,\ldots,\hat{\zeta}_N)$ (e.g., the OLS coefficients), we recover each sector's residual by re-arranging Eq. 17,

$$\hat{u}_{i,t}(\hat{\zeta}_i) \;=\; \Delta q_{i,t} - \widehat{\overline{\Delta q_i}} - \hat{\lambda}_i\eta_t + \hat{\zeta}_{ii}\Delta p_t, \tag{A-29}$$

form the instrument $\hat{u}_{S(-i),t}$ for each sector i, and update the elasticities via Eq. (A-28). We repeat until the estimates converge. The iteration is necessary because the first-stage residual $\hat{\varepsilon}_{i,t} = u_{i,t} - \zeta_j\Delta p_t$, which removes macro factors and fixed effects, is still affected by the endogenous price response $-\zeta_i\Delta p_t$; adding back the current guess $\hat{\zeta}_i\Delta p_t$ removes this contamination exactly at the fixed point $\hat{\zeta}_i \to \zeta_i$.

A.4 Partial and Total Elasticities: Mapping the Model to a Multi-Asset Demand System

This appendix maps the model of Section 2 into the multi-asset demand system of Chaudhary et al. (2024) and shows that the conditional schedules of Propositions 4 and 5 are the *total* elasticities implied by that mapping. The exercise formalizes the discussion of the empirical estimates in Section 5: the hedge fund's partial own- and cross-price elasticities in the repo market are of equal magnitude, specialness passes through fully into cash bond prices in equilibrium, and the total repo elasticity that survives the offset of these two forces is the cash-market depth parameter η.

Consider two assets: asset 1 is the repo contract, whose price is specialness χ, and asset 2 is the cash bond. Because the futures rate r_F is exogenous, the basis $x = r_F - r_T$ moves one-for-one with the bond's price, and we use x as the price of asset 2. Following Chaudhary et al. (2024), a sector's demand for asset 1 is

$$q_1 = -\zeta_{11}\,\chi + \zeta_{12}\,x + u_1, \tag{A-30}$$

where ζ_{11} is the partial own-price elasticity (the response to specialness holding the bond price fixed), ζ_{12} is the partial cross-price elasticity (the response to the bond price holding specialness fixed), and u_1 is a demand shifter. Imposing market clearing for asset 2 expresses its price as a function of χ, with slope given by the passthrough coefficient $P_{1\to 2} \equiv \partial x/\partial \chi$, and collapses the system to a single-asset demand curve with *total* elasticity

$$\tilde{\zeta}_1 \;=\; \zeta_{11} - \zeta_{12}\,P_{1\to 2}, \tag{A-31}$$

which measures the equilibrium response of q_1 to a change in specialness after the bond price has adjusted.

Under perfect competition, the hedge fund's demand taking both prices as given is zero when $x < \chi + mk$ and unbounded when $x > \chi + mk$. Its partial elasticities are therefore not finite objects. To characterize them, we allow for the fund to have finite arbitrage slope κ,

$$|D_H| = \kappa\,(x - \chi - mk)\,, \tag{A-32}$$

which arises when funds bear convex holding costs or internalize their price impact, as would be the case in a model with imperfect competition. The competitive benchmark of

Section 2.7 is the limit $\kappa \to \infty$. Demand (A-32) delivers the partial elasticities:

$$\zeta_{11}^H = \zeta_{12}^H = \kappa. \tag{A-33}$$

The two partial elasticities are the same because the trade is an arbitrage: only the spread $x - \chi$ enters the fund's profit, so a change in either price (holding the other fixed) moves demand by the same amount. If $\kappa \to \infty$, both partial elasticities diverge together: a large cross-price elasticity necessarily comes with an equally large own-price elasticity.

Cash-market clearing, equation (14), gives $x = \bar{x} - |D_H|/\eta$. Differentiating (A-32) and (14) with respect to χ and solving the two equations,

$$\frac{\partial |D_H|}{\partial \chi} = -\frac{\kappa\eta}{\kappa + \eta}, \qquad P_{1\to 2} = \frac{\partial x}{\partial \chi} = \frac{\kappa}{\kappa + \eta} \xrightarrow{\kappa\to\infty} 1. \tag{A-34}$$

Equivalently, in the notation of Chaudhary et al. (2024), $P_{1\to 2} = \zeta_{S,21}/\zeta_{S,22}$, where $\zeta_{S,21}$ is the aggregate cross-price response of cash-market demand to specialness and $\zeta_{S,22}$ the aggregate own-price response of the cash market. In our model, $\zeta_{S,21} = \kappa$: the hedge fund is the only investor active in both markets, so it alone carries the cross term. Similarly, $\zeta_{S,22} = \eta + \kappa$ sums preferred-habitat depth and the fund's own response. In the competitive limit, passthrough is complete: the zero-profit condition of Proposition 3 implies $dx = d\chi$, so a change in the cost of borrowing the bond is capitalized fully into its price, as in Duffie (1996). The passthrough is an equilibrium object that obtains from market clearing.

Combining (A-31), (A-33), and (A-34),

$$\tilde{\zeta}_1 = \zeta_{11}^H - \zeta_{12}^H P_{1\to 2} = \kappa - \frac{\kappa^2}{\kappa + \eta} = \frac{\kappa\eta}{\kappa + \eta} \xrightarrow{\kappa\to\infty} \eta. \tag{A-35}$$

Under perfect competition, both terms of the total elasticity grow without bound. Their difference, however—the response of the fund's repo quantity to a change in specialness after the bond price has adjusted—remains finite and equals the cash-market depth η: the fund's equilibrium response is limited not by its own price sensitivity but by the price impact its trading exerts through cash-market clearing. Equation (A-35) is exactly the slope of the conditional repo schedule in Proposition 5: the conditional schedule, obtained by substituting cash-market clearing into the fund's problem, is the model's counterpart of the total elasticity obtained by collapsing the two-asset system.

Three implications follow. First, since $\kappa\eta/(\kappa+\eta) \le \min\{\kappa, \eta\}$, the fund's total repo elasticity can never exceed the depth of the cash market, no matter how aggressively it trades. Second, the baseline knife-edge of Proposition 1 is nested as the degenerate limit $\eta \to 0$, in which the

total elasticity vanishes even though both partial elasticities remain unbounded: a measured elasticity near zero does not mean that the fund ignores prices. Third, the same derivation with the two markets interchanged—substituting repo-market clearing (11) into (A-32)—yields a total cash-market elasticity of $\kappa/(1+\kappa\phi) \to 1/\phi$, reproducing Proposition 4 and the cross-market asymmetry documented in Section 6.

Collapsing the system also transforms cash-market demand shocks into repo-market demand shifters. Substituting cash-market clearing into (A-32) gives, for finite κ,

$$|D_H| = \frac{\kappa\eta}{\kappa+\eta}\left(\bar{x} - \chi - mk\right), \qquad \frac{\partial |D_H|}{\partial \Delta_S} = \frac{\kappa}{\kappa+\eta} \xrightarrow{\kappa\to\infty} 1, \tag{A-36}$$

since the natural basis $\bar{x}$ carries the demand premium Δ_S/η. A cash-market imbalance therefore shifts the fund's repo demand one-for-one while leaving its price sensitivity unchanged. The coefficient $\kappa/(\kappa+\eta)$ is precisely the loading $\zeta_{12}/\zeta_{S,22}$ on the aggregate demand shifter that arises in Chaudhary et al. (2024) when a multi-asset system is collapsed to a single asset. This result provides the formal basis for the interpretation of the empirical findings in Section 5: time variation in specialness reflects shifts in the quantity of cash-market imbalance that hedge funds intermediate, not changes in their price elasticity.

Table A-1
Securities holdings and borrowing and lending flows on the repo market, by sector

In this table, we show the amount of securities held, and the amounts of securities lent and borrowed on the repo market. We report the average by sector at the yearly level, and restrict the sample to German government bonds and repo transactions that use them as collateral. The amounts borrowed and lent on the repo market represents daily flows. All quantities are expresse in billion of euros. eurosystem borrowing and lending figures are limited to cash-collateralized transactions. Holdings of dealers, banks (monetary financial institutions, MFI), and other financial institutions (OFI) are reported as one, as the holding dataset does not distinguish between them. Non-euro area holdings are computed as a residual of Bund. We exclude money market funds, households and non-financial corporations, as their participation to the repo market is minimal. Our sample extends from January 2021 to June 2024. Holdings data are from SHSS, repo volumes from SFTDS, and eurosystem repo lending data are from MOPDB. The repo contract sample includes only fixed-rate, single-collateral–backed repos.

Year	Area	Sector	Holdings	Borrowed	Lent
2021	**euro Area**	ICPF	121	0.1	0.1
		Bank	11	26.5	30.4
		OFI		1.8	0.7
		Dealer		35.0	33.3
		Hedge funds	117	0.5	0.1
		Government	172	3.9	11.6
		Central bank	554	0.0	3.0
	Non-euro Area		554	55.0	43.3
2022	**euro Area**	ICPF	121	0.1	0.4
		Bank	3	32.0	33.0
		OFI		3.6	3.9
		Dealer		53.4	62.5
		Hedge funds	148	0.7	1.2
		Government	169	3.7	19.1
		Central bank	655	0.0	9.0
	Non-euro Area		534	95.6	60.0
2023	**euro Area**	ICPF	141	0.0	0.6
		Bank	3	35.9	39.3
		OFI		4.4	5.7
		Dealer		57.3	57.8
		Hedge funds	178	1.7	2.0
		Government	203	6.3	19.1
		Central bank	621	0.0	4.0
	Non-euro Area		616	91.0	67.8
2024	**euro Area**	ICPF	144	0.2	0.4
		Bank	3	23.3	34.1
		OFI		3.7	3.4
		Dealer		48.6	45.1
		Hedge funds	178	1.4	0.8
		Government	208	6.6	11.6
		Central bank	577	0.0	0.9
	Non-euro Area		700	66.0	53.9

Table A-2
First stage of Optimal-GIV - other sectors

In this table, we show the coefficients resulting from the first stage of the O-GIV where sector i net securities borrowing volumes in the repo market are residualised against a broad set of common factors and sector fixed effects, see equation 22. The dependent variable is the percentage deviation of sector i ' net securities borrowing volumes from its 3-months moving average. The same transformation is applied to common factos, which are listed in Section 3. We include sectors whose outstanding positions make up more than 5% of the total. These sectors cumulatively represent 17% (9%) of the supply (demand) in the market. We allocate a sector to supply (demand) if the sector is a net lender (borrower) of securities, on average, over our sample period. *, **, and *** indicate that parameters are significantly different from zero at the 10%, 5%, or 1% level, respectively. Significance is based on autoregression and heteroskedasticity-robust Newey-West standard errors.

	ICPF EA	Hedge funds EA	MMF EA	OFI EA	Dealer EA	Central bank non-EA	Government non-EA	Bank non-EA	ICPF non-EA
OIS1M DFR spread	−0.005	−0.014	0.048*	0.020	0.001	−0.004	0.003	0.016	0.000
3M1M OIS spread	−0.037	0.042	0.050	−0.083	0.186***	0.107**	0.064*	0.043	−0.006
1Y3M OIS spread	−0.072*	−0.052	−0.060**	0.065***	0.042	0.002	0.009	−0.172***	−0.044**
KFW vs DE spread	0.357*	0.531*	0.011	0.490*	0.299	−0.149	−0.319	−0.050	0.031
DE vs FR spread	−1.092	−0.554	1.673*	0.013	−1.324	−2.183***	0.781	2.526***	−0.658**
OIS 1M	−0.015	0.003	−0.203**	−0.013	0.388***	0.040	0.078	−0.142**	−0.020
Slope YC	0.213*	−0.197	−0.596***	0.360	−0.283*	0.037	0.025	0.205	0.022
OID skew	0.478**	−0.368*	0.084	−0.309	0.513*	−0.147	−0.157	−0.388	−0.172
SMOVE	−2.602**	−1.466	−1.586*	−0.751	−2.482**	1.947**	1.387*	0.967	−0.475
Bund vol prem	−0.027	0.315**	−0.014	0.029	0.075	0.069	0.233	0.061	−0.019
NSS dev	−0.537	0.781	−0.324	−0.427	2.024***	0.566	0.402	−1.512***	0.694***
Bund OpenInt	0.054	−1.060	−0.266	−4.607***	1.559	−1.143	1.455	−3.741**	−0.939**
Cross-ccy basis	0.108	−0.009	0.090	0.471***	0.452***	0.082	0.007	−0.434***	−0.054
Duffie	−0.032	−0.001	0.031	−0.002	−0.014	−0.019	0.013	0.046	−0.009
Free float	0.441	16.224***	9.916*	−10.341*	14.309**	8.811**	−8.325	−13.369***	4.426***
ECB holdings	3.822	1.689	7.096	−0.009	8.229	−7.646	−16.932***	−11.275	3.404
DFA holdings	0.529	−13.165***	−6.940***	−7.235	−24.268***	−0.797	5.959	7.686***	−3.684***
Constant	−0.311	−0.021	0.085	0.224	−0.196	−0.349*	−0.134	−0.232	−0.145*
R-squared	0.26	0.23	0.23	0.20	0.29	0.22	0.19	0.34	0.39
Observations	823	823	823	823	823	823	823	823	823
Weights	5%	1%	0%	1%	5%	3%	1%	5%	5%
Side	Supply	Demand	Demand	Demand	Supply	Supply	Supply	Demand	Supply

Figure A-1. The Unreported Sector

In this figure, we show the characteristics of cleared trades for which one of the two counterparties is missing (the "missing" sector). In Panel A, we show the net securities lending of the unreported sector as a function of the asset's specialness. The binscatter is based on bond-day level observations. Panel B shows the aggregate net securities lending (securities lent minus securities borrowed) by the missing sector at the daily level. Data are sourced from SFTDS and are limited to fixed-rate single-collateral–backed repo transactions for German central government bonds, for the 2021–2024 sample.

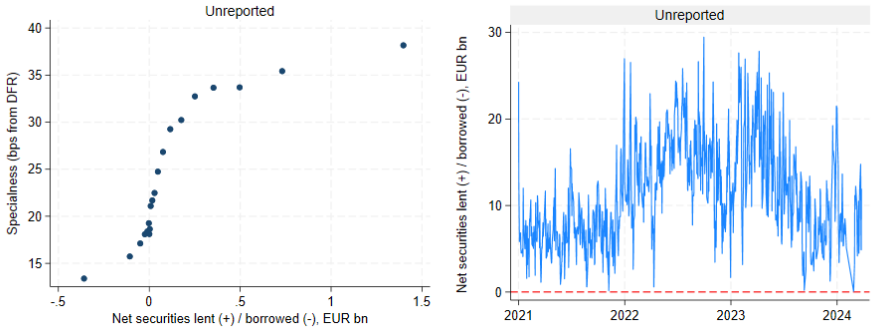

Figure A-2. Implications of SFTDS reporting rules

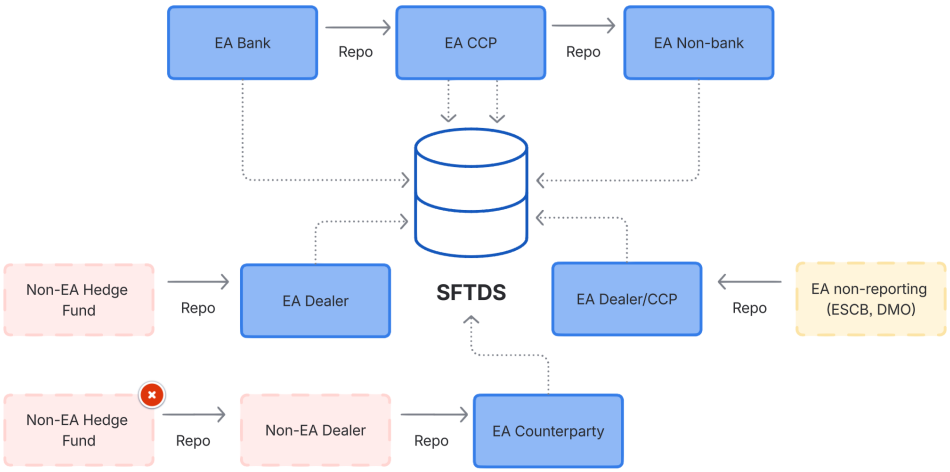

Figure A-3. Volumes of securities borrowed by location of hedge funds (IF)

In this figure, we show the distribution of securities borrowed by hedge fund based on the country of incorporation of the fund. Data are sourced from SFTDS and are limited to fixed-rate single-collateral–backed repo transactions for German central government bonds, for the 2023 sample.

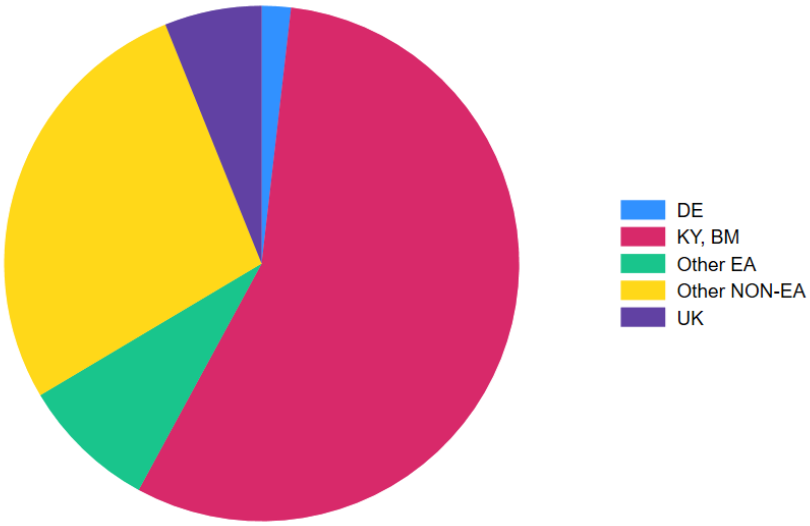

Figure A-4. Idiosyncratic shocks

In this figure, we show the daily time series of the residuals (i.e. idiosyncratic shocks) from the first stage of the O-GIV where sector i net securities borrowing volumes in the repo market are residualised against a broad set of common factors and sector fixed effects, see equation 22 and Table 2 and Table A-2. The dependent variable is the percentage deviation of sector i ' net securities borrowing volumes from its 3-months moving average. The same transformation is applied to common factos, which are listed in Section 3. We aggregate the shocks from all sectors by taking the daily sum.

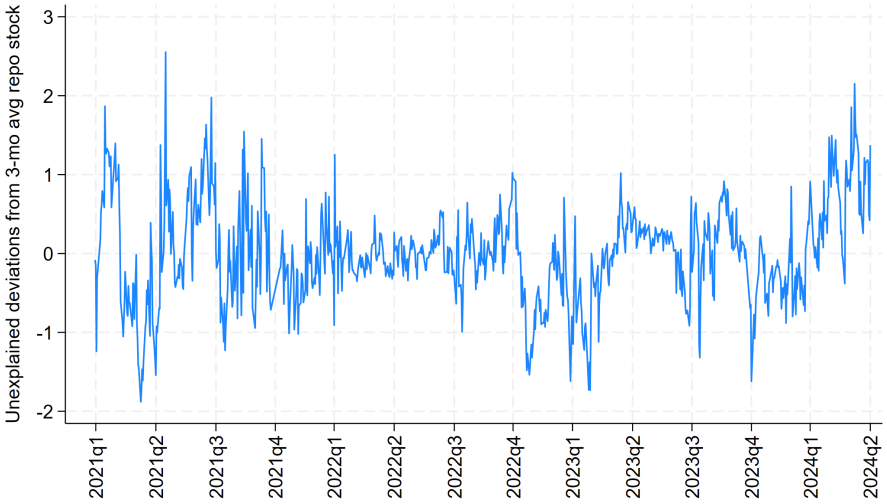

Figure A-5. Actual and predicted specialness deviations from 3-months moving average

In this figure, we show the strong positive correlation between the predicted values from our model $\Delta\hat{p}_t$ - calculated from Eq. 26 - and observed deviations Δp_t as evidence of the good fit of our empirical specification.

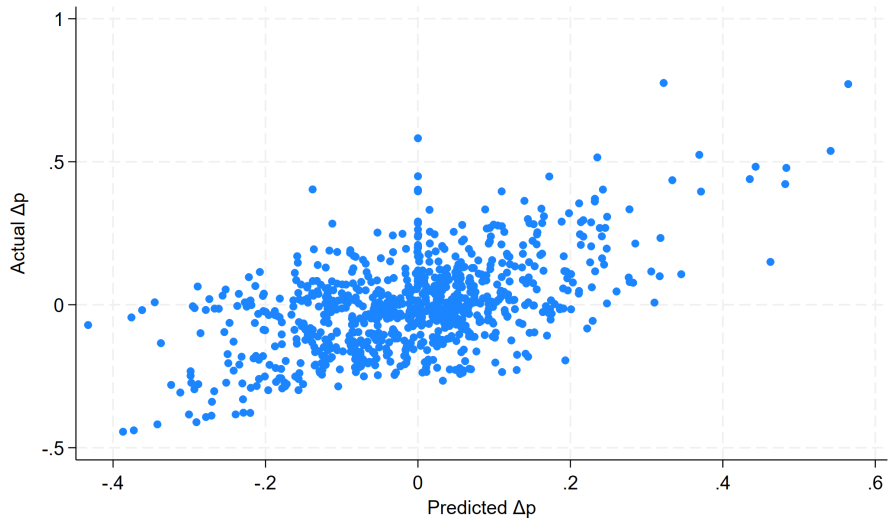

Figure A-6. Decomposition of repo specialness contributions by common factors

This figure shows the decomposition of the changes in the specialness spread among common factors. We define changes in specialness as the relative deviation from the three-month moving average, and changes in quantities accordingly. We average the daily contributions at the quarter level, and classify the factors into monetary policy, repo demand, repo supply, and bond supply factors. The factor decomposition for factor $\eta(k)$ is calculated as $\Delta\hat{p}_{\eta(k),t} = \frac{1}{\hat{\zeta}_S} S_i \hat{\lambda}_i(k) \eta_t(k)$, for and day-$t$, where $\hat{\zeta}_S$ is reported in Table 3. The sample period goes from March 2021 to June 2024 and the sample consists of repos with underlying collateral a German government bond. Repo data are from SFTDS, and eurosystem repo lending data are from MOPDB

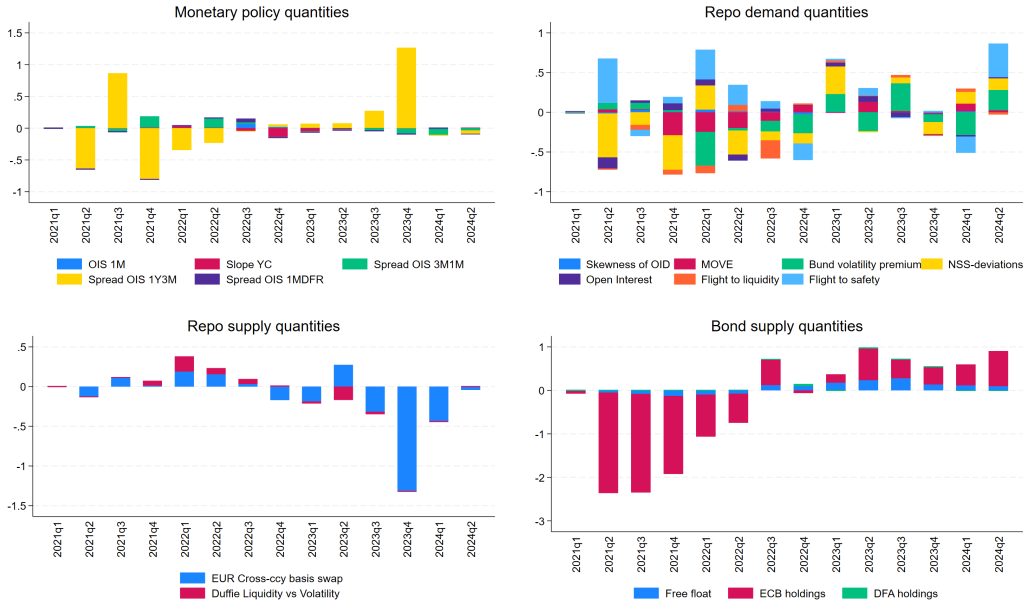

Figure A-7. Decomposition of repo specialness contributions of EA intermediaries at quarter-end

In this figure, we show the cyclicality of sectoral contribution to specialness within a quarter. That is, we calculate the contribution $\Delta\hat{p}_{i,t} = \frac{1}{\hat{\zeta}_S} S_i(\hat{\Delta q}_i + \hat{\lambda}_i \eta_t + \hat{u}_{i,t})$, for sector-$i$ and day-t, where $\hat{\zeta}_S$ is reported in Table 3, aggregate it by sector type, and average it by the number of days to the end of the quarter. We group investors in euro-Area intermediaries (banks and dealers) and the rest. We show the results for the period from March 2022 to June 2024. In this sample, euro-area intermediaries exited from special exemption period, which allowed them to exclude sizeable central bank exposures from the calculation of the leverage ratio, contributing to lower balance sheet costs. The sample consists of repos with underlying collateral a German government bond. Repo data are from SFTDS, and eurosystem repo lending data are from MOPDB.

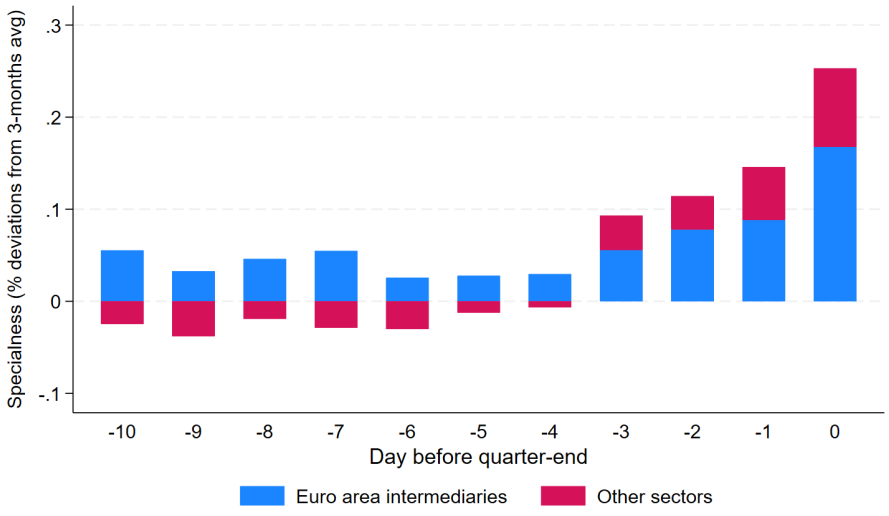

Figure A-8. Elasticity by sector: estimation with Bund yield

In this figure we compare the price-elasticity ζ_i for the largest participants to the repo market when the yield of the 10 year benchmark yield for the Bund is added to the macro-controls in the first stage to account for cash market clearing. We exclude sectors whose outstanding positions make up less than 5% of the total. These sectors cumulatively represent 83% (91%) of the supply (demand) in the market. The elasticities are computed from Eq. 25. Our data includes repo transactions that use German government bonds as collateral for the March 2021–June 2024 period. Repo data are from SFTDS, and eurosystem repo lending data are from MOPDB

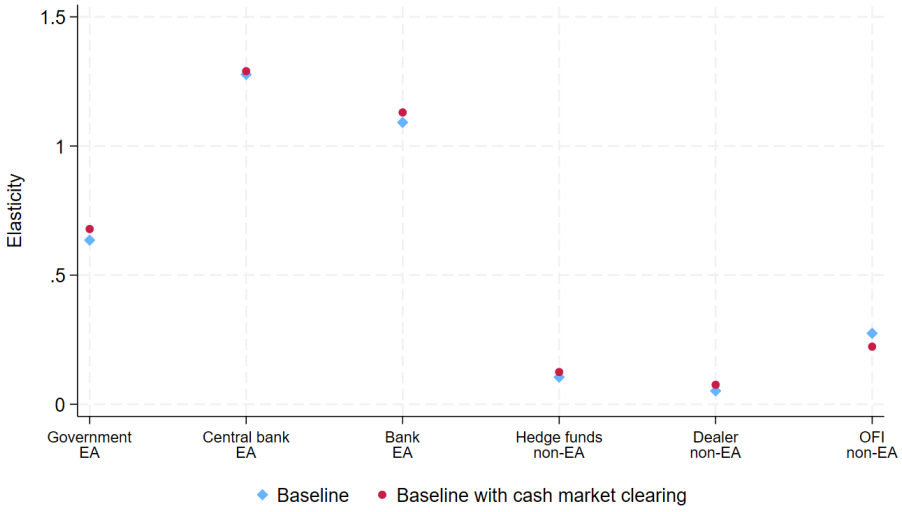

Figure A-9. Elasticity by sector: sample splits

In this figure we compare the price-elasticity ζ_i for the largest participants to the repo market using two sub-samples. The maturity split is based on the remaining maturity of the undelying Bund (i.e. above and below 5 year remaining maturity) and the specialness split is based on the share of the outstanding amount of the underlying bond held at the Eurosystem (i.e. above and below daily median holdings). We exclude sectors whose outstanding positions make up less than 5% of the total. These sectors cumulatively represent 83% (91%) of the supply (demand) in the market. The elasticities are computed from Eq. 25. Our data includes repo transactions that use German government bonds as collateral for the March 2021–June 2024 period. Repo data are from SFTDS, and eurosystem repo lending data are from MOPDB

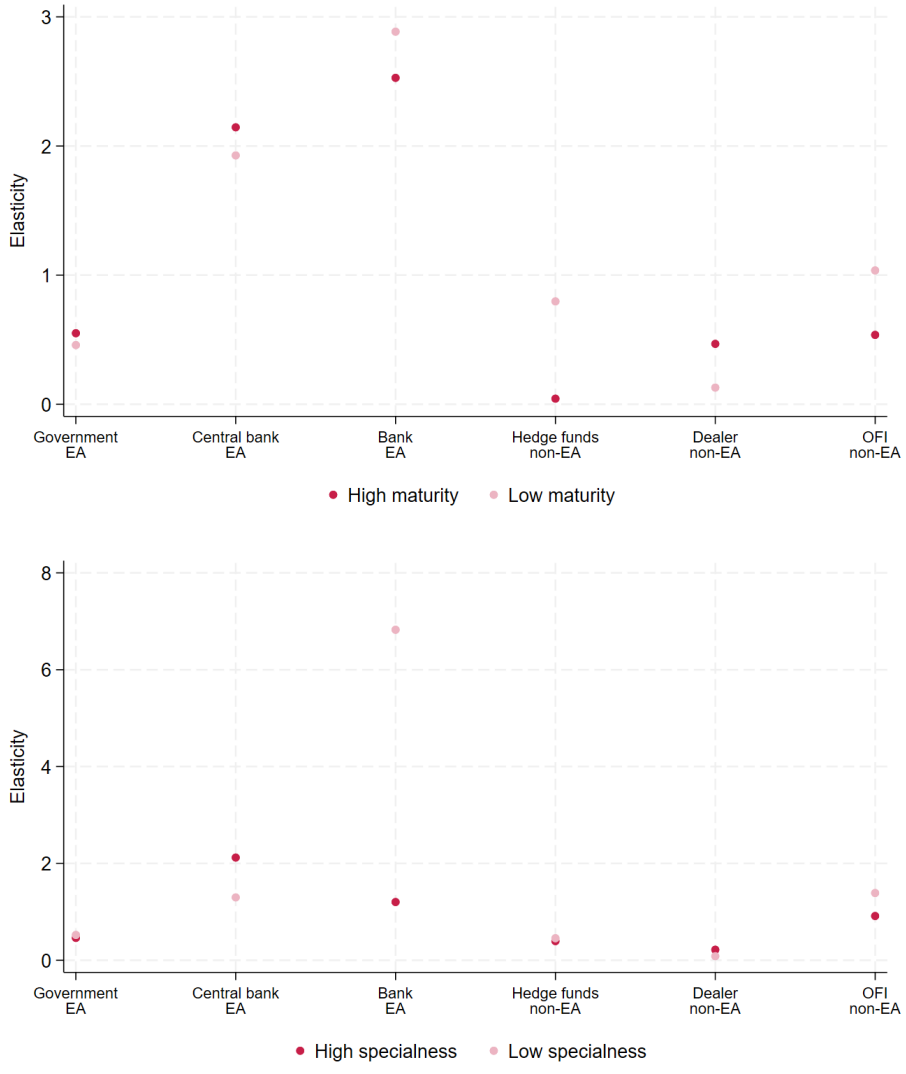

Figure A-10. Elasticity by sector: alternative transformation of quantities and prices

In this figure we compare the price-elasticity ζ_i for the largest participants to the repo market using two alternative transformation of prices (quantities).The baseline model' prices (quantities) are expressed as the percentage deviation from the 3-months moving average, see Eq. 20. We conduct two robustness tests by using the 1-month and 6-months percentage deviation. We exclude sectors whose outstanding positions make up less than 5% of the total. These sectors cumulatively represent 83% (91%) of the supply (demand) in the market. The elasticities are computed from Eq. 25. Our data includes repo transactions that use German government bonds as collateral for the March 2021–June 2024 period. Repo data are from SFTDS, and eurosystem repo lending data are from MOPDB

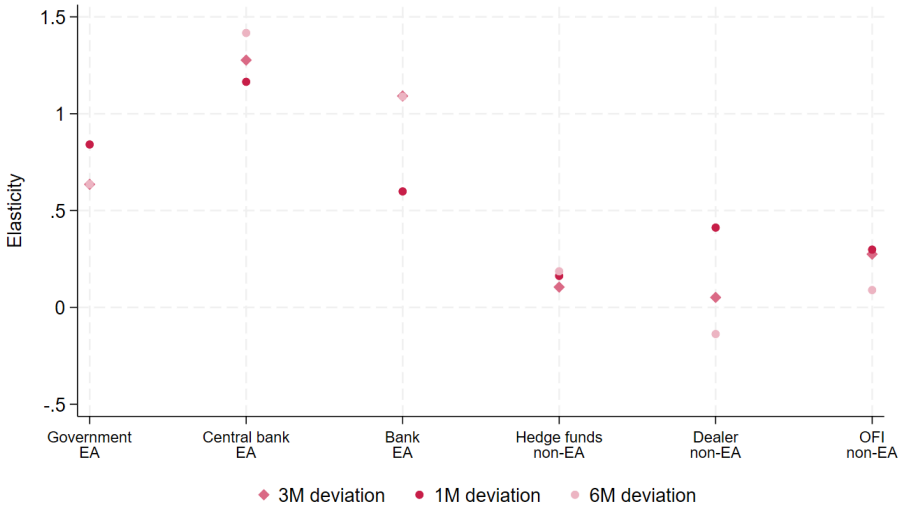

Figure A-11. Elasticity by sector: endogeneity of trading volumes

In this figure we compare the price-elasticity ζ_i of the baseline model and a model where we lag the moving average of prices (quantities) variables by 4 months, see Eq. 20. We exclude sectors whose outstanding positions make up less than 5% of the total. These sectors cumulatively represent 83% (91%) of the supply (demand) in the market. The elasticities are computed from Eq. 25. Our data includes repo transactions that use German government bonds as collateral for the March 2021–June 2024 period. Repo data are from SFTDS, and eurosystem repo lending data are from MOPDB

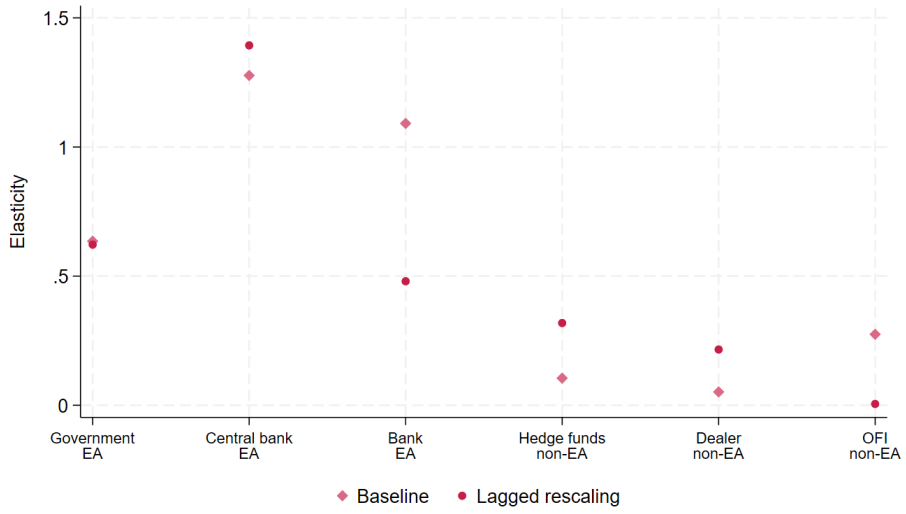

Stock Volatility and Index Fund Ownership

Shmuel Baruch* and Xiaodi Zhang†‡

Abstract

Low-cost index funds have broadened market participation, raising the concern that democratized access may make markets more fragile. We test this concern using the documented positive relation between index fund ownership and individual stock volatility. If that relation is driven by new investors entering through index products, their flows should increase the systematic component of stock volatility and broad-index volatility. We find little evidence of either. The relation between index ownership and volatility is concentrated in idiosyncratic stock volatility, and index-fund market share is not significantly related to S&P 500 or DJIA volatility. The evidence rules against a volatility footprint from broadened index-fund participation and is more consistent with internal reallocation by existing investors from individual-stock strategies toward index-centric strategies.

JEL Classification: G12, G14, G23

Keywords: Index fund ownership, Stock volatility, Index Volatility, Passive investing, Index funds

*Baruch is from Price College of Business, University of Oklahoma, Norman, OK 73019. E-mail: baruch@ou.edu

†Zhang is from the School of Finance, Shanghai University of Finance and Economics, Shanghai 200433, China. E-mail: zhang.xiaodi@sufe.edu.cn.

‡The authors declare that they have no relevant or material financial interests that relate to the research described in this paper. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

1 Introduction

As the index fund industry has grown into a widely used investment vehicle, it has attracted substantial participation.[1] Such activity can arise from two distinct sources: external flows to index funds (investors entering and exiting the market via index products, such as 401(k)s) and internal shifts to index strategies (existing participants reallocating between individual stock selection and index-centric approaches, such as factor investing).

An established empirical finding, documented by Ben-David, Franzoni, and Moussawi (2018) and Coles, Heath, and Ringgenberg (2022), is the positive link between individual stock return volatility and index fund ownership. What remains unsettled is whether this relation reflects external flows that expand the set of investors participating through index products, or instead reflects a change in how a given investor base participates in price discovery. This paper focuses on the testable implication of the external-flow explanation.

Our diagnostic logic follows from a basic economic observation about how index-linked demand is implemented. Changes in demand for index exposure are ultimately expressed through trading (and rebalancing) in the index basket. Consequently, index-linked trading pressure is concentrated at the aggregate (index) level. Thus, if the stock-level volatility effect is primarily caused by external flow shocks that inject noise through index products, that disturbance should not be confined to idiosyncratic stock movements: it should, first and foremost, load on the systematic component of constituent returns and, in the aggregate, on index volatility itself.

By contrast, internal reallocation has a different footprint. When sophisticated traders shift back and forth between direct stock selection and index-centric positions, they continue to participate in the price discovery of the market portfolio: explicitly when they trade index

[1] Recent evidence in Chinco and Sammon (2024) further shows that passive ownership can extend beyond registered index funds through internal indexing and related index-linked strategies.

instruments, and indirectly when they trade the underlying constituents. What changes is their direct participation in the price discovery of individual stocks. With fewer (or less intensive) stock-specific traders, stock-level prices can become noisier relative to fundamentals even if market-wide price discovery remains anchored. In this case, a rise in indexing can show up primarily as higher idiosyncratic volatility of constituents, with little or no corresponding change in either systematic volatility or index volatility.

This logic motivates a diagnostic empirical strategy: we examine not only whether index fund ownership predicts total stock volatility, but also *which component* of volatility moves (systematic vs. idiosyncratic) and whether the volatility of the aggregate index moves at all. These implications require no knife-edge assumptions about functional forms or distributions; they come from the institutional fact that indexing aggregates trades and from the economic distinction between (i) shocks that enter through the index basket and (ii) changes in stock-level participation in price discovery.

Our empirical strategy exploits the rich cross-sectional and time-series variation in index fund ownership generated by overlapping index memberships and heterogeneous index demand, rather than relying on infrequent index reconstitutions.[2] Using a large panel of U.S. stocks from 1980–2022, we first confirm the established positive relation between index fund ownership and individual-stock volatility.

We then decompose each stock's annual volatility into systematic and idiosyncratic components using a market-model regression and the corresponding sums-of-squares decomposition. The key diagnostic result is that the positive relation with indexing is concentrated in idiosyncratic volatility: higher lagged index fund ownership is positively related to subsequent idiosyncratic volatility, while we do not detect a statistically significant relation with systematic volatility.

[2]For example, Apple and Nike both belong to the price-weighted DJIA and the value-weighted S&P 500, but only Apple belongs to the QQQ. Money flowing in and out of the QQQ thus creates differences between the two stocks, which, after controlling for other explanatory variables, enable us to isolate the impact of changes in index fund ownership on stock volatility.

Finally, we test the aggregate implication directly. If external flows through index funds were the primary source of the stock-level volatility effect, we would expect a corresponding positive relation between index-fund market share and the volatility of broad indices. In annual time-series regressions for the S&P 500 and the DJIA (and their flagship ETFs), we find no such relationship.[3] Taken together, these results are difficult to reconcile with external flows as the primary driver of the stock-level volatility effect.

To keep the paper self-contained and to connect the above diagnostic to a transparent price-discovery mechanism, we present a simple model in the spirit of Treynor (1987). The model shows that, because demand shocks enter through the index basket, external flows should be reflected not only in constituent-stock volatility but also in systematic and index-level volatility. Moreover, the model predicts that, for a broad value-weighted index such as the S&P 500, an external-flow effect large enough to increase the variance of a median constituent should imply a much larger increase in the variance of the index fund. In our S&P 500 calibration, the model-implied index-fund variance increase is several orders of magnitude larger than the corresponding increase for the median constituent. Empirically, the predicted footprint is absent: index-fund market share is not significantly related to broad-index volatility; the estimated relation is far smaller than the model-implied benchmark from the S&P 500 calibration; and index ownership is not significantly related to systematic stock volatility.

The model clarifies what our tests can and cannot identify. They do not reveal the identities of marginal investors, but they do test whether the stock-level volatility relation is accompanied by the systematic and index-level volatility responses implied by external flows. The absence of statistically significant responses rules against the external-flow interpretation. At the

[3]Barberis, Shleifer, and Wurgler (2005) document a positive link between indexing and stock comovement. An increase in comovement, however, does not necessarily imply higher market-level volatility. Return comovement (R^2 or pairwise correlation) reflects the relative share of systematic variance versus idiosyncratic noise. As shown in Baruch and Zhang (2022), when passive investing reduces stock-specific price discovery without altering market-wide information aggregation, stock comovement increases while aggregate market volatility remains unchanged.

same time, the evidence is consistent with an internal-reallocation mechanism.

Our paper contributes to the broad literature on the consequences of index investing. For a comprehensive survey of this field, we refer the reader to Allard, Farkas, Pedio, Rubio, and Tonks (2025). Our work is specifically situated at the intersection of two key research areas.

First, we add to the literature that investigates the macroeconomic and market-wide consequences of index trading. For example, Sushko and Turner (2018) show that ETF trading can affect the returns of the aggregate bond index, and Anadu, Kruttli, McCabe, and Osambela (2020) examine how passive investment may affect broader financial market stability. Our contribution is to test whether the documented stock-level volatility relation has the footprint of external flows into indexed products. This distinction matters because external flows should affect both constituent stocks and the aggregate index, whereas internal shifts can increase stock-level volatility without producing a comparable increase in index-level volatility, a possibility formalized in the rational-expectations model of Baruch and Zhang (2022). The empirical pattern we document therefore speaks not only to whether indexing matters, but also to how it matters.

Second, our study contributes to the literature on equity mutual fund flows. While this literature often focuses on flow-performance sensitivity (see Dannhauser and Pontiff (2025) for a comprehensive study and Hortaçsu and Syverson (2004) for a focused analysis of S&P 500 funds), our findings offer a new perspective on what fund-flow explanations should imply in prices. We do not need to observe the identity of the marginal investor to test the external-flow volatility story, because that story has an observable implication: index-flow pressure should move the systematic component of constituent volatility.

2 A Model of the Price Discovery Process in the Presence of Indexers

We now develop a simple model that formalizes the diagnostic logic described in the introduction. Our approach extends the model of Treynor (1987) by incorporating both index and stock traders, which allows us to derive testable predictions. We consider an index consisting of K stocks, indexed by $k \in \{1, 2, \ldots, K\}$. Each constituent stock k trades at a current price p_k per share, with a future price v_k per share.

The index is defined as

$$\text{Index} = \frac{1}{\text{divisor}} \sum_{k=1}^{K} n_k p_k,$$

where the divisor is some constant. If n_k is the number of shares outstanding, then the index is value-weighted (e.g., S&P 500). If all the n_k's are equal (for example, $n_k = 1$), then the index is price-weighted (e.g., the DJIA).

An index fund that tracks the index issues shares. Each share trades at a current price p_F, with a future price v_F, and is backed by a portfolio of stocks. Specifically, there exists a *portfolio scaling parameter* θ such that each fund share represents ownership of a portfolio containing θn_k shares of stock k. Accordingly, the Net Asset Value (NAV) of the fund is given by

$$\text{NAV} = \theta \sum_{k=1}^{K} n_k p_k.$$

This setup ensures that the percentage return on the fund NAV equals the percentage change in the index. A reasonable choice for the portfolio scaling parameter is the reciprocal of the index's divisor, since in that case the fund's NAV equals the index.[4]

We make several simplifying assumptions: transaction costs, bid-ask spreads, and cre-

[4]The portfolio scaling parameter is not a term used in the industry. SPY's scaling parameter is 10% of the reciprocal of the divisor of the S&P 500, leading to a NAV that is exactly 10% of the index level. DIA's scaling parameter is 1% of the reciprocal of the divisor of the DJIA, leading to a NAV that is exactly 1% of the index level it tracks.

ation/redemption fees are absent, and fractional shares can be traded.

2.1 The Traders

There are I index traders, indexed by $i \in \{1, 2, \ldots, I\}$, who trade only shares of the index fund. There are also J stock traders, indexed by $j \in \{1, 2, \ldots, J\}$, who trade only shares of individual stocks.

Traders trade in proportion to what they perceive as asset misvaluation. Specifically, we assume that the ith index trader trades $\alpha(v_{Fi} - p_F)$ shares of the index fund, where v_{Fi} is the trader's subjective estimation of the future NAV of the fund share, and α is some positive constant, capturing the trading intensity. Thus, whenever an index trader perceives the fund to be undervalued, he buys, and whenever the trader perceives it to be overvalued, he sells fund shares. For simplicity, we assume that α is common to all index traders.

Similarly, the jth stock trader trades $\beta(v_{kj} - p_k)$ shares of stock k, where v_{kj} is the trader's subjective estimation of the future price of stock k, and β is the trading intensity. Thus, when a stock trader perceives a stock to be undervalued, he buys, and when he perceives it to be overvalued, he sells. Again, for simplicity, we assume that β is common to all stock traders and all stocks.

2.2 The price discovery process

Index funds exist in two forms: mutual funds and exchange-traded funds (ETFs). A mutual fund receives demand in the form of market orders (the demand is expressed in dollars or shares), and the fund managers determine the price and number of shares created or redeemed at NAV after the trading in the underlying assets has taken place. A static model like ours is not suitable to study price discovery by means of market orders. We therefore assume the

fund is organized as an ETF, and we explicitly model the authorized participants. These authorized participants are arbitrageurs who make profits from a discrepancy between the fund share price and its NAV and, in doing so, ensure that any difference between the price and the NAV is negligible.

Authorized participants may occasionally hold inventory for strategic purposes. However, for our analysis, we assume that their position is zero after the trade concludes, as they are neither investors nor speculators. Furthermore, to focus on how the price discovery transmits sentiments, we assume the authorized participants earn no arbitrage profit, and thus the fund shares trade at par with their NAV.

In our static model, there are three legs to the price-discovery process and they occur concurrently. For pedagogical reasons, we describe the process sequentially.

First, the share price of the fund, p_F, satisfies the market clearing condition:

$$0 = AP_F + \sum_{i=1}^{I} \alpha(v_{Fi} - p_F), \tag{1}$$

where AP_F represents the trading activity of authorized participants. The clearing condition states that the trade of the authorized participants offsets the index traders' order imbalance. In all likelihood, $AP_F \neq 0$. If it is positive, then the authorized participants redeem the AP_F shares from the fund issuer and receive $\theta n_k \times AP_F$ shares of each of the K stocks. They then sell those stock shares. If, on the other hand, AP_F is negative, the reverse happens and the authorized participants buy $\theta n_k \times AP_F$ shares of each of the K assets, and deliver this basket to the ETF issuer in exchange for AP_F shares of the fund. Either way, after the creation/redemption process is complete, the authorized participants hold zero shares of the index fund.

Second, the share prices of the underlying stocks satisfy the K clearing conditions

$$0 = -\theta n_k \times AP_F + \sum_{j=1}^{J} \beta(v_{kj} - p_k), \quad k = 1 \dots K. \tag{2}$$

The clearing conditions state that the trade of the authorized participants offsets the stock traders' order imbalance. Note that we put a minus sign in front of AP_F, because here the authorized participants undo the position they took in their ETF trading.

Finally, in the third leg, the current share price equals its NAV:

$$p_F = \theta n_1 p_1 + \theta n_2 p_2 + \ldots + \theta n_K p_K. \tag{3}$$

Inserting (3) into the clearing condition for the ETF shares, (1), solving for AP_F and then inserting the resulting AP_F into the underlying stocks clearing condition, (2), gives us a system of K equations for the equilibrium pricing of the K stocks:

$$0 = \theta n_k \sum_{i=1}^{I} \alpha \left(v_{Fi} - (\theta n_1 p_1 + \theta n_2 p_2 + \ldots + \theta n_K p_K) \right) + \sum_{j=1}^{J} \beta (v_{kj} - p_k), \quad k = 1 \ldots K \tag{4}$$

The system of equations is simultaneous, and the coupling arises from the presence of index traders. Even if the index traders themselves do not participate in stock trading, the authorized participants' trade conveys the index traders' sentiment. This system of equations reveals how index trading creates price linkages between seemingly unrelated stocks. The first term represents the index traders' influence on stock k's price, transmitted through authorized participants' arbitrage activities. The second term captures direct stock-specific sentiment. In equilibrium, each stock's price reflects both direct trading based on its own market sentiment and indirect pressure from index-level sentiment. This creates a mechanism of price comovement among index constituents that would not exist otherwise. When index traders become more numerous or active (higher I or α), the coupling between stocks strengthens, potentially causing correlated price movements even among fundamentally unrelated companies.[5]

[5]Barberis, Shleifer, and Wurgler (2005) show that stocks added to the S&P 500 have larger betas with the S&P 500, and they attribute their findings to market frictions or investor sentiment. Greenwood (2008) documents stronger comovement among stocks in the Nikkei index. This evidence further confirms that the extra comovement can result from commonality in investors' behavior. In addition, Cong, Huang, and Xu (2024) theoretically show that the introduction of composite securities increases comovement of underlying assets.

The model we have described so far is of subjective beliefs. Traders agree to disagree about pricing, meaning that every trader understands that the counterparty to the trade holds opposing views on the stock's value.

We refer to the difference between the actual future price and the stock traders' estimates as traders' errors. Specifically, for each of the J stock-traders and each of the K stocks, we define the trader's error as the deviation:

$$\underbrace{v_k}_{\text{future price}} - \underbrace{v_{kj}}_{\substack{\text{stock trader's} \\ \text{estimation of} \\ \text{future price}}} =: \underbrace{x_{kj}}_{\text{trader's error}}, \quad j = 1, \ldots, J, \; k = 1, \ldots, K. \tag{5}$$

Similarly, we refer to the difference between the unobserved actual NAV and the index traders' estimates as traders' error. Thus, for each of the I index-traders, we write

$$\underbrace{\sum_{h=1}^{K} \theta n_h v_h}_{\text{future NAV}} - \underbrace{v_{Fi}}_{\substack{\text{index trader's} \\ \text{estimation of} \\ \text{future NAV}}} =: \underbrace{x_{Fi}}_{\text{error}}, \; i = 1, \ldots, I. \tag{6}$$

We denote by

$$\bar{x}_k := \frac{1}{J} \sum_{j=1}^{J} x_{kj}$$

and

$$\bar{x}_F := \frac{1}{I} \sum_{i=1}^{I} x_{Fi}.$$

Let the dollar returns on the stocks be

$$r_k = v_k - p_k$$

and let the dollar return on the fund's NAV be

$$r_F = \sum_{h=1}^{K} \theta n_h r_h = \sum_{h=1}^{K} \theta n_h (v_h - p_h).$$

Next, we re-express the equilibrium pricing condition, (4), in terms of the dollar returns and errors rather than the asset prices and perceptions of future prices, which simplifies to

$$0 = \theta n_k \alpha I \left(r_F - \bar{x}_F \right) + \beta J \left(r_k - \bar{x}_k \right), \quad k = 1 \dots K \tag{7}$$

Multiplying (7) by θn_k and summing over k, we get a single linear equation in r_F and its solution is

$$r_F = \underbrace{\frac{I\alpha \sum_{k=1}^{K} \theta^2 n_k^2}{I\alpha \sum_{k=1}^{K} \theta^2 n_k^2 + J\beta}}_{=: a} \times \bar{x}_F + \underbrace{\frac{J\beta}{I\alpha \sum_{k=1}^{K} \theta^2 n_k^2 + J\beta}}_{= 1-a} \times \sum_{k=1}^{K} \theta n_k \bar{x}_k \tag{8}$$

That is, the dollar return on the fund's NAV is a weighted average of traders' errors.

After plugging r_F back into (7), we get a system of K linear equations in r_k that is no longer simultaneous, and therefore the solution of the system is straightforward. We have

$$\theta n_k r_k = \theta n_k \bar{x}_k + \frac{n_k^2}{\sum_{h=1}^{K} n_h^2} a \left(\bar{x}_F - \sum_{h=1}^{K} \theta n_h \bar{x}_h \right) \tag{9}$$

and adding up (9) recovers (8).

Equation (7) reveals the mechanism through which index trading affects individual stock returns. Each stock's return reflects both direct trading pressure from stock-specific sentiment ($\bar{x}_k$) and indirect pressure transmitted through index traders via the authorized participants' arbitrage activities. The term a (defined in (8)) captures the relative influence of index versus stock trading.

3 Index Trading and Volatility of Returns

We now adopt the perspective of an econometrician analyzing data, treating the traders' errors in valuations as random variables. This means it becomes meaningful to use the model

to compute the standard deviation (volatility) of the returns and relate it to the number of indexers in the economy.

We start by decomposing the traders' errors as

$$x_{kj} = \mu_k + \epsilon_{kj} \qquad \text{and} \qquad x_{Fi} = \sum_{k=1}^{K} \theta n_k \mu_k + \epsilon_{Fi} \tag{10}$$

On its own, this decomposition is vacuous. Nevertheless, it motivates the distributional assumptions we are about to impose. We interpret μ_k as representing common valuation errors affecting stock k that stem from two sources: exogenous information shocks (e.g., news events) and systematic biases in trader sentiment (e.g., media-driven optimism or pessimism that affects all traders). In particular, we do not assume that the μ_k's are statistically independent of each other, which allows the model to capture market risk or, more generally, factor risk.

We denote

$$\bar{\epsilon}_k := \frac{1}{J} \sum_{j=1}^{J} \epsilon_{kj}$$

and

$$\bar{\epsilon}_F := \frac{1}{I} \sum_{i=1}^{I} \epsilon_{Fi}.$$

We rewrite the dollar returns, (8) and (9), in terms of μs and ϵs:

$$r_F = \sum_{k=1}^{K} \theta n_k \mu_k + a\bar{\epsilon}_F + (1-a) \sum_{k=1}^{K} \theta n_k \bar{\epsilon}_k \tag{11}$$

$$\theta n_k r_k = \theta n_k (\mu_k + \bar{\epsilon}_k) + \frac{n_k^2}{\sum_{h=1}^{K} n_h^2} a \left(\bar{\epsilon}_F - \sum_{h=1}^{K} \theta n_h \bar{\epsilon}_h \right) \tag{12}$$

Before specifying our distributional assumptions, we connect the model to the empirical diagnostic. The literature has already established the stock-level relation: higher index fund ownership is associated with higher volatility in constituent stocks. The external-inflows

explanation implies more. If this stock-level pattern is driven by external inflows to index funds, the same channel should also increase the volatility of the index, and hence of the index fund, itself.

The reason is mechanical in our setup. Index-fund order imbalance is accommodated by authorized participants via creation/redemption, which is implemented by trading the underlying basket. As a result, any external-flow-driven distortion that shows up in the dollar returns of constituents is transmitted to, and aggregated in, the index-fund dollar return. For broad value-weighted indices, this aggregation also has a sharp magnitude implication: the implied increase in index-fund volatility is typically larger than the increase in the volatility of a median constituent.

Theorem 1 formalizes this implication under distributional assumptions chosen for tractability. These assumptions are meant to isolate the external-flow footprint, not to define external inflows narrowly. In particular, they do not preclude (i) the possibility that marginal index-fund participants become noisier with entry, nor (ii) the possibility that index-fund participants herd on a common sentiment. The theorem shows how index-fund volatility and constituent-stock volatility must change together under the external-inflows channel.

For any expression that depends on the index traders, we let Δ denote the forward difference operator. For example,

$$\Delta \operatorname{var}(\bar{\epsilon}_F) := \operatorname{var}\left(\frac{1}{I+1}\sum_{i=1}^{I+1}\epsilon_{Fi}\right) - \operatorname{var}\left(\frac{1}{I}\sum_{i=1}^{I}\epsilon_{Fi}\right).$$

We have the following.

Theorem 1. *Assume* $\operatorname{cov}(\bar{\epsilon}_F, \bar{\epsilon}_k) = 0$, $\operatorname{cov}(\bar{\epsilon}_F, \mu_k) = 0$, *and* $\operatorname{cov}(\bar{\epsilon}_k, \mu_h) = 0$ *for every* $h, k \in \{1, \ldots, K\}$.

Further assume that the random variables $\{\bar{\epsilon}_k\}_{k=1}^{K}$ *are independent and have common variance*

σ^2. *Then,*

$$\Delta \operatorname{var}(\theta n_k r_k) = \left(\frac{n_k^2}{\sum_{h=1}^{K} n_h^2} \right)^2 \Delta \operatorname{var}(r_F) \tag{13}$$

The proof of Theorem 1 is in Appendix B.

While Theorem 1 is derived in terms of dollar returns to preserve the linear tractability of the equilibrium mapping, empirical researchers conventionally test variances of percentage returns. Because dollar variances are not scale-invariant, one might be concerned that the dramatic difference between the volatility response of the index and the constituent stocks is merely an artifact of scaling. We translate the variance of dollar returns into the variance of percentage returns. From the perspective of the econometrician, the equilibrium price is a random variable reflecting fluctuating market sentiments. Therefore, the percentage return, $R = r/p$, is the ratio of two random variables.[6] Following the long-standing tradition in asset pricing of using first-order Taylor approximations to linearize returns around their unconditional means, akin to the return decomposition framework pioneered by Campbell and Shiller (1988), we map the unconditional variance of dollar returns to percentage returns.

Corollary 2 (Translation to Percentage Returns)**.** *Let $R_k = r_k/p_k$ and $R_F = r_F/p_F$ denote the percentage returns of the stock and the fund, respectively. Let $\bar{p}_k$ and $\bar{p}_F$ denote their unconditional mean prices. The increase in the percentage-return variance of the constituent stock is related to the increase in the percentage-return variance of the index by:*

$$\bar{w}_k^2 \Delta \operatorname{var}(R_k) \approx \left(\frac{n_k^2}{\sum_{h=1}^{K} n_h^2} \right)^2 \Delta \operatorname{var}(R_F) \tag{14}$$

where $\bar{w}_k = \frac{\theta n_k \bar{p}_k}{\bar{p}_F}$ is the unconditional expected weight of stock k in the index fund.

The proof is a straightforward application of the Delta method to the variances on each side of the identity in Equation (13) (i.e., a first-order Taylor expansion around the respective unconditional means $\bar{r}$ and $\bar{p}$) and is therefore omitted.

[6] Nowhere in our framework do we assume normality. Consequently, we are not forced to assign a price support that includes zero, ensuring the unconditional variance of the percentage return remains well-defined.

The analysis shows that under the external-inflows channel, an increase in constituent-stock volatility should be accompanied by an increase in index volatility, with the latter generally larger in magnitude. Specifically, using ETF holdings data from September 11, 2025, the theorem implies that the increase in the variance of SPY is 3,627 times the median increase in the variance of stocks in the index. This number arises because $n_k^2 \ll \sum_{h=1}^{K} n_h^2$.[7] In other words, we have

$$\Delta \operatorname{var}(r_k) > 0 \implies \Delta \operatorname{var}(r_k) \ll \Delta \operatorname{var}(r_F).$$

Corollary 2 shows that the dollar-return restriction in Theorem 1 survives the translation to percentage returns. Rearranging Equation (14) gives

$$\Delta \operatorname{var}(R_F) \approx \bar{w}_k^2 \left(\frac{\sum_{h=1}^{K} n_h^2}{n_k^2} \right)^2 \Delta \operatorname{var}(R_k). \tag{15}$$

Thus, the relevant stock-to-index variance multiplier is not any single factor in Equation (14), but the entire coefficient on $\Delta \operatorname{var}(R_k)$ after the equation is solved for $\Delta \operatorname{var}(R_F)$. For a broad value-weighted index, this coefficient is large for the median constituent, and more generally for most constituents: the aggregation term $\left(\sum_h n_h^2/n_k^2\right)^2$ dominates the attenuation from the squared portfolio weight $\bar{w}_k^2$. Hence, under the external-inflows channel, a positive increase in a constituent's percentage-return variance implies a positive, and typically much larger, increase in the index fund's percentage-return variance. For a price-weighted index, by contrast, the same coefficient is of order one for a representative constituent, so the model predicts a proportional rather than magnified increase in index-fund variance.

Therefore, if empirical data show increased volatility in R_k without a corresponding positive increase in the fund's percentage-return volatility, this evidence is difficult to reconcile with

[7] A closely related external-flows formulation treats new index-fund participation as exogenous, price-inelastic noise-trader demand for ETF shares. Such demand is accommodated through authorized participants' creation/redemption activity. In this variant, the flow shock loads directly on the index fund's dollar return, while its impact on each constituent stock's dollar return is scaled by the stock's index weight in the creation/redemption basket (proportional to n_k). As a result, the induced increase in constituent-stock dollar-return variance is accompanied by a (typically much larger, for broad value-weighted indices) increase in index dollar-return variance. Hence, the qualitative prediction behind Theorem 1 is unchanged if we model external flow as pure inelastic noise demand.

external index-flow pressure as the driver of the observed pattern. In our empirical tests, we find no statistical evidence of a proportional increase in DJIA volatility, nor of the larger increase in S&P 500 volatility predicted by the value-weighted case. More broadly, we find no evidence that index ownership is associated with higher market-level volatility. This makes the external-inflows hypothesis harder to reconcile with the data under our model's assumptions.

The evidence therefore rules against the volatility footprint implied by the external-flow explanation. At the same time, the observed pattern of higher stock-level volatility without higher systematic or index-level volatility can arise if the link between index-fund ownership and volatility primarily reflects existing investors substituting away from stock-specific price discovery toward index-centric strategies. This is the internal-reallocation mechanism studied in Baruch and Zhang (2022). To keep the paper self-contained, Appendix A records a simple partial-equilibrium benchmark for this consistency result, but the empirical test itself does not require us to identify reallocating investors directly.

4 Data and Empirical Analyses

This section presents the empirical analysis and interprets the results through the model's predictions. That said, we cannot directly observe whether variation in index fund ownership reflects changes in the investor base or strategic reallocation by existing investors; we can only regress volatility on the level of index fund ownership.

Our sample period spans from 1980 to 2022, giving us 43 years of data. We choose 1980 as our beginning year because Thomson Reuters S12 started to record mutual fund ownership data at the end of 1979. We collect stock data from CRSP and related accounting data from the CRSP-Compustat merged database. We estimate the regressions at the annual frequency.

To construct an annual proxy of index fund ownership for each constituent stock, we rely on the CRSP Mutual Fund and the Thomson Reuters S12 databases. From CRSP mutual fund data, we identify equity funds by restricting the CRSP objective code to "EDC" or "EDY". To select equity index funds, we either let the index fund flag be "D" or follow Appel, Gormley, and Keim (2016) by searching fund names that contain any of the following strings: {Index, Idx, Indx, Ind_ (where _ indicates a space), Russell, S & P, S and P, S&P, SandP, DOW, Dow, DJ, MSCI, Bloomberg, KBW, NASDAQ, NYSE, STOXX, FTSE, Wilshire, Morningstar, 100, 400, 500, 600, 900, 1000, 1500, 2000, or 5000}. This definition is designed to include open-end mutual funds, closed-end funds, and ETFs identified as index funds. Next, for each index fund, we use Thomson Reuters S12 to locate its quarterly stock ownership. With these data, we construct quarterly index fund ownership for each stock k by summing all shares held by index funds and dividing by shares outstanding in the corresponding quarter. Then our measure of index fund ownership of stock k in year t, Index-Fund-Ownership$_{k,t}$, is the average of quarterly ownership in year t.

To construct an annual measure of volatility, we use CRSP daily returns to compute the sum of squared deviations (SSD) for each stock k in each year t. Our measure of volatility is the annualized SSD, i.e. $\frac{252}{N_{k,t}-1}\,\text{SSD}_{k,t}$, where $N_{k,t}$ is the number of daily returns we have for stock k in year t. This measure is proportional to the conventional variance and therefore provides a legitimate proxy for volatility.

We run, at the annual frequency, the following regression:

$$\text{Volatility}_{k,t} = \alpha + \beta\ \text{Index-Fund-Ownership}_{k,t-1} + \text{Controls}_{k,t-1} + \varepsilon_{k,t} \tag{16}$$

where the control variables are lagged and include the lagged market capitalization (large-cap stocks typically have lower volatility), the firm's lagged leverage (higher leverage amplifies equity volatility, as per Modigliani-Miller Proposition 2), the lagged annual return (volatility increases following a negative return), and the lagged market-to-book value of equity (high

values suggest growth, and therefore are associated with higher volatility). We also control for lagged volatility and, following Ben-David, Franzoni, and Moussawi (2018), the lagged reciprocal of the price. Finally, we control for year and firm fixed effects. Table 1 Panel A provides the summary statistics of these stock-level variables; Appendix C provides detailed definitions.

We do not control for total institutional ownership because the relevant variation is the ownership share held through index funds rather than the total institutional presence in a stock. Conditioning on total institutional ownership would therefore absorb part of the shift toward indexed ownership that the paper is designed to study.

The results in Table 2 are consistent with the existing literature. Unlike reconstitution-based designs, our method exploits cross-sectional variation in index fund ownership generated by the multiplicity of overlapping indices: different stocks belong to different index sets.

To test whether external flows drive this baseline result, we decompose the stocks' return volatility into systematic and idiosyncratic volatilities. If external index-flow pressure drives the volatility-indexing relationship, we should observe a positive relation between stocks' systematic volatility and index fund ownership, not only a relation with idiosyncratic volatility.

To estimate the annual systematic and idiosyncratic volatilities, we do the following. We use CRSP daily stock returns and CRSP daily returns of the market portfolio, r_{mkt}. In each year and for each stock, we compute the ordinary least squares regression of the market model

$$r_{k,t,d} = \alpha_{k,t} + \beta_{k,t} \times r_{\text{mkt},t,d} + \varepsilon_{k,t,d}, \tag{17}$$

where d stands for a trading day in year t. We use the identity

$$\text{SSD} = \text{SSE} + \text{SSR}$$

for the decomposition, where SSE is the explained sum of squares and SSR is the residual sum of squares. Specifically, as for the stock's volatility regression in Table 2, we use $\frac{252}{N_{k,t}-1}\,\text{SSD}$

Table 1: Summary statistics

This table provides the summary statistics of variables used in the paper. See Appendix C for detailed variable definitions.

Panel A: Summary statistics of stock sample

Variable	Mean	P25	Median	P75	Std	Number
Index-Fund-Ownership	0.0329	0.0014	0.0081	0.0442	0.0508	156,207
Volatility	0.3763	0.0846	0.1832	0.4221	0.5415	156,548
Idiosyncratic-Volt	0.3367	0.0664	0.1509	0.3674	0.5190	156,544
Systematic-Volt	0.0400	0.0048	0.0162	0.0420	0.0661	156,544
MktCap($Million)	3618.02	120.42	465.92	1949.31	10149.17	155,892
Log(MktCap)	6.2263	4.7910	6.1440	7.5752	2.0322	155,892
Leverage	0.2142	0.0100	0.1416	0.3473	0.2290	152,955
1/Price	0.1928	0.0296	0.0557	0.1277	0.6492	156,350
Ret	0.1242	−0.2097	0.0606	0.3384	0.5793	156,275
Mkt-to-Book	3.2442	1.1984	1.9097	3.3735	4.4905	147,320

Panel B: Summary statistics of index sample

Variable	Mean	P25	Median	P75	Std	Number
Index-Fund-Market-Share	0.0414	0.0024	0.0312	0.0670	0.0421	43
Policy-Uncertainty	120.14	90.60	112.94	142.40	44.84	43
Log(Policy-Uncertainty)	4.7363	4.5065	4.7269	4.9586	0.3126	43
S&P 500						
Volatility	0.0325	0.0129	0.0216	0.0333	0.0324	43
Ret	0.0998	−0.00003	0.1278	0.2577	0.1628	43
S&P 500 ETF: SPY						
Volatility	0.0357	0.0124	0.0245	0.0495	0.0353	30
Ret	0.1100	0.0127	0.1428	0.2633	0.1809	30
DJIA						
Volatility	0.0308	0.0118	0.0215	0.0348	0.0307	43
Ret	0.0991	−0.0061	0.1185	0.2234	0.1423	43
DJIA ETF: DIA						
Volatility	0.0366	0.0135	0.0258	0.0473	0.0369	25
Ret	0.0931	0.0012	0.0981	0.2083	0.1523	25

as the dependent variable. In Table 3, for the stock's systematic volatility, we use $\frac{252}{N_{k,t}-2}$ SSE as the dependent variable, and for the stock's idiosyncratic volatility, we use $\frac{252}{N_{k,t}-2}$ SSR as the dependent variable.

The regression results in Table 3 are difficult to reconcile with the external-flow explanation. Index fund ownership is strongly related to idiosyncratic volatility, but not to systematic volatility. This pattern is consistent with internal reallocation, but our main conclusion is

Table 2: Volatility and Index Investment of Constituent Stocks

This table presents the regression results for the stock-level relationship between annual volatility and index investment. The dependent variable Volatility is the annualized sum of squared deviations of stock daily returns. Index-Fund-Ownership is the average index ownership of a stock in a year. MktCap is the market value of equity. Leverage is the long-term debt divided by the sum of long-term debt and market value of equity. 1/Price is the inverse of the stock price. Ret is the annual stock return. Mkt-to-Book is the market-to-book ratio. See Appendix C for detailed definitions. All regressions include firm and year fixed effects. Standard errors are clustered at both firm and year levels and p-values are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	Dependent variable: $\text{Volatility}_{k,t}$
$\text{Index-Fund-Ownership}_{k,t-1}$	0.3759^{***}
	(0.00)
$\log(\text{MktCap})_{k,t-1}$	-0.0342^{***}
	(0.00)
$\text{Leverage}_{k,t-1}$	0.2178^{***}
	(0.00)
$1/\text{Price}_{k,t-1}$	0.2092^{***}
	(0.00)
$\text{Ret}_{k,t-1}$	-0.0454^{**}
	(0.01)
$\text{Mkt-to-Book}_{k,t-1}$	0.0069^{***}
	(0.00)
$\text{Volatility}_{k,t-1}$	0.2479^{***}
	(0.00)
Firm FE	Yes
Year FE	Yes
Clustered SE	Yes
Observations	126,265
R^2	0.6578

the sharper rejection of the external-flow volatility footprint.

Having established that the pattern is concentrated in the idiosyncratic component, we conduct a final macro-level falsification test on the index itself. Our theory (specifically, Theorem 1) tells us that if the primary driver of the positive relation between stocks' volatility and index fund ownership is external index-flow pressure, then when we regress the volatility of the index itself on the index fund market share, we should find a similar positive relation. To that end, we test the relation with two indices, the S&P 500 and the DJIA. As a proxy for annual index fund market share, we use the assets under management of all equity index funds at the end of year t divided by the total market capitalization (i.e. sum of all equity values) at the end of year t as constructed by CRSP. In particular, we use the same value for

Table 3: Idiosyncratic and Systematic Volatility with Index Investment

This table presents the regression results for the stock-level relationship between index investment and either idiosyncratic or systematic volatility. Idiosyncratic-Volt is the annualized residual sum of squares (SSR) from the market model regression for each stock. Systematic-Volt is the corresponding annualized explained sum of squares. Index-Fund-Ownership is the average index ownership of a stock in a year. MktCap is the market value of equity. Leverage is the long-term debt divided by the sum of long-term debt and market value of equity. 1/Price is the inverse of a stock price. Ret is the annual stock return. Mkt-to-Book is the market value of equity divided by book value of equity. See Appendix C for detailed definitions. All regressions include firm and year fixed effects. Standard errors are clustered at both firm and year levels and p-values are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	Dependent variable:	
	(1) Idiosyncratic-Volt$_{k,t}$	(2) Systematic-Volt$_{k,t}$
Index-Fund-Ownership$_{k,t-1}$	0.2889^{***}	0.0606
	(0.00)	(0.34)
log(MktCap)$_{k,t-1}$	-0.0457^{***}	0.0078^{***}
	(0.00)	(0.00)
Leverage$_{k,t-1}$	0.1857^{***}	0.0247^{***}
	(0.00)	(0.00)
1/Price$_{k,t-1}$	0.2200^{***}	-0.0008
	(0.00)	(0.52)
Ret$_{k,t-1}$	-0.0499^{***}	0.0069^{**}
	(0.00)	(0.04)
Mkt-to-Book$_{k,t-1}$	0.0060^{***}	0.0007^{**}
	(0.00)	(0.03)
Idiosyncratic-Volt$_{k,t-1}$	0.2152^{***}	0.0134^{***}
	(0.00)	(0.00)
Systematic-Volt$_{k,t-1}$	0.1689	0.3139^{***}
	(0.20)	(0.00)
Firm FE	Yes	Yes
Year FE	Yes	Yes
Clustered SE	Yes	Yes
Observations	126,263	126,263
R^2	0.6566	0.6420

both indices.

As a measure of annual volatility, we use the same measure we used for stocks, namely the annualized SSD of the respective index. We run, at the annual frequency, the following regression:

$$\text{Volatility}_{\text{index/ETF},t} = \alpha + \beta\ \text{Index-Fund-Market-Share}_{t-1} + \text{Ret}_{\text{index/ETF},t-1} + \text{Log(Policy-Uncertainty)}_{t-1} + \text{S\&P 500 Dummy} + \varepsilon_{\text{index/ETF},t}, \quad (18)$$

where the lagged measure of policy uncertainty is as constructed by Baker, Bloom, and Davis (2016). We also include a dummy variable for the index, which serves as an index fixed effect. Table 1 Panel B provides the summary statistics of these index-level variables.

Table 4: Index Volatility and Index Fund Market Share

This table presents the regression results for the index-level relationship between annual volatility and index investment. The dependent variable Volatility is the annualized sum of squared deviations of daily returns of an index or its ETF. Index-Fund-Market-Share is the assets under management of all index funds divided by the market capitalization of equities at the end of a year. Ret is the annual return of an index or its ETF. Policy-Uncertainty is the economic policy uncertainty index. S&P500 Dummy is equal to 1 if the index is S&P 500 and 0 otherwise. See Appendix C for detailed definitions. Standard errors are clustered at the year level and p-values are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	Dependent variable:	
	(1) $\text{Volatility}_{\text{index},t}$	(2) $\text{Volatility}_{\text{ETF},t}$
$\text{Index-Fund-Market-Share}_{t-1}$	0.1582	0.1408
	(0.43)	(0.58)
Ret_{t-1}	−0.0095	−0.0163
	(0.75)	(0.63)
$\text{Log(Policy-Uncertainty)}_{t-1}$	−0.0079	−0.0138
	(0.78)	(0.66)
S&P500 Dummy	0.0018	0.0013
	(0.17)	(0.55)
Clustered SE	Yes	Yes
Observations	86	53
R^2	0.0291	0.0148

Table 4 shows no statistically significant relation between index fund market share and the volatility of the index. To explore this null result further, we examine the economic significance of the estimate in Column (1) of Table 4. The coefficient on lagged index fund market share is 0.1582 with a p-value of 0.43. This implies a 95% confidence interval of approximately $[-0.24, 0.56]$. This interval not only contains negative values, but its upper bound of 0.56 is economically modest. This finding contrasts with the prediction from our model's external-inflows scenario (Theorem 1). If external inflows were the driver, any increase in stock volatility must be accompanied by an increase in index volatility. Specifically, the theory predicts that the coefficient in the index volatility regression (Table 4) should be materially larger than the coefficient in the stock volatility regression (Table 2). Our empirical results show the opposite: the point estimate of 0.1582 is not only statistically

indistinguishable from zero, but it is actually smaller than the highly significant coefficient of 0.3759 we find for individual stocks. This evidence makes the external-inflows hypothesis harder to reconcile with the data under our model's assumptions.

5 Concluding Remarks

Prior studies document a positive association between indexation and stock volatility. This paper asks whether that pattern is driven by external flows into index products, and finds little evidence for that explanation.

While the data do not reveal the identity of the marginal participants, a price-discovery model allows us to test the indirect footprint of the external-flow explanation. External inflows into index products should primarily operate through index-level trading pressure and therefore should leave traces in systematic or market-level volatility, not only in idiosyncratic stock-level volatility. The empirical pattern we document is different: the association between index ownership and volatility is concentrated at the constituent level, while we find little evidence of a corresponding market-level response. This pattern is difficult to reconcile with external inflows from new investors as the primary driver of the stock-level volatility effect.

As a falsification test, we decompose stock volatility into its systematic and idiosyncratic components, and find that the association between higher levels of indexation and greater stock volatility is confined to the idiosyncratic part. This result is difficult to reconcile with the view that new entrants who participate through index products are the primary driver of the enhanced stock volatility.

Taken together, the evidence rules against a simple external-flow volatility story and is consistent with existing investors shifting to index-centric strategies. By doing so, they reduce their direct participation, contributing only indirectly to the price discovery of individual

stocks through authorized participants. This insight demonstrates that growth in index fund ownership is not monolithic: its consequences hinge on the underlying mechanism.

Greenwood and Sammon (2025) show that traditional index reconstitution events are losing power due to anticipation and diminished price effects. Our cross-sectional approach offers an alternative empirical design. Relative to reconstitution-based designs, this approach exploits broader cross-sectional and time-series variation in index exposure to study indexing's effects.

Appendix

A Internal Reallocation Benchmark

The main text uses the model to derive the systematic-volatility implication of external index-flow pressure. This appendix records a simple internal-reallocation benchmark showing that the opposite footprint can arise naturally when existing traders shift from individual stock picking to index tracking. The result is not needed to reject the external-flow volatility explanation. It simply shows that the empirical pattern we document, higher stock-level volatility without higher systematic or index-level volatility, is consistent with reduced stock-level price discovery.

To establish this consistency, our result hinges on three key assumptions. First, we specify a baseline relationship between the trading intensities of stock and index traders that ensures their aggregate market impact can be compared. These assumptions are stylized, but they allow the model to isolate the internal-reallocation channel and derive a sharp empirical prediction. Furthermore, this simplification is grounded in equilibrium theory. Baruch and Zhang (2022) study a rational expectations equilibrium with endogenous trading intensities and also find that internal reallocation leaves the price of the market portfolio unchanged. Here, in our partial equilibrium model, the assumption acts as a theoretical anchor. If, upon shifting, index traders trade at an intensity even lower (higher) than specified in the theorem, the resulting volatility of the index would be correspondingly lower (respectively higher). Our objective is simply to establish that internal reallocation is consistent with finding no relation between index volatility and index fund market share.

Second, we assume that when traders switch to indexing, they form their view on the overall market by aggregating the views they previously held on individual stocks. This seems reasonable. Under these first two assumptions, the aggregate perception of the index's future

value remains completely unchanged during the transition. As a result, both the equilibrium dollar return and the equilibrium price of the index are perfectly invariant to the internal shift.

Our third assumption is that trader-specific errors are uncorrelated across traders. This simply means that the component of a trader's error that is not explained by common market factors is unique to that trader. This assumption allows us to cleanly derive the result that as there are fewer stock traders contributing to the price discovery of individual stocks, the greater the volatility of the stock.

Formally, we have the following.

Theorem A.1. *Assume the following conditions hold:*

A) ***Equivalent Trading Intensity:*** *The trading intensities α (for index traders) and β (for stock traders) satisfy the relation $\alpha \sum_{h=1}^{K} \theta^2 n_h^2 = \beta$. When a trader switches, they adjust their intensity accordingly.*

B) ***Consistent Views:*** *A trader switching from stock picking to index tracking forms their new view on the fund's future value by aggregating their prior views on the constituent stocks:*

$$x_{F,I+1} = \sum_{k=1}^{K} \theta n_k x_{kJ} = \sum_{k=1}^{K} \theta n_k (\mu_k + \epsilon_{kJ}) \tag{A.1}$$

C) ***Trader-specific Errors:*** *For each stock trader j, let $z_{kj} = \theta n_k \epsilon_{kj} - \frac{n_k^2}{\sum_h n_h^2} \sum_h \theta n_h \epsilon_{hj}$. Then, $\text{var}\left(\frac{1}{J} \sum_{j=1}^{J} z_{kj}\right)$ is decreasing with J.*[8]

Then, a switch from stock trading to index trading increases individual stock volatility while having no impact on index volatility.

The proof of Theorem A.1 is in Appendix B.

[8]This assumption holds if the trader-specific errors, ϵ_{kj} are i.i.d.

B Proofs

Proof of Theorem 1. To prove the theorem, we only need to show that

$$\operatorname{var}(\theta n_k r_k) = \left(\frac{n_k^2}{\sum_{h=1}^K n_h^2}\right)^2 \operatorname{var}(r_F) + \text{ some terms that do not depend on index traders.} \tag{B.1}$$

We have

$$\operatorname{var}(r_F) \underset{(11)}{=} a^2 \operatorname{var}(\bar{\epsilon}_F) + \theta^2(a^2 - 2a)\sum_{h=1}^K n_h^2\sigma^2 + f, \tag{B.2}$$

where f absorbs all the terms that do not depend on I. Specifically,

$$f = \operatorname{var}\left(\sum_{k=1}^K \theta n_k \mu_k\right) + \operatorname{var}\left(\sum_{k=1}^K \theta n_k \bar{\epsilon}_k\right).$$

We also have

$$\begin{aligned}
\operatorname{var}(\theta n_k r_k) &\underset{(12)}{=} \left(\frac{n_k^2}{\sum_{h=1}^K n_h^2}\right)^2 a^2 \left(\operatorname{var}(\bar{\epsilon}_F) + \theta^2 \sum_{h=1}^K n_h^2\sigma^2\right) - 2\theta^2 a \frac{n_k^2}{\sum_{h=1}^K n_h^2} n_k^2\sigma^2 + g \\
&= \left(\frac{n_k^2}{\sum_{h=1}^K n_h^2}\right)^2 \left(a^2 \operatorname{var}(\bar{\epsilon}_F) + a^2\theta^2 \sum_{h=1}^K n_h^2\sigma^2 - 2a\theta^2 \sum_{h=1}^K n_h^2\sigma^2\right) + g \\
&= \left(\frac{n_k^2}{\sum_{h=1}^K n_h^2}\right)^2 \left(a^2 \operatorname{var}(\bar{\epsilon}_F) + (a^2 - 2a)\theta^2 \sum_{h=1}^K n_h^2\sigma^2\right) + g \\
&\underset{(B.2)}{=} \left(\frac{n_k^2}{\sum_{h=1}^K n_h^2}\right)^2 (\operatorname{var}(r_F) - f) + g \\
&= \left(\frac{n_k^2}{\sum_{h=1}^K n_h^2}\right)^2 \operatorname{var}(r_F) + \underbrace{g - \left(\frac{n_k^2}{\sum_{h=1}^K n_h^2}\right)^2 f}_{\text{does not depend on } I},
\end{aligned}$$

where g does not depend on I. Specifically,

$$g = \operatorname{var}(\theta n_k \mu_k) + \operatorname{var}(\theta n_k \bar{\epsilon}_k).$$

Thus, we have shown (B.1), and this concludes the proof.

■

Proof of Theorem A.1. We prove the theorem in two parts. First, we show that the switch has no impact on the fund's dollar return, r_F, and hence its volatility. Second, we show that the switch increases the volatility of individual stocks' dollar returns, r_k.

Part 1: Invariance of Index Return (r_F)

The first assumption on equivalent trading intensity simplifies the weight a in equation (8) to $a = I/(I+J)$. This means the fund's dollar return, r_F, becomes a simple average of the perceived future NAV across all $N = I + J$ market participants. The second assumption on consistent views ensures that when a trader switches strategies, their aggregate perception of the fund's future value does not change. Since the total number of traders is constant and the set of all traders' perceptions is unchanged, their average, r_F, is identical before and after the switch. Consequently, its variance, $\text{var}(r_F)$, is also unchanged.

Part 2: Increase in Stock Volatility ($\text{var}(r_k)$)

The dollar return for an individual stock k is given by:

$$\theta n_k r_k = \theta n_k(\mu_k + \bar{\epsilon}_k) + \frac{n_k^2}{\sum_{h=1}^K n_h^2}\left(r_F - \sum_{h=1}^K \theta n_h(\mu_h + \bar{\epsilon}_h)\right)$$

Rearranging this equation isolates the terms that depend on the number of stock traders, J. The terms μ_h are common errors and are unaffected by the switch. From Part 1, r_F is also invariant. Grouping all invariant terms together gives:

$$\theta n_k r_k - \underbrace{\left(\theta n_k \mu_k + \frac{n_k^2}{\sum_{h=1}^K n_h^2}\left(r_F - \sum_{h=1}^K \theta n_h \mu_h\right)\right)}_{\text{invariant terms}} = \theta n_k \bar{\epsilon}_k - \frac{n_k^2}{\sum_{h=1}^K n_h^2}\sum_{h=1}^K \theta n_h \bar{\epsilon}_h$$

By the theorem assumption, the variance of the right hand side is increasing as J decreases. It follows that the variance of $\theta n_k r_k$ must also increase as J decreases. This concludes the proof.

■

C Variable Definition

Variable	Definition
Stock Level	
Index-Fund-Ownership	The average of quarterly index ownership in a year of the individual stock. Quarterly index ownership is defined as shares held by equity index funds from Thomson Reuters S12 divided by common shares outstanding from CRSP at the end of each quarter.
Volatility	The annualized sum of squared deviations (SSD) of daily returns, $\frac{252}{N-1}$SSD, where N is the number of a stock's daily returns from CRSP in a year.
Idiosyncratic-Volt	The annualized residual sum of squares (SSR) from the market model regression using daily returns in a year from CRSP: $r = \alpha + \beta \times r_{\text{mkt}} + \varepsilon$, $\frac{252}{N-2}$SSR, where N is the number of a stock's daily returns r in a year and r_{mkt} is the daily value-weighted return of the market.
Systematic-Volt	The annualized explained sum of squares (SSE) from the market model regression using daily returns in a year from CRSP: $r = \alpha + \beta \times r_{\text{mkt}} + \varepsilon$, $\frac{252}{N-2}$SSE, where N is the number of a stock's daily returns r in a year and r_{mkt} is the daily value-weighted return of the market.
MktCap	Market value of equity (Compustat #199×#25) at the end of a year.
Leverage	Long-term debt (#9) divided by the sum of long-term debt (#9) and market value of equity (#199×#25) at the end of a year.
1/Price	The inverse of stock price (#199) at the end of a year.
Ret	The annual return constructed using monthly returns of the individual stock from CRSP in a year.
Mkt-to-Book	Market value of equity divided by book value of equity (#60) at the end of a year.
Index Level	
Index-Fund-Market-Share	The assets under management of all equity index funds at the end of year t divided by the total market capitalization. The total market capitalization is the sum of all equity values at the end of year t constructed by CRSP.
Volatility	The annualized sum of squared deviations constructed using daily returns of an index or its ETF, where the daily return of an index is defined as the percentage change of its price level and the daily return of an index ETF is from CRSP.
Ret	The annual return of an index constructed using its price levels or that of an index ETF constructed using monthly returns from CRSP.
Policy-Uncertainty	A normalized index of the volume of news articles discussing economic policy uncertainty constructed by Baker, Bloom, and Davis (2016). See www.PolicyUncertainty.com for details.

References

Allard, A., Farkas, M., Pedio, M., Rubio, S., and Tonks, I., 2025. A survey of the consequences of passive investment funds for financial markets. Working paper, University of Bristol.

Anadu, K., Kruttli, M., McCabe, P., and Osambela, E., 2020. The shift from active to passive investing: Risks to financial stability? Financial Analysts Journal 76 (4), 23–39.

Appel, I., Gormley, T., and Keim, D., 2016. Passive investors, not passive owners. Journal of Financial Economics 121 (1), 111–141.

Baker, S., Bloom, N., and Davis, S., 2016. Measuring Economic Policy Uncertainty. The Quarterly Journal of Economics 131 (4), 1593–1636.

Barberis, N., Shleifer, A., and Wurgler, J., 2005. Comovement. Journal of Financial Economics 75 (2), 283–317.

Baruch, S. and Zhang, X., 2022. The distortion in prices due to passive investing. Management Science 68 (8), 6219–6234.

Ben-David, I., Franzoni, F., and Moussawi, R., 2018. Do ETFs increase volatility? Journal of Finance 73 (6), 2471–2535.

Campbell, J.Y. and Shiller, R.J., 1988. The dividend-price ratio and expectations of future dividends and discount factors. Review of Financial Studies 1 (3), 195–228.

Chinco, A. and Sammon, M., 2024. The passive-ownership share is double what you think it is. Journal of Financial Economics 157: 103860.

Coles, J., Heath, D., and Ringgenberg, M., 2022. On index investing. Journal of Financial Economics 145 (3), 665–683.

Cong, L., Huang, S., and Xu, D., 2024. The rise of factor investing: "passive" security design and market implications. Technical report, National Bureau of Economic Research.

Dannhauser, C. and Pontiff, J., 2025. Flow. Working paper, Villanova University.

Greenwood, R., 2008. Excess comovement of stock returns: Evidence from cross-sectional variation in Nikkei 225 weights. The Review of Financial Studies 21 (3), 1153–1186.

Greenwood, R. and Sammon, M., 2025. The Disappearing Index Effect. Journal of Finance 80 (2), 657–698.

Hortaçsu, A. and Syverson, C., 2004. Product differentiation, search costs, and competition in the mutual fund industry: A case study of S&P 500 index funds. Quarterly Journal of Economics 119 (2), 403–456.

Sushko, V. and Turner, G., 2018. The implications of passive investing for securities markets. BIS Quarterly Review, March.

Treynor, J., 1987. Market Efficiency and the Bean Jar Experiment. Financial Analysts Journal 43 (3), 50–53.

Vijh, A., 1994. S&P 500 trading strategies and stock betas. Review of Financial Studies 7, 215–251.

Settlement Manipulation in Prediction Markets

David Dai[*] Ruizhe Jia[†] Shihao Yu[‡]

Current version: October 4, 2026

[*]Department of Management Science and Engineering, Stanford University; htdai@stanford.edu.
[†]Department of Management Science and Engineering, Stanford University; ruizhe@stanford.edu.
[‡]Lee Kong Chian School of Business, Singapore Management University; shihaoyu@smu.edu.sg.

Settlement Manipulation in Prediction Markets

Abstract

Prediction markets increasingly list contracts settling on an asset price that holders can move by trading the underlying. We build a model showing that such contracts transfer wealth from prediction-market liquidity traders to manipulators and harm price discovery in the underlying, even as it becomes more liquid. After the launch of Polymarket's five-minute Bitcoin contract, settlement-time spot order flow spikes, causing large price reversals after settlement. Manipulators capture a large amount of profit, mostly from retail. Manipulation is largely absent in the fifteen-minute contracts: lengthening the contract horizon removes it, providing the market-design remedy our model and evidence support.

1 Introduction

Prediction markets have grown with remarkable speed: combined monthly volume on Kalshi and Polymarket rose nearly fivefold in seven months, from under $5 billion in September 2025 to about $24 billion by April 2026.[1] Surprisingly, the fastest-growing of these contracts are not written on elections or sports, but on asset prices. Each is a binary contract: it pays $1 if the asset is above a chosen level at the contract's close, and $0 otherwise. These began in cryptocurrencies, where they are already large. On Kalshi alone, volume passed $1 billion a month for the first time in March 2026.[2] Now they are reaching mainstream finance: Nasdaq and Cboe have filed to list them on equity indices.[3] And more importantly, the contract horizon keeps shrinking: from a day, to an hour, to fifteen minutes, and on Polymarket to just five minutes.

Regulators have not merely permitted prediction contracts; they describe their benefits to the public. The CFTC tells investors that "a citrus farmer might buy a weather event contract to hedge against losses that might be caused by a sudden freeze," and that prediction markets "can sometimes forecast event outcomes better than polling or other forms of forecasting."[4] Academic work offers concrete support: the Iowa Electronic Markets forecast U.S. presidential vote shares to within about 1.5 percentage points, outperforming large-scale polls (Wolfers and Zitzewitz, 2004), and a prominent group of economists argues that prediction markets deliver "a lower prediction error than conventional forecasting methods" (Arrow et al., 2008). Even the worry that someone might manipulate the price has largely been set aside: a manipulator can increase price accuracy (Hanson and Oprea, 2009), and documented manipulation attempts reverted quickly (Rhode and Strumpf, 2008). If that logic carried over to asset prices, these contracts would make financial markets more informative, not less.

This paper shows the opposite for asset-price contracts: rather than making the underlying

[1] Pew Research Center, "Trading Volume on Prediction Markets Has Soared in Recent Months" (2026).

[2] CF Benchmarks, "Kalshi Leads Surging Crypto Event Contract Market" (2026).

[3] Nasdaq and Cboe each filed with the SEC in March 2026 to list binary (fixed-payout) options on equity indices.

[4] Commodity Futures Trading Commission, "Understanding Prediction Markets and Event Contracts" (investor advisory), `cftc.gov`.

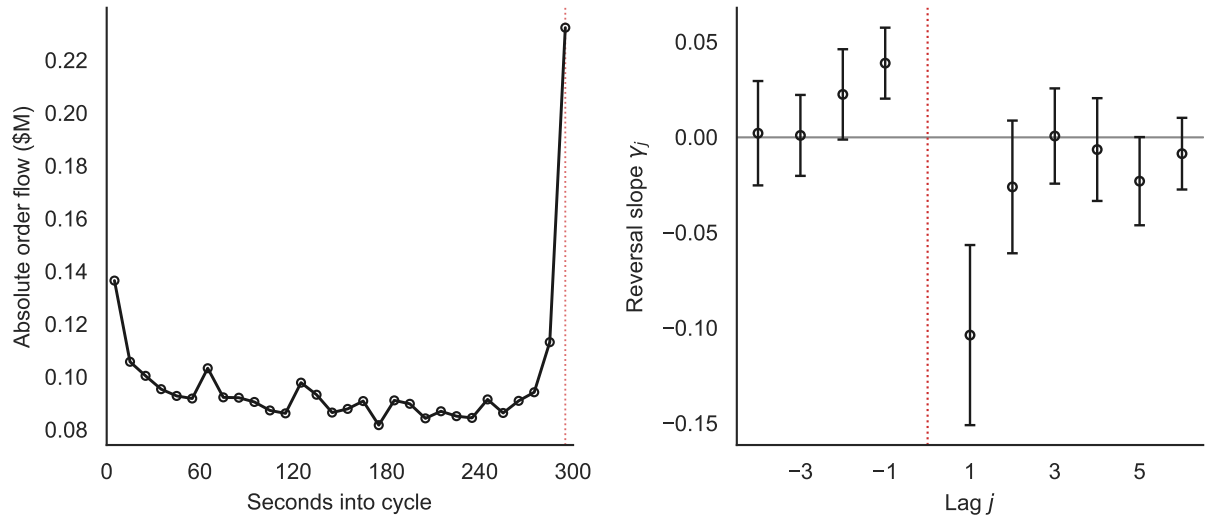

Figure 1. The footprint of settlement manipulation. This figure plots Binance BTC spot activity around the five-minute contract's close, in the post-launch period. *Left*: absolute order flow per five-minute cycle, folded onto the rolling clock; the dotted vertical line marks the final ten seconds before the close. *Right*: the post-settlement reversal slope γ_j at each lag j around the close, the slope of the lag-j ten-second return on the final-ten-second return. Section 6 details both constructions and establishes the patterns against pre-launch and across-moneyness benchmarks.

market more informative, they degrade its price discovery and transfer wealth from ordinary traders to those who manipulate them. The vulnerability is structural. An asset-price contract settles on a financial price, and that price can be moved by trading the underlying market itself, the cash-settlement manipulation of Kumar and Seppi (1992), now inside a retail prediction market with horizons as short as five minutes. U.S. law already requires every listed contract to be "not readily susceptible to manipulation";[5] we show that asset-price contracts raise exactly the concern that standard is meant to prevent. The footprint is in the data (Figure 1): after the five-minute contract went live on Polymarket, in the final seconds before settlement, Binance spot order flow spikes, resulting in price movements that revert shortly after settlement. Clearly, this is a transitory push to manipulate the spot price, not trading on information.

We develop a sequential game-theoretic model of the two markets. The spot market has its usual players (an informed trader who knows the asset's value and liquidity traders who do not) and now a manipulator, who has bought the contract and trades the spot just before settlement. Her order

[5]Commodity Exchange Act, Designated Contract Market Core Principle 3; reaffirmed for event contracts in CFTC Staff Advisory 26-08 (2026).

pools with the rest, and because she trades on no information, a late buy is now less likely to be the informed trader's. It therefore carries less information about the asset, so the settlement price reveals less of it and price discovery falls. Yet, by lowering the adverse selection the spot maker faces, it leads the spot maker to supply more liquidity, so the manipulation makes the spot market more liquid even as it makes the price less informative.

The same inability to tell who is trading also makes the liquidity traders in the prediction market earn worse-than-fair odds. A contract buy is public, but the spot maker cannot tell the manipulator from an ordinary contract buyer. It therefore puts some probability on a manipulator behind the trade, one who will push the spot toward her side at settlement. Anticipating that push, the spot maker shades its settlement-window quote against it, and the settlement tilts against the side the contract buyer took. A prediction-market liquidity trader who bought with no intent to manipulate is overpaying all the time for this contract. More importantly, we show that a prediction-market maker earns a rent as long as the bid and ask cannot cross. The prediction-market liquidity trader's loss is thus a transfer to the prediction-market maker and the manipulator.

These harms are not inevitable: both the evidence and the model point to one design lever, the contract's horizon. The model explains why the horizon matters. A short horizon leaves little room for the asset's own price discovery before settlement, so a fixed spot push is more likely to determine the contract's resolution. A longer horizon aggregates more ordinary spot-market information before the close, making the price paths in which a push is pivotal rarer and reducing the price-discovery loss on net. The remedy is direct: lengthen the horizon. In the data, the manipulation signature is sharp for the five-minute contract, but much attenuated for the fifteen-minute contract.

Empirically, we study Polymarket's Bitcoin five-minute up/down contract: a binary claim that pays \$1 if Bitcoin is higher at the close of a five-minute window than at its open, and \$0 otherwise. Every trade is recorded on a public blockchain, so each trade and trader is traceable, a transparency centralized exchanges and options venues do not offer. The contract resolves not on Polymarket but against a Chainlink oracle that averages the price across major spot exchanges. The price on Binance, the largest crypto exchange by volume, is an economically tight proxy for that benchmark:

its mid sits about two and a half basis points from the oracle and moves essentially one-for-one with it within each five-minute window, finishing on the same side of the strike as the resolution about 85% of the time, so a push that drives the Binance mid a few basis points past the strike reliably carries the resolution. We therefore measure the settlement-time flow on Binance. The contract's launch on February 12, 2026 offers a clean natural experiment: no five-minute contract existed before that date, so we compare three regimes: no such contract (P1), the 15-minute and 4-hour contracts only (P2), and the five-minute contract live (P3).

As soon as it launched, a relatively small prediction market began redirecting large spot orders to the settlement seconds. In the final ten seconds before each close, trading on Binance spikes: the magnitude of net order flow jumps, and volatility rises with it. The timing is no accident: the spike appears only after the five-minute launch (about 50% above the pre-launch level), and is significantly attenuated at the fifteen-minute horizon. The spike is sharpest where a push is pivotal: in the roughly 6% of cycles whose contract price still implies a near-even outcome just before the close, the near-settlement order-flow jump is about 3.9 times that in the rest. The clearest sign that this is not information is the reversal: within ten seconds the price reverts, by about a quarter in the near-even cycles and a tenth in the others. Real information would persist; the price impact of a manipulative push reverts. Last, the liquidity evidence is asymmetric across venues: near-even closes add depth on Binance, exactly where the model predicts an uninformed push should lower spot adverse selection, while Polymarket depth falls as binary-market makers face direct exposure to a flipping settlement.

We then identify the cycles most likely to contain a manipulative push. For each cycle we measure the size of the final-ten-second net order flow on Binance relative to the cycle's own typical flow. We classify the cycles in the top decile of this measure, about 1,600, as manipulated. Away from the close they are quieter than average, and they fall disproportionately in thin-liquidity hours: 56% in the overnight window and 44% on weekends, against 40% and 27% of normal cycles, where a given dollar of near-settlement flow moves the spot price, and therefore the resolution, the most.

Within these manipulated cycles, the pushes do more than cluster at the close: they change

which side wins. In cycles where the contract was priced near even, a push against the favored side flipped the winner 65% of the time, against 41% in ordinary trading. It works by dragging the Binance price, and thus the oracle price, across the strike. The stronger test is the cycles that were all but decided: even when the market gave one side a 90-to-100% chance just before the close, a push against it reversed the outcome 34% of the time, against 1% in cycles with no push. A bet the market treated as near-certain was overturned one time in three.

We next ask who gains from that and who pays for it. Because Polymarket settles on a public blockchain, every wallet's profit is visible, and we identify as manipulators the traders who gain in the manipulated cycles, an identification based on realized profits rather than a direct observation of intent. Just 821 of them fit, about one in three hundred of the roughly 243,000 who traded the contract, and they take $8.2 million in the pushed cycles while breaking even in the rest: the edge lives only where the manipulation does. Who pays divides just as sharply. Setting aside the market makers, who quote passively and end each cycle flat, almost none of the loss lands on them: 93% of it falls on retail. These retail accounts are the liquidity traders of the model, and the contract treats them accordingly: they are on the losing side in 65% of manipulated cycles, against 48% of normal ones. All of it comes from a single contract, on a single platform, over two months.

The natural innocent explanation is hedging: a prediction-market maker who has sold the binary contract offsets its risk by trading the spot. Two facts rule that out, and the clearer one is the state of the contract when the pushing happens. A binary contract's sensitivity to the spot price is large only when the price sits near the strike; once one side is almost certain to win, the contract barely moves with the spot, and the prediction-market maker is left with almost nothing to hedge. Yet those nearly decided cycles (where the market already gave one side a 90-to-100% chance) are exactly the ones a push still flips, one time in three. A trade placed when there is nothing to hedge cannot be a hedge; whatever hedging does occur can only come afterward, once the push has dragged the price back toward the strike and revived the exposure. The second fact is the timing. Even when there is something to hedge, a prediction-market maker would build its spot position gradually, as that exposure accumulates over the cycle. Instead the flow arrives in a single burst at the very

end: essentially all of the signed trading falls in the final fifty seconds, and almost none before. A position assembled only in the closing seconds is not a hedge accumulated over the cycle; it is a trade placed to move the settlement price.

The paper proceeds as follows. Section 2 positions the contribution. Section 3 describes the contracts and how they settle. Sections 4 and 5 develop the model and its predictions. Section 6 presents the testable predictions, the data, and the empirical results, and Section 7 discusses the market-design implications.

2 Related Literature

This paper joins three strands of literature that have largely developed apart: the economics of prediction markets, the manipulation of cash-settled derivatives, and the market microstructure of decentralized and crypto venues.

Prediction markets. Prediction markets are valued as information aggregators whose prices forecast events as accurately as expert judgment (Wolfers and Zitzewitz, 2004), a view recently reaffirmed for macroeconomic contracts on Kalshi (Diercks, Katz, and Wright, 2026). A long-standing concern is that such prices are vulnerable to manipulation, yet the literature has largely concluded that manipulation need not impair price accuracy. In a Kyle-style model, a manipulator's order flow raises the returns to informed trading and can even improve price accuracy (Hanson and Oprea, 2009); in laboratory markets, other traders anticipate and discount manipulators' biased orders, leaving accuracy intact (Hanson, Oprea, and Porter, 2006); and documented manipulation attempts in the field have produced only transient, mean-reverting price effects (Rhode and Strumpf, 2008). A growing empirical literature exploits transaction-level data from Polymarket and Kalshi: one strand examines who profits and who is informed, finding that a small group of sophisticated traders captures most of the gains while retail participants fund them (Akey, Grégoire, Harvie, and Martineau, 2026; Gómez-Cram, Guo, Jensen, and Kung, 2026; Della Vedova, 2026); a parallel

strand measures adverse selection and order flow (Bartlett and O'Hara, 2026; Tsang and Yang, 2026; Dubach, 2026) and screens for suspicious accounts (Mitts and Ofir, 2026).

To our knowledge this is the first study of ultra-short, oracle-settled prediction markets on an asset's price, and the first to identify settlement manipulation in them: a trader who holds the binary contract trades the underlying asset near settlement to move the price on which the contract pays. We show that this manipulation not only transfers wealth from liquidity traders to manipulators in the prediction market, but also harms price discovery in the underlying spot market.

Manipulation of cash-settled derivatives. At its core, our mechanism is the cash-settlement manipulation of Kumar and Seppi (1992): a trader takes a derivative position and then trades the underlying spot to move the cash-settlement price, profiting when the gain on the position exceeds the cost of the spot trade. Another theory paper studying manipulation in derivatives markets is Zhang (2022), in which manipulators can take endogenous derivative contract positions and leave every agent worse off. On the empirical side, the literature documents settlement and expiration distortions in equities and volatility benchmarks (Ni, Pearson, and Poteshman, 2005; Comerton-Forde and Putniņš, 2011; Griffin and Shams, 2018; Baltussen, Terstegge, and Whelan, 2023). Spatt (2014) surveys a variety of manipulation strategies. Finally, a design literature asks how to build benchmarks that resist manipulation (Duffie and Dworczak, 2021).

Our contribution relative to this work is twofold. First, on the theory side, we build a model of cash-settlement manipulation in a new market: ultra-short, oracle-settled prediction markets on an asset's price. This model is based on the sequential trading framework of Glosten and Milgrom (1985) and on the two-venue price-discovery structure of Zhu (2014). In this market, manipulation harms liquidity traders and degrades the underlying spot's price discovery. We further provide a market-design remedy: a longer contract horizon lowers both the manipulator's rents and the price-discovery loss. Second, on the empirical side, because this market settles on-chain—unlike the equity, futures, and volatility-benchmark venues where cash-settlement manipulation is otherwise studied—we observe trading at the level of pseudonymous wallets, which allows a more granular

analysis: we identify the likely manipulators and measure the wealth they extract.

Decentralized and crypto market microstructure. Finally, the paper contributes to the microstructure of decentralized and crypto markets by studying a new on-chain venue: a binary, oracle-settled prediction market. One strand studies the microstructure, price discovery, and quality of crypto trading venues, both centralized and on-chain (Makarov and Schoar, 2020; Lehar, Parlour, and Zoican, 2022; Klein, Kozhan, Viswanath-Natraj, and Wang, 2023; Aoyagi and Ito, 2025; Easley, O'Hara, Yang, and Zhang, 2026; Barbon and Ranaldo, 2026). A second studies liquidity provision and price discovery on decentralized exchanges and automated market makers (Hasbrouck, Rivera, and Saleh, 2024; Capponi and Jia, 2025; Lehar and Parlour, 2025; Capponi, Jia, and Yu, 2026; Hasbrouck, Saleh, and Rivera, 2026). A third studies value extraction through transaction ordering and front-running, and the design flaws and responses it has prompted (Daian et al., 2020; Park, 2023; Capponi, Jia, and Wang, 2025). A fourth concerns on-chain settlement and the price oracles that feed external data to smart contracts (Chiu and Koeppl, 2019; Cong, Prasad, and Rabetti, 2025). Our mechanism is nonetheless distinct from these on-chain frictions: Polymarket runs a central limit order book, and the manipulation operates on the centralized spot exchange where the contract settles, working through trading into the settlement window rather than transaction reordering or a compromised oracle.

3 Institutional Background

3.1 Prediction markets and asset-price contracts

A prediction market lets participants trade contracts that pay off on the realized outcome of a future event, so the contract's price can be read directly as the market's estimated probability that the event occurs. The established contracts settle on real-world events (elections, policy decisions, sporting results) and are valued as forecasters whose prices have matched or beaten conventional methods (Wolfers and Zitzewitz, 2004; Arrow et al., 2008). One desirable feature stands out: the settlement

event lies outside the traders' control (it is *exogenous* to them).

The fastest-growing asset-price contracts are the cryptocurrency "up/down" markets: a yes-or-no bet on the direction of a reference asset (most heavily Bitcoin) over a fixed horizon, paying off only if the asset ends the horizon above where it began. The horizons are short, from minutes to hours, and the contracts open and settle automatically on a rolling schedule, so a single asset generates a continuous stream of fresh binary markets through the day.[6]

The same structure is now reaching mainstream finance. In March 2026 Nasdaq filed with the U.S. Securities and Exchange Commission to list binary options on the Nasdaq-100 and Nasdaq-100 Micro indices, and Cboe is pursuing comparable "all-or-none" contracts on financial and economic events; like the crypto up/down markets, each pays a fixed amount if a stated price condition is met at expiry and nothing otherwise.[7] These proposals would carry the asset-price contract onto the most heavily traded U.S. equity indices, so the settlement vulnerability this paper documents in crypto would extend to far larger markets.

3.2 Venues

These contracts trade on venues that differ sharply in design. *Polymarket* is a decentralized, blockchain-based market: its contracts are created and settled automatically by code that runs on a public blockchain (on-chain smart contracts), every trade is recorded on that blockchain, and positions are held in a token pegged to the U.S. dollar (a dollar stablecoin) (Tsang and Yang, 2026; Dubach, 2026). It lists the widest menu of horizons, including five-minute, fifteen-minute, and four-hour up/down contracts on Bitcoin and several other coins. *Kalshi* is a centralized, federally regulated U.S. exchange (trades route through one operator rather than a blockchain) (Diercks, Katz, and Wright, 2026; Commodity Futures Trading Commission, 2026); over our sample it lists only a

[6]Within months of launch, Polymarket's five- and fifteen-minute crypto up/down markets traded over $4 billion cumulatively and roughly tripled platform daily volume; sources: Polymarket and Chainlink launch announcements, 2026.

[7]Nasdaq's proposed contracts are priced between one cent and one dollar on the Nasdaq-100 and Nasdaq-100 Micro indices (SEC rule-change filing, March 2026); Cboe has announced "all-or-none" event contracts targeting a 2026 launch, subject to regulatory approval. Sources: Nasdaq and Cboe SEC filings and announcements, 2026.

fifteen-minute crypto up/down contract, with no five-minute or four-hour equivalent; its hourly and daily crypto series are price-range markets, not binary up/down.

3.3 Polymarket

Each Polymarket crypto up/down contract pays $1 if the reference asset's price at the close exceeds its price at the open, and $0 otherwise. The contracts are created and settled programmatically on a rolling schedule, so a fresh five-minute Bitcoin contract opens every five minutes around the clock. The three horizons trade simultaneously and overlap on the clock (every fifteen-minute contract spans three consecutive five-minute contracts), a feature we exploit below.

Resolution. The five-minute, fifteen-minute, and four-hour contracts settle against a price produced by a decentralized oracle (an outside service that reports a reference price by combining data from many exchanges) rather than against any single exchange. The oracle, Chainlink, aggregates trade data from many centralized and decentralized venues (including the largest spot exchanges) and reduces them to one consensus print through a multi-stage procedure (Cong, Prasad, and Rabetti, 2025; Aspembitova and Bentley, 2022): several professional aggregators each compute a volume-weighted average price (a price that weights each trade by its size) across the venues they cover, independent node operators take the middle value (the median) across aggregators, and the network publishes the middle value across nodes.[8] The strike against which a contract resolves is the oracle price sampled at the contract's open, and the settlement value is the oracle price at its close; the "up" side wins if and only if the close print exceeds the strike.[9] (The one-hour and daily contracts, which fall outside our sample, instead reference a single exchange's spot price directly.) This aggregated construction is central to the manipulation question we study. Because

[8]For the fast markets Polymarket reads Chainlink Data Streams, a low-latency pull-based feed, rather than the standard on-chain push feed, so the settlement value can be sampled at the exact close; see Chainlink Data Feeds documentation, `https://docs.chain.link/data-feeds`.

[9]During our sample period between July 1, 2025 and April 8, 2026, Polymarket set the strike and the settlement value using snapshots of Chainlink's spot-price feeds. Beginning August 7, 2026, it switched to Chainlink-computed TWAPs at those same boundaries. The initial lookback was 30 seconds for five-minute markets and 60 seconds for fifteen-minute and four-hour markets. On August 14, the five-minute lookback also became 60 seconds (Polymarket, 2026).

the resolution price combines many exchanges into one agreed figure, moving it appears to require moving the price on many exchanges at once rather than on just one. Yet within a five-minute horizon, arbitrage (traders profiting from price gaps across exchanges) keeps the largest single venue and the combined price within a few basis points (hundredths of a percent) of each other. A spot-price move concentrated on one venue can therefore still push the combined price across the strike if the move is large enough relative to that narrow margin. Appendix D measures this gap and the residual disagreement directly.

Product rollout. The three Bitcoin contracts launched in sequence, and that staggered introduction is the backbone of our between-period comparison. The fifteen-minute contract began trading on October 9, 2025 and the four-hour contract on October 15, 2025; the five-minute contract was created on-chain in December 2025 but sat dormant, with no trading, until its public launch on February 12, 2026.[10] We therefore partition the sample into three periods: a pre-launch baseline before any of these contracts traded, an intermediate period with the fifteen-minute and four-hour contracts live, and a final period with all three (including the five-minute contract) trading together.

4 Theoretical Setup

Two markets operate in parallel. On a spot market a risky asset changes hands; on a separate venue a binary contract settles on that asset's price at the close of trading, with some noise. The settlement rule is what links them: the contract settles on a price the spot market itself discovers—and a trader who holds the contract can move that price by trading the asset. An uninformed manipulator turns this into a strategy. She takes a position in the contract, then trades the spot in the final moments before settlement to push its price toward the side she holds.

[10] Launch dates are identified from the on-chain trade record; see Section 6.2.

4.1 The Two Markets and the Contract

This section describes the two assets the model needs—a risky spot asset and a binary contract written on it—and the rule that settles the contract.

The first asset is the spot asset itself. The fundamental value of the spot asset is a random variable V with equal probability of being $-\sigma$ or $+\sigma$, drawn by nature at stage 0. $\sigma > 0$ is the volatility of the fundamental value and exogenous. One informed spot trader observes the realization. Everyone else holds the symmetric prior. Spot positions ultimately liquidate at the value of V.

The second asset is the binary contract—a claim on the asset's direction. Crucially, it does not settle on the realized value V but on the spot market's *posterior mean* of V at the end of the prediction horizon, that is, the value the spot market's own trading has revealed, perturbed by a small independent noise. This is the link the manipulator exploits: by moving the closing spot price, she moves the very quantity her contract settles on. Write h for the public order flow and $\bar{V}_h$ for the spot market's posterior mean of V given h (both made precise in Section 4.5).

For a final price path h, the binary contract pays

$$D_\varepsilon(h) = \begin{cases} 0, & \bar{V}_h + \varepsilon\xi < 0, \\ \frac{1}{2}, & \bar{V}_h + \varepsilon\xi = 0, \\ 1, & \bar{V}_h + \varepsilon\xi > 0, \end{cases} \tag{1}$$

where ξ is a random noise term uniformly distributed on $[-1, 1]$, independent of V and all orders, and $\varepsilon > 0$ is an exogenous parameter that controls the noise level.[11]

The payoff in (1) makes the contract a *binary* (or digital) option on the spot price: it pays out only when the settlement value $\bar{V}_h + \varepsilon\xi$ clears a fixed threshold, here zero. We therefore adopt standard option terminology throughout. We call that threshold the *strike*; the contract finishes

[11] In the no-noise limit $\varepsilon \downarrow 0$, the expected payoff jumps when $\bar{V}_h$ crosses zero and we might not have an equilibrium. The noise smooths the payoff and guarantees the existence of an equilibrium. In the large-noise limit $\varepsilon \uparrow \infty$, the binary contract decouples from the spot market and its expected payoff converges to $1/2$ regardless of the price path, which makes manipulation unprofitable.

in-the-money when the settlement value is above the strike (it pays 1), *out-of-the-money* when below it (it pays 0), and *at-the-money* when the price sits at the strike, where the outcome is still in doubt.

Why settle on a noisy price rather than on the spot price directly? In practice the settlement reference is not the manipulator's to dictate. The oracle that resolves the contract reads many exchanges at once—not just the one venue she can trade—and her push lands amid a stream of other traders' orders she does not control. A push therefore carries the settlement across the strike only with some probability, never for sure. The term $\varepsilon\xi$ captures exactly this residual uncertainty, and ε sets how much of it remains.

Let H denote the cumulative distribution function of ξ. The expected payoff of the PM contract given h is

$$\mathbb{E}\left[D_\varepsilon(h) \mid h\right] = H\left(\frac{\bar{V}_h}{\varepsilon}\right) = \begin{cases} 0, & \bar{V}_h \le -\varepsilon, \\ \frac{1}{2} + \frac{\bar{V}_h}{2\varepsilon}, & -\varepsilon < \bar{V}_h < \varepsilon, \\ 1, & \bar{V}_h \ge \varepsilon. \end{cases} \tag{2}$$

That is, the contract pays 1 if the final posterior mean is weakly above ε, 0 weakly below $-\varepsilon$, and interpolates linearly on the band of width 2ε around zero.

4.2 Timeline

The model has five types of participants—market makers on the spot and prediction-market venues, an informed spot trader, a manipulator, and liquidity traders—who act over five stages (Figure 2). Each order below is publicly observed. At each trading stage the venue's market maker has posted bid and ask quotes, and the arriving order executes against them. Later, Section 4.3 specifies how each participant behaves strategically; here we only lay out the sequence.

Stage 0 (Nature): Draws $V \in \{-\sigma, +\sigma\}$, observed only by the informed spot trader.

Stage 1 (Prediction market): One order Y_1^{P} arrives on the prediction market (PM), from the manipulator or a liquidity trader.

Stage 2 (Spot market): One order Y_2^{S} arrives on the spot market, from the informed trader or a liquidity trader.

Stage 3 (Spot market during settlement window): One order Y_3^{S} arrives on the spot market just before the close: the manipulator if she entered at stage 1, otherwise an ordinary spot order like stage 2's.

Stage 4 (Settlement): The binary contract settles on the spot market's posterior mean of V plus a small noise after stage 3 ends. All spot positions liquidate at the realized value of V.

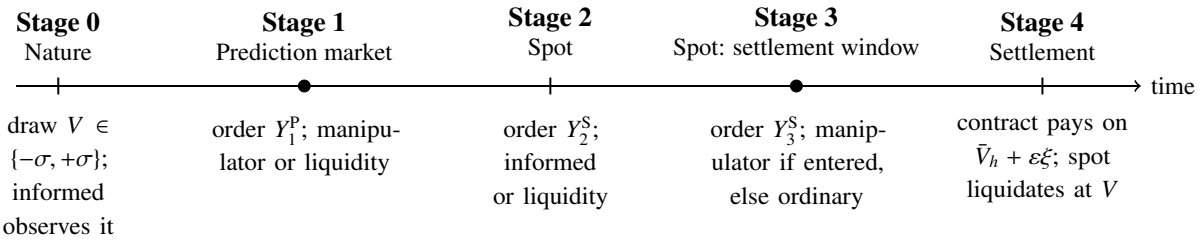

Figure 2. The five-stage timeline. Filled dots mark the manipulator's two decision points: entry on the prediction market (Stage 1) and the settlement-window spot trade (Stage 3).

4.3 Players

Market makers. A competitive, risk-neutral market maker stands on each venue and sees the same public information $\{\mathcal{F}_t\}$. Competition drives its expected profit to zero, so it quotes a break-even price—a one-unit bid and a one-unit ask equal to the asset's expected value given the order it faces. One qualification is required: the book cannot be crossed, so the ask never falls below the bid. Where the break-even quotes satisfy this, the maker posts them and earns nothing; where they would cross, it cannot—and the tightest feasible book, a zero spread, lets it earn a rent instead. It cannot distinguish informed or strategic orders from liquidity trades, so the quote conditions on the order alone. The spot maker's quotes are $\left(B_t^{\mathrm{S}}, A_t^{\mathrm{S}}\right)$: at stage 2, $\left(B_2^{\mathrm{S}}, A_2^{\mathrm{S}}\right)$; at stage 3, after stage-2 outcome h_2, (B_{h_2}, A_{h_2}). The prediction-market maker's stage-1 quotes are $\left(B_1^{\mathrm{P}}, A_1^{\mathrm{P}}\right)$.

Informed trader. The informed trader knows the realized V and trades to profit from it. She arrives with probability $\pi \in (0, 1)$ at each *ordinary* spot stage—stage 2 always, and stage 3 whenever

the manipulator has not entered—and follows the Glosten–Milgrom rule (Glosten and Milgrom, 1985),

$$\text{buy if } V > A_t^{\mathrm{S}}, \qquad \text{sell if } V < B_t^{\mathrm{S}}, \qquad \text{stay out otherwise.} \tag{3}$$

That is, she buys when the asset is worth more than the ask and sells when it is worth less. Because break-even quotes lie inside $[-\sigma, \sigma]$, one side always pays, so she never stays out.

Liquidity traders. Liquidity traders are uninformed and trade for reasons outside the model, buying or selling at random. One sits on each venue: on the prediction market a liquidity trader takes the stage-1 slot whenever the manipulator does not enter, and on the spot market a liquidity trader takes any ordinary spot slot the informed trader does not. Either way exactly one order prints per stage, so the stage-2 outcome is a buy or a sell.

Manipulator. The manipulator never observes V. She profits not from information but from the contract itself: by trading the spot in the settlement window she moves the price the contract pays on. At stage 1 she enters the prediction market with probability $p_{\mathrm{M}} \in [0, 1]$, buying or selling the contract; with probability $1 - p_{\mathrm{M}}$ a liquidity trader takes the stage-1 slot instead. The prediction-market maker sees only the order, not who sent it, so it cannot tell an entering manipulator from a liquidity trader. Conditioning on a stage-1 buy (the sell branch is symmetric), let $M_t(\delta) = \mathbb{P}(\Delta_t = \delta \mid \mathcal{F}_t)$ be the public belief that her contract inventory after stage t is $\delta \in \{-1, 0, +1\}$:

$$M_1(-1) = 0, \qquad M_1(0) = 1 - p_{\mathrm{M}}, \qquad M_1(+1) = p_{\mathrm{M}}.$$

She does not trade at stage 2, so $M_{3^-} = M_1$. What happens at stage 3 depends on whether she entered. If she did, she alone takes the settlement-window slot and chooses how hard to push: with probability $\alpha \in [0, 1]$ she buys, moving the close toward the side of the contract she holds, and otherwise does nothing. This manipulation probability is her second choice variable, alongside the entry probability; the equilibrium pins down both. If she did not enter, stage 3 is an ordinary spot

stage, exactly like stage 2: the informed trader with probability π, a liquidity trader otherwise.

4.4 Equilibrium

Definition 1 (Equilibrium)**.** An *equilibrium* is a manipulator strategy (her entry probability p_{M} and her manipulation probability α), a quoting rule for each market maker, and a system of beliefs, taking the informed trader's rule in (3) and liquidity-trader behavior as given, such that:

1. *sequential rationality*—the manipulator's trades maximize her expected profit given the makers' quotes and the beliefs;
2. *zero-profit quoting*—each maker's bid and ask equal the expected value conditional on the order it faces, subject to a non-crossing book (the ask at least the bid), so that where the unconstrained break-even quotes would cross, the maker quotes a zero spread and earns the resulting rent;
3. *belief consistency*—beliefs are formed from the strategies by Bayes' rule on every information set reached with positive probability.

Where condition (i) leaves more than one manipulation probability α consistent with the quotes, we select the smallest one that is stable to small perturbations of the conjecture, following Zhu (2014).

4.5 Price Paths and Notation

Each spot order has one of three outcomes—a buy $+$, a sell $-$, or no trade $\mathtt{0}$—and we write $O = \{+, -, \mathtt{0}\}$ for this set. At an ordinary spot stage some order always trades (no trade has probability zero), so the outcome is a buy or a sell; write $O_2 = \{+, -\}$ for this smaller set. A spot price path h records the realized outcomes—the stage-2 outcome h_2, and after the settlement window the pair h_2h_3. Given a price path h, the posterior probability of the up state is $\rho_h = \mathbb{P}(V = +\sigma \mid h)$;

because V takes only the two values $\pm\sigma$, this one number fixes the posterior mean and variance,

$$\bar{V}_h = \sigma\,(2\rho_h - 1)\,, \qquad \nu_h = \sigma^2(1 - (2\rho_h - 1)^2) = 4\sigma^2\rho_h\,(1 - \rho_h)\,.$$

Finally, two price-path probabilities recur. We write q_{h_2} for the probability that the stage-2 price path is h_2, and $q_{h_2h_3}$ for the *conditional* probability that the settlement window then produces outcome h_3, given h_2; their product $q_{h_2}q_{h_2h_3}$ is the probability of the full price path h_2h_3.

We restrict attention throughout to the empirically relevant region: adverse selection is moderate, and the settlement noise is small.

Assumption 1. $\varepsilon < \sigma\pi < \frac{1}{2}$.

5 Theoretical Results

This section solves the model and then interprets the equilibrium. We first solve the manipulator's settlement-window push for a fixed entry probability (Section 5.1). We then show that manipulator participation reduces settlement-time spot price discovery (Section 5.2), endogenize entry in the equilibrium and characterize the transfer and liquidity asymmetry it produces (Section 5.3), use a second ordinary spot stage to study the contract horizon as the design lever, and generalize the horizon lever to n ordinary spot stages (Section 5.4). Throughout, proofs are deferred to the appendix.

5.1 The Manipulator's Settlement-Window Push

The prediction-market contract settles on the spot price discovered at the close, so the manipulator, who holds the contract, can move what it pays by trading the spot asset.

She acts in the settlement window: stage 3, the spot trade in the final moments before the close. She reaches the window holding the prediction-market contract she bought at stage 1, and we condition throughout on her having bought it (a long position); the short-contract branch is

symmetric. Her one decision is whether to *push*—to buy the spot so as to lift the close toward the side she holds—and how often. The push has both a payoff and a price: it lifts the settlement-time spot price, and with it her contract's expected payoff, but she must pay the spot maker's ask, which exceeds the price the spot has discovered. Figure 3 lays out the resulting continuation game: the stage-2 spot order leaves an up or a down public posterior, after which the manipulator, if she entered, chooses whether to push.

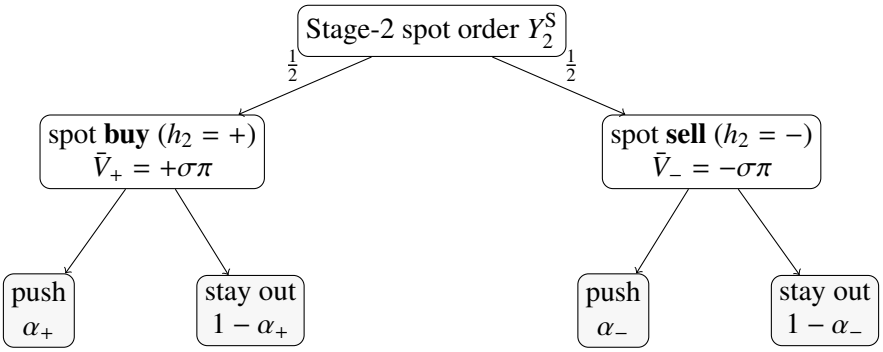

Figure 3. The stage-3 continuation game, conditional on a stage-1 manipulator buy on the prediction market. The stage-2 spot order produces an up posterior ($h_2 = +$) or a down posterior ($h_2 = -$), each with probability $\frac{1}{2}$. If the manipulator entered at stage 1 (probability p_{M}), she then either pushes the spot with the buy probability α_{h_2} derived below or stays out; otherwise the settlement-window order is an ordinary spot order (an informed trader or a liquidity trader), exactly as at stage 2.

The object we solve for is a *continuation equilibrium*: a set of stage-3 spot quotes and a stage-3 strategy for the manipulator that fit together. Three requirements make them fit: the spot maker quotes break-even prices, equal to the asset's expected value given how the manipulator is expected to trade, subject to a non-crossing book; the manipulator trades optimally given those prices; and the spot maker's beliefs follow Bayes' rule. Two probabilities drive the analysis. The first, p_{M}, is the probability the manipulator *enters* the contract at stage 1. The second, α, is the probability she *pushes* the spot at stage 3 once she has entered. We solve this continuation game with p_{M} held fixed, then let stage 1 choose p_{M} in Section 5.3.

We begin with the spot market's ordinary Glosten–Milgrom round at stage 2, whose only job is to fix the beliefs that stage 3 starts from. A single spot order arrives: with probability π the informed trader, who has seen the asset and buys when it is worth more than the ask and sells when it is worth less; otherwise a liquidity trader, who trades at random. The spot maker cannot tell the two apart, so

it quotes to break even, setting each quote to the asset's expected value given the order it sees. Only the informed trader's order carries news, so a buy raises the posterior that the asset is high and a sell lowers it, and the break-even quotes are the matching posterior means. Lemma 1 states these quotes and the posteriors stage 3 inherits, with the derivation in Appendix A.2.

Lemma 1 (Stage-2 quotes and posteriors)**.** *The stage-2 break-even quotes are the ask* $A_2^{\mathrm{S}} = \sigma\pi$ *and the bid* $B_2^{\mathrm{S}} = -\sigma\pi$*. Each outcome* $h_2 \in \{+, -\}$ *occurs with probability* $q_{h_2} = \frac{1}{2}$ *and leaves the spot posterior*

$$\rho_{h_2} = \frac{1 \pm \pi}{2}, \qquad \bar{V}_{h_2} = \pm\sigma\pi, \qquad \nu_{h_2} = \sigma^2\left(1 - \pi^2\right),$$

with the plus sign after a buy and the minus sign after a sell. The posterior mean $\bar{V}_{h_2}$ *is the price the spot market has discovered, and the residual variance* ν_{h_2} *is the uncertainty about the asset that remains.*

Stage 2 leaves the spot maker with the posterior ρ_{h_2} from Lemma 1, so stage 3 opens from that belief. The manipulator's stage-3 action remains to be pinned down. She would never sell the spot, because a sell moves the close downward, the wrong way for the long contract she holds, and so can only cost her. She is therefore left to buy, to stay out, or to randomize between the two, so her whole stage-3 strategy collapses to a single number, the probability α_{h_2} that she pushes (buys the spot) after stage-2 outcome h_2.

Before pricing the interior case, two values of this push probability can be set aside. Take $p_{\mathrm{M}} = 1$ first: the lone stage-3 trader is then the manipulator, who is uninformed about the asset, so her buy reveals nothing and the break-even ask equals the prior mean $\bar{V}_{h_2}$, so her push probability does not affect the price. Take $p_{\mathrm{M}} = 0$ next: she is absent from stage 3 altogether, so there is no push. We then fix $p_{\mathrm{M}} \in (0, 1)$ for the rest of the analysis.

Consider next the price the spot maker sets for the single spot order that arrives, an order whose sender it cannot see. A buy can come from any of three types. It can come from the informed trader, who buys only when the asset is truly high; from a liquidity trader, who buys at random; or from the manipulator, who pushes with probability α_{h_2} to lift the close. Of the three, only the informed

trader's buy carries information about the asset, because the liquidity trader and the manipulator both trade for reasons unrelated to value and are therefore uninformed. The spot maker cannot tell which type sent the order, so it prices against this pooled flow and sets each quote to break even, earning zero in expectation given the order it observes. Imposing that break-even condition pins the quotes down, as Proposition 1 records.

Proposition 1 (Stage-3 quotes and posteriors)**.** *Fix the stage-2 outcome $h_2 \in \{+, -\}$, an interior entry probability $p_{\mathrm{M}} \in (0, 1)$, and a conjectured stage-3 push probability $\alpha_{h_2} \in [0, 1]$. The unique break-even stage-3 quotes are*

$$A_+(\alpha_+; p_{\mathrm{M}}) = \frac{\sigma\pi\left[p_{\mathrm{M}}\alpha_+ + (1 - p_{\mathrm{M}})\right]}{p_{\mathrm{M}}\alpha_+ + (1 - p_{\mathrm{M}})\frac{1+\pi^2}{2}}, \qquad A_-(\alpha_-; p_{\mathrm{M}}) = \frac{-\sigma\pi\, p_{\mathrm{M}}\alpha_-}{p_{\mathrm{M}}\alpha_- + (1 - p_{\mathrm{M}})\frac{1-\pi^2}{2}},$$

$$B_+ = 0, \qquad B_- = \frac{-2\sigma\pi}{1 + \pi^2}.$$

The break-even ask balances the spot maker's expected gain when it sells to an uninformed buyer (a liquidity trader or the manipulator) against its expected loss when it sells to the informed trader. The corollary below traces how the ask responds as the manipulator is expected to push more aggressively.

Corollary 1 (Stage-3 quote monotonicity)**.** *The ask falls as the manipulator is expected to push harder: for each h_2, the ask $A_{h_2}(\alpha; p_{\mathrm{M}})$ is strictly decreasing in the conjectured push α and weakly decreasing in the entry probability p_{M} (strictly when $\alpha > 0$), while the bid B_{h_2} and the no-trade posterior $\bar{V}_{h_2\mathtt{0}} = \bar{V}_{h_2}$ are unaffected.*

The more often the manipulator is expected to push, the more uninformed buys are pooled in with the informed trader's, so any given buy is less likely to be hers. A buy therefore carries less adverse selection risk, and so the spot maker lowers its ask.

The question now is how hard the manipulator actually chooses to push. She pushes only when pushing pays, so we compare the payoff from buying with the payoff from staying out. The difference is the best-response gap $\Gamma_{h_2}(\alpha_{h_2})$, the rise in her contract's expected payoff from a higher

close net of the spot half-spread $A_{h_2} - \bar{V}_{h_2}$ she pays to push:

$$\Gamma_{h_2}(\alpha_{h_2}) = \underbrace{H\left(\frac{A_{h_2}(\alpha_{h_2})}{\varepsilon}\right) - \left[A_{h_2}(\alpha_{h_2}) - \bar{V}_{h_2}\right]}_{\substack{\text{expected payoff} \\ \text{from buying}}} - \underbrace{H\left(\frac{\bar{V}_{h_2}}{\varepsilon}\right)}_{\substack{\text{expected payoff} \\ \text{from staying out}}}$$
$$= \underbrace{H\left(\frac{A_{h_2}(\alpha_{h_2})}{\varepsilon}\right) - H\left(\frac{\bar{V}_{h_2}}{\varepsilon}\right)}_{\substack{\text{increase in expected} \\ \text{contract payoff from buying}}} - \underbrace{\left[A_{h_2}(\alpha_{h_2}) - \bar{V}_{h_2}\right]}_{\text{spot cost}}.$$

Note that the price the manipulator pays at stage-1 to enter the prediction market is sunk for the purpose of solving her stage-3 optimality, so it does not enter the gap.

In equilibrium the spot maker's conjecture must be borne out: the push the manipulator actually chooses must equal the one the spot maker priced against. Her α_{h_2} is therefore a *fixed point* of the gap. She stays out when pushing never pays ($\Gamma_{h_2}(0) \leq 0$), pushes for sure when it always pays ($\Gamma_{h_2}(1) \geq 0$), and mixes only at an interior probability where the two exactly tie ($\Gamma_{h_2}(\alpha_{h_2}) = 0$). Because both the gain and the cost shrink as the conjecture rises, the gap can cross zero more than once; following Zhu (2014), the equilibrium keeps the smallest *stable* fixed point, and we write $\mathcal{E}_{h_2}$ for the set of stable fixed points.

Proposition 2 (Selected stage-3 continuation equilibrium)**.** *Fix an interior entry probability $p_M \in (0, 1)$. The set of stable fixed points $\mathcal{E}_{h_2}$ is nonempty for each h_2, so the selected push $\alpha^*_{h_2} = \min \mathcal{E}_{h_2}$ is well defined.*

Existence is the standard fixed-point argument applied to the continuous gap Γ_{h_2}, and the equilibrium takes its smallest stable crossing. This pins down the continuation $\alpha^*_{h_2}$ for every entry probability p_M, which stage 1 then chooses in Section 5.3.

5.2 Price Discovery Falls

The manipulator's mere presence makes the settlement-time spot price strictly less informative, for any entry probability $p_M \in (0, 1)$.

To make "less informative" precise, we measure how much uncertainty about the asset's value V the spot price still leaves once stage 3 closes. This is the *expected residual variance* $\mathbb{E}[\nu_{h_2h_3}]$, the variance of V that remains after conditioning on the discovered price, averaged over outcomes:

$$\mathbb{E}[\nu_{h_2h_3}] = \sum_{h_2} q_{h_2} \sum_{h_3} q_{h_2h_3}\, \nu_{h_2h_3}.$$

By the law of total variance, this residual variance and the variance of the discovered price sum to the asset's prior variance, $\nu_\emptyset = \mathbb{E}[\nu_{h_2h_3}] + \mathrm{Var}\left(\bar{V}_{h_2h_3}\right)$, with the total fixed at $\nu_\emptyset = \sigma^2$. A price that tracks V closely leaves little residual variance, so a *higher* residual variance means a *less* informative price. The proposition below shows that any interior entry raises it above the no-entry benchmark $p_{\mathrm{M}} = 0$.[12]

Proposition 3 (Manipulator participation strictly harms price discovery)**.** *With informed traders present ($\pi > 0$), for every interior entry probability $p_{\mathrm{M}} \in (0, 1)$ and every conjectured profile of stage-3 push probabilities $(\alpha_{h_2})_{h_2 \in O_2}$,*

$$\mathbb{E}[\nu_{h_2h_3}](p_{\mathrm{M}}) > \mathbb{E}[\nu_{h_2h_3}](0).$$

The loss runs through two channels. The first is displacement: the manipulator occupies the stage-3 slot with probability p_{M}, thinning the informed order flow by the factor $1 - p_{\mathrm{M}}$, and this alone makes the loss strict even when she never pushes ($\alpha_{h_2} = 0$). The second is dilution: once she does push ($\alpha_{h_2} > 0$), her uninformed buys pool with the informed trader's, so a buy reveals less about the asset. The benchmark $p_{\mathrm{M}} = 0$ shuts down both.

This is the mirror image of Zhu's dark-pool result. In Zhu (2014), the second venue draws relatively uninformed flow away from the exchange, concentrating informed traders there and sharpening the exchange price. Here the linked contract feeds uninformed, settlement-motivated flow into the price-setting market, so the settlement-time spot price becomes less informative. A

[12] From the spot market's perspective, setting $p_{\mathrm{M}} = 0$ is equivalent to shutting down the prediction market, since it is the only interaction channel between the two.

derivative that settles on a market's price can therefore degrade the informativeness of the price it references.

5.3 The Manipulation Equilibrium

We then solve the stage-1 entry-and-pricing game. The manipulator chooses how often to enter the prediction market. The prediction-market maker, unable to tell the manipulator from an ordinary buyer, quotes a price. Two results emerge. The first is the distinctive feature of the model. The prediction market posts a *zero* bid-ask spread, yet the prediction-market maker still earns a positive rent while the manipulator earns a positive profit. The second is the welfare counterpart. A single uninformed outsider, the liquidity trader, funds both the rent and the profit.

To solve the game we step back to stage 1. Two kinds of buyer arrive there, and the prediction-market maker cannot tell them apart. One is the manipulator. When ordinary flow has gone against the side she holds, she pushes the spot at stage 3, so her contract is worth a high expected payoff. The other is an uninformed liquidity trader. She has no such push, so her contract may be worth far less. The prediction-market maker prices a single contract knowing it may be either type, and we work out three stage-1 objects in turn: the two types' expected payoffs, the break-even price the prediction-market maker posts, and the manipulator's entry decision.

Start with the prediction-market maker's problem. Fix the manipulator's entry probability $p_{\mathrm{M}} \in (0, 1)$. Both payoffs are built from the contract's settlement value H, defined in (2). There, $H(x)$ is the contract's expected payoff when the posterior mean sits x settlement-noise units above the strike, that is, the probability the contract settles in-the-money. If the manipulator enters, her expected PM contract payoff after stage-2 outcome h_2 is

$$\Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}}, h_2) = \alpha_{h_2}^{*} H\left(\frac{A_{h_2}^{*}}{\varepsilon}\right) + \left[1 - \alpha_{h_2}^{*}\right] H\left(\frac{\bar{V}_{h_2}}{\varepsilon}\right).$$

This averages her contract value across two events. When the stage-2 price path has gone against her, she pushes the spot at stage 3. When ordinary flow already carries the price her way, she does

not. The liquidity trader has no such ability, so she faces ordinary settlement, and her expected PM contract payoff after h_2 is

$$\Pi_{\mathrm{L}}^{+}(p_{\mathrm{M}}, h_2) = \left[\frac{1-\pi}{2} + \pi\rho_{h_2}\right] H\left(\frac{A_{h_2}^{*}}{\varepsilon}\right) + \left[\frac{1-\pi}{2} + \pi(1-\rho_{h_2})\right] H\left(\frac{B_{h_2}}{\varepsilon}\right),$$

where ρ_{h_2} is the stage-2 posterior that settlement favors her side. Averaging over the stage-2 outcome h_2 using the outcome probabilities q_{h_2} as weights gives the unconditional expected payoff for each stage-1 type:

$$\Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}}) = \sum_{h_2} q_{h_2} \Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}}, h_2), \qquad \Pi_{\mathrm{L}}^{+}(p_{\mathrm{M}}) = \sum_{h_2} q_{h_2} \Pi_{\mathrm{L}}^{+}(p_{\mathrm{M}}, h_2).$$

The competitive prediction-market maker sees only that someone wishes to buy, so it posts a single price that breaks even over the mix of the two types. That break-even ask is the weighted average of the two payoffs:

$$A_1^{\mathrm{P,BE}}(p_{\mathrm{M}}) = p_{\mathrm{M}} \Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}}) + (1 - p_{\mathrm{M}}) \Pi_{\mathrm{L}}^{+}(p_{\mathrm{M}}).$$

This break-even ask can fall strictly below $1/2$, which implies a break-even bid strictly above $1/2$ and therefore strictly above the break-even ask. But the prediction market cannot post a crossed book, so the posted ask is floored at $1/2$:

$$A_1^{\mathrm{P*}}(p_{\mathrm{M}}) = \max\left\{A_1^{\mathrm{P,BE}}(p_{\mathrm{M}}), \tfrac{1}{2}\right\}, \qquad B_1^{\mathrm{P*}}(p_{\mathrm{M}}) = 1 - A_1^{\mathrm{P*}}(p_{\mathrm{M}}).$$

The floor is a general property of the book, not a restriction we impose. When the break-even already exceeds $1/2$, the floor is slack, the posted ask equals the break-even ask, and the spread is positive, just as in an ordinary adverse-selection market. The striking case is the opposite one, and it is the one the model delivers: when the break-even sits below $1/2$, the non-crossing constraint binds, the prediction-market maker is pinned at $1/2$, and it keeps the gap $\frac{1}{2} - A_1^{\mathrm{P,BE}}$ as a strictly

positive rent on every contract sold, even though the posted spread is exactly zero.

Now turn to the manipulator's problem. Her profit from entering, paying the posted ask, and pushing at settlement is her expected contract payoff net of the spot-push cost and the price she paid:

$$S_{\mathrm{M}}^{+}(p_{\mathrm{M}}) = \left[\Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}}) - K^{+}(p_{\mathrm{M}})\right] - A_{1}^{\mathrm{P}*}(p_{\mathrm{M}}),$$

where K^{+} is the expected resources she spends moving the spot, summed over the two ordinary stage-2 outcomes $h_2 \in O_2 = \{+, -\}$:

$$K^{+}(p_{\mathrm{M}}) = \sum_{h_2 \in O_2} \underbrace{q_{h_2}\alpha_{h_2}^{*}}_{\substack{\text{probability} \\ \text{of manipulation}}} \underbrace{\left[A_{h_2}^{*} - \bar{V}_{h_2}\right]}_{\substack{\text{spot manipulation} \\ \text{cost}}} \geq 0.$$

She chooses her entry frequency p_{M} to maximize her ex-ante expected profit, the entry probability times the profit per entry:[13]

$$\bar{S}_{\mathrm{M}}^{+}(p_{\mathrm{M}}) = p_{\mathrm{M}}\, S_{\mathrm{M}}^{+}(p_{\mathrm{M}}), \qquad p_{\mathrm{M}}^{*} \in \arg\max_{p_{\mathrm{M}} \in [0,1]} \bar{S}_{\mathrm{M}}^{+}(p_{\mathrm{M}}).$$

For compactness, write the pure-push cutoff derived in Appendix A.5.2 as

$$p_0 := -\left(1 + \frac{2(\sigma\pi - \varepsilon)}{\varepsilon(1-\pi^2)(1-2\sigma\pi)}\right)^{-1}.$$

The following proposition characterizes the full equilibrium.

Proposition 4 (The manipulation equilibrium)**.** *Fix $\sigma > 0$, $\pi \in (0,1)$, and $\varepsilon > 0$ satisfying Assumption 1. The equilibrium of Definition 1 has:*

1. *optimal entry probability*

$$p_{\mathrm{M}}^{*} = \sqrt{\left(\frac{1-\pi^2}{1+\pi^2}\right)^2 + \frac{1-\pi^2}{1+\pi^2}p_0} - \frac{1-\pi^2}{1+\pi^2} \in (0, p_0);$$

[13] A maximizer exists whenever $\bar{S}_{\mathrm{M}}^{+}$ is continuous in p_{M}, which holds below, where $\bar{S}_{\mathrm{M}}^{+}$ is available in closed form. When the maximizer is not unique, we select the smallest one.

2. *push policy*

$$\alpha_{+}^{*} = 0, \qquad \alpha_{-}^{*} = 1;$$

3. *strictly positive expected profit*

$$\bar{S}_{\mathrm{M}}^{+}(p_{\mathrm{M}}^{*}) = \frac{(1 - 2\sigma\pi)\, p_{\mathrm{M}}^{*2}}{4p_0} > 0;$$

4. *zero-spread quotes for the prediction market:*

$$A_1^{\mathrm{P}*} = B_1^{\mathrm{P}*} = \frac{1}{2}.$$

The manipulator enters with an interior probability p_{M}^{*}. She enters often enough to profit, but not so often that a worse prediction-market ask and a smaller settlement gain from her stage-3 push erode her profit. Her push policy sits at a corner. She never pushes after a stage-2 buy ($\alpha_{+}^{*} = 0$), because ordinary flow has already carried the price her way, and there is nothing to manipulate. She always pushes after a stage-2 sell ($\alpha_{-}^{*} = 1$), because flow has gone against her side and a push can still move the settlement her way.

If both the prediction-market maker and the manipulator profit at a zero spread, a third party must fund them. That party is the uninformed liquidity trader. She buys at the same $1/2$ but has no settlement-window push of her own, so she expects to receive less than what she pays. Why does she lose in expectation, holding no information and no view? The pooling that does it sits at the *spot* maker, not the prediction-market maker. The manipulator is long the "up" contract and would lift the settlement price by buying the spot. A settlement-window buy becomes less informative, so by Corollary 1 the spot maker lowers its ask while the bid is unchanged. The result is a settlement bias *against* the "up" holder, including the liquidity trader even though she is innocent. She need not meet the manipulator, hold any view about the spot asset, or even know the spot market exists; simply holding the "up" contract exposes her to the spot-maker quote response. The harm is cross-market (Figure 4): a prediction-market position triggers a spot-maker quote response that

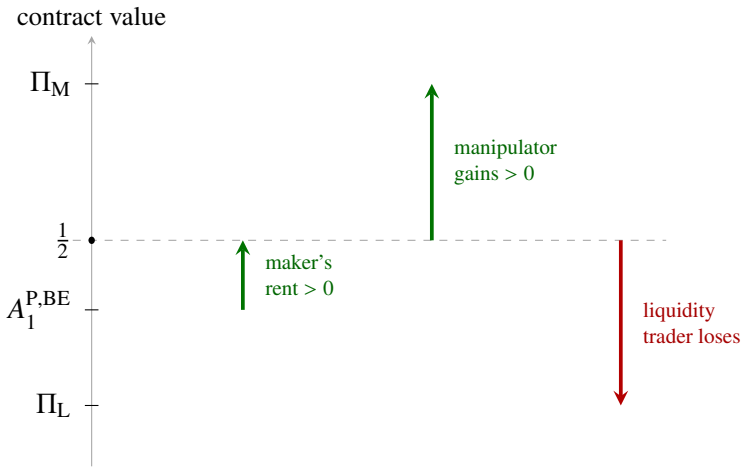

(a) The zero-spread rent: the break-even price lies below $1/2$, so the non-crossing floor pins the posted price at $1/2$ and the prediction-market maker keeps the gap as a positive rent.

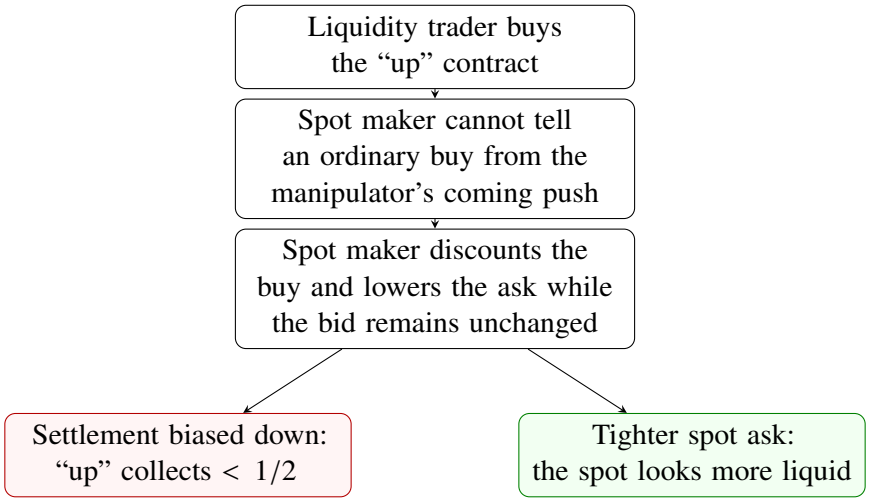

(b) The cross-market channel: a prediction-market "up" buy makes a settlement-window spot buy look possibly manipulative, so the spot maker lowers its ask; the settlement is biased against the "up" side (the holder collects $< 1/2$) while the spot ask tightens (more liquid).

Figure 4. The zero-spread rent and the cross-market channel. Panel (a) shows where the prediction-market maker's rent comes from; panel (b) shows how the uninformed liquidity trader's loss is inflicted through the spot maker's quotes.

tilts the settlement against that position, which harms the contract holder but makes the spot market more liquid.

The following proposition gives the accounting.

Proposition 5 (Liquidity-trader loss decomposition)**.** *At the equilibrium of Proposition 4, evaluated at the interior entry probability $p_M^* \in (0, p_0)$, the liquidity trader's ex-ante loss splits exactly into*

three components:

$$(1 - p_{\mathrm{M}}^{*})\, L^{+}(p_{\mathrm{M}}^{*}) = R_{\mathrm{P}}^{+*} + p_{\mathrm{M}}^{*} S_{\mathrm{M}}^{+}(p_{\mathrm{M}}^{*}) + p_{\mathrm{M}}^{*} K^{+}(p_{\mathrm{M}}^{*}),$$

where the prediction-market maker's rent, the liquidity trader's per-order loss, and the spot push cost are

$$R_{\mathrm{P}}^{+*} :- A_{1}^{\mathrm{P}*} - A_{1}^{\mathrm{P,BE}}(p_{\mathrm{M}}^{*}) = \frac{p_{\mathrm{M}}^{*}(\sigma\pi - \varepsilon)}{4\varepsilon} > 0,$$

$$L^{+}(p_{\mathrm{M}}^{*}) :- \frac{1}{2} - \Pi_{\mathrm{L}}^{+}(p_{\mathrm{M}}^{*}) = \frac{\left(1-\pi^{2}\right)\sigma\pi p_{\mathrm{M}}^{*}}{4\varepsilon\left[(1-\pi^{2}) + (1+\pi^{2})\, p_{\mathrm{M}}^{*}\right]} > 0,$$

$$K^{+}(p_{\mathrm{M}}^{*}) = \frac{1}{2}\,\frac{\sigma\pi\left(1-\pi^{2}\right)\left(1-p_{\mathrm{M}}^{*}\right)}{(1-\pi^{2}) + (1+\pi^{2})\, p_{\mathrm{M}}^{*}}.$$

The identity shows that the liquidity trader's expected loss funds the two rents and one cost on the right-hand side. The first rent is the prediction-market maker's R_{P}^{+*}, the gap between the posted $1/2$ and the break-even ask that the prediction-market maker keeps because it cannot quote any tighter under the non-crossing constraint. The second is the manipulator's ex-ante profit $p_{\mathrm{M}}^{*} S_{\mathrm{M}}^{+}$ while paying the same $1/2$. The third term, $p_{\mathrm{M}}^{*} K^{+}$, is the manipulator's ex-ante cost of moving the spot in the settlement-window push, which helps the spot maker finance the adverse-selection loss to the informed trader.

Having pinned down who pays, we next ask when manipulation incentives are strongest.

Corollary 2 (Comparative statics)**.** *The signs of the comparative statics are summarized in the table below.*

Quantity	*Symbol*	$\partial/\partial\varepsilon$	$\partial/\partial\sigma$	$\partial/\partial\pi$
Pure-push cutoff	p_0	>0	<0	<0
Optimal entry probability	p_{M}^*	>0	<0	<0
Stage-3 ask	$A_-^*\left(p_{\mathrm{M}}^*\right)$	<0	>0	>0
Spot manipulation cost at p_{M}^*	$A_-^*\left(p_{\mathrm{M}}^*\right)-\bar{V}_-$	<0	>0	>0
Ex-ante manipulation profit	$\bar{S}_{\mathrm{M}}^{+}\left(p_{\mathrm{M}}^*\right)$	>0	<0	<0

When settlement noise ε rises, the adverse outcome is easier to rescue, so the manipulator enters more often and p_{M}^* rises. Because she now enters more often, a stage-3 buy after a stage-2 sell is more likely to be manipulative and therefore less informative about the fundamental, so the stage-3 ask and the spot-push cost both fall. Even so, expected profit rises, because the gain from entering more often outweighs the lower profit per entry.

When fundamental volatility σ or information frequency π rises, the adverse outcome is harder to rescue. Each moves the stage-2 posterior further from the strike, so the settlement-window push is more costly. The manipulator enters less often, and expected profit falls.

The full comparative-statics derivations are in Appendix A.5.3.

5.4 The Horizon Lever

A venue cannot change a trader's signal or a manipulator's budget, but it can change one thing directly: the contract's *horizon*, the span of trading before the price settles. The horizon can be a design lever against manipulation. Each ordinary trading round before settlement lets the price absorb more information, so the posterior entering the settlement window sits further from the strike and a fixed-size push stays pivotal on a smaller set of potential paths. As the horizon grows, manipulation grows *rarer* and the settlement price *more* informative; lengthen the horizon without

bound and manipulation vanishes outright. We first solve the smallest extension, one extra round, where the mechanism is cleanest, then the general case of n ordinary rounds.

We lengthen the contract's horizon by inserting $n \geq 1$ ordinary trading rounds between the manipulator's entry and the settlement window. Each round is an ordinary Glosten–Milgrom round: an informed trader arrives with probability π and trades on her signal as in (3), and otherwise a liquidity trader buys or sells at random, so each round adds one buy + or one sell − to the public record. The manipulator still trades only in the settlement window. The baseline is the one-round case, $n = 1$.

A stage-2 price path with n rounds is denoted by a string of n buys or sells, $h_2 = h_{2.1} \cdots h_{2.n}$. What the market learns from it is summarized by a single number, the *order imbalance* $d\,(h_2)$, the number of buys minus the number of sells. The imbalance is a sufficient statistic for the posterior entering settlement: two price paths with the same imbalance leave the market at the same posterior, so a stage-2 price path can be summarized by d alone.

A fixed-size push moves the settlement price by a bounded amount, so it can change which side wins only when the price already sits near the strike—that is, only when the imbalance is near balance. Once the order flow has moved the imbalance far to one side, the outcome is settled and the manipulator stays out. The horizon works by making a near-balance imbalance rarer. Figure 5 lays out the timeline.

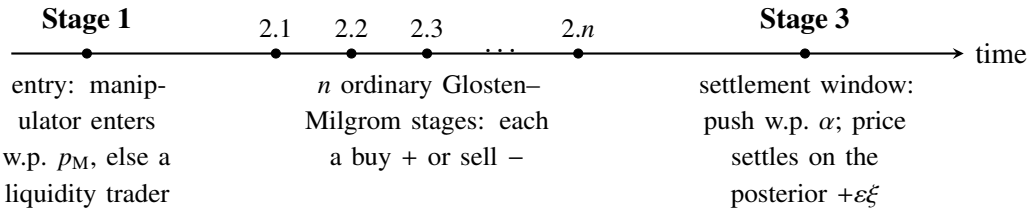

Figure 5. The extended game with n ordinary spot rounds between entry and settlement. The manipulator decides whether to enter at stage 1, n ordinary Glosten–Milgrom rounds then reveal order flow, and at stage 3 the manipulator (if she entered) pushes the spot price before the contract settles. The baseline (Section 5.3) is the case $n = 1$.

We solve the smallest case first—two rounds, where the few price paths can be enumerated and the mechanism is clearest—then the general n-round case. With two ordinary rounds the order flow can come out four ways. When the two orders agree—both buys or both sells—they move the

posterior far enough from the strike that the settlement is already decided. When they disagree—a buy and a sell—they cancel and leave the posterior at the strike, where a push can still change the settlement. So manipulation survives only on the two disagreeing paths. Figure 6 traces the four paths, and Appendix A.6.1 records the posteriors they reach.

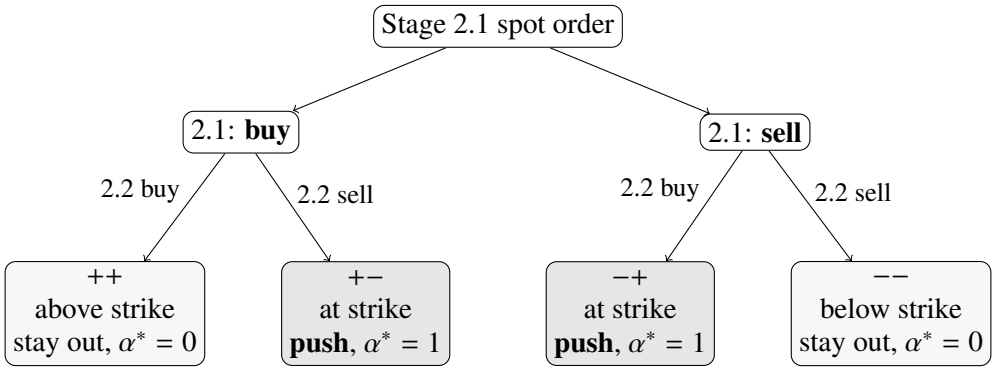

Figure 6. The lengthened game with two ordinary spot rounds before the settlement window, conditional on a stage-1 buy. The two ordinary orders (stage 2.1 then stage 2.2) produce four stage-2 price paths. The agreeing price paths ++ and −− move the posterior far enough from the strike that a push cannot change the settlement, so the manipulator stays out; the disagreeing price paths +− and −+ leave it at the strike, where the outcome stays undecided and she pushes.

The equilibrium of the two-round game is as follows.

Proposition 6 (Equilibrium with two ordinary spot rounds)**.** *Fix $\sigma > 0$, $\pi \in (0,1)$, and $\varepsilon > 0$ satisfying Assumption 1. The two-round game has a unique equilibrium (Definition 1), with a unique optimal entry probability $p_{\mathrm{M}}^{*(2)} \in (0,1)$. In this equilibrium:*

1. *she pushes only after the two ordinary orders disagree, and never after they agree:*

$$\alpha^*_{++} = \alpha^*_{--} = 0, \qquad \alpha^*_{+-} = \alpha^*_{-+} = 1;$$

2. *price discovery strictly improves relative to the baseline:*

$$\mathbb{E}\left[\nu^{(2)}_{h_2 h_3}\right]\left(p_{\mathrm{M}}^{*(2)}\right) < \mathbb{E}\left[\nu^{(1)}_{h_2 h_3}\right]\left(p_{\mathrm{M}}^{*(1)}\right).$$

The two-round case isolates the horizon lever in miniature: the extra ordinary round filters out the agreeing price paths, leaving only the disagreeing price paths vulnerable to manipulation. We

now solve the same logic for general n.

Proposition 7 (Equilibrium with $n > 2$ ordinary spot rounds)**.** *Fix $n > 2$, $\sigma > 0$, $\pi \in (0,1)$, and $\varepsilon > 0$ satisfying Assumption 1. The n-round game has a unique equilibrium (Definition 1), with a unique optimal entry probability $p_{\mathrm{M}}^{*(n)} \in (0,1)$. In this equilibrium:*

1. *she pushes only when the order imbalance is near balance, leaving the price close to the strike:*

$$\alpha_{h_2}^{*(n)} = \begin{cases} 1, & n \text{ is odd and } d(h_2) = -1, \\ 1, & n \text{ is even and } d(h_2) = 0, \\ 0, & \text{otherwise;} \end{cases}$$

2. *price discovery strictly improves with every added ordinary spot stage:*

$$\mathbb{E}\left[\nu_{h_2 h_3}^{(n)}\right]\left(p_{\mathrm{M}}^{*(n)}\right) < \mathbb{E}\left[\nu_{h_2 h_3}^{(n-1)}\right]\left(p_{\mathrm{M}}^{*(n-1)}\right).$$

3. *the entry probability differs by the parity of n:*

$$p_{\mathrm{M}}^{*(n)} = \begin{cases} p_{\mathrm{M}}^{*(1)}, & n \text{ is odd}, \\ p_{\mathrm{M}}^{*(2)}, & n \text{ is even}, \end{cases}$$

Two things matter. First, however long the horizon, a push pays on the price paths that leave the price closest to the strike but still on the wrong side of it: the balanced ($d = 0$) price paths when n is even and the one-step-below ($d = -1$) price paths when n is odd. Second, price discovery improves with every round: each additional round gives the market one more reading of the order flow, so the settlement price tracks the asset's value more closely, and the expected residual variance $\mathbb{E}[\nu_{h_2 h_3}^{(n)}]$ introduced in Section 5.2 falls. The entry probability itself does not settle down—it oscillates with the parity of the horizon, but that oscillation does not govern the contract's safety. What governs safety is how often a push can be pivotal at all.

That a push pays only on near-balanced price paths is why manipulation disappears as the horizon grows. The imbalance follows a random walk, moving ±1 each round (Figure 7). A push is pivotal only if the imbalance ends within one step of balance, but its dispersion grows with the horizon—the unconditional standard deviation of the imbalance is of order n—while the width of the pivotal band stays fixed. The probability of ending inside the band therefore falls to zero, and the manipulator's expected profit with it.

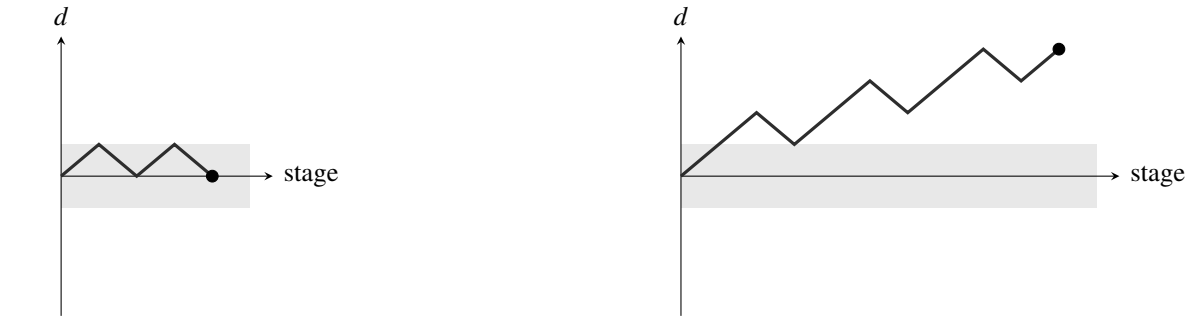

short horizon: the walk ends in the band ($|d| \leq 1$)—a push can still change the settlement

long horizon: the walk is likely to end far from balance, outside the band, so a push no longer affects the settlement

Figure 7. The pre-settlement order imbalance d as a ±1 random walk over the ordinary stages. The shaded strip is the pivotal band: a settlement-window push can change the settlement only when the walk ends within one step of balance ($|d| \leq 1$)—exactly $d = 0$ when n is even and $d = -1$ when n is odd. As the horizon n grows the walk spreads out (the imbalance has standard deviation of order n) while the band stays fixed, so the chance of ending inside it—the pivotal probability $m^{(n)}$—falls to zero.

Proposition 8 (Manipulation vanishes as the horizon grows)**.** *As the horizon grows, manipulation becomes both rarer and worthless. Write $m^{(n)}$ for the chance a push is still pivotal, $S_{\mathrm{M}}^{+(n)}$ for the manipulator's profit conditional on entry, and $\bar{S}_{\mathrm{M}}^{+(n)}$ for her unconditional expected profit at the optimal entry. All three vanish:*

$$\lim_{n\uparrow\infty} m^{(n)} = 0, \qquad \lim_{n\uparrow\infty} S_{\mathrm{M}}^{+(n)}\left(p_{\mathrm{M}}^{*(n)}\right) = 0, \qquad \lim_{n\uparrow\infty} \bar{S}_{\mathrm{M}}^{+(n)}\left(p_{\mathrm{M}}^{*(n)}\right) = 0.$$

The horizon is therefore a design lever for market safety. Lengthening a contract's horizon does not make manipulation impossible—a large enough push always moves the settlement price—but it makes a fixed-size push almost never pivotal. A contract is safe not when a manipulator cannot act,

but when enough price discovery runs before settlement that a single trader is unlikely to affect the outcome.

A robust settlement reference is a second, equally powerful lever. Our contract settles on the spot price at a single instant—the close—which is exactly what a final-seconds push can move. A reference that aggregates instead—averaged over a window, or struck in a closing auction that pools many participants' orders—forces the manipulator to move the whole aggregate rather than one print; a time-weighted oracle, for one, raises the manipulation cost roughly in proportion to the averaging window (Aspembitova and Bentley, 2022). Polymarket's switch to Chainlink TWAP feeds is a real-world instance of this lever (Polymarket, 2026). The two levers address manipulation from opposite sides: a longer horizon adds price discovery *before* settlement, while an aggregating reference spreads the settlement price *across* orders or time. The binary Nasdaq-100 index options the SEC approved in April 2026 (SR-MRX-2026-05) sit on the safe side of both: their expirations run from weekly to multi-year, and they settle on the index value struck in the Nasdaq Closing Cross, the exchange's closing auction, which pools the day's closing interest into one print and is correspondingly hard to push.[14]

6 Empirics

This section takes the model's predictions to the data. We first state what settlement manipulation should leave in the data according to the theory (Section 6.1). We then describe the sample (Section 6.2) and present the empirical evidence in two parts: Section 6.3 documents a near-settlement *push* in the spot market, defined as a burst of one-sided trading in the final seconds before a contract settles, and shows it bears the predicted signature: it appears only for the five-minute contract, concentrates in the cycles the prediction market has left undecided, causes price reversals

[14]Nasdaq MRX, LLC, *Order Granting Accelerated Approval of a Proposed Rule Change to Adopt New Options Rule 3B to List and Trade Binary Broad-Based Index Options*, Exchange Act Release No. 34-105342, File No. SR-MRX-2026-05 (Apr. 30, 2026). The settlement value is the Nasdaq-100 Index as reported at the conclusion of the Nasdaq Closing Cross (Nasdaq Equity 4, Rule 4757); the contracts are cash-settled, European-style binary options listed in weekly, monthly, quarterly, end-of-month, and long-term (12–60 month) expirations.

within seconds of settlement, and improves spot liquidity; Section 6.4 then shows the push is deliberate manipulation: it overturns resolutions even in near-certain cycles, transfers wealth from retail holders to a small cohort of manipulators, and cannot be explained solely by prediction-market makers' hedging.

6.1 Model Predictions

The first set of predictions concerns the footprint manipulation leaves in the market. If settlement manipulation exists, it must appear as a directional push in the final seconds before the close: a burst of one-sided order flow on the spot venue the oracle samples, large enough to move the spot price (Proposition 4). Because that flow carries no information about fundamentals, its price impact should be transitory: the move should revert once settlement has passed, because the settlement-window price is less informative (Proposition 3), in contrast to the permanent impact of informed trading.

The same lack of information content shapes liquidity on the spot venue. The manipulator's uninformed order pools with informed flow and lowers the adverse selection any individual trade carries, so spot market makers supply *more* liquidity into the close even as activity surges (Corollary 1). Finally, the entire footprint should be governed by whether a push can change anything: it concentrates in cycles the prediction market has left undecided near the close, and in the shortest-horizon contract, where the asset's own price discovery has had the least time to unfold and a fixed-size push is most likely to be pivotal (Proposition 7 and Proposition 8).

The second prediction concerns who pays. A successful push can overturn the resolution the market had priced, even in near-certain cycles we later identify in the data (Proposition 4). The contract is then mispriced relative to the manipulated resolution, and the loss falls on whoever holds the losing side at the close: liquidity traders, as predicted in the model. Manipulator profits are, to first order, liquidity traders' losses, according to the model's accounting identity (Proposition 5).

6.2 Data

The analysis combines three data sources: on-chain Polymarket trade records for the crypto up/down contracts, off-chain Polymarket order-book snapshots, and trades and quotes for BTC spot on Binance. Binance is the reference spot market throughout: it carries the largest BTC spot trading volume and arguably has the largest impact on the Chainlink reference price against which the Polymarket contracts resolve, so it is the natural venue on which a settlement manipulator would trade. Appendix D validates this choice directly, showing that the Binance midquote tracks the Chainlink resolution price to within a few basis points.

- *Polymarket on-chain trades.* Every order fill executed on Polymarket's exchange contracts settles on the Polygon blockchain, from which we extract the crypto up/down markets: over 60 million fills on the BTC five-minute contract alone. Each fill carries the block timestamp, the wallet identifiers on both sides, the asset on each side (one side is USDC, the other an Up- or Down-token), and the amounts deposited and received, which together recover the trade's price (in dollars per share, between 0 and 1), size, direction, and aggressor.[15] We link each fill to its market's cycle, strike, and realized resolution. Volume is measured two ways: *share volume*, the number of shares changing hands (each share pays $1 if its side wins, so the share count is the natural unit of exposure), and *dollar volume*, the cash transferred. We sign each fill by the aggressor's direction, expressed in the Up token (an aggressive Up buy or Down sell is bullish-Up), so net order flow over a window measures directional pressure on the contract.

- *Polymarket order book.* Polymarket's market-data feed delivers both full limit-order-book snapshots (the prevailing best bid and ask alongside several dozen price levels of resting depth on each side) and single-level updates whenever an existing level's resting size changes, for

[15] The on-chain log produces multiple records for each economic Polymarket trade: one aggregate record summarizing the aggressor's fill, plus one record for each resting order the aggressor consumes. We keep the single aggregate record per trade, so volume is never double-counted; on that record, the aggressor's deposited and received assets jointly identify both the contract side traded and the dollars exchanged. For wallet-level position and PnL tracking (Section 6.4.4), we instead sum the resting-side counterparty across both record classes, so each wallet's net position is recovered exactly once.

both tokens of each market. The feed covers roughly the final seven weeks of the sample. From it we measure the Up-token midquote, the bid–ask spread, and resting depth near the mid.

- *Binance BTC spot trades and order book.* The full tick-by-tick trade records and order-book snapshots for the reference asset. Each trade carries price, size, and aggressor side, from which we build signed spot order flow; each book snapshot carries the five best bid and ask levels with their resting quantities, timestamped at 100-millisecond precision, from which we build time-weighted spreads and depth. Both are aggregated to ten-second buckets aligned with the five-minute contract clock.

Sample window and rollout periods. The sample covers July 1, 2025 through April 8, 2026, starting about three months before the launch of the oracle-resolved Polymarket up/down contracts we study and ending about two months after the five-minute contract launched. We divide the sample into three periods based on the contract rollout schedule: the pre-launch baseline (P1), July 1 through October 8, 2025 (100 days); the intermediate period with the fifteen-minute and four-hour contracts live (P2), October 9, 2025 through February 11, 2026 (126 days); and the post-launch period with the five-minute contract also live (P3), February 12 through April 8, 2026 (56 days).[16]

Summary statistics. Table 1 reports summary statistics for the per-cycle trading activity over the five-minute contract's live window.[17]

Two observations stand out. First, the prediction market is small relative to its underlying spot: the median cycle trades about 105 thousand $1-face shares on Polymarket against $3.15 million

[16]Four narrow windows within the post-launch period require care: in three (late February 2026) the oracle-resolved Polymarket crypto contracts stop accepting trades while the rest of the platform trades normally (consistent with brief oracle-integration interruptions), and in a fourth (March 2026) the platform halts entirely. Cycles in these windows carry no on-chain Polymarket trades, so they drop out of every analysis that requires Polymarket data; the Binance-side panels retain them. On the spot side, we exclude October 10, 2025 (a market-wide flash crash) from the intra-cycle Binance figures, because a transient hundred-dollar spread during that episode would dominate a time-weighted average.

[17]Each five-minute contract is listed for trading several hours before its live resolution window opens, but about 93% of five-minute Bitcoin trades fall inside the live window (the 300 seconds between open and close), with the remainder split roughly evenly between pre-window positioning and the brief post-close interval before settlement. All cycle-level aggregates restrict to the live window.

Table 1. Per-cycle summary statistics, BTC five-minute contract. This table reports per-cycle trading activity on the BTC five-minute contract and the underlying Binance spot market, one observation per cycle in the contract's live window (February 12 – April 8, 2026); *N* is the number of cycles and the remaining columns report the cross-cycle distribution. The *Polymarket* rows aggregate the contract's own trades over the 300-second live window: *Trades* is the per-cycle taker trade count, *Volume* the number of shares traded (each carrying $1 of face value), and *Distinct traders* the number of unique wallets with a fill. The *Binance* rows aggregate the reference asset's spot tape over the same cycles: *Trades* and *Volume* (dollar volume) are per-cycle sums, and *Absolute return* is the absolute open-to-close change in the log midquote over the five-minute window (basis points).

	N	Mean	SD	Min	P25	Median	P75	Max
Polymarket BTC 5-minute contract								
Trades	16,073	3,912	1,228	2	3,084	3,811	4,642	12,622
Volume (1000 shares)	16,073	108.9	36.4	0.0	84.3	104.9	129.8	323.9
Distinct traders	16,073	1,689	374	4	1,467	1,685	1,918	3,332
Binance BTC spot								
Trades	16,128	15,793	14,039	765	7,131	11,676	19,290	223,213
Volume ($M)	16,128	5.08	7.01	0.17	1.79	3.15	5.86	284.69
Absolute return (bps)	16,128	10.16	11.81	0.00	2.88	6.58	13.51	376.29

of Binance volume. At first glance this asymmetry cuts against a manipulation story: if the prize pool is small relative to the market one must move, why bother? Two considerations restore the incentive. The gap is a factor of thirty, not several orders of magnitude, so the pool at stake is economically meaningful. More importantly, the five-minute-volume comparison is not the relevant one: the manipulator does not need to match five minutes of spot trading, only to deliver the final push, a burst of one-sided trading in the last seconds that carries the close across the strike. Second, ordinary five-minute spot moves are modest: the median absolute open-to-close return is 6.6 bps, with the inter-quartile range from 2.9 to 13.5 bps. A near-settlement spot push of a few basis points can therefore be large enough to clear the strike in a near-the-money cycle.

6.3 Does manipulation leave a footprint?

Section 6.1 tells us what to look for: if manipulation exists, the close of each cycle should carry a footprint with four parts: a burst of directional *order flow* (the push itself), an outsized *absolute return* (the push moves the price), improved spot *depth* (the push is uninformed, so spot makers supply more liquidity, not less), and a post-settlement *reversal* (for the same reason, the price impact does not last). We look for this footprint in three steps. We first show it emerges with the

launch of the five-minute contract. A launch coincidence alone cannot establish manipulation, so we then turn to the identification: within the post-launch period, the footprint concentrates in the cycles where a push pays off, the cycles still undecided near the close. We close with the model's horizon prediction: the footprint should weaken as the contract horizon lengthens, and indeed the fifteen-minute contract, live earlier and settling on the same rule, exhibits only a strongly attenuated version of it.

6.3.1 The footprint emerges with the five-minute launch

Near-settlement activity spike. We divide each five-minute cycle into 30 ten-second bins, indexed 0 through 29, and compute within each bin three of the footprint's four parts: absolute order flow, absolute return, and depth within \$0.10 of the midquote. For each bin we then average across all cycles in a period and compare the pre-launch baseline (P1) against the five-minute-live period (P3); the intermediate period (P2), in which only the fifteen-minute and four-hour contracts trade, returns when we study the fifteen-minute contract on its own (Section 6.3.3).

Figure 8 plots the resulting intra-cycle profiles. The pattern is striking. In bin 29, the last ten seconds of the cycle, all three metrics spike once the five-minute contract is live (the +5m series); in the pre-launch baseline, bin 29 looks no different from the rest of the cycle (we call the first 25 bins the *body*). The jump is confined to the final bin, so the contrast at the close is not a broad shift in Binance trading.[18]

We formalize the visual pattern with a pre/post comparison of the bin-29 activity, comparing the pre-launch baseline (P1) against the period in which the five-minute contract is live (P3):

$$\log y_{c,b} \;=\; \gamma_b + \alpha_d + \mu_h + \delta_{P3}\,\mathbf{1}\{b{=}29\}{\cdot}\mathbf{1}\{P3\} + \varepsilon_{c,b},$$

where $\log y_{c,b}$ is the log of the metric in cycle c and bin b, restricted to the body bins and the final-10s bin, $b \in \{0, \ldots, 24, 29\}$, so the comparison is clean between the near-settlement activity and the body.

[18] There are minor spikes at the minute-ends. They are likely driven by trading algorithms that react to the minute mark.

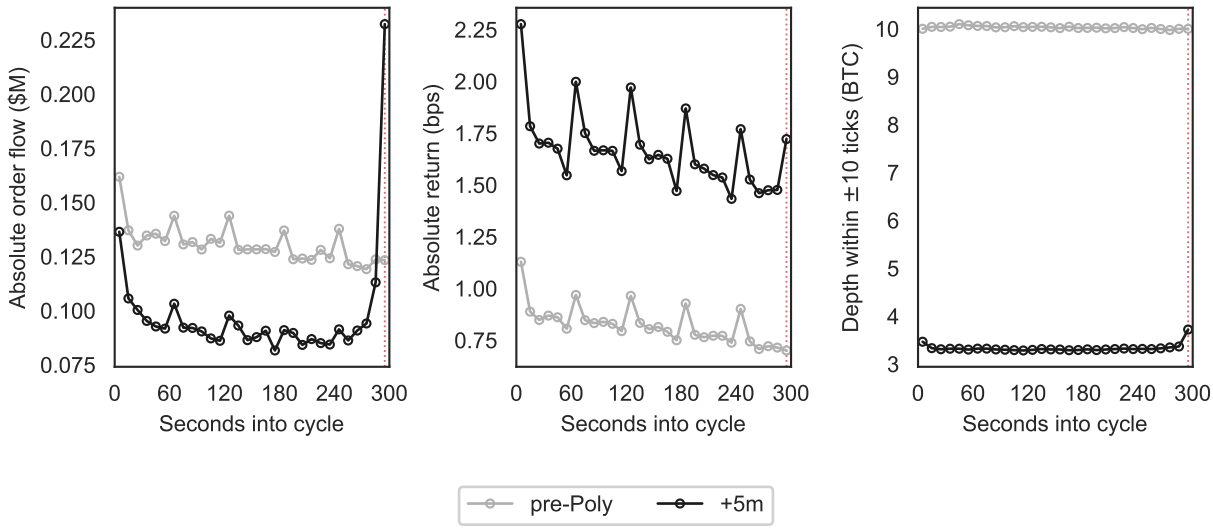

Figure 8. Binance BTC intra-cycle metrics per 5-min cycle, by period. This figure plots the intra-cycle profile of three Binance BTC metrics, by period. Each panel folds every 10s Binance BTC bucket onto its position in the rolling 5-minute clock cycle and overlays, for each period, the metric pooled over all 5-minute cycles in that period (see Section 6.3.1). *Left*: absolute order flow per cycle ($M), the per-cycle |net signed taker dollar flow| in the bin. *Center*: absolute 10s return (|Δ log mid|, bps). *Right*: time-weighted depth within ±10 ticks of the mid, cumulative resting size (bid + ask) at price levels within 10 ticks ($0.10) of the mid. The dotted vertical line marks the final 10s before close. Two periods are overlaid: pre-Polymarket (P1, Jul–Oct 2025) and the five-minute-live period (P3, Feb 12 – Apr 8, 2026); the intermediate period with only the fifteen-minute and four-hour contracts live (P2) is omitted.

α_d and μ_h are date and hour-of-day fixed effects; γ_b is a bin fixed effect that absorbs the pre-existing intra-cycle profile. The specification pools all cycles: we do not differentiate overlapping from non-overlapping closes here; the overlap split returns in the fifteen-minute analysis (Section 6.3.3) and, interacted with moneyness, in the within-P3 refinements of Appendix C. The coefficient δ_{P3} is the average bin-29-versus-body activity increase once the five-minute contract is live, relative to the pre-launch baseline.

Table 2 reports the estimates. Once the five-minute contract is live, the typical cycle's bin-29 absolute order flow is 50% higher relative to the body than before the launch ($\hat{\delta}_{P3} = 0.406$ log points), its absolute return is 15% higher (0.140), and its near-mid depth 5% higher (0.048), all significant at the 1% level. Three parts of the footprint thus appear at the close exactly when the five-minute contract begins to settle there.

Table 2. Near-settlement activity: pre-Polymarket vs five-minute period. This table reports cycle-level OLS estimates on the (cycle, bin) panel restricted to bins {0–24, 29} and to the pre-Polymarket period (P1, Jul–Oct 2025) and the five-minute-live period (P3, Feb 12 – Apr 8, 2026); the intermediate period in which only the fifteen-minute and four-hour contracts are live (P2) is excluded. The dependent variable is $\log y_{c,b}$, the log of the metric named in the column header: *Order flow*, per-cycle absolute net signed taker dollar flow ($|\sum_i d_i\, p_i a_i|$, $d = +1$ buy / -1 sell); *Abs. return*, absolute 10 s log return ($|\Delta \log \text{mid}|$, bps); *Depth*, time-weighted resting size within ±10 ticks of the mid. δ_{P3} is the coefficient on $\mathbf{1}\{b{=}29\} \cdot \mathbf{1}\{P3\}$: the bin-29-versus-body activity increase once the five-minute contract is live, relative to the pre-Polymarket baseline. Fixed effects are date, hour-of-day, and bin; the specification pools all cycles and does not differentiate overlap from non-overlap closes. Coefficients are log differences. One observation per (cycle, bin) cell. HC1 standard errors in parentheses are clustered at the cycle level. Significance: $^{*}\ p < 0.10$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$.

	(1) Order flow	(2) Abs. return	(3) Depth
δ_{P3}	0.406***	0.140***	0.048***
	(0.019)	(0.018)	(0.003)
Date FE	Yes	Yes	Yes
Hour-of-day FE	Yes	Yes	Yes
Bin FE	Yes	Yes	Yes
Observations	1,168,118	738,313	1,168,114
Clusters (cycle)	44,928	44,696	44,928

Post-settlement reversal. The fourth part of the footprint is the reversal: manipulation flow is uninformative, so its price impact should not last. We test this by regressing each lagged ten-second return around the settlement on the final-10s return r_{29} at the cycle level.

Separately at each lag j, we estimate

$$r_{j,c} \;=\; \alpha_d + \mu_h + \beta\, r_{29,c} + \delta_{P3}\, r_{29,c}\, \mathbf{1}\{P3\} + \varepsilon_{j,c},$$

on the pre-Polymarket (P1) and five-minute-live (P3) cycles. The dependent variable $r_{j,c}$ is the lag-j ten-second return of cycle c measured relative to its close: positive lags map to bins 0–5 of the next cycle (post-settlement), negative lags to bins 25–28 of the same cycle (the run-up into the close), and $j = 0$ is the bin-29 reference itself, omitted because its slope is 1 trivially. The regressor $r_{29,c}$ is the final-10s (bin-29) return. β is the pre-Polymarket reversal slope and δ_{P3} the change in that slope once the five-minute contract is live; α_d and μ_h are date and hour-of-day fixed effects, and standard errors are clustered by date. A reverting price is a negative slope; the key lag is $j = +1$, the first ten seconds after the close, where a transitory manipulation push predicts reversal.

Figure 9 overlays the reversal slope γ_j for the two periods. A sharp reversal appears in the first

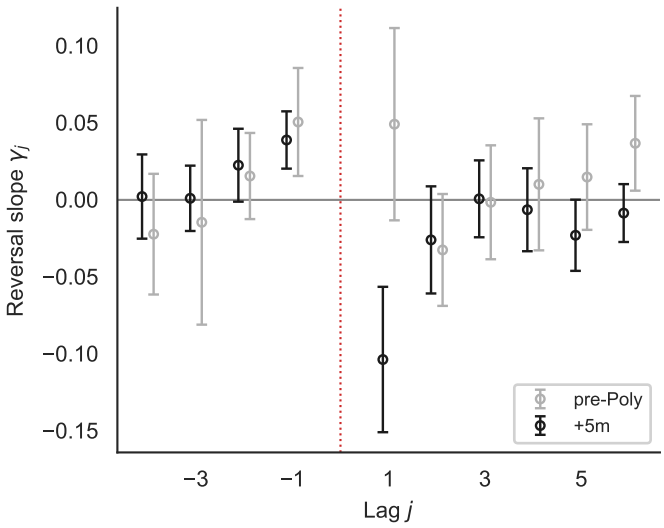

Figure 9. Post-settlement reversal slope γ_j: pre-Polymarket vs five-minute period. This figure plots, at each lag j, the slope γ_j from a per-period pooled OLS of the lag-j 10s return on the final-10s return r_{29} (date and hour-of-day fixed effects, standard errors clustered by date), overlaid for the pre-Polymarket period (P1) and the five-minute-live period (P3); the intermediate +15m/4h period (P2) is omitted. Markers are γ_j; error bars are ±1.96× cluster-robust SE; $j = 0$ is omitted. The dotted vertical marks the close. Negative j is bin $29 + j$ of the same cycle; positive j is bin $j - 1$ of the next cycle.

ten seconds after the close, and only with the five-minute contract: γ_{+1} is strongly negative once the contract is live and indistinguishable from zero in the pre-launch baseline, while at $j \geq +2$ the slopes are flat in both periods.

Table 3 confirms it. Against a pre-launch baseline with no reversal at the close ($\beta = 0.05$ at $j = +1$, indistinguishable from a random walk), the five-minute contract shifts the reversal slope by $\delta_{P3} = -0.15$, so the post-close slope under the five-minute contract is $\beta + \delta_{P3} = -0.10$ at $j = +1$: about a tenth of the final-10s spot move reverses within the next ten seconds. The shift is specific to the first post-close lag: δ_{P3} is flat at $j \geq +2$. Taken together, all four parts of the footprint appear with the five-minute contract, and only then. A before/after comparison, however, cannot on its own distinguish manipulation from any other change in close-time trading that coincides with the launch. The identification comes from the cross-section within the post-launch period, to which we now turn.

Table 3. Post-settlement reversal: pre-Polymarket vs five-minute period. This table reports the cycle-level OLS $r_j = \alpha_d + \mu_h + \beta\, r_{29} + \delta_{P3}\, r_{29}\mathbf{1}\{P3\} + \varepsilon$, estimated separately at each lag $j \in \{-3, -2, -1, +1, +2, +3\}$, on the pre-Polymarket period (P1, Jul 1 – Oct 8, 2025) and the five-minute-live period (P3, Feb 12 – Apr 8, 2026); the intermediate period in which only the fifteen-minute and four-hour contracts are live (P2) is excluded. β is the pre-Polymarket reversal slope on r_{29} and δ_{P3} its change once the five-minute contract is live (a negative slope is reversal). Date and hour-of-day fixed effects; HC1 standard errors in parentheses are clustered by date. Negative j = same-cycle bin $(29 + j)$; positive j = bin $(j - 1)$ of the next cycle. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Pre-close			Post-close		
	$j = -3$	$j = -2$	$j = -1$	$j = +1$	$j = +2$	$j = +3$
β	−0.014	0.015	0.051***	0.049	−0.033*	−0.002
	(0.034)	(0.014)	(0.018)	(0.032)	(0.019)	(0.019)
δ_{P3}	0.015	0.007	−0.012	−0.152***	0.007	0.002
	(0.036)	(0.019)	(0.020)	(0.040)	(0.026)	(0.023)
Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Hour-of-day FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	44,772					
Clusters (UTC date)	156					

6.3.2 The footprint concentrates in NTM cycles

The before/after evidence so far shows a four-part footprint that coincides with the five-minute launch, but a launch coincidence cannot by itself separate manipulation from any change in close-time trading. The identification comes from a cross-sectional cut *within* the five-minute period. As the model predicts, if the push is manipulation, all four parts of the footprint should be largest in cycles where the outcome is still undecided near the close — there a small spot push can flip the resolution, so the manipulator's payoff is greatest — and muted in cycles already decided, where a push changes nothing. Because this contrast is taken within a single period, it differences out any time trend common to the launch.

Near-the-money classification. We classify each post-launch five-minute cycle by its prediction-market price near settlement. Specifically, for each cycle, we take the Up-token price 10 seconds before the settlement and label the cycle *near-the-money* (NTM) if the Up-token price is between 0.40 and 0.60 (the contract is still a coin-flip, so a push across the strike can flip the outcome) and *far-from-the-money* (FFM) otherwise. About 6% of cycles are NTM.[19]

Near-settlement spike in NTM cycles. Figure 10 plots three intra-cycle metrics — absolute order flow, absolute return, and depth — on both venues, separately for NTM and FFM cycles. The figure carries two messages. First, the Binance spike concentrates in NTM cycles: order flow and absolute return rise at the close in both groups, but markedly more where the contract is still undecided and a push can flip the resolution. Second, liquidity diverges across the venues into the NTM close: Binance makers *add* near-mid depth while Polymarket makers *withdraw* it. Two secondary observations complete the picture. The binary's own price moves more in NTM cycles for a mechanical reason: a price near 0.5 must travel to 0 or 1 as the contract resolves, while a decided contract's price is already there. Its order flow, by contrast, spikes *less*: close-time flow on

[19]The ~4% of P3 cycles with no qualifying trade in the 10–12 s window are dropped as unclassified; these are systematically the sparse-trading cycles. The tight window trades coverage for freshness: a wider window would retain them but allow the classifying price to be stale by up to several minutes, mislabeling a cycle when the spot moves in the interim. Pre-launch cycles (P1, P2) have no on-chain 5m trades and are unclassified by construction.

the binary is dominated by the one-sided sweep of an already-decided contract, largest exactly when the outcome is no longer in play, and in contested cycles the directional flow appears on Binance instead.

To quantify the NTM-vs-FFM contrast on both venues, we estimate the following regression:

$$\log y_{c,b} \;=\; \alpha_d + \gamma_b + \mu_h + \beta_{\text{NTM}}\,\mathbf{1}\{\text{NTM}_c\} + \delta_{\text{NTM}}\,\mathbf{1}\{b{=}29\}{\cdot}\mathbf{1}\{\text{NTM}_c\} + \varepsilon_{c,b},$$

where α_d, μ_h, γ_b capture the date, hour-of-day, and bin fixed effects. The standard errors are clustered at the cycle level. The FFM bin-29 increase is absorbed into the bin fixed effect γ_{29}, so δ_{NTM} measures the additional bin-29-vs-body activity increase in NTM cycles over the FFM baseline.

Table 4 reports the regression results. There are two key observations. First, on Binance, the bin-29 order flow is $\hat{\delta}_{\text{NTM}} = 1.37$ log points larger in NTM than in FFM cycles (about $e^{1.37} \approx 3.9$ times the FFM one on the typical cycle), and the same NTM concentration holds for absolute return (0.56), both positive at the 1% level. Polymarket diverges on flow: its NTM order-flow effect is *negative* (−0.18), the binary's close spike being smaller in coin-flip cycles than in decided ones, where the one-sided sweep of the certain winner dominates. Its price-change effect is large and positive (2.62), but largely mechanical: resolution pulls a near-0.5 price to 0 or 1, while a decided contract's price is already there. Where the outcome is in play, the incremental directional trade concentrates on the spot venue, not on the binary itself.

Second, Binance and Polymarket see opposite patterns in liquidity: into the NTM close Binance liquidity improves ($\hat{\delta}_{\text{NTM}} = 0.17$) while Polymarket liquidity worsens ($\hat{\delta}_{\text{NTM}} = -1.02$). The liquidity divergence is exactly the model's prediction (Section 6.1). A maker quoting the Polymarket binary near the close becomes the counterparty to anyone about to push the spot across the strike: the manipulator buys the side he is about to make win, so the prediction-market maker who sold it to him is left on the losing side of the resolution. Polymarket makers therefore pull their book as the close nears, and pull it hardest in NTM cycles, where a flipping push is most likely. On Binance the

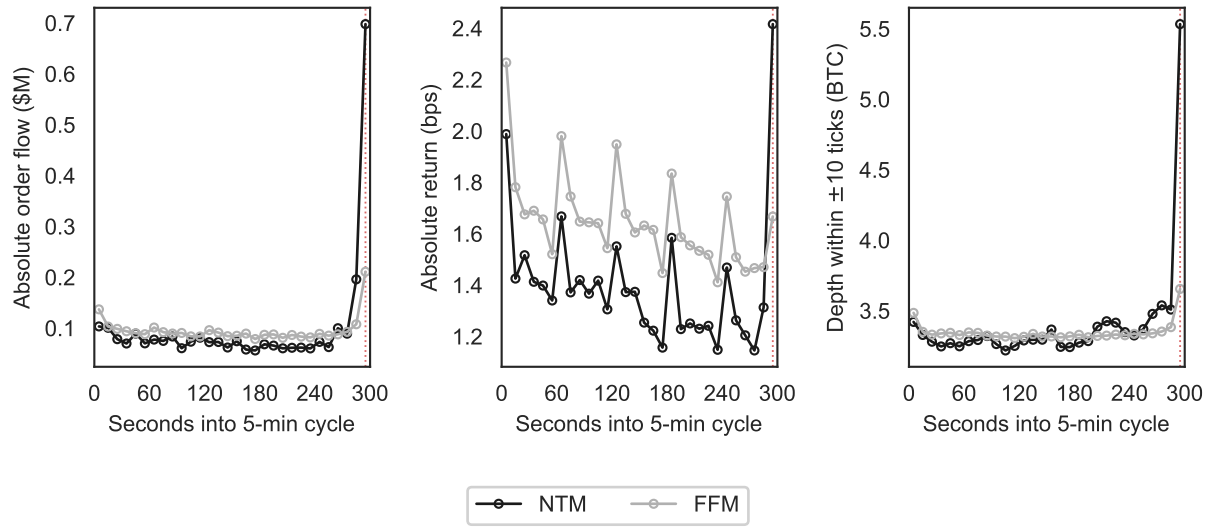

(a) Binance BTC spot market.

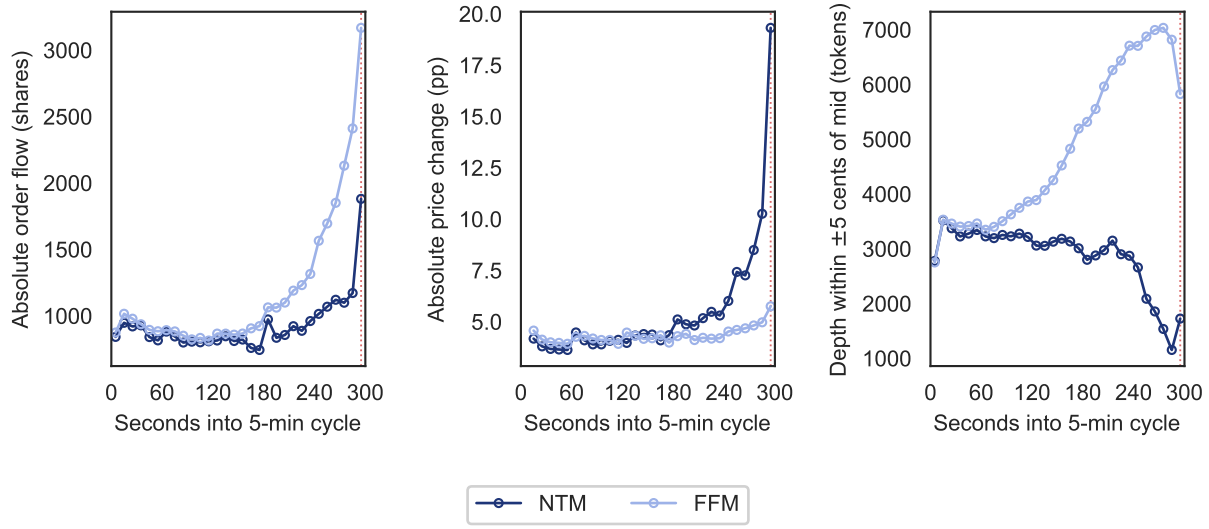

(b) Polymarket BTC 5-minute up/down contract.

Figure 10. Intra-cycle metrics per 5-min cycle, NTM vs. FFM (P3 only). This figure plots intra-cycle profiles on both venues, separately for NTM and FFM cycles. Each P3 5-minute cycle is classified NTM ($|p_{\mathrm{Up}} - 0.5| < 0.10$) or FFM from the latest taker-focused on-chain Polymarket Up-token trade at time-to-maturity TTM $\in [10, 12]$ s; each metric is pooled over all classified P3 5-minute cycles. Each 1×3 grid plots one metric's intra-cycle profile against seconds into the 5-min cycle; the dotted vertical marks the final 10s before close. **(a) Binance BTC spot**: absolute order flow per cycle ($M, |net signed taker dollar flow|), absolute 10s return ($|\Delta \log \mathrm{mid}|$, bps), time-weighted depth within ±10 ticks of the mid. **(b) Polymarket BTC 5-minute up/down contract**: absolute order flow per cycle (|net signed Up-token share flow|), absolute 10s probability change ($|\Delta p|$, pp), time-weighted depth within ±5 cents of the mid (cumulative Up-token resting size at price levels within $0.05 of the mid).

same push is uninformed and therefore generates only a *transitory* price impact. So the push flow is benign, and Binance makers supply *more* depth into the NTM close. Appendix C.1 refines the

Table 4. P3 near-settlement activity, NTM vs. FFM cycles, Binance and Polymarket. This table reports cycle-level OLS estimates on the (cycle, bin) panel restricted to bins {0–24, 29} and P3 cycles, $\log y_{c,b} = \alpha_d + \gamma_b + \mu_h + \beta_{\text{NTM}} \mathbf{1}\{\text{NTM}_c\} + \delta_{\text{NTM}} \mathbf{1}\{b{=}29\}\mathbf{1}\{\text{NTM}_c\} + \varepsilon_{c,b}$, with date, hour-of-day, and bin fixed effects. The first three columns report Binance spot, the last three the Polymarket 5-minute contract; within each venue the dependent variable $\log y_{c,b}$ is the log of the column metric. The activity column is *Order flow* on both venues — the per-cycle absolute net signed flow, in taker dollars on Binance ($|\sum_i d_i\, p_i a_i|$) and in Up-token shares on Polymarket. The second column per venue is *Abs. return* on Binance (absolute 10 s log return, $|\Delta \log \text{mid}|$, bps) and *Abs. price change* on Polymarket (absolute 10 s probability change, $|\Delta p|$, pp); the third is *Depth* (time-weighted resting size within a mid-anchored band: ±10 ticks of the mid on Binance, ±5 cents on Polymarket). A P3 cycle is NTM iff the 5m-contract Polymarket Up-token price is within 0.10 of 0.5 at TTM $\in [10, 12]$ s, FFM otherwise. δ_{NTM} is the bin-29-vs-body activity increase in NTM cycles over the FFM baseline (the latter absorbed in the bin FE); coefficients are log differences, so δ is a multiplicative effect of e^{δ} on the typical (geometric-mean) cycle. Appendix C.1 reports the same regression with the NTM cut refined into the full 15-min-overlap × per-contract-NTM split. One observation per (cycle, bin) cell. Polymarket order-book metrics (absolute return, depth) have smaller samples than order flow because the order-book archive begins later than the trade archive and outage days are excluded. HC1 standard errors in parentheses are clustered at the cycle level. Significance: $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$.

	Binance			Polymarket		
	Order flow	Abs. return	Depth	Order flow	Abs. price change	Depth
δ_{NTM}	1.366^{***}	0.555^{***}	0.173^{***}	-0.182^{***}	2.618^{***}	-1.016^{***}
	(0.086)	(0.052)	(0.022)	(0.044)	(0.089)	(0.052)
Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Hour-of-day FE	Yes	Yes	Yes	Yes	Yes	Yes
Bin FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	395,790	307,108	395,790	395,154	138,405	142,468
Clusters (cycle)	15,223	15,221	15,223	15,223	5,631	5,631

NTM cut by 15-min contract overlap.

Post-settlement price reversal in NTM cycles. In addition to the activity spike, NTM cycles see a more pronounced post-settlement price reversal. Using the same moneyness classifier, at each lag we estimate the following regression:

$$r_{j,c} \;=\; \alpha_d + \mu_h + \nu_g + \gamma_j\, r_{29,c} + \delta_{j,\text{NTM}}\, \mathbf{1}\{\text{NTM}_c\}\, r_{29,c} + \varepsilon_{j,c},$$

where γ_j is the FFM baseline reversal slope on r_{29} and $\delta_{j,\text{NTM}}$ is the NTM increment over it, so the NTM-cycle slope is $\gamma_j + \delta_{j,\text{NTM}}$ and $\delta_{j,\text{NTM}}$ is directly the NTM-minus-FFM reversal gap. α_d, μ_h, and ν_g are date, hour-of-day, and moneyness-group fixed effects; standard errors are clustered by date.

We report the regression results in Table 5 and plot the slopes in Figure 11. FFM cycles revert by about 10% ($\gamma_{+1} = -0.10$), and NTM cycles revert significantly harder: the NTM increment is

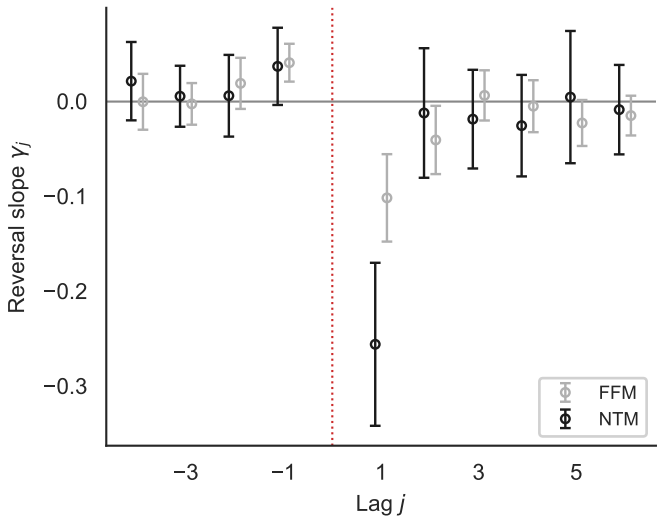

Figure 11. Within-P3 post-settlement reversal slope γ_j, by 5m-contract moneyness. This figure plots γ_j for each 5m-contract moneyness group (the FFM baseline and the NTM total $\gamma_j + \delta_{j,\mathrm{NTM}}$ of Table 5), on all P3 cycles split by 5m-contract Polymarket NTM (Up-token price within 0.10 of 0.5 near TTM = 10 s). Markers are γ_j; error bars are ±1.96× cluster-robust SE (clustered by date); $j = 0$ is omitted.

$\delta_{+1,\mathrm{NTM}} = -0.15$, so NTM cycles revert by about 25% in total. A transient push pays off only by moving a price the prediction market has not already decided, so the cycles that revert hardest are exactly the NTM ones, where the push can flip the resolution. As before, in Appendix C.1 we further partition the cycles by 15-minute overlap. We find that the reversal is deeper in cycles where the 5m and 15m Polymarket contracts settle at the same close.

Table 5. Within-P3 post-settlement reversal slope by 5m-contract moneyness. This table reports, at each lag $j \in \{-3, -2, -1, +1, +2, +3\}$, the two coefficients from a single cycle-level OLS on all P3 cycles (Feb 12 – Apr 8, 2026), $r_j = \alpha_d + \mu_h + \nu_g + \gamma_j r_{29} + \delta_{j,\mathrm{NTM}} \mathbf{1}\{\mathrm{NTM}\} r_{29} + \varepsilon$: γ_j is the FFM baseline reversal slope and $\delta_{j,\mathrm{NTM}}$ the NTM increment over it, so the NTM-cycle slope is $\gamma_j + \delta_{j,\mathrm{NTM}}$ and $\delta_{j,\mathrm{NTM}}$ is the NTM-minus-FFM gap; the stars test each coefficient against zero. A cycle is $\mathrm{NTM}_{5\mathrm{m}}$ iff the 5m-contract Polymarket Up-token price is within 0.10 of 0.5 near TTM $\in [10, 12]$ s, FFM otherwise. Date, hour-of-day, and moneyness-group fixed effects; standard errors in parentheses are clustered by date. Negative j = same-cycle bin $(29 + j)$; positive j = bin $(j - 1)$ of the next cycle. Appendix C.1 dissects these groups further by 15-minute overlap. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Pre-close			Post-close		
	$j = -3$	$j = -2$	$j = -1$	$j = +1$	$j = +2$	$j = +3$
$\mathrm{FFM}_{5\mathrm{m}}$	−0.002	0.019	0.041***	−0.101***	−0.040**	0.007
	(0.011)	(0.014)	(0.010)	(0.023)	(0.018)	(0.013)
$\delta_{\mathrm{NTM}_{5\mathrm{m}}}$	0.008	−0.013	−0.004	−0.154***	0.028	−0.025
	(0.017)	(0.024)	(0.021)	(0.038)	(0.037)	(0.026)
Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Hour-of-day FE	Yes	Yes	Yes	Yes	Yes	Yes
Cycle-group FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	15,168					
Clusters (UTC date)	56					

6.3.3 The manipulation footprint fades at the fifteen-minute horizon

The model makes one more prediction: the manipulation footprint should weaken as the settlement horizon lengthens, because a longer window gives the asset's own price discovery more time to unfold, so a fixed-size push is less likely to be pivotal and manipulation opportunities become rarer ex ante (Section 6.1). The fifteen-minute contract provides the test. It went live earlier (in period P2), resolves based on the same Chainlink oracle price, and, during P2, was the most heavily traded of the crypto up/down contracts (Appendix B), so a weak footprint cannot be attributed to a small prize.

Identification. Within the five-minute clock, the fifteen-minute (and four-hour) contract settles only at the closes that coincide with a 15-minute mark: the *overlapping* cycles, one in three; the other closes carry no fifteen-minute settlement. The fifteen-minute effect is therefore identified by contrasting overlapping with non-overlapping closes. One complication forces a richer design: 15-minute marks might be intrinsically special on the spot market (round-time execution algorithms cluster there), so a raw overlap contrast would confound the fifteen-minute settlement with this clock effect. We difference it against the same contrast in the pre-Polymarket baseline: on the (cycle, bin) panel over P1 and P2, restricted to bins {0–24, 29} as before, we estimate

$$\log y_{c,b} \;=\; \theta_{t_c,\,b} + \alpha_d + \mu_h + \delta_{P2}\,\mathbf{1}\{b{=}29\}{\cdot}\mathbf{1}\{P2\} + \delta_{P2,\,5\mathrm{m}\cap 15\mathrm{m}}\,\mathbf{1}\{b{=}29\}{\cdot}\mathbf{1}\{P2\}{\cdot}D_{5\mathrm{m}\cap 15\mathrm{m},c} + \varepsilon_{c,b},$$

where $D_{5\mathrm{m}\cap 15\mathrm{m},c}$ indicates an overlapping cycle. The specification modifies the activity regression of Section 6.3.1 in two ways. First, the bin fixed effect is replaced by a cycle-type×bin fixed effect $\theta_{t_c,\,b}$, which gives overlapping and non-overlapping cycles each their own intra-cycle profile and thereby absorbs the clock effect, whatever its shape. Second, the bin-29 shift is split by overlap status: δ_{P2} is the bin-29-versus-body increase in P2 at non-overlapping closes, a pure period effect, and $\delta_{P2,\,5\mathrm{m}\cap 15\mathrm{m}}$ is the *additional* increase at overlapping closes, the difference-in-differences (P2 − P1) × (overlap − non-overlap) that isolates the fifteen-minute settlement effect, netting out

both the period and the clock. The reversal regression of Section 6.3.1 is augmented in the same way. At each lag j, estimated separately on P1 and P2 cycles,

$$\begin{aligned} r_{j,c} \;=\; & \alpha_d + \mu_h + \beta\, r_{29,c} + \beta_{5\text{m}\cap 15\text{m}}\, r_{29,c}\, D_{5\text{m}\cap 15\text{m},c} \\ & + \delta_{P2}\, r_{29,c}\, \mathbf{1}\{P2\} + \delta_{P2\times(5\text{m}\cap 15\text{m})}\, r_{29,c}\, \mathbf{1}\{P2\}\cdot D_{5\text{m}\cap 15\text{m},c} + \varepsilon_{j,c}, \end{aligned}$$

with the lower-order level terms $D_{5\text{m}\cap 15\text{m},c}$ and $\mathbf{1}\{P2\}\cdot D_{5\text{m}\cap 15\text{m},c}$ included; the $\mathbf{1}\{P2\}$ level is absorbed by the date fixed effects. The four slopes mirror the activity design: β is the pre-Polymarket reversal slope at non-overlapping closes and $\beta_{5\text{m}\cap 15\text{m}}$ its overlap increment, the clock effect on the slope; δ_{P2} is the P2 slope change at non-overlapping closes, the period effect; and $\delta_{P2\times(5\text{m}\cap 15\text{m})}$ is the additional P2 change at overlapping closes, the difference-in-differences that gives the fifteen-minute settlement effect on the reversal slope.

Activity and reversal. Table 6 reports the activity difference-in-differences over P1 and P2, and Table 7 the reversal one. The fifteen-minute settlement increment $\delta_{P2,\,5\text{m}\cap 15\text{m}}$ is small on every dimension: an order-flow increase of 0.068 against the five-minute 0.406, a depth increase of 0.021 against 0.048, no detectable absolute-return effect, and no post-settlement reversal at $j = +1$. Smaller but not exactly zero is what the horizon mechanism predicts: a longer window shrinks the incentive rather than eliminating it.

Conditioning on moneyness sharpens the picture but does not change it: even the fifteen-minute NTM cycles show at most a modest activity increase and no reversal, far below the five-minute footprint (Appendix C.2). The horizon prediction is borne out: the footprint that arrives in full force with the five-minute contract, the shortest, most frequent, and most cheaply moved settlement, appears at the fifteen-minute horizon only in strongly attenuated form.

Table 6. Fifteen-minute settlement: near-settlement activity, pre-Polymarket vs P2 by overlap. This table reports cycle-level OLS estimates on the (cycle, bin) panel restricted to bins {0–24, 29} and to the pre-Polymarket period (P1, Jul–Oct 2025) and the +15m/4h period (P2, Oct 2025 – Feb 2026). δ_{P2} is the bin-29-versus-body activity increase in P2 non-overlapping cycles relative to P1; $\delta_{P2,\,5\text{m}\cap15\text{m}}$ is the *additional* increase at overlapping closes (those landing on a 15-min mark, where the fifteen-minute contract settles) — the difference-in-differences (P2 − P1) × (overlap − non-overlap) isolating the fifteen-minute settlement effect. Metrics and standard errors as in Table 2; fixed effects are date, hour-of-day, and cycle-type×bin. Coefficients are log differences. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1) Order flow	(2) Abs. return	(3) Depth
δ_{P2}	-0.062^{***}	-0.091^{***}	0.000
	(0.017)	(0.019)	(0.003)
$\delta_{P2,\,5\text{m}\cap15\text{m}}$	0.068^{**}	0.042	0.021^{***}
	(0.032)	(0.035)	(0.005)
Date FE	Yes	Yes	Yes
Hour-of-day FE	Yes	Yes	Yes
Cycle-type × Bin FE	Yes	Yes	Yes
Observations	1,684,789	1,071,666	1,684,782
Clusters (cycle)	64,800	64,391	64,800

Table 7. Fifteen-minute settlement: post-settlement reversal, pre-Polymarket vs P2 by overlap. This table reports, at each lag $j \in \{-3, \ldots, +3\}$, a cycle-level OLS over P1 and P2 cycles. β is the pre-Polymarket reversal slope on r_{29} at non-overlapping closes, $\beta_{5\text{m}\cap15\text{m}}$ its pre-Polymarket overlap increment (the clock effect at 15-min marks), δ_{P2} the P2 slope shift at non-overlapping closes, and $\delta_{P2\times(5\text{m}\cap15\text{m})}$ the P2 overlap increment — the fifteen-minute settlement effect, a difference-in-differences. A negative slope is reversal. Date and hour-of-day fixed effects; HC1 standard errors in parentheses clustered by date. Negative j = same-cycle bin $(29 + j)$; positive j = bin $(j - 1)$ of the next cycle. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Pre-close			Post-close		
	$j = -3$	$j = -2$	$j = -1$	$j = +1$	$j = +2$	$j = +3$
β	−0.034	0.000	0.036	0.073^{***}	0.000	0.010
	(0.047)	(0.018)	(0.022)	(0.028)	(0.019)	(0.018)
$\beta_{5\text{m}\cap15\text{m}}$	0.066	0.053^{*}	0.050	−0.081	-0.110^{**}	−0.039
	(0.047)	(0.030)	(0.034)	(0.077)	(0.045)	(0.047)
δ_{P2}	0.033	−0.017	0.015	-0.116^{***}	0.014	0.011
	(0.052)	(0.033)	(0.030)	(0.036)	(0.024)	(0.025)
$\delta_{P2\times(5\text{m}\cap15\text{m})}$	0.002	−0.058	-0.081^{*}	0.024	0.066	−0.005
	(0.085)	(0.046)	(0.043)	(0.087)	(0.060)	(0.059)
Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Hour-of-day FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	64,575					
Clusters (UTC date)	225					

6.4 Who wins and loses in manipulated cycles?

Section 6.3 shows that the five-minute launch coincides with a transitory final-10s spot push concentrated in the cycles the prediction market has left undecided: the manipulation footprint. This section tests the model's second prediction (Section 6.1): that the push transfers wealth, from the liquidity traders holding the losing side to the manipulator. We first single out the cycles carrying the largest manipulation pressure and document where they sit and whether the push succeeds; we then sort the traders active in them into types and measure who profits and who pays. Last, we rule out the alternative that the push is prediction-market makers hedging their binary exposure. Throughout this section, the sample is the five-minute-live period (P3), the same window as the within-period identification above.

6.4.1 Identification of manipulated cycles

We identify which cycles carry the largest manipulation pressure using a within-cycle, outcome-independent statistic:

$$PushIntensity_c \;=\; \frac{|\,\text{order flow}_{29,c}\,|}{\text{median}_{b\in\{0,\ldots,24\}}\,|\,\text{order flow}_{b,c}\,|},$$

where $\text{order flow}_{b,c}$ is the Binance taker buy minus sell dollar volume in bin b of cycle c. The numerator is the magnitude of the final-10s order flow; the denominator is the typical magnitude of order flow per body bin in the *same* cycle. This normalization has three useful properties: (i) it is dimensionless and bounded below by zero; (ii) it is invariant to the overall scale of activity in the cycle, so a quiet and an active cycle with the same push-vs-body asymmetry score equally; and (iii) it is constructed entirely from the cycle's own intra-cycle Binance order flow, with no input from the resolution price, the Polymarket Up-token price, or the strike, so a trader's PnL is not mechanically related to its score.

We label a cycle a *manipulated cycle* if its *PushIntensity* is at or above the 90th percentile of the cycle distribution, $PushIntensity \geq 16.11$. Of the 16,073 cycles with on-chain Polymarket trading,

this gives 1,613 manipulated cycles and 14,460 *normal cycles*.[20] The top-decile cut is a convention, chosen large enough that the manipulated-cycle sample carries meaningful statistical power (1,613 cycles × on average ∼1,500 traders per cycle) and small enough that the cycles selected are visibly distinct from the body of the distribution (the 90th-percentile *PushIntensity* of 16 is an order of magnitude larger than the median *PushIntensity* of ∼ 0.9, the latter close to what would be expected under no asymmetry between the final 10s and the body).

6.4.2 Characteristics of manipulated cycles

In Table 8, we contrast manipulated cycles with the normal ones along four dimensions: Binance activity, Polymarket activity, timing, and moneyness.

Because manipulated cycles are selected on *PushIntensity*, a large Binance spike among them is in part mechanical; the informative contrast is across venues. On Binance the bin-29 order flow averages $1.7M against a typical body bin of just $18k; on Polymarket the bin-29 order flow ($2.8k) sits much closer to its own body bin ($0.4k) and barely exceeds its normal-cycle level. The near-settlement push moves Binance spot without any comparable surge in Polymarket's own order flow, the first sign that the push is executed on the spot venue, matching the manipulation footprint.

Manipulated cycles are also *quieter* away from the close, not busier. The typical body-bin order flow is lower than in normal cycles on both venues, and the cycles fall disproportionately in thin-liquidity windows: 56% land in the Asia/overnight hours (vs 40% of normal cycles) and 44% on weekends (vs 27%), against only 23% in the deep European–US overlap (vs 35%). Yet despite this quieter backdrop the absolute Binance near-settlement order flow is an order of magnitude larger than in a normal cycle ($1.7M vs $68k). Manipulation thus concentrates precisely where the book is thinnest, the overnight and weekend hours, where a given dollar of near-settlement order flow moves the spot, and the Chainlink resolution it feeds, the most.

If the near-settlement push is manipulation, it should concentrate where the outcome is still in

[20] A further 55 cycles in the window, Polymarket outage periods and ultra-thin overnight cycles, carry a *PushIntensity* score from Binance order flow but have no Polymarket trades, and drop out of the trader-level analysis.

Table 8. Summary statistics for manipulated cycles vs normal cycles, P3 BTC 5m. This table reports summary statistics for manipulated cycles, the top decile by *PushIntensity* (*PushIntensity* $\geq$ 16.11), against normal cycles, the remaining 90%. The *Binance activity* and *Polymarket activity* panels report the same two flow rows in thousands of dollars (USDC on Polymarket): the absolute net order flow in bin 29 (the final 10 s), and the per-cycle median of the absolute net order flow across the body bins (bins 0–24, the pre-ramp window, the same body denominator that enters *PushIntensity*). Bin 29 is a single 10 s bin and the body figure is the cycle's typical body bin, so the two rows are on the same per-bin scale within each venue and across the two venues. *Binance activity* additionally reports the cross-cycle mean and median of *PushIntensity*. *Timing*: hour-of-day bucket (UTC) and weekday share. *Moneyness*: the share of cycles that are NTM (5m-contract Up-token price within 0.10 of 0.5 near TTM $\in$ [10, 12] s), computed over the cycles classifiable by that rule (99% of manipulated and 94% of normal cycles). All flow values are cross-cycle means.

	Manipulated cycles	Normal cycles
Number of cycles	1,613	14,460
Binance activity		
PushIntensity, mean	157.4	2.0
PushIntensity, median	57.8	0.9
Bin-29 |signed notional| ($k)	$1,714	$68
Body |signed notional|, median/bin ($k)	$18	$43
Polymarket activity		
Bin-29 |signed notional| ($k)	$2.83	$2.69
Body |signed notional|, median/bin ($k)	$0.39	$0.51
Timing		
Asia / overnight (21:00–07:00 UTC)	55.7%	39.9%
Europe, pre-US (07:00–13:00 UTC)	21.3%	25.5%
US & EU–US overlap (13:00–21:00 UTC)	23.1%	34.6%
Weekend (Sat–Sun)	43.6%	27.0%
Moneyness		
NTM share	18.7%	4.4%

play: the cycles in which moving the spot across the strike can actually flip the resolution. The *Moneyness* panel of Table 8 shows it does: among manipulated cycles 18.7% are NTM, against only 4.4% of normal cycles. Equivalently, a cycle that is NTM near the close is manipulated 33.0% of the time versus 9.1% for FFM cycles, a 3.6-fold higher rate. The two classifications are independent by construction: *PushIntensity* is built entirely from Binance signed flow and the moneyness cut entirely from the Polymarket Up-token price. So the concentration is not mechanical. It is the cross-sectional counterpart of the NTM activity and reversal results: the manipulation footprint is far more common exactly in the cycles where a push pays off.

6.4.3 Does manipulation succeed?

Does the push change the resolution? We classify each cycle by how confident Polymarket is just before the close: the favored-side price at TTM = 10 s, defined as $p_{\text{fav}} = \max(p_{\text{Up}},\ 1 - p_{\text{Up}})$, ranges from 0.5 (a coin flip) to 1 (a settled outcome). Within manipulated cycles, the push can go in the *same* direction as the favored side or in the *opposite* direction, toward the underdog. Table 9 reports, by confidence bin and push direction, how often the side the push backs wins the resolution; the same side's win rate in normal cycles, which carry no push, is the baseline. Because manipulated cycles are selected on the size of the bin-29 flow, not on its direction or the outcome, the comparison is not mechanical.

The clearest pattern is that opposite-direction manipulation works: a push against the favored side raises the underdog's win rate far above its no-push baseline in every confidence bin, most dramatically where Polymarket is most confident. In the [0.90, 1.00) bin the underdog wins just 1.0% of normal cycles but 34.2% of manipulated ones: a push against a side the market priced at $\geq 90\%$ reverses the resolution roughly a third of the time, versus once in a hundred absent a push. Same-direction pushes, by contrast, matter only near the money: they add ~ 18 pp to the favored side's win rate in the near-NTM bin but are flat, even marginally negative, once the favored side already wins the large majority of normal cycles, where there is little left to push for.

The second pattern concerns the direction of the pushes: the split between same- and opposite-direction pushes is close to even in *every* confidence bin, including the most lopsided one, where Polymarket prices the favored side at $\geq 90\%$ yet half the pushes still drive against it. Manipulators are not merely reinforcing near-certain outcomes; they routinely bet on overturning them, and the bet is attractive: an underdog priced below ten cents that wins a third of the time returns several times its stake in expectation. We quantify the realized profits below.

Table 9. Manipulated-cycle outcome rates by Polymarket-favored confidence, P3 BTC 5m. This table reports, by confidence bin and push direction, how often the side the push backs wins the resolution. Cycles are binned by the Polymarket-favored confidence at TTM = 10 s, $p_{\text{fav}} = \max(p_{\text{Up}}, 1 - p_{\text{Up}})$ (width 0.10, folded over Up/Down since the empirical Up-vs-Down asymmetry is negligible). Each bin has two rows for the two manipulation directions, defined relative to the Polymarket-favored side: *Same-direction*, where the bin-29 Binance push agrees with the favored side, and *Opposite-direction*, where it pushes against the favored side (toward the underdog). *# cycles* and *% cycles* give the number and share of the bin's manipulated cycles in each direction. *Win rate* is the rate at which the side the push targets, the favored side for same-direction, the underdog for opposite-direction, resolves as the winner among the manipulated cycles of that direction; *Win rate (normal cycles)* is the same side's unconditional win rate in the bin's normal cycles (which carry no push, hence no same/opposite split), serving as the no-manipulation baseline. Manipulated cycles are the top decile by *PushIntensity* (*PushIntensity* ≥ 16.11); normal cycles are the remaining 90%.

p_{fav} range	Manipulation direction	# cycles	% cycles	Win rate	Win rate (normal cycles)
[0.50, 0.60)	Same-direction	149	52%	76.5%	58.6%
	Opposite-direction	136	48%	65.4%	41.4%
[0.60, 0.70)	Same-direction	125	46%	76.0%	70.9%
	Opposite-direction	145	54%	68.3%	29.1%
[0.70, 0.80)	Same-direction	111	54%	82.9%	82.3%
	Opposite-direction	94	46%	54.3%	17.7%
[0.80, 0.90)	Same-direction	95	43%	92.6%	91.3%
	Opposite-direction	125	57%	54.4%	8.7%
[0.90, 1.00)	Same-direction	300	50%	96.7%	99.0%
	Opposite-direction	304	50%	34.2%	1.0%

6.4.4 Trader types and PnL

PnL calculation. To attribute profits and losses across trader types, we compute realized PnL at the wallet-cycle level. Each on-chain Polymarket fill identifies the two counterparties by Polygon wallet address; for each (wallet, cycle) we accumulate the wallet's signed cash flow and net Up/Down token positions from its fills, then settle the token positions at the Chainlink resolution price to obtain the cycle's realized PnL:

$$\text{PnL}_{w,c} \;=\; \Delta\text{cash}_{w,c} + N_{\text{Up},w,c}\,\mathbf{1}\{\text{Up wins}\} + N_{\text{Down},w,c}\,\mathbf{1}\{\text{Down wins}\},$$

where $\Delta\text{cash}_{w,c}$ is the wallet's USDC inflow during the cycle (positive when sells exceed buys, negative otherwise) and each token settles at \$1 for the winning side and \$0 for the losing side. Because each 5m market has its own token IDs, a wallet's pre-cycle position is identically zero and each cycle is a closed loop. This is *gross* PnL: exchange fees are not subtracted; the fee

schedule changes mid-sample, and net-of-fee PnL is left to future work.[21] Settlement uses the Chainlink resolution price where the on-chain archive has it; for the small fraction of cycles with a missing Chainlink print, we fall back to the last trade in [0, +30 s] after close, where Polymarket secondary trade convergence to the Chainlink outcome takes ∼10 s. The panel covers 16,073 cycles with on-chain Polymarket trading and 243,155 distinct wallets, which we treat as 243,155 traders (multi-wallet identity by a single trader is unobserved).

Trader-type identification. We partition traders into three mutually exclusive types. A trader is labeled a *manipulator* if it satisfies two criteria across the sample: it participated in at least five manipulated cycles, and its aggregate net PnL across those cycles is at least $2,000. The filter selects on participation and profitability in manipulated cycles, not on manipulation itself: a trader could in principle profit in those cycles by luck or by unrelated information.

A trader is labeled a *market maker* (MM) on the prediction-market contract if it does not pass the manipulator filter and satisfies three behavioral criteria, all defined on its trading across the sample. First, at least 85% of its fills have it on the resting (maker) side rather than crossing the spread. Second, conditional on being active in a cycle, it averages at least 50 fills per cycle, roughly one fill every six seconds. Third, its per-cycle median directional inventory at the close, defined as the absolute difference between its Up- and Down-token positions as a fraction of the gross shares it traded in that cycle, is at most 20%. The first two criteria capture passive execution and continuous quoting; the third separates liquidity provision from position-taking: an MM closes each cycle near delta-neutral rather than carrying a directional bet across resolution.

Every other trader is *retail/other*, the residual category. Under this classification the trader population breaks down to 821 manipulators (0.34% of the population), 227 MMs (0.09%), and 242,107 retail/other traders (99.57%).

[21] Fee-enabled crypto contracts charge only the aggressor: resting liquidity pays no exchange fee and is eligible for rebates funded out of taker-fee revenue. The taker fee is proportional to $p(1-p)$ in the contract price p, so it peaks at even odds, where short-horizon up/down contracts spend much of their lives. The schedule moved several times within our sample; the five-minute contracts were fee-enabled from launch, with a peak effective rate near 1.56% at even odds, raised to roughly 1.80% at the end of March 2026.

6.4.5 Profits and losses in manipulated cycles

Who profits and who pays in the manipulated cycles? Table 10 answers with four metrics per trader type, each split across manipulated and normal cycles: total net PnL, net PnL per cycle, the fraction of cycles in which the type loses money, and its share of the losses within a cycle. All four point to the same answer: the push transfers wealth from retail traders to manipulators, with market makers largely out of the way.

The PnL levels give the size of the transfer. Manipulators capture +\$8.2 M in aggregate across manipulated cycles, +\$5,096 per manipulated cycle. Retail traders fund nearly all of it: they lose −\$7.6 M, −\$4,715 per manipulated cycle. MMs largely dodge the transfer (−\$0.6 M, −\$381 per cycle), consistent with the Section 6.3.2 finding that MMs withdraw their book into the NTM close: they are rarely the counterparty because they do not quote in the cycles where the push is most likely.

The remaining two metrics show that the transfer is broad-based, not driven by a few large-loss cycles. Retail is net negative in 65% of manipulated cycles, against 48% of normal ones. The last metric works at the cycle level: the gross loss pool, defined as the sum across trader types of their net losses within a cycle, equals the sum of their gains by zero-sum, and a type's loss share is its loss divided by that pool. Weighted across cycles by pool size, retail bears 70.5% of manipulated-cycle losses versus 46.0% in normal cycles. The manipulator cohort still carries 20.2% of the manipulated-cycle pool even as the dominant net winner: that is the loss it books in the minority of cycles where its push fails, outweighed by its wins (in normal cycles, where it is more often a net loser, its share rises to 43.2%). The MM share is small in both regimes (9.3% and 10.8%): the loss flow runs from retail to manipulators directly.

Normal cycles show the reverse and confirm that the transfer is specific to manipulation. There, MMs collect the spread (+\$3.1 M, +\$215 per cycle), retail pays the ordinary cost of trading (−\$3.2 M, −\$221 per cycle), and manipulators break even (+\$0.1 M, +\$6 per cycle): the cohort that extracts \$5,096 per manipulated cycle has no edge in the remaining cycles. Retail's per-cycle loss is

Table 10. Trader-type PnL in manipulated vs normal cycles, P3 BTC 5m. This table reports four PnL metrics per trader type, split across manipulated and normal cycles. Trader types are mutually exclusive and exhaustive. *Manipulator*: traders with ≥5 manipulated cycles and ≥$2,000 aggregate manipulated-cycle PnL. *MM*: residual of that filter satisfying passive-rate ≥85%, ≥50 fills per cycle when active, and per-cycle median directional inventory ≤20% of gross shares traded. *Retail/other*: residual. Columns report four metrics, each split across manipulated (top-decile *PushIntensity*) and normal (other) cycles. *Total PnL* is the cohort's aggregate net PnL over all its (trader, cycle) observations of that type, in $M; within each cycle type it sums to zero across cohorts by construction. *PnL per cycle* is Total PnL divided by the number of cycles of that type (1,613 manipulated, 14,460 normal). *% losing cycles* is the fraction of cycles in which the cohort is net negative, where the (cohort, cycle) cell aggregates a cohort's PnL across its member traders within a cycle. *Loss share* is the cohort's share of the *gross* loss pool (the sum across cohorts of their losses within a cycle), aggregated across cycles with pool-size weights, so it sums to 100% down each regime, and a net-winning cohort still carries a positive share from the cycles in which it happens to lose.

	Total PnL		PnL per cycle		% losing cycles		Loss share	
	Manip.	Normal	Manip.	Normal	Manip.	Normal	Manip.	Normal
Manipulator	+$8.22M	+$0.09M	+$5,096	+$6	38.0%	62.4%	20.2%	43.2%
Market maker	−$0.62M	+$3.11M	−$381	+$215	58.6%	37.7%	9.3%	10.8%
Retail / other	−$7.61M	−$3.20M	−$4,715	−$221	65.3%	47.7%	70.5%	46.0%

roughly twenty times larger when a push is present, and its counterparty switches from the MMs earning the spread to the manipulators. The full trader-type attribution is robust to the manipulator and MM thresholds (Appendix E).

6.4.6 Ruling out the gamma-hedging channel

The leading innocent explanation for the bin-29 push is hedging: a Polymarket maker who has sold Up tokens through the body offsets the resulting short by trading BTC, and that hedge, not manipulation, is what moves the spot across the strike. The channel requires no intent. Retail buys Up tokens through the body; makers absorb the demand by selling Up; the accumulated short must be hedged in the spot; the hedge pushes the close across the strike, and the body-period Up buyers profit. Two facts rule this out. The first, and the sharper, is the state of the contract when the push lands: in the cycles a push most often flips, the maker has essentially nothing to hedge. The second is the timing of the spot flow.

There is nothing to hedge in the cycles a push flips. A binary contract's sensitivity to the spot price is a delta that collapses away from the strike. When the favored side is priced near 0.5 the contract behaves almost like the spot itself, and a maker who is short it carries a real directional

exposure; but once the price moves toward 0 or 1, the contract is all but settled, its value barely responds to a further spot move, and the short position it represents is close to delta-flat. A maker holding such a position has almost nothing to hedge. Yet these near-certain cycles are exactly the ones the push overturns. In the [0.90, 1.00) confidence bin — where Polymarket priced the favored side at ≥ 90% ten seconds before the close, so the binary's delta has all but vanished — a push against the favored side reverses the resolution 34.2% of the time, against 1.0% in cycles with no push (Table 9). A trade placed when there is nothing to hedge cannot be a hedge. Whatever hedging response the spot flow contains can only come *after* the push, once it has dragged the price back toward the strike and revived the very exposure the hedging story needs; it cannot be the push's cause. The hedging channel is coherent only in the small near-the-money slice where the binary's delta is still live (about 6% of cycles); it has no purchase in the near-certain cycles that the push flips one time in three.

The timing of the spot flow is wrong for a hedge. The second fact concerns the cycles where there *is* something to hedge, and asks when the spot trading arrives. Under dynamic hedging, the standard practice in liquid markets, a maker hedges each marginal share as the short accumulates: if the cohort's short builds steadily through the body, Binance buying should ramp in step. Figure 12 shows it does not. The Polymarket positions build continuously — the manipulator cohort's net Up holdings climb to roughly 8,800 shares by the close, and the MM cohort absorbs a short that grows to roughly 1,350 shares as the final ten seconds begin, a near-linear accumulation rather than a flat book that dumps at the end. The matching spot hedge never appears through the body. Cumulative Binance order flow stands at just −$46 k at second 215 of the 300-second cycle, essentially zero and, if anything, slightly against the eventual push direction; the buying arrives only in the final minute, and the final ten seconds alone carry roughly $1.7 M, essentially the cycle's entire net order flow in a single step. A position assembled only in the closing seconds is not a hedge accumulated against an exposure that built over the cycle; it is a single trade timed to move the settlement price.

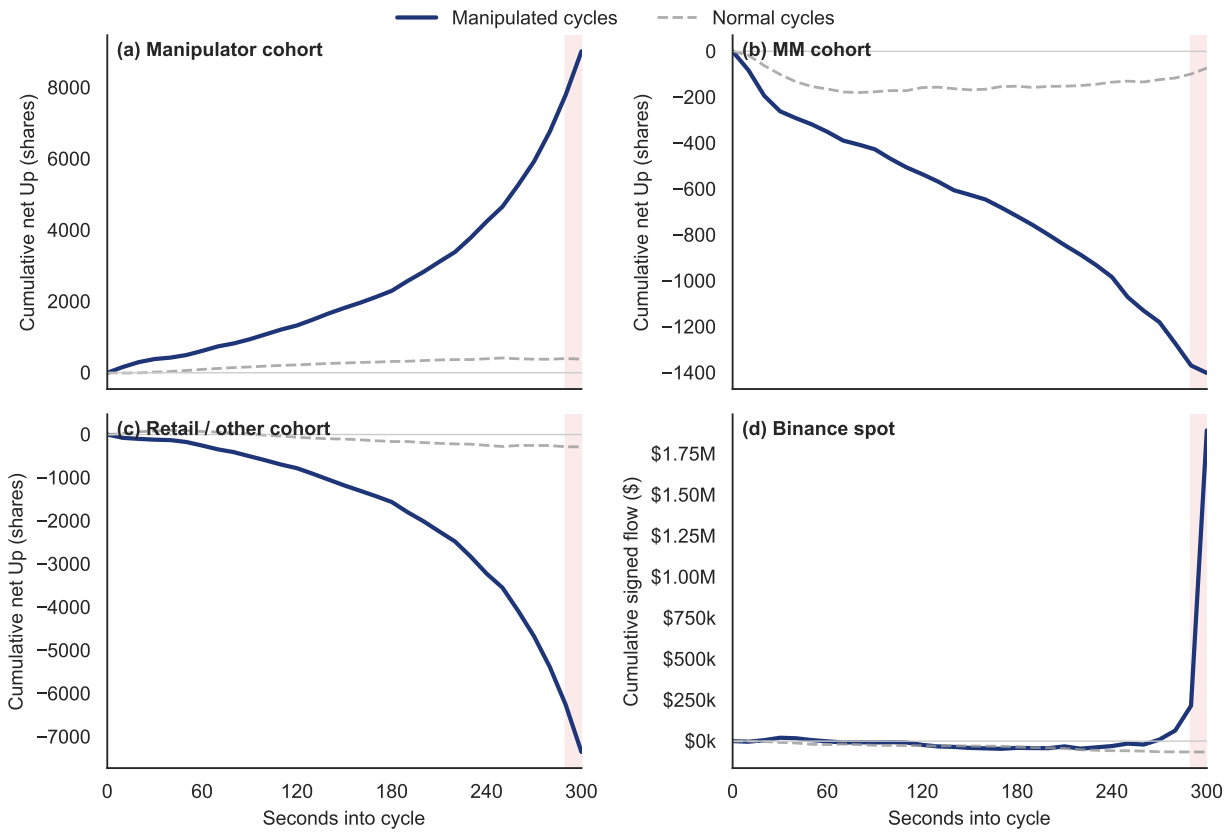

Figure 12. Cumulative net position trajectory across the cycle, push-up vs. normal cycles, P3 BTC 5m. This figure plots, in each panel, the mean per-cycle cumulative position (or order flow) against seconds into the cycle, starting from zero at the open. The solid line conditions on *push-up* cycles, manipulated cycles whose near-settlement push is upward (labeled "manipulated cycles" in the legend), so that positions are not averaged across the two push directions; push-down cycles are reported in Appendix F. Panels **(a)**–**(c)** plot the net Up-token holdings (shares) of the manipulator, MM, and retail/other cohorts; panel **(d)** plots the cumulative Binance spot order flow ($). Dashed gray lines: normal cycles as baselines. Pink shading: the final-10 s region (seconds 290–300).

The magnitudes rule out a deferred hedge. One variant survives the timing test: a maker who holds the naked delta through the body and hedges only as gamma surges at the very close would also trade in a single late burst. The magnitudes rule it out. A short binary position can lose at most its notional, so the MM cohort's roughly 1,350-share short puts about $1,400 per cycle at risk — three orders of magnitude below the $1.7 M spot trade the channel would attribute to it. No maker, let alone all 28 active in the average manipulated cycle, hedges a four-figure maximum loss with a seven-figure spot position. The more credible reading is that the bin-29 push is direct spot-side action, coordinated with the manipulator cohort's body-period positioning on Polymarket, with at most a residual hedging component in the bin-29 buying.

7 Conclusion

Prediction markets increasingly list contracts that settle on an asset price that holders can move by trading the underlying. A model and the on-chain record of Polymarket's five-minute Bitcoin contract agree: such a contract degrades price discovery in the underlying and transfers wealth from ordinary traders to the few who manipulate it, even as it makes the underlying more liquid.

The mechanism is a settlement-time push: a holder of the binary contract trades the underlying in the final seconds to move the settlement reference toward her side. The spot maker cannot separate this uninformed push from ordinary flow, so price discovery falls and wealth passes to the manipulator and, under a spread floor, to the prediction-market maker. The near-settlement spike appears only once the five-minute contract exists; the price move reverts within ten seconds; and the effect concentrates in the still-live cycles, where even a side priced above ninety percent is overturned about a third of the time, against one in a hundred otherwise. Some $8.2 million was transferred over two months, mostly from retail.

The harm can be fixed: it tracks the contract horizon, the remedy our model provides. A manipulator enters only when a feasible push can carry the price across the strike cheaply enough, so a longer horizon makes pivotal price paths rarer and can eliminate the incentive. The settlement rule supplies parallel levers, from averaging or randomizing the close to capping the binary position or taxing settlement-window spot trades. Polymarket's adoption of a TWAP benchmark provides a setting to study whether temporal averaging reduces settlement-time manipulation (Polymarket, 2026). Manipulation here is a settlement-design problem.

Appendix

A Omitted Proofs

A.1 Stage 2 Quotes and Posteriors

We prove Lemma 1 in this appendix.

Outcome probabilities. The informed trader enters with probability π. She buys if $V = +\sigma$ and sells if $V = -\sigma$ conditional on entry. A liquidity trader enters with probability $1 - \pi$. She buys or sells with equal probability conditional on entry. Therefore,

$$\mathbb{P}\left(Y_2^{\mathrm{S}} = + \mid V = \pm\sigma\right) = \frac{1 \pm \pi}{2}, \qquad \mathbb{P}\left(Y_2^{\mathrm{S}} = - \mid V = \pm\sigma\right) = \frac{1 \mp \pi}{2}.$$

Marginalizing over V gives

$$q_+ = q_- = \frac{1}{2}.$$

Posterior and quotes. Bayes' rule gives

$$\rho_+ = \mathbb{P}\left(V = +\sigma \mid Y_2^{\mathrm{S}} = +\right) = \frac{\mathbb{P}\left(Y_2^{\mathrm{S}} = + \mid V = +\sigma\right)\mathbb{P}\left(V = +\sigma\right)}{q_+} = \frac{\frac{1+\pi}{2} \cdot \frac{1}{2}}{\frac{1}{2}} = \frac{1+\pi}{2},$$

$$\rho_- = 1 - \rho_+ = \frac{1-\pi}{2}.$$

Break-even pricing requires

$$\frac{1-\pi}{2} A_2^{\mathrm{S}} = \frac{\pi}{2}\left(\sigma - A_2^{\mathrm{S}}\right)$$

where the left-hand side is the expected profit from selling to a liquidity trader and the right-hand side is the expected loss from selling to the informed trader. Solving gives

$$A_2^{\mathrm{S}} = \sigma\pi = \bar{V}_+.$$

By symmetry,

$$B_2^{\mathrm{S}} = -\sigma\pi = \bar{V}_-.$$

The quotes are indeed inside $(-\sigma, \sigma)$ for $\pi \in (0, 1)$, confirming that the informed trader will never stay out at stage 2.

Posterior variance. Direct calculation of $\nu_{h_2} = 4\sigma^2\rho_{h_2}(1 - \rho_{h_2})$ gives

$$\nu_+ = \nu_- = \sigma^2\left(1 - \pi^2\right).$$

A.2 Stage 3 Quotes and Posteriors

We prove Proposition 1 in this appendix.

State-conditional outcome probabilities. The manipulator enters with probability p_{M}. She buys with probability α_{h_2} and stays out with probability $1 - \alpha_{h_2}$ conditional on entry. The informed trader enters with probability $(1 - p_{\mathrm{M}})\pi$. She buys if $V = +\sigma$ and sells if $V = -\sigma$ conditional on entry. A liquidity trader enters with probability $(1 - p_{\mathrm{M}})(1 - \pi)$. She buys or sells with equal probability conditional on entry. Therefore,

$$\begin{aligned}
\mathbb{P}\left(Y_3^{\mathrm{S}} = + \mid V = \pm\sigma, Y_2^{\mathrm{S}} = h_2\right) &= p_{\mathrm{M}}\alpha_{h_2} + (1 - p_{\mathrm{M}})\frac{1 \pm \pi}{2}, \\
\mathbb{P}\left(Y_3^{\mathrm{S}} = \mathbb{0} \mid V = \pm\sigma, Y_2^{\mathrm{S}} = h_2\right) &= p_{\mathrm{M}}(1 - \alpha_{h_2}), \\
\mathbb{P}\left(Y_3^{\mathrm{S}} = - \mid V = \pm\sigma, Y_2^{\mathrm{S}} = h_2\right) &= (1 - p_{\mathrm{M}})\frac{1 \mp \pi}{2}.
\end{aligned}$$

Marginalizing over $V \mid Y_2^{\mathrm{S}} = h_2$ gives

$$
\begin{aligned}
q_{h_2+} &= p_{\mathrm{M}}\alpha_{h_2} + (1-p_{\mathrm{M}})\left[\frac{1-\pi}{2} + \pi\rho_{h_2}\right], \\
q_{h_2\emptyset} &= p_{\mathrm{M}}\left(1-\alpha_{h_2}\right), \\
q_{h_2-} &= (1-p_{\mathrm{M}})\left[\frac{1-\pi}{2} + \pi\left(1-\rho_{h_2}\right)\right].
\end{aligned}
$$

Break-even ask. A stage-3 buy comes from the manipulator or a liquidity trader, both uninformed, or from the informed trader, who buys only when $V = +\sigma$. Break-even pricing requires

$$
\left[p_{\mathrm{M}}\alpha_{h_2} + (1-p_{\mathrm{M}})\,\frac{1-\pi}{2}\right]\left(A_{h_2} - \bar{V}_{h_2}\right) = (1-p_{\mathrm{M}})\,\pi\rho_{h_2}\left(\sigma - A_{h_2}\right)
$$

where the left-hand side is the expected profit from selling to an uninformed buyer and the right-hand side is the expected loss from selling to the informed trader. Solving gives

$$
A_{h_2} = \sigma\,\frac{p_{\mathrm{M}}\alpha_{h_2}\left(2\rho_{h_2}-1\right) + \frac{1-p_{\mathrm{M}}}{2}\left[2\rho_{h_2}-1+\pi\right]}{p_{\mathrm{M}}\alpha_{h_2} + \frac{1-p_{\mathrm{M}}}{2}\left[1+\left(2\rho_{h_2}-1\right)\pi\right]}. \tag{4}
$$

Substituting the stage-2 posteriors $\rho_+ = \frac{1+\pi}{2}$ and $\rho_- = \frac{1-\pi}{2}$ gives

$$
A_+ = \frac{\sigma\pi\left[p_{\mathrm{M}}\alpha_+ + (1-p_{\mathrm{M}})\right]}{p_{\mathrm{M}}\alpha_+ + (1-p_{\mathrm{M}})\,\frac{1+\pi^2}{2}}, \qquad A_- = \frac{-\sigma\pi\,p_{\mathrm{M}}\alpha_-}{p_{\mathrm{M}}\alpha_- + (1-p_{\mathrm{M}})\,\frac{1-\pi^2}{2}}.
$$

Break-even bid. The manipulator never sells, so a stage-3 sell comes only from a liquidity trader or from the informed trader, who sells only when $V = -\sigma$. Break-even pricing requires

$$
(1-p_{\mathrm{M}})\,\frac{1-\pi}{2}\left(\bar{V}_{h_2} - B_{h_2}\right) = (1-p_{\mathrm{M}})\,\pi\left(1-\rho_{h_2}\right)\left(\sigma + B_{h_2}\right)
$$

where the left-hand side is the expected profit from buying from a liquidity seller and the right-hand side is the expected loss from buying from the informed trader. The common factor $(1-p_{\mathrm{M}})$ cancels,

so the bid is independent of α_{h_2} and p_{M}. Solving gives

$$B_{h_2} = \sigma \frac{(2\rho_{h_2} - 1) - \pi}{1 - (2\rho_{h_2} - 1)\pi}. \tag{5}$$

Substituting the stage-2 posteriors $\rho_+ = \frac{1+\pi}{2}$ and $\rho_- = \frac{1-\pi}{2}$ gives

$$B_+ = 0, \quad B_- = \frac{-2\sigma\pi}{1 + \pi^2}.$$

No-trade posterior. No-trade occurs only when the manipulator is present, and she is uninformed. When it has positive probability ($\alpha_{h_2} < 1$) the posterior equals the prior entering stage 3; at $\alpha_{h_2} = 1$ the event has zero probability and is assigned the same belief by convention. Hence

$$\mathbb{P}\left(V = \pm\sigma \mid Y_3^{\mathrm{S}} = \emptyset, Y_2^{\mathrm{S}} = h_2\right) = \mathbb{P}\left(V = \pm\sigma \mid Y_2^{\mathrm{S}} = h_2\right)$$

and

$$\bar{V}_{h_2\emptyset} = \bar{V}_{h_2} = \sigma\left(2\rho_{h_2} - 1\right),$$

constant in α_{h_2} and p_{M}.

Monotonicities. Recall the break-even pricing condition for the ask

$$\left[p_{\mathrm{M}}\alpha_{h_2} + (1 - p_{\mathrm{M}})\frac{1 - \pi}{2}\right]\left(A_{h_2} - \bar{V}_{h_2}\right) = (1 - p_{\mathrm{M}})\pi\rho_{h_2}\left(\sigma - A_{h_2}\right).$$

For $p_{\mathrm{M}} \in (0, 1)$ and $\pi \in (0, 1)$ both probability coefficients are strictly positive, so the balance forces $A_{h_2} \in \left(\bar{V}_{h_2}, \sigma\right)$, which we can also verify from the closed form above. The right-hand side does not contain α_{h_2}, while the left-hand side strictly increases in α_{h_2}. Raising α_{h_2} therefore moves the uninformed-profit side above the informed-loss side at the prevailing quote. Hence A_{h_2} has to strictly decrease to restore equality. For monotonicity in p_{M}, divide both sides by $1 - p_{\mathrm{M}} > 0$ to get the

following, then apply a similar argument.

$$\left[\frac{p_{\mathrm{M}}}{1-p_{\mathrm{M}}}\,\alpha_{h_2}+\frac{1-\pi}{2}\right]\left(A_{h_2}-\bar{V}_{h_2}\right)=\pi\rho_{h_2}\left(\sigma-A_{h_2}\right).$$

The bid and the no-trade posterior mean are independent of both p_{M} and α_{h_2}, as the closed forms above show.

A.3 Selected Stage-3 Continuation Equilibrium

Definition 2 (Stable stage-3 fixed point)**.** A stage-3 fixed point α_{h_2} is *stable* if for some $\eta>0$,

$$\left(\alpha-\alpha_{h_2}\right)\Gamma_{h_2}\left(\alpha\right)\leq 0\quad\text{for all }\alpha\in\left(\alpha_{h_2}-\eta,\alpha_{h_2}+\eta\right)\cap[0,1]\,. \tag{6}$$

Equivalently, $\Gamma_{h_2}\geq 0$ just left and ≤ 0 just right of α_{h_2}. This excludes strict crossings from below, where local pressure pushes the conjecture away. Write $\mathcal{E}_{h_2}$ for the set of stable fixed points.

We prove Proposition 2 in this appendix. The closed-form expression for A_{h_2} shows that it is continuous in α_{h_2}, which implies that Γ_{h_2} is continuous in α_{h_2} as well since H is continuous. Moreover, $H\left(\cdot/\varepsilon\right)$ is constant outside $(-\varepsilon,\varepsilon)$ and affine inside, and $A_{h_2}\left(\alpha_{h_2}\right)$ is continuous and strictly monotone, so Γ_{h_2} is strictly monotone on each of the at most three intervals on which the ask lies above, inside, or below $[-\varepsilon,\varepsilon]$.

Specifically, if the ask lies inside $[-\varepsilon,\varepsilon]$,

$$\Gamma'_{h_2}\left(\alpha_{h_2}\right)=\left(\frac{1}{2\varepsilon}-1\right)A'_{h_2}\left(\alpha_{h_2}\right)\neq 0$$

because $\varepsilon\neq\frac{1}{2}$.[22] If the ask lies above ε or below $-\varepsilon$, then

$$\Gamma'_{h_2}\left(\alpha_{h_2}\right)=-A'_{h_2}\left(\alpha_{h_2}\right)>0.$$

[22]The exclusion of the degenerate case $\varepsilon=\frac{1}{2}$ plays no role below, as the canonical regime of Proposition 4 has $\varepsilon<\sigma\pi<\frac{1}{2}$.

In particular, Γ_{h_2} has finitely many zeros. If $\Gamma_{h_2} \leq 0$ on a right-neighborhood of 0, then $\alpha_{h_2} = 0$ is stable. If $\Gamma_{h_2} \geq 0$ on a left-neighborhood of 1, then $\alpha_{h_2} = 1$ is stable. Otherwise neither holds, so every right-neighborhood of 0 contains a point with $\Gamma_{h_2} > 0$ and every left-neighborhood of 1 a point with $\Gamma_{h_2} < 0$; among the finitely many zeros lying between two such points, the first downcrossing has $\Gamma_{h_2} \geq 0$ immediately to its left and ≤ 0 immediately to its right, hence is stable. Therefore, there is at least one stable fixed point, which allows us to select the smallest $\alpha^*_{h_2} = \min \mathcal{E}_{h_2}$.

A.4 Price Discovery

We prove Proposition 3 in this appendix. By the law of total variance within each h_2,

$$\sum_{h_3} q_{h_2 h_3} \nu_{h_2 h_3} = \nu_{h_2} - \mathrm{Var}\left(\bar{V}_{h_2 h_3} \mid Y_2^{\mathrm{S}} = h_2\right),$$

where

$$\mathrm{Var}\left(\bar{V}_{h_2 h_3} \mid Y_2^{\mathrm{S}} = h_2\right) = \sum_{h_3} q_{h_2 h_3} \left(\bar{V}_{h_2 h_3} - \bar{V}_{h_2}\right)^2$$

is the dispersion of the stage-3 posterior mean within each h_2. Since stage 2 has no manipulator, its outcome probabilities q_{h_2} and conditional variances ν_{h_2} are all independent of p_{M}. Hence, it suffices to show, for each h_2, that

$$\mathrm{Var}\left(\bar{V}_{h_2 h_3} \mid Y_2^{\mathrm{S}} = h_2\right)(p_{\mathrm{M}}) < \mathrm{Var}\left(\bar{V}_{h_2 h_3} \mid Y_2^{\mathrm{S}} = h_2\right)(0).$$

By Proposition 1, $\bar{V}_{h_2 \mathtt{0}}(p_{\mathrm{M}}) = \bar{V}_{h_2}$, so the $h_3 = \mathtt{0}$ term does not contribute to the within-h_2 dispersion. A sell order never comes from the manipulator, so the posterior mean after a sell is unaffected by p_{M}, but its probability is thinned by the manipulator's presence:

$$\bar{V}_{h_2 -}(p_{\mathrm{M}}) = \bar{V}_{h_2 -}(0), \qquad q_{h_2 -}(p_{\mathrm{M}}) = (1 - p_{\mathrm{M}})\, q_{h_2 -}(0),$$

so the $h_3 = -$ term equals $(1 - p_{\mathrm{M}})$ times its benchmark value.

A buy order only moves the posterior mean when it comes from the informed trader:

$$
\begin{aligned}
& q_{h_2+}(p_{\mathrm{M}}) \left(\bar{V}_{h_2+}(p_{\mathrm{M}}) - \bar{V}_{h_2} \right) \\
=& \mathbb{E}\left[\left(V - \bar{V}_{h_2} \right) \mathbf{1}\{ Y_3^{\mathrm{S}} = + \} \mid Y_2^{\mathrm{S}} = h_2 \right] \\
=& \underbrace{\mathbb{E}\left[\left(V - \bar{V}_{h_2} \right) \mathbf{1}\{ Y_3^{\mathrm{S}} = + \text{ from the manipulator} \} \mid Y_2^{\mathrm{S}} = h_2 \right]}_{=0 \quad (\text{buy} \perp V)} \\
& + \underbrace{\mathbb{E}\left[\left(V - \bar{V}_{h_2} \right) \mathbf{1}\{ Y_3^{\mathrm{S}} = + \text{ from a liquidity trader} \} \mid Y_2^{\mathrm{S}} = h_2 \right]}_{=0 \quad (\text{buy} \perp V)} \\
& + \underbrace{\mathbb{E}\left[\left(V - \bar{V}_{h_2} \right) \mathbf{1}\{ Y_3^{\mathrm{S}} = + \text{ from the informed trader} \} \mid Y_2^{\mathrm{S}} = h_2 \right]}_{= (1 - p_{\mathrm{M}}) \pi \rho_{h_2} \left(\sigma - \bar{V}_{h_2} \right)} \\
=& (1 - p_{\mathrm{M}}) \, \pi \rho_{h_2} \left(\sigma - \bar{V}_{h_2} \right),
\end{aligned}
$$

which equals $(1 - p_{\mathrm{M}})$ times the benchmark quantity $q_{h_2+}(0) \left(\bar{V}_{h_2+}(0) - \bar{V}_{h_2} \right)$. Hence,

$$
\begin{aligned}
& q_{h_2+}(p_{\mathrm{M}}) \left(\bar{V}_{h_2+}(p_{\mathrm{M}}) - \bar{V}_{h_2} \right)^2 \\
=& (1 - p_{\mathrm{M}})^2 \frac{q_{h_2+}(0)}{q_{h_2+}(p_{\mathrm{M}})} \, q_{h_2+}(0) \left(\bar{V}_{h_2+}(0) - \bar{V}_{h_2} \right)^2 \\
\leq& (1 - p_{\mathrm{M}}) \, q_{h_2+}(0) \left(\bar{V}_{h_2+}(0) - \bar{V}_{h_2} \right)^2
\end{aligned}
$$

because $q_{h_2+}(p_{\mathrm{M}}) = p_{\mathrm{M}} \alpha_{h_2} + (1 - p_{\mathrm{M}}) \, q_{h_2+}(0) \geq (1 - p_{\mathrm{M}}) \, q_{h_2+}(0)$.

Summing the three outcomes,

$$
\mathrm{Var}\left(\bar{V}_{h_2 h_3} \mid Y_2^{\mathrm{S}} = h_2 \right)(p_{\mathrm{M}}) \leq (1 - p_{\mathrm{M}}) \, \mathrm{Var}\left(\bar{V}_{h_2 h_3} \mid Y_2^{\mathrm{S}} = h_2 \right)(0) \, .
$$

With $\pi > 0$ the informed trader moves the benchmark posterior, so

$$
\mathrm{Var}\left(\bar{V}_{h_2 h_3} \mid Y_2^{\mathrm{S}} = h_2 \right)(0) > 0.
$$

Therefore, for any $p_{\mathrm{M}} \in (0,1)$,

$$\mathrm{Var}\left(\bar{V}_{h_2 h_3} \mid Y_2^{\mathrm{S}} = h_2\right)(p_{\mathrm{M}}) < \mathrm{Var}\left(\bar{V}_{h_2 h_3} \mid Y_2^{\mathrm{S}} = h_2\right)(0)\,.$$

A.5 Canonical-Regime Equilibrium Calculations

We prove Proposition 4 in this appendix, which focuses on the canonical regime

$$\varepsilon < \sigma\pi < \frac{1}{2}.$$

A.5.1 Selected stage-3 continuation equilibrium

We find the selected continuation equilibrium at each price path by a three-step test of Proposition 2:

Step 1: Check whether $\alpha_{h_2} = 0$ is a stable fixed point.

Step 2: If $\alpha_{h_2} = 0$ fails, locate the interior fixed point by solving $\Gamma_{h_2}(\alpha_{h_2}) = 0$.

Step 3: If no interior fixed point exists either, the continuation equilibrium can only be $\alpha_{h_2} = 1$.

Favorable price path $h_2 = +$: pure stay-out equilibrium. For any $p_{\mathrm{M}} \in (0,1)$ and any $\alpha_+ \in [0,1]$, $A_+(\alpha_+; p_{\mathrm{M}}) \geq \bar{V}_+ = \sigma\pi > \varepsilon$, so $H(A_+(\alpha_+; p_{\mathrm{M}})/\varepsilon) = H\left(\bar{V}_+/\varepsilon\right) = 1$. Hence

$$\Gamma_+(\alpha_+; p_{\mathrm{M}}) = 1 - 1 - [A_+(\alpha_+; p_{\mathrm{M}}) - \sigma\pi] = -[A_+(\alpha_+; p_{\mathrm{M}}) - \sigma\pi] \leq 0.$$

Therefore, for any $p_{\mathrm{M}} \in (0,1)$, $\alpha_+^* = 0$ is a stable fixed point and is clearly the smallest. Intuitively, after a favorable stage-2 buy the posterior already sits above the settlement band ($\bar{V}_+ = \sigma\pi > \varepsilon$), so the contract pays 1 whether or not the manipulator buys again; a stage-3 buy cannot raise her settlement payoff and only incurs spot cost, so she stays out.

Unfavorable price path $h_2 = -$: mixed or pure push. Fix $p_M \in (0,1)$. At $\alpha_- = 0$, Proposition 1 gives $A_-(0; p_M) = 0$, while $\bar{V}_- = -\sigma\pi < -\varepsilon$. Therefore

$$\Gamma_-(0; p_M) = \frac{1}{2} - 0 - [0 - (-\sigma\pi)] = \frac{1}{2} - \sigma\pi > 0,$$

so $\alpha_- = 0$ is not a fixed point. For an interior fixed point with $A_- \in [-\varepsilon, \varepsilon]$, the fixed-point condition is

$$\frac{A_- + \varepsilon}{2\varepsilon} - (A_- + \sigma\pi) = 0,$$

which gives the candidate mixed-branch stage-3 ask

$$A_- = \frac{\varepsilon(2\sigma\pi - 1)}{1 - 2\varepsilon} \in (-\varepsilon, 0).$$

Inverting the closed-form expression in Proposition 1 gives the candidate mixed-branch intensity

$$\alpha_-(p_M) = \frac{1 - p_M}{p_M} \frac{\varepsilon\left(1 - \pi^2\right)(1 - 2\sigma\pi)}{2(\sigma\pi - \varepsilon)} > 0.$$

This candidate is admissible if and only if

$$\alpha_-(p_M) < 1 \quad \Longleftrightarrow \quad p_M > p_0 := \frac{\varepsilon\left(1 - \pi^2\right)(1 - 2\sigma\pi)}{2(\sigma\pi - \varepsilon) + \varepsilon\left(1 - \pi^2\right)(1 - 2\sigma\pi)}.$$

Since A_- strictly decreases in α_- by Corollary 1, we can verify that Γ_- crosses zero from positive to negative as α_- crosses this candidate from below to above. Thus, for $p_M \in (p_0, 1)$, this candidate is the selected continuation equilibrium. For $p_M \in (0, p_0]$, this candidate would be equal to or greater than 1, which is not admissible, so the selected continuation equilibrium is the endpoint $\alpha_-^* = 1$:

$$\alpha_-^*(p_M) = \begin{cases} 1, & 0 < p_M \le p_0, \\ \dfrac{1 - p_M}{p_M} \dfrac{\varepsilon\left(1 - \pi^2\right)(1 - 2\sigma\pi)}{2(\sigma\pi - \varepsilon)}, & p_0 < p_M < 1. \end{cases}$$

The corresponding equilibrium ask is

$$A_{-}^{*}(p_{\mathrm{M}}) = \begin{cases} -\dfrac{2\sigma\pi p_{\mathrm{M}}}{(1-\pi^2)+(1+\pi^2)\,p_{\mathrm{M}}}, & 0 < p_{\mathrm{M}} \le p_0, \\ -\dfrac{\varepsilon\,(1-2\sigma\pi)}{1-2\varepsilon}, & p_0 < p_{\mathrm{M}} < 1. \end{cases}$$

On both branches, this ask remains in the settlement band.[23] Therefore the settlement probability induced by a stage-3 push after an unfavorable price path is

$$H\left(A_{-}^{*}(p_{\mathrm{M}})/\varepsilon\right) = \begin{cases} \dfrac{1}{2} - \dfrac{\sigma\pi p_{\mathrm{M}}}{\varepsilon\left[(1-\pi^2)+(1+\pi^2)\,p_{\mathrm{M}}\right]}, & 0 < p_{\mathrm{M}} \le p_0, \\ \dfrac{\sigma\pi-\varepsilon}{1-2\varepsilon}, & p_0 < p_{\mathrm{M}} < 1. \end{cases}$$

The per-push spot cost after a stage-2 sell is

$$A_{-}^{*}(p_{\mathrm{M}}) - \bar{V}_{-} = \begin{cases} \dfrac{\sigma\pi\left(1-\pi^2\right)(1-p_{\mathrm{M}})}{(1-\pi^2)+(1+\pi^2)\,p_{\mathrm{M}}}, & 0 < p_{\mathrm{M}} \le p_0, \\ \dfrac{\sigma\pi-\varepsilon}{1-2\varepsilon}, & p_0 < p_{\mathrm{M}} < 1. \end{cases}$$

The four displayed quantities are continuous at $p_{\mathrm{M}} = p_0$, where the pure-push and mixed branches meet.

A.5.2 Optimal manipulator entry probability

Both the pricing of the binary contract and the manipulator's stage-1 entry problem use the selected continuation equilibrium from Appendix A.5.1. Combining Appendix A.1, the general ask and bid formulas (4) and (5), and Appendix A.5.1, the stage-3 asks, bids, and no-trade posterior means used

[23] On the mixed branch, the ask lies in $(-\varepsilon, 0)$ by Assumption 1. On the pure-push branch, it is clearly below zero, and it is easy to verify that the ask is strictly decreasing in p_{M}. The two branches meet at $p_{\mathrm{M}} = p_0$.

below are

$$A_+^* = \frac{2\sigma\pi}{1+\pi^2}, \qquad A_-^*(p_M) = \begin{cases} -\dfrac{2\sigma\pi p_M}{(1-\pi^2)+(1+\pi^2)p_M}, & 0 < p_M \le p_0, \\ -\dfrac{\varepsilon(1-2\sigma\pi)}{1-2\varepsilon}, & p_0 < p_M < 1, \end{cases}$$

$$B_+ = 0, \qquad B_- = -\frac{2\sigma\pi}{1+\pi^2},$$

$$\bar{V}_+ = \sigma\pi, \qquad \bar{V}_- = -\sigma\pi.$$

Pure-push branch $p_M \in (0, p_0]$. On this branch, $\alpha_+^* = 0$ and $\alpha_-^* = 1$. As a result, there is no spot cost after a stage-2 buy, while the spot cost after a stage-2 sell is $A_-^*(p_M) - \bar{V}_-$. The per-path expected binary-contract payoffs are:

$$\begin{aligned}
\Pi_M^+(p_M, +) &= \alpha_+^* H(A_+^*/\varepsilon) + (1-\alpha_+^*) H\left(\bar{V}_+/\varepsilon\right) = H\left(\bar{V}_+/\varepsilon\right) = 1, \\
\Pi_M^+(p_M, -) &= \alpha_-^* H(A_-^*(p_M)/\varepsilon) + (1-\alpha_-^*) H\left(\bar{V}_-/\varepsilon\right) = H(A_-^*(p_M)/\varepsilon), \\
\Pi_L^+(p_M, +) &= q_{++}(0) H(A_+^*/\varepsilon) + q_{+-}(0) H(B_+/\varepsilon) = \frac{1+\pi^2}{2}\cdot 1 + \frac{1-\pi^2}{2}\cdot\frac{1}{2} = \frac{3+\pi^2}{4}, \\
\Pi_L^+(p_M, -) &= q_{-+}(0) H(A_-^*(p_M)/\varepsilon) + q_{--}(0) H(B_-/\varepsilon) = \frac{1-\pi^2}{2} H(A_-^*(p_M)/\varepsilon),
\end{aligned}$$

where $q_{h_2 h_3}(0)$ denotes the conditional stage-3 outcome probabilities defined in Section 4.5 evaluated at $p_M = 0$.

The manipulator's expected binary-contract payoff net of expected spot manipulation cost is

$$
\begin{aligned}
&\Pi_{\mathrm{M}}^{+}\left(p_{\mathrm{M}}\right)-K^{+}\left(p_{\mathrm{M}}\right)\\
=&q_{+}\Pi_{\mathrm{M}}^{+}\left(p_{\mathrm{M}},+\right)+q_{-}\Pi_{\mathrm{M}}^{+}\left(p_{\mathrm{M}},-\right)-q_{+}\cdot 0-q_{-}\left[A_{-}^{*}\left(p_{\mathrm{M}}\right)-\bar{V}_{-}\right]\\
=&\frac{1}{2}+\frac{1}{2}H\left(A_{-}^{*}\left(p_{\mathrm{M}}\right)/\varepsilon\right)-\frac{1}{2}\left[A_{-}^{*}\left(p_{\mathrm{M}}\right)-\bar{V}_{-}\right]\\
=&\frac{1}{2}+\frac{1}{2}\cdot\frac{\varepsilon\left[\left(1-\pi^{2}\right)+\left(1+\pi^{2}\right)p_{\mathrm{M}}\right]-2\sigma\pi p_{\mathrm{M}}}{2\varepsilon\left[\left(1-\pi^{2}\right)+\left(1+\pi^{2}\right)p_{\mathrm{M}}\right]}-\frac{1}{2}\cdot\frac{\sigma\pi\left(1-\pi^{2}\right)\left(1-p_{\mathrm{M}}\right)}{\left(1-\pi^{2}\right)+\left(1+\pi^{2}\right)p_{\mathrm{M}}}\\
=&\frac{1}{2}+\frac{\left(1-\pi^{2}\right)\left(1-2\sigma\pi\right)\left(1-p_{\mathrm{M}}/p_{0}\right)}{4\left[\left(1-\pi^{2}\right)+\left(1+\pi^{2}\right)p_{\mathrm{M}}\right]}.
\end{aligned}
$$

The unconstrained shadow stage-1 ask obtained from the prediction-market maker's break-even condition is

$$
\begin{aligned}
&A_{1}^{\mathrm{P,BE}}\left(p_{\mathrm{M}}\right)\\
=&p_{\mathrm{M}}\Pi_{\mathrm{M}}^{+}\left(p_{\mathrm{M}}\right)+\left(1-p_{\mathrm{M}}\right)\Pi_{\mathrm{L}}^{+}\left(p_{\mathrm{M}}\right)\\
=&p_{\mathrm{M}}\left[q_{+}\Pi_{\mathrm{M}}^{+}\left(p_{\mathrm{M}},+\right)+q_{-}\Pi_{\mathrm{M}}^{+}\left(p_{\mathrm{M}},-\right)\right]+\left(1-p_{\mathrm{M}}\right)\left[q_{+}\Pi_{\mathrm{L}}^{+}\left(p_{\mathrm{M}},+\right)+q_{-}\Pi_{\mathrm{L}}^{+}\left(p_{\mathrm{M}},-\right)\right]\\
=&p_{\mathrm{M}}\left[\frac{1}{2}+\frac{1}{2}H\left(A_{-}^{*}\left(p_{\mathrm{M}}\right)/\varepsilon\right)\right]+\left(1-p_{\mathrm{M}}\right)\left[\frac{3+\pi^{2}}{8}+\frac{1-\pi^{2}}{4}H\left(A_{-}^{*}\left(p_{\mathrm{M}}\right)/\varepsilon\right)\right]\\
=&\frac{p_{\mathrm{M}}}{2}+\left(1-p_{\mathrm{M}}\right)\frac{3+\pi^{2}}{8}+\left[\frac{p_{\mathrm{M}}}{2}+\frac{\left(1-p_{\mathrm{M}}\right)\left(1-\pi^{2}\right)}{4}\right]\left[\frac{1}{2}-\frac{\sigma\pi p_{\mathrm{M}}}{\varepsilon\left[\left(1-\pi^{2}\right)+\left(1+\pi^{2}\right)p_{\mathrm{M}}\right]}\right]\\
=&\frac{1}{2}-\frac{p_{\mathrm{M}}\left(\sigma\pi-\varepsilon\right)}{4\varepsilon}<\frac{1}{2}.
\end{aligned}
$$

Mixed branch $p_{\mathrm{M}} \in (p_0, 1)$**.** On this branch, $\alpha_+^* = 0$ and $\alpha_-^*(p_{\mathrm{M}}) \in (0,1)$. The per-path expected binary-contract payoffs are:

$$\Pi_{\mathrm{M}}^+(p_{\mathrm{M}}, +) = 1 \qquad \text{(same as the pure-push branch)},$$

$$\Pi_{\mathrm{M}}^+(p_{\mathrm{M}}, -) = \alpha_-^*(p_{\mathrm{M}}) H\left(A_-^*(p_{\mathrm{M}})/\varepsilon\right) + \left(1 - \alpha_-^*(p_{\mathrm{M}})\right) H\left(\bar{V}_-/\varepsilon\right) = \alpha_-^*(p_{\mathrm{M}}) H\left(A_-^*(p_{\mathrm{M}})/\varepsilon\right),$$

$$\Pi_{\mathrm{L}}^+(p_{\mathrm{M}}, +) = \frac{3+\pi^2}{4} \qquad \text{(same as the pure-push branch)},$$

$$\Pi_{\mathrm{L}}^+(p_{\mathrm{M}}, -) = q_{-+}(0) H\left(A_-^*(p_{\mathrm{M}})/\varepsilon\right) + q_{--}(0) H\left(B_-/\varepsilon\right) = \frac{1-\pi^2}{2} H\left(A_-^*(p_{\mathrm{M}})/\varepsilon\right).$$

The manipulator's expected binary-contract payoff net of expected spot manipulation cost is

$$\begin{aligned}
&\Pi_{\mathrm{M}}^+(p_{\mathrm{M}}) - K^+(p_{\mathrm{M}}) \\
=&q_+ \Pi_{\mathrm{M}}^+(p_{\mathrm{M}}, +) + q_- \Pi_{\mathrm{M}}^+(p_{\mathrm{M}}, -) - q_+ \cdot 0 - q_- \alpha_-^*(p_{\mathrm{M}}) \left[A_-^*(p_{\mathrm{M}}) - \bar{V}_-\right] \\
=&\frac{1}{2} + \frac{\alpha_-^*(p_{\mathrm{M}})}{2} \frac{\sigma\pi - \varepsilon}{1-2\varepsilon} - \frac{\alpha_-^*(p_{\mathrm{M}})}{2} \frac{\sigma\pi - \varepsilon}{1-2\varepsilon} \\
=&\frac{1}{2}.
\end{aligned}$$

The unconstrained shadow stage-1 ask obtained from the prediction-market maker's break-even condition is

$$\begin{aligned}
&A_1^{\mathrm{P,BE}}(p_{\mathrm{M}}) \\
=&p_{\mathrm{M}} \Pi_{\mathrm{M}}^+(p_{\mathrm{M}}) + (1-p_{\mathrm{M}}) \Pi_{\mathrm{L}}^+(p_{\mathrm{M}}) \\
=&p_{\mathrm{M}} \left[q_+ \Pi_{\mathrm{M}}^+(p_{\mathrm{M}}, +) + q_- \Pi_{\mathrm{M}}^+(p_{\mathrm{M}}, -)\right] + (1-p_{\mathrm{M}}) \left[q_+ \Pi_{\mathrm{L}}^+(p_{\mathrm{M}}, +) + q_- \Pi_{\mathrm{L}}^+(p_{\mathrm{M}}, -)\right] \\
=&p_{\mathrm{M}} \left[\frac{1}{2} + \frac{\alpha_-^*(p_{\mathrm{M}})}{2} \frac{\sigma\pi - \varepsilon}{1-2\varepsilon}\right] + (1-p_{\mathrm{M}}) \left[\frac{3+\pi^2}{8} + \frac{\left(1-\pi^2\right)(\sigma\pi - \varepsilon)}{4(1-2\varepsilon)}\right] \\
=&\frac{p_{\mathrm{M}}}{2} + (1-p_{\mathrm{M}}) \frac{\varepsilon\left(1-\pi^2\right)(1-2\sigma\pi)}{4(1-2\varepsilon)} + (1-p_{\mathrm{M}}) \left[\frac{3+\pi^2}{8} + \frac{\left(1-\pi^2\right)(\sigma\pi - \varepsilon)}{4(1-2\varepsilon)}\right] \\
=&\frac{1}{2} - \frac{(1-p_{\mathrm{M}})\left(1-\pi^2\right)(1-2\sigma\pi)}{8} < \frac{1}{2}.
\end{aligned}$$

Non-crossing equilibrium ask and optimal entry probability. From the above calculations, $A_1^{\mathrm{P,BE}}(p_{\mathrm{M}}) < 1/2$ for any $p_{\mathrm{M}} \in (0,1)$. Therefore, the non-crossing equilibrium stage-1 ask is

$$A_1^{\mathrm{P}*}(p_{\mathrm{M}}) = \max\left\{A_1^{\mathrm{P,BE}}(p_{\mathrm{M}}), \frac{1}{2}\right\} = \frac{1}{2}$$

for any $p_{\mathrm{M}} \in (0,1)$. The manipulator's expected profit conditional on entering at stage 1 is

$$S_{\mathrm{M}}^{+}(p_{\mathrm{M}}) = \Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}}) - K^{+}(p_{\mathrm{M}}) - \frac{1}{2},$$

which yields

$$S_{\mathrm{M}}^{+}(p_{\mathrm{M}}) = \begin{cases} \dfrac{\left(1-\pi^2\right)(1-2\sigma\pi)(1-p_{\mathrm{M}}/p_0)}{4\left[(1-\pi^2)+(1+\pi^2)p_{\mathrm{M}}\right]}, & 0 < p_{\mathrm{M}} \le p_0, \\ 0, & p_0 < p_{\mathrm{M}} < 1. \end{cases}$$

The pure-push branch is positive when $0 < p_{\mathrm{M}} < p_0$ and vanishes at $p_{\mathrm{M}} = p_0$. Therefore, the manipulator maximizes the ex-ante profit

$$\bar{S}_{\mathrm{M}}^{+}(p_{\mathrm{M}}) = p_{\mathrm{M}} S_{\mathrm{M}}^{+}(p_{\mathrm{M}}) = \frac{p_{\mathrm{M}}\left(1-\pi^2\right)(1-2\sigma\pi)(1-p_{\mathrm{M}}/p_0)}{4\left[(1-\pi^2)+(1+\pi^2)p_{\mathrm{M}}\right]} \tag{7}$$

over $0 < p_{\mathrm{M}} < p_0$. The first-order condition reduces to

$$\left(1+\pi^2\right)p_{\mathrm{M}}^2 + 2\left(1-\pi^2\right)p_{\mathrm{M}} - \left(1-\pi^2\right)p_0 = 0.$$

The unique positive root is

$$p_{\mathrm{M}}^{*} = \sqrt{\left(\frac{1-\pi^2}{1+\pi^2}\right)^2 + \frac{1-\pi^2}{1+\pi^2}p_0} - \frac{1-\pi^2}{1+\pi^2},$$

which we can verify is the unique maximizer over $(0, p_0)$. At p_{M}^*,

$$\begin{aligned}
\alpha_+^*(p_{\mathrm{M}}^*) &= 0, \\
\alpha_-^*(p_{\mathrm{M}}^*) &= 1, \\
S_{\mathrm{M}}^+(p_{\mathrm{M}}^*) &= \frac{(1-2\sigma\pi)\, p_{\mathrm{M}}^*}{4p_0}, \\
\bar{S}_{\mathrm{M}}^+(p_{\mathrm{M}}^*) &= \frac{(1-2\sigma\pi)\, p_{\mathrm{M}}^{*2}}{4p_0}.
\end{aligned}$$

A.5.3 Comparative statics derivations

We verify the signs in Corollary 2 under the equilibrium of Proposition 4.

Pure-push cutoff p_0. We can sign the partial derivatives by inspection of the closed-form expression for p_0

$$p_0 = \left(1 + \frac{2(\sigma\pi - \varepsilon)}{\varepsilon(1-\pi^2)(1-2\sigma\pi)}\right)^{-1},$$

which gives

$$\frac{\partial p_0}{\partial \varepsilon} > 0, \qquad \frac{\partial p_0}{\partial \sigma} < 0, \qquad \frac{\partial p_0}{\partial \pi} < 0.$$

Optimal entry probability p_{M}^*. For $z \in (0,1)$, $p_0 \in (0,1)$, the map

$$(z, p_0) \mapsto \sqrt{z^2 + zp_0} - z = \frac{p_0}{\sqrt{1 + p_0/z} + 1}$$

strictly increases in both z and p_0. Since

$$p_{\mathrm{M}}^* = \sqrt{\left(\frac{1-\pi^2}{1+\pi^2}\right)^2 + \frac{1-\pi^2}{1+\pi^2} p_0} - \frac{1-\pi^2}{1+\pi^2},$$

and $\left(1-\pi^2\right)/\left(1+\pi^2\right)$ strictly decreases in π, we obtain

$$\frac{\partial p_{\mathrm{M}}^*}{\partial \varepsilon} > 0, \qquad \frac{\partial p_{\mathrm{M}}^*}{\partial \sigma} < 0, \qquad \frac{\partial p_{\mathrm{M}}^*}{\partial \pi} < 0.$$

Stage-3 ask $A_-^*\left(p_{\mathrm{M}}^*\right)$. Recall that

$$A_-^*\left(p_{\mathrm{M}}^*\right) = -\frac{2\sigma\pi p_{\mathrm{M}}^*}{\left(1-\pi^2\right)+\left(1+\pi^2\right)p_{\mathrm{M}}^*} = -\frac{2\sigma\pi}{\left(1-\pi^2\right)/p_{\mathrm{M}}^*+\left(1+\pi^2\right)} < 0.$$

This quantity depends on ε only through p_{M}^* and strictly decreases in p_{M}^* holding all other terms fixed. Since $\partial p_{\mathrm{M}}^*/\partial\varepsilon > 0$, we obtain

$$\frac{\partial A_-^*\left(p_{\mathrm{M}}^*\right)}{\partial\varepsilon} < 0.$$

For σ and π, using the closed-form expression for p_{M}^*, we have

$$\left(1-\pi^2\right)+\left(1+\pi^2\right)p_{\mathrm{M}}^* = \sqrt{\left(1-\pi^2\right)^2+\left(1-\pi^2\right)\left(1+\pi^2\right)p_0}$$

and therefore

$$A_-^*\left(p_{\mathrm{M}}^*\right) = -\frac{2\sigma\pi}{1+\pi^2}\left[1-\sqrt{\frac{1-\pi^2}{\left(1-\pi^2\right)+\left(1+\pi^2\right)p_0}}\right].$$

Substituting the closed-form expression for p_0 gives

$$\begin{aligned}
A_-^*\left(p_{\mathrm{M}}^*\right) &= -\frac{2\sigma\pi}{1+\pi^2}\left[1-\sqrt{1-\frac{\varepsilon\left(1+\pi^2\right)\left(1-2\sigma\pi\right)}{2\sigma\pi\left(1-2\varepsilon\right)}}\right] \\
&= -\frac{\varepsilon\left(1-2\sigma\pi\right)}{1-2\varepsilon}\left[1+\sqrt{1-\frac{\varepsilon\left(1+\pi^2\right)\left(1-2\sigma\pi\right)}{2\sigma\pi\left(1-2\varepsilon\right)}}\right]^{-1} \\
&= -\frac{\varepsilon\left(1-2\sigma\pi\right)}{1-2\varepsilon}\left[1+\sqrt{1-\frac{\varepsilon}{2\sigma\left(1-2\varepsilon\right)}\left[\pi+\frac{1}{\pi}-2\sigma\left(1+\pi^2\right)\right]}\right]^{-1}.
\end{aligned}$$

By inspection of the last expression, the square-root term strictly increases in both σ and π, while the positive fraction multiplying the leading negative sign strictly decreases in both σ and π. Therefore,

$$\frac{\partial A_-^*\left(p_{\mathrm{M}}^*\right)}{\partial\sigma} > 0, \qquad \frac{\partial A_-^*\left(p_{\mathrm{M}}^*\right)}{\partial\pi} > 0.$$

Spot manipulation cost at p_{M}^*. Since $\bar{V}_- = -\sigma\pi$, the spot manipulation cost is

$$A_-^*\left(p_{\mathrm{M}}^*\right) - \bar{V}_- = A_-^*\left(p_{\mathrm{M}}^*\right) + \sigma\pi.$$

Therefore,

$$\begin{aligned}
\frac{\partial}{\partial \varepsilon}\left[A_-^*\left(p_{\mathrm{M}}^*\right) - \bar{V}_-\right] &= \frac{\partial A_-^*\left(p_{\mathrm{M}}^*\right)}{\partial \varepsilon} < 0, \\
\frac{\partial}{\partial \sigma}\left[A_-^*\left(p_{\mathrm{M}}^*\right) - \bar{V}_-\right] &= \frac{\partial A_-^*\left(p_{\mathrm{M}}^*\right)}{\partial \sigma} + \pi > 0, \\
\frac{\partial}{\partial \pi}\left[A_-^*\left(p_{\mathrm{M}}^*\right) - \bar{V}_-\right] &= \frac{\partial A_-^*\left(p_{\mathrm{M}}^*\right)}{\partial \pi} + \sigma > 0.
\end{aligned}$$

Ex-ante manipulation profit $\bar{S}_{\mathrm{M}}^+\left(p_{\mathrm{M}}^*\right)$. Using the first-order condition from Appendix A.5.2 evaluated at p_{M}^*, we can verify that

$$\bar{S}_{\mathrm{M}}^+\left(p_{\mathrm{M}}^*\right) = \frac{1-2\sigma\pi}{4}\frac{p_{\mathrm{M}}^{*2}}{p_0} = \frac{1-2\sigma\pi}{4}\left(\frac{1+\pi^2}{1-\pi^2} + \frac{2}{p_{\mathrm{M}}^*}\right)^{-1}.$$

Combined with the comparative statics results for p_{M}^*, we obtain by inspection of the last expression that

$$\frac{\partial \bar{S}_{\mathrm{M}}^+\left(p_{\mathrm{M}}^*\right)}{\partial \varepsilon} > 0, \qquad \frac{\partial \bar{S}_{\mathrm{M}}^+\left(p_{\mathrm{M}}^*\right)}{\partial \sigma} < 0, \qquad \frac{\partial \bar{S}_{\mathrm{M}}^+\left(p_{\mathrm{M}}^*\right)}{\partial \pi} < 0.$$

Conditional manipulation profit $S_{\mathrm{M}}^+\left(p_{\mathrm{M}}^*\right)$. Let

$$T :- \sqrt{1 - \frac{\varepsilon\left(1+\pi^2\right)\left(1-2\sigma\pi\right)}{2\sigma\pi\left(1-2\varepsilon\right)}} = \sqrt{\frac{1-\pi^2}{\left(1-\pi^2\right)+\left(1+\pi^2\right)p_0}}.$$

In the canonical regime, $T \in \left(\sqrt{\left(1-\pi^2\right)/2}, 1\right)$. Using the closed-form expression for p_{M}^*, we have

$$\frac{p_{\mathrm{M}}^*}{p_0} = \frac{1-\pi^2}{1+\pi^2}\frac{1}{p_0}\left(\sqrt{1+\frac{1+\pi^2}{1-\pi^2}p_0} - 1\right) = \left(\sqrt{1+\frac{1+\pi^2}{1-\pi^2}p_0} + 1\right)^{-1} = \frac{T}{1+T}.$$

Therefore, the conditional manipulation profit at the optimal entry probability decomposes as

$$S_{\mathrm{M}}^{+}\left(p_{\mathrm{M}}^{*}\right) = \underbrace{\frac{1-2\sigma\pi}{4}}_{\text{prize-compression channel}} \cdot \underbrace{\frac{T}{1+T}}_{\text{entry-ratio channel}}.$$

Since $\partial p_0/\partial\varepsilon > 0$, the expression for T shows that T strictly decreases in ε. Hence

$$\frac{\partial S_{\mathrm{M}}^{+}\left(p_{\mathrm{M}}^{*}\right)}{\partial\varepsilon} < 0.$$

For σ and π, the two channels move in opposite directions. The prize-compression channel $(1-2\sigma\pi)/4$ strictly decreases in both σ and π. The entry-ratio channel $T/(1+T)$ strictly increases in T, and direct differentiation of T gives

$$\begin{aligned}
\frac{\partial T}{\partial\sigma} &= \frac{1-T^2}{2\sigma T\left(1-2\sigma\pi\right)} > 0, \\
\frac{\partial T}{\partial\pi} &= \frac{\left(1-T^2\right)\left(1-\pi^2+4\sigma\pi^3\right)}{2\pi T\left(1+\pi^2\right)\left(1-2\sigma\pi\right)} > 0.
\end{aligned}$$

Thus the entry-ratio channel strictly increases in both σ and π. Putting the two channels together,

$$\begin{aligned}
\frac{\partial S_{\mathrm{M}}^{+}\left(p_{\mathrm{M}}^{*}\right)}{\partial\sigma} &= \frac{1-T-4\sigma\pi T^2}{8\sigma T\left(1+T\right)}, \\
\frac{\partial S_{\mathrm{M}}^{+}\left(p_{\mathrm{M}}^{*}\right)}{\partial\pi} &= \frac{\left(1-T\right)\left(1-\pi^2+4\sigma\pi^3\right)-4\sigma\pi\left(1+\pi^2\right)T^2}{8\pi T\left(1+T\right)\left(1+\pi^2\right)}.
\end{aligned}$$

The denominators are strictly positive. Therefore,

$$\begin{aligned}
\frac{\partial S_{\mathrm{M}}^{+}\left(p_{\mathrm{M}}^{*}\right)}{\partial\sigma} > 0 \quad &\Longleftrightarrow \quad T < T_{\sigma}^{*} :- \frac{2}{1+\sqrt{1+16\sigma\pi}}, \\
\frac{\partial S_{\mathrm{M}}^{+}\left(p_{\mathrm{M}}^{*}\right)}{\partial\pi} > 0 \quad &\Longleftrightarrow \quad T < T_{\pi}^{*} :- \frac{2}{1+\sqrt{1+\frac{16\sigma\pi\left(1+\pi^2\right)}{1-\pi^2+4\sigma\pi^3}}}.
\end{aligned}$$

Moreover, $T_\pi^* < T_\sigma^*$, because

$$1 + \pi^2 > 1 - \pi^2 + 4\sigma\pi^3 \quad \Longleftrightarrow \quad \sigma\pi < \frac{1}{2}.$$

To state the sign conditions as settlement-noise cutoffs, note that $1 - T^2$ strictly decreases in T when $T \geq 0$, so for $j \in \{\sigma, \pi\}$,

$$\frac{\partial S_{\mathrm{M}}^{+}\left(p_{\mathrm{M}}^*\right)}{\partial j} > 0 \quad \Longleftrightarrow \quad T < T_j^* \quad \Longleftrightarrow \quad 1 - T^2 > 1 - T_j^{*2}.$$

Since

$$1 - T^2 = \frac{\varepsilon\left(1 + \pi^2\right)(1 - 2\sigma\pi)}{2\sigma\pi\,(1 - 2\varepsilon)}$$

strictly increases in ε, the last condition is equivalent to

$$\varepsilon > \varepsilon_j^* :- \frac{2\sigma\pi\left(1 - T_j^{*2}\right)}{\left(1 + \pi^2\right)(1 - 2\sigma\pi) + 4\sigma\pi\left(1 - T_j^{*2}\right)},$$

provided $\varepsilon_j^* < \sigma\pi$. If $\varepsilon_j^* \geq \sigma\pi$, the partial derivative remains negative throughout the canonical regime. Since $T_\pi^* < T_\sigma^*$, we have $\varepsilon_\pi^* > \varepsilon_\sigma^*$. That is, a positive partial derivative with respect to π, when it exists, requires higher settlement noise than one with respect to σ.

When ε is small, both $\partial S_{\mathrm{M}}^{+}/\partial\sigma$ and $\partial S_{\mathrm{M}}^{+}/\partial\pi$ are negative: at low settlement noise the entry-ratio channel $T/(1 + T)$ is already near its ceiling of $1/2$, so increasing σ or π mainly compresses the prize and conditional profit falls.

A.5.4 Liquidity-trader loss decomposition

We prove Proposition 5. At $p_{\mathrm{M}}^* \in (0, p_0)$, the selected continuation equilibrium after a stage-2 sell is on the pure-push branch. From Appendix A.5.2, the shadow stage-1 ask is

$$A_1^{\mathrm{P,BE}}\left(p_{\mathrm{M}}^*\right) = \frac{1}{2} - \frac{p_{\mathrm{M}}^*\,(\sigma\pi - \varepsilon)}{4\varepsilon}.$$

Since the posted non-crossing ask is $A_1^{\mathrm{P}*} = 1/2$, the prediction-market maker's expected rent per buy order is

$$R_{\mathrm{P}}^{+*} :- A_1^{\mathrm{P}*} - A_1^{\mathrm{P,BE}}\left(p_{\mathrm{M}}^*\right) = \frac{p_{\mathrm{M}}^*\left(\sigma\pi - \varepsilon\right)}{4\varepsilon} > 0.$$

The liquidity trader's payoff on the pure-push branch is

$$\Pi_{\mathrm{L}}^{+}\left(p_{\mathrm{M}}^*\right) = \frac{3+\pi^2}{8} + \frac{1-\pi^2}{4} H\left(A_{-}^{*}\left(p_{\mathrm{M}}^*\right)/\varepsilon\right),$$

with

$$H\left(A_{-}^{*}\left(p_{\mathrm{M}}^*\right)/\varepsilon\right) = \frac{1}{2} - \frac{\sigma\pi p_{\mathrm{M}}^*}{\varepsilon\left[\left(1-\pi^2\right) + \left(1+\pi^2\right) p_{\mathrm{M}}^*\right]}.$$

Therefore the expected loss per buy order for the liquidity trader is

$$L^{+}\left(p_{\mathrm{M}}^*\right) :- \frac{1}{2} - \Pi_{\mathrm{L}}^{+}\left(p_{\mathrm{M}}^*\right) = \frac{\left(1-\pi^2\right)\sigma\pi p_{\mathrm{M}}^*}{4\varepsilon\left[\left(1-\pi^2\right) + \left(1+\pi^2\right) p_{\mathrm{M}}^*\right]} > 0.$$

Finally, the accounting identity follows directly from the definitions:

$$\begin{aligned}
&R_{\mathrm{P}}^{+*} + p_{\mathrm{M}}^* S_{\mathrm{M}}^{+}\left(p_{\mathrm{M}}^*\right) + p_{\mathrm{M}}^* K^{+}\left(p_{\mathrm{M}}^*\right) \\
=&\frac{1}{2} - \left[p_{\mathrm{M}}^* \Pi_{\mathrm{M}}^{+}\left(p_{\mathrm{M}}^*\right) + \left(1-p_{\mathrm{M}}^*\right)\Pi_{\mathrm{L}}^{+}\left(p_{\mathrm{M}}^*\right)\right] + p_{\mathrm{M}}^*\left[\Pi_{\mathrm{M}}^{+}\left(p_{\mathrm{M}}^*\right) - K^{+}\left(p_{\mathrm{M}}^*\right) - \frac{1}{2}\right] + p_{\mathrm{M}}^* K^{+}\left(p_{\mathrm{M}}^*\right) \\
=&\left(1-p_{\mathrm{M}}^*\right)\left[\frac{1}{2} - \Pi_{\mathrm{L}}^{+}\left(p_{\mathrm{M}}^*\right)\right] \\
=&\left(1-p_{\mathrm{M}}^*\right) L^{+}\left(p_{\mathrm{M}}^*\right).
\end{aligned}$$

A.6 Two-Ordinary-Round Extension Calculations

This appendix derives the two-ordinary-round extension used in Section 5.4. The analysis is conditional on a stage-1 prediction-market buy; the sell branch is symmetric.

A.6.1 Stage 2 Posteriors

Stage-2 price path probabilities. Stages 2.1 and 2.2 are independent ordinary Glosten–Milgrom spot rounds except that the prior over V entering stage 2.2 is the posterior from stage 2.1. At each round, the informed trader enters with probability π. She buys if $V = +\sigma$ and sells if $V = -\sigma$ conditional on entry. A liquidity trader enters with probability $1 - \pi$. She buys or sells with equal probability conditional on entry. Therefore, for $j \in \{2.1, 2.2\}$,

$$\mathbb{P}\left(Y_j^{\mathrm{S}} = + \mid V = \pm\sigma\right) = \frac{1 \pm \pi}{2}, \qquad \mathbb{P}\left(Y_j^{\mathrm{S}} = - \mid V = \pm\sigma\right) = \frac{1 \mp \pi}{2}.$$

Conditional on V, the two ordinary spot orders are independent. Hence, for stage-2 price path $h_2 = h_{2.1}h_{2.2}$,

$$\mathbb{P}(h_2 = ++ \mid V = +\sigma) = \left(\frac{1+\pi}{2}\right)^2, \qquad \mathbb{P}(h_2 = ++ \mid V = -\sigma) = \left(\frac{1-\pi}{2}\right)^2,$$

$$\mathbb{P}(h_2 = -- \mid V = +\sigma) = \left(\frac{1-\pi}{2}\right)^2, \qquad \mathbb{P}(h_2 = -- \mid V = -\sigma) = \left(\frac{1+\pi}{2}\right)^2,$$

$$\mathbb{P}(h_2 = +- \mid V = +\sigma) = \mathbb{P}(h_2 = +- \mid V = -\sigma) = \frac{1-\pi^2}{4},$$

$$\mathbb{P}(h_2 = -+ \mid V = +\sigma) = \mathbb{P}(h_2 = -+ \mid V = -\sigma) = \frac{1-\pi^2}{4}.$$

Marginalizing over V gives

$$q_{++} = q_{--} = \frac{1+\pi^2}{4}, \qquad q_{+-} = q_{-+} = \frac{1-\pi^2}{4}.$$

Posteriors. Bayes' rule gives

$$\rho_{++} = \mathbb{P}(V = +\sigma \mid h_2 = ++) = \frac{\mathbb{P}(h_2 = ++ \mid V = +\sigma)\,\mathbb{P}(V = +\sigma)}{q_{++}} = \frac{\left(\frac{1+\pi}{2}\right)^2 \cdot \frac{1}{2}}{\frac{1+\pi^2}{4}} = \frac{(1+\pi)^2}{2\,(1+\pi^2)}$$

and similarly

$$\rho_{+-} = \rho_{-+} = \frac{1}{2}, \qquad \rho_{--} = \frac{(1-\pi)^2}{2(1+\pi^2)}.$$

Using $\bar{V}_{h_2} = \sigma(2\rho_{h_2} - 1)$ gives

$$\bar{V}_{++} = \frac{2\sigma\pi}{1+\pi^2}, \qquad \bar{V}_{+-} = \bar{V}_{-+} = 0, \qquad \bar{V}_{--} = -\frac{2\sigma\pi}{1+\pi^2}.$$

Posterior variance. Direct calculation of $\nu_{h_2} = 4\sigma^2\rho_{h_2}(1-\rho_{h_2})$ gives

$$\nu_{++} = \nu_{--} = \frac{\sigma^2\left(1-\pi^2\right)^2}{\left(1+\pi^2\right)^2}, \qquad \nu_{+-} = \nu_{-+} = \sigma^2.$$

A.6.2 Selected stage-3 continuation equilibrium

Recall that the gap between the expected continuation payoff from buying and that from staying out after stage-2 price path h_2 is

$$\Gamma_{h_2}(\alpha_{h_2}; p_{\mathrm{M}}) = H\left(\frac{A_{h_2}(\alpha_{h_2}; p_{\mathrm{M}})}{\varepsilon}\right) - H\left(\frac{\bar{V}_{h_2}}{\varepsilon}\right) - \left[A_{h_2}(\alpha_{h_2}; p_{\mathrm{M}}) - \bar{V}_{h_2}\right].$$

We find the selected continuation equilibrium at each stage-2 price path by the same three-step test of Appendix A.5.1. Throughout, stage-3 asks are obtained by substituting the relevant stage-2 posterior from Appendix A.6.1 into the general ask formula (4).

Favorable agreeing price path $h_2 = ++$: pure stay-out equilibrium. For any $p_{\mathrm{M}} \in (0,1)$ and any $\alpha_{++} \in [0,1]$, the stage-3 ask is at least the prior mean, so

$$A_{++}(\alpha_{++}; p_{\mathrm{M}}) \geq \bar{V}_{++} = \frac{2\sigma\pi}{1+\pi^2} > \sigma\pi > \varepsilon.$$

Therefore $H\left(A_{++}(\alpha_{++}; p_{\mathrm{M}})/\varepsilon\right) = H\left(\bar{V}_{++}/\varepsilon\right) = 1$. Hence

$$\Gamma_{++}(\alpha_{++}; p_{\mathrm{M}}) = 1 - 1 - \left[A_{++}(\alpha_{++}; p_{\mathrm{M}}) - \bar{V}_{++}\right] \leq 0.$$

Thus, for any $p_{\mathrm{M}} \in (0,1)$, $\alpha^*_{++} = 0$ is a stable fixed point and is clearly the smallest.

Unfavorable agreeing price path $h_2 = --$: pure stay-out equilibrium. For any $p_{\mathrm{M}} \in (0,1)$,

$$\bar{V}_{--} = -\frac{2\sigma\pi}{1+\pi^2} < -\sigma\pi < -\varepsilon.$$

At $\alpha_{--} = 0$, substituting $\rho_{--} = (1-\pi)^2 / \left(2\left(1+\pi^2\right)\right)$ into (4) gives

$$A_{--}(0; p_{\mathrm{M}}) = -\sigma\pi < -\varepsilon.$$

Since the ask strictly decreases in α_{--} and is at least the prior mean,

$$\bar{V}_{--} \leq A_{--}(\alpha_{--}; p_{\mathrm{M}}) \leq -\sigma\pi < -\varepsilon.$$

Consequently $H\left(A_{--}(\alpha_{--}; p_{\mathrm{M}})/\varepsilon\right) = H\left(\bar{V}_{--}/\varepsilon\right) = 0$, and

$$\Gamma_{--}(\alpha_{--}; p_{\mathrm{M}}) = 0 - 0 - \left[A_{--}(\alpha_{--}; p_{\mathrm{M}}) - \bar{V}_{--}\right] \leq 0.$$

Thus, for any $p_{\mathrm{M}} \in (0,1)$, $\alpha^*_{--} = 0$ is a stable fixed point and is clearly the smallest.

Disagreeing price paths $h_2 \in \{+-, -+\}$: pure push equilibrium. Fix $p_{\mathrm{M}} \in (0,1)$. For either disagreeing price path, $\bar{V}_{h_2} = 0$. At $\alpha_{h_2} = 0$, substituting $\rho_{h_2} = 1/2$ into (4) gives $A_{h_2}(0; p_{\mathrm{M}}) = \sigma\pi$, and therefore

$$\Gamma_{h_2}(0; p_{\mathrm{M}}) = 1 - \frac{1}{2} - [\sigma\pi - 0] = \frac{1}{2} - \sigma\pi > 0,$$

so $\alpha_{h_2} = 0$ is not a fixed point. For any $\alpha_{h_2} \in [0,1]$, the ask is

$$A_{h_2}(\alpha_{h_2}; p_{\mathrm{M}}) = \frac{\sigma\pi(1-p_{\mathrm{M}})}{1-p_{\mathrm{M}}+2p_{\mathrm{M}}\alpha_{h_2}} \in \left[\frac{\sigma\pi(1-p_{\mathrm{M}})}{1+p_{\mathrm{M}}}, \sigma\pi\right] \subset \left(0, \frac{1}{2}\right).$$

If $A_{h_2}(\alpha_{h_2}; p_{\mathrm{M}}) \geq \varepsilon$, then

$$\Gamma_{h_2}(\alpha_{h_2}; p_{\mathrm{M}}) = \frac{1}{2} - A_{h_2}(\alpha_{h_2}; p_{\mathrm{M}}) \geq \frac{1}{2} - \sigma\pi > 0.$$

If $0 < A_{h_2}(\alpha_{h_2}; p_{\mathrm{M}}) < \varepsilon$, then

$$\Gamma_{h_2}(\alpha_{h_2}; p_{\mathrm{M}}) = A_{h_2}(\alpha_{h_2}; p_{\mathrm{M}})\left(\frac{1}{2\varepsilon} - 1\right) > 0$$

because $\varepsilon < 1/2$. Thus no interior fixed point exists either, and the selected continuation equilibrium is the endpoint

$$\alpha^*_{+-} = \alpha^*_{-+} = 1.$$

A.6.3 Optimal manipulator entry probability

Both the pricing of the binary contract and the manipulator's stage-1 entry problem use the selected continuation equilibrium from Appendix A.6.2. Combining Appendix A.6.1, the general ask and bid formulas (4) and (5), and Appendix A.6.2, the stage-3 asks, bids, and no-trade posterior means used below are

$$A^*_{++}(p_{\mathrm{M}}) = \frac{\sigma\pi\left(3+\pi^2\right)}{1+3\pi^2}, \qquad A^*_{--}(p_{\mathrm{M}}) = -\sigma\pi, \qquad A^*_{+-}(p_{\mathrm{M}}) = A^*_{-+}(p_{\mathrm{M}}) = \frac{\sigma\pi(1-p_{\mathrm{M}})}{1+p_{\mathrm{M}}}.$$

$$B_{++} = \sigma\pi, \qquad B_{--} = -\frac{\sigma\pi\left(3+\pi^2\right)}{1+3\pi^2}, \qquad B_{+-} = B_{-+} = -\sigma\pi.$$

$$\bar{V}_{++} = \frac{2\sigma\pi}{1+\pi^2}, \qquad \bar{V}_{+-} = \bar{V}_{-+} = 0, \qquad \bar{V}_{--} = -\frac{2\sigma\pi}{1+\pi^2}.$$

For each stage-2 price path h_2, the per-path expected binary-contract payoffs are

$$\Pi^+_{\mathrm{M}}(p_{\mathrm{M}}, h_2) = \alpha^*_{h_2} H\left(A^*_{h_2}(p_{\mathrm{M}})/\varepsilon\right) + \left(1-\alpha^*_{h_2}\right) H\left(\bar{V}_{h_2}/\varepsilon\right),$$
$$\Pi^+_{\mathrm{L}}(p_{\mathrm{M}}, h_2) = q_{h_2+}(0)\, H\left(A^*_{h_2}(p_{\mathrm{M}})/\varepsilon\right) + q_{h_2-}(0)\, H\left(B_{h_2}/\varepsilon\right).$$

Substituting the selected continuation equilibrium and the stage-3 quotes gives

$$\begin{aligned}
&\Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}},++)=1, \qquad \Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}},--)=0,\\
&\Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}},+-)=\Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}},-+)=H\left(\frac{\sigma\pi(1-p_{\mathrm{M}})}{\varepsilon(1+p_{\mathrm{M}})}\right),\\
&\Pi_{\mathrm{L}}^{+}(p_{\mathrm{M}},++)=1, \qquad \Pi_{\mathrm{L}}^{+}(p_{\mathrm{M}},--)=0,\\
&\Pi_{\mathrm{L}}^{+}(p_{\mathrm{M}},+-)=\Pi_{\mathrm{L}}^{+}(p_{\mathrm{M}},-+)=\frac{1}{2}H\left(\frac{\sigma\pi(1-p_{\mathrm{M}})}{\varepsilon(1+p_{\mathrm{M}})}\right).
\end{aligned}$$

Aggregating over the four stage-2 price paths gives

$$\begin{aligned}
\Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}}) &= \frac{1+\pi^2}{4}+\frac{1-\pi^2}{2}H\left(\frac{\sigma\pi(1-p_{\mathrm{M}})}{\varepsilon(1+p_{\mathrm{M}})}\right),\\
\Pi_{\mathrm{L}}^{+}(p_{\mathrm{M}}) &= \frac{1+\pi^2}{4}+\frac{1-\pi^2}{4}H\left(\frac{\sigma\pi(1-p_{\mathrm{M}})}{\varepsilon(1+p_{\mathrm{M}})}\right),\\
K^{+}(p_{\mathrm{M}}) &= \frac{\left(1-\pi^2\right)\sigma\pi(1-p_{\mathrm{M}})}{2(1+p_{\mathrm{M}})}.
\end{aligned}$$

The unconstrained shadow stage-1 ask obtained from the prediction-market maker's break-even condition is

$$\begin{aligned}
&A_1^{\mathrm{P,BE}}(p_{\mathrm{M}})\\
=&p_{\mathrm{M}}\Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}})+(1-p_{\mathrm{M}})\Pi_{\mathrm{L}}^{+}(p_{\mathrm{M}})\\
=&\frac{1+\pi^2}{4}+\frac{\left(1-\pi^2\right)(1+p_{\mathrm{M}})}{4}H\left(\frac{\sigma\pi(1-p_{\mathrm{M}})}{\varepsilon(1+p_{\mathrm{M}})}\right)\\
=&\begin{cases}
\dfrac{1}{2}+\dfrac{\left(1-\pi^2\right)p_{\mathrm{M}}}{4}, & \dfrac{\sigma\pi(1-p_{\mathrm{M}})}{1+p_{\mathrm{M}}}\geq\varepsilon,\\[2ex]
\dfrac{1}{2}+\dfrac{\left(1-\pi^2\right)(1-p_{\mathrm{M}})(\sigma\pi-\varepsilon)}{8\varepsilon}, & \dfrac{\sigma\pi(1-p_{\mathrm{M}})}{1+p_{\mathrm{M}}}<\varepsilon,
\end{cases}
\end{aligned}$$

which is strictly greater than $1/2$ for any $p_{\mathrm{M}}\in(0,1)$. Therefore the non-crossing equilibrium stage-1

ask is just the break-even stage-1 ask:

$$A_1^{\mathrm{P}*}(p_{\mathrm{M}}) = \max\left\{A_1^{\mathrm{P,BE}}(p_{\mathrm{M}}), \frac{1}{2}\right\} = A_1^{\mathrm{P,BE}}(p_{\mathrm{M}})$$

for any $p_{\mathrm{M}} \in (0,1)$, and

$$B_1^{\mathrm{P}*}(p_{\mathrm{M}}) = 1 - A_1^{\mathrm{P}*}(p_{\mathrm{M}}).$$

The manipulator's expected profit conditional on entering at stage 1 is

$$\begin{aligned}
&S_{\mathrm{M}}^{+}(p_{\mathrm{M}}) \\
=&\Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}}) - K^{+}(p_{\mathrm{M}}) - A_1^{\mathrm{P}*}(p_{\mathrm{M}}), \\
=&\frac{1-\pi^2}{4}\left[(1-p_{\mathrm{M}})H\left(\frac{\sigma\pi(1-p_{\mathrm{M}})}{\varepsilon(1+p_{\mathrm{M}})}\right) - \frac{2\sigma\pi(1-p_{\mathrm{M}})}{1+p_{\mathrm{M}}}\right] \\
=&\frac{1-\pi^2}{4}\begin{cases}
\dfrac{(1-p_{\mathrm{M}})(1+p_{\mathrm{M}}-2\sigma\pi)}{1+p_{\mathrm{M}}}, & 0 < p_{\mathrm{M}} \le \dfrac{\sigma\pi-\varepsilon}{\sigma\pi+\varepsilon}, \\
\dfrac{(1-p_{\mathrm{M}})[\varepsilon(1+p_{\mathrm{M}})+\sigma\pi(1-p_{\mathrm{M}})-4\varepsilon\sigma\pi]}{2\varepsilon(1+p_{\mathrm{M}})}, & \dfrac{\sigma\pi-\varepsilon}{\sigma\pi+\varepsilon} < p_{\mathrm{M}} < 1.
\end{cases}
\end{aligned}$$

Therefore, the manipulator maximizes the ex-ante profit

$$\bar{S}_{\mathrm{M}}^{+}(p_{\mathrm{M}}) = p_{\mathrm{M}} S_{\mathrm{M}}^{+}(p_{\mathrm{M}}) = \frac{1-\pi^2}{4}\Phi(p_{\mathrm{M}};\sigma,\pi,\varepsilon),$$

where

$$\Phi(p_{\mathrm{M}};\sigma,\pi,\varepsilon) :- p_{\mathrm{M}}\left[(1-p_{\mathrm{M}})H\left(\frac{\sigma\pi(1-p_{\mathrm{M}})}{\varepsilon(1+p_{\mathrm{M}})}\right) - \frac{2\sigma\pi(1-p_{\mathrm{M}})}{1+p_{\mathrm{M}}}\right]. \tag{8}$$

Equivalently,

$$\Phi(p_{\mathrm{M}};\sigma,\pi,\varepsilon) = \begin{cases}
\dfrac{p_{\mathrm{M}}(1-p_{\mathrm{M}})(1+p_{\mathrm{M}}-2\sigma\pi)}{1+p_{\mathrm{M}}}, & 0 < p_{\mathrm{M}} \le \dfrac{\sigma\pi-\varepsilon}{\sigma\pi+\varepsilon}, \\
\dfrac{p_{\mathrm{M}}(1-p_{\mathrm{M}})[\varepsilon(1+p_{\mathrm{M}})+\sigma\pi(1-p_{\mathrm{M}})-4\varepsilon\sigma\pi]}{2\varepsilon(1+p_{\mathrm{M}})}, & \dfrac{\sigma\pi-\varepsilon}{\sigma\pi+\varepsilon} < p_{\mathrm{M}} < 1.
\end{cases}$$

The two branches agree at the cutoff $(\sigma\pi - \varepsilon) / (\sigma\pi + \varepsilon)$, so Φ is continuous on $(0, 1)$. The first branch is strictly positive because $1 + p_{\mathrm{M}} - 2\sigma\pi > 1 - 2\sigma\pi > 0$. On the second branch, the bracketed factor has slope $\varepsilon - \sigma\pi < 0$ in p_{M}, so it is bounded below by its value at $p_{\mathrm{M}} = 1$, namely $2\varepsilon(1 - 2\sigma\pi) > 0$. Therefore, Φ is strictly positive on $(0, 1)$. If we define the boundary values

$$\Phi(0; \sigma, \pi, \varepsilon) = \lim_{p_{\mathrm{M}} \downarrow 0^{+}} \Phi(p_{\mathrm{M}}; \sigma, \pi, \varepsilon) = 0, \qquad \Phi(1; \sigma, \pi, \varepsilon) = \lim_{p_{\mathrm{M}} \uparrow 1^{-}} \Phi(p_{\mathrm{M}}; \sigma, \pi, \varepsilon) = 0,$$

consistent with the convention that $\bar{S}^{+}_{\mathrm{M}}(0) = \bar{S}^{+}_{\mathrm{M}}(1) = 0$, then Φ is continuous on $[0, 1]$. Therefore, a maximizer exists and must be an interior point in $(0, 1)$. To show that it is unique, we can equivalently work with $\ln \Phi$, whose branchwise first-order derivative is

$$\frac{\partial \ln \Phi}{\partial p_{\mathrm{M}}} = \begin{cases} \dfrac{1}{p_{\mathrm{M}}} - \dfrac{1}{1 - p_{\mathrm{M}}} + \dfrac{1}{1 + p_{\mathrm{M}} - 2\sigma\pi} - \dfrac{1}{1 + p_{\mathrm{M}}}, & 0 < p_{\mathrm{M}} < \dfrac{\sigma\pi - \varepsilon}{\sigma\pi + \varepsilon}, \\ \dfrac{1}{p_{\mathrm{M}}} - \dfrac{1}{1 - p_{\mathrm{M}}} + \dfrac{\varepsilon - \sigma\pi}{\varepsilon(1 + p_{\mathrm{M}}) + \sigma\pi(1 - p_{\mathrm{M}}) - 4\varepsilon\sigma\pi} - \dfrac{1}{1 + p_{\mathrm{M}}}, & \dfrac{\sigma\pi - \varepsilon}{\sigma\pi + \varepsilon} < p_{\mathrm{M}} < 1. \end{cases}$$

The branchwise second-order derivative is

$$\frac{\partial^2 \ln \Phi}{\partial p_{\mathrm{M}}^2} = \begin{cases} -\dfrac{1}{p_{\mathrm{M}}^2} - \dfrac{1}{(1 - p_{\mathrm{M}})^2} - \dfrac{1}{(1 + p_{\mathrm{M}} - 2\sigma\pi)^2} + \dfrac{1}{(1 + p_{\mathrm{M}})^2} < 0, & 0 < p_{\mathrm{M}} < \dfrac{\sigma\pi - \varepsilon}{\sigma\pi + \varepsilon}, \\ -\dfrac{1}{p_{\mathrm{M}}^2} - \dfrac{1}{(1 - p_{\mathrm{M}})^2} - \dfrac{(\sigma\pi - \varepsilon)^2}{\left[\varepsilon + \sigma\pi - 4\varepsilon\sigma\pi - (\sigma\pi - \varepsilon) p_{\mathrm{M}}\right]^2} & \dfrac{\sigma\pi - \varepsilon}{\sigma\pi + \varepsilon} < p_{\mathrm{M}} < 1. \\ + \dfrac{1}{(1 + p_{\mathrm{M}})^2} < 0, & \end{cases}$$

On the first branch, the inequality follows because $0 < 1 + p_{\mathrm{M}} - 2\sigma\pi < 1 + p_{\mathrm{M}}$. On the second branch, the inequality follows because $0 < 1 - p_{\mathrm{M}} < 1 + p_{\mathrm{M}}$. Thus $\partial \ln \Phi / \partial p_{\mathrm{M}}$ is strictly decreasing on both branches. At $(\sigma\pi - \varepsilon) / (\sigma\pi + \varepsilon)$, the left derivative exceeds the right derivative:

$$\lim_{p_{\mathrm{M}} \uparrow \left(\frac{\sigma\pi - \varepsilon}{\sigma\pi + \varepsilon}\right)_{-}} \frac{\partial \ln \Phi}{\partial p_{\mathrm{M}}} - \lim_{p_{\mathrm{M}} \downarrow \left(\frac{\sigma\pi - \varepsilon}{\sigma\pi + \varepsilon}\right)_{+}} \frac{\partial \ln \Phi}{\partial p_{\mathrm{M}}} = \frac{(\sigma\pi + \varepsilon)^2}{4\varepsilon\sigma\pi(1 - \sigma\pi - \varepsilon)} > 0.$$

Therefore, $\ln \Phi$ (and equivalently Φ) has a unique maximizer on $(0, 1)$. It is either the unique

interior zero of the first-order derivative or the kink $(\sigma\pi - \varepsilon) / (\sigma\pi + \varepsilon)$. Let $p_{\mathrm{M}}^{*(2)}$ denote this unique maximizer. It remains to show that it is strictly greater than the corresponding maximizer in the baseline model, which we denote by $p_{\mathrm{M}}^{*(1)}$.

Bounding the closed-form expression for $p_{\mathrm{M}}^{*(1)}$ in Proposition 4, we have

$$0 < p_{\mathrm{M}}^{*(1)} = \sqrt{\left(\frac{1-\pi^2}{1+\pi^2}\right)^2 + \frac{1-\pi^2}{1+\pi^2}p_0} - \frac{1-\pi^2}{1+\pi^2} < \sqrt{1 + \frac{\varepsilon(1-2\sigma\pi)}{2(\sigma\pi-\varepsilon)+\varepsilon(1-2\sigma\pi)}} - 1 < \sqrt{2} - 1,$$

where we used the fact that $p_{\mathrm{M}}^{*(1)}$ strictly increases in both $\left(1-\pi^2\right) / \left(1+\pi^2\right)$ and p_0 as established in Appendix A.5.3, $0 < \left(1-\pi^2\right) / \left(1+\pi^2\right) < 1$, and

$$0 < p_0 < \frac{\varepsilon(1-2\sigma\pi)}{2(\sigma\pi-\varepsilon)+\varepsilon(1-2\sigma\pi)} < 1.$$

Because the first-order derivative of $\ln\Phi$ strictly decreases on each branch and jumps downward at $(\sigma\pi - \varepsilon) / (\sigma\pi + \varepsilon)$, it suffices to show that it is still positive at the sharper upper bound on $p_{\mathrm{M}}^{*(1)}$ displayed above. If that upper bound lies weakly below $(\sigma\pi - \varepsilon) / (\sigma\pi + \varepsilon)$, only the first branch is relevant. It is strictly decreasing on $(0, 1)$ and equals $\left(\sqrt{2} - 2\sigma\pi\right)^{-1} > 0$ at $p_{\mathrm{M}} = \sqrt{2} - 1$, so it must be positive at the sharper upper bound. If instead that upper bound lies strictly above $(\sigma\pi - \varepsilon) / (\sigma\pi + \varepsilon)$, it suffices to check the second branch at that value. At the sharper upper-bound value,

$$p_{\mathrm{M}}^2 + 2p_{\mathrm{M}} = \frac{\varepsilon(1-2\sigma\pi)}{2(\sigma\pi-\varepsilon)+\varepsilon(1-2\sigma\pi)} \iff \varepsilon = \frac{2\sigma\pi\left(p_{\mathrm{M}}^2 + 2p_{\mathrm{M}}\right)}{(1+2\sigma\pi)\left(p_{\mathrm{M}}^2 + 2p_{\mathrm{M}}\right) + 1 - 2\sigma\pi}.$$

Substituting this expression into the third term in the second branch gives

$$\frac{\varepsilon - \sigma\pi}{\varepsilon(1+p_{\mathrm{M}}) + \sigma\pi(1-p_{\mathrm{M}}) - 4\varepsilon\sigma\pi} = \frac{p_{\mathrm{M}}^2 + 2p_{\mathrm{M}} - 1}{(1+p_{\mathrm{M}})\left(p_{\mathrm{M}}^2 + 4p_{\mathrm{M}} + 1\right)}.$$

The remaining terms in the second branch satisfy

$$\frac{1}{p_{\mathrm{M}}}-\frac{1}{1-p_{\mathrm{M}}}-\frac{1}{1+p_{\mathrm{M}}}=\frac{-p_{\mathrm{M}}^{2}-2p_{\mathrm{M}}+1}{p_{\mathrm{M}}\left(1-p_{\mathrm{M}}\right)\left(1+p_{\mathrm{M}}\right)}.$$

Therefore the second branch at the sharper upper-bound value equals

$$\frac{p_{\mathrm{M}}^{2}+2p_{\mathrm{M}}-1}{1+p_{\mathrm{M}}}\left[\frac{1}{p_{\mathrm{M}}^{2}+4p_{\mathrm{M}}+1}-\frac{1}{p_{\mathrm{M}}\left(1-p_{\mathrm{M}}\right)}\right]=-\frac{\left(2p_{\mathrm{M}}+1\right)\left(p_{\mathrm{M}}^{2}+2p_{\mathrm{M}}-1\right)}{p_{\mathrm{M}}\left(1-p_{\mathrm{M}}\right)\left(p_{\mathrm{M}}^{2}+4p_{\mathrm{M}}+1\right)}>0,$$

with p_{M} evaluated at the sharper upper-bound value. The expression is positive because this sharper upper bound lies in $\left(0,\ \sqrt{2}-1\right)$, so $p_{\mathrm{M}}^{2}+2p_{\mathrm{M}}-1<0$ whereas all other factors are positive. Therefore, $\bar{S}_{\mathrm{M}}^{+}\left(p_{\mathrm{M}}\right)$ (under the two-ordinary-round extension) is still increasing at $p_{\mathrm{M}}^{*(1)}$, which implies

$$p_{\mathrm{M}}^{*(2)}>p_{\mathrm{M}}^{*(1)}.$$

A.6.4 Price discovery

We prove the price-discovery part of Proposition 6. By the law of total variance, for $n=1,2$,

$$\mathbb{E}\left[\nu_{h_2h_3}^{(n)}\right]\left(p_{\mathrm{M}}\right)=\nu_{\emptyset}-\mathrm{Var}\left(\bar{V}_{h_2h_3}^{(n)}\right)\left(p_{\mathrm{M}}\right)=\sigma^{2}-\mathbb{E}\left[\left(\bar{V}_{h_2h_3}^{(n)}\left(p_{\mathrm{M}}\right)\right)^{2}\right],$$

because $\mathbb{E}\left[\bar{V}_{h_2h_3}^{(n)}\left(p_{\mathrm{M}}\right)\right]=\mathbb{E}\left[V\right]=0$ by the law of iterated expectations. As in Appendix A.4, lower residual variance is equivalent to greater dispersion of the final posterior mean.

We first compute the no-manipulation benchmarks. In the baseline model ($n=1$), when $p_{\mathrm{M}}=0$, stages 2 and 3 consist of two ordinary spot rounds and thus share the same event structure as stage 2 under the two-ordinary-round extension, which allows us to reuse the results from Appendix A.6.1. The two agreeing price paths ($h_2=h_3=+$ or $-$) each have probability $\left(1+\pi^{2}\right)/4$ and posterior means $\pm 2\sigma\pi/\left(1+\pi^{2}\right)$, while the two disagreeing price paths ($h_2=+, h_3=-$ or $h_2=-, h_3=+$)

have final posterior mean zero. Hence

$$\mathbb{E}\left[\nu^{(1)}_{h_2h_3}\right](0) = \sigma^2 - 2 \cdot \frac{1+\pi^2}{4}\left(\frac{2\sigma\pi}{1+\pi^2}\right)^2 = \frac{\sigma^2\left(1-\pi^2\right)}{1+\pi^2}.$$

Under the two-ordinary-round extension ($n = 2$), when $p_{\mathrm{M}} = 0$, stages 2.1, 2.2, and 3 are three consecutive ordinary spot rounds. We can calculate the price path probabilities and posteriors following the same approach as in Appendix A.6 with an additional ordinary spot round. The all-buy and all-sell price paths each have probability $\left(1+3\pi^2\right)/8$ and posterior means $\pm\sigma\pi\left(3+\pi^2\right)/\left(1+3\pi^2\right)$. The three ordered price paths with two buys and one sell have posterior mean $\sigma\pi$, the three ordered price paths with one buy and two sells have posterior mean $-\sigma\pi$, and each of these six mixed ordered price paths has probability $\left(1-\pi^2\right)/8$. Therefore

$$\mathbb{E}\left[\nu^{(2)}_{h_2h_3}\right](0) = \sigma^2 - 2 \cdot \frac{1+3\pi^2}{8}\left(\frac{\sigma\pi\left(3+\pi^2\right)}{1+3\pi^2}\right)^2 - 6 \cdot \frac{1-\pi^2}{8}(\sigma\pi)^2 = \frac{\sigma^2\left(1-\pi^2\right)^2\left(1+2\pi^2\right)}{1+3\pi^2}.$$

Hence the no-manipulation gain from adding the extra ordinary round is

$$\mathbb{E}\left[\nu^{(1)}_{h_2h_3}\right](0) - \mathbb{E}\left[\nu^{(2)}_{h_2h_3}\right](0) = \frac{\sigma^2\pi^2\left(1-\pi^2\right)\left(1+\pi^2+2\pi^4\right)}{\left(1+\pi^2\right)\left(1+3\pi^2\right)} > 0.$$

We next compute the increase in $\mathbb{E}\left[\nu^{(2)}_{h_2h_3}\right]$ caused by the selected manipulation continuation for a fixed $p_{\mathrm{M}} \in (0,1)$. Applying the law of total variance to the dispersion of the final posterior mean as in Appendix A.4, we have

$$\mathrm{Var}\left(\bar{V}^{(2)}_{h_2h_3}\right)(p_{\mathrm{M}}) = \mathrm{Var}\left(\bar{V}^{(2)}_{h_2}\right) + \sum_{h_2} q_{h_2}\,\mathrm{Var}\left(\bar{V}^{(2)}_{h_2h_3} \mid Y^{\mathrm{S}}_2 = h_2\right)(p_{\mathrm{M}}).$$

$\mathrm{Var}\left(\bar{V}^{(2)}_{h_2}\right)$ and q_{h_2} are independent of p_{M}. Therefore, the increase in $\mathbb{E}\left[\nu^{(2)}_{h_2h_3}\right]$ equals the weighted

loss of posterior-mean dispersion conditional on stage-2 price path h_2:

$$\begin{aligned}&\mathbb{E}\left[\nu^{(2)}_{h_2h_3}\right](p_{\mathrm{M}}) - \mathbb{E}\left[\nu^{(2)}_{h_2h_3}\right](0)\\ =&\sum_{h_2} q_{h_2}\left[\mathrm{Var}\left(\bar{V}^{(2)}_{h_2h_3} \mid Y_2^{\mathrm{S}} = h_2\right)(0) - \mathrm{Var}\left(\bar{V}^{(2)}_{h_2h_3} \mid Y_2^{\mathrm{S}} = h_2\right)(p_{\mathrm{M}})\right].\end{aligned}$$

At the two agreeing price paths $h_2 \in \{++, --\}$, $q_{++} = q_{--} = \left(1+\pi^2\right)/4$, and, as we will prove below, the bracket equals

$$\frac{\sigma^2\pi^2 p_{\mathrm{M}}\left(1-\pi^2\right)^3}{\left(1+\pi^2\right)^2\left(1+3\pi^2\right)}.$$

At the two disagreeing price paths $h_2 \in \{+-, -+\}$, $q_{+-} = q_{-+} = \left(1-\pi^2\right)/4$, and, as we will prove below, the bracket equals

$$\frac{2\sigma^2\pi^2 p_{\mathrm{M}}}{1+p_{\mathrm{M}}}.$$

Therefore,

$$\mathbb{E}\left[\nu^{(2)}_{h_2h_3}\right](p_{\mathrm{M}}) - \mathbb{E}\left[\nu^{(2)}_{h_2h_3}\right](0) = \sigma^2\pi^2 p_{\mathrm{M}}\left(1-\pi^2\right)\left[\frac{\left(1-\pi^2\right)^2}{2\left(1+\pi^2\right)\left(1+3\pi^2\right)} + \frac{1}{1+p_{\mathrm{M}}}\right],$$

which strictly increases in p_{M} by inspection.[24] Therefore, for every $p_{\mathrm{M}} \in (0,1)$,

$$\begin{aligned}&\mathbb{E}\left[\nu^{(1)}_{h_2h_3}\right](0) - \mathbb{E}\left[\nu^{(2)}_{h_2h_3}\right](p_{\mathrm{M}})\\ =&\left[\mathbb{E}\left[\nu^{(1)}_{h_2h_3}\right](0) - \mathbb{E}\left[\nu^{(2)}_{h_2h_3}\right](0)\right] - \left[\mathbb{E}\left[\nu^{(2)}_{h_2h_3}\right](p_{\mathrm{M}}) - \mathbb{E}\left[\nu^{(2)}_{h_2h_3}\right](0)\right]\\ =&\sigma^2\pi^2\left(1-\pi^2\right)(1-p_{\mathrm{M}})\left[\frac{\left(1-\pi^2\right)^2}{2\left(1+\pi^2\right)\left(1+3\pi^2\right)} + \frac{1}{2\left(1+p_{\mathrm{M}}\right)}\right]\\ >&0.\end{aligned}$$

[24] It is straightforward to verify that the expression's left limit at $p_{\mathrm{M}} = 1$ equals the gain from adding the extra ordinary round in the no-manipulation benchmark $\mathbb{E}\left[\nu^{(1)}_{h_2h_3}\right](0) - \mathbb{E}\left[\nu^{(2)}_{h_2h_3}\right](0)$, as expected, since when $p_{\mathrm{M}} = 1$, stage 3 is fully occupied by the manipulator and does not contribute to price discovery at all.

Finally, since $p_{\mathrm{M}}^{*(1)}, p_{\mathrm{M}}^{*(2)} \in (0,1)$,

$$\mathbb{E}\left[\nu_{h_2h_3}^{(2)}\right]\left(p_{\mathrm{M}}^{*(2)}\right) < \mathbb{E}\left[\nu_{h_2h_3}^{(1)}\right](0) < \mathbb{E}\left[\nu_{h_2h_3}^{(1)}\right]\left(p_{\mathrm{M}}^{*(1)}\right),$$

where the second inequality follows from Proposition 3. This proves the price-discovery part of Proposition 6.

We now verify the claims about $\mathrm{Var}\left(\bar{V}_{h_2h_3}^{(2)} \mid Y_2^{\mathrm{S}} = h_2\right)(0) - \mathrm{Var}\left(\bar{V}_{h_2h_3}^{(2)} \mid Y_2^{\mathrm{S}} = h_2\right)(p_{\mathrm{M}})$ below.

Agreeing price paths $h_2 \in \{++, --\}$**.** At $h_2 = ++$, the stage-2 posterior mean is $\bar{V}_{++} = 2\sigma\pi/\left(1+\pi^2\right)$. When $p_{\mathrm{M}} = 0$, stage 3 is another ordinary spot round, so the conditional probability of a stage-3 buy is

$$q_{+++}(0) = \rho_{++}\frac{1+\pi}{2} + (1-\rho_{++})\frac{1-\pi}{2} = \frac{1+3\pi^2}{2\left(1+\pi^2\right)}$$

and the conditional probability of a stage-3 sell is

$$q_{++-}(0) = 1 - q_{+++}(0) = \frac{1-\pi^2}{2\left(1+\pi^2\right)}.$$

Bayes' rule gives the posterior mean after a stage-3 buy

$$\begin{aligned}
&\bar{V}_{+++}^{(2)}(0)\\
=&\sigma\frac{\mathbb{P}(h_2 = ++, h_3 = + \mid V = +\sigma)\,\mathbb{P}(V = +\sigma) - \mathbb{P}(h_2 = ++, h_3 = + \mid V = -\sigma)\,\mathbb{P}(V = -\sigma)}{\mathbb{P}(h_2 = ++, h_3 = + \mid V = +\sigma)\,\mathbb{P}(V = +\sigma) + \mathbb{P}(h_2 = ++, h_3 = + \mid V = -\sigma)\,\mathbb{P}(V = -\sigma)}\\
=&\sigma\frac{\left(\frac{1+\pi}{2}\right)^3\frac{1}{2} - \left(\frac{1-\pi}{2}\right)^3\frac{1}{2}}{\left(\frac{1+\pi}{2}\right)^3\frac{1}{2} + \left(\frac{1-\pi}{2}\right)^3\frac{1}{2}}\\
=&\sigma\frac{(1+\pi)^3 - (1-\pi)^3}{(1+\pi)^3 + (1-\pi)^3}\\
=&\frac{\sigma\pi\left(3+\pi^2\right)}{1+3\pi^2}
\end{aligned}$$

and similarly the posterior mean after a stage-3 sell

$$\bar{V}^{(2)}_{++-}(0) = \sigma \frac{(1+\pi)^2(1-\pi) - (1-\pi)^2(1+\pi)}{(1+\pi)^2(1-\pi) + (1-\pi)^2(1+\pi)} = \sigma\pi.$$

Therefore the posterior-mean dispersion conditional on stage-2 price path $h_2 = ++$ is

$$\begin{aligned}
&\mathrm{Var}\left(\bar{V}^{(2)}_{++h_3} \mid Y_2^{\mathrm{S}} = ++\right)(0) \\
=&\frac{1+3\pi^2}{2(1+\pi^2)}\left[\frac{\sigma\pi\left(3+\pi^2\right)}{1+3\pi^2} - \frac{2\sigma\pi}{1+\pi^2}\right]^2 + \frac{1-\pi^2}{2(1+\pi^2)}\left[\sigma\pi - \frac{2\sigma\pi}{1+\pi^2}\right]^2 \\
=&\frac{\sigma^2\pi^2\left(1-\pi^2\right)^3}{(1+\pi^2)^2(1+3\pi^2)}.
\end{aligned}$$

When $p_{\mathrm{M}} \in (0,1)$, the selected continuation equilibrium has $\alpha^*_{++} = 0$, so the manipulator does not pool with ordinary buy or sell orders. When she occupies stage 3 with probability p_{M}, she always produces a no-trade outcome which does not move the posterior and hence contributes zero to $\mathrm{Var}\left(\bar{V}^{(2)}_{++h_3} \mid Y_2^{\mathrm{S}} = ++\right)(p_{\mathrm{M}})$. Otherwise, with probability $1 - p_{\mathrm{M}}$, the manipulator is absent from stage 3 and it becomes an ordinary spot round. Thus

$$\mathrm{Var}\left(\bar{V}^{(2)}_{++h_3} \mid Y_2^{\mathrm{S}} = ++\right)(p_{\mathrm{M}}) = (1-p_{\mathrm{M}})\frac{\sigma^2\pi^2\left(1-\pi^2\right)^3}{(1+\pi^2)^2(1+3\pi^2)}.$$

The same calculation applies at $h_2 = --$ by symmetry:

$$\mathrm{Var}\left(\bar{V}^{(2)}_{--h_3} \mid Y_2^{\mathrm{S}} = --\right)(p_{\mathrm{M}}) = (1-p_{\mathrm{M}})\frac{\sigma^2\pi^2\left(1-\pi^2\right)^3}{(1+\pi^2)^2(1+3\pi^2)}.$$

Subtracting the two expressions verifies the claim.

Disagreeing price paths $h_2 \in \{+-, -+\}$**.** At $h_2 = +-$, the stage-2 posterior mean is $\bar{V}_{+-} = 0$. When $p_{\mathrm{M}} = 0$, stage 3 is another ordinary spot round, so the conditional probability of a stage-3 buy

is

$$q_{+-+}(0) = \rho_{+-}\frac{1+\pi}{2} + (1-\rho_{+-})\frac{1-\pi}{2} = \frac{1}{2}$$

and the conditional probability of a stage-3 sell is

$$q_{+--}(0) = 1 - q_{+-+}(0) = \frac{1}{2}.$$

Bayes' rule gives the posterior mean after a stage-3 buy

$$\begin{aligned}
&\bar{V}^{(2)}_{+-+}(0)\\
=&\sigma\frac{\mathbb{P}(h_2 = +-, h_3 = + \mid V = +\sigma)\,\mathbb{P}(V = +\sigma) - \mathbb{P}(h_2 = +-, h_3 = + \mid V = -\sigma)\,\mathbb{P}(V = -\sigma)}{\mathbb{P}(h_2 = +-, h_3 = + \mid V = +\sigma)\,\mathbb{P}(V = +\sigma) + \mathbb{P}(h_2 = +-, h_3 = + \mid V = -\sigma)\,\mathbb{P}(V = -\sigma)}\\
=&\sigma\frac{\left(\frac{1+\pi}{2}\right)^2\frac{1-\pi}{2}\frac{1}{2} - \left(\frac{1-\pi}{2}\right)^2\frac{1+\pi}{2}\frac{1}{2}}{\left(\frac{1+\pi}{2}\right)^2\frac{1-\pi}{2}\frac{1}{2} + \left(\frac{1-\pi}{2}\right)^2\frac{1+\pi}{2}\frac{1}{2}}\\
=&\sigma\frac{(1+\pi)^2(1-\pi) - (1-\pi)^2(1+\pi)}{(1+\pi)^2(1-\pi) + (1-\pi)^2(1+\pi)}\\
=&\sigma\pi,
\end{aligned}$$

and similarly the posterior mean after a stage-3 sell

$$\bar{V}^{(2)}_{+--}(0) = \sigma\frac{(1+\pi)(1-\pi)^2 - (1-\pi)(1+\pi)^2}{(1+\pi)(1-\pi)^2 + (1-\pi)(1+\pi)^2} = -\sigma\pi.$$

Therefore the posterior-mean dispersion conditional on stage-2 price path $h_2 = +-$ is

$$\mathrm{Var}\left(\bar{V}^{(2)}_{+-h_3} \mid Y_2^S = +-\right)(0) = \frac{1}{2}[\sigma\pi - 0]^2 + \frac{1}{2}[-\sigma\pi - 0]^2 = \sigma^2\pi^2.$$

When $p_M \in (0,1)$, the selected continuation equilibrium has $\alpha^*_{+-} = 1$ and equilibrium ask and bid as given in Appendix A.6.3:

$$A^*_{+-}(p_M) = \frac{\sigma\pi(1-p_M)}{1+p_M}, \qquad B_{+-} = -\sigma\pi.$$

The conditional probability of a stage-3 buy is

$$q_{+-+}(p_M) = p_M + (1-p_M)\, q_{+-+}(0) = \frac{1+p_M}{2}$$

and the conditional probability of a stage-3 sell is

$$q_{+--}(p_M) = 1 - q_{+-+}(p_M) = \frac{1-p_M}{2}.$$

A no-trade outcome at stage-3 must come from the manipulator, so it does not move the posterior and contributes zero to $\mathrm{Var}\left(\bar{V}^{(2)}_{+-h_3} \mid Y^S_2 = +-\right)(p_M)$. Therefore,

$$\mathrm{Var}\left(\bar{V}^{(2)}_{+-h_3} \mid Y^S_2 = +-\right)(p_M) = \frac{1+p_M}{2}\left[\frac{\sigma\pi(1-p_M)}{1+p_M} - 0\right]^2 + \frac{1-p_M}{2}\left[-\sigma\pi - 0\right]^2 = \sigma^2\pi^2\frac{1-p_M}{1+p_M}.$$

The same calculation applies at $h_2 = -+$ by symmetry:

$$\mathrm{Var}\left(\bar{V}^{(2)}_{-+h_3} \mid Y^S_2 = -+\right)(p_M) = \sigma^2\pi^2\frac{1-p_M}{1+p_M}.$$

Subtracting the two expressions verifies the claim.

A.7 n-Ordinary-Round Extension Calculations

We prove Proposition 7 in this appendix. The analysis is conditional on a stage-1 prediction-market buy; the sell branch is symmetric.

A.7.1 Stage 2 Posteriors

Fix $n \geq 1$. A stage-2 price path is a sequence of n outcomes:

$$h_2 = h_{2.1} \cdots h_{2.n} \in \{+, -\}^n .$$

As we will see, the relevant sufficient statistic for the stage-2 posterior that enters stage 3 is the order imbalance

$$d(h_2) = \sum_{j=1}^{n} \mathbf{1}\left\{h_{2.j} = +\right\} - \sum_{j=1}^{n} \mathbf{1}\left\{h_{2.j} = -\right\} = 2\sum_{j=1}^{n} \mathbf{1}\left\{h_{2.j} = +\right\} - n,$$

which takes values in the set

$$\mathcal{D}_n = \{-n, -n+2, \ldots, n-2, n\}.$$

For $k \in \mathcal{D}_n$, let

$$b(k) = \frac{n+k}{2}, \qquad s(k) = \frac{n-k}{2},$$

so that stage-2 price paths with imbalance k have $b(k)$ buys and $s(k)$ sells. Conditional on $V = +\sigma$, any individual stage-2 price path with imbalance k has probability

$$\left(\frac{1+\pi}{2}\right)^{b(k)} \left(\frac{1-\pi}{2}\right)^{s(k)},$$

whereas conditional on $V = -\sigma$, it has probability

$$\left(\frac{1-\pi}{2}\right)^{b(k)} \left(\frac{1+\pi}{2}\right)^{s(k)}.$$

Averaging over the symmetric prior on V, any individual stage-2 price path h_2 with $d(h_2) = k$ has probability

$$q_{h_2} = \frac{(1+\pi)^{b(k)}(1-\pi)^{s(k)} + (1-\pi)^{b(k)}(1+\pi)^{s(k)}}{2^{n+1}}.$$

Thus the total probability of all stage-2 price paths with imbalance k is

$$q_{d=k} :- \sum_{h_2 : d(h_2)=k} q_{h_2} = \binom{n}{b(k)} \frac{(1+\pi)^{b(k)} (1-\pi)^{s(k)} + (1-\pi)^{b(k)} (1+\pi)^{s(k)}}{2^{n+1}}. \tag{9}$$

Since $b(-k) = s(k) = n - b(k)$ and similarly $s(-k) = b(k) = n - s(k)$, we can verify that

$$q_{d=-k} = q_{d=k}. \tag{10}$$

Let

$$r :- \frac{1+\pi}{1-\pi}$$

denote the likelihood ratio of a buy order in an ordinary spot round: a buy is r times as likely when $V = +\sigma$ as when $V = -\sigma$. By symmetry, the likelihood ratio of a sell order is r^{-1}. Since the n ordinary spot rounds are conditionally independent given V, Bayes' rule gives directly, for any stage-2 price path h_2 with $d(h_2) = k$,

$$\begin{aligned}
&\rho_{h_2} \\
=&\mathbb{P}\left(V = +\sigma \mid Y_2^{\mathrm{S}} = h_2\right) \\
=&\frac{\mathbb{P}\left(Y_2^{\mathrm{S}} = h_2 \mid V = +\sigma\right)\mathbb{P}(V = +\sigma)}{\mathbb{P}\left(Y_2^{\mathrm{S}} = h_2 \mid V = +\sigma\right)\mathbb{P}(V = +\sigma) + \mathbb{P}\left(Y_2^{\mathrm{S}} = h_2 \mid V = -\sigma\right)\mathbb{P}(V = -\sigma)} \\
=&\frac{\left(\frac{1+\pi}{2}\right)^{b(k)}\left(\frac{1-\pi}{2}\right)^{s(k)}\frac{1}{2}}{\left(\frac{1+\pi}{2}\right)^{b(k)}\left(\frac{1-\pi}{2}\right)^{s(k)}\frac{1}{2} + \left(\frac{1-\pi}{2}\right)^{b(k)}\left(\frac{1+\pi}{2}\right)^{s(k)}\frac{1}{2}} = \frac{r^k}{1+r^k},
\end{aligned}$$

where the last equality cancels common factors and uses $b(k) - s(k) = k$. This shows that ρ_{h_2} depends on h_2 only through $k = d(h_2)$, which proves that $d(h_2)$ is indeed a sufficient statistic for the stage-2 posterior. We write this common value as

$$\rho_{d=k} = \frac{r^k}{1+r^k}.$$

The posterior mean and variance are therefore

$$\bar{V}_{d=k} = \sigma\left(2\rho_{d=k} - 1\right) = \sigma\frac{r^k - 1}{r^k + 1},$$

and

$$\nu_{d=k} = 4\sigma^2\rho_{d=k}\left(1 - \rho_{d=k}\right) = \frac{4\sigma^2 r^k}{\left(1 + r^k\right)^2}.$$

Thus every stage-2 price path with the same imbalance induces the same posterior:

$$\rho_{h_2} = \rho_{d=d(h_2)}, \qquad \bar{V}_{h_2} = \bar{V}_{d=d(h_2)}, \qquad \nu_{h_2} = \nu_{d=d(h_2)}.$$

A.7.2 Selected stage-3 continuation equilibrium

Substituting $\rho_{d=k}$ into the general ask and bid formulas (4) and (5) gives

$$A_{d=k}\left(\alpha_{d=k}; p_{\mathrm{M}}\right) = \sigma\frac{p_{\mathrm{M}}\alpha_{d=k}\left(\frac{r^k-1}{r^k+1}\right) + \frac{1-p_{\mathrm{M}}}{2}\left(\frac{r^k-1}{r^k+1} + \pi\right)}{p_{\mathrm{M}}\alpha_{d=k} + \frac{1-p_{\mathrm{M}}}{2}\left(1 + \frac{r^k-1}{r^k+1}\pi\right)}, \qquad B_{d=k} = \sigma\frac{r^{k-1} - 1}{r^{k-1} + 1}.$$

At imbalance k, the gap between the expected continuation payoff from buying and that from staying out is

$$\Gamma_{d=k}\left(\alpha_{d=k}; p_{\mathrm{M}}\right) = H\left(\frac{A_{d=k}\left(\alpha_{d=k}; p_{\mathrm{M}}\right)}{\varepsilon}\right) - H\left(\frac{\bar{V}_{d=k}}{\varepsilon}\right) - \left[A_{d=k}\left(\alpha_{d=k}; p_{\mathrm{M}}\right) - \bar{V}_{d=k}\right].$$

Favorable imbalances $k \geq 1$: pure stay-out equilibrium. For any $p_{\mathrm{M}} \in (0, 1)$, any $\alpha_{d=k} \in [0, 1]$, and any $k \geq 1$, the stage-3 ask is at least the prior mean, so

$$A_{d=k}\left(\alpha_{d=k}; p_{\mathrm{M}}\right) \geq \bar{V}_{d=k} \geq \bar{V}_{d=1} = \sigma\pi > \varepsilon.$$

Therefore $H\left(A_{d=k}\left(\alpha_{d=k}; p_{\mathrm{M}}\right)/\varepsilon\right) = H\left(\bar{V}_{d=k}/\varepsilon\right) = 1$. Hence

$$\Gamma_{d=k}\left(\alpha_{d=k}; p_{\mathrm{M}}\right) = 1 - 1 - \left[A_{d=k}\left(\alpha_{d=k}; p_{\mathrm{M}}\right) - \bar{V}_{d=k}\right] \leq 0.$$

Thus, for any $p_{\mathrm{M}} \in (0,1)$ and any $k \geq 1$, $\alpha^*_{d=k} = 0$ is a stable fixed point and is clearly the smallest.

Hopeless imbalances $k \leq -2$: pure stay-out equilibrium. For any $p_{\mathrm{M}} \in (0,1)$ and any $k \leq -2$,

$$A_{d=k}(0; p_{\mathrm{M}}) = \sigma \frac{r^{k+1} - 1}{r^{k+1} + 1} \leq \sigma \frac{r^{-2+1} - 1}{r^{-2+1} + 1} = \bar{V}_{d=-1} = -\sigma\pi < -\varepsilon.$$

Since the ask strictly decreases in $\alpha_{d=k}$ and is at least the prior mean,

$$\bar{V}_{d=k} \leq A_{d=k}(\alpha_{d=k}; p_{\mathrm{M}}) \leq A_{d=k}(0; p_{\mathrm{M}}) < -\varepsilon.$$

Consequently $H\left(A_{d=k}(\alpha_{d=k}; p_{\mathrm{M}})/\varepsilon\right) = H\left(\bar{V}_{d=k}/\varepsilon\right) = 0$, and

$$\Gamma_{d=k}(\alpha_{d=k}; p_{\mathrm{M}}) = 0 - 0 - \left[A_{d=k}(\alpha_{d=k}; p_{\mathrm{M}}) - \bar{V}_{d=k}\right] \leq 0.$$

Thus, for any $p_{\mathrm{M}} \in (0,1)$ and any $k \leq -2$, $\alpha^*_{d=k} = 0$ is a stable fixed point and is clearly the smallest.

Boundary imbalance $k = 0$: pure push equilibrium. $k = 0$ can occur only when n is even. Fix $p_{\mathrm{M}} \in (0,1)$. Here

$$\bar{V}_{d=0} = 0,$$

and the stage-3 ask is

$$A_{d=0}(\alpha_{d=0}; p_{\mathrm{M}}) = \frac{\sigma\pi(1 - p_{\mathrm{M}})}{1 - p_{\mathrm{M}} + 2p_{\mathrm{M}}\alpha_{d=0}}.$$

At $\alpha_{d=0} = 0$,

$$A_{d=0}(0; p_{\mathrm{M}}) = \sigma\pi > \varepsilon,$$

so

$$\Gamma_{d=0}(0; p_{\mathrm{M}}) = 1 - \frac{1}{2} - \sigma\pi = \frac{1}{2} - \sigma\pi > 0.$$

Thus $\alpha_{d=0} = 0$ is not a fixed point. For any $\alpha_{d=0} \in [0, 1]$, the ask is

$$A_{d=0}(\alpha_{d=0}; p_{\mathrm{M}}) \in \left[\frac{\sigma\pi(1 - p_{\mathrm{M}})}{1 + p_{\mathrm{M}}}, \sigma\pi\right] \subset \left(0, \frac{1}{2}\right).$$

If $A_{d=0}(\alpha_{d=0}; p_{\mathrm{M}}) \geq \varepsilon$, then

$$\Gamma_{d=0}(\alpha_{d=0}; p_{\mathrm{M}}) = \frac{1}{2} - A_{d=0}(\alpha_{d=0}; p_{\mathrm{M}}) \geq \frac{1}{2} - \sigma\pi > 0.$$

If $0 < A_{d=0}(\alpha_{d=0}; p_{\mathrm{M}}) < \varepsilon$, then

$$\Gamma_{d=0}(\alpha_{d=0}; p_{\mathrm{M}}) = A_{d=0}(\alpha_{d=0}; p_{\mathrm{M}})\left(\frac{1}{2\varepsilon} - 1\right) > 0$$

because $\varepsilon < 1/2$. Thus no interior fixed point exists either, and the selected continuation equilibrium is the endpoint

$$\alpha^*_{d=0} = 1.$$

Boundary imbalance $k = -1$: mixed or pure push. $k = -1$ can occur only when n is odd. Fix $p_{\mathrm{M}} \in (0, 1)$. Here

$$\bar{V}_{d=-1} = -\sigma\pi,$$

and

$$A_{d=-1}(\alpha_{d=-1}; p_{\mathrm{M}}) = \frac{-\sigma\pi p_{\mathrm{M}}\alpha_{d=-1}}{p_{\mathrm{M}}\alpha_{d=-1} + (1 - p_{\mathrm{M}})\frac{1-\pi^2}{2}}.$$

This is exactly the unfavorable-path continuation problem from Appendix A.5.1. Therefore, with

$$p_0 = \frac{\varepsilon\left(1 - \pi^2\right)(1 - 2\sigma\pi)}{2(\sigma\pi - \varepsilon) + \varepsilon(1 - \pi^2)(1 - 2\sigma\pi)},$$

the selected continuation equilibrium is

$$\alpha^*_{d=-1}(p_M) = \begin{cases} 1, & 0 < p_M \le p_0, \\ \dfrac{1-p_M}{p_M}\dfrac{\varepsilon\left(1-\pi^2\right)(1-2\sigma\pi)}{2(\sigma\pi-\varepsilon)}, & p_0 < p_M < 1. \end{cases}$$

The corresponding equilibrium ask is

$$A^*_{d=-1}(p_M) = \begin{cases} -\dfrac{2\sigma\pi p_M}{(1-\pi^2)+(1+\pi^2)p_M}, & 0 < p_M \le p_0, \\ -\dfrac{\varepsilon(1-2\sigma\pi)}{1-2\varepsilon}, & p_0 < p_M < 1. \end{cases}$$

Thus $d = -1$ is the unique pivotal imbalance when n is odd.

A.7.3 Optimal manipulator entry probability

Both the pricing of the binary contract and the manipulator's stage-1 entry problem use the selected continuation equilibrium from Appendix A.7.2. Appendix A.7.1 shows that the stage-2 posterior depends on h_2 only through the imbalance $d(h_2)$, so we can aggregate over imbalances rather than individual price paths. Combining Appendix A.7.1, the general ask and bid formulas (4) and (5), and Appendix A.7.2, we obtain the following stage-3 asks, bids, and no-trade posterior means for a stage-2 price path h_2 with $d(h_2) = k$:

$$A^*_{d=k}(p_M) = \begin{cases} \sigma\dfrac{r^{k+1}-1}{r^{k+1}+1} = \bar{V}_{d=k+1}, & k \ge 1 \text{ or } k \le -2, \\ \dfrac{\sigma\pi(1-p_M)}{1+p_M}, & k = 0, \\ -\dfrac{2\sigma\pi p_M}{(1-\pi^2)+(1+\pi^2)p_M}, & k = -1 \text{ and } 0 < p_M \le p_0, \\ -\dfrac{\varepsilon(1-2\sigma\pi)}{1-2\varepsilon}, & k = -1 \text{ and } p_0 < p_M < 1, \end{cases}$$

$$B_{d=k} = \sigma \frac{r^{k-1} - 1}{r^{k-1} + 1} = \bar{V}_{d=k-1}.$$

$$\bar{V}_{d=k} = \sigma \frac{r^{k} - 1}{r^{k} + 1}.$$

For each stage-2 price path h_2 with $d(h_2) = k$, the expected binary-contract payoffs are

$$\Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}}, k) = \alpha_{d=k}^{*}(p_{\mathrm{M}}) H\left(A_{d=k}^{*}(p_{\mathrm{M}})/\varepsilon\right) + \left(1 - \alpha_{d=k}^{*}(p_{\mathrm{M}})\right) H\left(\bar{V}_{d=k}/\varepsilon\right),$$

$$\Pi_{\mathrm{L}}^{+}(p_{\mathrm{M}}, k) = q_{d=k,+}(0) H\left(A_{d=k}^{*}(p_{\mathrm{M}})/\varepsilon\right) + q_{d=k,-}(0) H\left(B_{d=k}/\varepsilon\right).$$

Even n. Suppose n is even. Substituting the selected continuation equilibrium and the stage-3 quotes gives

$$\Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}}, k) = \begin{cases} 1, & k \geq 2, \\ H\left(\dfrac{\sigma\pi(1 - p_{\mathrm{M}})}{\varepsilon(1 + p_{\mathrm{M}})}\right), & k = 0, \\ 0, & k \leq -2, \end{cases}$$

$$\Pi_{\mathrm{L}}^{+}(p_{\mathrm{M}}, k) = \begin{cases} 1, & k \geq 2, \\ \dfrac{1}{2} H\left(\dfrac{\sigma\pi(1 - p_{\mathrm{M}})}{\varepsilon(1 + p_{\mathrm{M}})}\right), & k = 0, \\ 0, & k \leq -2. \end{cases}$$

The probability of imbalance $k = 0$ is $q_{d=0}$. By (10), the positive and negative tails have equal probability, so

$$\sum_{k \geq 2} q_{d=k} = \sum_{k \leq -2} q_{d=k} = \frac{1 - q_{d=0}}{2}.$$

Aggregating over k gives

$$\begin{aligned}
\Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}}) &= \frac{1-q_{d=0}}{2} + q_{d=0} H\left(\frac{\sigma\pi(1-p_{\mathrm{M}})}{\varepsilon(1+p_{\mathrm{M}})}\right), \\
\Pi_{\mathrm{L}}^{+}(p_{\mathrm{M}}) &= \frac{1-q_{d=0}}{2} + \frac{q_{d=0}}{2} H\left(\frac{\sigma\pi(1-p_{\mathrm{M}})}{\varepsilon(1+p_{\mathrm{M}})}\right), \\
K^{+}(p_{\mathrm{M}}) &= q_{d=0}\frac{\sigma\pi(1-p_{\mathrm{M}})}{1+p_{\mathrm{M}}}.
\end{aligned}$$

The unconstrained shadow stage-1 ask obtained from the prediction-market maker's break-even condition is

$$\begin{aligned}
&A_1^{\mathrm{P,BE}}(p_{\mathrm{M}}) \\
=&p_{\mathrm{M}}\Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}}) + (1-p_{\mathrm{M}})\Pi_{\mathrm{L}}^{+}(p_{\mathrm{M}}) \\
=&\frac{1-q_{d=0}}{2} + \frac{q_{d=0}(1+p_{\mathrm{M}})}{2} H\left(\frac{\sigma\pi(1-p_{\mathrm{M}})}{\varepsilon(1+p_{\mathrm{M}})}\right) \\
=&\begin{cases}
\dfrac{1}{2} + \dfrac{q_{d=0}p_{\mathrm{M}}}{2}, & \dfrac{\sigma\pi(1-p_{\mathrm{M}})}{1+p_{\mathrm{M}}} \geq \varepsilon, \\
\dfrac{1}{2} + \dfrac{q_{d=0}(1-p_{\mathrm{M}})(\sigma\pi-\varepsilon)}{4\varepsilon}, & \dfrac{\sigma\pi(1-p_{\mathrm{M}})}{1+p_{\mathrm{M}}} < \varepsilon.
\end{cases}
\end{aligned}$$

which is strictly greater than $1/2$ for any $p_{\mathrm{M}} \in (0,1)$. Therefore the non-crossing equilibrium stage-1 ask is just the break-even stage-1 ask:

$$A_1^{\mathrm{P}*}(p_{\mathrm{M}}) = \max\left\{A_1^{\mathrm{P,BE}}(p_{\mathrm{M}}), \frac{1}{2}\right\} = A_1^{\mathrm{P,BE}}(p_{\mathrm{M}}),$$

and

$$B_1^{\mathrm{P}*}(p_{\mathrm{M}}) = 1 - A_1^{\mathrm{P}*}(p_{\mathrm{M}})$$

for any $p_{\mathrm{M}} \in (0,1)$.

The manipulator's expected profit conditional on entering at stage 1 is

$$\begin{aligned}
&S_{\mathrm{M}}^{+}(p_{\mathrm{M}}) \\
=&\Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}}) - K^{+}(p_{\mathrm{M}}) - A_1^{\mathrm{P}*}(p_{\mathrm{M}}) \\
=&\frac{q_{d=0}}{2}\left[(1-p_{\mathrm{M}})\,H\left(\frac{\sigma\pi(1-p_{\mathrm{M}})}{\varepsilon(1+p_{\mathrm{M}})}\right) - \frac{2\sigma\pi(1-p_{\mathrm{M}})}{1+p_{\mathrm{M}}}\right].
\end{aligned}$$

Hence

$$\bar{S}_{\mathrm{M}}^{+}(p_{\mathrm{M}}) = p_{\mathrm{M}} S_{\mathrm{M}}^{+}(p_{\mathrm{M}}) = \frac{q_{d=0}}{2}\Phi(p_{\mathrm{M}};\sigma,\pi,\varepsilon),$$

where Φ is exactly the same function as defined in (8) for the manipulator's optimization problem under the two-ordinary-round extension. Thus every even n has the same optimal entry probability as the two-ordinary-round extension:

$$p_{\mathrm{M}}^{*(n)} = p_{\mathrm{M}}^{*(2)} \quad \text{if } n \text{ is even.}$$

Restoring horizon superscripts, since $q_{d=0} = m^{(n)}$ for even n and $m^{(2)} = \left(1-\pi^2\right)/2$, where $m^{(n)}$ is defined in Section 5.4 and $m^{(2)}$ in Section 5.4, the preceding formulas imply

$$\begin{aligned}
S_{\mathrm{M}}^{+(n)}\left(p_{\mathrm{M}}^{*(n)}\right) &= \frac{m^{(n)}}{m^{(2)}} S_{\mathrm{M}}^{+(2)}\left(p_{\mathrm{M}}^{*(2)}\right), \\
\bar{S}_{\mathrm{M}}^{+(n)}\left(p_{\mathrm{M}}^{*(n)}\right) &= \frac{m^{(n)}}{m^{(2)}} \bar{S}_{\mathrm{M}}^{+(2)}\left(p_{\mathrm{M}}^{*(2)}\right)
\end{aligned}$$

for even n.

Odd n. Suppose n is odd. Substituting the selected continuation equilibrium and the stage-3 quotes gives

$$\Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}},k)=\begin{cases}1, & k\geq 1,\\ \alpha_{d=-1}^{*}(p_{\mathrm{M}})\,H\left(A_{d=-1}^{*}(p_{\mathrm{M}})/\varepsilon\right), & k=-1,\\ 0, & k\leq -3,\end{cases}$$

$$\Pi_{\mathrm{L}}^{+}(p_{\mathrm{M}},k)=\begin{cases}1, & k\geq 3,\\ \dfrac{3+\pi^{2}}{4}, & k=1,\\ \dfrac{1-\pi^{2}}{2}H\left(A_{d=-1}^{*}(p_{\mathrm{M}})/\varepsilon\right), & k=-1,\\ 0, & k\leq -3.\end{cases}$$

The imbalances $k=\pm 1$ have probabilities $q_{d=\pm 1}$. By (10), $q_{d=1}=q_{d=-1}$ and the positive and negative tails have equal probability, so

$$\sum_{k\geq 3}q_{d=k}=\sum_{k\leq -3}q_{d=k}=\frac{1}{2}-q_{d=-1}.$$

Aggregating over k gives

$$\begin{aligned}\Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}})&=\frac{1}{2}+q_{d=-1}\alpha_{d=-1}^{*}(p_{\mathrm{M}})\,H\left(A_{d=-1}^{*}(p_{\mathrm{M}})/\varepsilon\right),\\ \Pi_{\mathrm{L}}^{+}(p_{\mathrm{M}})&=\frac{1}{2}+q_{d=-1}\left(1-\pi^{2}\right)\left[\frac{1}{2}H\left(A_{d=-1}^{*}(p_{\mathrm{M}})/\varepsilon\right)-\frac{1}{4}\right],\\ K^{+}(p_{\mathrm{M}})&=q_{d=-1}\alpha_{d=-1}^{*}(p_{\mathrm{M}})\left[A_{d=-1}^{*}(p_{\mathrm{M}})+\sigma\pi\right].\end{aligned}$$

The unconstrained shadow stage-1 ask obtained from the prediction-market maker's break-even condition is

$$
\begin{aligned}
&A_1^{\mathrm{P,BE}}(p_{\mathrm{M}}) \\
=&p_{\mathrm{M}}\Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}}) + (1-p_{\mathrm{M}})\,\Pi_{\mathrm{L}}^{+}(p_{\mathrm{M}}) \\
=&\frac{1}{2} + q_{d=-1}\left[p_{\mathrm{M}}\alpha_{d=-1}^{*}(p_{\mathrm{M}})\,H\left(A_{d=-1}^{*}(p_{\mathrm{M}})/\varepsilon\right) + (1-p_{\mathrm{M}})\left(1-\pi^2\right)\left[\frac{1}{2}H\left(A_{d=-1}^{*}(p_{\mathrm{M}})/\varepsilon\right) - \frac{1}{4}\right]\right] \\
=&\begin{cases}
\dfrac{1}{2} - q_{d=-1}\dfrac{p_{\mathrm{M}}(\sigma\pi-\varepsilon)}{2\varepsilon}, & 0 < p_{\mathrm{M}} \le p_0, \\[2ex]
\dfrac{1}{2} - q_{d=-1}\dfrac{(1-p_{\mathrm{M}})\left(1-\pi^2\right)(1-2\sigma\pi)}{4}, & p_0 < p_{\mathrm{M}} < 1.
\end{cases}
\end{aligned}
$$

which is strictly less than $1/2$ for any $p_{\mathrm{M}} \in (0,1)$. Therefore the non-crossing equilibrium stage-1 ask and bid are respectively

$$
A_1^{\mathrm{P*}}(p_{\mathrm{M}}) = \max\left\{A_1^{\mathrm{P,BE}}(p_{\mathrm{M}}), \frac{1}{2}\right\} = \frac{1}{2}, \qquad B_1^{\mathrm{P*}}(p_{\mathrm{M}}) = 1 - A_1^{\mathrm{P*}}(p_{\mathrm{M}}) = \frac{1}{2}
$$

for any $p_{\mathrm{M}} \in (0,1)$.

The manipulator's expected profit conditional on entering at stage 1 is

$$
S_{\mathrm{M}}^{+}(p_{\mathrm{M}}) = \Pi_{\mathrm{M}}^{+}(p_{\mathrm{M}}) - K^{+}(p_{\mathrm{M}}) - \frac{1}{2},
$$

which yields

$$
S_{\mathrm{M}}^{+}(p_{\mathrm{M}}) = \begin{cases}
q_{d=-1}\dfrac{\left(1-\pi^2\right)(1-2\sigma\pi)\left(1-p_{\mathrm{M}}/p_0\right)}{2\left[\left(1-\pi^2\right)+\left(1+\pi^2\right)p_{\mathrm{M}}\right]}, & 0 < p_{\mathrm{M}} \le p_0, \\[2ex]
0, & p_0 < p_{\mathrm{M}} < 1.
\end{cases}
$$

The pure-push branch is positive when $0 < p_{\mathrm{M}} < p_0$ and vanishes at $p_{\mathrm{M}} = p_0$. Therefore, the

manipulator maximizes the ex-ante profit

$$\bar{S}^{+}_{\mathrm{M}}(p_{\mathrm{M}}) = p_{\mathrm{M}} S^{+}_{\mathrm{M}}(p_{\mathrm{M}}) = q_{d=-1} \frac{p_{\mathrm{M}}\left(1-\pi^2\right)(1-2\sigma\pi)\left(1-p_{\mathrm{M}}/p_0\right)}{2\left[(1-\pi^2)+(1+\pi^2)\,p_{\mathrm{M}}\right]}$$

over $0 < p_{\mathrm{M}} < p_0$. This equals $2q_{d=-1}$ times the manipulator's ex-ante profit (7) in the baseline model. Thus every odd n has the same optimal entry probability as the baseline model:

$$p^{*(n)}_{\mathrm{M}} = p^{*(1)}_{\mathrm{M}} \quad \text{if } n \text{ is odd.}$$

Since $p^{*(n)}_{\mathrm{M}} = p^{*(1)}_{\mathrm{M}} \in (0, p_0)$, the optimum lies on the pure-push branch. Restoring horizon superscripts, since $q_{d=-1} = m^{(n)}$ for odd n and $m^{(1)} = 1/2$, where $m^{(n)}$ is defined in Section 5.4 and $m^{(1)}$ in Section 5.4, the preceding formulas imply

$$\begin{aligned} S^{+(n)}_{\mathrm{M}}\left(p^{*(n)}_{\mathrm{M}}\right) &= \frac{m^{(n)}}{m^{(1)}} S^{+(1)}_{\mathrm{M}}\left(p^{*(1)}_{\mathrm{M}}\right), \\ \bar{S}^{+(n)}_{\mathrm{M}}\left(p^{*(n)}_{\mathrm{M}}\right) &= \frac{m^{(n)}}{m^{(1)}} \bar{S}^{+(1)}_{\mathrm{M}}\left(p^{*(1)}_{\mathrm{M}}\right) \end{aligned}$$

for odd n. The same observation pins down

$$\alpha^{*(n)}_{d=-1}\left(p^{*(n)}_{\mathrm{M}}\right) = 1 \quad \text{if } n \text{ is odd.}$$

A.7.4 Price discovery

Fix $n > 2$. In this subsection only, write $q^{(n)}_{d=k}$ for the imbalance-k probability in (9) computed under the n-ordinary-round extension to emphasize the dependence on n. For any stage-2 price path h_2 with $d(h_2) = k$ and any $p^{(n)}_{\mathrm{M}} \in (0,1)$, also write $q_{d=k,h_3}\left(p^{(n)}_{\mathrm{M}}\right)$, $\bar{V}_{d=k,h_3}\left(p^{(n)}_{\mathrm{M}}\right)$, and $v_{d=k,h_3}\left(p^{(n)}_{\mathrm{M}}\right)$ for the corresponding conditional stage-3 outcome probability, posterior mean, and posterior variance.

Then

$$\mathbb{E}\left[v^{(n)}_{h_2 h_3}\right]\left(p^{(n)}_{\mathrm{M}}\right) = \sum_{k \in \mathcal{D}_n} q^{(n)}_{d=k} \sum_{h_3 \in \{+,0,-\}} q_{d=k,h_3}\left(p^{(n)}_{\mathrm{M}}\right) v_{d=k,h_3}\left(p^{(n)}_{\mathrm{M}}\right)$$

by the law of total expectation. To prove the price-discovery part of Proposition 7, it suffices to compare horizon n to horizon $n-1$ through the no-manipulation $(n-1)$-stage benchmark. We show, for every $p_{\mathrm{M}}^{(n)}, p_{\mathrm{M}}^{(n-1)} \in (0,1)$,

$$\mathbb{E}\left[\nu_{h_2h_3}^{(n)}\right]\left(p_{\mathrm{M}}^{(n)}\right) < \mathbb{E}\left[\nu_{h_2h_3}^{(n-1)}\right](0) < \mathbb{E}\left[\nu_{h_2h_3}^{(n-1)}\right]\left(p_{\mathrm{M}}^{(n-1)}\right).$$

At the no-manipulation benchmark $p_{\mathrm{M}}^{(n-1)} = 0$, the $(n-1)$-ordinary-round model has $n-1$ ordinary spot rounds in stage 2 and one ordinary spot round in stage 3, and thus stages 2 and 3 together share the same event structure as stage 2 under the n-ordinary-round extension. Hence

$$\mathbb{E}\left[\nu_{h_2h_3}^{(n-1)}\right](0) = \sum_{k\in\mathcal{D}_n} q_{d=k}^{(n)} \nu_{d=k}.$$

For each $k \in \mathcal{D}_n$, the law of total variance gives

$$\sum_{h_3\in\{+,\mathbb{0},-\}} q_{d=k,h_3}\left(p_{\mathrm{M}}^{(n)}\right)\nu_{d=k,h_3}\left(p_{\mathrm{M}}^{(n)}\right) = \nu_{d=k} - \mathrm{Var}\left(\bar{V}_{d=k,h_3} \mid d(h_2)=k\right)\left(p_{\mathrm{M}}^{(n)}\right).$$

Since $\rho_{d=k} \in (0,1)$,

$$q_{d=k,-}\left(p_{\mathrm{M}}^{(n)}\right) = \left(1-p_{\mathrm{M}}^{(n)}\right)\left[\frac{1-\pi}{2} + \pi\left(1-\rho_{d=k}\right)\right] > 0,$$

$$\bar{V}_{d=k,-}\left(p_{\mathrm{M}}^{(n)}\right) - \bar{V}_{d=k} = B_{d=k} - \bar{V}_{d=k} = -\frac{\sigma\pi\left[1-(2\rho_{d=k}-1)^2\right]}{1-(2\rho_{d=k}-1)\pi} < 0.$$

Therefore

$$\mathrm{Var}\left(\bar{V}_{d=k,h_3} \mid d(h_2)=k\right)\left(p_{\mathrm{M}}^{(n)}\right) \geq q_{d=k,-}\left(p_{\mathrm{M}}^{(n)}\right)\left[\bar{V}_{d=k,-}\left(p_{\mathrm{M}}^{(n)}\right) - \bar{V}_{d=k}\right]^2 > 0.$$

Taking the $q^{(n)}_{d=k}$-weighted average of the law-of-total-variance identity over $k \in \mathcal{D}_n$ gives

$$\mathbb{E}\left[\nu^{(n)}_{h_2h_3}\right]\left(p^{(n)}_{\mathrm{M}}\right) = \sum_{k\in\mathcal{D}_n} q^{(n)}_{d=k}\left[\nu_{d=k} - \mathrm{Var}\left(\bar{V}_{d=k,h_3} \mid d\left(h_2\right) = k\right)\left(p^{(n)}_{\mathrm{M}}\right)\right].$$

Each $q^{(n)}_{d=k}$ is strictly positive, and the previous display shows that each subtracted posterior-mean dispersion term is strictly positive, so

$$\mathbb{E}\left[\nu^{(n)}_{h_2h_3}\right]\left(p^{(n)}_{\mathrm{M}}\right) < \sum_{k\in\mathcal{D}_n} q^{(n)}_{d=k}\nu_{d=k} = \mathbb{E}\left[\nu^{(n-1)}_{h_2h_3}\right](0).$$

Finally, although Proposition 3 is stated for the baseline model, its proof in Appendix A.4 conditions only on the stage-2 posterior and thus still applies to the $(n-1)$-ordinary-round extension, giving, for every $p^{(n-1)}_{\mathrm{M}} \in (0,1)$,

$$\mathbb{E}\left[\nu^{(n-1)}_{h_2h_3}\right]\left(p^{(n-1)}_{\mathrm{M}}\right) > \mathbb{E}\left[\nu^{(n-1)}_{h_2h_3}\right](0).$$

Chaining the inequalities gives the stated result.

B Polymarket Contract Trading Volume

This appendix documents the trading volume of the Polymarket crypto up/down contracts that the main analysis studies, across the six coins (BTC, ETH, SOL, XRP, DOGE, BNB) and the main contract durations (five-minute, fifteen-minute, four-hour). The contracts carry tens of millions of dollars of daily volume; the five-minute contract is the most heavily traded and BTC dominates the cross-section.

By coin. Figure A1a stacks daily Polymarket crypto up/down USDC volume by coin, summed across all contract durations, from the earliest Polymarket trade. Volume grows from a few million dollars a day in October 2025 to roughly $30–40M by March 2026, with a visible step up when the five-minute contract launches in February. BTC is the large majority of the total throughout, ETH a distant second, and SOL/XRP/DOGE/BNB thin slivers, consistent with the manipulation footprint being concentrated in the coin whose pool is deepest.

By contract duration. Figure A1b stacks daily BTC volume by contract duration. The fifteen-minute contract carries most of the volume through the pre-five-minute period and the four-hour contract a thin sliver; the five-minute layer appears at its February launch and quickly becomes the largest, raising total daily BTC volume to roughly $30M.

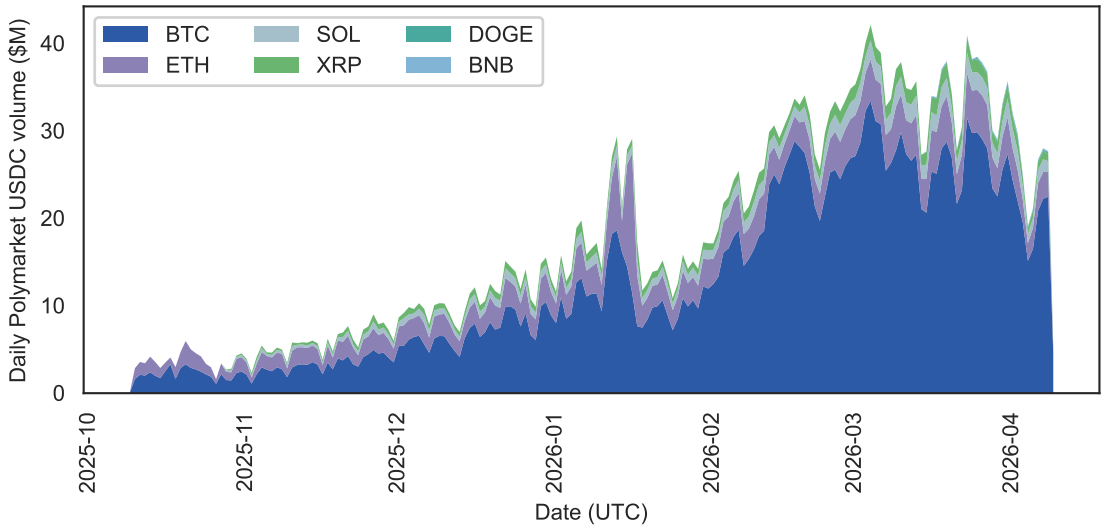

(a) All coins, by coin.

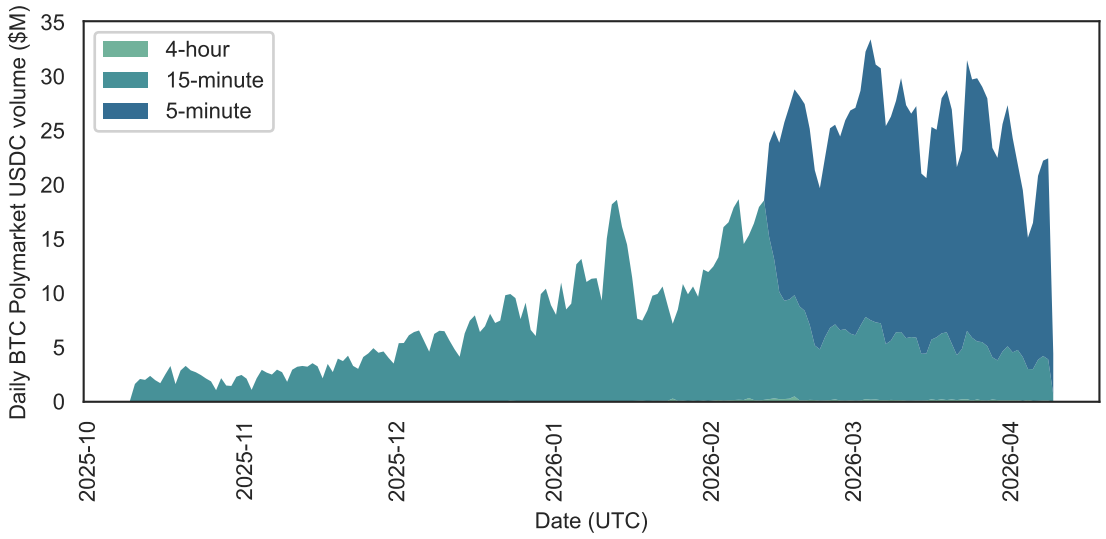

(b) BTC, by contract duration.

Figure A1. Daily Polymarket crypto up/down USDC volume. This figure plots stacked daily taker USDC notional on the Polymarket crypto up/down contracts, from the earliest Polymarket trade (the October 2025 fifteen-minute / four-hour launch). **(a)** All six coins, summed across the three durations (five-minute, fifteen-minute, four-hour), stacked by coin. **(b)** BTC only, stacked by contract duration (four-hour, fifteen-minute, five-minute, from bottom to top).

C Overlap and Moneyness Refinements

The two subsections below refine the cycle-level partition under which the main-body activity and reversal results are estimated. The first refines the within-P3 activity regression (Table 4) and reversal regression (Table 5) into the full 15-min-overlap × per-contract-NTM groups. The second applies the moneyness cut to the fifteen-minute contract alone in P2, before the five-minute launch, asking whether the NTM concentration already appears at the longer horizon.

Proof of Proposition 8. Since $1-\pi^2 \in (0,1)$ and $\binom{n}{\lfloor n/2 \rfloor} \leq 2^n$,

$$0 \leq m^{(n)} \leq \left(1-\pi^2\right)^{\lfloor n/2 \rfloor}.$$

The limit of $m^{(n)}$ then follows from the squeeze theorem:

$$\lim_{n \uparrow \infty} \left(1-\pi^2\right)^{\lfloor n/2 \rfloor} = 0.$$

Appendix A.7.3 shows that the optimal profit objects are the corresponding lower-horizon objects scaled by the pivotal probabilities:

$$S_{\mathrm{M}}^{+(n)}\left(p_{\mathrm{M}}^{*(n)}\right) = \begin{cases} \dfrac{m^{(n)}}{m^{(1)}} S_{\mathrm{M}}^{+(1)}\left(p_{\mathrm{M}}^{*(1)}\right), & n \text{ is odd}, \\ \dfrac{m^{(n)}}{m^{(2)}} S_{\mathrm{M}}^{+(2)}\left(p_{\mathrm{M}}^{*(2)}\right), & n \text{ is even}, \end{cases}$$

and similarly

$$\bar{S}_{\mathrm{M}}^{+(n)}\left(p_{\mathrm{M}}^{*(n)}\right) = \begin{cases} \dfrac{m^{(n)}}{m^{(1)}} \bar{S}_{\mathrm{M}}^{+(1)}\left(p_{\mathrm{M}}^{*(1)}\right), & n \text{ is odd}, \\ \dfrac{m^{(n)}}{m^{(2)}} \bar{S}_{\mathrm{M}}^{+(2)}\left(p_{\mathrm{M}}^{*(2)}\right), & n \text{ is even}. \end{cases}$$

Clearly, as $n \uparrow \infty$, both profit objects converge to zero along with $m^{(n)}$. □

C.1 Activity and reversal by overlap and moneyness

This subsection refines the within-P3 NTM results, the activity regression of Table 4 and the reversal regression of Table 5, by whether the cycle's close also coincides with a 15-minute mark, where the 15m Polymarket contract settles alongside the 5m contract. The cycle groups are written in set notation: 5m ∩ 15m denotes overlapping cycles, where both the 5m and 15m Polymarket contracts settle at the close (four of twelve per hour), and 5m \ 15m non-overlapping cycles, where only the 5m contract settles. NTM is per contract from the Polymarket Up-token price near time-to-maturity TTM ∈ [10, 12] s: $\mathrm{NTM}^{\mathrm{poly}}_{5\mathrm{m}}$ iff $|p_{Up,5\mathrm{m}} - 0.5| < 0.10$, $\mathrm{NTM}^{\mathrm{poly}}_{15\mathrm{m}}$ iff $|p_{Up,15\mathrm{m}} - 0.5| < 0.10$ (defined only for overlapping cycles). Together these yield six mutually exclusive groups: 5m \ 15m split by 5m-contract NTM, and 5m ∩ 15m split into neither-NTM, $\mathrm{NTM}_{5\mathrm{m}}$-only, $\mathrm{NTM}_{15\mathrm{m}}$-only, and both-NTM.

Activity. Restricting to P3 and to bins in {0–24, 29}, the regression is

$$
\begin{aligned}
\log y_{c,b} \;=\; & \alpha_d + \gamma_b + \mu_h + \delta_{5\mathrm{m}\cap 15\mathrm{m}}\,\mathbf{1}\{b{=}29\}\cdot D_{5\mathrm{m}\cap 15\mathrm{m},c} \\
& + \delta_{\mathrm{NTM}_{5\mathrm{m}},\,5\mathrm{m}\setminus 15\mathrm{m}}\,\mathbf{1}\{b{=}29\}\cdot D_{5\mathrm{m}\setminus 15\mathrm{m},c}\cdot\mathbf{1}\{\mathrm{NTM}^{\mathrm{poly}}_{5\mathrm{m},c}\} \\
& + \delta_{\mathrm{NTM}_{5\mathrm{m}},\,5\mathrm{m}\cap 15\mathrm{m}}\,\mathbf{1}\{b{=}29\}\cdot D_{5\mathrm{m}\cap 15\mathrm{m},c}\cdot\mathbf{1}\{\mathrm{NTM}^{\mathrm{poly}}_{5\mathrm{m},c}\} \\
& + \delta_{\mathrm{NTM}_{15\mathrm{m}},\,5\mathrm{m}\cap 15\mathrm{m}}\,\mathbf{1}\{b{=}29\}\cdot D_{5\mathrm{m}\cap 15\mathrm{m},c}\cdot\mathbf{1}\{\mathrm{NTM}^{\mathrm{poly}}_{15\mathrm{m},c}\} \\
& + \delta_{\mathrm{NTM}_{5\mathrm{m}}\times\mathrm{NTM}_{15\mathrm{m}},\,5\mathrm{m}\cap 15\mathrm{m}}\,\mathbf{1}\{b{=}29\}\cdot D_{5\mathrm{m}\cap 15\mathrm{m},c}\cdot\mathbf{1}\{\mathrm{NTM}^{\mathrm{poly}}_{5\mathrm{m},c}\}\cdot\mathbf{1}\{\mathrm{NTM}^{\mathrm{poly}}_{15\mathrm{m},c}\} \\
& + (\text{level effects}) + \varepsilon_{c,b},
\end{aligned}
$$

estimated with date, hour-of-day, and bin fixed effects and cycle-clustered standard errors. The baseline, absorbed into the bin fixed effect γ_{29}, is the 5m \ 15m × FFM group; each δ is the bin-29-vs-body increase in its group over that baseline, $\delta_{5\mathrm{m}\cap 15\mathrm{m}}$ in overlapping cycles where neither contract is NTM, $\delta_{\mathrm{NTM}_{5\mathrm{m}},\,5\mathrm{m}\setminus 15\mathrm{m}}$ the 5m-contract NTM effect in non-overlapping cycles, $\delta_{\mathrm{NTM}_{5\mathrm{m}},\,5\mathrm{m}\cap 15\mathrm{m}}$ the 5m-contract NTM effect in overlapping cycles, and $\delta_{\mathrm{NTM}_{15\mathrm{m}},\,5\mathrm{m}\cap 15\mathrm{m}}$ the incremental 15m-contract NTM effect in overlapping cycles. The final term, $\delta_{\mathrm{NTM}_{5\mathrm{m}}\times\mathrm{NTM}_{15\mathrm{m}},\,5\mathrm{m}\cap 15\mathrm{m}}$, is the both-NTM interaction:

it saturates the four overlap-moneyness cells, so the both-NTM cell is a free parameter rather than the additive sum of the two single-contract effects.

In the order-flow column of Table A1, the overlap-without-NTM effect is small ($\hat{\delta}_{5\text{m}\cap 15\text{m}} = 0.21$), while the 5m-contract NTM effect is large in both the non-overlap group ($\hat{\delta}_{\text{NTM}_{5\text{m}},\,5\text{m}\setminus 15\text{m}} = 1.32$) and the overlap group ($\hat{\delta}_{\text{NTM}_{5\text{m}},\,5\text{m}\cap 15\text{m}} = 1.39$), and the 15m-contract NTM coefficient adds a further $\hat{\delta}_{\text{NTM}_{15\text{m}},\,5\text{m}\cap 15\text{m}} = 1.12$ inside the overlap group. The both-NTM interaction is positive on top of that: in cycles NTM by *both* contracts, the bin-29 spike exceeds the sum of the two single-contract effects by a further $\hat{\delta}_{\text{NTM}_{5\text{m}}\times\text{NTM}_{15\text{m}}} = 1.14$ log points (significant at the 5% level despite resting on only 11 such cycles), so the both-NTM overlap close sums all four terms, 0.21 + 1.39 + 1.12 + 1.14 = 3.87 log points ($e^{3.87} \approx 48\times$ on the typical cycle) above the 5m \ 15m FFM baseline. This is what the model predicts: with two pools winnable at one close, a push is worth more than the sum of its parts. The same ordering holds for absolute return, and in the Polymarket columns the depth coefficients are again negative (prediction-market makers withdraw). The activity surge thus compounds with both moneyness and contract overlap: it is largest exactly where two contracts settle at one close and the outcome is still live by both.

Reversal. The reversal counterpart estimates, at each lag, a baseline slope plus four additive increments over the same six groups: the 5m \ 15m × FFM group is the baseline slope γ_j on r_{29}, and we add $\delta_{5\text{m}\cap 15\text{m}}$ (overlapping close, neither contract NTM), $\delta_{\text{NTM}_{5\text{m}},\,5\text{m}\setminus 15\text{m}}$ (5m-contract NTM in non-overlapping cycles), $\delta_{\text{NTM}_{5\text{m}},\,5\text{m}\cap 15\text{m}}$ (5m-contract NTM in overlapping cycles), and $\delta_{\text{NTM}_{15\text{m}},\,5\text{m}\cap 15\text{m}}$ (incremental 15m-contract NTM in overlapping cycles), each as a slope increment on r_{29}. The nine overlapping cycles NTM by both contracts load on both overlap-NTM increments; unlike the saturated activity regression, nine cycles cannot support a free interaction slope.

Figure A2 plots these slopes by group. At lag $j = +1$ in Table A2, the moneyness increments carry the result. Against the non-overlapping FFM baseline ($\gamma_{+1} = -0.08$), a non-overlapping cycle that is $\text{NTM}_{5\text{m}}$ adds a further −0.16, about tripling the baseline reversal, and within overlapping cycles, NTM by either contract adds about −0.22 ($\delta_{\text{NTM}_{5\text{m}},\,5\text{m}\cap 15\text{m}} = -0.23$ and $\delta_{\text{NTM}_{15\text{m}},\,5\text{m}\cap 15\text{m}} =$

Table A1. Within-P3 bin-29 activity increase by 15-min overlap and per-contract Polymarket NTM, Binance and Polymarket. This table reports cycle-level OLS estimates on the (cycle, bin) panel restricted to bins {0–24, 29} and P3 cycles, with date, hour-of-day, and bin fixed effects; dependent variable $\log y_{c,b}$, the log of the column metric. Coefficients are log differences (δ is a multiplicative e^{δ} effect on the typical geometric-mean cycle). Set notation for the cycle group: 5m ∩ 15m = 5m cycles whose close coincides with a 15-min boundary (both the 5m and 15m contracts settle); 5m \ 15m = the other eight of twelve per hour (only the 5m contract settles). NTM is per contract from the Polymarket Up-token price near TTM ∈ [10, 12] s: $\text{NTM}^{\text{poly}}_{5\text{m}}$ iff $|p_{Up,5\text{m}} - 0.5| < 0.10$, $\text{NTM}^{\text{poly}}_{15\text{m}}$ iff $|p_{Up,15\text{m}} - 0.5| < 0.10$ (overlapping cycles only). The five reported coefficients are $\mathbf{1}\{b{=}29\}$ interactions; the 5m \ 15m× FFM bin-29 increase is the baseline, absorbed in the bin FE. The $\delta_{\text{NTM}_{5\text{m}}\times\text{NTM}_{15\text{m}}}$ interaction saturates the four overlap-moneyness cells, so each cell's bin-29 increase is a free parameter rather than an additive constraint; on the Binance order-flow sample the cells hold 586 (5m \ 15m × $\text{NTM}_{5\text{m}}$), 211 ($\text{NTM}_{5\text{m}}$-only), 132 ($\text{NTM}_{15\text{m}}$-only), and 11 (both-NTM) cycles, so the interaction is estimated from few cycles. One observation per (cycle, bin) cell. HC1 standard errors in parentheses are clustered at the cycle level. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Binance			Polymarket		
	Order flow	Abs. return	Depth	Order flow	Abs. price change	Depth
$\delta_{5\text{m}\cap 15\text{m}}$	0.212***	0.109***	0.013*	0.163***	0.768***	0.026
	(0.043)	(0.034)	(0.008)	(0.032)	(0.153)	(0.047)
$\delta_{\text{NTM}_{5\text{m}},\,5\text{m}\setminus 15\text{m}}$	1.324***	0.555***	0.154***	−0.090*	2.896***	−0.980***
	(0.105)	(0.063)	(0.026)	(0.052)	(0.110)	(0.064)
$\delta_{\text{NTM}_{5\text{m}},\,5\text{m}\cap 15\text{m}}$	1.393***	0.479***	0.168***	−0.307***	1.879***	−1.146***
	(0.185)	(0.111)	(0.041)	(0.099)	(0.176)	(0.113)
$\delta_{\text{NTM}_{15\text{m}},\,5\text{m}\cap 15\text{m}}$	1.123***	0.435***	0.205***	0.167	0.251	−0.234
	(0.228)	(0.145)	(0.062)	(0.134)	(0.517)	(0.211)
$\delta_{\text{NTM}_{5\text{m}}\times\text{NTM}_{15\text{m}},\,5\text{m}\cap 15\text{m}}$	1.143**	0.457*	0.259	0.290	−1.331	−0.060
	(0.583)	(0.277)	(0.262)	(0.341)	(1.097)	(0.324)
Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Hour-of-day FE	Yes	Yes	Yes	Yes	Yes	Yes
Bin FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	352,370	272,785	352,370	351,791	122,958	126,552
Clusters (cycle)	13,553	13,551	13,553	13,553	5,001	5,001

−0.22), all significant; the deepest-reverting groups, overlap × NTM, revert about five times as hard as the baseline. That the $\text{NTM}_{5\text{m}}$ and $\text{NTM}_{15\text{m}}$ increments are almost identical says each contract's moneyness is on its own enough to mark a high-reversal cycle, consistent with the manipulator-selection reading: the payoff per dollar of push is larger when two contracts settle at the same close. The overlap level increment ($\delta_{5\text{m}\cap 15\text{m}} = -0.08$) we treat as a descriptive control rather than a settlement effect: 15-minute marks attract round-time algorithmic flow of their own (Section 6.3.3). The manipulation reading therefore rests on the moneyness increments, which are defined by the live contracts' own prices and have no clock analogue.

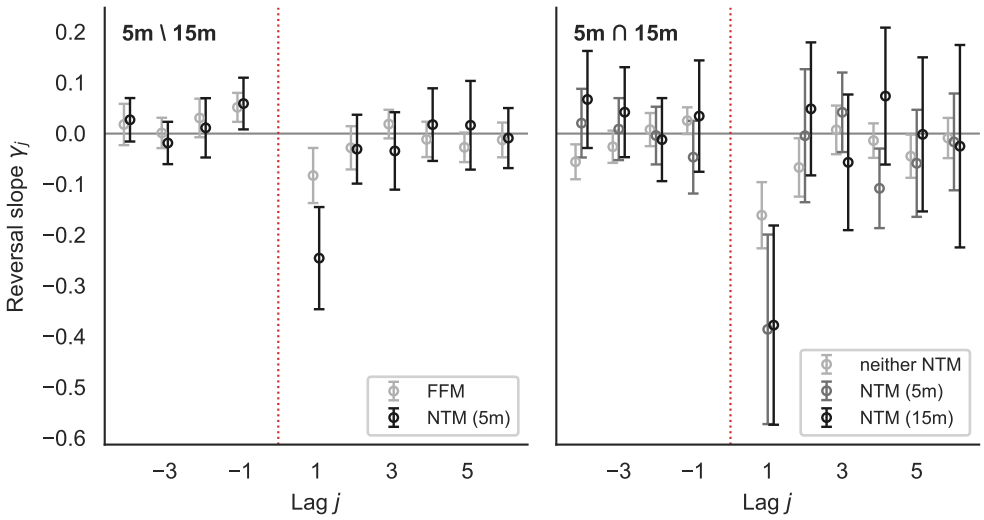

Figure A2. Within-P3 post-settlement reversal slope γ_j, by cycle group. This figure plots γ_j for each cycle group (the underlying regression is Table A2): left panel non-overlapping (5m \ 15m) cycles, right panel overlapping (5m ∩ 15m) cycles, each split by per-contract Polymarket NTM. Markers are γ_j; error bars are ±1.96× cluster-robust SE (clustered by date); $j = 0$ is omitted. The both-NTM overlap group (nine cycles) is omitted.

C.2 Moneyness for the fifteen-minute contract alone

Section 6.3.3 finds only a strongly attenuated fifteen-minute footprint, but that test is unconditional: it averages over all overlapping closes, most of which are far from the money by the time the contract settles. Within P3, the footprint concentrates in NTM cycles (Section 6.3.2), so a genuine fifteen-minute effect could hide in the unconditional average, diluted by the FFM majority. This appendix therefore applies the moneyness cut to the fifteen-minute contract on its own. We take P2, where the fifteen-minute (and four-hour) contract is live but the five-minute contract is not, restrict to overlapping cycles, those whose close lands on a 15-min mark, where the fifteen-minute contract settles, and classify them NTM vs FFM by the *fifteen-minute* contract's Polymarket Up-token price within 0.10 of 0.5 near TTM ∈ [10, 12] s, the same moneyness rule applied to the contract that is actually live. If the fifteen-minute horizon supported manipulation, its footprint should surface here, in exactly the cycles where a push could flip its resolution.

Table A3 reports the result, and it is largely a null. Panel A shows a modest bin-29 activity increase in NTM cycles — about 0.24 in order flow and 0.30 in absolute return, roughly 27% and

Table A2. Within-P3 post-settlement reversal slope by 15-min overlap and per-contract moneyness. This table reports, at each lag $j \in \{-3, -2, -1, +1, +2, +3\}$, the coefficients from a single cycle-level OLS on P3 cycles (Feb 12 – Apr 8, 2026): the 5m \ 15m × FFM group is the baseline reversal slope γ_j on r_{29}, and the four reported δ are additive slope increments over it, $\delta_{5\text{m}\cap 15\text{m}}$ (overlapping close), $\delta_{\text{NTM}_{5\text{m}},\,5\text{m}\setminus 15\text{m}}$ and $\delta_{\text{NTM}_{5\text{m}},\,5\text{m}\cap 15\text{m}}$ (5m-contract NTM in non-overlapping and overlapping cycles), and $\delta_{\text{NTM}_{15\text{m}},\,5\text{m}\cap 15\text{m}}$ (incremental 15m-contract NTM in overlapping cycles); the stars test each coefficient against zero. Groups use set notation: 5m \ 15m = non-overlapping cycles (only the 5m contract settles), 5m ∩ 15m = overlapping cycles whose close coincides with a 15-min mark (the 5m and 15m contracts settle simultaneously). Within each, cycles are split by per-contract Polymarket NTM, defined by the Up-token price near TTM ∈ [10, 12] s: $\text{NTM}_{5\text{m}}^{\text{poly}}$ iff $|p_{Up,5\text{m}} - 0.5| < 0.10$, $\text{NTM}_{15\text{m}}^{\text{poly}}$ iff $|p_{Up,15\text{m}} - 0.5| < 0.10$. The nine overlapping cycles NTM by *both* contracts load on both the $\text{NTM}_{5\text{m}}$ and $\text{NTM}_{15\text{m}}$ overlap increments — their slope is constrained to the additive sum (unlike the saturated activity Table A1; nine cycles cannot support a free interaction slope). Date, hour-of-day, and cycle-group fixed effects; standard errors in parentheses are clustered by date. Sample is P3 cycles with the relevant Polymarket classifications available: 5m \ 15m cycles need the 5m classifier, 5m ∩ 15m cycles need both. Negative j = same-cycle bin $(29 + j)$; positive j = bin $(j - 1)$ of the next cycle. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Pre-close			Post-close		
	$j = -3$	$j = -2$	$j = -1$	$j = +1$	$j = +2$	$j = +3$
5m \ 15m, FFM (baseline)	0.001	0.031	0.052***	−0.083***	−0.028	0.019
	(0.015)	(0.019)	(0.015)	(0.028)	(0.022)	(0.014)
$\delta_{5\text{m}\cap 15\text{m}}$	−0.027	−0.023	−0.026	−0.078**	−0.039	−0.011
	(0.022)	(0.025)	(0.021)	(0.033)	(0.037)	(0.028)
$\delta_{\text{NTM}_{5\text{m}},\,5\text{m}\setminus 15\text{m}}$	−0.020	−0.019	0.007	−0.163***	−0.003	−0.053
	(0.022)	(0.032)	(0.028)	(0.044)	(0.039)	(0.039)
$\delta_{\text{NTM}_{5\text{m}},\,5\text{m}\cap 15\text{m}}$	0.035	−0.012	−0.072**	−0.225**	0.062	0.035
	(0.032)	(0.032)	(0.035)	(0.097)	(0.069)	(0.040)
$\delta_{\text{NTM}_{15\text{m}},\,5\text{m}\cap 15\text{m}}$	0.068	−0.020	0.009	−0.216**	0.115	−0.064
	(0.045)	(0.043)	(0.054)	(0.106)	(0.071)	(0.072)
Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Hour-of-day FE	Yes	Yes	Yes	Yes	Yes	Yes
Cycle-group FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations			13,507			
Clusters (UTC date)			56			

35% more close-time activity — both significant at the 5% level, but no effect on depth. These magnitudes are a small fraction of the five-minute effect in P3: differencing the log coefficients, the P3 NTM order-flow effect of Table 4 is about $e^{1.37-0.24} \approx 3.1\times$ larger, and P3 carries the depth and reversal signatures that P2 lacks. Panel B shows the reversal channel is absent altogether: the FFM baseline reverts at the first post-close lag (the general overlap reversal, present once the fifteen-minute contract is live), but the NTM increment $\delta_{\text{NTM},+1}$ is small, wrong-signed, and far from significant. The thin NTM sample (about 230 cycles) is the binding constraint on Panel B. Taken together, the fifteen-minute contract alone generates at most a modest moneyness-graded

Table A3. Fifteen-minute-contract moneyness in P2 (BTC). This table reports near-settlement activity (Panel A) and post-settlement reversal (Panel B) in P2 overlapping cycles, split by fifteen-minute-contract Polymarket NTM. Sample: P2 cycles (Oct 9, 2025 – Feb 11, 2026; fifteen-minute and four-hour contracts live, five-minute not yet) whose close coincides with a 15-min mark, classified NTM iff the fifteen-minute-contract Up-token price is within 0.10 of 0.5 near TTM $\in [10, 12]$ s, FFM otherwise. *Panel A:* the bin-29 NTM activity increase δ_{NTM} (the coefficient on $\mathbf{1}\{b{=}29\} \cdot \mathbf{1}\{\text{NTM}\}$) for each metric, log-OLS on the (cycle, bin) panel restricted to bins {0–24, 29} with date, hour-of-day, and bin fixed effects; standard errors clustered at the cycle level. Metrics as in Table 4; coefficients are log differences, so δ is a multiplicative e^{δ} effect. *Panel B:* the post-settlement reversal slope by lag j, the FFM baseline slope γ_j on r_{29} and the NTM increment $\delta_{\text{NTM},j}$, with date, hour-of-day, and group fixed effects; standard errors clustered by date; a negative slope is reversal. Negative j = same-cycle bin $(29 + j)$; positive j = bin $(j - 1)$ of the next cycle. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Near-settlement activity

	Order flow	Abs. return	Depth
δ_{NTM}	0.243**	0.303**	−0.015
	(0.116)	(0.144)	(0.017)
Date FE	Yes	Yes	Yes
Hour-of-day FE	Yes	Yes	Yes
Bin FE	Yes	Yes	Yes
Observations	158,570	102,329	158,570
NTM cycles		229	
FFM cycles		5,870	

Panel B: Post-settlement reversal

	Pre-close			Post-close		
	$j = -3$	$j = -2$	$j = -1$	$j = +1$	$j = +2$	$j = +3$
FFM γ_j	0.121	−0.026	0.015	−0.137***	−0.011	−0.029
	(0.108)	(0.036)	(0.028)	(0.045)	(0.051)	(0.037)
$\delta_{\text{NTM},\,j}$	−0.082	−0.129*	−0.004	0.076	−0.184	0.011
	(0.115)	(0.077)	(0.061)	(0.234)	(0.115)	(0.071)
Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Hour-of-day FE	Yes	Yes	Yes	Yes	Yes	Yes
Cycle-group FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations			6,069			
Clusters (UTC date)			125			

activity increase, with neither the liquidity withdrawal-versus-provision pattern nor the deeper NTM reversal that define the P3 footprint. Even in the cycles where a push could flip the fifteen-minute resolution, the footprint barely appears: the fifteen-minute window is too long for the push to be reliably pivotal, and the near-settlement effect switches on with the five-minute contract.

Table A4. Binance spot midquote vs. the Chainlink resolution price, BTC. This table compares the Binance midquote with the Chainlink price that settles each market, one observation per Polymarket BTC market with both the Chainlink strike (fixed at window open) and the Chainlink settlement print (at window close) known. *Median* $Basis_t$ is the median per-market (Chainlink − Binance mid)/Binance mid at window open (bps). *SD* $\Delta Basis_t$ is the standard deviation of the open-to-close change $\Delta Basis_t \equiv Basis_{\text{close}} - Basis_{\text{open}}$ (bps). *Strike-side agreement* is the fraction of markets on which the Binance mid finishes on the same side of the strike as the realized resolution, $\text{sign}(\text{mid}_{\text{close}} - \text{strike}) = \text{sign}(\text{settlement} - \text{strike})$.

Interval	n	Median $Basis_t$	SD $\Delta Basis_t$	Strike-side agreement
5m	5,553	−2.45	1.44	84.8%
15m	1,855	−2.49	2.24	91.3%

D Binance Spot Tracks the Chainlink Resolution Price

The 5m / 15m / 4h crypto up/down contracts resolve on a Chainlink VWAP price, not on Binance. Every Binance-based result in the main body therefore rests on the Binance midquote being a faithful proxy for the price that actually settles each market. We check this directly on the subset of BTC markets for which the Chainlink strike (fixed at window open) and the Chainlink settlement print (at window close) are both known; coverage is concentrated in the Mar 16 – Apr 10, 2026 window. For each such market we read the Binance midquote at the window open and close from the 10-second midquote series via a strictly causal match: each bucket stores the *last* book mid in $[t, t + 10\,\text{s})$, so its effective timestamp is the bucket end, and we match each window boundary to the last book mid of the 10-second bucket ending at or before that instant.

Define, per market, the Chainlink–Binance basis at time t (with mid_t the Binance midquote):

$$Basis_t = (\text{Chainlink}_t - \text{mid}_t)/\text{mid}_t.$$

The market's realized Up/Down outcome is sign(settlement − strike); the directional call a manipulator targets is whether the Binance mid finishes above or below the strike, $\text{sign}(\text{mid}_{\text{close}} - \text{strike})$, the side a final-10s push tries to land on.

Table A4 reports the comparison. The basis is a stable level offset of ≈ -2.5 bps and is essentially flat across horizons, so it cancels in the open-to-close return comparison: the within-window change

in the basis has standard deviation of only $\approx$ 1.4–2.2 bps, small relative to a typical window move (5.8 bps at 5m, 10.8 bps at 15m), so within each window the Binance mid and the Chainlink price move essentially together.

Where Binance and Chainlink *can* part company is right at the strike. Because Binance sits $\approx$ 2.5 bps above Chainlink, the Binance mid can finish just above the strike while Chainlink, and hence the resolution, is still below it. The strike-side agreement in the last column is therefore below one: 84.8% at 5m and 91.3% at 15m, rising with horizon as the typical move grows relative to the fixed $\approx$ 2.5 bps basis. The residual disagreement is confined to genuine coin-toss windows: sorting the BTC 5m markets into quintiles by the size of the Binance window move, agreement is 61% in the smallest quintile ($|\text{move}| < 1.8$ bps), 87% in the middle, and 100% once the move exceeds $\approx$ 13 bps. That $\approx$ 2.5 bps band is precisely the margin a spot push must clear: to flip the resolution a manipulator must push Binance *past* the strike by more than the basis, not merely to it; a push that only reaches the strike still fails on the $\sim$ 15% of 5m windows where the basis leaves Chainlink on the other side. A push well past the strike, however, reliably carries the resolution, which makes spot-market pressure a coherent channel for moving the prediction-market outcome.

E Manipulator Cohort Threshold Sensitivity

The manipulator cohort definition in Section 6.4.4 uses two ad hoc thresholds: a minimum manipulated-cycle participation count of 5 and a minimum aggregate manipulated-cycle PnL of \$2,000. Table A5 re-derives the *full* trader-type attribution at each cell of a 4×5 grid of $(n_{\text{cycle,min}}, \$_{\text{min}})$ thresholds: the manipulator cohort is recomputed, the MM cohort is re-derived as the behavioral residual (same passive-rate, fills-per-active-cycle, and inventory criteria as the main text), and retail/other is the final residual, so the entire partition responds to the threshold. For each cell the three panels report the three cohort metrics of Table 10: net PnL (Panel A), the fraction of cycles in which each cohort is net negative (Panel B), and the pool-weighted loss share (Panel C), each split across manipulated and normal cycles. Within each regime the three cohort PnLs sum to zero and the three loss shares to 100% by construction.

All three metrics are stable across the twenty cells. The manipulated-cycle wealth transfer runs directly from retail to manipulators: manipulators net between +\$4.1 M and +\$10.0 M and retail is the near-exact mirror (−\$3.6 M to −\$9.3 M); retail is on the losing side of 64–66% of manipulated cycles and bears 62–74% of the manipulated-cycle loss pool at every threshold. The MM cohort is a small bystander in manipulated cycles: its net PnL holds near −\$0.6 M throughout and its loss share stays modest (8–16%). Its behavioral profile — passive, delta-neutral quoting — is orthogonal to the manipulator's directional-PnL threshold, so moving the threshold barely reclassifies it. In normal cycles the manipulator cohort's PnL is near zero at moderate cells and only mildly negative at the tightest activity cuts ($n_{\text{cycle,min}} \geq 10$), confirming that manipulators do not systematically profit in normal cycles. The core finding — a large, retail-funded manipulated-cycle profit for the manipulator cohort — is robust to both threshold choices.

Table A5. Trader-type attribution sensitivity to manipulator thresholds. This table re-derives the full trader-type partition at each cell of a 4×5 threshold grid: the manipulator cohort is the set of wallets with $n_{\text{cycle,p}} \geq n_{\text{cyc,min}}$ and net manipulated-cycle PnL $\geq \$_{\text{min}}$; the MM cohort is the behavioral residual passing the same passive-rate, fills-per-active-cycle, and inventory criteria as the main text; retail/other is the final residual. Each panel reports one trader-type metric of Table 10 for all three cohorts, split across manipulated and normal cycles: *Panel A*, aggregate net PnL ($M), plus the manipulator cohort size (Wallets); *Panel B*, the fraction of cycles in which the cohort is net negative; *Panel C*, the cohort's share of the gross loss pool, aggregated across cycles with pool-size weights. Within each regime the three cohort PnLs sum to zero and the three loss shares to 100%.

Panel A. Net PnL ($ M)

			Manipulated cycles ($M)			Normal cycles ($M)		
$n_{\text{cyc,min}}$	$\$_{\text{min}}$	Wallets	Manip.	MM	Retail	Manip.	MM	Retail
3	$500	2,276	+9.98	−0.65	−9.33	−0.92	+3.01	−2.08
3	$1,000	1,456	+9.41	−0.64	−8.76	−0.19	+3.02	−2.83
3	$2,000	857	+8.56	−0.62	−7.94	+0.28	+3.11	−3.39
3	$5,000	414	+7.27	−0.59	−6.68	+0.68	+3.21	−3.90
3	$10,000	221	+5.89	−0.55	−5.34	+0.31	+3.40	−3.70
5	$500	2,158	+9.56	−0.65	−8.91	−0.94	+3.01	−2.07
5	$1,000	1,389	+9.03	−0.64	−8.39	−0.30	+3.02	−2.72
5	$2,000	821	+8.22	−0.62	−7.61	+0.09	+3.11	−3.20
5	$5,000	397	+6.99	−0.59	−6.40	+0.52	+3.21	−3.73
5	$10,000	210	+5.65	−0.55	−5.10	+0.16	+3.40	−3.56
10	$500	1,949	+8.63	−0.65	−7.98	−1.43	+3.01	−1.58
10	$1,000	1,275	+8.16	−0.64	−7.52	−0.89	+3.02	−2.13
10	$2,000	751	+7.41	−0.62	−6.79	−0.53	+3.11	−2.58
10	$5,000	356	+6.27	−0.59	−5.68	−0.04	+3.21	−3.18
10	$10,000	181	+5.01	−0.55	−4.46	−0.31	+3.40	−3.09
20	$500	1,673	+7.37	−0.65	−6.72	−1.43	+3.01	−1.58
20	$1,000	1,130	+6.99	−0.64	−6.35	−1.05	+3.02	−1.97
20	$2,000	672	+6.33	−0.62	−5.71	−0.75	+3.11	−2.36
20	$5,000	315	+5.30	−0.59	−4.71	−0.33	+3.21	−2.88
20	$10,000	151	+4.12	−0.55	−3.57	−0.54	+3.40	−2.86

Table A5. Trader-type attribution sensitivity to manipulator thresholds *(continued).*

Panel B. % losing cycles

		Manipulated cycles			Normal cycles		
$n_{\text{cyc,min}}$	$\$_{\text{min}}$	Manip.	MM	Retail	Manip.	MM	Retail
3	\$500	35.8%	58.8%	66.3%	62.2%	37.6%	43.8%
3	\$1,000	36.8%	58.8%	66.0%	61.8%	37.7%	45.4%
3	\$2,000	38.1%	58.6%	65.4%	62.3%	37.7%	47.7%
3	\$5,000	38.2%	58.7%	65.7%	60.5%	37.5%	53.4%
3	\$10,000	40.2%	57.6%	65.6%	66.6%	37.2%	55.7%
5	\$500	36.0%	58.8%	65.8%	62.3%	37.6%	43.7%
5	\$1,000	36.9%	58.8%	65.7%	61.9%	37.7%	45.4%
5	\$2,000	38.0%	58.6%	65.3%	62.4%	37.7%	47.7%
5	\$5,000	38.2%	58.7%	65.5%	60.6%	37.5%	53.4%
5	\$10,000	40.3%	57.6%	65.2%	66.6%	37.2%	55.7%
10	\$500	36.7%	58.8%	65.0%	62.3%	37.6%	43.8%
10	\$1,000	37.2%	58.8%	65.1%	62.0%	37.7%	45.4%
10	\$2,000	38.4%	58.6%	64.8%	62.6%	37.7%	47.9%
10	\$5,000	38.6%	58.7%	65.2%	60.7%	37.5%	53.3%
10	\$10,000	40.7%	57.6%	64.8%	66.9%	37.2%	55.5%
20	\$500	37.2%	58.8%	64.4%	62.3%	37.6%	44.7%
20	\$1,000	37.3%	58.8%	64.4%	61.9%	37.7%	46.3%
20	\$2,000	38.7%	58.6%	64.4%	62.5%	37.7%	48.4%
20	\$5,000	38.5%	58.7%	65.0%	60.9%	37.5%	53.5%
20	\$10,000	40.7%	57.6%	64.3%	67.0%	37.2%	55.6%

Panel C. Loss share

		Manipulated cycles			Normal cycles		
$n_{\text{cyc,min}}$	$\$_{\text{min}}$	Manip.	MM	Retail	Manip.	MM	Retail
3	\$500	17.8%	8.4%	73.8%	46.3%	9.1%	44.6%
3	\$1,000	18.7%	8.6%	72.6%	44.9%	9.6%	45.5%
3	\$2,000	19.7%	9.1%	71.2%	43.1%	10.6%	46.3%
3	\$5,000	20.3%	10.3%	69.4%	40.0%	12.4%	47.5%
3	\$10,000	20.1%	12.5%	67.3%	37.0%	15.3%	47.7%
5	\$500	18.4%	8.6%	73.0%	46.2%	9.3%	44.5%
5	\$1,000	19.3%	8.8%	71.9%	44.8%	9.8%	45.4%
5	\$2,000	20.2%	9.3%	70.5%	43.2%	10.8%	46.0%
5	\$5,000	20.7%	10.6%	68.7%	40.1%	12.7%	47.3%
5	\$10,000	20.6%	12.8%	66.6%	37.1%	15.5%	47.4%
10	\$500	19.1%	9.2%	71.6%	46.5%	9.9%	43.6%
10	\$1,000	19.9%	9.5%	70.5%	45.3%	10.4%	44.3%
10	\$2,000	20.8%	10.1%	69.1%	43.6%	11.6%	44.8%
10	\$5,000	21.5%	11.4%	67.1%	40.4%	13.5%	46.1%
10	\$10,000	21.5%	13.9%	64.7%	37.3%	16.3%	46.3%
20	\$500	19.8%	10.5%	69.7%	45.9%	10.7%	43.4%
20	\$1,000	20.5%	10.8%	68.7%	44.9%	11.2%	43.9%
20	\$2,000	21.3%	11.5%	67.2%	43.3%	12.4%	44.3%
20	\$5,000	22.2%	13.0%	64.8%	40.2%	14.3%	45.4%
20	\$10,000	22.2%	16.1%	61.7%	37.2%	17.1%	45.7%

F Push-Down Cohort Flow

The cohort-flow figure in Section 6.4.6 (Figure 12) conditions on *push-up* cycles, manipulated cycles whose final-10s spot push is upward, so that the directional accumulation reads as a positive ramp. This appendix repeats the same construction on the *push-down* mirror (manipulated cycles whose push is downward) in Figure A3.

The pattern is the sign-flipped mirror with the identical timing signature. The manipulator cohort ramps smoothly into a net *short* Up position (buying the Down token), reaching $\sim -9{,}500$ shares by close; MMs and retail take the long-Up side; and the Binance order flow is flat through the body and then steps down by $\sim$ \$1.75 M in the final 10 s (cumulative $\sim$ $-$\$2 M by close), the down-side counterpart of the +\$1.7 M up-push. The mirror is not exact in magnitude: both the final-10s Binance step ($\sim$ $-$\$1.75 M vs $\sim$ +\$1.7 M) and the manipulator cohort's cumulative directional position ($\sim -9{,}500$ vs $\sim$ +8,800 shares) are marginally larger on the down side. But the linear-ramp-then-late-step timing is the same in both directions, confirming that the timing argument against the gamma-hedging channel (Section 6.4.6) is not an artifact of conditioning on the up direction.

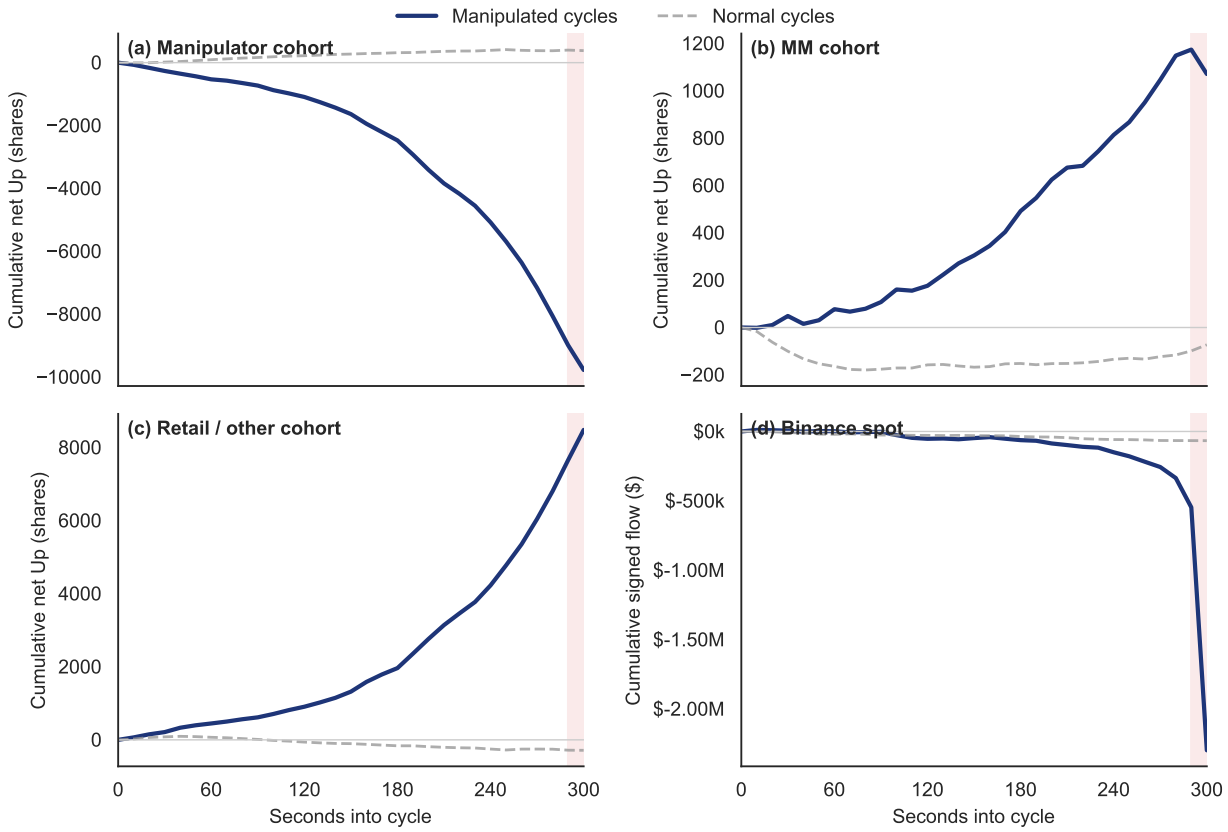

Figure A3. Cumulative net position trajectory across the cycle, push-down vs. normal cycles, P3 BTC 5m. This figure plots the push-down counterpart of Figure 12; the solid line conditions on push-down cycles (labeled "manipulated cycles" in the legend), construction otherwise identical. Panels **(a)**–**(c)** plot the net Up-token holdings (shares) of the manipulator, MM, and retail/other cohorts; panel **(d)** plots the cumulative Binance spot order flow ($). Dashed gray lines: normal cycles. Pink shading: the final-10 s region (seconds 290–300).

References

Akey, Grégoire, Harvie, and Martineau (2026). *Who Wins and Who Loses in Prediction Markets? Evidence from Polymarket*. Working Paper 6443103. SSRN.

Aoyagi, Jun and Yuki Ito (2025). "Coexisting Exchange Platforms: Limit Order Books and Automated Market Makers". In: *Journal of Political Economy Microeconomics*.

Arrow, Kenneth J. et al. (2008). "The Promise of Prediction Markets". In: *Science* 320.5878, pp. 877–878.

Aspembitova and Bentley (2022). "Oracles in Decentralized Finance: Attack Costs, Profits and Mitigation Measures". In: *Entropy* 24.1, p. 60.

Baltussen, Guido, Julian Terstegge, and Paul Whelan (2023). *The Derivative Payoff Bias*. Working Paper 4562800. SSRN.

Barbon, Andrea and Angelo Ranaldo (2026). "On the Quality of Cryptocurrency Markets: Centralized versus Decentralized Exchanges". In: *Management Science*.

Bartlett, Robert P. and Maureen O'Hara (2026). *Adverse Selection in Prediction Markets: Evidence from Kalshi*. Working Paper 6615739. SSRN.

Capponi, Agostino and Ruizhe Jia (2025). "Liquidity Provision on Blockchain-Based Decentralized Exchanges". In: *Review of Financial Studies* 38.10, pp. 3040–3085.

Capponi, Agostino, Ruizhe Jia, and Kanye Ye Wang (2025). "Maximal Extractable Value and Allocative Inefficiencies in Public Blockchains". In: *Journal of Financial Economics* 172, p. 104132.

Capponi, Agostino, Ruizhe Jia, and Shihao Yu (2026). "Price Discovery on Decentralized Exchanges". In: *Review of Financial Studies*. Forthcoming; SSRN 4236993.

Chiu, Jonathan and Thorsten V. Koeppl (2019). "Blockchain-Based Settlement for Asset Trading". In: *Review of Financial Studies* 32.5, pp. 1716–1753.

Comerton-Forde, Carole and Tālis J. Putniņš (2011). "Measuring Closing Price Manipulation". In: *Journal of Financial Intermediation* 20.2, pp. 135–158.

Commodity Futures Trading Commission (Mar. 2026). *Staff Advisory on Event Contracts*. CFTC Staff Letter 26-08. Commodity Futures Trading Commission.

Cong, Prasad, and Rabetti (2025). *Financial and Informational Integration Through Oracle Networks*. NBER Working Paper 33639. National Bureau of Economic Research.

Daian et al. (2020). "Flash Boys 2.0: Frontrunning, Transaction Reordering, and Consensus Instability in Decentralized Exchanges". In: *IEEE Symposium on Security and Privacy (S&P)*.

Della Vedova (2026). *Who Profits from Prediction Markets? Execution, not Information*. Working Paper 6191618. SSRN.

Diercks, Katz, and Wright (2026). *Kalshi and the Rise of Macro Markets*. NBER Working Paper 34702. National Bureau of Economic Research.

Dubach (2026). *The Anatomy of a Decentralized Prediction Market: Microstructure Evidence from the Polymarket Order Book*. Working Paper arXiv:2604.24366. arXiv.

Duffie, Darrell and Piotr Dworczak (2021). "Robust Benchmark Design". In: *Journal of Financial Economics* 142.2, pp. 775–802.

Easley, O'Hara, Yang, and Zhang (2026). "Microstructure and Market Dynamics in Crypto Markets". In: *Journal of Financial Markets*. SSRN 4814346.

Glosten, Lawrence R. and Paul R. Milgrom (1985). "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders". In: *Journal of Financial Economics* 14.1, pp. 71–100.

Gómez-Cram, Guo, Jensen, and Kung (2026). *Prediction Market Accuracy: Crowd Wisdom or Informed Minority?* Working Paper 6617059. SSRN.

Griffin, John M. and Amin Shams (2018). "Manipulation in the VIX?" In: *Review of Financial Studies* 31.4, pp. 1377–1417.

Hanson, Robin and Ryan Oprea (2009). "A Manipulator Can Aid Prediction Market Accuracy". In: *Economica* 76.302, pp. 304–314.

Hanson, Robin, Ryan Oprea, and David Porter (2006). "Information Aggregation and Manipulation in an Experimental Market". In: *Journal of Economic Behavior and Organization* 60.4, pp. 449–459.

Hasbrouck, Joel, Thomas J. Rivera, and Fahad Saleh (2024). "An Economic Model of a Decentralized Exchange with Concentrated Liquidity". In: *Management Science*.

Hasbrouck, Joel, Fahad Saleh, and Thomas Rivera (2026). "The Need for Fees at a DEX: How Increases in Fees Can Increase DEX Trading Volume". In: *Management Science*. Accepted.

Klein, Frederik, Roman Kozhan, Ganesh Viswanath-Natraj, and Junxuan Wang (2023). *Price Discovery in Cryptocurrencies: Trades versus Liquidity Provision*. Working Paper. SSRN 4642411.

Kumar, Praveen and Duane J. Seppi (1992). "Futures Manipulation with "Cash Settlement"". In: *Journal of Finance* 47.4, pp. 1485–1502.

Lehar, Alfred and Christine A. Parlour (2025). "Decentralized Exchange: The Uniswap Automated Market Maker". In: *Journal of Finance* 80.1, pp. 321–374.

Lehar, Alfred, Christine A. Parlour, and Marius Zoican (2022). *Liquidity Fragmentation on Decentralized Exchanges*. Working Paper. SSRN 4267429.

Makarov, Igor and Antoinette Schoar (2020). "Trading and Arbitrage in Cryptocurrency Markets". In: *Journal of Financial Economics* 135.2, pp. 293–319.

Mitts and Ofir (2026). *From Iran to Taylor Swift: Informed Trading in Prediction Markets*. Working Paper 6426778. SSRN.

Ni, Sophie X., Neil D. Pearson, and Allen M. Poteshman (2005). "Stock Price Clustering on Option Expiration Dates". In: *Journal of Financial Economics* 78.1, pp. 49–87.

Park, Andreas (2023). "The Conceptual Flaws of Decentralized Automated Market Making". In: *Management Science* 69.11, pp. 6731–6751.

Polymarket (Aug. 7, 2026). *Chainlink TWAP Resolution for Crypto Up/Down Markets*. Predictions Changelog. URL: https://docs.polymarket.com/changelog/predictions (visited on 09/02/2026).

Rhode, Paul W. and Koleman S. Strumpf (2008). *Manipulating Political Stock Markets: A Field Experiment and a Century of Observational Data*. NBER Working Paper 14377. National Bureau of Economic Research.

Spatt, Chester (2014). "Security market manipulation". In: *Annu. Rev. Financ. Econ.* 6.1, pp. 405–418.

Tsang, Kwok Ping and Zichao Yang (2026). *The Anatomy of a Blockchain Prediction Market: Polymarket in the 2024 U.S. Presidential Election*. Working Paper arXiv:2603.03136. arXiv.

Wolfers, Justin and Eric Zitzewitz (2004). "Prediction Markets". In: *Journal of Economic Perspectives* 18.2, pp. 107–126.

Zhang, Anthony Lee (2022). "Competition and Manipulation in Derivative Contract Markets". In: *Journal of Financial Economics* 144.2, pp. 396–413.

Zhu, Haoxiang (2014). "Do Dark Pools Harm Price Discovery?" In: *Review of Financial Studies* 27.3, pp. 747–789.

Information in Fixed-Price Trading

Patrick Blonien* Thomas Ernst† Alexander Ober‡

October 4, 2026

ABSTRACT

Fixed-price venues, such as dark pools, are commonly thought to hinder the incorporation of information into prices. We show that fixed-price venues can produce information beyond what is contained in prices. In a dynamic model of strategic traders with heterogeneous information, cross-period price impact from a fixed-price session deters informed trading on signals at the fixed price, while inventory-based demand is not deterred. Aggregate quantities help to reveal inventories, which the market-clearing price does not fully separate from the fundamental. This information reduces adverse selection and improves price informativeness in subsequent price-clearing sessions. Anticipating the fixed-price session, traders delay trade, reducing liquidity, but improved inventory inference can outweigh this cost, increasing welfare when information is sufficiently asymmetric and the market is large.

*David A. Tepper School of Business, Carnegie Mellon University; e-mail: pblonien@andrew.cmu.edu.
†Robert H. Smith School of Business, University of Maryland.
‡Kellogg School of Management, Northwestern University.

I. Introduction

Prices aggregate information. This is the fundamental insight of Hayek (1945), who argues that market prices communicate dispersed information, and of Grossman (1976), who formalizes how equilibrium prices reflect investors' heterogeneous private information. It is natural to then think that a trading mechanism without a price of its own cannot reveal any further information. Following this logic, practitioners, regulators, and academics have long worried that trading venues where prices are fixed, such as dark pools[1], detract from the information production that prices provide in lit markets (e.g., Securities and Exchange Commission, 2010, Degryse et al., 2014, Hatheway et al., 2017, Buti et al., 2017). The studies that find dark trading need not harm price discovery rely on sorting: dark pools draw the uninformed away from the exchange, so the lit price becomes more informative because of who leaves it, not because dark volume carries information of its own (Zhu, 2013, Comerton-Forde and Putniņš, 2015). In both views, the fixed-price trading venue itself produces no information. We show that it can, and we study its dynamic implications for exchange trade.

A fixed price rules out learning from the price, but not learning from trade in that type of trading mechanism. By observing the quantity that is allocated in the session, each trader can infer the aggregate quantity demanded. Traders therefore learn the information contained in the quantity that clears at the fixed price. The key mechanism is that, when traders are large, rational, and strategic, trades affect beliefs and therefore future prices. So, while there is no contemporaneous price impact, by construction of trading at a fixed price, there is cross-period (intertemporal) price impact, since traders update their beliefs after observing the aggregate quantity demanded in the fixed-price session.

This cross-period price impact is large. Because there is no contemporaneous price impact, a trader is tempted to deviate by trading more aggressively on their private information than the others expect. The other traders still believe that everyone is playing the equilibrium strategy, so the deviation distorts their beliefs. In equilibrium, the cross-period price impact that would follow such a deviation must be large enough that trading at worse prices in the future costs the trader more than they gain from trading aggressively in the fixed-price session. As a result, the informed trade little on their signals in the fixed-price session. Inventory-driven trade is not deterred in the same way, because traders do not unwind it in

[1]Other examples of fixed-price trading mechanisms are workups in the interdealer Treasury and index CDS markets and matching sessions in corporate bonds and swaps.

the next session, so a price that moves against them costs them little. The aggregate quantity is therefore disproportionately a signal about inventories. The market-clearing price, on the other hand, places relatively more weight on the fundamental value of the asset. Information regarding counterparties' inventories is useful because it reduces adverse selection in future price-clearing sessions. Traders face less price impact in subsequent price-clearing sessions, because large demands are more likely to be attributed to offloading large inventories than to exploiting large private signals.

To formalize this argument, we study a dynamic model of strategic trade in which large traders differ in the quality of their information and a fixed-price session may follow a price-clearing session. We study a model with two types of traders, informed and uninformed, who trade to manage inventories and differ in the precision of their private signals about the value of the asset. Trade takes place over two periods. Traders first trade in a price-clearing session, which may be followed by a fixed-price session that executes at the prior market-clearing price, and then trade again in a second price-clearing session. Dynamic models of imperfect competition typically abstract from heterogeneous information (e.g., Vayanos, 1999, Du and Zhu, 2017, Antill and Duffie, 2020, Chen and Duffie, 2021). With heterogeneous information, beliefs depend on the entire history of prices, and a trader's demand today changes what others infer in every future session, so learning is non-Markovian, and the equilibrium cannot be found by standard Bellman equations (Rostek et al., 2025). The fixed-price session adds to this difficulty. We solve the model by extending the non-recursive approach of Rostek et al. (2025), in which the cross-period price and demand impacts are constants in a linear-quadratic setting and are determined jointly as a fixed point.

The model yields three main results. First, the aggregate quantity submitted in the fixed-price session is informative relative to the prevailing market-clearing price. Given the market-clearing price, the aggregate quantity submitted is almost entirely a signal about inventories. In our baseline parameterization, it reduces the uninformed's uncertainty about the value of the asset by roughly 80%. This reduction in uncertainty reduces adverse selection and price impact in the subsequent price-clearing session. Second, the fixed-price session does introduce some inefficiencies. Traders who anticipate the fixed-price session delay part of their inventory adjustment by shading their demands in the price-clearing session, which lowers liquidity there whether or not the fixed-price session takes place (Antill and Duffie, 2020, Blonien, 2026). Third, the information acquired in the fixed-price session can outweigh

the cost of this delay, making the addition of a fixed-price session welfare improving. The benefits and costs are shared by different traders. The uninformed gain because they learn from the aggregate quantity, while the informed lose part of their trading profits. The fixed-price session raises aggregate welfare when the uninformed's gains outweighs the informed's losses. This requires sufficient asymmetric information for learning from quantities to reduce, and a market large enough that the strategic delay in price-clearing sessions preceding the fixed-price session is small.

Related Literature

Our paper contributes to three strands of the literature. The first studies how trade is affected by the addition of fixed-price trading sessions, sometimes referred to as size discovery sessions. Duffie and Zhu (2017) show that a fixed-price session held before trading begins improves allocative efficiency, and Antill and Duffie (2020) show that when trade is continuous and a fixed-price session may occur in the future, traders delay trade in anticipation, which lowers allocative efficiency. Blonien (2026) shows that the net effect depends on the frequency of trade and the degree of adverse selection. In each of these papers, the fixed-price session reallocates inventory but reveals no information relative to the prevailing market-clearing price. These studies only have a single type of trader who is ex-ante identical, so the price already reveals everything traders can learn from each other. We show that when traders differ in the quality of their information, the aggregate quantity is informative, which gives the fixed-price session a further welfare benefit. Empirically, the quantities traded at a fixed price do contain information. Fleming and Nguyen (2018) find that shocks to order flow in U.S. Treasury workups (fixed-price sessions) explain 6–8% of the variation in returns on benchmark notes and about 10% of the variation in the level of the yield curve. They also find that the information in workup order flow differs from the information in the trades that set the price: information related to dealers' proprietary client order flow is more likely to show up in workup trades, while information from public announcements is more likely to show up in pre-workup trades. This is broadly consistent with the separation of trade into fundamental and non-fundamental components that we find.

The second strand studies the information content that market participants learn from quantities. Blume et al. (1994) and Schneider (2009) show that trading volume carries information that cannot be deduced from the price, because it reveals the precision of traders'

information. Ganguli and Yang (2009) and Manzano and Vives (2011) study markets in which the price reflects a second source of variation, the supply of the asset or traders' exposure to aggregate risk, about which traders are privately informed. In our model, the second source of variation is inventories, and the aggregate quantity is a public signal about them that lets traders separate the two components, fundamental and non-fundamental, of the information contained in the price. This signal arises endogenously from trade at a fixed price, because the cross-period price impact deters the informed from trading on their signals in the fixed-price session. Our mechanism also differs from the literature on the choice between exchanges and fixed-price venues (Hendershott and Mendelson, 2000, Zhu, 2013, Buti et al., 2017, Menkveld et al., 2017). In Zhu (2013), a dark pool improves price discovery because the informed self-select onto the exchange. In our model, the sessions are sequential, every trader participates in both, and the information arises from the quantity submitted at the fixed price.

The third strand studies dynamic trade under imperfect competition. Strategic traders split their orders over time and across venues to reduce price impact, which can reduce liquidity in any one session (Vayanos, 1999, Du and Zhu, 2017, Sannikov and Skrzypacz, 2016, Chen and Duffie, 2021, Blonien and Ober, 2026). With private information, traders also learn from prices, and price impact reflects adverse selection (Kyle, 1989, Vives, 2011, Rostek and Weretka, 2012). Dynamic models with traders who differ in the quality of their information are difficult to solve, because a trade today changes what others infer in every future session. Rostek et al. (2025) develop a non-recursive approach to this problem and show that making a statistic of past trades public introduces dynamic adverse selection, which can raise price impact and lower welfare. We apply their approach to a market in which the public statistic is generated by a fixed-price session, and we show that intertemporal price impact modifies traders' demands in fixed-price sessions, which implies the aggregate quantity is informative.

II. The Model

In this section, we set up a dynamic model of trade among a finite number of rational large traders. Because trading behavior today affects future beliefs, the resulting equilibrium is neither stationary nor Markovian (Rostek et al. (2025)). The way trading behavior affects future beliefs will also differ by trading mechanism. We study a market with two types of traders who both experience inventory shocks, but the informed traders also observe a noisy

private signal of a dividend's liquidation value, while uninformed traders do not. We assume there are three trading sessions across two periods.[2]

Figure 1 shows the setup of the model. There are two periods during which trade occurs, denoted period 1 and period 2. As we will see later, there will be potentially two successive trading opportunities in period 1. Each period, traders exchange a divisible asset that pays a dividend of v_t upon liquidation per unit of the asset. Before trade in the first period, the liquidating value is uncertain and equal to $\tilde{v}_1 \sim N(0, \sigma^2)$.[3] Before the second period of trade, the liquidating value will be shocked to $\tilde{v}_2 \sim N(v_1, \sigma_\eta^2)$.

There are two types of traders. There are informed traders ($N \in \mathbb{Z}_{\geq 0}$) who receive private independent noisy signals about the liquidation value of the asset at a given time, while uninformed traders ($M \in \mathbb{Z}_{\geq 0}$) do not, but will learn about the liquidation value through market outcomes. Before trade each period, the private noisy signal informed traders receive is $S_{i,t} = v_t + \epsilon_{i,t}$, where $\epsilon_{i,t} \overset{\text{iid}}{\sim} N(0, \sigma_\epsilon^2)$. Throughout, we write $\tau \in \{I, U\}$ for a trader's type, with $-\tau$ denoting the opposite type. For trade to occur in the equilibria we consider, we require $N + M > 2$.

In the model, traders dynamically manage inventory positions. Each trader, i, has an inventory position at the beginning of period t of z_t^i. Before each trading session, each trader receives a random endowment shock to their inventory, denoted $h_t^i \sim N(0, \sigma_h^2)$. For simplicity, assume all traders have an initial inventory of zero before the first shock. In each period, traders engage in trade to manage their inventory positions z_t^i. Traders are risk-neutral but face quadratic costs for holding inventory, representing risk management and collateral considerations. Chen and Duffie (2021), Antill and Duffie (2020), Duffie and Zhu (2017), Du and Zhu (2017), Sannikov and Skrzypacz (2016), Rostek and Weretka (2012), Vives (2011), Blonien and Ober (2026), Blonien (2026) and Chen (2022) all use a similar quadratic holding cost. More generally, including these exogenous inventory costs is a reduced-form approach to modeling incentives to risk share. For simplicity, we focus on the case where the marginal holding cost for all traders is governed by the same parameter, $\gamma > 0$.[4]

We describe the trading sessions further below. With probability $1 - \delta$ the asset will liquidate after period 1, and otherwise, with probability δ, the asset will liquidate after

[2]The model can be extended to have more types of traders and additional trading periods, though such extensions greatly complicate computation and obscure the main economic insights.

[3]Adding a non-zero mean has no qualitative effects.

[4]This assumption can also be relaxed, though it is not the focus of our paper.

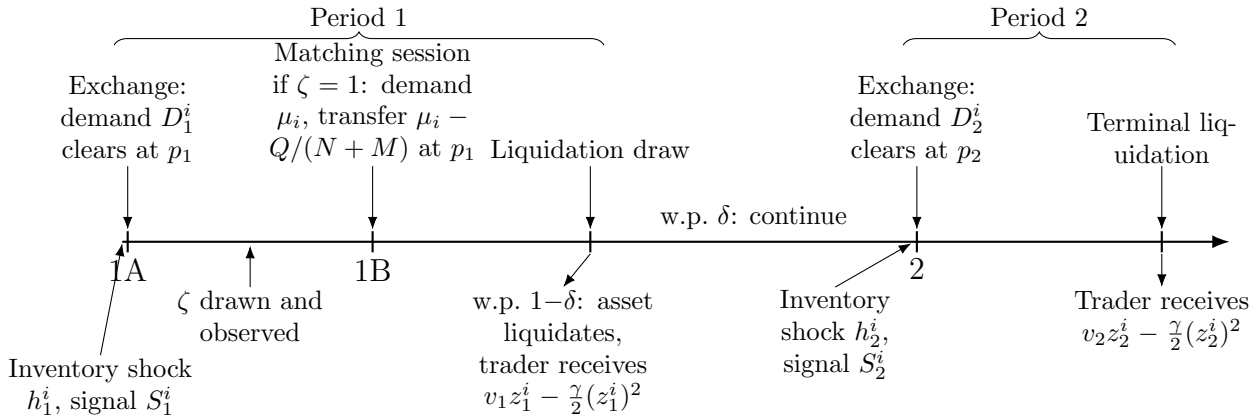

Figure 1: Timeline of the model

The model begins with each trader receiving a random inventory shock and a private signal about the asset's value. Period 1 trade takes place in two stages. In stage 1A, traders submit demand schedules to an exchange that clears at a single price. Nature then reveals whether a matching session occurs, which happens with probability θ. If it does, in stage 1B each trader submits a demand, the aggregate is published, and each receives their demand net of the average at the stage 1A price. After period 1 trade, the asset liquidates with probability $1-\delta$, and traders receive its value net of a quadratic holding cost on their inventory. Otherwise, at period 2 each trader receives a new inventory shock and signal and trades in a second exchange session, after which the asset liquidates and traders receive its value net of the holding cost on their final positions.

period 2. Upon liquidation in period t, each trader receives the value of their position minus their holding costs, $v_t\, z_{i,t} - \frac{\gamma}{2}\, z_{i,t}^2$.

There are three trading sessions across the two periods. The first is a price-clearing session at the start of period 1, which we call stage 1A. The second is a matching (quantity-clearing) session that may take place immediately after, which we call stage 1B. The third is a second price-clearing session at period 2. Each price-clearing session is modeled as a uniform-price double auction (Klemperer and Meyer, 1989, Kyle, 1989, Vives, 2011). Specifically, each trader submits a demand schedule $D_t^i : \mathbb{R} \to \mathbb{R}$ that is a mapping of price to demand, $p \mapsto D_t^i(p)$, which specifies how many units they wish to buy or sell at each possible price. The market-clearing price, p_t, is the price that sets the net demand to zero,

$$\sum_{i=1}^{N+M} D_t^i(p_t) = 0. \tag{1}$$

After a session, each trader pays the equilibrium price times the amount of the asset they were allocated, $p_t \times D_t^i(p_t)$, and is allocated $D_t^i(p_t)$ units of the asset. If they are selling at the equilibrium price, $D_t^i(p_t) < 0$, then trader i receives the price times the amount of the

asset they sell. Modeling trade as a double auction, rather than through a limit-order book, preserves analytical tractability while capturing the essential mechanism of price impact. After the stage 1A exchange, trader i's inventory is $\hat{z}_1^i = h_1^i + D_1^i(p_1)$.

After the stage 1A exchange clears, nature draws $\zeta \in \{0, 1\}$, with $\Pr(\zeta = 1) = \theta$, independently of all other random variables, and ζ is observed by all traders. If $\zeta = 1$, a fixed-price matching session takes place at stage 1B. Unlike an exchange, the matching session does not discover a price. All quantity traded in the matching session settles at the market clearing stage 1A price, p_1. Importantly, regardless of the quantity demanded or cleared, the price per unit is *fixed* (frozen) at p_1. Each trader submits a single quantity, μ_i, which we refer to as their matching session demand. The allocation mechanism, which can be referred to as proportional rationing, allocates to each trader their demand net of the average demand,

$$D_{SD}^i = \mu_i - \frac{Q}{N+M}. \tag{2}$$

A trader who demands more than the average buys in the matching session, and a trader who demands less than the average sells. By construction, $\sum_j D_{SD}^j = 0$, so the matching session is budget-balanced and requires no outside liquidity. After the matching session, trader i's inventory is $z_1^i = \hat{z}_1^i + D_{SD}^i$. If $\zeta = 0$, no trade occurs at stage 1B, so that $D_{SD}^i \equiv 0$ and $z_1^i = \hat{z}_1^i$. The matching session is a stylized version of size discovery mechanisms, such as workup in the U.S. Treasury market and block-crossing sessions in equity dark pools, which allow large quantities to change hands at a price determined elsewhere (Duffie and Zhu, 2017, Antill and Duffie, 2020). Because the price is fixed at p_1, a trader's matching session demand has no impact on the price at which they trade. All traders can infer aggregate demand Q from their quantity demanded versus allocated, and traders use this to update their beliefs about the fundamental value and aggregate inventory positions, affecting trading strategies in period 2.

If the asset does not liquidate after period 1, each trader is subject to a second inventory shock, h_2^i, the liquidation value of the asset is shocked, and all informed traders receive a new signal, S_2^i. Their inventory before period 2 exchange trade is $\hat{z}_2^i = z_1^i + h_2^i$, and their inventory after it is $z_2^i = \hat{z}_2^i + D_2^i(p_2)$. The asset then liquidates.

Before any trade, each trader's expected utility is given by

$$\mathbb{E}\Big[- p_1\big(D_1^i + D_{SD}^i\big) + (1-\delta)\Big(v_1 z_1^i - \frac{\gamma}{2}\big(z_1^i\big)^2\Big) + \delta\Big(- p_2 D_2^i + v_2 z_2^i - \frac{\gamma}{2}\big(z_2^i\big)^2\Big)\Big]. \tag{3}$$

The first term represents the payment for trade at period 1. Trade in the exchange and in the matching session are paired together, since both settle at p_1. The next term applies if liquidation happens after period 1, which occurs with probability $1-\delta$. The first part gives the liquidation payoff, and the second accounts for the holding cost of the trader's position. The third term is the payoff if liquidation does not happen after period 1 and, therefore, occurs after period 2. Within this third term, the first term represents the payment for trade in the period 2 exchange trade session, the second is the liquidation payoff, and the final term is the holding cost of the trader's final position.

Expanding equation (3), we can express the endogenous inventory process in terms of exogenous inventory shocks and traded quantities. Each trader's optimization problem is therefore

$$\max_{\substack{D_1^i,\, D_2^i:\mathbb{R}\to\mathbb{R} \\ \mu_i\in\mathbb{R}}} \mathbb{E}\Bigg[-p_1\big(D_1^i(p_1)+D_{SD}^i\big)+(1-\delta)\Big(v_1\big(h_1^i+D_1^i(p_1)+D_{SD}^i\big) \\ -\frac{\gamma}{2}\big(h_1^i+D_1^i(p_1)+D_{SD}^i\big)^2\Big)+\delta\Big(-p_2D_2^i(p_2)+v_2\big(h_1^i+D_1^i(p_1)+D_{SD}^i+h_2^i+D_2^i(p_2)\big) \\ -\frac{\gamma}{2}\big(h_1^i+D_1^i(p_1)+D_{SD}^i+h_2^i+D_2^i(p_2)\big)^2\Big)\Bigg], \tag{4}$$

where $D_{SD}^i=\mu_i-Q/(N+M)$ if $\zeta=1$ and $D_{SD}^i=0$ if $\zeta=0$. Traders choose demand schedules, $D_1^i(\cdot)$ and $D_2^i(\cdot)$, in the two price-clearing sessions, which are evaluated at the market clearing prices p_1 and p_2. Traders only choose a single quantity, μ_i, in the matching session, since trade settles at the prevailing exchange trade price p_1. The period 1 demand schedule is chosen before ζ is realized, while the matching session demand and the period 2 demand schedule are chosen after.

In the first trading session, the information set of trader i of an informed trader is $\mathcal{F}_i^{(1),I}=\{S_1^i,h_1^i,p_1\}$, where the price is included because traders submit demand schedules that condition on the clearing price, and $\mathcal{F}_i^{(1),U}=\{h_1^i,p_1\}$ for uninformed traders. At stage 1B, traders only learn whether the matching session takes place, so their information set is $\mathcal{F}_i^{(\mathrm{SD}),\tau}=\mathcal{F}_i^{(1),\tau}$. In particular, traders submit their matching session demand before observing Q. In the second trading session, informed trader i observes $\mathcal{F}_i^{(2),I}=\{S_1^i,S_2^i,h_1^i,h_2^i,p_1,p_2,Q\}$, where Q is included only if the matching session took place, and uninformed trader i observes $\mathcal{F}_i^{(2),U}=\{h_1^i,h_2^i,p_1,p_2,Q\}$. The aggregate matching session demand is therefore a second public signal, alongside the two market-clearing prices.

This will play an important role because traders are trying to learn two dimensions in the model: the asset's value and the distribution of inventory across types. While the asset's liquidating value is common to all, how inventory is spread across the types also affects the market-clearing price and therefore trade terms in the second period.

III. Equilibrium

We study a type-symmetric linear Perfect Bayesian Nash Equilibrium. The linear-quadratic specification implies that traders' optimal demand schedules and matching session demands are linear functions of their information and inventories. We solve for the equilibrium by starting with the period 2 price-clearing session and working backward, first to the matching session and then to the period 1 price-clearing session. At period 1, traders recognize that their demand influences not only the current price, but also their allocation in the matching session, the aggregate statistic Q, and the equilibrium price and demand functions at period 2. At stage 1B, traders recognize that their matching session demand influences the period 2 price through Q. Because traders of the same type are ex-ante identical, the equilibrium coefficients that determine their trading behavior are the same within type, and we denote them with a superscript τ. Individual realized demands, however, still vary with the realization of each trader's signals and inventory shocks. Appendix A contains the full derivations.

A. Trading in the Second Session

Starting from the period 2 price-clearing session, each trader solves the following optimization problem:

$$D_2^i(p) \in \underset{D_2^i}{\operatorname{argmax}}\ \mathbb{E}\Big[-pD_2^i + v_2\big(\hat{z}_2^i + D_2^i\big) - \frac{\gamma}{2}\big(\hat{z}_2^i + D_2^i\big)^2 \ \Big|\ \mathcal{F}_i^{(2),\tau}\Big], \tag{5}$$

subject to the market-clearing condition $D_2^i(p) + \sum_{j\neq i} D_2^j(p) = 0$, where p is the market-clearing price. Each trader can therefore be viewed as trading against a residual demand curve. In equilibrium, they will solve for the optimal demand schedule given their beliefs and conditional on all other traders playing their equilibrium demand schedules. An equilibrium condition will be that their beliefs and actions must be mutually consistent. Define trader i's price impact as $\lambda_2^\tau = \frac{\mathrm{d}p_2}{\mathrm{d}D_2^i}$, which is a constant in equilibrium that we pin down later. Differentiating with respect to D_2^i, and recognizing that the trader's demand moves the

price as they trade against the residual demand curve, the optimal demand schedule satisfies the first-order condition

$$\mathbb{E}\big[v_2 - \gamma\big(\hat{z}_2^i + D_2^i\big) \mid \mathcal{F}_i^{(2),\tau}\big] = p + \lambda_2^\tau D_2^i \tag{6}$$

at each price p. The left-hand side is the expected marginal value of an additional unit of the asset, which is its expected value net of the marginal holding cost. The right-hand side is the marginal cost of buying that unit. The trader pays the price, p, for the additional unit, and because their demand moves the price by λ_2^τ, they also pay more for the D_2^i units they were already demanding. Similarly, a trader who sells an additional unit lowers the price they receive on all of the units they sell. In a competitive market, $\lambda_2^\tau = 0$, and the trader trades until the expected marginal value of the asset equals the price. Under imperfect competition, the second term leads traders to stop short of this position, whether they are buying or selling. Because traders submit demand schedules, (6) holds at every price, and the conditional expectation conditions on the price, since the price carries information about the value and distribution of the asset. Solving (6) for D_2^i, the optimal demand schedule is

$$D_2^i(p) = \chi_2^\tau\left(\frac{1}{\gamma}\big(\mathbb{E}[v_2 \mid \mathcal{F}_i^{(2),\tau}] - p\big) - \hat{z}_2^i\right), \qquad \chi_2^\tau := \frac{\gamma}{\gamma + \lambda_2^\tau}. \tag{7}$$

This expression admits a natural interpretation. First, the term inside the large parentheses is the demand schedule a trader would submit in a competitive market, where price impact is zero. The first part is how much a trader tilts their demand to take advantage of their information, relative to the price, scaled by the marginal cost of holding inventory. They also sell off their current inventory position. The prefactor, χ_2^τ, is always between zero and one, as price impact is weakly positive in equilibrium, and shows how strategic traders scale down their demand as they weigh price impact, λ_2^τ, against their holding cost, γ. In competitive markets, this scaling factor is one. We refer to χ_2^τ as the period 2 trade intensity of type τ. In equilibrium, it is the fraction of the gap between a trader's current and competitive inventory position that the trader closes, and a larger value is associated with more efficient trade. In the literature, this scaling is also known as bid shading.

Importantly, the conditional expectation $\mathbb{E}[v_2 \mid \mathcal{F}_i^{(2),\tau}]$ is a linear function of both signals, both prices, both inventory shocks, and, if the matching session took place, the aggregate matching session demand Q (see equation (28)). For an informed trader with a perfect signal, the weight on all variables except S_2^i is zero. In contrast, the uninformed trader places weight on each of them. They use the two prices to learn about the value of the asset,

their own inventory shocks to remove the effects of their own orders out of those prices, and, as we will see later, Q to better separate the part of the prices that is driven by inventories from the part driven by the value of the asset. We analyze this last channel, which is what distinguishes our model from a sequence of exchange sessions, in Section IIIE.2.

Given the optimal demand schedules, the market-clearing price is

$$\begin{aligned} p_2 =& \frac{\frac{N}{\gamma+\lambda_2^I}}{\frac{N}{\gamma+\lambda_2^I}+\frac{M}{\gamma+\lambda_2^U}}\left(\frac{1}{N}\sum_{j=1}^{N}\mathbb{E}\big[v_2 \mid \mathcal{F}_j^{(2),I}\big]-\gamma\,\frac{1}{N}\sum_{j=1}^{N}\hat{z}_2^j\right) \\ &+\frac{\frac{M}{\gamma+\lambda_2^U}}{\frac{N}{\gamma+\lambda_2^I}+\frac{M}{\gamma+\lambda_2^U}}\left(\frac{1}{M}\sum_{j=1}^{M}\mathbb{E}\big[v_2 \mid \mathcal{F}_j^{(2),U}\big]-\gamma\,\frac{1}{M}\sum_{j=1}^{M}\hat{z}_2^j\right). \end{aligned} \tag{8}$$

The price is only implicitly defined by (8), since the conditional expectations themselves depend on p_2. The equilibrium price thus represents a weighted average of each type's average valuation of the asset, net of the marginal holding cost of the type's average inventory. Traders who face smaller price impacts receive larger weights in this average because they trade more aggressively and therefore have a stronger effect on prices.

Finally, market clearing implies that the set of price impacts $\{\lambda_2^I, \lambda_2^U\}$ must be consistent with the inverse slope of the residual supply curve faced by each trader.[5] That is, $\lambda_2^\tau = -\big(\sum_{j\neq i}\frac{\partial D_2^j}{\partial p_2}\big)^{-1}$, which for an informed trader is

$$\lambda_2^I = \frac{\gamma}{(N-1)\,\chi_2^I\big(1-c_{p_2,2}^I\big)+M\,\chi_2^U\big(1-c_{p_2,2}^U\big)}, \tag{9}$$

where $c_{p_2,2}^\tau$ is the weight a type τ trader's conditional expectation places on p_2. The expression for λ_2^U is symmetric. The denominator of (9), divided by γ, is the sum of the slopes of the other traders' demand schedules, reflecting how much their total demand falls when the price rises by one unit. Each of these slopes, $\chi_2^\tau\big(1-c_{p_2,2}^\tau\big)/\gamma$, has two parts. The first is the trade intensity, so a trader's price impact is smaller when there are more traders on the other side of their order and when those traders trade more aggressively. Because each trader's trade intensity depends on their own price impact, the price impacts of the two types are determined jointly. The second part reflects learning from the price. A higher price could mean that the value of the asset, v_2, is higher, but it could also mean that other traders have

[5] From the market-clearing condition, demand by trader i at price p must be that $D_t^i + \sum_{j\neq i} D_t^j(p) = 0$. Then, differentiating with respect to the price gives $\frac{\partial D_t^i}{\partial p} + \sum_{j\neq i}\frac{\partial D_t^j(p)}{\partial p} = 0$. Solving for $\frac{\partial p}{\partial D_t^i}$, we get that $\lambda_t^\tau := \frac{\partial p}{\partial D_t^i} = -\big(\sum_{j\neq i}\frac{\partial D_t^j(p)}{\partial p}\big)^{-1}$.

negative inventory positions and are buying to move them back toward zero, which reduces their holding costs. When the price increases, a trader who learns from the price attributes part of the increase to a higher v_2, and so they reduce their demand by less. When adverse selection is greater, traders place more weight on the price when estimating the asset's value, so a higher price makes them less willing to sell, and a buyer must move the price further to fill an order.

The matching session improves liquidity in period 2 through this mechanism. Together with p_1, Q helps the uninformed separate the value of the asset from the inventories of other traders in period 1, so they enter period 2 with less uncertainty about v_2. Because their beliefs are more precise, a change in p_2 moves their valuation of the asset by less. They therefore are willing to sell more to a trader whose order pushes up the price, and price impact in period 2 falls.

B. The Matching Session

If $\zeta = 1$, a matching session occurs and each trader demands μ_i given their information set $\mathcal{F}_i^{(\mathrm{SD}),\tau} = \mathcal{F}_i^{(1),\tau}$. The matching session demand has no impact on p_1. The matching session demand affects the period 2 environment in two ways: it alters the inventory the trader carries into period 2, and it alters the aggregate demand Q that the other traders use to update their beliefs. Therefore, the demanded quantity in period 1B will have a cross-period price impact on the clearing price in period 2. Denote the resulting cross-period price impact by $\lambda_{\mathrm{SD},2}^{\tau} = \frac{\mathrm{d}p_2}{\mathrm{d}\mu_i}$. Each trader's optimization problem at stage 1B is

$$\mu_i \in \underset{\mu_i}{\operatorname{argmax}}\ \mathbb{E}\Big[-p_1 D_{SD}^i + (1-\delta)\Big(v_1\big(\hat{z}_1^i + D_{SD}^i\big) - \frac{\gamma}{2}\big(\hat{z}_1^i + D_{SD}^i\big)^2\Big) + \delta\, V_{i,2} \;\Big|\; \mathcal{F}_i^{(\mathrm{SD}),\tau}\Big], \quad (10)$$

where $V_{i,2} = -p_2 D_2^i + v_2 z_2^i - \frac{\gamma}{2}(z_2^i)^2$ is the period 2 continuation value. After some simplification, the first-order condition can be written as

$$\mathbb{E}\big[v_1 - \gamma\big(\hat{z}_1^i + D_{SD}^i\big) \;\big|\; \mathcal{F}_i^{(\mathrm{SD}),\tau}\big] = \Omega_{\mathrm{SD}}^{\tau}\, p_1 + \big(1 - \Omega_{\mathrm{SD}}^{\tau}\big)\, \mathbb{E}\big[p_2 \;\big|\; \mathcal{F}_i^{(\mathrm{SD}),\tau}\big], \quad (11)$$

where

$$\Omega_{\mathrm{SD}}^{\tau} = \frac{1}{1 - \delta\big[1 - (1-\chi_2^{\tau})^2\big] - \delta\,(1-\chi_2^{\tau})(2-\chi_2^{\tau})\,\dfrac{(N+M)\,\lambda_{\mathrm{SD},2}^{\tau}}{(N+M-1)\,\lambda_2^{\tau}}}. \quad (12)$$

The left-hand side of (11) is the expected value of the asset net of the marginal holding cost. The right-hand side is an effective price, a weighted combination of the price at which the

unit is bought, p_1, and the price the trader expects at period 2. In a competitive market, $\chi_2^\tau = 1$ and $\Omega_{\text{SD}}^\tau = \frac{1}{1-\delta}$, so the effective price is $\frac{p_1 - \delta\,\mathbb{E}[p_2]}{1-\delta}$, the cost of a unit bought at p_1 and resold at p_2 if the asset does not liquidate. Under imperfect competition, Ω_{SD}^τ deviates from $\frac{1}{1-\delta}$ for two reasons. First, the trader will only partially unwind their position at period 2 due to bid shading in response to the price impact, since $\chi_2^\tau < 1$. Second, the trader's matching session demand moves the period 2 price through the cross-period price impact, $\lambda_{\text{SD},2}^\tau$.

Solving (11) for the matching session allocation, the expected allocation can be written as the expected competitive allocation plus an additional term due to imperfect competition,

$$\mathbb{E}\big[D_{SD}^i \;\big|\; \mathcal{F}_i^{(\text{SD}),\tau}\big] = \underbrace{\frac{1}{\gamma}\left(\mathbb{E}[v_1 \mid \mathcal{F}_i^{(\text{SD}),\tau}] - \frac{p_1 - \delta\,\mathbb{E}[p_2 \mid \mathcal{F}_i^{(\text{SD}),\tau}]}{1-\delta}\right) - \hat{z}_1^i}_{\text{competitive demand}} + \underbrace{\frac{1}{\gamma}\Big(\frac{1}{1-\delta} - \Omega_{\text{SD}}^\tau\Big)\big(p_1 - \mathbb{E}[p_2 \mid \mathcal{F}_i^{(\text{SD}),\tau}]\big)}_{\text{imperfect-competition correction}}. \tag{13}$$

Unlike the exchange demand in (7), the competitive part of the matching session allocation is not scaled down by a trade intensity. The trader's demand does not move the price at which they trade, so they aim for their full desired position given the effective price. Only the correction for the cross-period price impact in period 2 remains. Since the allocation is the trader's demand net of the average demand, the matching session demand itself satisfies

$$\frac{N+M-1}{N+M}\,\mu_i = \frac{1}{\gamma}\Big(\mathbb{E}[v_1 \mid \mathcal{F}_i^{(\text{SD}),\tau}] - \big(\Omega_{\text{SD}}^\tau\, p_1 + \big(1 - \Omega_{\text{SD}}^\tau\big)\big)\,\mathbb{E}[p_2 \mid \mathcal{F}_i^{(\text{SD}),\tau}]\Big) - \hat{z}_1^i + \mathbb{E}\left[\frac{1}{N+M}\sum_{j\neq i}\mu_j \;\middle|\; \mathcal{F}_i^{(\text{SD}),\tau}\right]. \tag{14}$$

The last term is the trader's forecast of the average demand of their rivals. A trader who expects their rivals to demand a lot must demand more themselves to receive the same allocation. Two features therefore distinguish a matching session demand from a demand schedule in a price-clearing session. First, there is no trade intensity scaling down the demand. Because trade settles at p_1, the trader's matching session demand does not move the price at which they trade. Imperfect competition still affects this target, but only through period 2. The weight Ω_{SD}^τ depends on the period 2 trade intensity and on how the matching session demand moves the period 2 price through a cross-period price impact, $\lambda_{\text{SD},2}^\tau$. Second,

because each trader's allocation is their demand net of the average demand, the demand contains the trader's forecast of their rivals' matching session demands. Both features are central to why and about what Q is informative, which we discuss in Section IIIE.2.[6]

C. Trading in the First Session

In period 1A, traders submit demand schedules before learning whether the matching session will take place.[7] The decision at date 1 depends on two "branches" of the model, which are characterized by the realization of ζ. We first describe the period 1 first-order condition on each branch separately, and then combine them. In each case, traders act strategically, recognizing that their current demand influences not only the current price but also future prices and demands. Importantly, because of the linear-quadratic setting we consider, these marginal impacts are constant, keeping this non-Markovian model tractable. We denote the cross-period price impact by $\lambda^{\tau}_{1,2} = \frac{\mathrm{d}p_2}{\mathrm{d}D^i_1}$, the cross-period demand impact by $\frac{\mathrm{d}D^i_2}{\mathrm{d}D^i_1}$, and the trader's current price impact by $\lambda^{\tau}_1 = \frac{\mathrm{d}p_1}{\mathrm{d}D^i_1}$.

No matching session. Suppose first that there is no matching session. Each trader's optimization problem at period 1 is then

$$D^i_1(p) \in \underset{D^i_1}{\operatorname{argmax}}\ \mathbb{E}\Big[-pD^i_1 + (1-\delta)\Big(v_1\big(h^i_1 + D^i_1\big) - \frac{\gamma}{2}\big(h^i_1 + D^i_1\big)^2\Big) + \delta\, V_{i,2} \ \Big|\ \mathcal{F}^{(1),\tau}_i\Big], \quad (15)$$

where $V_{i,2}$ is the period 2 continuation value defined in Section IIIB. Unlike at period 2, the trader's demand now also affects this continuation value. An additional unit bought at period 1 is carried into period 2, where the trader sells part of it. Implicitly differentiating

[6] Matching session demands are determined only up to a common shift that depends on public information. If every trader adds the same amount $f(p_1)$ to their demand, then each allocation D^i_{SD} is unchanged and Q shifts by the publicly known amount $(N+M)\,f(p_1)$, so it carries the same information. Allocations, prices, price impacts, beliefs, and welfare are therefore the same for every such shift. Within the linear equilibria we study, $f(p_1) = k_0 + k\,p_1$ for constants k_0 and k. We set $k_0 = 0$, since all random variables have mean zero. Only the loadings of μ_i and of period 2 demand on p_1, and the variance of Q, depend on k, and we report them for a single normalization, in which we choose k so that the informed's matching session demand places no weight on p_1.

[7] If the matching session took place for certain, $\theta = 1$, units bought in the price-clearing session and in the matching session both settle at p_1, so they are perfect substitutes. Trading in the price-clearing session creates price impact, which the matching session avoids, so traders would shift their trade to the matching session. The price-clearing session would then no longer determine a price, and the matching session would have no price at which to trade. We therefore focus on $\theta \in [0, 1)$.

the period 2 first-order condition with respect to period 1 demand,

$$\frac{\mathrm{d}D_2^i}{\mathrm{d}D_1^i} = -\frac{\gamma + \lambda_{1,2}^\tau}{\gamma + \lambda_2^\tau}. \tag{16}$$

The amount the trader sells at period 2 depends on the period 2 holding cost, the period 2 price impact, and how the period 1 trade affects the period 2 price. The cross-period price impact, $\lambda_{1,2}^\tau$, is determined by differentiating period 2 market clearing with respect to period 1 demand (Appendix AE.4). It captures that a period 1 trade changes the inventories that the trader and their counterparties carry into period 2, as well as the period 1 price, which rivals use to form their period 2 beliefs.

After some simplification, the period 1 first-order condition can be written as

$$\kappa_v^\tau \,\mathbb{E}\big[v_1 - \gamma\big(h_1^i + D_1^i\big) \;\big|\; \mathcal{F}_i^{(1),\tau}\big] + \big(1 - \kappa_v^\tau\big)\,\mathbb{E}\big[p_2 \;\big|\; \mathcal{F}_i^{(1),\tau}\big] = p + \lambda_1^\tau D_1^i, \tag{17}$$

where

$$\kappa_v^\tau := (1-\delta) + \delta\,(1-\chi_2^\tau)\left(1 + \frac{\mathrm{d}D_2^i}{\mathrm{d}D_1^i} - \frac{\lambda_{1,2}^\tau}{\lambda_2^\tau}\right). \tag{18}$$

The right-hand side of (17) is the marginal cost of the purchase of a unit, given by the price, p, plus the increase in the price paid on the trader's other units, $\lambda_1^\tau D_1^i$. The left-hand side is the marginal benefit, a weighted average of the unit's marginal value after period 1 and the price the trader expects at period 2. In a competitive market, $\kappa_v^\tau = 1 - \delta$. With probability $1 - \delta$, the asset liquidates and the unit is worth its marginal value, and with probability δ, the trader resells the unit at p_2. Under imperfect competition, κ_v^τ also depends on how much of the unit the trader unwinds at period 2 and on how their period 1 demand affects the period 2 price. Solving (17) for D_1^i, the period 1 demand schedule is

$$D_1^i(p) = \chi_1^\tau\left(\frac{1}{\gamma}\left(\mathbb{E}\big[v_1 \;\big|\; \mathcal{F}_i^{(1),\tau}\big] - \frac{p - \big(1-\kappa_v^\tau\big)\,\mathbb{E}\big[p_2 \;\big|\; \mathcal{F}_i^{(1),\tau}\big]}{\kappa_v^\tau}\right) - h_1^i\right), \tag{19}$$

where $\chi_1^\tau := \frac{\gamma\kappa_v^\tau}{\gamma\kappa_v^\tau + \lambda_1^\tau}$. This form is analogous to the period 2 demand in equation (7). The term inside the large parentheses is the competitive demand at an effective price that combines p and the expected period 2 price. The outer fraction, χ_1^τ, is the period 1 trade intensity, which moderates how aggressively the trader trades. Without price impact, $\kappa_v^\tau = 1 - \delta$, $\chi_1^\tau = 1$, and the effective price is $\frac{p - \delta\,\mathbb{E}[p_2]}{1-\delta}$, as in the matching session.

Both informed and uninformed traders form conditional expectations that depend on their signal, if they receive one, their inventory shock, and the market-clearing price. The uninformed traders infer information about both the fundamental value and the distribution

of inventories across trader types from the observed price. In contrast, informed traders with signals use the price disproportionately more to update their beliefs about the inventories of other traders. These beliefs inform traders' expectations of the price that they will encounter in the second trading session.

Probabilistic matching session When the matching session takes place with probability θ, the period 1 first-order condition weights the marginal expected utilities on the two branches by $1-\theta$ and θ,

$$(1-\theta)\Big(\kappa_v^\tau\,\mathbb{E}\big[v_1-\gamma\big(h_1^i+D_1^i\big)\;\big|\;\mathcal{F}_i^{(1),\tau}\big]+\big(1-\kappa_v^\tau\big)\,\mathbb{E}\big[p_2\;\big|\;\mathcal{F}_i^{(1),\tau},\zeta=0\big]-p_1-\lambda_1^\tau D_1^i\Big)+\theta\,G_{\mathrm{SD}}^i=0, \tag{20}$$

where the first term is the marginal expected utility from (17) and G_{SD}^i is the marginal expected utility on the matching session branch, derived in Appendix AD. The second term does not depend on D_1^i. On the matching session branch, the trader re-optimizes their position in the matching session, so an additional unit bought in the price-clearing session can be sold back in the matching session. The session in which the trader purchases a unit is not irrelevant. A unit purchased in the price-clearing session affects the period 2 price in a different way than a unit purchased in the matching session, and it increases p_1, which applies to all of the trader's period 1 trade. Period 1 demand is determined by the first term alone, which receives weight $1-\theta$. Solving (20) for D_1^i, the resulting demand schedule is

$$D_1^i(p_1)=\frac{\chi_1^\tau-\theta}{1-\theta}\left(\frac{1}{\gamma}\Big(\mathbb{E}\big[v_1\;\big|\;\mathcal{F}_i^{(1),\tau}\big]-\bar{p}_1^{\,i}\Big)-h_1^i\right), \tag{21}$$

where $\bar{p}_1^{\,i}$ is an effective price that combines p_1 and the trader's forecasts of p_2 on both of the two branches (Appendix AD). This is of the same form as (19), but the anticipation of the matching session lowers the period 1 trade intensity from χ_1^τ to $\frac{\chi_1^\tau-\theta}{1-\theta}$.[8] With probability θ, a unit bought in the price-clearing session can be undone in the matching session, so it only adds price impact. Traders postpone part of their inventory adjustment to the matching session, even though it takes place only with probability θ, which harms liquidity in the period 1 price-clearing session (Antill and Duffie, 2020). As $\theta\to 1$, the branch without a matching session receives no weight, and the period 1 demand schedule is no longer determined, which is the degeneracy described above. When $\theta=0$, (21) reduces to (19).

[8] This requires $\theta<\chi_1^\tau$, which holds in every parameterization we study.

D. *Equilibrium*

We solve for a type-symmetric linear Perfect Bayesian equilibrium, which is characterized by the following proposition:

PROPOSITION 1: *A Perfect Bayesian linear equilibrium requires that:*

1. *Given the price impacts* λ_1^τ, λ_2^τ, $\lambda_{1,2}^\tau$, *and* $\lambda_{\mathrm{SD},2}^\tau$, *and the cross-period demand impacts* $\frac{\mathrm{d}D_2^i}{\mathrm{d}D_1^i}$, $\frac{\mathrm{d}D_2^i}{\mathrm{d}\mu_i}$, *and* $\frac{\mathrm{d}\mu_i}{\mathrm{d}D_1^i}$, *exchange demands and matching session demands are optimal given individuals' information sets, and exchange prices are chosen to set aggregate demands to 0,*
2. *Given the price and demand impacts and equilibrium demands and prices, beliefs* $\mathbb{E}[v_1 \mid \mathcal{F}_i^{(1),\tau}]$, $\mathbb{E}[p_2 \mid \mathcal{F}_i^{(1),\tau}]$, $\mathbb{E}[v_2 \mid \mathcal{F}_i^{(2),\tau}]$, *and the forecasts of rivals' matching session demands,* $\mathbb{E}\big[\frac{1}{N+M}\sum_{j\neq i}\mu_j \mid \mathcal{F}_i^{(\mathrm{SD}),\tau}\big]$, *are determined by Bayesian updating,*
3. $\{\lambda_1^\tau, \lambda_2^\tau\}$ *are determined to ensure price impacts are consistent with residual demand curves, while* $\{\lambda_{1,2}^\tau, \lambda_{\mathrm{SD},2}^\tau\}$ *and the cross-period demand impacts are determined by implicit differentiation of the period 2 market-clearing condition and of the period 2 and matching session first-order conditions.*

This characterization is an application of Rostek et al. (2025) to the case in which a matching session is modeled. The model quickly becomes complex because a trader's equilibrium price impact depends not only on how traders of the same type react and learn, but also on the responses of traders of the other type, both in the current and in future sessions. The matching session adds to this complexity in three ways. It creates a second branch at period 2, on which traders observe aggregate quantity demanded, Q, and hold different inventories. It creates a new cross-period price impact, $\lambda_{\mathrm{SD},2}^\tau$. And, it requires each trader to forecast the demands of their rivals, because their allocation depends on the realization of Q, which is not observed when the demands are submitted. With two trader types, the following twelve price impacts are jointly determined in equilibrium:

$$\big\{\lambda_1^\tau,\ \lambda_2^{\tau,\mathrm{SD}},\ \lambda_2^{\tau,\mathrm{no}},\ \lambda_{1,2}^{\tau,\mathrm{SD}},\ \lambda_{1,2}^{\tau,\mathrm{no}},\ \lambda_{\mathrm{SD},2}^\tau\big\}_{\tau\in\{I,U\}}, \tag{22}$$

where the superscripts SD and no denote the matching session ($\zeta = 1$) and no matching session ($\zeta = 0$) branches. The period 1 price impact is common to both branches, since period 1 demand is submitted before ζ is realized. This intertemporal and cross-type dependence creates a nonlinear, non-recursive fixed-point problem for the price impacts, the strategy

coefficients, and the projection coefficients. As a result, we must rely on numerical methods to solve the model (Appendix AH). When $\theta = 0$, the model reduces to a two-period model of Rostek et al. (2025), which we use as a benchmark throughout.

E. Equilibrium Properties

We now illustrate the economic forces in the model. Unless otherwise stated, we use the following parameterization throughout:

$$N = 20, \quad M = 20, \quad \gamma = 2, \quad \delta = 0.5, \quad \theta = 0.1,$$
$$\sigma^2 = 1, \quad \sigma_\eta^2 = 0.1, \quad \sigma_{\epsilon,I}^2 = .25, \quad \sigma_h^2 = 1.$$

Setting $\sigma_{\epsilon,I}^2 = 0.25$ means that informed traders observe a noisy signal of v_t at each period, while uninformed traders have no private signal and learn about the value of the asset only from the prices and, when the matching session takes place, from Q. The value $\sigma^2 = 1$ normalizes the prior variance of the value of the asset, and $\sigma_\eta^2 = 0.1$ ensures that the period 2 signals carry meaningful new information. With $\delta = 0.5$, the asset liquidates with probability one half between the two periods, and $\theta = 0.1$ means the matching session happens 10% of the time. The market has twenty traders of each type, and inventory shocks have the same variance as the value of the asset, $\sigma_h^2 = \sigma^2 = 1$. We compare the equilibrium to the benchmark without a matching session, in which $\theta = 0$.

Table I reports the equilibrium demand coefficients in an economy with and without the matching session. Informed traders trade on their signal, while the uninformed, who have no private signal, place no weight on it. Both types sell off most of their inventory shock in period 1. Anticipating the matching session, traders postpone part of their inventory adjustment to it, as described in Section IIIC, resulting in the loadings on the inventory shock being slightly smaller in magnitude for both types.

The matching session demand loads negatively on the inventory shock, so a trader who receives a larger inventory shock submits a lower matching session demand. Because each trader receives their demand net of the average demand, traders with above-average inventories sell in the matching session, and traders with below-average inventories buy. The matching session therefore moves traders part of the way toward the cross-sectional average inventory. For the informed, the loading of -0.1205 is close to the $1 + (-0.8645) = 0.1355$ units of each inventory shock that remain after the period 1 exchange. Because the price is

Table IEquilibrium Demand Coefficients
We report the loadings of the period 1 demand, the matching session demand, and the period 2 demand on each element of the trader's information set, by type. The columns labeled "Matching" report the equilibrium with a probabilistic matching session, $\theta = 0.1$, with the matching session demand and period 2 demand evaluated on the matching session branch ($\zeta = 1$). The columns labeled "None" report the benchmark without a matching session, $\theta = 0$. The matching session demand is written in terms of (S_1^i, h_1^i, p_1). Its loading on p_1 and the matching session branch period 2 loading on p_1 are reported for the normalization described in Section IIIB. The table uses the baseline parameterization.

	Informed		Uninformed	
Loading	Matching	None	Matching	None
Period 1 demand, D_1^i				
S_1^i	+0.3138	+0.3080	—	—
h_1^i	−0.8645	−0.8780	−0.8435	−0.8566
p_1	−0.3422	−0.3398	−0.0903	−0.0942
Matching session demand, μ_i				
S_1^i	+0.0161	—	—	—
h_1^i	−0.1205	—	−0.1161	—
p_1	+0.0000	—	+0.0133	—
Period 2 demand, D_2^i				
S_1^i	−0.3201	−0.1456	—	—
S_2^i	+0.1477	+0.2413	—	—
p_1	+1.1182	+0.3389	+1.1970	+0.2901
p_2	−0.4373	−0.4244	−0.4127	−0.3013
h_1^i	−0.0151	−0.1112	−0.0385	−0.1020
h_2^i	−0.9680	−0.9618	−0.9666	−0.9541
Q	−0.5297	—	−0.7678	—

fixed, the informed use the matching session to shed most of their remaining excess inventory, as in Antill and Duffie (2020) and Blonien (2026), where matching sessions perfectly reallocate the excess inventory. Notice, however, that the loading on their private signal is fairly low, especially relative to the loading in period 1A. As we will see, the cross-period demand impact from the matching session and period 2 price will be substantial, to discipline informed traders from trading too aggressively on their private information.

The matching session also changes period 2 demand. Without it, period 2 demand places

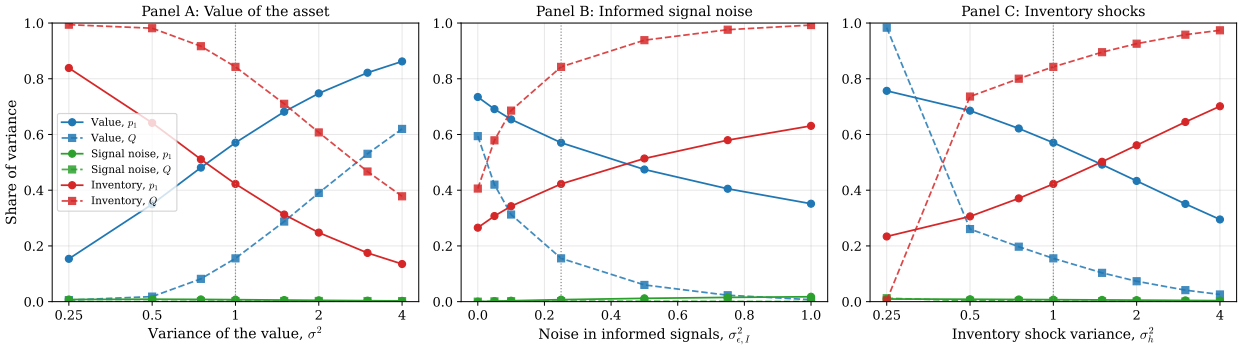

Figure 2. Variance Decomposition of the Public Signals
We plot the share of the variance of the period 1 price, p_1 (solid), and of the aggregate matching session demand, Q (dashed), attributable to the value of the asset, v_1, to the noise in traders' signals, $\{\epsilon_{i,1}\}_i$, and to traders' inventory shocks, $\{h_1^i\}_i$. Panel A varies the variance of the value of the asset, σ^2, Panel B varies the noise in the informed's signals, $\sigma^2_{\epsilon,I}$, and Panel C varies the variance of inventory shocks, σ^2_h. Dotted vertical lines mark the baseline parameterization, and all other parameters are at their baseline values. For Q, we decompose the part that loads on traders' signals and inventory shocks, which does not depend on the normalization described in Section IIIB.

substantial weight on the period 1 inventory shock, as traders finish offloading positions they did not offload at period 1. With it, that weight is close to zero, since the matching session has already offloaded most of traders' excess inventory. At period 2, the uninformed also place a large negative weight on Q. A high aggregate matching session demand indicates that traders are, on average, short inventory. Part of the period 1 price therefore reflected buying pressure from inventories rather than the value of the asset, and the uninformed lower their valuation. Because the signals of the informed are noisy in the baseline parameterization, the informed also learn from Q, and they place a negative weight on it as well. If the informed observed the value of the asset exactly, they would have nothing to learn from Q. Their period 2 demand would still depend on Q, but only because their allocation in the matching session, and therefore the inventory they carry into period 2, depends on it.

Figure 2 shows that variation in the aggregate matching session demand Q is primarily due to inventory shocks, while the price in period 1A is more of a function of the private signals in the aggregate. At the benchmark parameterization, inventory shocks account for 42% of the variance of the period 1 price, but 84% of the variance of Q.

Table IIEquilibrium Price Impacts
We report the own-period and cross-period price impacts, by type. λ_1^τ is the period 1 price impact, which is common to both branches. $\lambda_2^{\tau,\mathrm{SD}}$ and $\lambda_2^{\tau,\mathrm{no}}$ are the period 2 price impacts on the matching session and no matching session branches, and $\lambda_{1,2}^{\tau,\mathrm{SD}}$ and $\lambda_{1,2}^{\tau,\mathrm{no}}$ are the corresponding cross-period price impacts of period 1 demand on the period 2 price. $\lambda_{\mathrm{SD},2}^\tau$ is the cross-period price impact of the matching session demand on the period 2 price. The column labeled "None" reports the benchmark without a matching session, $\theta = 0$. The table uses the baseline parameterization.

Price impact	Matching ($\theta = 0.1$)	None ($\theta = 0$)	Difference
λ_1^I	0.1204	0.1199	+0.0005
λ_1^U	0.1168	0.1165	+0.0003
$\lambda_2^{I,\mathrm{SD}}$	0.0604	0.0710	−0.0106
$\lambda_2^{U,\mathrm{SD}}$	0.0603	0.0704	−0.0101
$\lambda_2^{I,\mathrm{no}}$	0.0690	0.0710	−0.0020
$\lambda_2^{U,\mathrm{no}}$	0.0685	0.0704	−0.0019
$\lambda_{1,2}^{I,\mathrm{SD}}$	+1.7465	+0.0345	+1.7120
$\lambda_{1,2}^{U,\mathrm{SD}}$	+1.7175	+0.0317	+1.6858
$\lambda_{1,2}^{I,\mathrm{no}}$	+0.0375	+0.0345	+0.0031
$\lambda_{1,2}^{U,\mathrm{no}}$	+0.0346	+0.0317	+0.0029
$\lambda_{\mathrm{SD},2}^I$	−1.5466	—	—
$\lambda_{\mathrm{SD},2}^U$	−1.5305	—	—

E.1. Price Impacts

Table II reports the price impacts with and without the matching session. Period 1 price impact is less than 1% higher when the matching session is available. Anticipating the matching session, traders shade their period 1 demand schedules, but with forty traders the period 1 exchange is fairly liquid, and the bid shading is small in magnitude. Period 2 price impact is lower on both branches than in the benchmark without a matching session. With a matching session, traders learn from Q, which allows the uninformed to better disentangle the source of period 2 demand, and they trade more aggressively. Period 2 price impact is also slightly lower even without a matching session. Anticipating the matching session, traders postpone part of their inventory-driven trade, but not their information-driven trade. The informed's ratio of signal to inventory loadings at period 1 rises, from 0.35 to 0.36, so the period 1 price is more informative about the value of the asset, and the uninformed face

less adverse selection at period 2.

The cross-period price impacts are much larger when there is a matching session. The impact of the matching session demand on the period 2 price, $\lambda^{\tau}_{\text{SD},2}$, is about 25 times as large as the period 2 price impact itself. Both result from the information content of Q for the uninformed. A trader who increases their matching session demand increases Q. The uninformed interpret the higher aggregate demand as evidence that traders are short inventory, lower their valuation of the asset, and sell at period 2, which lowers the period 2 price. The trader also carries more inventory into period 2, which they sell as well. The cross-period price impact of period 1 demand, $\lambda^{\tau}_{1,2}$, is also large when there is a matching session. Suppose a trader buys an additional unit at period 1, which increases p_1. Without Q, the uninformed cannot discern whether a higher p_1 reflects a higher value of the asset or traders who are short inventory, so they only partly attribute the price change to the value of the asset. Knowing that with Q, they can tell the two apart, they attribute more of the higher price to the value of the asset. They raise their valuation and buy at period 2, which increases p_2. The effect is even larger because of the matching session. After buying the additional unit, the trader holds more inventory, and a trader with more inventory demands less in the matching session, so Q falls. The uninformed interpret the lower Q as a signal that traders hold more inventory than they thought. More inventory should push the period 1 price down, yet p_1 went up. The uninformed conclude that the value of the asset must be even higher, which increases p_2 further. Without a matching session, cross-period price impacts are small.

Because trade in the matching session settles at a fixed price, the price impact of the matching session demand on the period 2 price, $\lambda^{\tau}_{\text{SD},2}$, is the only price impact a trader faces there. Beyond the holding cost of the inventory they receive, it is the only thing that restrains an aggressive matching session demand. This deterrent falls mostly on the informed when they trade on their signals. The informed unwind part of their signal-driven position at period 2, so a matching session demand that moves the period 2 price moves it against their period 2 trade. In contrast, after the matching session, period 2 demand places almost no weight on the period 1 inventory shock, so the deterrent has little effect on inventory-driven demand. In unreported results, if traders ignored this cross-period price impact when choosing their matching session demand, the informed would place about two and a half times as much weight on their signal. The price impact of the matching session demand

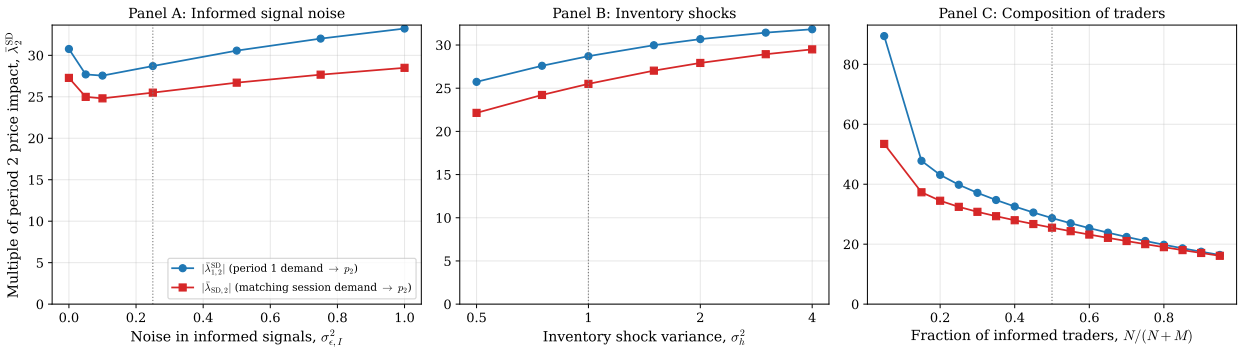

Figure 3. Cross-Period Price Impacts
We plot the population-weighted cross-period price impacts on the matching session branch as multiples of the population-weighted period 2 price impact, $\bar{\lambda}_2^{SD}$, with $\theta = 0.1$. The blue line reports the price impact of period 1 demand on the period 2 price, $|\bar{\lambda}_{1,2}^{SD}|$, and the red line reports the price impact of the matching session demand on the period 2 price, $|\bar{\lambda}_{SD,2}|$. Population weights are $N/(N+M)$ for informed and $M/(N+M)$ for uninformed traders. Panel A varies the noise in the informed's signals, $\sigma^2_{\epsilon,I}$, Panel B varies the variance of inventory shocks, σ^2_h, and Panel C varies the fraction of informed traders, $N/(N+M)$, holding $N+M=40$ fixed. Dotted vertical lines mark the baseline parameterization, and all other parameters are at their baseline values.

and the information in Q therefore reinforce each other. Because Q is mostly a signal about inventories, a deviation in the matching session moves beliefs and the period 2 price by a lot, and because it does, the informed trade little on their signals there, which keeps Q a signal about inventories.

Figure 3 shows how the cross-period price impacts vary with the parameters of the model. We scale each by the period 2 price impact, so that it measures how much more a deviation moves the period 2 price through the beliefs of other traders than a period 2 trade does. In every parameterization, both cross-period price impacts are several times as large as the period 2 price impact, so the deterrent against deviating in the matching session is large. The size of the deterrent depends on how much the uninformed rely on the public signal. Apart from when the informed's signals are nearly exact, both cross-period price impacts increase with the noise in the informed's signals. When the informed know less, they also learn from Q and p_1, so more traders react to any deviation. Both decrease with the fraction of informed traders, since there are fewer uninformed traders whose period 2 demands respond

to Q and p_1. With two informed traders, the cross-period price impact of period 1 demand is about 90 times the period 2 price impact. Larger inventory shocks increase the cross-period price impacts. The price impact of the matching session demand increases with the variance of inventory shocks, since Q is a signal about inventories and the uninformed place more weight on it when inventory shocks are larger. The price impact of period 1 demand is nearly unchanged, but the period 2 price impact falls, since more of period 2 trade is driven by inventories rather than information.

When the cross-period price impact of period 1 demand is large, traders shade their period 1 demand more in anticipation of the matching session. As Table IV shows, the increase in period 1 price impact from the matching session is largest when there are few informed traders, which coincides with when the cross-period price impact of period 1 demand is largest. Cross-period price impact reduces aggressiveness in period 1, which increases price impact during period 1 itself. However, with forty traders, the increase is small in every parameterization.

E.2. Why the Matching Session Is Informative

The period 1 price combines two sources of information: the value of the asset and the inventories of other traders. To see this, consider the case in which informed traders observe the fundamental value exactly, $\sigma^2_{\epsilon,I} = 0$, and uninformed traders have no private signal, and write the period 1 demand as $D^i_1 = a^\tau_{1,S} S^i_1 + a^\tau_{1,h} h^i_1 + a^\tau_{1,p} p_1$. Market clearing then gives

$$p_1 = \ell_v \, v_1 + \ell_I \, H_I + \ell_U \, H_U, \qquad \ell_v := \frac{N \, a^I_{1,S}}{A}, \qquad \ell_\tau := \frac{a^\tau_{1,h}}{A}, \tag{23}$$

where $H_\tau = \sum_{j \in \tau} h^j_1$ is the aggregate inventory shock of type τ and $A := -\big(N a^I_{1,p} + M a^U_{1,p}\big) > 0$ is the aggregate slope of period 1 demand. In equilibrium, demand increases in the informed's signal and decreases in each trader's inventory shock, as the period 1 demand schedule (19) suggests and Table I confirms, so that $\ell_v > 0$ and $\ell_\tau < 0$. An uninformed trader who observes a low price cannot discern whether the value of the asset is low or whether other traders hold large inventories. A second price-clearing session would not help to separate the components of the price, since its price aggregates traders' valuations in the same way and so reveals nearly the same combination of fundamental value and inventories.[9]

[9] In a numerical counterfactual in which the matching session is replaced by a second uniform-price double auction at stage 1B, the two prices are nearly collinear, and a linear equilibrium fails to exist unless new

The matching session's demands, on the other hand, weight value and inventory differently from the price. Let m_S^τ and m_h^τ denote the loadings of the matching session demand on the signal and on the inventory shock, as reported in Table I. Removing the part of Q that is already known from the price, the new information in Q is

$$Q^\perp = r_I\, H_I + r_U\, H_U, \qquad r_\tau := m_h^\tau - \frac{m_S^I}{a_{1,S}^I}\, a_{1,h}^\tau. \tag{24}$$

Given the price, Q is a pure inventory signal. This calculation requires only public information and knowledge of the equilibrium strategies. If matching session demands were simply rescaled copies of period 1 demands, then $r_\tau = 0$ and Q would reveal no new information. Matching session demands instead tilt toward inventory for two reasons. First, because trade settles at a fixed price, traders use the matching session to shed most of the inventory left after period 1, while their signal enters only through the gap between their bid-shaded period 1 demand and their desired position. Second, as discussed in Section IIIE.1, the price impact of the matching session demand deters the informed from trading on their signals. Without this deterrent, the informed would place about two and a half times as much weight on their signal in the matching session.

The uninformed therefore have two public signals and three unknowns. Q tells them about the inventories of other traders, and once they account for those inventories, the period 1 price tells them about the value of the asset. Revelation would be exact if the two types loaded on their inventories in the same way, $r_I/r_U = \ell_I/\ell_U$. Because the types load somewhat differently, some uncertainty remains, but the matching session removes most of it: at the baseline, the conditional variance of the period 2 value for the uninformed falls from 0.51 to 0.11 (Table III). In this special case, only the uninformed learn, since the informed already know the value. In the baseline parameterization, the signals of the informed are noisy, so the informed also learn from Q, as the negative weight on Q in Table I shows.

inventory shocks arrive between the two sessions. When it exists, the second price is nearly uninformative.

IV. Welfare and the Informational Role of the Matching Session

A. *Welfare*

We now analyze the welfare implications of the model. Specifically, we define aggregate welfare as the sum of traders' ex-ante expected payoffs. Because payments for trade sum to zero across traders by market clearing and budget balance, and because $\sum_i z_1^i = \sum_i h_1^i$ and $\sum_i z_2^i = \sum_i (h_1^i + h_2^i)$ are independent of the value of the asset, neither payments nor the liquidation value of positions contribute to aggregate welfare. Aggregate welfare is therefore the negative of expected holding costs. Aggregate welfare, normalized by the total number of traders, is therefore

$$W = -\frac{\gamma}{2}\left(\frac{N}{N+M}\times\mathbb{E}_I\left[(1-\delta)\left(z_1^i\right)^2+\delta\left(z_2^i\right)^2\right]+\frac{M}{N+M}\times\mathbb{E}_U\left[(1-\delta)\left(z_1^i\right)^2+\delta\left(z_2^i\right)^2\right]\right)]. \quad (25)$$

Since traders within a type are ex-ante identical, the efficient allocation splits the aggregate inventory equally, and welfare is higher when inventories are more evenly distributed across traders at the end of each period. With a probabilistic matching session, $W = (1-\theta)\, W^{\text{no}} + \theta\, W^{\text{SD}}$, where W^{SD} and W^{no} denote welfare on the matching session and no matching session branches. Payments net to zero only in the aggregate. Within a type, expected payments need not be zero, since informed traders trade against uninformed traders at favorable prices. We therefore define the welfare of a trader of type τ, W^τ, as their ex-ante expected payoff in (3), which includes their transfers from trading and the liquidation values of their positions.

The matching session affects welfare through three channels. First, it equalizes inventories: the matching session moves each trader toward the cross-sectional average inventory at the end of period 1, which lowers holding costs directly. Second, it provides information: the published aggregate Q improves the uninformed's estimate of the value of the asset at period 2, which lowers adverse selection and period 2 price impact, as shown in Section IIIE.1. Third, it distorts period 1 trade: anticipating the matching session, traders shade their period 1 demand schedules, which raises period 1 price impact and leaves traders with larger inventory imbalances whenever the matching session does not take place.

Table III reports welfare and price informativeness with and without the matching session. When there is a matching session, the conditional variance of the period 2 value for the uninformed falls from 0.514 to 0.110, reflecting the substantial information that Q adds

Table IIIWelfare and Price Informativeness
We report welfare by type and in aggregate, on each branch and combined, with a probabilistic matching session, $\theta = 0.1$, and without a matching session, $\theta = 0$. The row labeled "combined" is $(1-\theta)\,W^{\text{no}} + \theta\,W^{\text{SD}}$. $\text{Var}(v_2 \mid \mathcal{F}_U^{(2)})$ is the conditional variance of the period 2 value for an uninformed trader given their period 2 information set, and $\text{Var}(v_2 \mid p_2)$ is the conditional variance of the period 2 value given the period 2 price. The table uses the baseline parameterization.

	Matching ($\theta = 0.1$)	None ($\theta = 0$)	Difference
$W^{I,\text{SD}}$	-0.0143	-0.0299	$+0.0155$
$W^{U,\text{SD}}$	-0.0901	-0.1151	$+0.0250$
$W^{I,\text{no}}$	-0.0338	-0.0299	-0.0039
$W^{U,\text{no}}$	-0.1146	-0.1151	$+0.0006$
W^{I} (combined)	-0.0318	-0.0299	-0.0019
W^{U} (combined)	-0.1121	-0.1151	$+0.0030$
W^{SD}	-0.05223	-0.07250	$+0.02027$
W^{no}	-0.07416	-0.07250	-0.00165
W (combined)	-0.07196	-0.07250	$+0.00054$
$\text{Var}(v_2 \mid \mathcal{F}_U^{(2)})^{\text{SD}}$	0.1095	0.5138	-0.4042
$\text{Var}(v_2 \mid \mathcal{F}_U^{(2)})^{\text{no}}$	0.5034	0.5138	-0.0104
$\text{Var}(v_2 \mid p_2)^{\text{SD}}$	0.2693	0.5465	-0.2772
$\text{Var}(v_2 \mid p_2)^{\text{no}}$	0.5403	0.5465	-0.0062

to their period 2 information set. The period 2 price also becomes substantially more informative. Conditional on the matching session taking place, welfare is much higher for both types, and the uninformed capture a larger fraction of the aggregate gains. The large welfare gain on conditional on a matching session occurring is realized with probability $\theta = 0.1$. When there is no matching session, which occurs with probability $1-\theta = 0.9$, period 1 trade is distorted by the anticipation of the matching session, and welfare is lower than without a matching session. With forty traders, the distortion is small, and the loss when there is no matching session does not outweigh the gain on when a matching session occurs. Welfare differs across the two types. Combined across branches, the welfare of the uninformed increases by 0.003, while the welfare of the informed decreases by 0.002. Anticipating the matching session, the informed trade slightly more aggressively on their signals at period 1, so the period 1 price is more informative on both branches. This lowers the trading profits

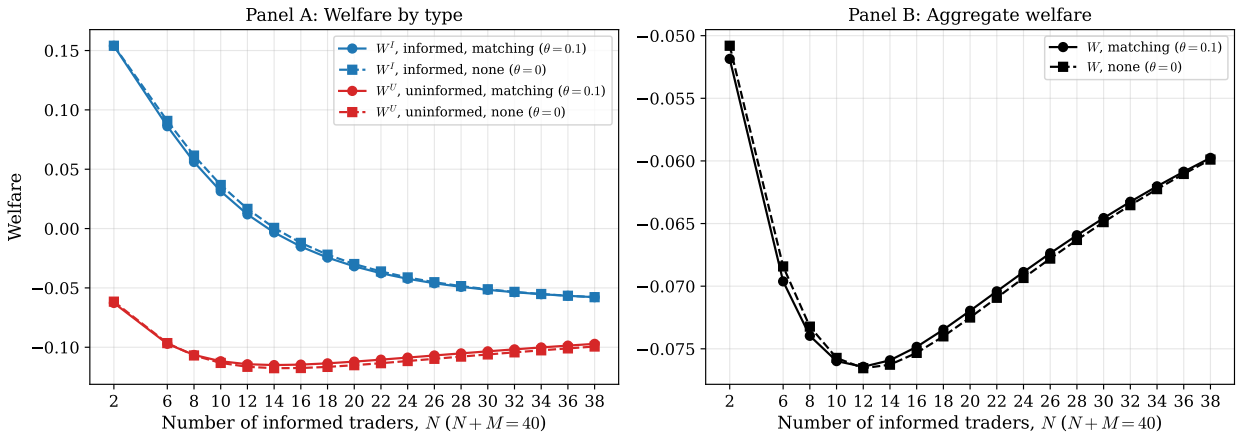

Figure 4. Welfare as a Function of the Composition of Traders
We plot welfare by type (Panel A) and aggregate welfare (Panel B) as a function of the number of informed traders, N, holding $N + M = 40$ fixed. Solid lines report the combined welfare with a probabilistic matching session, $\theta = 0.1$, and dashed lines report the benchmark without a matching session, $\theta = 0$. All other parameters are at the baseline parameterization.

of the informed and raises the welfare of the uninformed even when the matching session does not take place. The informed are better off than the uninformed with and without the matching session, even though they bear higher holding costs, because they earn trading profits at the expense of the uninformed.

B. Comparative Statics

We now study how the effects of the matching session vary with the parameters of the model. Figure 4 and Panel A of Table IV vary the composition of informed and uninformed traders, holding the total number of traders fixed. The matching session lowers period 2 price impact and the conditional variance of the period 2 value for the uninformed when there is a matching session. The net welfare effect is negative with ten or fewer informed traders and positive with twelve or more, and it is largest when the market is evenly split. With few informed traders, the matching session is most informative, but the increase in period 1 price impact from the anticipation of the matching session is also largest, and the loss when there is no matching session outweighs any gains. With more informed traders, period 1 price impact barely increases, and the gain when there is no matching session

dominates. Aggregate welfare is lowest when about a third of traders are informed, both with and without a matching session. Welfare differs across the two types. The welfare of the informed falls as the number of informed traders increases, as they compete away their trading profits, while the welfare of the uninformed is lowest when about fourteen traders are informed. The matching session increases the welfare of the uninformed whenever at least eight traders are informed, and it lowers the welfare of the informed for nearly every composition, with the largest loss when about eight traders are informed.

Panel B varies the probability of a matching session. The welfare effect is positive for every value of θ. It increases in θ at first, but the gain flattens and peaks at $\theta = 0.25$, since the increase in period 1 price impact grows faster than linearly in θ. Period 1 price impact increases in θ, since the more likely the matching session, the more traders shade their period 1 demand schedules in anticipation.

Panel C varies the variance of the innovation in the value of the asset between periods. As the innovation becomes larger, the information gap between the informed and uninformed grows from the matching session to date 2 trade. The reduction in the conditional variance of the period 2 value shrinks. The welfare gain from the matching session reduces as a function of the variance of the innovations.

Panel D varies the probability that the asset survives to period 2. The welfare gain increases in δ. As period 2 receives more weight, the information in Q becomes more important as the liquidation probability before it can be exploited decreases. The incentive to postpone trade to the matching session also grows, but with many traders the increase in period 1 price impact remains small.

Panel E varies the variance of inventory shocks. The welfare gain is largest at the benchmark parameterization and smaller for both smaller and larger inventory shocks. With larger inventory shocks, Q is more informative, and the learning component rises at first. The equalization component, however, is negative, and it becomes more negative as inventory shocks grow. The transfer is more valuable when traders have larger imbalances to equalize, but there is also more inventory adjustment to postpone, and the cost of this delay, which is paid whenever the matching session does not take place, grows by more. For large enough inventory shocks, this outweighs the additional learning, and the net gain shrinks.

Finally, Panel F varies the noise in the uninformed's signals. As the uninformed's signals become noisier, the matching session is more informative, and the reduction in their condi-

tional variance from observing Q grows. The welfare gain increases as well. The matching session improves welfare in every parameterization in Table IV, except when few traders are informed, where the increase in period 1 price impact from the anticipation of the matching session is largest.

V. Conclusion

This paper studies a dynamic model of strategic traders with heterogeneous information, in which a matching session, where further trade can take place at the fixed prior market-clearing price, may follow the first exchange session. The resulting cross-period price and demand impacts imply that the model is substantially more challenging to solve than stationary models of strategic trade. To address these challenges, we employ the methods of Rostek et al. (2025) in order to tractably compute the cross-period demand and price impacts.

The results show that matching sessions play an important informational role. The matching session demands weighs inventories and the value of the asset differently from the exchange price. It therefore provides information beyond the price, reducing adverse selection and price impact in the subsequent exchange session and improving price informativeness. However, the anticipation of the matching session leads traders to postpone trade from the preceding exchange session. When there are enough traders, this distortion is small, and the matching session increases welfare.

Table IVComparative Statics

We report the change in each variable from the benchmark without a matching session to the equilibrium with a probabilistic matching session, varying one parameter at a time. $\Delta\bar{\lambda}_1$ is the change in the population-weighted period 1 price impact, and $\Delta\bar{\lambda}_2^{\text{SD}}$ is the change in the population-weighted period 2 price impact on the matching session branch. $\Delta\text{Var}(v_2 \mid \mathcal{F}_U^{(2)})^{\text{SD}}$ and $\Delta\text{Var}(v_2 \mid p_2)^{\text{SD}}$ are the changes in the conditional variances of the period 2 value on the matching session branch. ΔW is the change in combined aggregate welfare. The number of informed traders is varied holding $N + M = 40$ fixed. Dashes indicate parameterizations for which we find no equilibrium with a matching session. All other parameters are at the baseline parameterization.

Value	$\Delta\bar{\lambda}_1$	$\Delta\bar{\lambda}_2^{\text{SD}}$	$\Delta\text{Var}(v_2 \mid \mathcal{F}_U^{(2)})^{\text{SD}}$	$\Delta\text{Var}(v_2 \mid p_2)^{\text{SD}}$	ΔW
Panel A: Number of informed traders, N					
8	+0.0047	−0.0087	−0.755	−0.614	−0.00073
16	+0.0010	−0.0108	−0.493	−0.365	+0.00048
24	+0.0002	−0.0094	−0.337	−0.211	+0.00048
32	+0.0002	−0.0069	−0.247	−0.122	+0.00027
Panel B: Probability of a matching session, θ					
0.025	+0.0001	−0.0103	−0.404	−0.277	+0.00015
0.1	+0.0004	−0.0103	−0.404	−0.277	+0.00054
0.2	+0.0012	−0.0103	−0.404	−0.277	+0.00086
0.3	+0.0026	−0.0103	−0.404	−0.277	+0.00085
Panel C: Variance of the value innovation, σ_η^2					
0.05	+0.0004	−0.0099	−0.418	−0.292	+0.00059
0.2	+0.0004	−0.0092	−0.378	−0.254	+0.00042
0.5	+0.0005	−0.0048	−0.328	−0.218	+0.00018
Panel D: Probability the asset survives to period 2, δ					
0.5	+0.0004	−0.0103	−0.404	−0.277	+0.00054
0.7	+0.0009	−0.0083	−0.392	−0.270	+0.00081
Panel E: Inventory shock variance, σ_h^2					
0.5	+0.0005	−0.0163	−0.278	−0.212	+0.00039
1	+0.0004	−0.0103	−0.404	−0.277	+0.00054
4	+0.0006	−0.0033	−0.694	−0.311	+0.00010
Panel F: Noise in uninformed signals, $\sigma_{\epsilon,U}^2$					
0.5	+0.0003	−0.0068	−0.080	−0.095	+0.00018
4	+0.0004	−0.0109	−0.274	−0.209	+0.00043
16	+0.0004	−0.0108	−0.361	−0.255	+0.00051
∞	+0.0004	−0.0103	−0.404	−0.277	+0.00054

VI. Works Cited

Antill, Samuel, and Darrell Duffie, 2020, Augmenting Markets with Mechanisms, *The Review of Economic Studies* 88, 1665–1719.

Blonien, Patrick, 2026, Size Discovery in Slow Markets, *SSRN* .

Blonien, Patrick, and Alexander Ober, 2026, Is 24/7 Trading Better?, *SSRN* .

Blume, Lawrence, David Easley, and Maureen O'Hara, 1994, Market statistics and technical analysis: The role of volume, *Journal of Finance* 49, 153–181.

Buti, Sabrina, Barbara Rindi, and Ingrid M. Werner, 2017, Dark pool trading strategies, market quality, and welfare, *Journal of Financial Economics* 124, 244–265.

Chen, Daniel, 2022, Optimal Design of a Financial Exchange, *SSRN* .

Chen, Daniel, and Darrell Duffie, 2021, Market fragmentation, *American Economic Review* 111, 2247–2274.

Comerton-Forde, Carole, and Tālis J. Putniņš, 2015, Dark trading and price discovery, *Journal of Financial Economics* 118, 70–92.

Degryse, Hans, Frank de Jong, and Vincent van Kervel, 2014, The Impact of Dark Trading and Visible Fragmentation on Market Quality, *Review of Finance* 19, 1587–1622.

Du, Songzi, and Haoxiang Zhu, 2017, What is the Optimal Trading Frequency in Financial Markets?, *The Review of Economic Studies* 84, 1606–1651.

Duffie, Darrell, and Haoxiang Zhu, 2017, Size Discovery, *The Review of Financial Studies* 30, 1095–1150.

Fleming, Michael J, and Giang Nguyen, 2018, Price and Size Discovery in Financial Markets: Evidence from the U.S. Treasury Securities Market, *The Review of Asset Pricing Studies* 9, 256–295.

Ganguli, Jayant Vivek, and Liyan Yang, 2009, Complementarities, multiplicity, and supply information, *Journal of the European Economic Association* 7, 90–115.

Grossman, Sanford J., 1976, On the efficiency of competitive stock markets where trades have diverse information, *Journal of Finance* 31, 573–585.

Hatheway, Frank, Amy Kwan, and Hui Zheng, 2017, An empirical analysis of market segmentation on U.S. equity markets, *Journal of Financial and Quantitative Analysis* 52, 2399–2427.

Hayek, Friedrich A., 1945, The use of knowledge in society, *American Economic Review* 35, 519–530.

Hendershott, Terrence, and Haim Mendelson, 2000, Crossing networks and dealer markets: Competition and performance, *Journal of Finance* 55, 2071–2115.

Klemperer, Paul D., and Margaret A. Meyer, 1989, Supply function equilibria in oligopoly under uncertainty, *Econometrica* 57, 1243–1277.

Kyle, Albert S., 1989, Informed Speculation with Imperfect Competition, *The Review of Economic Studies* 56, 317–355.

Manzano, Carolina, and Xavier Vives, 2011, Public and private learning from prices, strategic substitutability and complementarity, and equilibrium multiplicity, *Journal of Mathematical Economics* 47, 346–369.

Menkveld, Albert J., Bart Zhou Yueshen, and Haoxiang Zhu, 2017, Shades of darkness: A pecking order of trading venues, *Journal of Financial Economics* 124, 503–534.

Rostek, Marzena, and Marek Weretka, 2012, Price inference in small markets, *Econometrica* 80, 687–711.

Rostek, Marzena J, Ji Hee Yoon, and Xian Wu, 2025, Dynamic imperfectly competitive markets: A non-recursive approach Working Paper.

Sannikov, Yuliy, and Andrzej Skrzypacz, 2016, Dynamic Trading: Price Inertia and Front-Running, *SSRN* .

Schneider, Jan, 2009, A rational expectations equilibrium with informative trading volume, *Journal of Finance* 64, 2783–2805.

Securities and Exchange Commission, 2010, Concept release on equity market structure, Technical Report Release No. 34-61358; File No. S7-02-10, Securities and Exchange Commission, 75 Fed. Reg. 3594.

Vayanos, Dimitri, 1999, Strategic Trading and Welfare in a Dynamic Market, *The Review of Economic Studies* 66, 219–254.

Vives, Xavier, 2011, Strategic supply function competition with private information, *Econometrica* 79, 1919–1966.

Zhu, Haoxiang, 2013, Do Dark Pools Harm Price Discovery?, *The Review of Financial Studies* 27, 747–789.

Appendix

A. FOC Derivations and Matching Conditions

This appendix derives the equilibrium FOCs and matching conditions that pin down the strategy coefficients of the linear strategies (70), (51), and (35). We first state the trader's problem. The derivation then proceeds backward through the timeline: period 2 (terminal exchange), matching session stage 1B, then period 1 (initial exchange). We describe the equilibrium assuming there is a matching session for certain. The probabilistic matching session is an extension.

A. The Trader's Problem

Throughout, v_t denotes the common value of the asset at period t, which follows a random walk, p_t the period t exchange price, D_t^i trader i's period t exchange demand, S_t^i trader i's private signal about v_t, and h_t^i trader i's period t inventory shock. At stage 1B, trader i submits a matching session demand μ_i, the aggregate demand $Q = \sum_j \mu_j$ is published, and trader i receives the matching session transfer $D_{SD}^i = \mu_i - Q/(N+M)$, their demand net of the average demand, at the period 1 price p_1. Inventories evolve as

$$\hat{z}_1^i = h_1^i + D_1^i, \qquad z_1^i = \hat{z}_1^i + D_{SD}^i,$$
$$\hat{z}_2^i = z_1^i + h_2^i, \qquad z_2^i = \hat{z}_2^i + D_2^i,$$

so that $\hat{z}_t^i$ is the inventory immediately before, and z_t^i the inventory after, the last trading session of period t. There are N informed traders and M uninformed traders. When a formula applies to either type, n_τ denotes the number of traders of type $\tau \in \{I, U\}$, so that $n_I = N$ and $n_U = M$, and $-\tau$ denotes the type other than τ.

Trader i of type τ maximizes

$$\mathbb{E}\Big[- p_1\big(D_1^i + D_{SD}^i\big) + (1-\delta)\Big(v_1\big(\hat{z}_1^i + D_1^i + D_{SD}^i\big) - \frac{\gamma}{2}\big(\hat{z}_1^i + D_1^i + D_{SD}^i\big)^2\Big) + \delta\Big(- p_2 D_2^i + v_2\big(\hat{z}_2^i + D_2^i\big) - \frac{\gamma}{2}\big(\hat{z}_2^i + D_2^i\big)^2\Big)\Big]. \quad (26)$$

Holding inventory z costs $\frac{\gamma}{2}z^2$, so the marginal holding cost is γz. Denote the ex-post realized continuation value as

$$V_{i,2}(\hat{z}_2^i, D_2^i, p_2) = -p_2 D_2^i + v_2\big(\hat{z}_2^i + D_2^i\big) - \frac{\gamma}{2}\big(\hat{z}_2^i + D_2^i\big)^2.$$

Therefore, before period 2, traders maximize

$$\mathbb{E}\Big[-p_1\big(D_1^i + D_{SD}^i\big) + (1-\delta)\Big(v_1 z_1^i - \frac{\gamma}{2}\big(z_1^i\big)^2\Big) + \delta V_{i,2}(\hat{z}_2^i, D_2^i, p_2)\Big].$$

Trader i of type τ has information set $\mathcal{F}_i^{(1),\tau}$ at period 1, $\mathcal{F}_i^{(\mathrm{SD}),\tau}$ at stage 1B, and $\mathcal{F}_i^{(2),\tau}$ at period 2, defined in (37) and (27). Because traders submit demand schedules, the period t information set includes the clearing price p_t.

Define the own-period price impacts of trader i of type τ to be the constants $\lambda_1^\tau := \frac{\mathrm{d}p_1}{\mathrm{d}D_1^i}$ and $\lambda_2^\tau := \frac{\mathrm{d}p_2}{\mathrm{d}D_2^i}$, and the cross-period price impacts to be the constants $\lambda_{1,2}^\tau := \frac{\mathrm{d}p_2}{\mathrm{d}D_1^i}$ and $\lambda_{\mathrm{SD},2}^\tau := \frac{\mathrm{d}p_2}{\mathrm{d}\mu_i}$. These are determined in Appendix AE.4.

B. period 2 FOC

At period 2, trader i of type τ has information set

$$\mathcal{F}_i^{(2),\tau} = \big\{S_1^i,\ S_2^i,\ p_1,\ p_2,\ h_1^i,\ h_2^i,\ Q\big\}, \tag{27}$$

where $Q = \sum_j \mu_j$ is the aggregate demand in the matching session.

Conjecture that trader i of type τ has linear beliefs at period 2. The projection of v_2 onto $\mathcal{F}_i^{(2),\tau}$ is

$$\mathbb{E}[v_2 \mid \mathcal{F}_i^{(2),\tau}] = c_{S_1,2}^\tau S_1^i + c_{S_2,2}^\tau S_2^i + c_{p_1,2}^\tau p_1 + c_{p_2,2}^\tau p_2 + c_{h_1,2}^\tau h_1^i + c_{h_2,2}^\tau h_2^i + c_{Q,2}^\tau Q, \tag{28}$$

a function of seven coefficients.

FOC and demand schedule. Differentiating (26) with respect to D_2^i and taking the conditional expectation given $\mathcal{F}_i^{(2),\tau}$, the common factor δ drops out and the FOC at period 2 is

$$\mathbb{E}\big[v_2 - \gamma\big(\hat{z}_2^i + D_2^i\big) \;\big|\; \mathcal{F}_i^{(2),\tau}\big] = p_2 + \lambda_2^\tau D_2^i. \tag{29}$$

The price-impact term $\lambda_2^\tau D_2^i$ captures the trader's own marginal effect on p_2 via market clearing. Solving (29) for D_2^i:

$$\implies D_2^i(p_2) = \chi_2^\tau\left(\frac{1}{\gamma}\big(\mathbb{E}[v_2 \mid \mathcal{F}_i^{(2),\tau}] - p_2\big) - \hat{z}_2^i\right), \qquad \chi_2^\tau := \frac{\gamma}{\gamma + \lambda_2^\tau}. \tag{30}$$

The trader's period 2 demand has two components: a valuation-price gap ($\mathbb{E}[v_2 \mid \mathcal{F}_i^{(2),\tau}] - p_2$), scaled by the marginal cost of holding more inventory, and an exit from their current inventory position ($-\hat{z}_2^i$). The term in the large parentheses is the demand a trader would submit in a perfectly competitive market. It is scaled by $\chi_2^\tau \in [0,1]$, which bid shades

demand due to imperfect competition.

Market clearing and own-period price impact. Aggregating $\sum_j D_2^j(p_2) = 0$, the period 2 clearing price is

$$p_2^* = \frac{\sum_\tau n_\tau \chi_2^\tau \left(\frac{1}{\gamma}\overline{V}_2^\tau - \overline{\hat{z}}_2^\tau\right)}{\sum_\tau n_\tau \chi_2^\tau \frac{1}{\gamma}(1 - c_{p_2,2}^\tau)}, \tag{31}$$

where $\overline{V}_2^\tau$ is the average across type-τ traders of their conditional valuation $\mathbb{E}[v_2 \mid \mathcal{F}_i^{(2),\tau}] - c_{p_2,2}^\tau p_2$ (stripping out the indirect dependence on p_2 through the conditional expectation), and $\overline{\hat{z}}_2^\tau$ is the type-average inventory.

The own-period price impact is a type-specific fixed point, $\lambda_2^\tau \equiv \frac{\mathrm{d}p_2}{\mathrm{d}D_2^i} = -\left(\sum_{j\neq i} \frac{\partial D_2^j}{\partial p_2}\right)^{-1}$, which is

$$\lambda_2^I = \frac{\gamma}{(N-1)\chi_2^I(1 - c_{p_2,2}^I) + M\chi_2^U(1 - c_{p_2,2}^U)}, \tag{32}$$

$$\lambda_2^U = \frac{\gamma}{N\chi_2^I(1 - c_{p_2,2}^I) + (M-1)\chi_2^U(1 - c_{p_2,2}^U)}, \tag{33}$$

which, after substituting $\chi_2^\tau = \gamma(\gamma + \lambda_2^\tau)^{-1}$, is a non-linear system in $(\lambda_2^I, \lambda_2^U)$.

Demand slope with price feedback. Because $\mathbb{E}[v_2 \mid \mathcal{F}_i^{(2),\tau}]$ depends on p_2 through $c_{p_2,2}^\tau$, the trader's effective demand slope on p_2 is

$$b_2^\tau := \left|\frac{\partial D_2^i}{\partial p_2}\right| = \chi_2^\tau \frac{1 - c_{p_2,2}^\tau}{\gamma}. \tag{34}$$

The numerator $1 - c_{p_2,2}^\tau$ measures the trader's net sensitivity to a unit increase in p_2: the cost of buying rises by 1, but the conditional expectation of v_2 rises by $c_{p_2,2}^\tau$ (the price carries information).

Strategy matching conditions. The period 2 strategy of trader i of type τ is the linear demand

$$D_2^i = \sum_\xi a_{2,\xi}^\tau \, \xi = a_{2,S_1}^\tau S_1^i + a_{2,S_2}^\tau S_2^i + a_{2,h_1}^\tau h_1^i + a_{2,h_2}^\tau h_2^i + a_{2,p_1}^\tau p_1 + a_{2,p_2}^\tau p_2 + a_{2,Q}^\tau Q, \tag{35}$$

where ξ ranges over the $\mathcal{F}^{(2)}$-basis directions $\{S_1, S_2, h_1, h_2, p_1, p_2, Q\}$. Substituting into (29) and matching coefficients yields

$$a^\tau_{2,\xi} = \chi^\tau_2 \left(\frac{1}{\gamma} \big(c^\tau_{\xi,2} - \mathbf{1}\{\xi = p_2\} \big) - z^\tau_\xi \right), \tag{36}$$

where z^τ_ξ is the loading of $\hat{z}^i_2$ on direction ξ in the $\mathcal{F}^{(2)}$ basis.

The z^τ_ξ loadings depend on the design. Letting $\hat{z}^i_2 = z^i_1 + h^i_2$, we have:

No Matching Session. Without matching session, $z^i_1 = h^i_1 + D^i_1$, so

$$\begin{aligned} z^\tau_{S_1} &= a^\tau_{1,S}, & z^\tau_{S_2} &= 0, \\ z^\tau_{h_1} &= 1 + a^\tau_{1,h}, & z^\tau_{h_2} &= 1, \\ z^\tau_{p_1} &= a^\tau_{1,p}, & z^\tau_{p_2} &= 0, \\ z^\tau_Q &= 0. \end{aligned}$$

Matching Session. With a matching session, $z^i_1 = h^i_1 + D^i_1 + \mu_i - Q/(N+M)$. Let $\mu_i = b^\tau_z \hat{z}^i_1 + b^\tau_S S^i_1 + b^\tau_h h^i_1$. Using z^i_1's primitive representation, the coefficient of z^i_1 on direction ξ in the $\mathcal{F}^{(2)}$ basis yields

$$\begin{aligned} z^\tau_{S_1} &= a^\tau_{1,S}\,(1 + b^\tau_z) + b^\tau_S, & z^\tau_{S_2} &= 0, \\ z^\tau_{h_1} &= (1 + a^\tau_{1,h})(1 + b^\tau_z) + b^\tau_h, & z^\tau_{h_2} &= 1, \\ z^\tau_{p_1} &= a^\tau_{1,p}(1 + b^\tau_z), & z^\tau_{p_2} &= 0, \\ z^\tau_Q &= -1/(N+M). \end{aligned}$$

Equations for intuition. Define the period 2 weight of type τ as

$$\omega^\tau_2 := \frac{\frac{n_\tau}{\gamma + \lambda^\tau_2}}{\frac{M}{\gamma + \lambda^U_2} + \frac{N}{\gamma + \lambda^I_2}},$$

and let $\overline{\hat{z}}^\tau_2$ denote the type-average inventory, as in (31). Aggregating (30) across traders, the period 2 clearing price can equivalently be written (implicitly, since the conditional expectations depend on p_2) as

$$p^*_2 = \sum_{\tau \in \{I,U\}} \omega^\tau_2 \left(\frac{1}{n_\tau} \sum_{j \in \tau} \mathbb{E}[v_2 \mid \mathcal{F}^{(2),\tau}_j] - \gamma\, \overline{\hat{z}}^\tau_2 \right),$$

where $j \in \tau$ denotes the traders of type τ. The equilibrium demand of trader i of type τ is then

$$D_2^i(p_2^*) = \chi_2^\tau \left[\frac{1}{\gamma} \left(\mathbb{E}[v_2 \mid \mathcal{F}_i^{(2),\tau}] - \sum_{\tau' \in \{I,U\}} \omega_2^{\tau'} \frac{1}{n_{\tau'}} \sum_{j \in \tau'} \mathbb{E}[v_2 \mid \mathcal{F}_j^{(2),\tau'}] \right) - \left(\hat{z}_2^i - \sum_{\tau' \in \{I,U\}} \omega_2^{\tau'} \overline{\hat{z}}_2^{\tau'} \right) \right].$$

The post-trade inventory is then

$$\hat{z}_2^i + D_2^i(p_2^*) = \frac{1}{\gamma + \lambda_2^\tau} \left(\mathbb{E}[v_2 \mid \mathcal{F}_i^{(2),\tau}] - \sum_{\tau' \in \{I,U\}} \omega_2^{\tau'} \frac{1}{n_{\tau'}} \sum_{j \in \tau'} \mathbb{E}[v_2 \mid \mathcal{F}_j^{(2),\tau'}] \right) + \frac{\lambda_2^\tau}{\gamma + \lambda_2^\tau} \hat{z}_2^i + \frac{\gamma}{\gamma + \lambda_2^\tau} \sum_{\tau' \in \{I,U\}} \omega_2^{\tau'} \overline{\hat{z}}_2^{\tau'}.$$

C. Matching Session FOC (Stage 1B)

At stage 1B, trader i submits a matching session demand μ_i given information set

$$\mathcal{F}_i^{(\mathrm{SD}),\tau} = \mathcal{F}_i^{(1),\tau} = \{S_1^i,\ h_1^i,\ p_1\}, \tag{37}$$

where p_1 is the just-cleared stage-1A price. Then simultaneously, the sum $Q = \sum_j \mu_j$ is published and the matching session transfers

$$D_{SD}^i = \mu_i - Q/(N+M), \tag{38}$$

so that each trader receives their demand net of the average demand. The transfer is settled at the stage-1A price p_1, so that the post-matching session inventory is

$$z_1^i = \hat{z}_1^i + D_{SD}^i = h_1^i + D_1^i + \mu_i - Q/(N+M). \tag{39}$$

By construction, $\sum_j D_{SD}^j = 0$, so the matching session mechanism is budget-balanced.

Marginal effects of μ_i. Differentiating the various relevant equilibrium objects with respect to μ_i:

$$\frac{\mathrm{d}z_1^i}{\mathrm{d}\mu_i} = \frac{N+M-1}{N+M}, \tag{40}$$

$$\frac{\mathrm{d}\hat{z}_2^i}{\mathrm{d}\mu_i} = \frac{\mathrm{d}z_1^i}{\mathrm{d}\mu_i} = \frac{N+M-1}{N+M}, \tag{41}$$

$$\frac{\mathrm{d}D_{SD}^i}{\mathrm{d}\mu_i} = \frac{N+M-1}{N+M}, \tag{42}$$

$$\frac{\mathrm{d}p_2}{\mathrm{d}\mu_i} = \lambda_{\mathrm{SD},2}^\tau, \tag{43}$$

$$\frac{\mathrm{d}D_2^i}{\mathrm{d}\mu_i} = -\chi_2^\tau\Big(\frac{N+M-1}{N+M} + \frac{\lambda_{\mathrm{SD},2}^\tau}{\gamma}\Big). \tag{44}$$

The result (44) follows from implicit differentiation of the period 2 FOC (Appendix AE.1).

FOC. At stage 1B, the period 1 price p_1 has already cleared, so the matching session demand has no impact on p_1. Traders choose μ_i to maximize

$$\mathbb{E}\Big[- p_1 D_{SD}^i + (1-\delta)\Big(v_1(\hat{z}_1^i + D_{SD}^i) - \frac{\gamma}{2}\big(\hat{z}_1^i + D_{SD}^i\big)^2\Big) + \delta V_{i,2}(\hat{z}_2^i, D_2^i, p_2) \;\Big|\; \mathcal{F}_i^{(\mathrm{SD}),\tau}\Big],$$

hence their FOC is

$$\mathbb{E}\Big[(1-\delta)\,\frac{N+M-1}{N+M}\,\big(v_1 - \gamma(\hat{z}_1^i + D_{SD}^i)\big) + \delta\,\frac{\mathrm{d}V_{i,2}}{\mathrm{d}\mu_i} \;\Big|\; \mathcal{F}_i^{(\mathrm{SD}),\tau}\Big] = \frac{N+M-1}{N+M}\,p_1. \tag{45}$$

Using the period 2 FOC, $\mathbb{E}[v_2 - \gamma z_2^i - p_2 \mid \mathcal{F}_i^{(2),\tau}] = \lambda_2^\tau D_2^i$, the marginal continuation value is

$$\begin{aligned}
\frac{\mathrm{d}V_{i,2}}{\mathrm{d}\mu_i} &= -\frac{\mathrm{d}p_2}{\mathrm{d}\mu_i}D_2^i - p_2\frac{\mathrm{d}D_2^i}{\mathrm{d}\mu_i} + v_2\Big(\frac{N+M-1}{N+M} + \frac{\mathrm{d}D_2^i}{\mathrm{d}\mu_i}\Big) - \gamma\big(\hat{z}_2^i + D_2^i\big)\Big(\frac{N+M-1}{N+M} + \frac{\mathrm{d}D_2^i}{\mathrm{d}\mu_i}\Big) \\
&= \big(v_2 - \gamma z_2^i\big)\Big(\frac{N+M-1}{N+M} + \frac{\mathrm{d}D_2^i}{\mathrm{d}\mu_i}\Big) - p_2\frac{\mathrm{d}D_2^i}{\mathrm{d}\mu_i} - \lambda_{\mathrm{SD},2}^\tau D_2^i \\
&= \big(v_2 - \gamma z_2^i - p_2\big)\left(\frac{N+M-1}{N+M} + \frac{\mathrm{d}D_2^i}{\mathrm{d}\mu_i} - \frac{\lambda_{\mathrm{SD},2}^\tau}{\lambda_2^\tau}\right) + \frac{N+M-1}{N+M}\,p_2 \\
&= \big(v_2 - \gamma z_1^i - p_2\big)(1-\chi_2^\tau)\left(\frac{N+M-1}{N+M} + \frac{\mathrm{d}D_2^i}{\mathrm{d}\mu_i} - \frac{\lambda_{\mathrm{SD},2}^\tau}{\lambda_2^\tau}\right) + \frac{N+M-1}{N+M}\,p_2.
\end{aligned}$$

By the law of iterated expectations, we've implicitly eliminated mean 0 terms: the third equality uses the period 2 FOC, and the fourth uses the period 2 demand (30) together with $\hat{z}_2^i = z_1^i + h_2^i$ and $\mathbb{E}[h_2^i \mid \mathcal{F}_i^{(\mathrm{SD}),\tau}] = 0$. Note, therefore, that the FOC at stage 1B actually implies

$$\begin{aligned}
\frac{N+M-1}{N+M}\,p_1 = (1-\delta)\,\frac{N+M-1}{N+M}\,\mathbb{E}\big[v_1 - \gamma(\hat{z}_1^i + D_{SD}^i) \;\big|\; \mathcal{F}_i^{(\mathrm{SD}),\tau}\big] \\
+ \delta\,(1-\chi_2^\tau)\Big(\frac{N+M-1}{N+M} + \frac{\mathrm{d}D_2^i}{\mathrm{d}\mu_i} - \frac{\lambda_{\mathrm{SD},2}^\tau}{\lambda_2^\tau}\Big)\,\mathbb{E}\big[v_2 - \gamma(\hat{z}_1^i + D_{SD}^i) - p_2 \;\big|\; \mathcal{F}_i^{(\mathrm{SD}),\tau}\big] \\
+ \delta\,\frac{N+M-1}{N+M}\,\mathbb{E}\big[p_2 \;\big|\; \mathcal{F}_i^{(\mathrm{SD}),\tau}\big].
\end{aligned}$$

Using the random walk identity $\mathbb{E}[v_2 \mid \mathcal{F}_i^{(\mathrm{SD}),\tau}] = \mathbb{E}[v_1 \mid \mathcal{F}_i^{(\mathrm{SD}),\tau}]$ and combining terms, the FOC can be rewritten as

$$\mathbb{E}\big[v_1 - \gamma(\hat{z}_1^i + D_{SD}^i) \;\big|\; \mathcal{F}_i^{(\mathrm{SD}),\tau}\big] = \Omega_{\mathrm{SD}}^\tau\,p_1 + (1-\Omega_{\mathrm{SD}}^\tau)\,\mathbb{E}[p_2 \mid \mathcal{F}_i^{(\mathrm{SD}),\tau}], \tag{46}$$

where

$$\Omega^{\tau}_{\mathrm{SD}} = \frac{1}{1-\delta\big[1-(1-\chi^{\tau}_2)^2\big]-\delta\,(1-\chi^{\tau}_2)(2-\chi^{\tau}_2)\,\dfrac{(N+M)\,\lambda^{\tau}_{\mathrm{SD},2}}{(N+M-1)\,\lambda^{\tau}_2}}. \tag{47}$$

Solving for D^i_{SD}, we get

$$\implies \mathbb{E}\big[D^i_{SD}\;\big|\;\mathcal{F}^{(\mathrm{SD}),\tau}_i\big] = \frac{1}{\gamma}\Big(\mathbb{E}[v_1-p_2\mid\mathcal{F}^{(\mathrm{SD}),\tau}_i]-\Omega^{\tau}_{\mathrm{SD}}\big(p_1-\mathbb{E}[p_2\mid\mathcal{F}^{(\mathrm{SD}),\tau}_i]\big)\Big)-\hat{z}^i_1. \tag{48}$$

This expected allocation can be rewritten as the expected competitive allocation plus an additional term due to the imperfect competition,

$$\begin{aligned}\mathbb{E}\big[D^i_{SD}\;\big|\;\mathcal{F}^{(\mathrm{SD}),\tau}_i\big] = {}& \underbrace{\frac{1}{\gamma}\left(\mathbb{E}[v_1\mid\mathcal{F}^{(\mathrm{SD}),\tau}_i]-\frac{p_1-\delta\,\mathbb{E}[p_2\mid\mathcal{F}^{(\mathrm{SD}),\tau}_i]}{1-\delta}\right)-\hat{z}^i_1}_{\text{competitive demand}}\\ &+\underbrace{\frac{1}{\gamma}\Big(\frac{1}{1-\delta}-\Omega^{\tau}_{\mathrm{SD}}\Big)\big(p_1-\mathbb{E}[p_2\mid\mathcal{F}^{(\mathrm{SD}),\tau}_i]\big)}_{\text{imperfect-competition correction}}.\end{aligned} \tag{49}$$

The quantity submitted in the matching session is then

$$\mu_i = \frac{1}{\gamma}\Big(\mathbb{E}[v_1\mid\mathcal{F}^{(\mathrm{SD}),\tau}_i]-\Omega^{\tau}_{\mathrm{SD}}p_1-(1-\Omega^{\tau}_{\mathrm{SD}})\,\mathbb{E}[p_2\mid\mathcal{F}^{(\mathrm{SD}),\tau}_i]\Big)-\hat{z}^i_1+\frac{1}{N+M}\,\mathbb{E}\big[Q\;\big|\;\mathcal{F}^{(\mathrm{SD}),\tau}_i\big]. \tag{50}$$

Workup matching conditions. Conjecture the linear matching session demand

$$\mu_i = b^{\tau}_z\,\hat{z}^i_1 + b^{\tau}_S\,S^i_1 + b^{\tau}_h\,h^i_1. \tag{51}$$

The pre-matching session inventory $\hat{z}^i_1$ has $\mathcal{F}^{(\mathrm{SD})}$-basis loadings

$$z^{\tau}_{S_1} = a^{\tau}_{1,S},\quad z^{\tau}_{h_1} = 1+a^{\tau}_{1,h},\quad z^{\tau}_{p_1} = a^{\tau}_{1,p}. \tag{52}$$

Define the FOC-implied numerator ν^{τ}_{ξ} on each direction $\xi\in\{S,h,p_1\}$, that is, the loading on ξ of the term in parentheses in (48) after applying the random walk identity:

$$\nu^{\tau}_S = c^{\tau}_{S,1}-(1-\Omega^{\tau}_{\mathrm{SD}})\rho^{\tau}_{S,2}, \tag{53}$$

$$\nu^{\tau}_h = c^{\tau}_{h,1}-(1-\Omega^{\tau}_{\mathrm{SD}})\rho^{\tau}_{h,2}, \tag{54}$$

$$\nu^{\tau}_{p_1} = c^{\tau}_{p_1,1}-\Omega^{\tau}_{\mathrm{SD}}-(1-\Omega^{\tau}_{\mathrm{SD}})\rho^{\tau}_{p_1,2}. \tag{55}$$

The c projections are loadings of $\mathbb{E}[v_1\mid\mathcal{F}^{(\mathrm{SD}),\tau}_i]$, and the ρ projections are loadings of $\mathbb{E}[p_2\mid\mathcal{F}^{(\mathrm{SD}),\tau}_i]$.

Since $D^i_{SD}=\mu_i-Q/(N+M)$, the left-hand side of (48) is $\frac{N+M-1}{N+M}\,\mu_i-\mathbb{E}\big[\frac{1}{N+M}\sum_{j\neq i}\mu_j\;\big|$

$\mathcal{F}_i^{(\mathrm{SD}),\tau}\big]$. The matching condition on each direction $\xi \in \{S, h, p_1\}$ of $\mathcal{F}_i^{(\mathrm{SD}),\tau}$ is

$$\frac{N+M-1}{N+M}\big(b_z^\tau z_\xi^\tau + b_\xi^\tau \mathbf{1}\{\xi \neq p_1\}\big) - \pi_\xi^\tau = \frac{\nu_\xi^\tau}{\gamma} - z_\xi^\tau, \tag{56}$$

where π_ξ^τ is the rival-aggregate projection $\pi_\xi^\tau =$ loading of $\frac{1}{N+M}\sum_{j\neq i}\mu_j^*$ on direction ξ. The three matching equations (56) pin down the three coefficients (b_z, b_S, b_h):

$$\frac{N+M-1}{N+M}\, b_z^\tau\, a_{1,p}^\tau - \pi_{p_1}^\tau = \frac{\nu_{p_1}^\tau}{\gamma} - a_{1,p}^\tau, \tag{57}$$

$$b_\xi^\tau = \frac{N+M}{N+M-1}\Big(\frac{\nu_\xi^\tau}{\gamma} - z_\xi^\tau + \pi_\xi^\tau\Big) - b_z^\tau\, z_\xi^\tau, \qquad \xi \in \{S, h\}. \tag{58}$$

The b_z coefficient is determined by matching p_1.

D. period 1 FOC

At period 1, trader i submits a demand schedule $D_1^i(p)$ prior to observing p_1. The realized $D_1^i = D_1^i(p_1)$ becomes a function of $\mathcal{F}_i^{(1),\tau} = \{S_1^i, h_1^i, p_1\}$ after the stage-1A market clears.

FOC. Both D_1^i and D_{SD}^i are settled at the period 1 price. At period 1, traders maximize

$$\mathbb{E}\Big[-p_1\big(D_1^i + D_{SD}^i\big) + (1-\delta)\Big(v_1 z_1^i - \frac{\gamma}{2}\big(z_1^i\big)^2\Big) + \delta V_{i,2}(\hat{z}_2^i, D_2^i, p_2) \;\Big|\; \mathcal{F}_i^{(1),\tau}\Big],$$

hence their FOC is

$$\begin{aligned}\mathbb{E}\Big[(1-\delta)\big(v_1 - \gamma z_1^i\big)\Big(1 + \frac{\mathrm{d}D_{SD}^i}{\mathrm{d}D_1^i}\Big) + \delta\,\frac{\mathrm{d}V_{i,2}}{\mathrm{d}D_1^i} \;\Big|\; \mathcal{F}_i^{(1),\tau}\Big] \\ = p_1\Big(1 + \frac{\mathrm{d}D_{SD}^i}{\mathrm{d}D_1^i}\Big) + \lambda_1^\tau\,\mathbb{E}\big[D_1^i + D_{SD}^i \;\big|\; \mathcal{F}_i^{(1),\tau}\big].\end{aligned} \tag{59}$$

Going back to the continuation value,

$$\begin{aligned}\frac{\mathrm{d}V_{i,2}}{\mathrm{d}D_1^i} &= -\frac{\mathrm{d}p_2}{\mathrm{d}D_1^i}D_2^i - p_2\frac{\mathrm{d}D_2^i}{\mathrm{d}D_1^i} + \big(v_2 - \gamma(\hat{z}_2^i + D_2^i)\big)\Big(1 + \frac{\mathrm{d}D_{SD}^i}{\mathrm{d}D_1^i} + \frac{\mathrm{d}D_2^i}{\mathrm{d}D_1^i}\Big) \\ &= \big(v_2 - \gamma z_2^i\big)\Big(1 + \frac{\mathrm{d}D_{SD}^i}{\mathrm{d}D_1^i} + \frac{\mathrm{d}D_2^i}{\mathrm{d}D_1^i}\Big) - p_2\frac{\mathrm{d}D_2^i}{\mathrm{d}D_1^i} - \lambda_{1,2}^\tau D_2^i \\ &= \big(v_2 - p_2 - \gamma z_2^i\big)\left(1 + \frac{\mathrm{d}D_{SD}^i}{\mathrm{d}D_1^i} + \frac{\mathrm{d}D_2^i}{\mathrm{d}D_1^i} - \frac{\lambda_{1,2}^\tau}{\lambda_2^\tau}\right) + p_2\Big(1 + \frac{\mathrm{d}D_{SD}^i}{\mathrm{d}D_1^i}\Big) \\ &= \big(v_2 - p_2 - \gamma z_1^i\big)(1-\chi_2^\tau)\left(1 + \frac{\mathrm{d}D_{SD}^i}{\mathrm{d}D_1^i} + \frac{\mathrm{d}D_2^i}{\mathrm{d}D_1^i} - \frac{\lambda_{1,2}^\tau}{\lambda_2^\tau}\right) + p_2\Big(1 + \frac{\mathrm{d}D_{SD}^i}{\mathrm{d}D_1^i}\Big).\end{aligned}$$

As at stage 1B, by the law of iterated expectations we've implicitly eliminated mean 0 terms: the third equality uses the period 2 FOC, and the fourth uses the period 2 demand

(30) together with $\hat{z}_2^i = z_1^i + h_2^i$ and $\mathbb{E}[h_2^i \mid \mathcal{F}_i^{(1),\tau}] = 0$. Define

$$\kappa_v^\tau \equiv (1-\delta)\Big(1 + \frac{\mathrm{d}D_{SD}^i}{\mathrm{d}D_1^i}\Big) + \delta(1-\chi_2^\tau)\Big(1 + \frac{\mathrm{d}D_{SD}^i}{\mathrm{d}D_1^i} + \frac{\mathrm{d}D_2^i}{\mathrm{d}D_1^i} - \frac{\lambda_{1,2}^\tau}{\lambda_2^\tau}\Big). \tag{60}$$

Using the random walk identity $\mathbb{E}[v_2 \mid \mathcal{F}_i^{(1),\tau}] = \mathbb{E}[v_1 \mid \mathcal{F}_i^{(1),\tau}]$, $z_1^i = h_1^i + D_1^i + D_{SD}^i$, and collecting terms, the FOC at period 1 can be rewritten as

$$\mathbb{E}\big[v_1 - \gamma z_1^i \,\big|\, \mathcal{F}_i^{(1),\tau}\big] = \Omega_1^\tau\, p_1 + (1-\Omega_1^\tau)\,\mathbb{E}\big[p_2 \,\big|\, \mathcal{F}_i^{(1),\tau}\big] + \frac{\lambda_1^\tau}{\kappa_v^\tau}\,\mathbb{E}\big[D_1^i + D_{SD}^i \,\big|\, \mathcal{F}_i^{(1),\tau}\big], \tag{61}$$

where

$$\Omega_1^\tau := \frac{1}{\kappa_v^\tau}\Big(1 + \frac{\mathrm{d}D_{SD}^i}{\mathrm{d}D_1^i}\Big). \tag{62}$$

This has the same form as the period 2 FOC (29) and the stage 1B FOC (46): the expected marginal value of the period 1 position equals an effective price, a weighted combination of p_1 and the expected period 2 price with weight Ω_1^τ on p_1, plus the price impact on the total period 1 trade, since both D_1^i and D_{SD}^i settle at p_1. Ω_1^τ is the period 1 counterpart of $\Omega_{\mathrm{SD}}^\tau$ in (46).

No matching session. Without a matching session, $D_{SD}^i \equiv 0$, so $\frac{\mathrm{d}D_{SD}^i}{\mathrm{d}D_1^i} = 0$, $\Omega_1^\tau = 1/\kappa_v^\tau$, and $z_1^i = h_1^i + D_1^i$. Solving (61) for D_1^i:

$$\Longrightarrow\; D_1^i(p_1) = \chi_1^\tau\bigg(\frac{1}{\gamma}\Big(\mathbb{E}\big[v_1 \,\big|\, \mathcal{F}_i^{(1),\tau}\big] - \Omega_1^\tau\, p_1 - (1-\Omega_1^\tau)\,\mathbb{E}\big[p_2 \,\big|\, \mathcal{F}_i^{(1),\tau}\big]\Big) - h_1^i\bigg), \tag{63}$$

where $\chi_1^\tau := \gamma\kappa_v^\tau/(\gamma\kappa_v^\tau + \lambda_1^\tau)$, and κ_v^τ, the price impacts, and the projection of p_2 are those of the no-matching-session branch. As in (30) and (48), the demand is the competitive demand, a net benefit scaled by the marginal holding cost minus the current inventory, bid shaded by $\chi_1^\tau \in [0,1]$. Without price impact, $\kappa_v^\tau = 1-\delta$, $\chi_1^\tau = 1$, and the effective price is $\big(p_1 - \delta\mathbb{E}[p_2]\big)/(1-\delta)$, the cost of a unit bought at p_1 and resold at p_2 if the asset does not liquidate.

Matching session. With a matching session, the stage 1B FOC (46) gives the left-hand side of (61), and (48) gives $h_1^i + D_1^i + \mathbb{E}[D_{SD}^i \mid \mathcal{F}_i^{(1),\tau}] = \frac{1}{\gamma}\big(\mathbb{E}[v_1 \mid \mathcal{F}_i^{(1),\tau}] - \Omega_{\mathrm{SD}}^\tau p_1 - (1-\Omega_{\mathrm{SD}}^\tau)\mathbb{E}[p_2 \mid \mathcal{F}_i^{(1),\tau}]\big)$. Substituting both, D_1^i drops out and the FOC becomes

$$0 = \kappa_v^\tau\big(\Omega_{\mathrm{SD}}^\tau - \Omega_1^\tau\big)\Big(p_1 - \mathbb{E}\big[p_2 \,\big|\, \mathcal{F}_i^{(1),\tau}\big]\Big)$$

$$-\frac{\lambda_1^\tau}{\gamma}\Big(\mathbb{E}\big[v_1 \mid \mathcal{F}_i^{(1),\tau}\big] - \Omega_{\mathrm{SD}}^\tau p_1 - (1-\Omega_{\mathrm{SD}}^\tau)\,\mathbb{E}\big[p_2 \mid \mathcal{F}_i^{(1),\tau}\big] - \gamma h_1^i\Big), \quad (64)$$

where these quantities are in the matching-session branch. The first term is the difference between the two effective prices, the exchange's Ω_1^τ and the matching session's $\Omega_{\mathrm{SD}}^\tau$, applied to the expected price change. The period 1 exchange demand does not appear because, on this branch, a unit bought in the exchange and a unit bought in the matching session are perfect substitutes.Both affect z_1^i one for one and both settle at p_1, and the matching session demand re-optimizes z_1^i whatever D_1^i was. Only the total period 1 trade is pinned down, by (48). What remains in (64) is the expected marginal value of shifting one unit from the matching session to the exchange, which is nonzero because the two routes carry different price impacts (λ_1^τ and $\lambda_{1,2}^\tau$ for the exchange, $\lambda_{\mathrm{SD},2}^\tau$ for the matching session). With a matching session for certain, the period 1 FOC therefore does not determine D_1^i. This is the degeneracy when $\theta = 1$ discussed in the main text.

Probabilistic matching session. The period 1 demand is submitted before it is known whether the matching session occurs. Let $\zeta = 1$ if it occurs and $\zeta = 0$ otherwise. The period 1 FOC weights the marginal expected utilities on the two branches by $1-\theta$ and θ. $(1-\theta)$ times κ_v^τ times the difference between the left- and right-hand sides of (61) on the no-matching-session branch, plus θ times the right-hand side of (64), equals zero. Since (64) does not involve D_1^i, the combined FOC has the same form as (61) on the no-matching-session branch, with combined coefficients:

$$\bar\kappa_v^\tau\,\mathbb{E}\big[v_1 - \gamma(h_1^i + D_1^i) \mid \mathcal{F}_i^{(1),\tau}\big] - \lambda_1^\tau D_1^i = (1-\theta)\,\kappa_v^\tau\Big(\Omega_1^\tau\, p_1 + (1-\Omega_1^\tau)\,\mathbb{E}\big[p_2 \mid \mathcal{F}_i^{(1),\tau}, \zeta = 0\big]\Big)$$
$$-\frac{\theta\,\lambda_1^\tau}{\gamma}\left(\Big(\Omega_{\mathrm{SD}}^\tau + \frac{\gamma\kappa_v^\tau}{\lambda_1^\tau}\big(\Omega_{\mathrm{SD}}^\tau - \Omega_1^\tau\big)\Big)\,p_1 + \Big(1 - \Omega_{\mathrm{SD}}^\tau - \frac{\gamma\kappa_v^\tau}{\lambda_1^\tau}\big(\Omega_{\mathrm{SD}}^\tau - \Omega_1^\tau\big)\Big)\,\mathbb{E}\big[p_2 \mid \mathcal{F}_i^{(1),\tau}, \zeta = 1\big]\right), \quad (65)$$

where

$$\bar\kappa_v^\tau := (1-\theta)\,\kappa_v^\tau - \theta\,\frac{\lambda_1^\tau}{\gamma}, \quad (66)$$

and in each term κ_v^τ and Ω_1^τ are evaluated on the same branch as the conditional expectation in that term. The no-matching-session branch contributes κ_v^τ to the expected marginal value and its effective price from (61). The matching-session branch contributes $-\lambda_1^\tau/\gamma$, because on that branch an extra unit of exchange demand only increases price impact, and its effective

price is (64) rearranged as a weighted combination of p_1 and the expected period 2 price. In both parentheses the weights sum to one.

The right-hand side of (65) is $\bar{\kappa}_v^\tau$ times a weighted average of p_1 and the two branch forecasts of p_2. Collecting the weight on p_1, and using $\kappa_v^\tau \Omega_1^\tau = 1$ on the no-matching-session branch,

$$\bar{\Omega}_1^\tau := \frac{1}{\bar{\kappa}_v^\tau}\left[(1-\theta) - \theta\Big(\kappa_v^\tau\big(\Omega_{\mathrm{SD}}^\tau - \Omega_1^\tau\big) + \frac{\lambda_1^\tau}{\gamma}\,\Omega_{\mathrm{SD}}^\tau\Big)\right], \tag{67}$$

with the matching-session branch quantities inside the parentheses. The remaining weight $1-\bar{\Omega}_1^\tau$ multiplies a forecast of p_2 that averages cross the two branches,

$$\begin{aligned}\bar{\kappa}_v^\tau\left(1-\bar{\Omega}_1^\tau\right)\bar{\mathbb{E}}\big[p_2 \,\big|\, \mathcal{F}_i^{(1),\tau}\big] &:= (1-\theta)\,(\kappa_v^\tau - 1)\,\mathbb{E}\big[p_2 \,\big|\, \mathcal{F}_i^{(1),\tau}, \zeta = 0\big] \\ &\quad + \theta\Big(\kappa_v^\tau\big(\Omega_{\mathrm{SD}}^\tau - \Omega_1^\tau\big) - \frac{\lambda_1^\tau}{\gamma}\big(1-\Omega_{\mathrm{SD}}^\tau\big)\Big)\mathbb{E}\big[p_2 \,\big|\, \mathcal{F}_i^{(1),\tau}, \zeta = 1\big], \end{aligned} \tag{68}$$

where κ_v^τ in the first term is from the no-matching-session branch and the objects in the second term are from the matching-session branch. The two coefficients on the right sum to $\bar{\kappa}_v^\tau(1-\bar{\Omega}_1^\tau)$, so $\bar{\mathbb{E}}[p_2 \mid \mathcal{F}_i^{(1),\tau}]$ is a weighted average of the two branch's forecasts. Solving (65) for D_1^i:

$$\implies D_1^i(p_1) = \frac{\gamma\bar{\kappa}_v^\tau}{\gamma\bar{\kappa}_v^\tau + \lambda_1^\tau}\left(\frac{1}{\gamma}\Big(\mathbb{E}\big[v_1 \,\big|\, \mathcal{F}_i^{(1),\tau}\big] - \bar{\Omega}_1^\tau\, p_1 - \big(1-\bar{\Omega}_1^\tau\big)\,\bar{\mathbb{E}}\big[p_2 \,\big|\, \mathcal{F}_i^{(1),\tau}\big]\Big) - h_1^i\right). \tag{69}$$

This is the same form as (30), (48), and (63). There's a bid-shading weight, which is χ_1^τ with $\bar{\kappa}_v^\tau$ in place of κ_v^τ, times the competitive demand at an effective price that weights p_1 by $\bar{\Omega}_1^\tau$ and the forecast of p_2 by $1-\bar{\Omega}_1^\tau$. When $\theta = 0$, $\bar{\kappa}_v^\tau = \kappa_v^\tau$, $\bar{\Omega}_1^\tau = \Omega_1^\tau$, the forecast is the no-matching-session forecast, and (69) reduces to (63).

Matching coefficients. Conjecturing the linear period 1 demand schedule

$$D_1^i = a_{1,S}^\tau S_1^i + a_{1,h}^\tau h_1^i + a_{1,p}^\tau p_1, \tag{70}$$

substituting the projections (82)–(83) into (69), and matching coefficients on each direction $\xi \in \{S, h, p_1\}$,

$$a_{1,\xi}^\tau = \frac{\gamma\bar{\kappa}_v^\tau}{\gamma\bar{\kappa}_v^\tau + \lambda_1^\tau}\left(\frac{1}{\gamma}\Big(c_{\xi,1}^\tau - \bar{\Omega}_1^\tau\,\mathbf{1}\{\xi = p_1\} - \big(1-\bar{\Omega}_1^\tau\big)\,\bar{\rho}_{\xi,2}^\tau\Big) - \mathbf{1}\{\xi = h\}\right), \tag{71}$$

where $\bar{\rho}_{\xi,2}^\tau$ is the loading of $\bar{\mathbb{E}}[p_2 \mid \mathcal{F}_i^{(1),\tau}]$ on ξ, the average of the two branches' $\rho_{\xi,2}^\tau$ with the weights in (68). This has the same structure as (53) with $\bar{\Omega}_1^\tau$ in place of $\Omega_{\mathrm{SD}}^\tau$.

E. Cross-Period Demand Directions and Cross-Period Price Impacts

To evaluate the value of a deviation today, we need to account for the adjustment of future trades while remaining optimal. Thus, we differentiate the later FOCs with respect to the earlier decisions. Given present and future price impacts, the responsiveness of future demands to deviations of decisions today solves a system of equations.

E.1. Matching session → period 2 Impacts

The following cross-period impact is derived by differentiating the period 2 FOC (29) with respect to μ_i. Differentiating, we get

$$-\gamma\Big(\frac{N+M-1}{N+M}+\frac{\mathrm{d}D_2^i}{\mathrm{d}\mu_i}\Big)=\lambda^\tau_{\mathrm{SD},2}+\lambda^\tau_2\,\frac{\mathrm{d}D_2^i}{\mathrm{d}\mu_i}$$

$$\implies \frac{\mathrm{d}D_2^i}{\mathrm{d}\mu_i}=-\frac{\gamma\frac{N+M-1}{N+M}+\lambda^\tau_{\mathrm{SD},2}}{\gamma+\lambda^\tau_2}. \tag{72}$$

E.2. period 1 exchange to period 2 impacts

A period 1 deviation also affects allocation through rivals' matching session demands, which respond to the deviation's effect on p_1. Since $z_1^i=\hat{z}_1^i+\mu_i-Q/(N+M)$,

$$\frac{\mathrm{d}z_1^i}{\mathrm{d}D_1^i}=1+\frac{\mathrm{d}\mu_i}{\mathrm{d}D_1^i}-\frac{1}{N+M}\frac{\mathrm{d}Q}{\mathrm{d}D_1^i},\qquad \frac{\mathrm{d}D^i_{SD}}{\mathrm{d}D_1^i}=\frac{\mathrm{d}\mu_i}{\mathrm{d}D_1^i}-\frac{1}{N+M}\frac{\mathrm{d}Q}{\mathrm{d}D_1^i}, \tag{73}$$

where

$$\frac{\mathrm{d}Q}{\mathrm{d}D_1^i}=\frac{\mathrm{d}\mu_i}{\mathrm{d}D_1^i}+(n_\tau-1)\,b_z^\tau\,a_{1,p}^\tau\,\lambda_1^\tau+n_{-\tau}\,b_z^{-\tau}\,a_{1,p}^{-\tau}\,\lambda_1^\tau. \tag{74}$$

Differentiating the period 2 FOC with respect to D_1^i:

$$0=-\gamma\Big(1+\frac{\mathrm{d}\mu_i}{\mathrm{d}D_1^i}-\frac{1}{N+M}\frac{\mathrm{d}Q}{\mathrm{d}D_1^i}+\frac{\mathrm{d}D_2^i}{\mathrm{d}D_1^i}\Big)-\lambda^\tau_{1,2}-\lambda^\tau_2\,\frac{\mathrm{d}D_2^i}{\mathrm{d}D_1^i} \tag{75}$$

$$\implies \frac{\mathrm{d}D_2^i}{\mathrm{d}D_1^i}=-\frac{\gamma\left(1+\frac{\mathrm{d}\mu_i}{\mathrm{d}D_1^i}-\frac{1}{N+M}\frac{\mathrm{d}Q}{\mathrm{d}D_1^i}\right)+\lambda^\tau_{1,2}}{\gamma+\lambda^\tau_2}. \tag{76}$$

E.3. period 1 exchange → SD Impacts

The period 1 exchange → SD impacts are given by differentiating the matching session FOC (45) with respect to the exchange decision at period 1:

$$\begin{aligned} 0 = & -(1-\delta)\,\gamma\Big(1-\frac{1}{N+M}\Big)\Big(1+\frac{\mathrm{d}\mu_i}{\mathrm{d}D_1^i}-\frac{1}{N+M}\frac{\mathrm{d}Q}{\mathrm{d}D_1^i}\Big)-\lambda_1^\tau\Big(1-\frac{1}{N+M}\Big) \\ & -\delta\,\gamma\Big(1-\frac{1}{N+M}+\frac{\mathrm{d}D_2^i}{\mathrm{d}\mu_i}\Big)\Big(1+\frac{\mathrm{d}\mu_i}{\mathrm{d}D_1^i}-\frac{1}{N+M}\frac{\mathrm{d}Q}{\mathrm{d}D_1^i}+\frac{\mathrm{d}D_2^i}{\mathrm{d}D_1^i}\Big) \\ & -\delta\,\lambda_{\mathrm{SD},2}^\tau\frac{\mathrm{d}D_2^i}{\mathrm{d}D_1^i}-\delta\,\frac{\mathrm{d}D_2^i}{\mathrm{d}\mu_i}\,\lambda_{1,2}^\tau. \end{aligned} \tag{77}$$

E.4. Own- and Cross-Period Price Impacts

Own-period price impacts need to satisfy

$$\lambda_t^\tau \equiv \frac{\mathrm{d}p_t}{\mathrm{d}D_t^i} = -\Big(\sum_{j\neq i}\frac{\partial D_t^j}{\partial p_t}\Big)^{-1} = \frac{1}{(n_\tau-1)\,b_t^\tau + n_{-\tau}\,b_t^{-\tau}}, \tag{78}$$

where b_2^τ is the period 2 demand slope of (34) and $b_1^\tau := |a_{1,p}^\tau| = -a_{1,p}^\tau$ is the period 1 demand slope on p_1 (in equilibrium $a_{1,p}^\tau < 0$).

Both cross-period price impacts are obtained by implicit differentiation of period 2 market clearing. For a period 1 deviation by trader i,

$$\lambda_{1,2}^\tau \equiv \frac{\mathrm{d}p_2}{\mathrm{d}D_1^i} = -\Big(\sum_{j\neq i}\frac{\partial D_2^j}{\partial p_2}\Big)^{-1}\Big(\frac{\mathrm{d}D_2^i}{\mathrm{d}D_1^i}+\sum_{j\neq i}\frac{\partial D_2^j}{\partial \hat{z}_2^j}\frac{\mathrm{d}\hat{z}_2^j}{\mathrm{d}D_1^i}+\sum_{j\neq i}\frac{\partial D_2^j}{\partial \mathcal{F}_j^{(2)}}\frac{\mathrm{d}\mathcal{F}_j^{(2)}}{\mathrm{d}D_1^i}\Big),$$

and analogously for a stage 1B deviation in the matching session demand μ_i. From the period 2 demand (30), a rival j of type τ' has $\frac{\partial D_2^j}{\partial p_2} = \frac{c_{p_2,2}^{\tau'}-1}{\gamma+\lambda_2^{\tau'}}$ and $\frac{\partial D_2^j}{\partial \hat{z}_2^j} = \frac{-\gamma}{\gamma+\lambda_2^{\tau'}}$, and responds to the public information in $\mathcal{F}_j^{(2)}$ that the deviation affects, p_1 and Q, with slopes $\frac{c_{p_1,2}^{\tau'}}{\gamma+\lambda_2^{\tau'}}$ and $\frac{c_{Q,2}^{\tau'}}{\gamma+\lambda_2^{\tau'}}$. The following is a common denominator:

$$\mathcal{D}_2^\tau \;:=\; \sum_{j\neq i}\frac{\partial D_2^j}{\partial p_2} \;=\; (n_\tau-1)\,\frac{c_{p_2,2}^{\tau}-1}{\gamma+\lambda_2^\tau} \;+\; n_{-\tau}\,\frac{c_{p_2,2}^{-\tau}-1}{\gamma+\lambda_2^{-\tau}}. \tag{79}$$

period 1 exchange →2 impact ($t=1$).

$$\lambda_{1,2}^\tau \;=\; -\frac{1}{\mathcal{D}_2^\tau}\left(\frac{\mathrm{d}D_2^i}{\mathrm{d}D_1^i}+(n_\tau-1)\,\frac{-\gamma}{\gamma+\lambda_2^\tau}\left(\frac{\partial D_1^\tau}{\partial p_1}\,\lambda_1^\tau + b_z^\tau a_{1,p}^\tau\lambda_1^\tau - \frac{1}{N+M}\frac{\mathrm{d}Q}{\mathrm{d}D_1^i}\right)\right.$$

$$+n_{-\tau}\,\frac{-\gamma}{\gamma+\lambda_2^{-\tau}}\left(\frac{\partial D_1^{-\tau}}{\partial p_1}\,\lambda_1^{\tau}+b_z^{-\tau}a_{1,p}^{-\tau}\lambda_1^{\tau}-\frac{1}{N+M}\frac{\mathrm{d}Q}{\mathrm{d}D_1^i}\right)$$
$$+\frac{(n_\tau-1)\,c_{p_1,2}^{\tau}\lambda_1^{\tau}+(n_\tau-1)\,c_{Q,2}^{\tau}\,\mathrm{d}Q/\mathrm{d}D_1^i}{\gamma+\lambda_2^{\tau}}+\frac{n_{-\tau}\,c_{p_1,2}^{-\tau}\lambda_1^{\tau}+n_{-\tau}\,c_{Q,2}^{-\tau}\,\mathrm{d}Q/\mathrm{d}D_1^i}{\gamma+\lambda_2^{-\tau}}\Bigg), \quad (80)$$

where $\partial D_1^{\tau'}/\partial p_1$ denotes the period 1 demand slope of a type-τ' rival, which equals $a_{1,p}^{\tau'}$ under (70). At $t=1$, each rival j's period 2 projection onto Q uses (74).

Matching session → period 2 impact (stage 1B).

$$\lambda_{\mathrm{SD},2}^{\tau} \;=\; -\frac{1}{\mathcal{D}_2^{\tau}}\Bigg(\frac{\mathrm{d}D_2^i}{\mathrm{d}\mu_i}+(n_\tau-1)\,\frac{\gamma}{(N+M)(\gamma+\lambda_2^{\tau})}+n_{-\tau}\,\frac{\gamma}{(N+M)(\gamma+\lambda_2^{-\tau})}$$
$$+\frac{(n_\tau-1)\,c_{Q,2}^{\tau}}{\gamma+\lambda_2^{\tau}}+\frac{n_{-\tau}\,c_{Q,2}^{-\tau}}{\gamma+\lambda_2^{-\tau}}\Bigg). \quad (81)$$

For both cross-period impacts, the first term is ($\mathrm{d}D_2^i/\mathrm{d}D_1^i$ for $\lambda_{1,2}$, $\mathrm{d}D_2^i/\mathrm{d}\mu_i$ for $\lambda_{\mathrm{SD},2}$). The second and third terms capture rivals' exchange-response through (for $\lambda_{1,2}$) p_1 or (at stage 1B) the SD-inventory channel $\mathrm{d}\hat{z}_2^j/\mathrm{d}\mu_i=-1/(N+M)$. The final two fractions collect the $c_{p_1,2}$ (only for $\lambda_{1,2}$) and $c_{Q,2}$ terms summed over the $(n_\tau-1)$ same-type and $n_{-\tau}$ opposite-type counterparties.

F. Learning (Bayesian Projections)

All projections are linear-Gaussian and computed by the standard projection theorem $\mathbb{E}[X|Y]=\mathrm{Cov}(X,Y)\,\mathrm{Var}(Y)^{-1}\,Y$ (or its multivariate generalization), where the joint covariance matrices are derived from the equilibrium primitive loadings of all state variables. The projections are type-specific because traders of types I and U have different signal precisions $\sigma_{\varepsilon,I}^2$ vs. $\sigma_{\varepsilon,U}^2$.

F.1. period 1

The trader observes $\mathcal{F}_i^{(1),\tau}=\{S_1^i,h_1^i,p_1\}$. Conjecture that trader i of type τ has linear beliefs at period 1. Two projections are needed:

$$\mathbb{E}[v_1\mid\mathcal{F}_i^{(1),\tau}]=c_{S,1}^{\tau}\,S_1^i+c_{h,1}^{\tau}\,h_1^i+c_{p_1,1}^{\tau}\,p_1, \quad (82)$$
$$\mathbb{E}[p_2\mid\mathcal{F}_i^{(1),\tau}]=\rho_{S,2}^{\tau}\,S_1^i+\rho_{h,2}^{\tau}\,h_1^i+\rho_{p_1,2}^{\tau}\,p_1. \quad (83)$$

The v_1 projections are used in the period 1 FOC (61) and the matching FOC (71). The p_2 projections are used in (61).

F.2. Workup Stage 1B

At stage 1B, $\mathcal{F}_i^{(\mathrm{SD}),\tau} = \mathcal{F}_i^{(1),\tau}$, so projections coincide with the period 1 projections (82)–(83). The key new object is the rival-aggregate projection

$$\pi_\xi^\tau := \text{loading on } \xi \text{ of } \mathbb{E}\left[\frac{1}{N+M}\sum_{j\neq i}\mu_j^* \;\middle|\; \mathcal{F}_i^{(\mathrm{SD}),\tau}\right], \tag{84}$$

for $\xi \in \{S, h, p_1\}$. These rival projections appear in the matching equations (57)–(58) and represent the information that trader i extracts about average rival matching session demands from their information set.

F.3. period 2

At period 2 the trader's information set expands to

$$\mathcal{F}_i^{(2),\tau} = \mathcal{F}_i^{(\mathrm{SD}),\tau} \cup \{S_2^i, h_2^i, p_2, [Q]\}, \tag{85}$$

where Q is included in Design 2 only. The relevant projection is

$$\mathbb{E}[v_2 \mid \mathcal{F}_i^{(2),\tau}] = c_{S_1,2}^\tau S_1^i + c_{S_2,2}^\tau S_2^i + c_{h_1,2}^\tau h_1^i + c_{h_2,2}^\tau h_2^i + c_{p_1,2}^\tau p_1 + c_{p_2,2}^\tau p_2 + c_{Q,2}^\tau Q. \tag{86}$$

G. *Summary of demands and prices*

For trader i of type τ, the equilibrium strategies and inventories are

$$\begin{aligned}
D_1^i &= a_{1,S}^\tau S_1^i + a_{1,h}^\tau h_1^i + a_{1,p}^\tau p_1 \\
\mu_i &= b_z^\tau\, \hat{z}_1^i + b_S^\tau\, S_1^i + b_h^\tau\, h_1^i \\
D_{SD}^i &= \mu_i - Q/(N+M), \qquad Q = \sum_j \mu_j \\
D_2^i &= a_{2,S_1}^\tau S_1^i + a_{2,S_2}^\tau S_2^i + a_{2,h_1}^\tau h_1^i + a_{2,h_2}^\tau h_2^i + a_{2,p_1}^\tau p_1 + a_{2,p_2}^\tau p_2 + a_{2,Q}^\tau Q \\
\hat{z}_1^i &= h_1^i + D_1^i \\
z_1^i &= \hat{z}_1^i + D_{SD}^i \\
\hat{z}_2^i &= z_1^i + h_2^i \\
z_2^i &= \hat{z}_2^i + D_2^i
\end{aligned}$$

where the strategy coefficients satisfy the matching conditions (71), (56), and (36). Let τ_j denote the type of trader j. Market clearing, $\sum_j D_t^j = 0$, then gives the prices

$$p_1 = -\frac{1}{\sum_j a_{1,p}^{\tau_j}} \sum_j \left(a_{1,S}^{\tau_j} S_1^j + a_{1,h}^{\tau_j} h_1^j \right)$$

$$p_2 = -\frac{1}{\sum_j a_{2,p_2}^{\tau_j}} \sum_j \left(a_{2,S_1}^{\tau_j} S_1^j + a_{2,S_2}^{\tau_j} S_2^j + a_{2,h_1}^{\tau_j} h_1^j + a_{2,h_2}^{\tau_j} h_2^j + a_{2,p_1}^{\tau_j} p_1 + a_{2,Q}^{\tau_j} Q \right)$$

Substituting recursively, every equilibrium object (inventories, prices, matching session demands) is a linear function of the primitive shocks, which is what the projections of Appendix AF use.

H. Fixed-Point Problem

The equilibrium projections must be consistent with the joint Gaussian distribution induced by the equilibrium strategies. In our implementation, this consistency is achieved via the iterative fixed-point procedure.

(1) Start from an initial guess of the strategy coefficients $(a_1^\tau, a_2^\tau, b^\tau)$ and price impacts $(\lambda_1^\tau, \lambda_2^\tau, \lambda_{1,2}^\tau, \lambda_{\mathrm{SD},2}^\tau)$.
(2) Build the joint primitive distribution: each equilibrium object (inventories, prices, matching session demands) is a linear function of the primitive shocks; their loadings are computed from the current strategy coefficients.
(3) Compute the projections $c_{\xi,t}^\tau, \rho_{\xi,2}^\tau, \pi_\xi^\tau$ via the Gaussian projection theorem applied to the joint primitive distribution.
(4) update strategies via the matching conditions (36), (56), and (71); update λ's via the price-impact fixed points (32)–(33), the matching session cross-period equation (81), and the 3×3 joint solve combining (75), (80), and (77).
(5) Iterate until convergence.

At convergence, the strategy and price-impact updates are consistent with the projections they induce.

The Limits of Voluntary Transparency in Competitive Markets*

Terrence Hendershott
Haas School of Business
University of California, Berkeley

Dan Li
Board of Governors of the Federal Reserve System

Dmitry Livdan
Haas School of Business
University of California, Berkeley

Norman Schürhoff
Faculty of Business and Economics and Swiss Finance Institute
University of Lausanne

October 4, 2026

*Terrence Hendershott is with the Walter A. Haas School of Business, University of California, Berkeley, Berkeley CA 94720. Tel.: (510) 643-0619; fax: (510) 643-1412; e-mail: hender@haas.berkeley.edu. He provides expert witness services to a variety of clients. He has taught a course for a financial institution that engages in liquidity provision and high-frequency trading activity. He gratefully acknowledges support from the Norwegian Finance Initiative. Dan Li is with the Board of Governors of the Federal Reserve System, Washington, DC. E-mail: dan.li@frb.gov. She has no conflict of interest to disclose. The views of this paper do not represent the views of the Federal Reserve Board. Dmitry Livdan is with the Walter A. Haas School of Business, University of California, Berkeley, Berkeley CA 94720. fax: (510) 207-3759; e-mail: livdan@haas.berkeley.edu. Livdan is also a Research Fellow of the CEPR and has no conflict of interest to disclose. Norman Schürhoff is with the Faculty of Business and Economics and Swiss Finance Institute, University of Lausanne, CH-1015 Lausanne, Switzerland. Tel.: +41 (21) 692-3447; fax: +41 (21) 692-3435; e-mail: norman.schuerhoff@unil.ch. Schürhoff is also a Research Fellow of the CEPR. He gratefully acknowledges research support from the Swiss Finance Institute and has no conflict of interest to disclose.

The Limits of Voluntary Transparency in Competitive Markets

Abstract

Many financial assets trade infrequently and are difficult to value. Can price transparency arise voluntarily among competitors, or must it be mandated? In collateralized loan obligation markets, dealers publicly share attributed, nonbinding price estimates, known as price talks, before competing in sealed-bid auctions. We model price talk as costly revelation of a common-value signal before idiosyncratic values are realized. We show that dealers with unfavorable information talk because their chance of winning is low and revealing private information raises it. Dealers with favorable information stay silent because revelation would dissipate their bidding advantage. In 53,000 auctions with 430,000 price talks, the talk consensus predicts transaction prices, yet silent dealers win the majority of auctions, and prices are further above the consensus when nontalkers win, a "silence premium." Voluntary transparency produces informative public valuations while maintaining competition among silent dealers. In counterfactuals, mandating disclosure would mainly redistribute surplus from dealers to sellers.

JEL Classification. G12, G14, G23, D44, D82

Key words. Over-the-counter markets, dealer intermediation, pre-trade disclosure, transparency, auctions, CLOs

1 Introduction

Many financial assets, including structured credit securities, are complex and difficult to value. Unlike in equity or Treasury markets, where frequent trading aggregates dispersed information into observable prices, transactions in these securities occur infrequently and provide little price guidance for other securities. Valuing a collateralized loan obligation (CLO) tranche is costly because it requires assessing the underlying loans and the contractual rules that allocate their cash flows. CLO debt tranches are also less liquid and more costly to hold than corporate bonds (Cordell, Roberts, and Schwert, 2023). In the secondary market for these securities, price information is nevertheless produced and shared before trade by the prospective buyers themselves. Investors sell through "bids-wanted-in-competition" (BWIC) auctions, and before submitting sealed, binding bids, dealers can publish attributed, nonbinding price estimates, known as "price talks," observed by all market participants, including the rival dealers who bid in the subsequent auction. Why do dealers help competitors value an asset they are trying to buy themselves?

This paper shows how competition determines why dealers talk, what information becomes public, and who bids most aggressively for the asset. The answer starts from the observation that, when dealers publish price talks, the audience for that disclosure includes rival bidders. Revealing an unfavorable signal raises the dealer's chance of winning, because it removes informational asymmetry between dealer and rivals. While private, the signal enters only the dealer's own bid. Once public, it enters the posterior reflected in all bids. Revealing a favorable signal dissipates the informational advantage the signal provides in bidding. A dealer thus trades off a higher chance of winning against the value of retaining private information. The net gain from disclosure declines with the favorability of the signal. An idiosyncratic value, arising from inventory capacity or client orders, leaves the dealer a reason to buy even after common-value information becomes public.

In equilibrium, dealers with sufficiently unfavorable signals and low talking costs disclose their information ("talkers"), give up their informational advantage, and compete on idiosyncratic values. Hence, a dealer who talks does not automatically lose the auction. Dealers with favorable signals or high talking costs remain silent ("nontalkers"), and both their common-value information and their idiosyncratic value continue to affect their bidding. Public disclosures are informative but selected toward less favorable price estimates. The market appears to talk down the asset even though information revelation is truthful and the equilibrium is noncollusive. Market participants account for this selection when interpreting both talk and silence. Partial information revelation is a noncooperative equilibrium and arises endogenously in this market. The partial talk we observe persists not because it serves dealers collectively but because dealers cannot commit to collective silence. Breaking such an agreement is profitable for precisely the dealers who talk in equilibrium.

This sorting is the opposite of the canonical unraveling prediction in disclosure games. In Grossman (1981) and Milgrom (1981), the discloser gains when the audience believes their news is favorable, so high types disclose to separate from low types, silence signals bad news, and skeptical

inference leads to full revelation. Disclosure costs (Verrecchia, 1983) and uncertainty about whether the sender is informed (Dye, 1985) explain why unraveling stops, but preserve its direction, with types above a threshold disclosing. In our setting, the audience consists of rival bidders, whose more favorable beliefs can make the asset more expensive to acquire. Tan (2016) establishes the related incentive to fully reveal unfavorable common-value information while retaining rents from private values. Our focus is the competition that results when disclosure remains partial, which is what we observe in the data. Talkers share their valuation information, silent dealers retain additional private signals, and the two groups compete in the same auction.

Our data make the disclosure decision itself observable. We study more than 53,000 BWIC auctions and 430,000 dealer price talks from 2012 through 2022, sourced from Empirasign and matched to TRACE transaction prices. To our knowledge, these are the first data in which individual dealers' voluntary disclosure decisions are observed before trade and linked to winner identities and transaction prices. Existing evidence on transparency in fixed-income markets largely concerns mandated publication after trade. Four facts stand out. Price talk is common, with 93.6% of auctions receiving at least one talk and eight dealers talking on average. Price talk is informative, with consensus surprises predicting price surprises with a coefficient of 0.95. Transaction prices lie inside the range of talks, above the consensus and below the most optimistic talk, and bids are far less dispersed than talks. And nontalkers win 55% of auctions. An informative public consensus produced by the bidders themselves coexists with silent dealers who win the majority of auctions.

The winner's type is itself news about value. Silent dealers pay prices further above the consensus than talkers do, a differential we call the "silence premium."[1] In regressions with day fixed effects, the price premium, defined as the transaction price minus the consensus, is 8 basis points (bp) larger when a nontalker wins, and the gap between the transaction price and the highest individual talk narrows by 26 bp. Both effects are larger in riskier tranches, where withheld information is most valuable. The silence premium also survives within dealer. The same dealer pays a silence premium of 10 bp when winning silently, rising from under 3 bp for AAA/AA tranches to 19 bp for equity tranches, so persistent dealer heterogeneity does not explain the result. The pooled estimates are significant at the 1% level.

The question the silence premium raises is who keeps the value of the withheld information. If silent winners captured it as an information rent, the premium would appear in the winning margin, defined as the winning bid minus the cover bid. It does not. The winning margin is no wider when a nontalker wins, because the cover bid adjusts upward alongside the winning bid. Retained information is competed into the price rather than captured by the dealers, so the silent side of the market remains competitive. The silence premium hence measures the market's shared retained information, not the private knowledge of a single dealer.

[1]The silence premium is consistent with favorable retained common-value information but does not by itself identify the direction of disclosure sorting, because silent winners can also be selected for high idiosyncratic values. Section 4 separates these channels.

The frequent allocation to silent dealers suggests a prima facie regulatory role for mandating transparency. However, in counterfactual analysis mandated full disclosure is mostly a redistribution from dealers to sellers and not an efficiency enhancement. The information a mandate would publish is already largely reflected in the price. The information rent that mandated disclosure competes away in CLOs is the first-price bid shading of silent dealers. It is largest for the highest-signal silent winners, so the burden of a mandate falls on the best-informed dealers. This rent is also the return to being informed. Endogenous information acquisition (Grossman and Stiglitz, 1980), participation, and forced-talk credibility therefore all cut against mandated disclosure.

On allocations, silence does create inefficiency. Because the most optimistic dealers stay silent and reduce their bids against the winner's curse, a talker with a lower posterior valuation can win against a silent dealer with a higher one. The asset's value is predominantly common, so reallocating it across dealers changes realized surplus only through the small idiosyncratic components, and the numerical counterfactual bounds these misallocation costs as small. This is in contrast to settings in which transparency mainly improves efficiency. Seller-provided information improves the allocation (Ganuza, 2006; Ganuza and Penalva, 2010), benchmark transparency reduces search frictions (Duffie, Dworczak, and Zhu, 2017), and mandated post-trade transparency lowered investors' trading costs in corporate bond markets (Bessembinder, Maxwell, and Venkataraman, 2006; Edwards, Harris, and Piwowar, 2007; Goldstein, Hotchkiss, and Sirri, 2007).

The model formalizes this mechanism and disciplines the counterfactual. Each dealer observes a noisy private signal about the common value of the security and draws a private cost of providing price talk. The idiosyncratic value is realized only after the disclosure decision, so a dealer cannot condition disclosure on knowing that their idiosyncratic value is high. A dealer talks when their cost is below the net gain from talking. Under conditions maintained throughout, talking occurs only below an upper signal boundary and, even there, only for sufficiently low cost draws. Every signal type stays silent with positive probability. Silence therefore does not reveal the signal, while a talk reveals it exactly, and talkers have less favorable signals than silent dealers on average.

The disclosure stage then determines the information structure of the auction. In the baseline model, talkers' signals are public, so their bids differ only through idiosyncratic values, and outbidding a talker carries no additional news about the common value. Nontalkers retain private signals, so outbidding a nontalker is bad news about the common value. The winner's-curse adjustment therefore depends on the information retained by silent rivals and on what the disclosure stage has made public. Talkers bid on the public information and their own idiosyncratic value, and the highest talker bid lies below a finite ceiling. A silent dealer whose signal and idiosyncratic value place their bid below that ceiling can be outbid by a talker. A silent dealer with a sufficiently favorable signal bids above the entire talker range and can lose only to other silent dealers.

The resulting three-region structure is testable. BWIC lists identify dealers who received the list but remained silent on a given security. Talkers win against such present-and-silent dealers in 24% of dealer-security pairs, and in 35% of the auctions in which such dealers are present, which

rejects a two-region model in which every silent dealer outbids the talker range.[2] Price talk is also associated with fewer failed auctions and higher execution probability, margins that lie outside the baseline model and are examined with a reserve extension in Section 6.

The results identify dealers as a private source of public valuations and show the limits of this arrangement. Relative to established results, the paper makes three departures. The sorting of voluntary disclosure reverses the unraveling benchmark, so silence is good news rather than bad news. The winner's curse is asymmetric and determined by the disclosure stage rather than by an exogenous information structure. And the disclosure decision is observed directly in the data rather than inferred from prices. The dealers whose signals would move the consensus most, those with favorable information, have the strongest incentive to stay silent. Voluntary disclosure therefore improves the information available to the market without approaching full revelation, and evaluating mandatory disclosure requires accounting for how it would change competition between talkers and nontalkers. This consideration is relevant to current regulatory efforts to improve disclosure in leveraged-loan and CLO markets (IOSCO, 2024).

Related literature. Our paper contributes to four literatures. First, we contribute to work on OTC market structure (Duffie, Gârleanu, and Pedersen, 2005; Lagos and Rocheteau, 2007; Hendershott and Madhavan, 2015; Hugonnier, Lester, and Weill, 2020; Di Maggio, Kermani, and Song, 2017; Li and Schürhoff, 2019; Hollifield, Neklyudov, and Spatt, 2017). This literature emphasizes dealers' roles in search, inventory absorption, matching, and balance sheet intermediation. We show that dealers also provide an information service. Before competing for the asset, they selectively reveal private information that rivals incorporate into their bids. This disclosure aggregates dispersed signals and improves price discovery, and it is distinct from the liquidity provision and matching functions emphasized in existing models of OTC intermediation.

Second, we contribute to the voluntary-disclosure literature (Grossman, 1981; Milgrom, 1981; Verrecchia, 1983; Dye, 1985; Matthews and Postlewaite, 1985). This literature explains incomplete disclosure through frictions that stop unraveling while preserving its direction. Our setting reverses the direction because disclosure is addressed to rival bidders, as discussed above. Heterogeneous talking costs generate coexistence of talkers and silent dealers at the same signal and play a role analogous to uncertainty about informedness in Dye (1985). Silence therefore does not identify a high signal. A dealer may stay silent because their information is favorable or because talking is costly.

The closest theoretical parallels are models in which disclosure changes the behavior of a strategic counterparty. Glode, Opp, and Zhang (2018) study coarse disclosure when revelation affects screening and information rents. Li, Liu, and Zhao (2022) show that an informed bidder may withhold favorable information because disclosure changes a rival bidder's participation through

[2]The rejection maintains that present-and-silent dealers submit bids. Because bid submission is not observed, the evidence is also consistent with nonparticipation by some present-and-silent dealers.

the winner's curse. Goldstein, Xiong, and Yang (2025) show that a coarsely informed investor may share information with a better-informed investor because the recipient trades against the error in the sender's signal, reducing the sender's price impact. Our mechanism is different. Each dealer chooses disclosure separately and immediately before a competitive auction, and the disclosure stage determines the information structure of bidding, which side faces the winner's curse, and how strongly. In the extension with seller reserves, disclosure also changes the seller's posterior and required price, which affects execution.

Third, we contribute to auction theory (Milgrom and Weber, 1982; Persico, 2000; Ganuza and Penalva, 2010; Campbell and Levin, 2000; Kagel and Levin, 2016) by studying bidder-initiated information sharing before a common-value auction. Existing work analyzes information supplied by the seller or auctioneer (Ganuza, 2006; Bergemann, Heumann, Morris, Sorokin, and Winter, 2022), experimental disclosure by bidders (Dufwenberg and Gneezy, 2002), and incentives to reveal common-value information when bidder values also contain private components (Tan, 2016). Seller-supplied information changes not only expected revenue but also which bidder wins. Board (2009) terms this the allocation effect. In our model the release decision belongs to the bidders and is selected on the content of the signal, so the allocation consequences of transparency are equilibrium outcomes rather than design choices. Tan (2016) is particularly close in showing that a bidder may reveal a low common-value signal because doing so lowers rivals' bids and increases rents earned from the private-value component. Our mechanism differs in two respects. Idiosyncratic values are realized only after the disclosure decision, so dealers cannot condition disclosure on favorable private-value realizations. In addition, heterogeneous disclosure costs generate partial disclosure and a coexistence region in which dealers with the same signal may either talk or remain silent. This coexistence prevents silence from perfectly partitioning the signal space and creates overlap between the subsequent bid distributions of talkers and silent dealers. Boyarchenko, Lucca, and Veldkamp (2021) study information sharing among dealers in Treasury auctions. We instead observe the disclosure decision dealer by dealer and link it directly to subsequent bidding, winner identity, and transaction prices.

Finally, we contribute to the market-transparency literature (Madhavan, 1995; Pagano and Röell, 1996; Boehmer, Saar, and Yu, 2005; Bessembinder, Maxwell, and Venkataraman, 2006; Edwards, Harris, and Piwowar, 2007; Goldstein, Hotchkiss, and Sirri, 2007; Asquith, Covert, and Pathak, 2013). Much of this literature studies mandated post-trade transparency and asks how publication of completed trades affects liquidity and price discovery. Our setting concerns voluntary information release before dealers compete for the asset. We show that such pre-trade transparency can arise endogenously while remaining selectively incomplete. Disclosure is informative and is associated with stronger price discovery, but competitive incentives limit how much favorable information dealers are willing to reveal. The model therefore explains why voluntary transparency can improve the information available to the market without producing complete revelation. Mandating additional pre-trade disclosure would improve allocations, but the numerical comparison in

Section 7 suggests the welfare gains would be small relative to the redistribution involved.

2 Institutional Setting and Data

2.1 BWIC auctions

Collateralized loan obligations (CLOs) are structured credit vehicles that pool leveraged loans and issue tranched securities. The U.S. CLO market exceeded $1 trillion in outstanding securities by 2022. Valuing a CLO tranche requires assessing the underlying loans and the contractual rules that allocate their cash flows. CLO debt tranches are more costly to hold than corporate bonds (Cordell, Roberts, and Schwert, 2023), and the secondary market is characterized by opacity and infrequent trading. Dealers play a central role in intermediating trades.

The primary trading mechanism in the secondary market is the bids-wanted-in-competition (BWIC) auction. Sellers circulate a BWIC as a list of securities identified by CUSIP, and each security on the list is auctioned separately. A BWIC i is an email-based first-price sealed-bid auction with a distinctive two-stage structure. There are dealers $j = 1, \ldots, N$, and the total number N of dealers bidding is not observed. In the first stage, before the bidding deadline, a subset $\mathcal{T}_i$ of the dealers voluntarily provide price talks S_{ij} to the public, which are non-binding estimates of fair value expressed as a percentage of par. The broker disseminates all talks to market participants. Talks are attributed in that recipients see which dealer provided each estimate, and dealer-level records allow a dealer's talk history to be scrutinized ex post against realized transaction prices. A dealer whose talks prove systematically slanted loses credibility with the very counterparties they trade with every day, and market participants indicate that brokers can drop dealers from future BWIC distribution lists for misbehavior. Providing a talk is not free and incurs cost κ_{ij}. Because the estimate is attributed and archived, it must be defensible ex post, which absorbs analyst time and desk attention that vary with staffing, security coverage, and the other demands on the desk that day, and it reveals information to competitors. The remaining dealers $\mathcal{S}_i$ remain silent but can still bid. The number of talks $N_i = |\mathcal{T}_i|$ is observed, while the number of silent bidders is not. Let the average (consensus) talk be $\bar{S}_i = N_i^{-1} \sum_{j \in \mathcal{T}_i} S_{ij}$, and let $S_i^{1:N_i} = \max_{j \in \mathcal{T}_i} S_{ij}$ denote the most optimistic talk. In the second stage, dealers submit binding bids by the deadline (typically the same day). The highest bid wins and determines the transaction price P_i. The second-highest bid, known as the cover C_i, is typically disclosed after the auction. The winning margin is the price minus the cover, $P_i - C_i$. The seller is not obligated to sell and may reject all bids, in which case the BWIC fails and the security does not trade.

2.2 Data

We combine data from multiple sources. Empirasign is a commercial data provider that collects data on the universe of BWIC auctions, including dealer-level price talks, auction outcomes, and

Table 1: Summary Statistics

This table reports summary statistics for the sample of BWIC auctions from 2012–2022. Dealers active per list $\hat{N}_i$ counts the distinct identified dealers that talk or win on any CUSIP of the list containing auction i.

	Mean	Median	SD	P25	P75	Obs
Panel A. Auction-Level						
Has price talk (%)	93.6	100	24.5	100	100	53,052
Number of talks (N_i)	8.01	8	4.50	5	11	53,052
Dealers active per list ($\hat{N}_i$)	10.55	11	4.44	7	14	52,485
Transaction price P_i (% par)	94.17	98.57	11.44	94.00	99.96	39,284
Winning margin (pp of par)	0.24	0.10	0.42	0.02	0.31	30,451
Cover disclosed (%, cond. on trade)	77.9	100	41.5	100	100	42,611
BWIC failure rate (%)	19.7	0	39.8	0	0	53,052
List size (no. of CUSIPs)	8.37	7	6.98	4	10	53,052
Panel B. By Tranche Seniority						
	AAA	AA	A	BBB	BB	B-Eq
Has price talk (%)	88.1	93.7	95.0	96.4	97.2	97.0
Failure rate (%)	8.2	12.6	15.3	20.2	25.9	34.2
BWICs	12,039	4,745	4,902	8,170	11,354	5,869

cover prices when disclosed. The regulatory version of TRACE provides transaction prices for executed trades. Bloomberg provides CLO characteristics such as ratings and collateral features. Market conditions data, including VIX and credit spreads, come from FRED.

Our sample covers 2012 through 2022 and includes 53,052 BWIC auctions and over 430,000 dealer-level price talks. Table 1 reports summary statistics. Talk is pervasive. A price talk occurs in 93.6% of auctions, and the average auction draws eight talks. The dealers active per list bound participation beyond the talkers, counting distinct dealers observed talking or winning anywhere on the same list, while the total number of bidders in an auction remains unobserved. Transaction prices average 94.2% of par with sizable dispersion reflecting heterogeneity in tranche quality and market conditions. The average winning margin is 24 basis points (bp) against a median of 10 bp, so the winning bid typically exceeds the cover by little. The seller discloses the cover after 78% of executed trades. One in five BWICs fails, and the median list contains seven CUSIPs.

Panel B of Table 1 shows substantial heterogeneity across tranche seniority. The share of auctions with talk rises from 88.1% for AAA to 97% for the junior tranches, so talk is most prevalent where valuation uncertainty is greatest. Auction failure rates increase monotonically with risk, from 8.2% for AAA to over 34% for B and equity tranches, the combined category labeled B-Eq in the table.

2.3 Four facts about voluntary transparency

Four facts characterize voluntary transparency in this market.

Fact 1. Price talk is widespread but not universal. In our sample, 93.6% of auctions receive

at least one price talk and the average auction features eight talking dealers (Table 1), out of 79 identified dealers that talk or win in the sample, 67 of whom talk at least once. Participation in talk rises with valuation difficulty, from 88% of auctions for AAA tranches to 97% for B and equity tranches. Talk is common but not universal. The average list has 10.6 active dealers against 8.0 talks per auction, and the total number of bidders is unobserved.

Fact 2. The transaction price lies inside the range of talks, above their mean and below their maximum. Panel A of Table 2 shows that the transaction price exceeds the consensus talk by 19 bp on average and exceeds it in 62% of auctions. The same price lies about 1.0 percentage point below the most optimistic talk $S_i^{1:N_i}$ and below it in roughly 82% of auctions. Bids are also far less dispersed than talks, with a bid dispersion of 0.226% of par against a talk dispersion of 0.482% (Table B.3). Figure 1 shows this anatomy for eight representative auctions. In each auction shown, nine to eleven dealers publicly post attributed estimates for one security, and their estimates disagree by tens of bp for senior tranches and by a few hundred bp for junior ones. This dispersion gradient previews the seniority patterns that Section 4 exploits. In seven of the eight auctions the price lies inside the talk range. In the bottom row a nontalker wins at a price above every disclosed talk. This crossing event occurs in 22.3% of nontalker-won auctions (Table 2), and Fact 4 returns to it. The averages themselves are location facts rather than evidence of selection. The winning bid is the maximum of the submitted bids while the consensus is an average of talks, so a positive average gap arises mechanically even if talks were unbiased forecasts of value. The inferential weight in Section 4 therefore rests on how these gaps differ by winner type, a difference the order-statistic effect alone cannot produce.

Fact 3. Price talk is informative. Measuring informativeness requires a pre-auction benchmark that excludes any information from dealer disclosures. We construct a reference price R_i with machine-learning algorithms that predict the transaction price using only pre-auction observable information. In our base specification, we train a neural network on features including tranche seniority, vintage, age, yield-to-maturity, deal size, and market conditions (VIX, credit spreads), estimated with five-fold cross-validation to generate out-of-sample predictions. The benchmark achieves an out-of-sample R^2 of 0.695.[3] The benchmark defines price and consensus surprises,

$$\tilde{P}_i \equiv P_i - R_i, \qquad \tilde{\bar{S}}_i \equiv \bar{S}_i - R_i. \tag{1}$$

Panel B of Table 2 shows that consensus surprises predict price surprises. The regression

$$\tilde{P}_i = \alpha_t + \beta \cdot \tilde{\bar{S}}_i + \epsilon_i \tag{2}$$

[3]Because yield-to-maturity is computed from pre-auction evaluated prices, the benchmark summarizes observable market information rather than fundamentals alone, which is the relevant conditioning set for measuring the incremental information in talks. Cross-validation folds drawn at the observation level can place the same security in training and validation samples, so grouping folds by CUSIP and rolling or time-block validation are natural robustness designs.

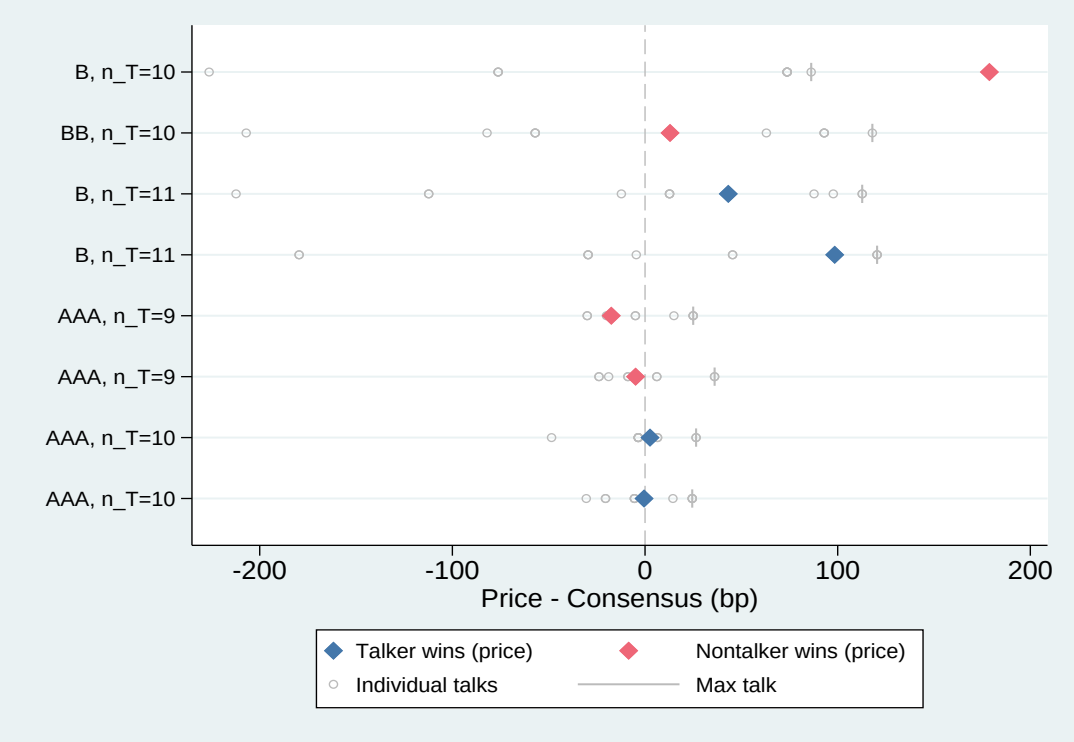

Figure 1: Anatomy of BWIC auctions. Each row is one auction, centered on its consensus $\bar{S}_i$ (dashed line at zero). Open circles are individual talks, and diamonds are transaction prices for talker-won (blue) and nontalker-won (red) auctions. Rows are labeled by seniority and number of talks.

yields $\hat{\beta} = 0.947$ with $R^2 = 0.958$, indicating near-unity pass-through from consensus to prices.[4] Figure 2 shows the relation as a binscatter. The coefficient is lower when nontalkers win ($\beta = 0.935$ vs. 0.962), and R^2 drops from 0.971 to 0.952, consistent with private information entering prices through bidding rather than through talk.

Fact 4. Nontalkers win the majority of auctions and pay more. Panel A of Table 2 shows that nontalkers win 55.0% of auctions, and transaction prices are 0.40 percentage points of par higher when they win. The price premium, the gap between the transaction price and the consensus, is 9.3 bp larger in nontalker-won auctions ($p < 0.01$). This difference is the silence premium. Under the canonical unraveling logic, silence is bad news and silent winners should pay less relative to public benchmarks, a silence discount. The data instead show a premium.

These facts pose a puzzle. If disclosure is informative and benefits price discovery, why do informed dealers stay silent? And if silence is so profitable, why does anyone voluntarily disclose

[4] A caveat applies to surprise regressions (Tables 6 and 7). Because the same estimated benchmark R_i is subtracted from both the dependent variable and the regressor, benchmark measurement error is common to both sides, which can bias the coefficient toward one and inflate R^2. The headline coefficient and fit should therefore be read as describing variation in benchmark-adjusted prices, not purely in transaction prices. A level specification avoids this common component. Regressing P_i on $\bar{S}_i$ and R_i as separate regressors, with day fixed effects, gives a consensus coefficient of 0.977 (s.e. 0.003) and a benchmark coefficient of 0.013 (0.004), with $R^2 = 0.986$. The consensus coefficient is 0.990 in talker-won and 0.967 in nontalker-won auctions, and the interaction with the nontalker indicator is -0.013 (0.003).

Table 2: Four Facts about Voluntary Transparency

Panel A reports where the transaction price falls relative to the talk distribution and compares auction outcomes by winner type. The price premium is the difference between the transaction price and the average talk, and the max-talk gap is the difference between the transaction price and the highest talk $S_i^{1:N_i}$ (both in percentage points of par). Panel B reports regressions of price surprises on consensus surprises, both measured relative to the ML benchmark price. Standard errors clustered by security in parentheses. ***, **, * denote significance at 1%, 5%, 10%.

Panel A. Price Location and Auction Outcomes by Winner Type

	All	Talker wins	Nontalker wins	Difference
Share of auctions (%)		45.0	55.0	
Transaction price (% of par)	94.17	93.94	94.34	0.40***
				(0.08)
Price premium (Price – Consensus)	0.19	0.14	0.23	0.09***
				(0.01)
Max-talk gap (Price – Highest talk)	−1.00	−1.19	−0.84	0.34***
				(0.02)
Share with Price > Consensus (%)	62.2	60.9	63.3	
Share with Price > Highest talk (%)	18.2	13.2	22.3	

Panel B. Consensus Predicts Transaction Prices, Less So When Nontalker Wins

	Dep. var. Price Surprise $\tilde{P}_i$			
	All	Talker wins	Nontalker wins	Difference
Consensus surprise $\tilde{\tilde{S}}_i$	0.947***	0.962***	0.935***	-0.027***
	(0.003)	(0.004)	(0.004)	(0.006)
R^2	0.958	0.971	0.952	
Obs	37,026			

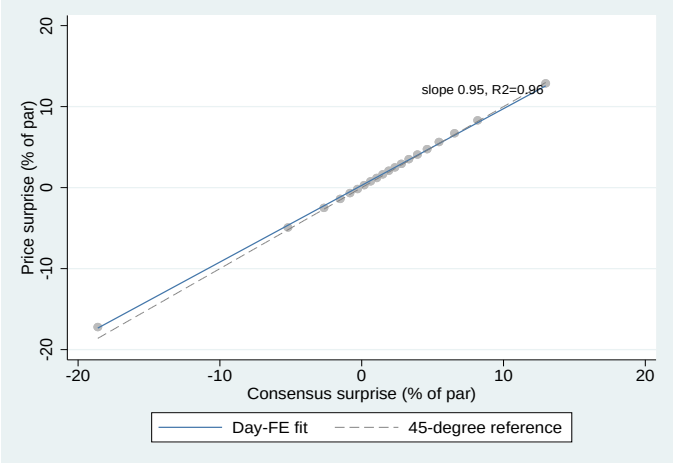

Figure 2: Consensus surprise and price surprise. Binscatter of the price surprise $\tilde{P}_i$ against the consensus surprise $\tilde{\tilde{S}}_i$, in percent of par, in twenty equal-count bins. The solid line is the day-fixed-effects fit from Table 2, Panel B, and the dashed line is the 45-degree line. Both surprises are measured against the same benchmark R_i (Fact 3).

their valuation to competitors? The following section develops a model that resolves this puzzle.

3 Model

Price talk changes both the information available before bidding and the information revealed by bids. A silent dealer retains a private signal about the security and an idiosyncratic value from acquiring it. Different combinations of these components can produce the same bid. A talker makes their signal public, so a strictly increasing bid reveals their remaining private valuation component. The model therefore allows pooling of valuation types among silent dealers and separation among talkers through disclosed signals and bids.

We study how this difference affects disclosure, bidding, and the information conveyed by the winner's identity. The disclosure informs future competitors. Under the maintained ranking of continuation payoffs, dealers with favorable signals prefer to retain that information, while those with less favorable signals are more willing to talk. Given a signal, heterogeneous costs of talking make disclosure partial. Idiosyncratic values realized after disclosure give talkers a reason to bid aggressively even after their signals become public. The analysis combines the winner's curse in common-value auctions (Milgrom and Weber, 1982) with information disclosure prior to bidding (Dufwenberg and Gneezy, 2002; Tan, 2016; Boyarchenko, Lucca, and Veldkamp, 2021). The equilibrium results state the continuation selection and disclosure conditions on which they depend.

3.1 Setup

Consider a BWIC auction with $N \geq 2$ risk-neutral dealers indexed by $j \in \{1, \ldots, N\}$. The security has a common value v. A dealer who wins at price P receives $v + u_j - P$, where u_j is their idiosyncratic value from acquiring the security. A dealer who does not win receives zero from the auction. Talking costs are incurred separately, regardless of the auction outcome.

The common value has prior $v \sim \mathcal{N}(\mu_0, \sigma_v^2)$, with $\sigma_v^2 > 0$. Dealer j observes privately

$$s_j = v + \varepsilon_j, \qquad \varepsilon_j \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma_\varepsilon^2), \tag{3}$$

where $\sigma_\varepsilon^2 > 0$ and the errors are independent of v. Write $\tau_v = 1/\sigma_v^2$ and $\tau_\varepsilon = 1/\sigma_\varepsilon^2$ for the prior and signal precisions. Gaussian signals make the posterior calculations explicit. The common component makes dealers' signals affiliated. A favorable signal raises a dealer's estimate of the asset, but winning also reveals that their bid exceeded other privately informed dealers' bids. This inference creates the winner's curse.

Before disclosure, the dealer j also observes their private talking cost

$$\kappa_j \overset{\text{iid}}{\sim} U[0, K], \qquad \kappa_j \perp (v, \varepsilon_1, \ldots, \varepsilon_N, u_1, \ldots, u_N), \tag{4}$$

where $K > 0$. Write $G_\kappa(x) = \min\{1, \max\{0, x/K\}\}$ for the cost cdf on the whole real line. The cost represents the attention required by analysts and desks to produce an attributed price estimate,

together with the reputational exposure associated with it. These demands can vary across dealers with the same information because their coverage, staffing, and other obligations differ. Whenever the gain from talking lies between zero and K, some dealers with a given signal disclose, and others remain silent. Cost heterogeneity therefore limits what observers can infer from silence, much as uncertainty about informedness does in Dye (1985).

Dealers simultaneously choose $T_j \in \{\text{Talk}, \text{Silent}\}$ without observing one another's choices. A talk mechanically reveals the signal, so the message m_j is equal to s_j whenever $T_j = \text{Talk}$. Exact disclosure is a maintained communication technology. Attributed talks and repeated interaction motivate this restriction, but truthful reporting is not derived from an unrestricted message game. The choice studied here is whether to release information that rivals will then use in the auction (Tan, 2016; Dufwenberg and Gneezy, 2002).

After disclosure choices and talks become public, each dealer privately observes $u_j = \bar{u} z_j$, where

$$z_j \overset{\text{iid}}{\sim} F_0 \text{ on } [0,1], \qquad (z_1, \ldots, z_N) \perp (v, \varepsilon_1, \ldots, \varepsilon_N, \kappa_1, \ldots, \kappa_N). \tag{5}$$

The cdf F_0 has a continuous density bounded above and away from zero on $[0,1]$. Thus $F_u(t) = F_0(t/\bar{u})$ on $[0,\bar{u}]$, with the usual extensions outside that interval. Let $\mu_u = \mathbb{E}[u_j] = \bar{u}\mathbb{E}[z_j]$ denote the mean idiosyncratic value. We take $\bar{u} > 0$ to be small relative to common-value uncertainty and impose $\bar{u} < K$. Under the continuation conditions below, this inequality ensures that every signal type retains a positive probability of silence.

The timing represents a client order, inventory requirement, or other benefit from acquiring the security that arrives between price talk and the bidding deadline. A dealer chooses disclosure using (s_j, κ_j) and later bids using the information they retain and the realized u_j. If u_j were known before disclosure, the decision to talk would generally depend on both valuation components. The present timing isolates disclosure of common-value information from subsequent dealer-specific demand.

These two sources of heterogeneity have distinct effects. Talking costs determine which dealers disclose at a given signal. Idiosyncratic values determine how aggressively dealers subsequently bid. Once a dealer talks, their signal is public and u_j is their remaining private valuation component. At $\bar{u} = 0$, Lemma A5 rules out an equilibrium in which talkers use a common deterministic bid, win with positive probability, and face a continuous strictly increasing silent bidding strategy. Positive $\bar{u}$ gives talkers private variation in their willingness to pay and can support overlap with silent bids. The lemma concerns a deterministic common-value benchmark.

The active set of N dealers is publicly known, and every dealer submits a bid. Let $\mathcal{T}$ and $\mathcal{S}$ denote the sets of talkers and silent dealers, with $n_T = |\mathcal{T}|$ and $n_S = |\mathcal{S}| = N - n_T$. After disclosure, dealers therefore know both sets and their sizes.[5] Dealers submit sealed bids $b_j \in \mathbb{R}$. The highest bidder acquires the security and pays their bid, so $P = \max_j b_j$. Independent random priorities resolve ties. The baseline has no reserve. Section 6 adds a seller acceptance decision.

[5]The active set is not observed in the data. Empirical talk counts consequently do not identify the number of silent competitors.

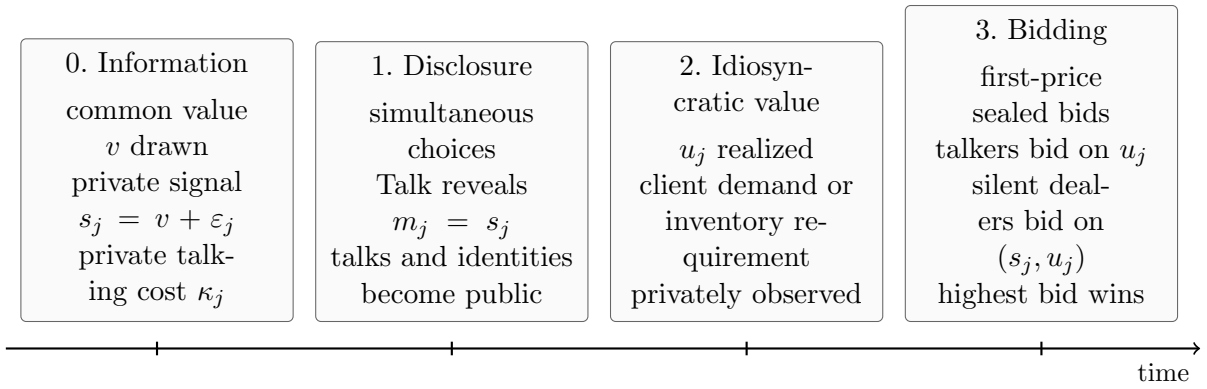

Figure 3: Model timing. Dealers observe their signals s_j and talking costs κ_j and decide whether to talk. Idiosyncratic values u_j are then realized, and dealers submit sealed bids. The common value v is never observed.

Let $\mathcal{M} = \{m_j : j \in \mathcal{T}\}$ denote the disclosed signals. When $n_T \geq 1$, their consensus is

$$\bar{S} = \frac{1}{n_T} \sum_{j \in \mathcal{T}} s_j. \tag{6}$$

The full public history is $\mathcal{I} = (\mathcal{M}, \mathcal{T}, \mathcal{S})$. Dealer identities support attribution of the talks. Under symmetry, the particular names carry no further information about valuations once the partition and disclosed values are given.

We study Perfect Bayesian Equilibria in symmetric cost-threshold disclosure strategies and selected continuation bidding equilibria. A disclosure threshold $\bar{\kappa} : \mathbb{R} \to \mathbb{R}$ prescribes talking when $\kappa_j \leq \bar{\kappa}(s_j)$, up to probability-zero indifference events. Write $p_T(s) = G_\kappa(\bar{\kappa}(s))$ for the propensity to talk and $w(s) = 1 - p_T(s)$ for the silence weight. We also use p for a conjectured disclosure policy. Beliefs follow Bayes' rule after disclosure, and continuation bids must be global best responses. After a forced verifiable announcement, including a report assigned zero probability by the conjectured policy, beliefs condition on the announced signal and apply the conjectured silence weights to the remaining silent dealers. A disclosure equilibrium requires the policy induced by these continuation values to equal the conjectured policy. Existence and regularity of the selected continuations and the monotonicity conditions on disclosure are maintained as stated in Assumptions A1 and A7 in Appendix A.

The values of the talks and the decision to remain silent convey different information. Using the disclosed signal values alone gives the Gaussian reference posterior

$$\mu_T = \frac{\tau_v \mu_0 + n_T \tau_\varepsilon \bar{S}}{\tau_v + n_T \tau_\varepsilon}, \qquad \sigma^2_{v|\mathcal{M}} = \frac{1}{\tau_v + n_T \tau_\varepsilon}. \tag{7}$$

If $n_T = 0$, these quantities equal μ_0 and $1/\tau_v$, with the talk contribution set to zero. Adding a silent

dealer's own signal gives the corresponding reference mean

$$\mu_S(s, \mathcal{M}) = \frac{\tau_v \mu_0 + n_T \tau_\varepsilon \bar{S} + \tau_\varepsilon s}{\tau_v + (n_T + 1)\tau_\varepsilon} = \mu_S(0, \mathcal{M}) + \gamma' s, \tag{8}$$

where $\gamma' = \tau_\varepsilon / [\tau_v + (n_T + 1)\tau_\varepsilon]$. These expressions describe the contribution of the signal values. The full posteriors also incorporate the disclosure decisions.

Conditional on v, a dealer remains silent with probability $q(v) = \int_{\mathbb{R}} w(t) f_\varepsilon(t - v)\, dt$, where f_ε is the Gaussian error density. Each silent dealer contributes this likelihood to the public posterior. The talk probabilities evaluated at already disclosed signals are constants with respect to v and cancel from Bayes' rule. Lemma A2 therefore gives

$$f(v \mid \mathcal{I}) \propto f(v \mid \mathcal{M}) q(v)^{n_S}. \tag{9}$$

For a silent dealer who also observes s, the full private posterior is proportional to $f(v \mid \mathcal{M}) f_\varepsilon(s - v) q(v)^{n_S - 1}$. Their own silence adds no information once their signal is known.

In the selective-disclosure equilibria characterized below, p_T is nonincreasing. The silence weight is then nondecreasing, and higher common values make silence more likely. Consequently, silence raises the public assessment of the asset. Write $\widetilde{\mu}_T = \mathbb{E}[v \mid \mathcal{I}]$ for the full public posterior mean. When $n_S > 0$ and the monotone silence weight is nonconstant on a set of positive measure, $\widetilde{\mu}_T > \mu_T$.

Under the maintained payoff bounds, Proposition 2 gives $p_T(s) \leq \mu_u / K \leq \bar{u} / K < 1$. Thus $w(s)$ and $q(v)$ are at least $1 - \mu_u / K > 0$. Silence remains possible at every signal and preserves full support. It raises the relative likelihood of favorable information while allowing a dealer with an unfavorable signal to remain silent because talking is costly. With a common disclosure cost and single crossing, silence would instead identify an upper signal region. All bidding payoffs below use the full information $\mathcal{I}$. The Gaussian reference quantities remain useful for connecting consensus to the empirical specifications. Appendix A.1 develops these posterior comparisons.

3.2 Equilibrium

We begin with bidding after a given disclosure history and then return to the decision to talk. The maintained continuation selection specifies globally optimal bids at every history needed to evaluate that decision. Conditional on this selection, disclosure determines the information available in the auction, and auction payoffs determine the incentive to disclose. Our results describe equilibria in which, after any pattern of talk and silence, dealers bid optimally given what they have learned. We assume such bidding equilibria exist rather than prove it from the model's primitives, as discussed after Proposition 1.

Fix a history $\mathcal{I}$ with $n_T \geq 1$ and $n_S \geq 2$. A talker bids $\beta^T(u; \mathcal{I})$ which is continuous and strictly increasing in $u \in (0, \bar{u}]$. Let $B_T = \max_{j \in \mathcal{T}} \beta^T(u_j; \mathcal{I})$ be the highest talker bid, let $H(\cdot \mid \mathcal{I})$ be its cdf, and let $\bar{b}(\mathcal{I}) = \beta^T(\bar{u}; \mathcal{I})$ be its finite upper endpoint. Conditional on $\mathcal{I}$, talker bids depend only

on independent idiosyncratic values. Thus B_T is independent of the common value and of silent dealers' private information.[6]

A silent dealer bids $b^S(s, u; \mathcal{I})$, which is continuous and strictly increasing in both arguments on the maintained interior domain. Define its zero-value slice by $b^S(s; \mathcal{I}) = b^S(s, 0; \mathcal{I})$. For $u > 0$, the pass-through of the idiosyncratic value is

$$\lambda^S(s, u; \mathcal{I}) = \frac{b^S(s, u; \mathcal{I}) - b^S(s, 0; \mathcal{I})}{u}. \tag{10}$$

The exact decomposition is therefore

$$b^S(s, u; \mathcal{I}) = b^S(s, 0; \mathcal{I}) + \lambda^S(s, u; \mathcal{I})u. \tag{11}$$

The coefficient can depend on both private components and on public information. Evaluating the zero-value slice sets only the focal dealer's idiosyncratic value to zero. Rivals retain their randomly drawn idiosyncratic values.

Pooling and separation. The silent bid combines two sources of willingness to pay. Fix an interior type (s_0, u_0) and its bid $b_0 = b^S(s_0, u_0; \mathcal{I})$. Continuity and strict increase in u imply that, for every s sufficiently close to s_0, there is a unique nearby $u(s)$ with $b^S(s, u(s); \mathcal{I}) = b_0$. The function $u(s)$ is continuous and decreases as the signal rises. A stronger assessment of the common value can therefore offset a smaller idiosyncratic value, leaving the bid unchanged.

This informational pooling does not require a mass point in the bids. Conditional on any signal, a strict increase in u means that at most one idiosyncratic value submits a specified bid. The idiosyncratic-value distribution is atomless, so the unconditional bid distribution is also atomless wherever this strict monotonicity holds on the full support. An econometrician who sees the bid and the relevant public history identifies $s = m_j$ and $u = (\beta^T(\cdot; \mathcal{I}))^{-1}(b_j)$ from the talker's observed signal and bid, separating their valuation components. This characterization is conditional on the stated continuation regularity. For both groups, talkers and silent dealers, to be present on the equilibrium path, talking must occur with positive probability. The disclosure result below states the additional conditions.

Competition between talkers and silent dealers. A talker must account for the information conveyed by defeating silent dealers. Let $B_S = \max_{k \in \mathcal{S}} b^S(s_k, u_k; \mathcal{I})$ be the highest silent bid. At a bid with positive clearing probability, the talker's win-conditional common-value mean is

$$\Psi^{\text{ex}}(b; \mathcal{I}) = \mathbb{E}[v \mid \mathcal{I}, B_S < b]. \tag{12}$$

[6] Statements about the zero-probability talker type $u = 0$ are interpreted through the limit as $u \downarrow 0$.

Defeating other talkers adds no information about v. The talker's expected margin conditional on winning is therefore $\Psi^{\text{ex}}(b;\mathcal{I}) + u - b$. For each silent rival, the signal required to beat that rival's bid depends on their idiosyncratic value. The exact winning event integrates over those different signal cutoffs, as shown in Appendix A.2.4.

To isolate the informational advantage of retaining a signal, consider the comparison event in which all silent signals lie below a common cutoff x. If $S_{\mathcal{S}}^{\max}$ is their maximum, define

$$\psi(x;\mathcal{I}) = \mathbb{E}[v \mid \mathcal{I}, S_{\mathcal{S}}^{\max} < x]. \tag{13}$$

The event is unfavorable news about the common value. Its conditional mean increases with x and lies below $\widetilde{\mu}_T$. Learning that the highest silent signal equals x instead gives $\varphi(x;\mathcal{I}) = \mathbb{E}[v \mid \mathcal{I}, S_{\mathcal{S}}^{\max} = x]$, defined through the conditional density. By exchangeability, this is also the assessment of a silent dealer with signal x who learns that all other silent signals are lower. The Gaussian likelihood ordering implies $\varphi(x;\mathcal{I}) > \psi(x;\mathcal{I})$.

The difference $\varphi - \psi$ measures the value of knowing the boundary signal rather than knowing only that all signals fall below it. This comparison helps explain why a talker needs idiosyncratic value to compete. The exact auction still uses Ψ^{ex}, because winning depends on rivals' bids and their idiosyncratic values, rather than on a common signal cutoff.

When signals are very unfavorable, a silent dealer's informational advantage disappears. Knowing that every silent signal lies below a very low threshold tells a bidder almost as much as knowing the highest silent signal exactly. The difference between the two estimates, $\varphi - \psi$, shrinks in proportion to how far the signal lies below the talkers' average, at rate $\sigma^2_{v|\mathcal{M}}/(\mu_T - s)$. The estimate ψ still moves with the threshold, so silence remains informative even in this range. Appendix A.2.7 proves these results for any fixed disclosure policy under which dealers stay silent with positive probability.

The link to the bids uses the maintained corridor $\psi(s;\mathcal{I}) \le b^S(s;\mathcal{I}) \le \varphi(s;\mathcal{I})$ on the overlap region in Assumption A1. Combined with the support result below, the shrinking wedge gives $b^S(s;\mathcal{I}) - \psi(s;\mathcal{I}) \to 0$. The corridor is a restriction on the exact continuation, rather than an identity between a bid-conditioned and a signal-conditioned posterior.

At higher signals, a silent bid reaches the top of the talker support. Define the zero-value entry margin $s^\times(\mathcal{I})$ by

$$b^S\big(s^\times(\mathcal{I});\mathcal{I}\big) = \bar{b}(\mathcal{I}). \tag{14}$$

For general u, define $s^\times(u;\mathcal{I})$ through $b^S(s^\times(u;\mathcal{I}), u;\mathcal{I}) = \bar{b}(\mathcal{I})$. A higher idiosyncratic value weakly lowers the signal needed to clear the entire talker support.

Proposition 1 (Characterization of monotone continuation equilibria)**.** *Fix a continuous disclosure policy satisfying $0 \le p(s) \le \bar{u}/K < 1$ and a public history $\mathcal{I}$ with $n_T \ge 1$ and $n_S \ge 2$. In any continuation equilibrium satisfying Assumption A1, the following statements hold.*

1. *The highest talker bid has an atomless distribution* $H(\cdot \mid \mathcal{I})$ *with support* $(-\infty, \bar{b}(\mathcal{I})]$*, and* $\beta^T(u;\mathcal{I}) \to -\infty$ *as* $u \downarrow 0$*.*

2. *For every* $u \in [0, \bar{u}]$*,* $b^S(s,u;\mathcal{I}) \to -\infty$ *as* $s \to -\infty$*. The zero-value slice satisfies* $b^S(s;\mathcal{I}) - \psi(s;\mathcal{I}) \to 0$*.*

3. *For every* $u \in [0, \bar{u}]$*, the entry margin* $s^\times(u;\mathcal{I})$ *is finite and unique. If* $s < s^\times(u;\mathcal{I})$*, the probability of being outbid by the talker side is* $1 - H(b^S(s,u;\mathcal{I}) \mid \mathcal{I}) > 0$*. If* $s \geq s^\times(u;\mathcal{I})$*, that probability is zero.*

Appendix A.2.8 establishes the lower-support result using deviations against the exact rival bid distribution. A proposed finite lowest talker bid leaves a profitable deviation for silent types near the corresponding boundary. The argument retains all rivals' random idiosyncratic values. The compatibility conclusion additionally uses the maintained corridor.

Continuation selection. Cases outside $n_T \geq 1$ and $n_S \geq 2$ require separate treatment. With no silent dealers, the continuation is a symmetric private-value auction with the public common-value mean added to valuations. With one silent dealer, their bid depends on their posterior expected value plus u, and it may reach a plateau at $\bar{b}$ under the endpoint condition in Remark A2. With no talkers, dealers retain both private components. The support and strict monotonicity statements above are not imposed on these cases by analogy.

The disclosure analysis needs a bidding equilibrium after every possible pattern of talk and silence, including patterns that arise only if a dealer deviates from its disclosure choice. Holding the other side's bidding fixed, Appendix A.2.13 shows that best responses can be chosen consistently. It does not show that an equilibrium exists when both sides adjust at once, so we assume one. Appendices A.2.14 and A.2.17 support this assumption with three special cases: an exact solution with two silent dealers and no idiosyncratic values, an approximate solution when idiosyncratic values are small, and an exact solution when bids are restricted to a finite grid. None of these alone proves that the assumed equilibrium exists in the full model.

Disclosure incentives. Before idiosyncratic values are realized, a dealer with signal s compares the expected gross continuation values $\bar{\Pi}^T(s;p)$ and $\bar{\Pi}^S(s;p)$ from talking and remaining silent when rivals use p. Each value integrates over rivals' signals and disclosure decisions and all idiosyncratic values, using the selected continuation after each resulting history. The gross gain from talking is

$$\Delta(s;p) = \bar{\Pi}^T(s;p) - \bar{\Pi}^S(s;p). \tag{15}$$

A dealer talks when $\kappa_j \leq \Delta(s_j;p)$. Their probability of talking is therefore $\mathcal{B}(p)(s) = G_\kappa(\Delta(s;p)) = \min\{1, \max\{0, \Delta(s;p)/K\}\}$. A symmetric disclosure equilibrium is a fixed point $p^* = \mathcal{B}(p^*)$.

Retaining favorable information can be valuable because the recipients of disclosure are future competitors. A favorable talk raises their common-value assessments. Silence lets the dealer use their signal privately in the auction. Under the continuation-dominance condition in Assumption A7, the gain from silence rises sufficiently strongly that $\Delta(s;p)$ decreases on its positive region. Appendix A.3.3 develops the signal and disclosure-switch comparisons behind this condition. The required ranking concerns exact continuation payoffs, including changes in rival bids and disclosure configurations.

The maintained talker-margin conditions give $0 \leq \Pi^T(u) \leq u$ at every history. Averaging over the focal idiosyncratic value and disclosure histories therefore gives $0 \leq \bar{\Pi}^T(s;p) \leq \mu_u$. Optimized silent payoffs are nonnegative, so $\Delta(s;p) \leq \mu_u$ and $\mathcal{B}(p)(s) \leq \mu_u/K < 1$. This sharper bound lies within the policy band $[0, \bar{u}/K]$ used in the appendix. Dealers with the same signal can take different actions when the gain from talking is positive, because sufficiently high costs still induce silence. A favorable signal raises the attraction of retaining information, while a high talking cost can induce silence even at an unfavorable signal.

Let $\mathcal{P}_{\bar{u}}$ be the set of all nonincreasing policies $p : \mathbb{R} \to [0, \bar{u}/K]$, ordered pointwise. It includes discontinuous policies. Assumption A7 requires the positive part of $\Delta(\cdot;p)$ to be nonincreasing and the best response to preserve the order of policies. Its stricter comparison on the positive region makes the canonical cutoff $\bar{\kappa}(s) = \Delta(s;p^*)$ strictly decreasing wherever positive. Let S denote a generic dealer signal. A nontrivial equilibrium is one in which talking occurs with positive probability.

Proposition 2 (Selective disclosure under continuation and order conditions)**.** *Let $\kappa \sim U[0,K]$, $\bar{u} < K$, and maintain Assumptions A1 and A7. The map $\mathcal{B}(p)(s) = G_\kappa(\Delta(s;p))$ is an order-preserving self-map of $\mathcal{P}_{\bar{u}}$ with least and greatest fixed points. At every fixed point p^*, a dealer talks exactly when $\kappa \leq \Delta(s;p^*)$, up to indifference, and $0 \leq p^*(s) \leq \mu_u/K < 1$. If the positive-talk region is nonempty, it equals $(-\infty, \bar{s}^*)$ for a finite $\bar{s}^*$. If both disclosure events have positive probability and p^* is nonconstant on a set of positive signal probability, then*

$$\mathbb{E}[S \mid \text{Talk}] < \mathbb{E}[S] < \mathbb{E}[S \mid \text{Silent}]. \tag{16}$$

Appendix A.3.4 proves Proposition 2. The domain $\mathcal{P}_{\bar{u}}$ is a complete lattice, and the maintained order condition permits Tarski's theorem to be applied. We use the greatest fixed point for the comparative statics. Its nontriviality is an additional requirement. A sufficient condition is $\Delta(s;0) > 0$ at some signal, which makes the greatest fixed point at least $\mathcal{B}(0)$ and yields positive talk probability on an interval. Whenever the positive-talk region is nonempty, its finite upper boundary follows for every fixed point, not only the greatest one.[7]

[7]Continuity of the gain and its strict decline on the positive region give $\{s : \Delta(s;p^*) > 0\} = (-\infty, \bar{s}^*)$. This statement does not require that Δ have no further zeros outside that region. Continuous monotone policies need not form a complete lattice. The contraction argument instead uses a closed set under the sup norm.

Under these conditions, the equilibrium combines selection across signals with partial disclosure within a signal. Talkers have less favorable signals on average, but silence does not reveal a dealer's signal. In the continuation auction, the silent bid can pool different valuation combinations, while the talk and bid jointly separate a talker's type. The disclosure fixed point and the continuation regularity are both needed for this equilibrium interpretation.

An equilibrium is locally unique if a small change in how often dealers expect others to talk changes the gain from talking by less than the spread of talking costs, K. Under this condition, no other equilibrium lies nearby, and repeated best responses converge back to the equilibrium. We do not derive this condition from the model's primitives; we impose it as an assumption. Appendix A.3.4 breaks the condition into its components, and Corollary A6 gives a sufficient condition for existence and uniqueness.

Bidding. Conditional on a selected disclosure equilibrium, we next characterize competition at a realized public history. Remaining silent does not by itself ensure that a dealer can outbid every talker. The disclosure boundary $\bar{s}^*$ determines where talking can occur, while the entry margin $s^\times(u;\mathcal{I})$ determines whether a silent dealer's bid clears the talker support. The following proposition distinguishes talker bids from silent bids that remain exposed to talker competition and silent bids that defeat the entire talker side.

Proposition 3 (Three bid distributions)**.** *Maintain Assumptions A1 and A7, the stated continuation selection, and $\bar{u} < K$. Fix an equilibrium history $\mathcal{I}$ with $n_T \geq 1$, $n_S \geq 2$, and a continuous equilibrium policy. Then the following statements hold.*

1. *The highest talker bid has an atomless distribution with support $(-\infty, \bar{b}(\mathcal{I})]$.*
2. *Silent types with $s < s^\times(u;\mathcal{I})$ bid below $\bar{b}(\mathcal{I})$ and can be outbid by the talker side with positive probability.*
3. *Silent types with $s \geq s^\times(u;\mathcal{I})$ bid at or above $\bar{b}(\mathcal{I})$ and cannot be outbid by a talker.*

The zero-value boundary is $s^\times(\mathcal{I}) = s^\times(0;\mathcal{I})$. A talker wins with positive probability conditional on $\mathcal{I}$ if and only if $\Pr(B_S < \bar{b}(\mathcal{I}) \mid \mathcal{I}) > 0$.

The support result and independence of B_T from silent private information yield this classification. It does not impose an ordering between the disclosure boundary and the entry margin. Proposition A19 establishes an ordering in the separate hard-truncation benchmark. In the weighted cost-threshold model, Figure 4 illustrates $\bar{s}^* < s^\times(\mathcal{I})$ without imposing that ordering generally.

Winner identity provides information even when the transaction price combines private information, idiosyncratic values, and bid shading. Let $\mathcal{W}_T$ denote a talker win and $\mathcal{W}_S$ a silent-dealer win. Write $\succeq_{\text{MLR}}$ for monotone likelihood-ratio dominance.

Proposition 4 (Winner identity and the common value)**.** *Fix a public history containing both groups. Suppose talker bids depend only on their independent idiosyncratic values, silent bids are nondecreasing in their signals at every idiosyncratic value, and cross-group ties have probability zero. If both winner events have positive probability, then*

$$v \mid (\mathcal{W}_S, \mathcal{I}) \succeq_{\text{MLR}} v \mid \mathcal{I} \succeq_{\text{MLR}} v \mid (\mathcal{W}_T, \mathcal{I}). \tag{17}$$

The corresponding mean inequalities are strict whenever $\Pr(\mathcal{W}_T \mid v, \mathcal{I})$ *is nonconstant in* v *on a set of positive posterior probability.*

Under the strictness condition, the mean comparison is

$$\mathbb{E}[v \mid \mathcal{W}_S, \mathcal{I}] > \mathbb{E}[v \mid \mathcal{I}] > \mathbb{E}[v \mid \mathcal{W}_T, \mathcal{I}]. \tag{18}$$

One sufficient condition is that talker bids place positive probability on interior silent signal cutoffs for a positive-measure set of silent idiosyncratic values, with silent bids strictly increasing at those cutoffs. Appendix A.3.6 gives the proof.

Higher common values shift the silence-conditioned signal distribution upward. Silent bids then become harder to beat, whereas talker bids vary only with their idiosyncratic values. A silent-dealer win is therefore favorable news about v. This comparison uses the exact bid distributions and the silence-conditioned likelihood. It does not require a common signal cutoff or imply an ordering of transaction prices relative to consensus. That price ordering additionally depends on idiosyncratic values and equilibrium bid shading.

Economic intuition. Figure 4 combines the disclosure decision with the subsequent bidding decision. Consider its illustrated ordering $\bar{s}^* < s^\times$ at a history with $n_T \geq 1$ and $n_S \geq 2$, and suppress $\mathcal{I}$ in the notation. The plotted silent bid sets the focal dealer's idiosyncratic value to zero. Rival idiosyncratic values remain random throughout.

Below $\bar{s}^*$, dealers with the same signal can enter bidding on either side. Those with costs below $\Delta(s; p^*)$ talk, while those with higher costs retain their signals. A silent dealer in this region can therefore have unfavorable information, despite silence being favorable news on average. Their zero-value bid lies below $\bar{b}$, so they can lose to a talker with a sufficiently high idiosyncratic value. They can also win against talkers with lower idiosyncratic values if their bid defeats the other silent dealers. Once disclosure is complete, the talking cost is sunk and does not enter either side's bidding decision.

For $\bar{s}^* \leq s < s^\times$, the dealer remains silent, up to indifference at zero cost, but their zero-value bid still lies within the talker support. Retaining information is at least as profitable as disclosing it even when disclosure is free. This preference does not imply that they find it profitable to outbid every talker. The value of keeping a signal private and the payment required to defeat talkers

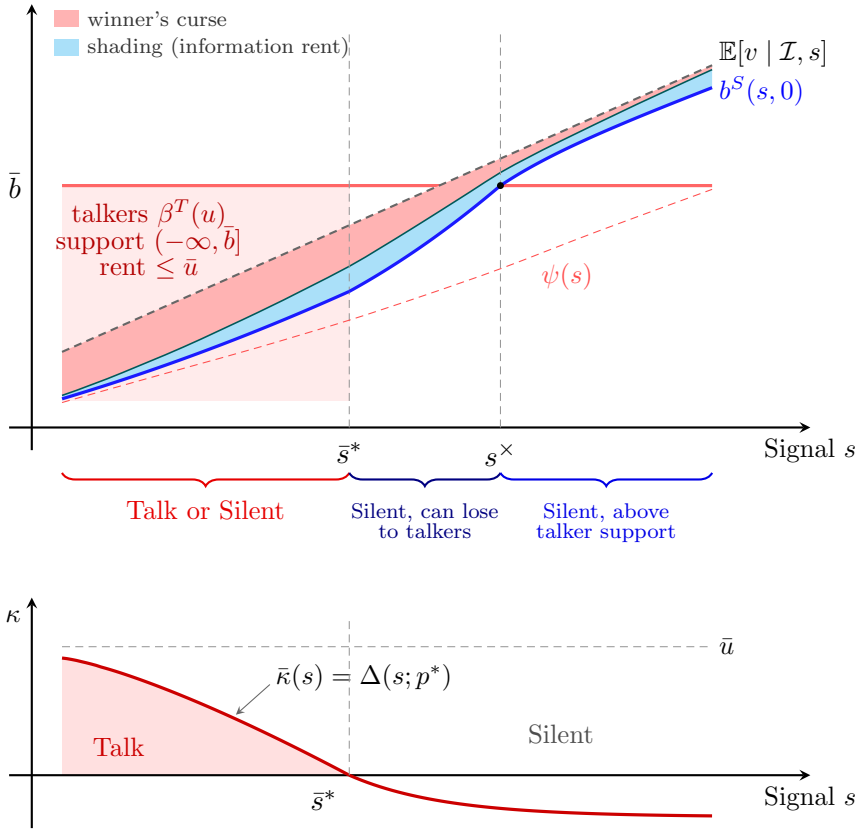

Figure 4: Disclosure and bidding in equilibrium. Top panel: the zero-value silent bid $b^S(s,0)$ (blue), the silent dealer's posterior mean $\mathbb{E}[v \mid \mathcal{I}, s]$ (dashed), and the win-conditional mean (teal). The upper shaded wedge is the winner's-curse discount and the lower wedge is bid shading. Talker bids lie in the shaded region below $\bar{b}$. The silent bid reaches $\bar{b}$ at the entry margin $s^\times$. Bottom panel: dealers with talking cost below the cutoff $\bar{\kappa}(s) = \Delta(s; p^*)$ talk. The cutoff falls to zero at $\bar{s}^*$. The figure assumes $\bar{s}^* < s^\times$ for illustration.

with high idiosyncratic values are different considerations. This middle region therefore allows informative silence to coexist with talker wins.

Above $s^{\times}$, even a silent dealer with $u = 0$ bids more than $\bar{b}$. They defeat every talker, but must still compete with the other silent dealers. A small increase in their bid can obtain additional wins only against those rivals. The talkers continue to affect public beliefs and the boundary of this region, and a sufficiently large downward deviation would again expose them to their bids. Clearing the talker support thus changes marginal competition without making the dealer certain to win.

Positive idiosyncratic values shift this classification through $s^{\times}(u)$. A stronger client or inventory motive raises the silent bid and weakly lowers the signal needed to defeat every talker. A dealer with $s < s^{\times}(0)$ can consequently lie above their own crossing after observing u. Even a high-cost dealer below the disclosure boundary can clear the talker support if their idiosyncratic value raises their bid sufficiently. The disclosure boundary concerns information release before u is known, while the bidding boundary depends on both private valuation components.

The first-price payment rule explains why silent dealers respond to the distribution of talker bids without necessarily matching their sensitivity to u. Let $G^{\text{ex}}(b \mid s)$ be the probability that bid b defeats every rival silent dealer, conditional on the focal signal and public history. Let $V^{\text{ex}}(s, b)$ be the expected common value conditional on that event. These objects integrate over the rivals' signals and idiosyncratic values. Since talker bids are independent of the common value and silent private information conditional on the history, the silent dealer's expected auction payoff is

$$U^S(b; s, u) = H(b)G^{\text{ex}}(b \mid s)\big[V^{\text{ex}}(s, b) + u - b\big]. \tag{19}$$

A higher bid raises the winning probability but also raises the payment in every state in which the dealer would already have won. A higher u increases the value of additional wins, while leaving this payment cost in place.

When silent dealers compete with one another, a better signal can offset a smaller idiosyncratic value. The winner therefore need not have the highest signal, and losing need not mean that a rival had better information about the asset. A rival with a higher idiosyncratic value may submit the higher bid. This explains why the clearing probabilities in (19) use the full bid distribution rather than a single signal cutoff. It also gives the bidding interpretation of pooling. Different combinations of information and idiosyncratic value can generate the same bid even though, at a fixed signal, the bid rises with u. This pooling concerns inference from the bid. Rivals cannot observe that bid and adjust their own bids within the simultaneous auction.

The distinction between pooling and separation consequently does not impose an ordering of bid responses to idiosyncratic values. Both sides' bids increase with u in the maintained continuation class, but their relative sensitivity depends on the winning probabilities and payoff curvature they face.[8] A talker's disclosed signal and bid identify their valuation components because only u remains

[8]Holding information and rival strategies fixed, either dealer's payoff can be written as $U(b, u) = A(b) + uQ(b)$,

private at the bidding stage. That informational property does not require a steeper response to u than a silent dealer's.

3.3 Predictions

The model connects dealers' voluntary talk to bidding, prices, and winner selection. The posterior formulas and symmetry within the talker side give direct restrictions. Predictions for transaction-price differences and comparative statics additionally depend on the relative importance of common-value information, idiosyncratic values, and bid shading. Appendix A.4 develops the empirical mappings.

1. Talk–bid crossing and silence premium. A silent-dealer win reveals favorable common-value information relative to the public history. When this information has a sufficiently large effect on bids compared with idiosyncratic values and shading, the transaction price exceeds the consensus in silent-won auctions, a "silence premium." This region motivates a larger $P_i - \bar{S}_i$ for silent wins and a smaller $P_i - S_i^{1:N_i}$ for talker wins, which we test in Section 4.1.

2. Pooling of silent dealers. In the three-region equilibrium of Proposition 3, silence pools favorable signals retained strategically with unfavorable signals retained because talking is costly. A talker can therefore win while silent dealers are present, and a talker win does not guarantee that the silent side was empty or weak. The presence test in Section 4.3 checks this implication against the alternative that silence indicates strength.

3. Compression toward the entry margin. More talks diminish silence, all else equal. The winning silent signal is the maximum over fewer draws, so it lies closer to the entry margin $s^{\times}$, where the premium is smallest. The silence premium therefore declines as the number of talks rises. Section 4.4 tests this comparative static.

4. Belief convergence and consensus predictiveness. In the Gaussian reference calculation, uncertainty falls to $\mathrm{Var}(v \mid \mathcal{M}) = 1/(\tau_v + n_T\tau_\varepsilon)$ as more signals become public. Disclosure gives dealers a common set of information on which to base their bids. This learning force is reflected in the bid dispersion relative to the dispersion of preceding talks.

The Gaussian reference posterior places weight $\beta(n_T) = n_T\tau_\varepsilon/(\tau_v + n_T\tau_\varepsilon)$ on consensus. This weight increases with the number and precision of disclosed signals. Consensus should therefore be informative about transaction prices when bids respond to the disclosed information. Prices load on the consensus more than on the winner's own talk. Under the maintained assumption that the posterior weights every disclosed signal equally, prices load on the winner's own talk only through the consensus. A silent winner's bid contains a retained signal that is absent from the talks, so consensus explains less of the price variation in silent-won auctions. Section 4.5 tests the aggregation of information.

where $A(b)$ is the payoff at zero idiosyncratic value and $Q(b)$ is the winning probability. At a differentiable interior optimum with $U_{bb} < 0$ and $Q'(b) > 0$, implicit differentiation gives $db^*/du = -Q'(b^*)/U_{bb}(b^*, u) > 0$. This expression does not rank the two sides' responses.

4 Empirical Evidence

This section tests the model's predictions. We begin by asking where transaction prices lie relative to the disclosed talks and how this gap depends on which type of dealer wins. The model predicts that silent dealers retain valuable private information which allows them to bid higher than talk consensus. This "crossing test" provides direct evidence combining prices and talk data on the selective retention of favorable information. We then explore within-dealer variation to quantify the silence premium. Finally, we provide evidence on the three-region structure of Proposition 3 and the prediction that talkers/nontalkers do not always lose/win when nontalkers/talkers are present. We conclude with evidence on the entry-margin comparative static and on information aggregation.

4.1 The crossing test

In the model, where the transaction price lies relative to the disclosed talks depends on who wins. A nontalker win reveals favorable private information that was deliberately withheld, so the winning bid reflects information that the talks do not. In the parameter region where common-value information dominates idiosyncratic values (Section 3.3), the model implies the following.

Prediction 1. *The transaction price exceeds the disclosed talks more when a nontalker wins than when a talker wins.*

We measure this with two gap measures that apply to both winner types. The *price premium* is $P_i - \bar{S}_i$, the price minus the consensus from (6). The *max-talk crossing* is $P_i - S_i^{1:N_i}$, the price minus the most optimistic talk, so named because a positive value means the price has crossed above the entire disclosed range. Both gaps are measured in percentage points of par. The winning dealer is observed in the matched auction records (Section 2), and we set $\text{Nontalker}_i = 1$ when they provided no talk on the security, and $\text{Nontalker}_i = 0$ otherwise. The crossing test asks whether both gaps are larger in nontalker-won auctions, which operationalizes Prediction 1.

Table 3 presents the results. Panel A reports summary statistics. The average transaction price exceeds the consensus by 22.7 basis points (bp) in nontalker-won auctions compared to 13.5 bp in talker-won auctions, a statistically significant difference of 9.3 bp ($t = 7.64$). The max-talk crossing provides an even sharper pattern, and the levels are also informative. In talker-won auctions, the average price falls 1.19 percentage points (pp) below the highest talk. In nontalker-won auctions, prices are higher and the gap narrows by 34.4 bp ($t = 16.9$) to 0.84 pp. The remaining columns show that the crossing is not due to outliers. Prices exceed the highest talk nearly twice as often when a nontalker wins (22.3% versus 13.2%), a difference of 9.1 pp.

Figure 5 shows both distributions by winner type. The plots show kernel density estimates of the price premium left and the max-talk crossing right, in bp, for talker-won (blue) and nontalker-won (red) auctions. The upward shift when a nontalker wins runs through the body of each distribution

Table 3: The Crossing Test

This table tests whether transaction prices deviate differently from the pre-auction consensus and from the highest individual talk depending on winner type. The price premium is $P_i - \bar{S}_i$. The max-talk crossing is $P_i - S_i^{1:N_i}$. Panel A reports summary statistics. Panel B reports regressions on a nontalker winner indicator with day fixed effects and standard errors clustered by security. ***, **, * denote significance at 1%, 5%, 10%.

Panel A. Summary Statistics

	(1) Mean	(2) Median	(3) % Above Zero	(4) Std. Dev.	(5) Obs
Price − Consensus ($P_i - \bar{S}_i$)					
Talker wins	0.135***	0.101	60.9%	1.095	16,718
	(0.008)				
Nontalker wins	0.227***	0.120	63.3%	1.213	20,423
	(0.008)				
Difference	0.093***	0.019	2.4 pp		
	(0.012)				
Price − Highest Talk ($P_i - S_i^{1:N_i}$)					
Talker wins	−1.188***	−0.450	13.2%	1.995	16,718
	(0.015)				
Nontalker wins	−0.844***	−0.268	22.3%	1.913	20,423
	(0.013)				
Difference	0.344***	0.182	9.1 pp		
	(0.020)				

Panel B. Silence Premium

	(1) All	(2) AAA/AA	(3) A/BBB	(4) BB/B	(5) Equity
Dep. var. Price − Consensus					
Nontalker wins	0.081***	0.029***	0.098***	0.118***	0.132**
	(0.012)	(0.008)	(0.020)	(0.031)	(0.060)
Dep. var. Price − Highest Talk					
Nontalker wins	0.262***	0.059***	0.178***	0.212***	0.196**
	(0.021)	(0.011)	(0.033)	(0.049)	(0.094)
Day FE	Yes	Yes	Yes	Yes	Yes
R^2 (premium)	0.191	0.423	0.359	0.327	0.380
R^2 (crossing)	0.207	0.400	0.405	0.385	0.409
Obs	37,026	12,906	9,361	9,742	4,061

rather than the tails, and it is particularly pronounced for the max-talk crossing on the right where the distribution is less left-tailed, more concentrated, and more right-tailed.

Table 3, Panel B confirms these patterns in a regression framework. For each outcome $Y_i \in \{P_i - \bar{S}_i,\ P_i - S_i^{1:N_i}\}$ we estimate

$$Y_i \;=\; \alpha_t + \delta\, \text{Nontalker}_i + \epsilon_i, \tag{20}$$

where α_t is a trade-day fixed effect absorbing market-wide price movements. Standard errors are

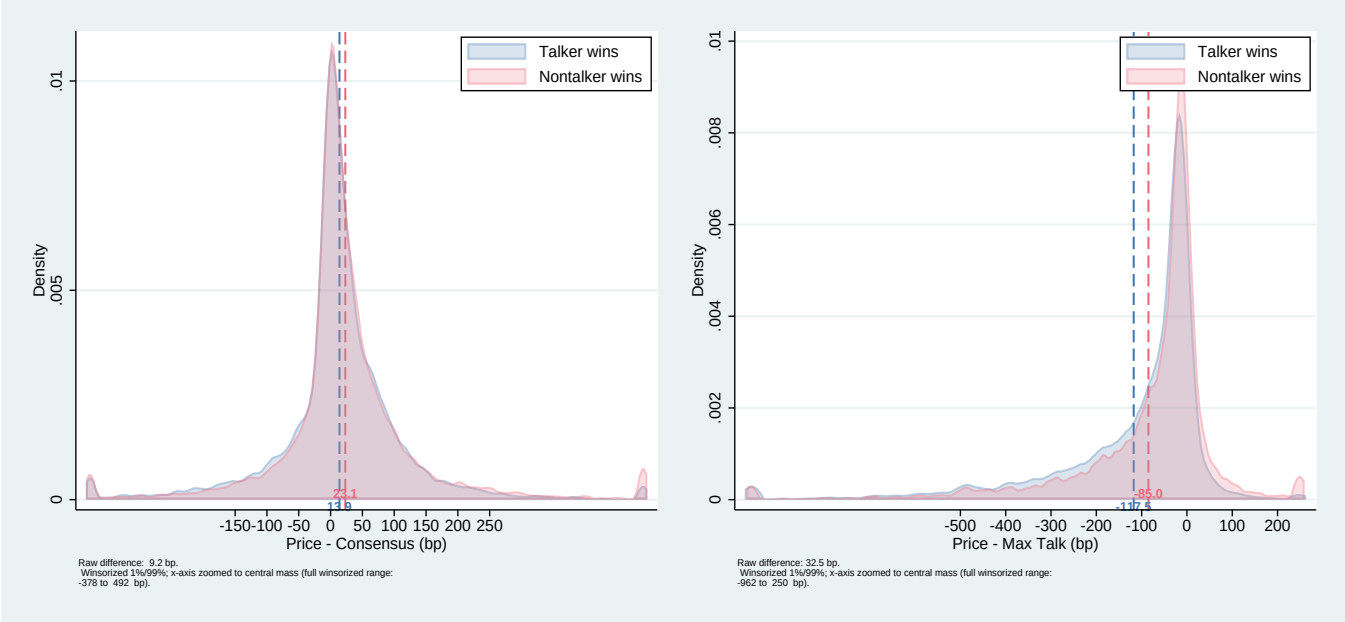

Figure 5: Crossing distributions by winner type. Kernel densities of the price premium $P_i - \bar{S}_i$ (left) and the max-talk crossing $P_i - S_i^{1:N_i}$ (right), in bp, for talker-won (blue) and nontalker-won (red) auctions. Dashed lines mark group means, winsorized at 1% and 99%, so they differ slightly from Table 3.

clustered by security. The coefficient δ compares auctions won by the two dealer types on the same day and is the regression counterpart of the Panel A differences. The silence premium is the population version of this comparison,

$$\text{Silence premium} \;\equiv\; \mathbb{E}\big[\, P_i - \bar{S}_i \mid \text{Nontalker}_i = 1 \,\big] - \mathbb{E}\big[\, P_i - \bar{S}_i \mid \text{Nontalker}_i = 0 \,\big], \tag{21}$$

defined analogously for the max-talk crossing. The Panel A differences estimate the silence premium directly, and $\hat{\delta}$ estimates it holding the trade day fixed. The estimate is $\hat{\delta} = 0.081$, or 8.1 bp ($p < 0.01$) for the price premium and $\hat{\delta} = 0.262$, or 26.2 bp ($p < 0.01$) for the max-talk crossing.

We split by tranche rating across columns (2)–(5). Consistent with the value of withheld information rising with fundamental uncertainty, the silence premium increases monotonically with credit risk, from 2.9 bp for AAA/AA to 13.2 bp for equity. The differential on the max-talk crossing is much larger outside the senior tranches, rising from 5.9 bp for AAA/AA to 17.8–21.2 bp for the junior groups. Junior tranches have greater valuation uncertainty and larger information rents from silence, so the hidden signal of a nontalker winner produces a larger deviation between the consensus and the price.

These estimates are conservative with respect to dealer asymmetry. If more optimistic talks predict winning among talkers, the talker-win benchmark is selected upward relative to the symmetric model (Section 4.6), in which the expected winning talk equals the consensus. Winner selection on optimistic talks raises both $P_i - \bar{S}_i$ and $P_i - S_i^{1:N_i}$ in talker-won auctions and therefore shrinks both

measured differentials. This auction-level selection is distinct from persistent dealer heterogeneity, which the within-dealer design in Section 4.2 absorbs. Dealer fixed effects do not remove selection on the realized talk, so the two corrections are complementary. Section 4.5 re-estimates the crossing net of this selection (equation (26)). The adjustment raises the silence premium from 8.1 to 13.2 bp and the max-talk differential from 26.2 to 31.4 bp.

4.2 The within-dealer silence premium

The crossing results could reflect persistent dealer heterogeneity rather than strategic disclosure if certain dealers systematically win at a higher price premium regardless of whether they talk. To address this concern, we restrict the sample to the auctions whose winning dealer also wins other auctions with opposite talk status, which are 95% of all auctions and ensures within-dealer variation in disclosure behavior. We then re-estimate the crossing regressions with a winning-dealer fixed effect $\gamma_{j(i)}$ added to equation (20),

$$Y_i \;=\; \alpha_t + \gamma_{j(i)} + \delta\,\text{Nontalker}_i + \epsilon_i, \tag{22}$$

so that δ estimates the silence premium of equation (21) exclusively from comparing the same dealer's outcomes across talk statuses.

Figure 6 plots the within-dealer seniority gradient of the price premium. Table B.2 in the appendix reports the estimates for both outcomes. Panel A shows that the same dealer achieves a price premium 9.9 bp larger when winning as a nontalker ($p < 0.01$). This estimate is larger than the 8.1 bp from equation (20), which compares across dealers. The shift in the price premium in Figure 5 is therefore not driven by which dealers win but by whether the winner talked. As in Table 3, the premium increases monotonically with credit risk, from 2.6 bp for AAA/AA to 18.8 bp for equity (all significant at $p < 0.01$). Panel B shows that the same dealer's winning price sits 17.6 bp closer to the highest talk when winning as a nontalker (again at $p < 0.01$). This is about two-thirds of the 26.2 bp across-dealer estimate, so part of the max-talk differential reflects which dealers win. The level and the seniority slope of the price premium survive within dealer. The max-talk effect is similarly concentrated outside the most senior tranches, though the equity estimate is imprecise.

Together these results are consistent with the strategic value of silence depending on the informativeness of the private information rather than solely on persistent dealer characteristics. Dealer fixed effects rule out explanations based solely on time-invariant dealer heterogeneity. They do not rule out time-varying dealer states such as inventory, client orders, funding or hedging needs (the model's u_j is precisely such a state), security selection, or dealer-security comparative advantage. Dealer-by-year and dealer-by-seniority fixed effects address part of this margin and leave the price premium essentially unchanged, at 9.2 bp and 9.5 bp (both $p < 0.01$). The max-talk crossing is

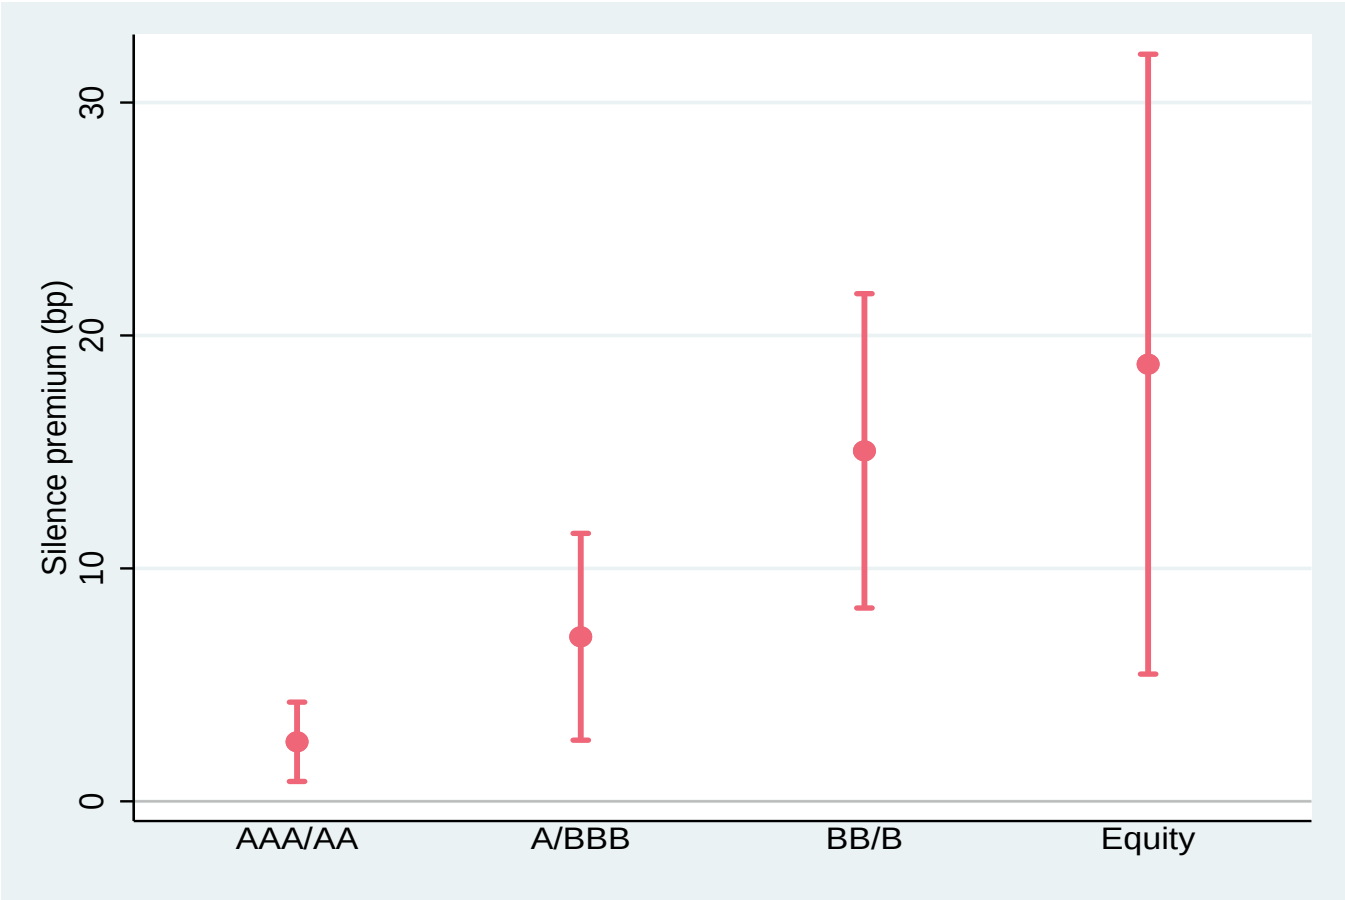

Figure 6: Within-dealer silence premium by seniority. Each point is the Nontalker$_i$ coefficient for the price premium $P_i - \bar{S}_i$ from the within-dealer regressions in Table B.2, Panel A, columns (2) to (5), in bp. Whiskers are 95% confidence intervals based on security-clustered standard errors.

23.3 bp with dealer-by-year and 15.0 bp with dealer-by-seniority fixed effects (both $p < 0.01$).[9]

4.3 Evidence on information pooling from three-region structure

The crossing test in Section 4.1 and the within-dealer silence premium in Section 4.2 establish that nontalker winners pay more than talker winners, consistent with selective retention of favorable information. These tests concern the direction of selection into silence. They leave open what silence itself means. If every silent dealer retained a signal strong enough to outbid the talker side, silence would guarantee bidding strength, and a dealer who is present but silent would never lose to a talker. Silence would then disclose strength rather than conceal it, the informational advantage

[9]The within-dealer sample contains 36 winning dealers. Because talk status and pricing behavior may be correlated across repeated observations for the same dealer, we also cluster standard errors two-way by dealer and security. The standard error of the price premium rises from 1.4 to 1.9 bp and that of the max-talk crossing from 2.4 to 4.9 bp, and clustering by dealer and date gives nearly identical values. With 36 dealer clusters, we also report wild-cluster bootstrap p-values (Webb weights, 9,999 replications, clustered by dealer), which are 0.000 for the price premium and 0.003 for the max-talk crossing. By seniority, the bootstrap p-values for the price premium are 0.057 (AAA/AA), 0.006 (A/BBB), 0.002 (BB/B), and 0.013 (equity). For the max-talk crossing they are 0.024, 0.012, 0.000, and 0.384.

that motivates silence would not survive, and silence would unravel.

A silent dealer with a favorable signal earns an information rent only if silence does not reveal the signal, and silence conceals it only because it is partially pooled. In the equilibrium of Proposition 3, dealers with weak signals also stay silent when talking is costly. They bid inside the talker range and lose to talkers with high idiosyncratic values. These weak silent dealers form the crowd in which dealers with favorable signals can hide. The prediction below tests for this crowd. A talker win against a present-and-silent dealer shows that silence does not guarantee a strong signal.

Prediction 2. *Talkers win with positive probability against dealers who are present and silent.*

We exploit the structure of BWIC lists to provide evidence on the existence of these weakly competitive silent dealers, region 2 of Proposition 3.

Identification through list-level dealer presence. The key challenge is identifying dealers who are present but choose to remain silent on a given security. While the full set of bidders is unobservable, presence on the list is the relevant notion, and the list structure identifies a subset of the dealers who received the BWIC. Each BWIC list contains multiple securities offered simultaneously. A dealer who submits a price talk on any security on a given list has demonstrably received that list and is active in that auction round. For every other security on the same list where this dealer does not talk, we can infer that the dealer was present (they saw the security) but chose to remain silent.[10] An important limitation is that a dealer may receive the list and talk on no security in it. Such dealers leave no visible indication of presence through talk, so the identified present-and-silent pairs understate presence. Winner identities provide a second indication of presence. A dealer who wins any security on the list has also received it, whether or not they talk anywhere on it, and we report this extension below.

Consider dealer j who talks on CUSIP A but is silent on CUSIP B, where both appear on the same list. If a talker wins CUSIP B, then the dealer j was present, silent, and outbid by a talker, provided they submitted a bid. If silence guaranteed bidding strength and bidding were universal, this could not happen. Silence would signal favorable information and imply a bid above every talker bid, so a silent dealer would always outbid the talker side. In the equilibrium of Proposition 3, dealer j could instead have a signal below the entry margin $s^{\times}$, where they bid inside the talker side's bid range and lose to high talker draws despite being silent. With cost-driven silence, they may also be silent with a modest signal because of a high cost draw, in which case they bid modestly and lose to talkers for the same reason.

Both channels generate present-and-silent dealers who are outbid by talkers. Conditional on present-and-silent dealers bidding, the frequency of such occurrences provides a direct measure of

[10]The construction identifies dealers who received the list and remained silent on a given CUSIP, not that they evaluated or bid on it. The model assumes every dealer submits a bid, so a silent dealer is a silent bidder. Bid submission is unobserved in the data, so the test maintains that present-and-silent dealers submit bids.

the prevalence of region 2. But because bid submission is not observed, the same frequency is also consistent with nonparticipation by some present-and-silent dealers.

The presence test. We construct the test as follows. From the dealer-level talk data (430,921 observations, of which 427,514 carry an identified dealer), we identify 83,348 unique (dealer, list) pairs. For each pair, we expand to all securities on that list, producing 440,488 (dealer, security, list) observations. Of these, 13,000 are cases where the dealer was active on the list but did not talk on the security (present and silent). Restricting to securities that received at least one talk ($n_T \geq 1$) and that traded, so that the winner's talk status is observed, we have 7,784 present-and-silent observations.

Table 4, Panel A reports the fraction of present-and-silent dealer-security pairs in which a talker won the auction. If silence guaranteed bidding strength and present-and-silent dealers bid, this fraction should be zero. We find that talkers win 24.1% of the time a silent dealer is present (standard error 1.5 pp, two-way clustered by list and dealer). Conditional on the maintained participation assumption, this significant conditional talker-win rate rejects the alternative in which every silent dealer outbids the talker side. Without that assumption, it constitutes evidence consistent with the intermediate region.

Two caveats apply. First, the rate is dealer-security-pair weighted. An auction with several identified present-and-silent dealers contributes several observations, so 24.1% is neither the fraction of relevant auctions won by talkers nor a direct estimate of the prevalence of the intermediate region. Weighting each auction equally, talkers win 35.2% of the 3,413 traded auctions with at least one identified present-and-silent dealer, and weighting each list equally gives 33.7%. The auction-weighted rate is higher because talker-won auctions tend to have fewer identified present-and-silent dealers. Second, the observations are not independent, since several come from the same auction, list, and dealer. We therefore report standard errors clustered by list and, two-way, by list and dealer. The rate is far from zero under every weighting and clustering choice.

Next, we progressively tighten the definition of "present" to address the concern that a dealer active on senior tranches might not have information relevant for junior tranches on the same list. In the same-rating restriction, we require that the dealer talked on another security of the same rating on the same list. Even within this restriction, talkers win 19.6% of the time ($N = 4{,}832$). In the same-deal restriction, we require that the dealer talk on another tranche of the same CLO deal. The rate is 32.0% ($N = 175$), though the sample is small. Tightening the definition does not drive the rate toward zero, suggesting that the presence of region 2 dealers is not driven by informational distance between securities. Conversely, widening presence to dealers who win any other security on the list raises the rate to 39.0% ($N = 79{,}739$, standard error 1.0 pp), so the talk-based construction is conservative.

Panel A also reports the test by seniority across columns (2)–(5). For the baseline specification, the talker-won rate ranges from 18.0% (A/BBB) to 31.6% (equity). For the same-deal restriction,

the rate is lowest for BB/B (10.0%), though the samples are small. The fact that region 2 behavior is prevalent across the seniority spectrum rules out the possibility that non-competitive silence is confined to a particular asset class.

Cover position relative to the talk range. The presence test rests on the maintained assumption that present-and-silent dealers submit bids. The cover bid, the second-highest bid in each auction, does not require this assumption, because a disclosed cover is an actual submitted bid. Panel B of Table 4 therefore provides corroborating, though less direct, evidence from bid positions. We classify the cover into three categories relative to the range of disclosed talks, including below all talks, within the talk range, and above all talks. If silence guaranteed bidding strength, the top two bids would both be silent bids whenever at least two silent dealers are present, and disclosed covers would routinely lie above all talks.

Panel B, column (1) shows that this condition rarely holds. Covers are above all talks in 11.7% of auctions, inside the talk range in 75.7%, and below all talks in 12.6%, so the two sides' bid distributions overlap the disclosed range rather than sit above it.[11]

In talker-won auctions (column 2), the cover falls below all talks 10.0% of the time. In these auctions, the cover lies in the range associated with weakly competitive bidding (although the identity and talk status of the runner-up are not observed). When the cover does fall below all talks, the median shortfall is 18 bp, and the mean is 47 bp, indicating that these runners-up are well below the talk range rather than marginally below.

In nontalker-won auctions (column 3), the cover exceeds all talks 14.4% of the time. In these auctions, both the winning bid and the cover exceed every disclosed talk, the pattern the model associates with competitive silent bidding (region 3), though the runner-up could also be a talker with a high idiosyncratic value. In unreported results, we find these patterns are consistent across seniority groups.

Same dealer in multiple regions. Panel C of Table 4 addresses whether the patterns above reflect dealer-level heterogeneity (some dealers always talk, others always stay silent and compete) or status changes by the same dealer across securities. We classify each dealer-security observation by observable status only: *(i)* whether the dealer talked, *(ii)* was present and silent while a talker won, or *(iii)* won as a nontalker.[12] We then ask how many dealers appear in multiple statuses across

[11] The cover's position does not identify the runner-up's region, because the runner-up's identity and talk status are unobserved. Winner's-curse shading lets a talker bid below their own talk, and a high idiosyncratic value lets a talker bid above all talks, so neither a low nor a high cover pins down the runner-up's type. The classification records where the second-highest bid falls relative to the disclosed information, which the model links probabilistically to the bidding regions. Panel B also conditions on the cover being disclosed, which occurs for 78% of trades and is positively associated with talk intensity (Table B.1), so disclosure selection can tilt the position distribution.

[12] Here, a present-and-silent appearance is inferred from presence rather than from an observed losing bid. Winning as a nontalker does not identify region 3, since an intermediate (region-2) silent dealer can win when the talker side's draw and all rival silent bids happen to be low. Even observing $P_i > \max_j S_{ij}$ would be only suggestive, because talks are not bids.

Table 4: Evidence from the Three-Region Structure

This table tests for the existence of silent, partially competitive dealers (region 2 of Proposition 3). Panel A uses list-level dealer activity to identify present-and-silent dealer-security pairs and reports the fraction in which a talker won the auction. If every silent dealer outbid the talker side and present-and-silent dealers submit bids, this fraction should be zero. The sample is traded auctions with at least one talk, so that the winner's talk status is observed. The any-activity specification identifies a dealer as present if they talked on any security on the same list. The auction-weighted row weights each auction with at least one identified present-and-silent dealer equally. The same-rating specification requires a talk on another security of the same rating on the list. The same-deal specification requires a talk on another tranche of the same CLO deal. The talk-or-win specification also counts as present a dealer who wins another security on the list, excluding the security the dealer itself won. Standard errors two-way clustered by list and dealer in parentheses, number of observations in brackets. Panel B reports the position of the cover bid relative to the range of disclosed talks (below all talks, within the talk range, above all talks). Panel C reports whether a dealer ever talks, is ever present and silent when a talker wins, or ever wins as a nontalker.

Panel A. Talker-Won Rate Among Present-and-Silent Pairs

	(1) All	(2) AAA/AA	(3) A/BBB	(4) BB/B	(5) Equity
Any activity on list	0.241	0.253	0.180	0.196	0.316
	(0.015)				
	[7,784]	[1,277]	[1,712]	[2,308]	[2,487]
Auction-weighted	0.352	0.365	0.333	0.374	0.337
	[3,413]	[687]	[778]	[918]	[1,030]
Same rating	0.196	0.253	0.139	0.162	0.274
	(0.019)				
	[4,832]	[942]	[1,209]	[1,727]	[954]
Same deal	0.320	0.341	0.357	0.100	0.379
	(0.059)				
	[175]	[44]	[14]	[30]	[87]
Talk or win on list	0.390	0.357	0.402	0.450	0.346
	(0.010)				
	[79,739]	[29,287]	[22,817]	[18,241]	[9,394]

Panel B. Cover Position Relative to the Talk Range

	(1) All	(2) Talker-won	(3) Nontalker-won
Above all talks	11.7%	8.1%	14.4%
Within talk range	75.7%	81.9%	71.0%
Below all talks	12.6%	10.0%	14.6%
Obs	21,463	9,258	12,205

Panel C. Dealer-Level Observable Statuses

	(1) Dealers	(2) Share
Total dealers in sample	67	
Ever talks	67	100%
Ever present and silent when a talker wins	40	59.7%
Ever wins as a nontalker	21	31.3%
In all three observable statuses	21	31.3%

Table 5: Compression of the Silence Premium Toward the Entry Margin

This table examines how the silence premium varies with proximity to the entry margin. Panel A interacts the nontalker indicator with the number of talks. Panel B sorts nontalker-won auctions by the decile of the winning margin, the distance between the two highest observed bids, and reports the average price premium within each bin. Negative values of the margin are due to imperfect matching (see Section 5.1). All regressions include day fixed effects with standard errors clustered by security. ***, **, * denote significance at 1%, 5%, 10%.

Panel A. Nontalker Effect by Number of Talkers

	Interaction		By bin			
	(1) Premium	(2) Crossing	(3) $N_i = 1$–3	(4) $N_i = 4$–6	(5) $N_i = 7$–9	(6) $N_i \geq 10$
Nontalker$_i$	0.176***	0.382***	0.229***	0.108***	0.091***	0.054***
	(0.038)	(0.057)	(0.069)	(0.038)	(0.024)	(0.015)
N_i	0.019***	−0.063***				
	(0.004)	(0.006)				
Nontalker$_i \times N_i$	−0.009**	−0.022***				
	(0.004)	(0.006)				
Day FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.192	0.217	0.384	0.343	0.237	0.238
Obs	37,026	37,026	4,678	5,669	9,354	16,670

Panel B. Price Premium by Winning Margin Decile (Nontalker Wins Only)

Winning margin decile	1 (low)	2–3	4–5	6–7	8–9	10 (high)
Price premium (bp)	14.8	10.7	18.3	26.0	38.6	57.3
Avg. winning margin (bp)	−21.2	1.6	5.7	15.2	43.4	122.9
Obs	1,689	3,385	3,369	3,380	3,375	1,688

different securities.

The 67 dealers in this sample talk at least once by construction. Of these, 40 (59.7%) are at least once present and silent on a security where a talker won. Twenty-one dealers (31.3%) win at least once as nontalkers on a list where they talk elsewhere. Every dealer who ever wins as a nontalker is also present and silent when a talker wins on other securities. That is, 21 dealers, nearly a third of the sample, appear in all three observable statuses. The same dealer talks on some securities, is present and silent on others when a talker wins, and wins silently on yet others.

This is consistent with dealers changing disclosure and winning status across securities, as signal-dependent behavior under Proposition 3 would generate. The dealer's action depends on their private signal relative to the cutoff $\bar{s}^*$ and the realization-specific entry margin $s^\times(\mathcal{M})$, not on a fixed dealer characteristic. Because the entry margin moves with the realized talks, the same dealer would naturally migrate across regions from one security to the next.

4.4 The silence premium and the entry margin

The crossing coefficient compares prices in nontalker-won auctions with prices in talker-won auctions. It therefore reflects the average winning silent signal, which lies in the upper tail of the signal

distribution. By contrast, the model ties the premium to the distance between the winning silent signal and the entry margin $s^{\times}$ in Figure 4, the signal at which a silent bid clears the top of the talker support and the premium is smallest. The winning silent signal is the maximum over the silent side, so a thinner silent side draws it toward the margin and compresses the premium. When more dealers talk, fewer remain silent, holding the set of active dealers fixed. The model implies the following.

Prediction 3. *The silence premium declines as the number of talks rises.*

We examine in three ways how the estimated coefficient varies with proximity to this entry-margin benchmark. First, we interact the nontalker indicator with the number of talks N_i, estimating

$$Y_i \;=\; \alpha_t + \delta\,\text{Nontalker}_i + \varphi\,N_i + \psi\,\text{Nontalker}_i \times N_i + \epsilon_i \tag{23}$$

for the same two outcomes as in equation (20), with day fixed effects and security-clustered standard errors as before. The interaction coefficient ψ captures how the silence premium changes with each additional talk and is the empirical counterpart of Prediction 3. The coefficient φ controls for the association between talk intensity and prices that is common to both winner types, so that ψ compares slopes. In the model, where the active set N is fixed and known, the premium should converge toward its entry-margin value as the number of talks rises.[13]

Table 5, Panel A is consistent with this prediction. The interaction is negative and significant for both the price premium with $\psi = -0.009$ ($p < 0.05$) in column (1) and the max-talk crossing with $\psi = -0.022$ ($p < 0.01$) in column (2).

Second, columns (3)–(6) split the sample into bins of the number of talks and re-estimate equation (20) for the price premium within each bin. The nontalker coefficient falls monotonically across the bins, from 22.9 bp with one to three talks to 5.4 bp with ten or more. The bins impose no functional form, so the monotone decline shows that the negative interaction in columns (1) and (2) is not an artifact of the linearity restriction. It also shows that the compression operates throughout the talk distribution rather than at a single threshold, as the order-statistic mechanism implies.

The bins trace a decomposition of the silence premium into the information rent earned by signals above the entry margin and a floor reflecting the dispersion of idiosyncratic values on the talker side. Talk competes the rent away and leaves the floor. In the limit of many talks, the differential estimates the entry-margin value of the premium, which equals the price paid by the marginal silent winner relative to the talker-won benchmark. The last bin places this floor at 5.4 bp, about a quarter of the 22.9 bp differential when few dealers talk, consistent with the positive limit the model implies. The marginal silent winner bids at the top of the talker range, while the

[13] In the data, total participation is unobserved, so a higher number of talks need not mean fewer silent bidders. It may instead reflect a larger active dealer set, safer or more liquid securities, or differences in list composition. We therefore read Panel A as reduced-form heterogeneity consistent with the entry-margin mechanism rather than as a direct test.

average winning talker bids below it. Hence, the limiting differential is the expected shortfall of the winning talker bid from that ceiling. The shortfall varies only with the talkers' idiosyncratic values, because the marginal winner pays the entire ceiling to win and competition at the margin leaves no information rent.

Third, we proxy proximity to the entry margin directly at the auction level. Panel B sorts nontalker-won auctions by the winning margin, the price minus the cover, and reports the average price premium $P_i - \bar{S}_i$ within each decile group. A small nonnegative margin implies that the winner barely outbid the runner-up. Panel B shows the price premium rises across the deciles of the margin, from 14.8 bp at the bottom to 57.3 bp at the top, consistent with an increasing $b^S(s)$. The low decile groups place the near-margin price premium at 10–18 bp. Net of the 13.5 bp talker-won average, they imply a near-margin differential close to zero, which is of the same order as the 5.4 bp floor from the talk side.[14]

The three comparisons agree on the mechanism. The silence premium is an information rent when the silent side is thick. It compresses toward a small positive floor that the dispersion of idiosyncratic values leaves when talk thins the silent side. Section 5.1 shows the within-auction counterpart, where competition among the silent bidders themselves transfers the rent to the seller. We next explore the informativeness of price talk.

4.5 Informativeness of price talk

The preceding tests show that silence carries information. This subsection asks whether talk still aggregates dispersed information. Public disclosures are selected toward unfavorable signals, so the disclosed pool could be misleading. We start with descriptive evidence on learning from dealer talk, and then test whether prices load on the consensus and whether its predictiveness falls when a nontalker wins.

4.5.1 Bid revisions

If price talk facilitates Bayesian learning, *(i)* dealers revise their bids toward the emerging consensus, and *(ii)* bid dispersion should be smaller than talk dispersion. Figure 7 shows average winning-bid revisions $P_i - S_i$ conditional on the winner's talk lying below or above the consensus (left panel) and compares bid dispersion with talk dispersion (right panel). Table B.3 in the appendix reports the underlying values.

The evidence in Figure 7 is suggestive of these predictions. The left panel conditions on the dealer winning and shows that winners whose talk is below consensus revise upward by 47 bp on

[14]The diagnostic is indirect. A small nonnegative margin identifies a close auction but not a winner near the entry margin $s^\times$, because a similarly strong signal held by another silent bidder also produces a small margin. The bottom decile has a negative average margin, which cannot arise under the institutional definition of the cover and reflects the matching frictions of Section 5.1, so the lower-bound interpretation attaches to the low nonnegative range. Identifying the winning bidder's equilibrium information rent would require full bid data, runner-up identities, or structural estimation.

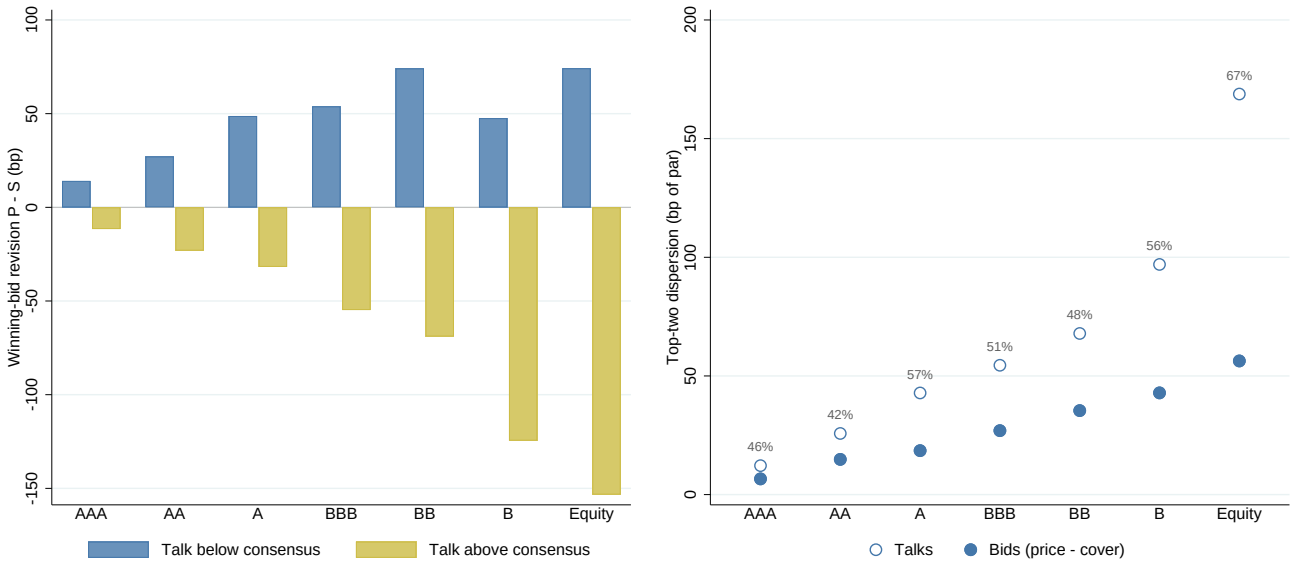

Figure 7: Winning-bid revisions and dispersion. Left: average winning-bid revision $P_i - S_i$ for winners whose talk lies below (blue) or above (yellow) the consensus. Right: top-two dispersion of talks (open) and bids (filled), with the percentage reduction labeled. Underlying values are in Table B.3.

average, while those above consensus revise downward by 51 bp, a roughly symmetric pattern with larger revisions for junior tranches. Because a below-consensus talker must generally raise their bid substantially to win while an above-consensus talker can shade downward and still win, winner selection can generate movement toward consensus without Bayesian updating.

The right panel compares the within-auction dispersion of talks with the within-auction dispersion of bids. Because only the two highest bids are observed, the transaction price and the cover, we measure dispersion on both sides by the same statistic, the gap between the two highest observations. Bid dispersion is the winning margin $P_i - C_i$, and talk dispersion is the gap between the highest and the second-highest talk, both winsorized at the 1st and 99th percentiles and computed on auctions with a disclosed cover. Bid dispersion of 0.226% of par is roughly half of talk dispersion of 0.482%, with the reduction ranging from 43–46% for the senior tranches to 67% for equity.[15]

Winning bids thus move toward the consensus, and the two observed bids disperse far less than the talks, consistent with dealers updating toward the pooled public information.

4.5.2 Information aggregation

In the model, talks are truthful, observers account for the endogenous selection into silence, disclosed signals enter the public posterior with equal weight, and retained signals enter prices only through

[15]Using the same top-two statistic on both sides removes the mechanical difference in the number of observations per auction, but not every order-statistic effect. The top two talks are drawn from the N_i talkers, whereas the top two bids are drawn from all bidders, so the two gaps differ in their underlying populations. A dealer-level learning regression of $b_{ij} - S_{ij}$ on $\bar{S}_{i,-j} - S_{ij}$ provides a direct test, but this test would require losing bids which are nonpublic.

the winning bid of a silent dealer. The model implies the following.

Prediction 4. *Despite selective disclosure, price talk aggregates dispersed information. Prices load on the talk consensus rather than on the winner's own talk, and the consensus is less predictive when a nontalker wins.*

We test both parts in turn.

Price decomposition. Table 2 already establishes that consensus surprises predict price surprises, and regression (2) finds a near-unity pass-through from consensus to prices. The model is more specific. When a talker wins, the equilibrium price reflects the consensus and not directly any individual signal. To check this, we measure all variables relative to the machine-learning benchmark R_i of Section 2.3, as in equation (1), and decompose the price surprise $\tilde{P}_i = P_i - R_i$ into contributions from the winner's own talk $\tilde{S}_i$ and the other talkers' average $\tilde{\bar{S}}_{-i}$,

$$\tilde{P}_i = \alpha_t + \beta_{\text{own}}\, \tilde{S}_i + \beta_{\text{others}}\, \tilde{\bar{S}}_{-i} + \epsilon_i, \tag{24}$$

where α_t is a trade-day fixed effect. Standard errors are clustered by security. The sample is restricted to talker wins, unlike regression (2) which uses all BWICs.

Table 6 reports the results for talker-won auctions. Prices load heavily on others' information with $\hat{\beta}_{\text{others}} = 0.754$ ($p < 0.01$) compared to the winner's own talk with $\hat{\beta}_{\text{own}} = 0.273$ ($p < 0.01$). About three-quarters of price variation reflects the information of other dealers, consistent with the first part of Prediction 4.[16] Dealers' own revealed information via their talks matters for their bids and ultimately prices. Other dealers' talks, however, matter more in aggregate. Prices hence aggregate not only the winner's information but also other dealers' revealed information.

The seniority splits across columns (2)–(5) reveal meaningful heterogeneity. Mezzanine tranches (A/BBB and BB/B) show the lowest own-talk contribution at 17% each, while equity tranches show the highest at 37%. The decomposition explains nearly all price surprise variation, with R^2 exceeding 0.97 across all groups.

A sharper test of the first part of Prediction 4 regresses the price surprise on the consensus and the winner's own talk jointly,

$$\tilde{P}_i \;=\; \alpha_t + \beta\, \tilde{\bar{S}}_i + \gamma\, \tilde{S}_i + \epsilon_i, \tag{25}$$

within talker-won auctions, with day fixed effects and security-clustered standard errors. Under the equal-weight assumption, the winner's own talk carries no weight beyond its share in the consensus, so $\gamma = 0$ at every number of talks, and the null is free of the pooling confound. Panel B of Table 6

[16] This regression is evidence that other dealers' talks matter more in aggregate, not as a direct test of the $1{:}(n_T - 1)$ decomposition. The model's prediction holds conditional on n_T, with the own-signal weight equal to β/n_T. Because n_T varies across auctions, the pooled coefficients confound the weight with the distribution of n_T (an own-talk share of 27% corresponds to roughly four talkers, whereas the average auction has about eight). Equation (25) provides the direct test, with a null that is free of this confound.

rejects the null. The winner's own talk carries an additional weight of $\hat{\gamma} = 0.231$ (s.e. 0.026) beyond the consensus, whose coefficient falls to $\hat{\beta} = 0.749$. The rejection runs in the direction of the talker asymmetry documented in Section 4.6. The extra own-talk weight is largest for equity (0.335) and insignificant for A/BBB (0.016). Across bins of the number of talks, $\hat{\gamma}$ is -0.09 (s.e. 0.23) with one to three talks, 0.18 (0.13) with four to six, 0.30 (0.04) with seven to nine, and 0.18 (0.02) with ten or more, so the rejection is not confined to auctions with few talks. Total pass-through $\hat{\beta} + \hat{\gamma} = 0.980$ is close to, though statistically below, one. The consensus-only restriction $\gamma = 0$ gives $\hat{\beta} = 0.965$ on the same sample, with almost the same fit ($R^2 = 0.971$ against 0.973), so the consensus remains the dominant determinant of the price.

Restricting equation (25) to unit total pass-through, $\beta + \gamma = 1$, also serves the crossing test. A direct adjustment for the dealer asymmetry in Section 4.1 estimates, within talker-won auctions,

$$P_i - \bar{S}_i \;=\; \alpha_t + \rho\,(S_i - \bar{S}_i) + \epsilon_i, \tag{26}$$

where we measure the winning talker's disclosed talk S_i relative to the consensus $\bar{S}_i$, as in Table B.3, and cluster standard errors by security. The coefficient ρ measures how much of the winner's optimism relative to the consensus passes into the price. Panel C gives $\hat{\rho} = 0.164$ (s.e. 0.023), close to $\hat{\gamma}$ in Panel B, as the restriction implies. The winning talk lies on average 32 bp above the consensus (median 13 bp), so winner selection raises the average talker-won premium by about 5 bp.

The fitted selection component $\hat{\rho}\,(S_i - \bar{S}_i)$ is then subtracted from the outcome in talker-won auctions, and equation (20) is re-estimated on the adjusted outcome, so the crossing is measured net of winner selection on optimistic talks. The adjustment raises the silence premium from 8.1 to 13.2 bp and the max-talk differential from 26.2 to 31.4 bp (both $p < 0.01$). Using the seniority-specific $\hat{\rho}$, the adjusted silence premium is 4.1 bp for AAA/AA, 11.2 bp for A/BBB, 14.9 bp for BB/B, and 30.9 bp for equity, against 2.9, 9.8, 11.8, and 13.2 bp unadjusted. The unadjusted estimates in Table 3 are therefore conservative.

Nontalker information weakens consensus predictiveness. Table 2, Panel B, establishes that the consensus predicts prices for both winner types and less so when a nontalker wins. Panel A of Table 7 extends this comparison by seniority, estimating the predictiveness regression of equation (2) separately for talker-won and nontalker-won auctions within each seniority group. The predictiveness gap between winner types widens with credit risk, from 1.0 percentage point of pass-through for AAA/AA to 3.7 for equity. This is consistent with silent dealers' information being more relevant the lower the rating.

Panel B formalizes the comparison, estimating

$$\tilde{P}_i \;=\; \alpha_t + \beta\,\tilde{\bar{S}}_i + \delta\,\text{Nontalker}_i + \eta\,\tilde{\bar{S}}_i \times \text{Nontalker}_i + \epsilon_i, \tag{27}$$

with standard errors clustered by security. The interaction coefficient η captures how the predictiveness of the consensus changes when a nontalker wins. The indicator coefficient δ captures the associated level shift.

Panel B adds the nontalker indicator and its interaction with the consensus surprise to the pooled predictiveness regression. Price surprises are 12.7 bp higher when nontalkers win ($p < 0.01$), so the silence premium survives the benchmark adjustment and the consensus control. The silence premium δ increases with credit risk, from 3.7 bp for AAA/AA to 15.2 bp for equity. The interaction term $\eta = -0.025$ ($p < 0.01$) indicates that consensus is 2.5 percentage points less predictive when nontalkers win, consistent with the second part of Prediction 4. The interaction grows in magnitude and significance from close to zero for AAA/AA to 3.5 pp for equity, consistent with the model's prediction that the information rent from silence is greatest where fundamental uncertainty is highest.

4.6 Dealer asymmetry

In the model, after disclosure talkers share their common-value information and differ privately only in independent idiosyncratic values. This symmetry across dealers is maintained for tractability throughout. It delivers the symmetric bidding schedules and the posterior formulas, and it carries a sharp implication. Conditional on a public history and a talker win, each talker is equally likely to be the winner. Proposition A20 therefore implies that the winning talker's signal rank is discrete uniform on $\{1, \ldots, n_T\}$ and that their expected disclosed signal equals consensus. Corollary A7 describes the corresponding selection on idiosyncratic value. Section A.4 of the appendix states the exact discrete distribution of the winning rank and the tests that respect it.

In the data, the winning talker's rank is not uniform. More optimistic talks predict winning. In the 16,249 talker-won auctions with at least two talks, the winning talker's normalized rank averages 0.594 (s.e. 0.002) against the null of one-half. The most optimistic talker wins 18.4% of these auctions, against 11.9% under the discrete uniform law, and the winning talk exceeds the consensus by 29 bp on average. In a dealer-level linear probability model with auction fixed effects, a talk one percentage point above the consensus raises the probability of winning by 1.7 percentage points (s.e. 0.1). The discrete-uniform null is rejected at every talk count from two to twelve (Appendix A.4.6). The rejection is expected, because dealers differ in coverage, clientele, and balance sheets. This fact rejects ex-ante symmetry but cannot separate three channels: *(i)* dealers may differ in signal precision or idiosyncratic values, *(ii)* part of the idiosyncratic value may be known before disclosure and correlated with the talk, or *(iii)* talkers may retain common-value information that the talk does not reveal.

The evidence in this section rests on results that do not require symmetry. The crossing and presence tests use the direction of sorting into silence and the within-dealer comparison, and winner selection on optimistic talks biases the crossing toward zero (Section 4.1). The equal-weight price

decomposition in Section 4.5 does require symmetry and is therefore descriptive.

The evidence in this section supports Predictions 1–4. Prices exceed talks more when nontalkers win, within dealer as well as across dealers. Talkers win against present-and-silent dealers, so silence pools weak signals with strong ones due to talking costs. This partial pooling stops standard unraveling (Grossman, 1981; Milgrom, 1981). The silence premium compresses in the all-talk limit toward a floor that captures the dispersion in idiosyncratic values. Price talk aggregates dispersed information as the consensus predicts prices at near-unity pass-through despite strategic selection of who talks. In turn, the consensus loses predictive power exactly when a silent winner's retained information enters the price. The market thus learns from the talks despite their selection, and what it does not learn enters through winning bids. The next section asks what this effect does to market quality.

5 Market Quality Implications

Section 4 establishes that silence carries information into prices. This section shows how strategic silence affects market quality. Two margins, prices and allocations, organize our analysis. On prices, nontalker wins occur at prices exceeding public valuations. To understand whether the winning dealer keeps this premium as an information rent or pays it to the seller, we explore whether the premium widens the winning margin between the top two bids or is absorbed into both the price and cover levels. On allocations, we ask what the winner's type reveals about the use of information and of idiosyncratic values, separating the informational content from allocations and welfare. We then ask what dealers' voluntary initiation of public price talk implies about collusion, given the evidence that silent winners do not capture significant information rents.

5.1 Winning margin and seller revenue

Seller revenue equals the price at which the auction awards the asset to the highest bidder. Section 4 already shows that, conditional on the winner's type, silence raises the price the seller receives by 8 bp on average, and by up to 13 bp in the junior tranches. But how is the value of retained information divided between the winning dealer and the seller? The division depends on whether silence leaves the market's hidden information as the winner's private monopoly or exposes it to competition among the silent. The part the winner keeps is an information rent, while the part the winner pays is the premium and accrues to the seller as revenue.[17]

The winning bid exceeds two references, the cover C_i and the consensus $\bar{S}_i$, and the identity

$$\underbrace{P_i - \bar{S}_i}_{\text{price premium}} = \underbrace{P_i - C_i}_{\text{winning margin}} + \underbrace{C_i - \bar{S}_i}_{\text{cover premium}} \tag{28}$$

[17]The comparison across disclosure regimes, where revenue under selective disclosure is set against full disclosure and no disclosure, is explored in Section 7.

splits the price premium into two components, and differencing across winner types splits the silence premium the same way. The cover premium is revenue that rival bids secure for the seller, since bidders must beat that cover to win. If it is sensitive to nontalkers winning, the silence premium is revenue to the seller secured by competition among informed silent dealers. The winning margin is the amount the winner pays beyond that lower bound. If it is sensitive to nontalkers winning, the silence premium rests on the winning bid alone and reflects the threat of informed competition.

The premium split in (28) ties back to the rent-versus-revenue division. The winner's rent is their valuation net of the price,

$$\text{Dealer rent}_i \;=\; V_i - P_i \;=\; \underbrace{(V_i - C_i)}_{\text{valuation over cover}} \;-\; \underbrace{(P_i - C_i)}_{\text{winning margin}}\,, \tag{29}$$

where $V_i = v + u_{j(i)}$ is the asset's value to the winning dealer $j(i)$. The margin is observable while the first term is not, but an informed cover bounds it because a runner-up holding correlated favorable information bids close to what the asset is worth to the winner. Realized competition therefore both delivers the premium to the seller and delivers that the rent left with the winner is small.

If the favorable signal the winner retains were exclusive, rival bids would rest on public information, and equilibrium bidding would let the informed dealer win at a small margin over the uninformed rivals while keeping the advantage as an information rent (Engelbrecht-Wiggans, Milgrom, and Weber, 1983; Hendricks and Porter, 1988). The expected rent in (29) then appears in no bid. Both the price and the cover would lie at their talker-won levels, and the silence premium would be close to zero. The seller gains from silence only if the winner pays for the advantage, and the winner pays only under competition. In the model, signals are correlated through the common value, so other silent dealers hold favorable signals in the states in which the winning nontalker does and bid on them. Outbidding an informed rival requires the winner to pay part of the advantage. The silence premium is this competed-away rent. It appears in the cover component, because the informed runner-up sets the price the winner must beat. This is the linkage mechanism of Milgrom and Weber (1982). Affiliated signals tie rival bids to the winner's information, and the tie limits the rent the winner can protect by shading.

The winning margin, measured against the disclosed cover, is the observable that separates the two cases.[18] Idiosyncratic values disperse bids under any disclosure regime, so the margin is positive in talker wins as well. The comparison across winner types nets out this common dispersion, and the nontalker coefficient on the margin estimates the margin component of the silence premium.

Table 8 reports the winning margin by winner type, unconditionally, by seniority, and by talk intensity, together with the regression counterpart. The raw comparison in Panel A is misleading.

[18]The data contain a small number of negative matched values (see Table 5, Panel B), which cannot arise under the institutional definition of the cover. They reflect TRACE–Empirasign matching frictions, including timing differences between the reported trade and the auction, trade corrections, and disclosed covers that need not correspond to the executed second-highest bid. We therefore read the empirical variable as the winning margin measured with matching error.

Unconditionally, the winning margin is 24.7 bp for talker wins compared to 23.3 for nontalker wins, a difference of only -1.4 bp ($p < 0.01$). This comparison confounds winner type with tranche composition, because talker wins are disproportionately concentrated in junior tranches where the margin is mechanically larger. Within the seniority groups, the margin is modestly larger for nontalker wins for AAA/AA ($+1.4$ bp, $p < 0.01$) and A/BBB ($+2.0$ bp, $p < 0.05$). For BB/B, the difference is small and insignificant (0.9 bp). For equity, the sign reverses (-3.3 bp, $p < 0.10$). This suggests that nontalker wins do not widen the winning margin. The within-seniority differences are an order of magnitude below the 9–41 bp margin levels and well below the silence premium, so the raw comparisons point to the cover component.

Panel B estimates the margin component in the specification of equation (23), with the winning margin $P_i - C_i$ as the outcome. The pooled nontalker coefficient on the margin is -2.5 bp ($p < 0.01$). Much of this effect reflects the tranche composition documented in Panel A, because talker wins concentrate in the junior tranches where margins are larger. Once controls for the number of talks are included and the sample is split by seniority, the coefficient is 0.7 bp or less for the investment-grade tranches and modestly negative for the junior tranches.

The identity (28) then assigns the silence premium to the cover component. On the common sample of 29,178 auctions with a matched cover, and with day fixed effects, the nontalker coefficient is 6.0 bp on the price premium, -2.2 bp on the winning margin, and 8.2 bp on the cover premium $C_i - \bar{S}_i$. Controlling for the number of talks and seniority, the three coefficients are 7.9, -0.9, and 8.8 bp. Relative to talker wins on the same day, the cover premium in nontalker-won auctions is therefore at least as large as the full silence premium. Hence, the runner-up bid carries the retained information as much as the winning bid does, which suggests it is shared information. The winner retains no measurable advantage over the nearest rival, and the premium the seller collects is a transfer, that is, a rent competed away by the silent rivals.

The same competition appears in Section 4.4, where additional talks thin the silent side and compress the premium toward the 5.4 bp floor that the dispersion of idiosyncratic values leaves. The flat margin is the bid-spacing counterpart of that result.[19] Seller revenue rises under a nontalker win while the winning margin stays fixed, because competition among silent dealers raises the entire top of the bid distribution. Overall, the silence premium is largely rent that realized competition among informed silent bidders transfers from the winning dealer to the seller, and not the exclusivity of private information or the threat of competition itself. Silence conceals information from the public but does not eliminate dealer competition.

[19] Two qualifications attach to the decomposition. The margin regression runs on the subsample with matched covers, so the components need not sum exactly. And the margin bounds the winner's advantage relative to the runner-up bid only. A component of retained information that lifts the two highest bids together leaves the margin unchanged, so the winner's equilibrium information rent is not identified, as noted in Section 4.4.

5.2 Allocations and the use of information

Strategic silence weakens the link between allocations and dealers' private information. In the equilibrium characterized in Section 4.3, talkers win in states where every silent dealer's signal lies below the entry margin $s^{\times}$ and the talker side draws a high idiosyncratic value u. Because high-signal dealers optimally remain silent, talker wins necessarily occur when no sufficiently strong private information is present to overturn the public consensus. As a consequence, the winning bidder in talker-won auctions need not be the dealer with the highest posterior valuation. In particular, a silent dealer with a moderately more optimistic posterior valuation may choose not to bid aggressively due to the winner's-curse adjustment, allowing a talker with a lower posterior valuation, based only on public information, to win the asset. Thus, price talk can reduce the winning probabilities of the highest-signal bidders.[20]

Allocative efficiency and the use of information rank the two winner types in opposite directions. Because the price is a transfer between seller and winner, the efficient owner is the dealer with the highest idiosyncratic value. Talker wins sort well on this margin. Once their signals are public, talkers bid monotonically in their idiosyncratic values, so a talker win goes to the talker with the highest such value. Since idiosyncratic values are realized after disclosure, this sorting is unrelated to the low silent signals that make talker wins possible. A nontalker win instead selects the winner on the retained signal, and because silent bids respond only weakly to idiosyncratic values, the winning nontalker's value is close to a random draw. On the information margin the ranking reverses. A nontalker win goes to the dealer with the most favorable retained information, whereas a talker win occurs only when no retained signal is strong enough to overturn the consensus. Nontalker wins thus use the market's information but not idiosyncratic values, and talker wins do the reverse. The talker wins against present-and-silent dealers documented in Section 4.3 are accordingly not allocation failures. The policy comparison in Section 7 quantifies the split. Mandating disclosure moves every auction to sorting on idiosyncratic values and perfects the allocation, banning talk moves every auction to sorting on signals and worsens it, and both effects are of the order of the idiosyncratic component.

These two margins leave a tension. Price talk aggregates dispersed information and, as we will see in Section 6, is associated with lower trade failure rates and a higher likelihood of transaction completion (see Table B.1). The same disclosure that improves execution weakens the informational content of allocations, because the dealers whose information would sort allocations best are the ones who stay silent. What strategic silence does to outcomes is therefore two-sided. The remaining

[20] This is an informational statement, not a welfare ranking. With $u_j = 0$ the asset has the same realized value v to every dealer, so allocating it to a talker rather than a more optimistic silent dealer does not reduce realized surplus. The silent dealer has a higher posterior, not a higher realized valuation. With $u_j > 0$ the efficient allocation assigns the asset to the dealer with the highest u_j, and a dealer bidding on information alone can displace it, and conversely. This establishes that surplus misallocation is possible, not how it compares across disclosure regimes. Section 7 makes the comparison numerically. Selective disclosure lowers the expected misallocation relative to banning talk, raises it relative to mandating full disclosure, and the surplus at stake is small in both comparisons.

market-quality question concerns conduct, and we turn to it next.

5.3 Public price talk and collusion

Dealers' voluntary initiation of public price talk raises a question about collusion. Consider a cartel that already shares information privately and coordinates bids to maximize joint profits. Its private arrangement allows members to pool their assessments, decide who acquires the security, and avoid competing against one another. A reason to introduce public talk must therefore explain the benefit from making information available to the seller or outside bidders.[21]

A cartel benchmark. A benchmark illustrates the cost of full public revelation. Suppose all bidders belong to a risk-neutral cartel that pools its information and submits one effective bid. Conditional on the seller's initial public information, let μ_C be the cartel's common-value estimate and let $m = \mathbb{E}[\mu_C]$ be the seller's estimate. Apply the calibrated reserve rule from Section 6, initially setting idiosyncratic values, the reservation premium, and the seller shock to zero. The seller sets the reserve before observing the bid. At reserve m, the cartel buys when $\mu_C > m$ and otherwise forgoes trade, earning expected profit $\mathbb{E}[(\mu_C - m)^+]$. Full public revelation before the reserve is set makes the seller's estimate equal to μ_C and reduces the cartel's expected acquisition profit to zero.

The cartel earns an information rent by choosing when to buy from a less informed seller. Favorable private information lets it acquire the security at a reserve below its own assessment. Unfavorable information lets it avoid the purchase. The comparison between private and public information sharing relates to the reduction in an informed bidder's rent when their information becomes public in Milgrom and Weber (1982). Private information sharing helps cartel members use this option while preserving their information advantage over the seller, but full public revelation removes it. Full revelation also weakly reduces cartel profits when the benchmark includes independent idiosyncratic values and the seller shock.[22]

Full revelation and voluntary partial disclosure act differently, however. A cartel with unfavorable information may benefit from revealing it in a particular auction if the seller then lowers an otherwise prohibitive reserve. Whether the cartel benefits from introducing this option before learning its information also depends on what the seller infers from silence.[23]

The commitment-to-silence problem. The competitive model supplies an individual incentive for public communication. A dealer with unfavorable information may disclose and then compete

[21] This comparison takes the cartel's internal reporting and enforcement problems as resolved. McAfee and McMillan (1992) study transfers and internal allocation in bidding rings. In common-value auctions, Hendricks, Porter, and Tan (2008) show that efficient collusion can conflict with truthful information sharing when the reserve binds.

[22] Let $U = \max_i u_i$ be the idiosyncratic value of the cartel member receiving the asset. Under the independence conditions in Section 6, expected profit before full revelation is $\mathbb{E}[\phi_\xi(\mu_C - m + U)]$ and afterward is $\mathbb{E}[\phi_\xi(U)]$. The optimized value ϕ_ξ is convex because it is the supremum of affine trading payoffs and the zero payoff from nonparticipation. Since $\mathbb{E}[\mu_C - m] = 0$ and U is independent of common-value information, Jensen's inequality conditional on U gives $\mathbb{E}[\phi_\xi(\mu_C - m + U)] \geq \mathbb{E}[\phi_\xi(U)]$. The comparison holds cartel coordination, the reserve rule, and the absence of outside bidders fixed.

[23] Inference from withheld information is central to voluntary-disclosure models (Milgrom, 1981). The full-revelation benchmark consequently does not rank all institutions allowing partial public talk.

on their idiosyncratic value. With a reserve, they may also gain by correcting an optimistic seller assessment. These incentives arise from their own expected payoff and need not maximize the joint profits of dealers. This gap between individual and joint incentives is a commitment problem.

A dealer's gain from talking is at most their idiosyncratic rent, $\Delta(s; p^*) \leq \bar{u}$, and talking consumes the cost κ_j. Rival dealers bear two consequences. After an unfavorable disclosure, the rise in the talker's chance of winning comes at their expense and adds nothing to joint dealer profits. The disclosed signal also becomes public information, and each remaining silent dealer's posterior then places less weight on their own signal (this is the force behind the rent reduction in Milgrom and Weber (1982) noted above). The regime comparison in Section 7 bounds what the disclosure margin can transfer. Mandating full talk instead of banning it lowers dealer surplus per auction while raising total surplus by less.

Dealers as a group would therefore plausibly prefer a commitment to collective silence, which would in addition save the talking costs.[24] Such a commitment is not available to competitive dealers. Whenever talking against silent rivals is profitable at some signal, that is $\Delta(s; 0) > 0$, a dealer with a sufficiently low talking cost gains from breaking collective silence (Lemma A40). This is the same payoff condition that supports a nontrivial equilibrium in Proposition 2. The condition for equilibrium talk to exist is thus the condition under which collective silence cannot be sustained.

The equilibrium itself needs no such binding. It is self-enforcing under unilateral deviations by construction, and under the contraction condition of Proposition A18 it is also locally stable, since best-response iteration returns to the fixed point after small perturbations of the talk policy. The joint deviation to collective silence instead fails the internal consistency requirement of coalition-proof Nash equilibrium (Bernheim, Peleg, and Whinston, 1987), because the deviating coalition recreates the incentive to talk inside itself. In this sense dealers can sustain no self-enforcing agreement beyond the noncooperative outcome within a single auction.[25] A cartel would supply exactly the missing commitment and would internalize the rents that a member's talk dissipates. Public talk that tracks individual incentives at the expense of joint dealer profits is consequently easier to reconcile with competition than with coordination (subject to the reserve motive above).

The scope of self-enforcement. First, self-enforcement concerns the decision to talk itself. The model treats a talk as exact disclosure of the signal (Section 3.1). Without that verifiability, truthful reporting in a single auction would require separate discipline. In the data, talks are attributed and can be compared with subsequent transaction prices, so dealers who interact repeatedly can build reputations for accuracy. The informativeness of talk (Section 4.5) indicates that this discipline succeeds. The same logic bounds what mandating talk can achieve. Forcing talk does not force its

[24] The numerical comparison brackets the voluntary regime rather than validating its dealer surplus level, so this preference is plausible rather than established.

[25] Coalition-proofness is defined for complete-information games. We use it as an interpretive benchmark for the deviation logic rather than as a formal refinement of the incomplete-information game. With multiple equilibria, an agreement to coordinate on another equilibrium is self-enforcing by construction. The claim therefore concerns agreements outside the equilibrium set, collective silence in particular (Lemma A40).

informativeness, so the credibility of forced talk depends on the same verification and reputation discipline.

Second, the static analysis also leaves a repeated-game route to collective silence. BWIC auctions recur, talks are public and attributed, and a deviation from a silence agreement would be observable, so sufficiently patient dealers could in principle sustain collective silence through future punishment (Fudenberg and Maskin, 1986). The data speak against this route. Talk occurs in 93.6% of auctions across the full sample period. Sustaining silence would require coordinating many dealers (over 60 in our sample) whose participation shifts from list to list, against a deviation gain that is immediate and private and a punishment that is diffuse. The prevalence of talk suggests that this discipline is absent in practice.

Third, while dealers' voluntary initiation of price talk therefore does not support collusion, public communication could help a cartel monitor conduct or enforce an agreement when actions and sales are privately observed (Harrington and Skrzypacz, 2011). Applying that explanation requires specifying what public talk accomplishes beyond the cartel's existing private arrangement. Changes in seller participation or outside bidding could provide other benefits. The empirical distinction turns on how publicity changes behavior and payoffs, rather than on the presence of communication alone.

The winning-margin evidence in Section 5.1 bears on this distinction. A cartel that designates the winner suppresses the competing bids, so coordination predicts a low cover and a wide winning margin in the auctions it controls. The data show the opposite pattern. The cover rises with the price and the margin stays fixed, as realized competition among bidders predicts. Cover disclosure provides a further check. Among traded auctions with at least one talk, the cover is disclosed after 83.2% of talker wins and 77.5% of nontalker wins. With day fixed effects the difference is -4.3 pp (s.e. 0.4). It falls to -0.8 pp (0.4) once the number of talks and seniority are controlled, and to -1.5 pp (0.3) with list fixed effects, so silent wins are not associated with a material withholding of the auction outcome. Further distinctions, using persistent dealer-pair links built from lagged talks, reciprocity and rotation in wins, and interactions of talk with link strength, are left for future work.

6 Seller Reserves and Auction Failures

The baseline model conditions on execution. This extension adds a seller who learns from the talks and requires a minimum price, which connects disclosure to execution outcomes beyond price, including auction failure.

6.1 The reserve mechanism

Price talk can change the price a seller requires and other auction outcomes. A dealer with unfavorable information may find that even their best bid falls short of an optimistic seller's reserve. Revealing the signal can lower the reserve and make trade possible. This gives disclosure a benefit beyond its effect on competition among bidders.[26]

To illustrate this mechanism, we retain exact disclosure as the communication technology. The seller observes the public history $\mathcal{I}$ and forms the full posterior mean $\widehat{v}(\mathcal{I}) = \mathbb{E}[v \mid \mathcal{I}]$, including information from silence. For a reservation premium ω, consider the reserve rule

$$R(\mathcal{I}) = \widehat{v}(\mathcal{I}) + \omega. \tag{30}$$

Trade occurs when the highest submitted bid reaches the reserve. The rule adjusts the seller's required price to their information about the asset.[27]

Additional information improves the accuracy of this calibrated reserve. If $\mathcal{I}'$ refines $\mathcal{I}$, the law of total variance gives

$$\mathbb{E}\Big[\big(v - \mathbb{E}[v \mid \mathcal{I}']\big)^2\Big] \leq \mathbb{E}\big[(v - \mathbb{E}[v \mid \mathcal{I}])^2\big]. \tag{31}$$

The same inequality applies to the reserve's squared error relative to $v + \omega$. Whether better information reduces failures depends on how the reserve compares with bids. Disclosure can affect both.

First hold the bidding environment fixed to isolate learning by the seller. Suppose they receive n exogenously supplied Gaussian signals, giving posterior mean μ_n and variance $V_n = (\tau_v + n\tau_\varepsilon)^{-1}$. Write the highest bid as $B = v + X$, where $X \sim \mathcal{N}(m_X, \sigma_X^2)$, $\sigma_X^2 > 0$, is independent of v and the seller's signals. Let $q_n = \Pr(B < \mu_n + \omega)$ be the failure probability, and let Φ denote the standard normal cdf.

Proposition 5 (Reserve information and failure)**.** *Suppose the seller uses $R_n = \mu_n + \omega$ and the highest bid satisfies $B = v + X$, with $X \sim \mathcal{N}(m_X, \sigma_X^2)$ independent of v and the signals used to construct μ_n. Then*

$$q_n = \Phi\left(\frac{\omega - m_X}{\sqrt{\sigma_X^2 + V_n}}\right). \tag{32}$$

If $m_X > \omega$, then $q_{n+1} < q_n < 1/2$ and the reduction $q_n - q_{n+1}$ strictly decreases with n. If $m_X = \omega$, then $q_n = 1/2$ for every n. If $m_X < \omega$, then $q_{n+1} > q_n > 1/2$.

When bids exceed the full-information reserve on average, some failures arise because the seller

[26] In Goldstein, Xiong, and Yang (2025), a less-informed investor may share information so that a better-informed counterparty trades against the sender's signal error and reduces the sender's price impact. Here disclosure changes rival bidding and the seller's posterior reservation value.

[27] Equation (30) minimizes the expected squared deviation from the full-information reservation value $v+\omega$. It need not maximize the seller's expected payoff from selling or retaining the asset. Appendix A.4.8 proves the calibration property, and Appendix A.4.10 considers the seller's optimal acceptance threshold.

overestimates the asset. More information reduces those errors and lowers the failure rate, with diminishing reductions as their information improves. When bids instead fall short of the full-information reserve on average, more information makes the seller less willing to accept, and failures increase. The proposition identifies the region in which reserve learning facilitates trade. Appendix A.4.9 gives the proof.

We next allow dealers to choose what the seller learns. The reserve is now $R = \widehat{v} + \omega + \xi_s$, where the seller's private shock ξ_s is independent of dealer information, idiosyncratic values, and randomization. Dealers know its cdf F_ξ but do not observe the realization before bidding. They choose whether to disclose before observing their independent idiosyncratic values, as in the baseline. This extension permits nonparticipation with payoff zero, in addition to bids at any real value.[28]

A one-dealer trading problem gives a bound used below. Let w be the dealer's expected valuation relative to the seller's common-value estimate, and let z be their bid relative to that estimate. Define their optimized payoff and its expectation after disclosure by

$$\phi_\xi(w) = \sup_{z \in \mathbb{R}} F_\xi(z - \omega)(w - z)^+, \qquad \pi_T = \mathbb{E}_u[\phi_\xi(u)]. \tag{33}$$

Here $x^+ = \max\{x, 0\}$. The cdf gives the probability of acceptance, and the positive part allows the dealer to forgo a negative-margin trade. The seller-shock conditions ensure $0 < \pi_T < \infty$.

With one potential dealer, disclosure makes their estimate $\mu_s = \mathbb{E}[v \mid s]$ public. Their expected gross payoff from talking is then the constant π_T. After silence, the seller uses their posterior mean m, while the dealer retains μ_s. The dealer's gross payoff is $\mathbb{E}_u[\phi_\xi(\mu_s - m + u)]$. An optimistic dealer benefits from a seller estimate below their own. A pessimistic dealer can benefit from correcting an estimate above their own.

For $K > \pi_T$, a one-dealer disclosure equilibrium exists. Within each equilibrium, the gain from talking strictly decreases with the signal and crosses zero once, at s^* satisfying $\mathbb{E}[v \mid s^*] = m$. Below this cutoff, dealers disclose when the gain exceeds their cost. Above it, they remain silent. Talking remains partial in the lower tail, where its probability converges to $\pi_T/K < 1$. The benefit from correcting the reserve is bounded, so sufficiently high costs deter disclosure even at very unfavorable signals. Appendix A.4.11 constructs this equilibrium.

With competing dealers, disclosure informs rivals as well as the seller. We state the result for the following continuation class. Let $\mathcal{G}_i$ denote dealer i's information before bidding, including the public history and their private information. When randomization is used, order nonparticipation below every submitted bid.

Assumption 1 (Reserve continuation)**.** *At every history used to compare disclosure choices, dealers choose globally optimal actions with nonparticipation available. The conditional distribution of*

[28] Maintain a continuous, full-support F_ξ with $\mathbb{E}[\xi_s^-] < \infty$, where $\xi_s^- = \max\{-\xi_s, 0\}$. At every history used to compare disclosure choices, including an off-path exact disclosure, beliefs use the Gaussian signal model, the disclosed values, and the remaining dealers' silence probabilities. A silent dealer also conditions on their own signal. Appendix A.4.12 gives a separate result when the seller shock has a finite lower endpoint with an atom.

a silent dealer's action is nondecreasing in the first-order stochastic order in their signal and idiosyncratic value. A talker's action depends on their remaining private information only through their idiosyncratic value and independent private randomization. Every submitted bid satisfies $b_i \leq \mathbb{E}[v \mid \mathcal{G}_i] + u_i$*, including bids with zero probability of an accepted win. Independent random priorities resolve ties.*

The assumption permits randomized bidding and atoms. For a bid with positive probability of an accepted win, the cap follows from global optimality and the winner's-curse comparison. Extending the cap to bids that never result in trade is a selection restriction. Let $\Pi^{T,R}(s)$ be the expected optimized gross payoff from disclosure before idiosyncratic values and rival disclosure decisions are realized, and let $p^*(s)$ be the equilibrium disclosure probability. We use s_- and s_+ for lower and upper signal thresholds and π_0 for a positive lower bound on talking payoffs in the lower tail.

Proposition 6 (Disclosure with competing dealers)**.** *Maintain the Gaussian primitives, the disclosure-cost and seller-shock conditions above, and* $K > \pi_T$*. For every finite* $N \geq 2$*, every disclosure equilibrium satisfying Assumption 1 obeys*

$$\Pi^{T,R}(s) \leq \pi_T, \qquad 0 \leq p^*(s) \leq \frac{\pi_T}{K} < 1. \tag{34}$$

There exists a finite s_+ *such that* $p^*(s) = 0$ *for almost every* $s \geq s_+$*. If* $N = 2$ *and* F_u *is non-degenerate, there exist* $\pi_0 > 0$ *and finite* $s_- < s_+$ *such that* $\Pi^{T,R}(s) \geq \pi_0$ *for almost every* $s \leq s_-$ *and*

$$p^*(s) \geq \frac{\pi_0}{2K} > 0 \qquad \text{for almost every } s \leq s_-. \tag{35}$$

The upper bound follows from the information conveyed by beating silent rivals. Winning is more likely when their signals are unfavorable, which lowers the disclosing dealer's common-value assessment conditional on winning. Their payoff cannot exceed the corresponding one-dealer value π_T. Because silence permits a nonnegative payoff, $K > \pi_T$ bounds disclosure probabilities below one. This bound is obtained before imposing a lower bound on the probability of silence.

At sufficiently favorable signals, retaining information becomes more valuable than disclosure. On histories in which rivals remain silent, an optimistic dealer can submit a bid that clears their bid caps and the seller's reserve with a payoff that grows with their signal. Disclosure eliminates that information advantage and yields a bounded payoff. The upper signal tail therefore contains only silence, up to null sets.

With two dealers, sufficiently unfavorable signals give the reverse comparison. After silence, the seller's reserve remains high relative to the dealer's assessment and their optimized payoff converges to zero. Disclosure aligns the seller's and the disclosing dealer's common-value estimates. A favorable idiosyncratic-value difference then lets the disclosing dealer strictly clear their rival's bid cap and earn a positive margin on an event with positive probability. The resulting talking payoff is bounded away from zero in the lower signal tail. Dealers with sufficiently low talking costs disclose there.

These conclusions allow atoms in bids. Independent priorities preserve the winner's-curse comparison at ties, and the lower-tail deviation strictly clears the rival's cap. Appendix A.4.12 supplies the proof and the bounded-lower-support extension. The proposition remains conditional on a disclosure equilibrium satisfying the stated continuation conditions.[29]

The two reserve propositions address different questions. Proposition 5 identifies when additional information reduces failures in a fixed bidding environment. Proposition 6 establishes incentives to supply unfavorable information in the stated equilibrium class. It does not impose a global monotone disclosure policy or order equilibrium failure rates by the realized number of talkers.

Conditional on public information and submitted bids, let B be the highest bid, with $B = -\infty$ if nobody submits one. Acceptance has probability $F_\xi(B - \widehat{v} - \omega)$. Disclosure can therefore improve execution even if it lowers B, provided it lowers the seller's estimate by more. Bad news can facilitate trade by reducing the price the seller requires.

Auctions with different talk counts also differ in the signals that induced disclosure and in the ensuing bids. The reserve-learning mechanism consequently supports an economic interpretation of the association between talk and execution in Table B.1, rather than a causal count comparison by itself. Fewer failures at a higher talk count require an ordering of the conditional distribution of $B - \widehat{v} - \omega$.[30]

6.2 Evidence on talk and execution

The reserve mechanism motivates reduced-form evidence on execution. We report within-list associations between talk intensity and market outcomes and emphasize at the outset that none of the estimates here or in Appendix B are causal. When sellers distribute lists containing multiple securities, dealers who choose to respond typically provide talks on many or all CUSIPs, reflecting administrative convenience and attention allocation rather than purely security-specific information. This institutional feature generates cross-CUSIP dependence in talk provision within a list.

For an outcome indicator O_i equal to one when the trade fails or, separately, when the cover is disclosed, we estimate the within-list linear probability model

$$O_i \;=\; \alpha_t + \alpha_{\ell(i)} + \lambda\, N_i + X_i'\theta + \epsilon_i, \tag{36}$$

where $\ell(i)$ is the BWIC list, N_i is the number of talks, and X_i collects the security controls of Table B.1. An additional talker is associated with a 6.0 percentage point lower probability of trade failure and a 3.3 percentage point higher probability that the cover is disclosed (Table B.1, OLS

[29] The appendix constructs a two-talker continuation under separate smoothness conditions and an unfloored one-talker continuation when the seller shock has an atom at a finite lower endpoint. General one-talker existence with full lower support, all-silent continuation existence, measurable selection across histories, and the competing-dealer disclosure fixed point remain to be established.

[30] Appendix A.4.13 gives the exact failure identity conditional on talk counts and dealer-side controls. First-order stochastic dominance of $B - \widehat{v} - \omega$ at a higher count is sufficient for fewer failures. The disclosure-incentive proposition does not establish that dominance.

columns). These associations plausibly overstate any effect of talk. CUSIPs that are ex ante easier to trade may both attract more talks and exhibit lower failure rates. The cover-disclosure association suggests that pre-trade information aggregation and post-trade transparency are complements. When more information is revealed before the auction, sellers are more likely to disclose outcomes afterward.

Appendix B reports a leave-one-out reweighting of the same within-list variation, in which N_i is projected on the average number of talks on the other CUSIPs in the same list. Because list fixed effects absorb the list total and list size, the leave-one-out measure is an exact within-list transformation of N_i itself. The estimator therefore delivers a weighted within-list association, not an instrumental-variables design, and we present it in the appendix to avoid inviting an IV interpretation. The reweighted estimates are much smaller in magnitude than OLS (-0.8 percentage points for trade failure and 2.8 percentage points for cover disclosure per additional talker), reinforcing the caution that the OLS associations should not be read causally. An exclusion-type concern applies to any list-based design. Dealers who talk on other securities in the same list may also be more likely to bid on CUSIP i, creating a direct competition channel.

The three-region evidence in Section 4.3 indicates that such cross-CUSIP participation is empirically relevant since dealers who are active on other securities on a list demonstrably receive the list and are present on CUSIP i, and may bid on it even when silent (bid submission itself is unobserved). While list fixed effects absorb the overall level of dealer participation at the list level, they cannot fully eliminate this within-list reallocation of competitive pressure.

The evidence is consistent with price talk enhancing market quality through fewer trade failures and greater transparency, particularly in settings with greater informational frictions, though the estimates do not establish this causally. At the same time, the list-based design highlights an important feature of OTC markets, which is that dealer attention and participation are jointly determined across securities, implying that information provision and competition are inherently interconnected.

7 Policy Analysis

The evidence establishes that voluntary price talk is informative but selectively incomplete. The natural policy question is whether disclosure in this market should be mandated instead. The question has direct counterparts in current regulatory discussions of leveraged-loan and CLO transparency (IOSCO, 2024). This section compares three disclosure regimes in a numerical counterfactual built on the model. Under the voluntary regime, dealers talk selectively and the auction mixes talkers and silent dealers as in Section 3. Under mandated disclosure, every dealer reveals their signal before bidding. Under banned disclosure, no dealer talks. The comparison quantifies what mandating transparency would change, by how much, and for whom.

7.1 Design of the counterfactual

The counterfactual uses the numerical configuration $\mu_0 = 0$, $\sigma_v = 1$, $\sigma_\varepsilon = 0.8$, $N = 8$, and idiosyncratic values uniform on $[0, \bar{u}]$ with $\bar{u} = 0.02$. We simulate 40,000 auctions. The same draws of the common value, the signals, and the idiosyncratic values are used in every regime, so regime differences are not driven by simulation noise and the common value cancels from surplus differences.

The three regimes are implemented as follows. Under mandated disclosure, all N signals are public, the posterior mean of v is common knowledge, and the remaining private heterogeneity is the idiosyncratic value. The continuation is a symmetric independent-private-value first-price auction, the equilibrium bid is the public posterior mean plus $\frac{N-1}{N}u_j$, and the dealer with the highest u_j wins. Under banned disclosure, the continuation is a symmetric common-value first-price auction. We solve the bidding schedule by best-response iteration until it reproduces itself, verify that dealer rents are nonnegative, and note that the dealer with the highest signal wins. The voluntary regime uses the hard-truncation illustration of the disclosure stage, in which dealers with signals below the cutoff talk. It is an illustration rather than an equilibrium special case of the baseline, which never assigns probability one to talking at any signal (Appendix A.1). Talkers bid $\beta^T(u)$ and silent dealers bid the benchmark schedule $b^{S,\mathrm{bm}}$. The silent-side schedule is validated for bid orderings and win shares rather than for absolute bid levels. We therefore report allocation-based statistics for the voluntary regime and base the revenue accounting on the comparison between the mandated and banned regimes, whose bid levels are exact and certified, with the voluntary regime lying between these two benchmarks in its disclosure content.

The accounting per auction is as follows. The seller's revenue is the transaction price P. The winning dealer w earns $v + u_w - P$ and losing dealers earn zero. Total surplus is $v + u_w$. Because the value is predominantly common, the efficient allocation assigns the asset to the dealer with the highest idiosyncratic value, and the allocative loss is

$$\mathcal{L} \;=\; \max_j u_j - u_w, \tag{37}$$

which is zero exactly when the allocation is efficient. By construction $\mathcal{L} \le \bar{u}$, so the entire efficiency stake of disclosure policy is of the order of the idiosyncratic value, while revenue differences can be of the order of the dispersion of common-value information.

7.2 Results

Table 9 reports the regime comparison. Three findings stand out.

First, the efficiency stakes are small in every comparison. Mandating disclosure perfects the allocation, raising the probability that the highest-value dealer wins from 0.27 to 1, yet the associated surplus gain is 0.0069 relative to the voluntary regime and 0.0078 relative to the banned

regime. Both gains are of the order of $\bar{u} = 0.02$, as equation (37) requires. Selective voluntary talk also answers the allocation question left open in Section 5. It lowers the expected allocative loss relative to banning talk, from 0.0078 to 0.0069, because talkers have made their signals public and therefore sort on idiosyncratic value, and it raises the loss relative to mandating, which removes informational sorting entirely.

Second, the revenue effect of full disclosure is larger than the efficiency effect. In the comparison between the mandated and banned regimes, where both bid levels are exact, mandating raises seller revenue by 0.0735 per auction, roughly 9% of the dispersion $\sigma = 0.78$ of conditional-mean values, while total surplus rises by only 0.0078. Dealer surplus falls by 0.0657, and the revenue gain is about nine times the efficiency gain. The rent that mandated disclosure competes away is the first-price bid shading of silent dealers, which is largest for silent winners with the most favorable signals. This transfer is the linkage principle of Milgrom and Weber (1982) operating at the regime level. Public revelation of affiliated information ties the price to information beyond the winner's own, and Section 5.1 documents the same linkage within the selective regime, where informed covers tie the price to rival signals. In this market, a transparency mandate therefore mainly transfers surplus to sellers and does little for efficiency.

Third, the price-informativeness column requires care. It measures the information content of the transaction price alone. Under banned disclosure the price tracks the common value closely because every bid is a monotone function of the bidder's signal. Under mandated disclosure the price reveals the full-information posterior. Under voluntary talk the transaction price alone is least informative, because the winner is often a dealer whose bid moves with idiosyncratic value rather than with news about the common value. The column nevertheless understates public information in the voluntary regime, which also publishes the talks themselves. Empirically, the consensus of disclosed talks explains most of the variation in benchmark-adjusted prices (Section 2.3). The voluntary regime concentrates price discovery in the disclosure stage rather than in the transaction price.

7.3 Margins outside the counterfactual

The counterfactual holds fixed several margins that all cut against mandated disclosure. Signal precision is exogenous in the model. Mandating disclosure competes away the informational rents that pay for producing signals in the first place, so endogenous information acquisition would move the full-disclosure outcome toward less informed markets (Grossman and Stiglitz, 1980). Participation is also fixed. Dealers who lose their rents under a mandate may stop attending or stop investing in coverage of hard-to-value securities. Credibility is a further concern. Voluntary talks are attributed, archived, and disciplined by the reputational mechanism described in Section 2, and a dealer who prefers silence has no report to slant. A mandate forces reports from exactly those dealers who prefer not to reveal, and no comparable mechanism disciplines the content of forced

reports (Section 5.3). Finally, the baseline model conditions on execution. The reserve extension in Section 6 and the association between talk and lower failure rates (Table B.1) suggest that disclosure also affects whether trade occurs at all, a margin the accounting above does not capture.

The comparison also clarifies what banning talk would forfeit. The banned regime loses little measured surplus, but it eliminates the public valuation stage entirely, moves allocations further from idiosyncratic values, and removes the pre-trade signal that is empirically associated with execution. The counterfactual therefore supports only a narrow case for intervention. Mandated pre-trade disclosure would improve allocations and raise seller revenue relative to a no-talk benchmark, but the efficiency gain is smaller than the redistribution, and the information-production margins outside the model weigh against a mandate.

8 Conclusion

This paper identifies a private source of public valuations and a limit to voluntary transparency. In the CLO BWIC market, dealers with the most favorable private information stay silent because disclosure would dissipate the informational advantage their signals provide in bidding. Those with less favorable information disclose because their chance of winning on private information alone is low and talking raises it. The result is a selective disclosure equilibrium in which publicly shared information systematically underrepresents the most optimistic assessments of value.

The crossing evidence supports this mechanism. Transaction prices deviate further above the pre-auction consensus when silent dealers win, a silence premium that increases monotonically in tranche risk. Within-dealer estimates show that the pattern is not explained solely by persistent, time-invariant dealer heterogeneity. List-level evidence is consistent with an intermediate silent region. Talkers win against identifiable present-and-silent dealers in 24% of cases, under the maintained assumption that present-and-silent dealers submit bids.

Supporting evidence indicates that price talks are informative despite being strategically selective. Consensus explains a large share of variation in benchmark-adjusted prices, winning dealers revise bids toward the emerging consensus, and bid dispersion is roughly half of talk dispersion, subject to the winner-selection and order-statistic caveats discussed in the text. Retained information is associated with price levels further above public benchmarks, without a correspondingly wider winning margin, because the cover bid adjusts upward alongside the winning bid.

These findings have implications for market design. Voluntary transparency can emerge absent regulatory mandates, but competitive incentives limit the completeness of voluntary disclosure. The dealers whose information would contribute most to the informational content of prices are precisely those with the strongest incentive to remain silent. This creates a wedge between the private optimality of silence and the informational value of disclosure that voluntary mechanisms do not close. The numerical comparison in Section 7 shows that mandated disclosure perfects allocations and raises seller revenue, but the efficiency gain is smaller than the transfer from dealers to sellers,

and the information-production margins outside the model weigh against a mandate. Neither full revelation nor collective silence is self-enforcing among competitors. Dealers with favorable signals gain by withholding them, and dealers with unfavorable signals gain by breaking any agreement to stay silent (Section 5.3). The observed partial talk is the noncooperative equilibrium between these poles. It produces an informative public consensus, and competition among silent dealers passes the value of their retained information to sellers. Where that boundary lies between what is revealed and what remains hidden therefore depends on market design as well as on dealers' incentives.

References

Asquith, Paul, Thomas R. Covert, and Parag A. Pathak (2013). *The Effects of Mandatory Transparency in Financial Market Design: Evidence from the Corporate Bond Market*. NBER Working Paper 19417. National Bureau of Economic Research. DOI: 10.3386/w19417.

Bergemann, Dirk, Tibor Heumann, Stephen Morris, Constantine Sorokin, and Eyal Winter (2022). "Optimal Information Disclosure in Classic Auctions". In: *American Economic Review: Insights* 4.3, pp. 371–388. DOI: 10.1257/aeri.20210504.

Bernheim, B. Douglas, Bezalel Peleg, and Michael D. Whinston (1987). "Coalition-Proof Nash Equilibria I. Concepts". In: *Journal of Economic Theory* 42.1, pp. 1–12. DOI: 10.1016/0022-0531(87)90099-8.

Bessembinder, Hendrik, William Maxwell, and Kumar Venkataraman (2006). "Market Transparency, Liquidity Externalities, and Institutional Trading Costs in Corporate Bonds". In: *Journal of Financial Economics* 82.2, pp. 251–288. DOI: 10.1016/j.jfineco.2005.10.002.

Board, Simon (2009). "Revealing Information in Auctions: The Allocation Effect". In: *Economic Theory* 38.1, pp. 125–135. DOI: 10.1007/s00199-006-0177-4.

Boehmer, Ekkehart, Gideon Saar, and Lei Yu (2005). "Lifting the Veil: An Analysis of Pre-trade Transparency at the NYSE". In: *Journal of Finance* 60.2, pp. 783–815. DOI: 10.1111/j.1540-6261.2005.00746.x.

Boyarchenko, Nina, David O. Lucca, and Laura Veldkamp (2021). "Taking Orders and Taking Notes: Dealer Information Sharing in Treasury Auctions". In: *Journal of Political Economy* 129.2, pp. 607–645. DOI: 10.1086/711954.

Campbell, Colin M. and Dan Levin (2000). "Can the Seller Benefit from an Insider in Common-Value Auctions?" In: *Journal of Economic Theory* 91.1, pp. 76–98.

Cordell, Larry, Michael R. Roberts, and Michael Schwert (2023). "CLO Performance". In: *Journal of Finance* 78.3, pp. 1235–1278. DOI: 10.1111/jofi.13224.

Di Maggio, Marco, Amir Kermani, and Zhaogang Song (2017). "The Value of Trading Relations in Turbulent Times". In: *Journal of Financial Economics* 124.2, pp. 266–284.

Duffie, Darrell, Piotr Dworczak, and Haoxiang Zhu (2017). "Benchmarks in Search Markets". In: *Journal of Finance* 72.5, pp. 1983–2044.

Duffie, Darrell, Nicolae Gârleanu, and Lasse Heje Pedersen (2005). "Over-the-Counter Markets". In: *Econometrica* 73.6, pp. 1815–1847. DOI: 10.1111/j.1468-0262.2005.00639.x.

Dufwenberg, Martin and Uri Gneezy (2002). "Information Disclosure in Auctions: An Experiment". In: *Journal of Economic Behavior & Organization* 48.4, pp. 431–444. DOI: 10.1016/S0167-2681(01)00235-9.

Dye, Ronald A. (1985). "Disclosure of Nonproprietary Information". In: *Journal of Accounting Research* 23.1, pp. 123–145.

Edwards, Amy K., Lawrence E. Harris, and Michael S. Piwowar (2007). "Corporate Bond Market Transaction Costs and Transparency". In: *Journal of Finance* 62.3, pp. 1421–1451.

Engelbrecht-Wiggans, Richard, Paul R. Milgrom, and Robert J. Weber (1983). "Competitive Bidding and Proprietary Information". In: *Journal of Mathematical Economics* 11.2, pp. 161–169.

Fudenberg, Drew and Eric Maskin (1986). "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information". In: *Econometrica* 54.3, pp. 533–554. DOI: 10.2307/1911307.

Ganuza, Juan-José (2006). "Information Revelation in Auctions". In: *Journal of Economic Theory* 127.1, pp. 111–130. DOI: 10.1016/j.jet.2004.12.002.

Ganuza, Juan-Jose and Jose Penalva (2010). "Signal Ordering Based on Dispersion and the Supply of Private Information in Auctions". In: *Econometrica* 78.3, pp. 1007–1030.

Glode, Vincent, Christian C. Opp, and Xingtan Zhang (2018). "Voluntary Disclosure in Bilateral Transactions". In: *Journal of Economic Theory* 175, pp. 652–688. DOI: 10.1016/j.jet.2018.02.005.

Goldstein, Itay, Yan Xiong, and Liyan Yang (2025). "Information Sharing in Financial Markets". In: *Journal of Financial Economics* 163, p. 103967. DOI: 10.1016/j.jfineco.2024.103967.

Goldstein, Michael A., Edith S. Hotchkiss, and Erik R. Sirri (2007). "Transparency and Liquidity: A Controlled Experiment on Corporate Bonds". In: *Review of Financial Studies* 20.2, pp. 235–273. DOI: 10.1093/rfs/hhl020.

Grossman, Sanford J. (1981). "The Informational Role of Warranties and Private Disclosure about Product Quality". In: *Journal of Law and Economics* 24.3, pp. 461–483.

Grossman, Sanford J. and Joseph E. Stiglitz (1980). "On the Impossibility of Informationally Efficient Markets". In: *American Economic Review* 70.3, pp. 393–408.

Harrington, Joseph E. and Andrzej Skrzypacz (2011). "Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices". In: *American Economic Review* 101.6, pp. 2425–2449.

Hendershott, Terrence and Ananth Madhavan (2015). "Click or Call? Auction versus Search in the Over-the-Counter Market". In: *Journal of Finance* 70.1, pp. 419–447. DOI: 10.1111/jofi.12164.

Hendricks, Kenneth and Robert H. Porter (1988). "An Empirical Study of an Auction with Asymmetric Information". In: *American Economic Review* 78.5, pp. 865–883.

Hendricks, Kenneth, Robert Porter, and Guofu Tan (2008). "Bidding Rings and the Winner's Curse". In: *RAND Journal of Economics* 39.4, pp. 1018–1041.

Hollifield, Burton, Artem Neklyudov, and Chester Spatt (2017). "Bid-Ask Spreads, Trading Networks, and the Pricing of Securitizations". In: *Review of Financial Studies* 30.9, pp. 3048–3085.

Hugonnier, Julien, Benjamin Lester, and Pierre-Olivier Weill (2020). "Frictional Intermediation in Over-the-Counter Markets". In: *Review of Economic Studies* 87.2, pp. 714–752.

IOSCO (June 2024). *Leveraged Loans and CLOs Good Practices for Consideration*. Final Report FR03/2024. International Organization of Securities Commissions.

Kagel, John H. and Dan Levin (2016). "Auctions: A Survey of Experimental Research, 1995–2010". In: *The Handbook of Experimental Economics, Volume 2*. Ed. by John H. Kagel and Alvin E. Roth. Princeton University Press, pp. 563–637.

Lagos, Ricardo and Guillaume Rocheteau (2007). "Search in Asset Markets: Market Structure, Liquidity, and Welfare". In: *American Economic Review* 97.2, pp. 198–202.

Li, Dan and Norman Schürhoff (2019). "Dealer Networks". In: *Journal of Finance* 74.1, pp. 91–144. DOI: 10.1111/jofi.12728.

Li, Jing, Tingjun Liu, and Ran Zhao (2022). "Strategic Nondisclosure in Takeovers". In: *The Accounting Review* 97.4, pp. 345–370. DOI: 10.2308/TAR-2020-0114.

Madhavan, Ananth (1995). "Consolidation, Fragmentation, and the Disclosure of Trading Information". In: *Review of Financial Studies* 8.3, pp. 579–603. DOI: 10.1093/rfs/8.3.579.

Matthews, Steven and Andrew Postlewaite (1985). "Quality Testing and Disclosure". In: *RAND Journal of Economics* 16.3, pp. 328–340. DOI: 10.2307/2555561.

McAfee, R. Preston and John McMillan (1992). "Bidding Rings". In: *American Economic Review* 82.3, pp. 579–599.

Milgrom, Paul R. (1981). "Good News and Bad News: Representation Theorems and Applications". In: *Bell Journal of Economics* 12.2, pp. 380–391.

Milgrom, Paul and Robert J. Weber (1982). "A Theory of Auctions and Competitive Bidding". In: *Econometrica* 50.5, pp. 1089–1122. DOI: 10.2307/1911865.

Pagano, Marco and Ailsa Röell (1996). "Transparency and Liquidity: A Comparison of Auction and Dealer Markets with Informed Trading". In: *Journal of Finance* 51.2, pp. 579–611. DOI: 10.1111/j.1540-6261.1996.tb02695.x.

Persico, Nicola (2000). "Information Acquisition in Auctions". In: *Econometrica* 68.1, pp. 135–148.

Tan, Xu (2016). "Information Revelation in Auctions with Common and Private Values". In: *Games and Economic Behavior* 97, pp. 147–165. DOI: 10.1016/j.geb.2016.04.006.

Verrecchia, Robert E. (1983). "Discretionary Disclosure". In: *Journal of Accounting and Economics* 5, pp. 179–194. DOI: 10.1016/0165-4101(83)90011-3.

Appendix

Appendix A Proofs

A.1 Information and beliefs after disclosure

This subsection derives the beliefs used in the continuation auction. We distinguish the information contained in the values of the disclosed talks from the information contained in the decision to remain silent. The baseline uses the uniform talking-cost distribution in Section 3 together with $\bar{u} < K$. The hard-truncation specification used for several auxiliary calculations below is a separate benchmark in which silence identifies an upper set of signals. It is not an equilibrium special case of the baseline under $\bar{u} < K$, because the baseline never assigns probability one to talking at any signal.

Let $\mathcal{T}$ and $\mathcal{S}$ denote the realized sets of talkers and silent dealers, with $n_T = |\mathcal{T}|$ and $n_S = |\mathcal{S}|$. Let $\mathcal{M} = \{m_j : j \in \mathcal{T}\}$ denote the disclosed signals and define the full public information after disclosure by

$$\mathcal{I} \equiv (\mathcal{M}, \mathcal{T}, \mathcal{S}). \tag{A.1}$$

Because dealers are ex ante symmetric and their signals are exchangeable conditional on v, dealer identities carry no payoff-relevant information beyond the realized partition and the disclosed signal values. Identity disclosure matters institutionally because the talks are attributed, which supports the maintained truthfulness assumption in Section 3.

Let $p_T(s)$ denote the equilibrium probability that a dealer with signal s talks, and define the corresponding silence weight

$$w(s) \equiv 1 - p_T(s). \tag{A.2}$$

Conditional on a signal s, the disclosure decision depends on the independent talking-cost draw and therefore contains no additional information about v. Conditional on v, however, the probability of silence is

$$q(v) \equiv \Pr(T_j = \text{Silent} \mid v) = \int_{\mathbb{R}} w(t) f_\varepsilon(t - v)\, dt, \tag{A.3}$$

where $f_\varepsilon(x) = \frac{1}{\sigma_\varepsilon}\phi\left(\frac{x}{\sigma_\varepsilon}\right)$ is the density of the signal error.

For $n_T \geq 1$, write $\bar{S} = n_T^{-1} \sum_{j \in \mathcal{T}} m_j$. When $n_T = 0$, every occurrence of $n_T \bar{S}$ means the empty sum, equal to zero. We first record the Gaussian posterior obtained from the disclosed signal values alone.

Lemma A1 (Talk-based posteriors)**.** *Conditional on the disclosed signals $\mathcal{M}$, the common value satisfies*

$$v \mid \mathcal{M} \sim \mathcal{N}\left(\mu_T, \sigma^2_{v|\mathcal{M}}\right), \quad \mu_T = \frac{\tau_v \mu_0 + n_T \tau_\varepsilon \bar{S}}{\tau_v + n_T \tau_\varepsilon}, \quad \sigma^2_{v|\mathcal{M}} = \frac{1}{\tau_v + n_T \tau_\varepsilon}. \tag{A.4}$$

Equivalently, $\mu_T = \alpha \mu_0 + \beta \bar{S}$, where

$$\alpha = \frac{\tau_v}{\tau_v + n_T \tau_\varepsilon}, \quad \beta = \frac{n_T \tau_\varepsilon}{\tau_v + n_T \tau_\varepsilon}, \quad \alpha + \beta = 1.$$

Conditional on $\mathcal{M}$ and an additional private signal s, the common value satisfies

$$v \mid (\mathcal{M}, s) \sim \mathcal{N}\big(\mu_S(s,\mathcal{M}), \sigma_S^2(\mathcal{M})\big), \quad \mu_S(s,\mathcal{M}) = \frac{\tau_v\mu_0 + n_T\tau_\varepsilon\bar{S} + \tau_\varepsilon s}{\tau_v + (n_T+1)\tau_\varepsilon}, \tag{A.5}$$

with

$$\frac{1}{\sigma_S^2(\mathcal{M})} = \tau_v + (n_T+1)\tau_\varepsilon. \tag{A.6}$$

Writing

$$\mu_S(s,\mathcal{M}) = \alpha'\mu_0 + \beta'\bar{S} + \gamma' s,$$

the coefficients are

$$\alpha' = \frac{\tau_v}{\tau_v + (n_T+1)\tau_\varepsilon}, \quad \beta' = \frac{n_T\tau_\varepsilon}{\tau_v + (n_T+1)\tau_\varepsilon}, \quad \gamma' = \frac{\tau_\varepsilon}{\tau_v + (n_T+1)\tau_\varepsilon}.$$

Finally, before conditioning on a dealer's disclosure decision, an unrevealed signal satisfies

$$s_k \mid \mathcal{M} \sim \mathcal{N}(\mu_T, \sigma_s^2), \quad \sigma_s^2 = \sigma_{v|\mathcal{M}}^2 + \sigma_\varepsilon^2.$$

Proof. The prior density of v is $f_v(v) \propto \exp\{-\frac{\tau_v}{2}(v-\mu_0)^2\}$. Conditional on v, a disclosed signal s_k has density $f_\varepsilon(s_k - v) \propto \exp\{-\frac{\tau_\varepsilon}{2}(s_k - v)^2\}$. Hence, conditional on the n_T disclosed signals,

$$f(v \mid \mathcal{M}) \propto \exp\left\{-\frac{1}{2}\left[\tau_v(v-\mu_0)^2 + \tau_\varepsilon\sum_{k\in\mathcal{T}}(s_k - v)^2\right]\right\}.$$

Expanding the two quadratic terms gives $\tau_v(v-\mu_0)^2 = \tau_v v^2 - 2\tau_v\mu_0 v + \tau_v\mu_0^2$ and

$$\tau_\varepsilon\sum_{k\in\mathcal{T}}(s_k - v)^2 = \tau_\varepsilon\sum_{k\in\mathcal{T}} s_k^2 - 2\tau_\varepsilon v\sum_{k\in\mathcal{T}} s_k + n_T\tau_\varepsilon v^2.$$

Collecting all terms that depend on v yields

$$-\frac{1}{2}\left[(\tau_v + n_T\tau_\varepsilon)v^2 - 2v\left(\tau_v\mu_0 + \tau_\varepsilon\sum_{k\in\mathcal{T}} s_k\right)\right].$$

Since $\sum_{k\in\mathcal{T}} s_k = n_T\bar{S}$, completing the square gives posterior precision $\tau_v + n_T\tau_\varepsilon$ and posterior mean

$$\frac{\tau_v\mu_0 + n_T\tau_\varepsilon\bar{S}}{\tau_v + n_T\tau_\varepsilon}.$$

This proves (A.4).

A dealer who additionally observes their own signal s contributes one further likelihood factor

$$\exp\left\{-\frac{\tau_\varepsilon}{2}(s-v)^2\right\}.$$

Repeating the preceding calculation therefore changes the posterior precision from $\tau_v + n_T\tau_\varepsilon$ to $\tau_v + (n_T+1)\tau_\varepsilon$ and changes the linear term in v from $\tau_v\mu_0 + n_T\tau_\varepsilon\bar{S}$ to $\tau_v\mu_0 + n_T\tau_\varepsilon\bar{S} + \tau_\varepsilon s$. Equations (A.5) and (A.6) follow.

For $n_T \geq 1$, the relative weights on the consensus satisfy

$$\beta' = \frac{n_T \tau_\varepsilon}{\tau_v + (n_T + 1)\tau_\varepsilon} < \frac{n_T \tau_\varepsilon}{\tau_v + n_T \tau_\varepsilon} = \beta < 1.$$

When $n_T = 0$, both consensus coefficients are zero. Treating n_T temporarily as continuous,

$$\frac{\partial \beta}{\partial n_T} = \frac{\tau_v \tau_\varepsilon}{(\tau_v + n_T \tau_\varepsilon)^2} > 0 \text{ and } \frac{\partial \beta'}{\partial n_T} = \frac{\tau_\varepsilon (\tau_v + \tau_\varepsilon)}{[\tau_v + (n_T + 1)\tau_\varepsilon]^2} > 0.$$

Both β and β' converge to one as $n_T \to \infty$.

Finally, $s_k = v + \varepsilon_k$. Conditional on $\mathcal{M}$, the variables v and ε_k remain independent because ε_k is independent of the prior and of the disclosed dealers' signal errors. Therefore

$$\mathbb{E}[s_k \mid \mathcal{M}] = \mathbb{E}[v \mid \mathcal{M}] + \mathbb{E}[\varepsilon_k] = \mu_T$$

and

$$\mathrm{Var}(s_k \mid \mathcal{M}) = \mathrm{Var}(v \mid \mathcal{M}) + \mathrm{Var}(\varepsilon_k) = \sigma^2_{v|\mathcal{M}} + \sigma^2_\varepsilon.$$

Normality follows because $v \mid \mathcal{M}$ and ε_k are independent normal random variables. □

The preceding lemma deliberately ignores the information contained in the disclosure decisions. We now derive the full equilibrium posterior.

Lemma A2 (Posterior after selective disclosure)**.** *Fix a symmetric talk-propensity function p_T and its silence weight $w = 1 - p_T$. Conditional on the realized public information $\mathcal{I} = (\mathcal{M}, \mathcal{T}, \mathcal{S})$,*

$$f(v \mid \mathcal{I}) = \frac{f(v \mid \mathcal{M}) q(v)^{n_S}}{\int_{\mathbb{R}} f(\tilde{v} \mid \mathcal{M}) q(\tilde{v})^{n_S} \, d\tilde{v}}. \tag{A.7}$$

If w is nondecreasing, then $q(v)$ is nondecreasing in v. If w is nonconstant on a set reached with positive probability by the signal distribution, then $q(v)$ is strictly increasing. Consequently, for $n_S > 0$, $v \mid \mathcal{I} \succeq_{\mathrm{MLR}} v \mid \mathcal{M}$, with strict ordering under the preceding nondegeneracy condition.

For a silent dealer j with exact private signal $s_j = s$,

$$f(v \mid s, \mathcal{I}) = \frac{f(v \mid s, \mathcal{M}) q(v)^{n_S - 1}}{\int_{\mathbb{R}} f(\tilde{v} \mid s, \mathcal{M}) q(\tilde{v})^{n_S - 1} \, d\tilde{v}}. \tag{A.8}$$

Their own silence event provides no information about v beyond the exact signal s, because its probability conditional on s is $w(s)$, which is independent of v.

Proof. Conditional on v, dealer j receives signal s_j with density $f_\varepsilon(s_j - v)$. If they talk and reveals $m_j = s_j$, the joint likelihood contribution of observing both the talk decision and the disclosed value m_j is $p_T(m_j) f_\varepsilon(m_j - v)$. The factor $p_T(m_j)$ depends on the observed signal but not on v. It therefore affects the normalizing constant of the posterior but not its shape as a function of v.

For a dealer who remains silent, the signal is unobserved. Integrating over all possible signal realizations gives the likelihood contribution

$$\int_{\mathbb{R}} w(t) f_\varepsilon(t - v) \, dt = q(v).$$

Conditional independence of signals and talking costs across dealers therefore gives

$$f(\mathcal{I} \mid v) = \prod_{j \in \mathcal{T}} \left[p_T(m_j) f_\varepsilon(m_j - v) \right] q(v)^{n_S}.$$

Bayes' rule implies

$$f(v \mid \mathcal{I}) \propto f_v(v) \prod_{j \in \mathcal{T}} p_T(m_j) f_\varepsilon(m_j - v) q(v)^{n_S}.$$

The factors $\prod_{j \in \mathcal{T}} p_T(m_j)$ are constant in v and can be absorbed into the normalizing constant. Hence

$$f(v \mid \mathcal{I}) \propto f_v(v) \prod_{j \in \mathcal{T}} f_\varepsilon(m_j - v) q(v)^{n_S}.$$

The first two factors are proportional to $f(v \mid \mathcal{M})$ by Lemma A1. Normalizing gives (A.7).

To establish monotonicity of q, write $q(v) = \mathbb{E}[w(v + \varepsilon)]$ where $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$. If $v' > v$ and w is nondecreasing, then $w(v' + \varepsilon) \geq w(v + \varepsilon)$ for every realization of ε. Taking expectations gives $q(v') \geq q(v)$. If w is nonconstant over a region crossed with positive probability when the signal is shifted from v to v', the inequality is strict.

The likelihood ratio of the full public posterior to the talk-only posterior is

$$\frac{f(v \mid \mathcal{I})}{f(v \mid \mathcal{M})} = C(\mathcal{I}) q(v)^{n_S},$$

where $C(\mathcal{I}) > 0$ does not depend on v. Because q is nondecreasing, this ratio is nondecreasing in v. This is precisely the monotone-likelihood-ratio ordering $v \mid \mathcal{I} \succeq_{\text{MLR}} v \mid \mathcal{M}$. The ordering is strict whenever $n_S > 0$ and q is strictly increasing.

Now fix a silent dealer j and condition additionally on their exact signal s. Their own signal contributes $f_\varepsilon(s - v)$ to the likelihood. Their own silence decision contributes the factor $w(s)$. Once s is observed, $w(s)$ is constant in v and therefore cancels from Bayes' rule. The remaining $n_S - 1$ silent dealers contribute $q(v)^{n_S - 1}$. Thus

$$f(v \mid s, \mathcal{I}) \propto f_v(v) \prod_{k \in \mathcal{T}} f_\varepsilon(m_k - v) f_\varepsilon(s - v) q(v)^{n_S - 1}.$$

The first three terms are proportional to $f(v \mid s, \mathcal{M})$, which gives (A.8). □

Lemma A2 gives the first formal statement of anti-unraveling. When the equilibrium talk propensity decreases with the signal, the silence weight w increases with the signal. A high common value shifts private signals upward and therefore makes silence more likely. Observing silence consequently shifts the posterior toward higher values of v.

We retain $\mu_T = \mathbb{E}[v \mid \mathcal{M}]$ and $\mu_S(s, \mathcal{M}) = \mathbb{E}[v \mid s, \mathcal{M}]$ as *talk-based location parameters*. They are Gaussian and affine in the observed signals. The equilibrium posteriors in (A.7) and (A.8) generally are not Gaussian. All payoff-relevant continuation objects below condition on the full public information $\mathcal{I}$.

Hard-truncation benchmark. For comparison, the auxiliary hard-truncation benchmark imposes

$$p_T(s) = \mathbf{1}\{s \leq \bar{s}\}, \quad w(s) = \mathbf{1}\{s > \bar{s}\}.$$

This benchmark is useful because it creates a finite lower support point for silent signals. Under the baseline restriction $\bar{u} < K$, it is not generated by the equilibrium cost-threshold policy. Equation (A.3) then becomes

$$q(v) = \int_{\bar{s}}^{\infty} f_\varepsilon(t - v)\, dt = 1 - \Phi\left(\frac{\bar{s} - v}{\sigma_\varepsilon}\right) = \Phi\left(\frac{v - \bar{s}}{\sigma_\varepsilon}\right). \tag{A.9}$$

Thus the hard-truncation posterior is obtained directly from (A.7). We use this benchmark below only where it permits an analytical characterization that is not available in closed form for a general silence weight.

The next result separates the information contained in the values of the talks from the information contained in the incidence of disclosure.

Proposition A1 (Content and incidence of disclosure)**.** *Fix a symmetric nonincreasing talk-propensity function p_T.*

(i) Fix (n_T, n_S) and two talk vectors $\mathcal{M}$ and $\mathcal{M}'$ containing the same number of disclosed signals. If $\sum_{j\in\mathcal{T}} m'_j > \sum_{j\in\mathcal{T}} m_j$, then $v \mid (\mathcal{M}', n_T, n_S) \succeq_{\text{MLR}} v \mid (\mathcal{M}, n_T, n_S)$.

(ii) Before talk contents are observed, let n_T denote the number of talkers among the N dealers. If p_T is nonconstant, then for each $k = 0, \ldots, N-1$, $v \mid n_T = k \succeq_{\text{MLR}} v \mid n_T = k+1$.

(iii) In the hard-truncation benchmark with cutoff $\bar{s}$, fix the information supplied by all other dealers. For any disclosed signal $m \le \bar{s}$, $v \mid \{j \text{ silent}\} \succeq_{\text{MLR}} v \mid \{m_j = m\}$.

Proof. For part (i), conditional on the disclosure counts, Lemma A2 implies

$$f(v \mid \mathcal{M}, n_T, n_S) \propto f_v(v) \prod_{j\in\mathcal{T}} f_\varepsilon(m_j - v) q(v)^{n_S}.$$

For the normal signal density,

$$\prod_{j\in\mathcal{T}} f_\varepsilon(m_j - v) \propto \exp\left\{-\frac{\tau_\varepsilon}{2}\sum_{j\in\mathcal{T}}(m_j - v)^2\right\}.$$

Expanding the sum gives

$$\sum_{j\in\mathcal{T}}(m_j - v)^2 = \sum_{j\in\mathcal{T}} m_j^2 - 2v\sum_{j\in\mathcal{T}} m_j + n_T v^2.$$

Hence

$$f(v \mid \mathcal{M}, n_T, n_S) \propto f_v(v) \exp\left\{\tau_\varepsilon v \sum_{j\in\mathcal{T}} m_j - \frac{n_T \tau_\varepsilon}{2} v^2\right\} q(v)^{n_S},$$

where all terms independent of v have been absorbed into the normalizing constant.

Taking the ratio of the posterior under $\mathcal{M}'$ to the posterior under $\mathcal{M}$ gives

$$\frac{f(v \mid \mathcal{M}', n_T, n_S)}{f(v \mid \mathcal{M}, n_T, n_S)} = C \exp\left\{\tau_\varepsilon v \left(\sum_j m'_j - \sum_j m_j\right)\right\},$$

where $C > 0$ is independent of v. The exponent has a strictly positive slope in v under the stated inequality. The ratio is therefore strictly increasing in v, which proves the MLR ordering.

For part (ii), define the probability that a dealer talks conditional on the common value by

$$\pi(v) \equiv \Pr(T_j = \text{Talk} \mid v) = \int_{\mathbb{R}} p_T(t) f_\varepsilon(t - v)\, dt = \mathbb{E}[p_T(v + \varepsilon)].$$

Because p_T is nonincreasing, $v' > v$ implies $p_T(v' + \varepsilon) \leq p_T(v + \varepsilon)$ for every ε. Hence $\pi(v)$ is nonincreasing in v, and it is strictly decreasing when p_T is nonconstant on a region receiving positive signal probability.

Conditional on v, the dealers' disclosure decisions are independent because their signals and talking costs are independent conditional on v. Therefore

$$n_T \mid v \sim \text{Binomial}(N, \pi(v)) \Longrightarrow \Pr(n_T = k \mid v) = \binom{N}{k} \pi(v)^k [1 - \pi(v)]^{N-k}.$$

For adjacent counts,

$$\frac{\Pr(n_T = k \mid v)}{\Pr(n_T = k + 1 \mid v)} = \frac{\binom{N}{k}}{\binom{N}{k+1}} \frac{1 - \pi(v)}{\pi(v)} = \frac{k + 1}{N - k} \frac{1 - \pi(v)}{\pi(v)}.$$

Since $\pi(v)$ is decreasing, $(1 - \pi(v))/\pi(v)$ is increasing. The ratio is therefore increasing in v. Bayes' rule implies

$$\frac{f(v \mid n_T = k)}{f(v \mid n_T = k + 1)} = C_k \frac{\Pr(n_T = k \mid v)}{\Pr(n_T = k + 1 \mid v)},$$

where $C_k > 0$ is independent of v. Hence $v \mid n_T = k$ MLR-dominates $v \mid n_T = k + 1$.

For part (iii), in the hard-truncation benchmark, the likelihood contribution of silence is

$$q(v) = \Phi\left(\frac{v - \bar{s}}{\sigma_\varepsilon}\right),$$

while observing an exact talk m contributes $f_\varepsilon(m - v) = \frac{1}{\sigma_\varepsilon} \phi\left(\frac{m - v}{\sigma_\varepsilon}\right)$. It is enough to show that $q(v)/f_\varepsilon(m - v)$ is strictly increasing in v.

Define

$$z \equiv \frac{v - \bar{s}}{\sigma_\varepsilon}, \quad \lambda(z) \equiv \frac{\phi(z)}{\Phi(z)}.$$

Then

$$\frac{d}{dv} \log q(v) = \frac{1}{\sigma_\varepsilon} \lambda(z).$$

Also,

$$\log f_\varepsilon(m - v) = \text{const} - \frac{(m - v)^2}{2\sigma_\varepsilon^2},$$

so

$$\frac{d}{dv} \log f_\varepsilon(m - v) = \frac{m - v}{\sigma_\varepsilon^2}.$$

Therefore

$$\frac{d}{dv}\log\frac{q(v)}{f_\varepsilon(m-v)}=\frac{1}{\sigma_\varepsilon}\left[\lambda(z)+\frac{v-m}{\sigma_\varepsilon}\right].$$

Because any disclosed signal satisfies $m\le\bar{s}$,

$$\frac{v-m}{\sigma_\varepsilon}\ge\frac{v-\bar{s}}{\sigma_\varepsilon}=z.$$

Hence

$$\frac{d}{dv}\log\frac{q(v)}{f_\varepsilon(m-v)}\ge\frac{1}{\sigma_\varepsilon}[\lambda(z)+z].$$

It remains to show that $\lambda(z)+z>0$ for every z. Since $\Phi(z)>0$,

$$\lambda(z)+z>0\quad\Longleftrightarrow\quad\phi(z)+z\Phi(z)>0.$$

For $z\ge0$ this is immediate. For $z<0$, the normal Mills inequality gives $\Phi(z)<\frac{\phi(z)}{-z}$. Multiplying by $-z>0$ yields $-z\Phi(z)<\phi(z)$, or equivalently $\phi(z)+z\Phi(z)>0$. Thus, the likelihood ratio is strictly increasing in v, which establishes the claimed MLR ordering. □

Parts (i) and (ii) apply directly to the cost-threshold equilibrium. Part (iii) is deliberately stated only for hard truncation. With a general silence weight, silence need not dominate every individual disclosed signal in the MLR order. The robust implications are that a higher consensus is favorable news at fixed disclosure counts and that, before contents are observed, fewer talks are favorable news whenever the talk propensity decreases with the signal.

A.1.1 Order statistics for silent competitors

We next define the order-statistic objects used in the continuation auction. The derivation is given explicitly because silent dealers' signals are not independent after conditioning on the public information. They remain correlated through the latent common value.

For later use, define the weighted signal mass

$$Q(x,v)\equiv\int_{-\infty}^{x}w(t)f_\varepsilon(t-v)\,dt.\tag{A.10}$$

Since $q(v)=Q(\infty,v)$, the conditional distribution of a signal given v and the event that its dealer is silent is

$$F_S(x\mid v)\equiv\Pr(s_k\le x\mid v,T_k=\text{Silent})=\frac{Q(x,v)}{q(v)}.\tag{A.11}$$

The signal density exists almost everywhere. At those points,

$$f_S(x\mid v)=\frac{w(x)f_\varepsilon(x-v)}{q(v)}.$$

Fix a silent dealer j with private signal $s_j=s$. Let

$$S_{(1)}\equiv\max_{k\in\mathcal{S}\setminus\{j\}}s_k$$

denote the highest signal among their $n_S - 1$ silent competitors. The maximum of an empty set is $-\infty$. For $n_S = 1$, set $G \equiv 1$ and $g \equiv 0$. The density formula below applies to $n_S \geq 2$, almost everywhere in the cutoff, and pointwise wherever w is continuous. Define

$$G(x \mid s, \mathcal{I}) \equiv \Pr(S_{(1)} \leq x \mid s, \mathcal{I}), \quad g(x \mid s, \mathcal{I}) \equiv \frac{\partial}{\partial x} G(x \mid s, \mathcal{I}). \tag{A.12}$$

Conditional on $(v, \mathcal{I}, s)$, the $n_S - 1$ rival silent signals are independent draws from (A.11). Hence

$$\Pr(S_{(1)} \leq x \mid v, s, \mathcal{I}) = \left[\frac{Q(x, v)}{q(v)}\right]^{n_S - 1}.$$

Integrating over dealer j's posterior for v gives

$$G(x \mid s, \mathcal{I}) = \int_{\mathbb{R}} \left[\frac{Q(x, v)}{q(v)}\right]^{n_S - 1} f(v \mid s, \mathcal{I})\, dv.$$

Using (A.8),

$$f(v \mid s, \mathcal{I}) = \frac{f(v \mid s, \mathcal{M}) q(v)^{n_S - 1}}{\int f(\tilde{v} \mid s, \mathcal{M}) q(\tilde{v})^{n_S - 1}\, d\tilde{v}}.$$

The powers of $q(v)$ cancel. Therefore

$$G(x \mid s, \mathcal{I}) = \frac{\displaystyle\int_{\mathbb{R}} f(v \mid s, \mathcal{M}) Q(x, v)^{n_S - 1}\, dv}{\displaystyle\int_{\mathbb{R}} f(v \mid s, \mathcal{M}) q(v)^{n_S - 1}\, dv}. \tag{A.13}$$

Because

$$\frac{\partial Q(x, v)}{\partial x} = w(x) f_\varepsilon(x - v),$$

differentiating (A.13) under the integral sign gives

$$g(x \mid s, \mathcal{I}) = \frac{(n_S - 1) w(x) \displaystyle\int_{\mathbb{R}} f(v \mid s, \mathcal{M}) f_\varepsilon(x - v) Q(x, v)^{n_S - 2}\, dv}{\displaystyle\int_{\mathbb{R}} f(v \mid s, \mathcal{M}) q(v)^{n_S - 1}\, dv}. \tag{A.14}$$

The two arguments of G play different roles. The first, x, is the rival-signal threshold associated with a contemplated bid. The second, s, is the dealer's true private signal and changes their posterior over v, which in turn changes their assessment of the strength of rival silent dealers.

Define the corresponding reverse hazard rate by

$$K(x \mid s, \mathcal{I}) \equiv \frac{g(x \mid s, \mathcal{I})}{G(x \mid s, \mathcal{I})}, \tag{A.15}$$

whenever $G(x \mid s, \mathcal{I}) > 0$. When a scalar first-order condition uses this signal density, its local regularity condition is $g(s \mid s, \mathcal{I}) > 0$. The one-silent configuration uses its separate best-response

condition below.

For the talker side, let

$$S_{\mathcal{S}}^{\max} \equiv \max_{k \in \mathcal{S}} s_k$$

denote the highest signal among all n_S silent dealers. If $n_S = 0$, set $G_0 \equiv 1$ and $g_0 \equiv 0$. The density calculation applies to $n_S \geq 1$, with the same almost-everywhere convention as above. Define

$$G_0(x \mid \mathcal{I}) \equiv \Pr(S_{\mathcal{S}}^{\max} \leq x \mid \mathcal{I}), \quad g_0(x \mid \mathcal{I}) \equiv \frac{\partial}{\partial x} G_0(x \mid \mathcal{I}). \tag{A.16}$$

Conditional on v and $\mathcal{I}$,

$$\Pr(S_{\mathcal{S}}^{\max} \leq x \mid v, \mathcal{I}) = \left[\frac{Q(x,v)}{q(v)}\right]^{n_S}.$$

Using (A.7) and again cancelling the powers of $q(v)$ gives

$$G_0(x \mid \mathcal{I}) = \frac{\int_{\mathbb{R}} f(v \mid \mathcal{M}) Q(x,v)^{n_S}\, dv}{\int_{\mathbb{R}} f(v \mid \mathcal{M}) q(v)^{n_S}\, dv}. \tag{A.17}$$

Differentiating gives

$$g_0(x \mid \mathcal{I}) = \frac{n_S w(x) \int_{\mathbb{R}} f(v \mid \mathcal{M}) f_\varepsilon(x - v) Q(x,v)^{n_S - 1}\, dv}{\int_{\mathbb{R}} f(v \mid \mathcal{M}) q(v)^{n_S}\, dv}. \tag{A.18}$$

Equations (A.13)–(A.18) make explicit why the silent-rival order statistics are not ordinary independent-signal order statistics after public information is observed. Conditional on v, silent signals are independent. After v is integrated out, the common-value uncertainty induces dependence among them. A silent dealer's own signal changes the posterior distribution of v and therefore changes the conditional distribution of their rivals, which is why $G(x \mid s, \mathcal{I})$ depends separately on the rival threshold x and the bidder's own signal s. A talker has no private signal at the bidding stage, so $G_0(x \mid \mathcal{I})$ depends only on public information.

A.2 Analysis of the bidding stage

We study the continuation auction conditional on the public information $\mathcal{I}$ generated by disclosure. A continuation equilibrium consists of symmetric strategies within each disclosure group that are global best responses given $\mathcal{I}$. Talkers retain private information only about their idiosyncratic values. Silent dealers retain both their signals and their idiosyncratic values. The exact silent strategy is therefore $b^S(s, u; \mathcal{I})$, and its zero-value slice is $b^S(s; \mathcal{I}) \equiv b^S(s, 0; \mathcal{I})$. Evaluating this slice does not set rivals' idiosyncratic values to zero.

We first distinguish pooling of valuation types from atoms in bids. We then establish posterior comparisons, exact clearing probabilities, rent bounds, and bid-support properties. The scalar equations concern an auxiliary auction in which every silent dealer has zero idiosyncratic value. A

separate two-dealer all-silent construction gives an exact zero-width continuation and a globally valid approximation at small positive width. A finite-grid construction also gives measurable equilibrium choices, uniform bid moments, and a bound on policy-induced regret. The exact positive-width selection required for disclosure retains its conditional scope.

Fix the disclosed signals $\mathcal{M}$. Write $f_M(v) = f(v \mid \mathcal{M})$ for the Gaussian posterior density, with mean μ_T and variance $\sigma^2_{v|\mathcal{M}} > 0$. As in Section A.1, $w(t) = 1 - p(t)$ is the silence weight, $Q(x, v) = \int_{-\infty}^{x} w(t) f_\varepsilon(t - v)\, dt$, and $q(v) = Q(\infty, v)$. Conditional on v and silence, a signal has density $w(t) f_\varepsilon(t - v)/q(v)$. Throughout, histories have positive likelihood under the specified information structure. Section 3.2 specifies beliefs after an off-path exact disclosure. Bids range over the real line and independent random priorities resolve ties. When $\bar{u} > 0$, F_u is continuous with full support on $[0, \bar{u}]$. Formulas involving f_u additionally require a continuous positive density on $(0, \bar{u})$.

A.2.1 Pooling and separation

A valuation type at the bidding stage is the pair (s, u). The talking cost is sunk and does not affect the continuation bid. The following result concerns the information conveyed by a bid and the public disclosure history. It does not require the silent strategy to depend on a scalar valuation index.

Lemma A3 (Pooling and separation of valuation types)**.** *Fix a public history. Suppose $b^S(s, u)$ is continuous on an open type rectangle and strictly increasing in u for every s. Through each interior type (s_0, u_0) there is a continuous curve $(s, u(s))$ of distinct valuation types submitting $b^S(s_0, u_0)$. If b^S is also strictly increasing in s, then $u(s)$ is strictly decreasing. If the conditional idiosyncratic-value law given s is atomless and the strict increase in u holds on its full support, the silent bid distribution is atomless. If a talker's signal is disclosed exactly and their bid is strictly increasing in u on its support, the disclosed signal and bid jointly identify their valuation type.*

Proof. Choose $u_- < u_0 < u_+$ inside the rectangle and put $b_0 = b^S(s_0, u_0)$. Strict increase gives $b^S(s_0, u_-) < b_0 < b^S(s_0, u_+)$. Continuity preserves these strict inequalities for all s in a neighborhood of s_0. The intermediate value theorem gives a solution $u(s) \in (u_-, u_+)$ to $b^S(s, u(s)) = b_0$, and strict increase in u makes it unique. For any convergent sequence of such signals, compactness of $[u_-, u_+]$ gives convergent subsequences of their idiosyncratic values. Continuity makes every subsequential limit solve the limiting bid equation. Uniqueness then gives convergence of the entire sequence, proving continuity of $u(s)$.

If $s_2 > s_1$ and $u(s_2) \geq u(s_1)$, strict signal monotonicity and idiosyncratic-value monotonicity would give $b^S(s_2, u(s_2)) > b^S(s_1, u(s_1))$, a contradiction. Hence the curve decreases. For atomlessness, fix a bid b and signal s. At most one idiosyncratic value submits b, so its conditional probability is zero. Integration over the signal law gives zero probability of an atom at b. Finally, equality of two observed talker pairs first implies equality of their disclosed signals. Strict increase of the talking bid then implies equality of their idiosyncratic values. □

Silent pooling concerns different combinations of common-value information and idiosyncratic value. It is compatible with a continuous distribution of observed bids. Talking separates these valuation components through the pair of observed actions. The pair need not identify the sunk talking cost. Neither the lemma nor atomlessness establishes that both disclosure actions occur in equilibrium with positive probability.

A.2.2 A preliminary ordering result

Let $S_{\mathcal{S}}^{\max}$ denote the maximum of the n_S silent signals and let $S_{(1)}$ denote the maximum excluding a specified silent dealer j. When $n_S = 1$, the event $S_{(1)} < x$ is vacuous. Write $G_0(x \mid \mathcal{I})$ and $G(x \mid s, \mathcal{I})$ for the corresponding maximum distributions, and g_0 and g for their densities when a maximum contains at least one signal. Define $K_0 = g_0/G_0$ and $K = g/G$ wherever their denominators are positive.

The common-value mean after all silent signals fall below a common cutoff is given by (13). It is defined when the conditioning event has positive probability. If there are no silent dealers, the event is vacuous, and the mean is the public posterior mean. The boundary-state mean is

$$\varphi(x; \mathcal{I}) \equiv \mathbb{E}[v \mid \mathcal{I}, S_{\mathcal{S}}^{\max} = x]. \tag{A.19}$$

Exact-signal conditioning is interpreted through the corresponding density. The following comparisons apply at interior points where $w(x) > 0$ and the conditioning distributions are nondegenerate.

Lemma A4 (Own-signal versus rival-signal conditioning)**.** *Fix $\mathcal{I}$ with $n_S \geq 1$ and an interior silent signal x satisfying the stated conditions. Then*

(a) $\mathbb{E}[v \mid s_j = x, \mathcal{I}, S_{(1)} < x] = \varphi(x; \mathcal{I})$.

(b) $\varphi(x; \mathcal{I}) > \psi(x; \mathcal{I})$.

Proof. Conditional on v, the likelihood of the event in part (a) is proportional to $w(x) f_\varepsilon(x - v) Q(x, v)^{n_S - 1}$. The likelihood that the maximum of all n_S silent signals equals x is n_S times this expression. The factor is independent of v, so Bayes' rule gives identical posteriors.

The likelihood ratio of the boundary posterior to the posterior conditional on all silent signals being below x is proportional to $f_\varepsilon(x - v)/Q(x, v)$. Its reciprocal satisfies

$$\frac{Q(x, v)}{f_\varepsilon(x - v)} = \int_0^\infty w(x - r) \exp\left(-\frac{v - x}{\sigma_\varepsilon^2} r - \frac{r^2}{2\sigma_\varepsilon^2} \right) dr. \tag{A.20}$$

For every $r > 0$, the integrand is strictly decreasing in v wherever its weight is positive. The posterior ratio is therefore strictly increasing. Strict monotone likelihood-ratio dominance gives the mean inequality. □

The function ψ concerns a common signal cutoff. It is not generally the common-value assessment of a talker who clears the actual silent bids. Differentiating its truncated-mean representation gives $\psi' = K_0(\varphi - \psi) > 0$ wherever the densities are continuous. As $x \to \infty$, it converges to $\widetilde{\mu}_T \equiv \mathbb{E}[v \mid \mathcal{I}]$.

A.2.3 Why talkers require idiosyncratic value

In the pure common-value limit $\bar{u} = 0$, all talkers have the same payoff-relevant information. For this limit, suppose the silent signal support is an interval, with positive density in its interior and continuous conditional means. Both full-support silence weights and the hard-truncation benchmark satisfy this support condition.

Lemma A5 (No pure talker bid in the common-value limit)**.** *Fix $\mathcal{I}$ with $n_T \geq 1$ and $n_S \geq 1$, and suppose $\bar{u} = 0$. Under the stated support conditions, there is no continuation equilibrium in which*

all talkers submit the same deterministic bid, the talker side wins with positive probability, and silent dealers use a continuous strictly increasing schedule.

Proof. Suppress $\mathcal{I}$ and denote the proposed common talker bid by b^T. Optimized gross payoffs are nonnegative. Indeed, the expected negative part of the payoff from a bid b is at most $\mathbb{E}[(b-v)_+]$, which tends to zero as $b \to -\infty$ by finite first moments. Thus the supremum over bids is at least zero, without requiring a finite bid that never wins.

Positive probability of a talker win implies that some silent bids lie below b^T. Some must also lie above it. Otherwise, the talker side always clears silence, and nonnegative talker payoff gives $b^T \leq \widetilde{\mu}_T$. The silent dealer's posterior mean $\mathbb{E}[v \mid s, \mathcal{I}]$ is strictly increasing in s under the Gaussian likelihood and averages to $\widetilde{\mu}_T$ over their conditional signal distribution. Hence a positive-probability set of silent types has posterior mean strictly above b^T. Such a type earns zero at their assigned bid below b^T and gains strictly by bidding just above b^T, which clears every rival. This is a contradiction. Tie-breaking among the identical talker bids does not change their common-value assessment.

Continuity, strict monotonicity, and the interval support now give a unique interior σ with $b^S(\sigma) = b^T$. Types just above σ win with positive probability. Nonnegative payoff and continuity imply $\varphi(\sigma) \geq b^T$. Types just below σ earn zero and must not gain by bidding just above b^T. Letting their signals approach σ gives the reverse inequality. Thus

$$b^T = \mathbb{E}[v \mid s_j = \sigma, \mathcal{I}, S_{(1)} < \sigma] = \varphi(\sigma; \mathcal{I}). \tag{A.21}$$

A talker who wins conditions on all silent signals being below σ, so their common-value assessment is $\psi(\sigma; \mathcal{I}) < b^T$ by Lemma A4. A talker with positive winning probability therefore earns a strictly negative payoff, contradicting global optimality. □

The bid needed to defeat the marginal silent dealer reflects the common value at the boundary of the talker's winning event. The talker values the asset using the average over that event, which is lower. Idiosyncratic value can cover this information gap. The lemma excludes the specified pure strategies. Randomization is discussed below after deriving the scalar first-order condition.

A.2.4 Exact clearing probabilities

Fix $n_T \geq 1$ and $n = n_S \geq 2$. Suppose silent bids are continuous, strictly increasing in the signal, and nondecreasing in u. Suppress $\mathcal{I}$ when no confusion arises. A rival's signal cutoff at bid b is $\chi(b,u) = \sup\{t : b^S(t,u) < b\}$, with infinite endpoints allowed. Set $Q(-\infty, v) = 0$ and $Q(+\infty, v) = q(v)$, with $\sup \varnothing = -\infty$. The joint likelihood that one dealer is silent and bids below b is

$$C(b,v) = \int_0^{\bar{u}} Q(\chi(b,u), v)\, dF_u(u). \tag{A.22}$$

Conditional on silence and v, their clearing probability is $C(b,v)/q(v)$. The weighted Gaussian signal density has a monotone likelihood ratio in its signal as v increases. Every silence-conditioned cutoff probability is therefore nonincreasing in v, and integration over u preserves this ordering. The conclusion concerns $C(b,v)/q(v)$. It does not require the unnormalized likelihood $C(b,v)$ itself to be monotone in v.

Let $\mathcal{C}_b^{\text{ex}} = \{b^S(s_k, u_k) < b \text{ for every } k \in \mathcal{S}\}$ and let $Q_T(b) = \Pr(\mathcal{C}_b^{\text{ex}} \mid \mathcal{I})$. The talker's common-value assessment is

$$\Psi^{\text{ex}}(b; \mathcal{I}) = \mathbb{E}[v \mid \mathcal{I}, \mathcal{C}_b^{\text{ex}}]. \tag{A.23}$$

Conditional independence given v and cancellation of the silence likelihoods give

$$Q_T(b) = \frac{\int f_M(v)C(b,v)^n\,dv}{\int f_M(v)q(v)^n\,dv}, \qquad \Psi^{\mathrm{ex}}(b) = \frac{\int v f_M(v)C(b,v)^n\,dv}{\int f_M(v)C(b,v)^n\,dv}. \tag{A.24}$$

All bid-conditioned means below are defined where the corresponding clearing probability is positive.

For a silent dealer with signal s, let $G^{\mathrm{ex}}(b \mid s)$ be the probability of clearing their $n-1$ silent rivals and let $V^{\mathrm{ex}}(s,b)$ be the associated common-value mean. Writing $f_s(v) = f(v \mid s, \mathcal{M})$ gives

$$G^{\mathrm{ex}}(b \mid s) = \frac{\int f_s(v)C(b,v)^{n-1}\,dv}{\int f_s(v)q(v)^{n-1}\,dv}, \qquad V^{\mathrm{ex}}(s,b) = \frac{\int v f_s(v)C(b,v)^{n-1}\,dv}{\int f_s(v)C(b,v)^{n-1}\,dv}. \tag{A.25}$$

Let H be the cdf of the highest talker bid and h its density where it exists. Conditional on $\mathcal{I}$, talker values are independent of v and silent signals. The silent dealer's exact fixed-bid payoff is therefore $H(b)G^{\mathrm{ex}}(b \mid s)[V^{\mathrm{ex}}(s,b) + u - b]$. Conditional on any idiosyncratic value, the signal has a density and the signal bid is strictly increasing. Hence a silent bid has no atoms, and integration over idiosyncratic values preserves atomlessness. We use the talker formulas on atomless portions of its bid distribution.

At a differentiable interior silent optimum with positive winning probability, write $g^{\mathrm{ex}} = \partial_b G^{\mathrm{ex}}$. Differentiating the exact payoff gives

$$0 = \left(\frac{h(b)}{H(b)} + \frac{g^{\mathrm{ex}}(b \mid s)}{G^{\mathrm{ex}}(b \mid s)}\right)[V^{\mathrm{ex}}(s,b) + u - b] + \partial_b V^{\mathrm{ex}}(s,b) - 1. \tag{A.26}$$

The support argument in Appendix A.2.8 uses these exact clearing objects. Rival strategies are held fixed in the bid derivatives. The clearing likelihood C integrates cutoffs from the full strategy $b^S(t,u)$. Consequently, this condition does not close as an equation for the slice $b^S(s,0)$ alone.

A.2.5 Maintained continuation-equilibrium class

Write $\bar{b} = \beta^T(\bar{u}; \mathcal{I})$ for the finite highest talker bid. For an assigned talker bid in the image of the zero-value silent slice, let $x(u)$ solve $b^S(x(u),0) = \beta^T(u)$. Define the exact and common-cutoff margins by $\rho^{\mathrm{ex}}(u) = \Psi^{\mathrm{ex}}(\beta^T(u)) + u - \beta^T(u)$ and $\rho(u) = \psi(x(u)) + u - \beta^T(u)$, respectively. The second margin is an information comparison, even though $x(u)$ is obtained from the exact zero-value slice.

The auxiliary scalar auction below sets all silent dealers' idiosyncratic values to zero and retains talker heterogeneity. Denote its silent schedule by $b^{S,\mathrm{bm}}$. This schedule and its continuation payoffs are the benchmark objects used in the maintained compact-set approximation. They need not equal the zero-value slice and payoffs of the exact game.

Assumption A1 (Monotone-equilibrium regularity)**.** *For sufficiently small $\bar{u} > 0$, the continuation games used to evaluate disclosure admit finite-value equilibria in globally optimal strategies. A measurable selection across the policies and public histories used in the disclosure comparison is maintained. In configurations with $n_T \geq 1$ and $n_S \geq 2$, the talker strategy is continuous and strictly increasing on $(0, \bar{u}]$, with finite $\bar{b}$. The silent strategy is continuous and strictly increasing in the signal and in u on the relevant interior domain. Its zero-value slice eventually exceeds $\bar{b}$. The highest talker bid has an atomless distribution.*

On the overlap region where $H(b^S(s,0)) > 0$ and $b^S(s,0) \leq \bar{b}$, the exact zero-value slice satisfies $\psi(s) \leq b^S(s,0) \leq \varphi(s)$. In these mixed configurations with $n_S \geq 2$, every positive-winning assigned

talker bid satisfies $0 \le \rho^{\mathrm{ex}}(u) \le \rho(u)$*. For configurations with one silent dealer, the positive-winning talker margins satisfy the maintained bound* $0 \le \rho^{\mathrm{ex}}(u) \le u$*. Boundary configurations otherwise use the best-response conditions in Remark A2.*

When the auxiliary scalar approximation is used, exact silent bids and continuation payoffs are within $O(\bar{u})$ *of their benchmark counterparts uniformly on compact signal sets. The denominator of the scalar inside-support slope equation is positive on the branch where it is used. In the talker-selection argument, the exact slope field is locally Lipschitz in the bid and has a positive denominator around the candidate schedules.*

The corridor and margin comparison are maintained restrictions. With one silent dealer, the rent restriction is stated directly because matching every talker bid to the zero-value silent slice is not guaranteed. The invariant-band result below retains that restriction in this boundary configuration. The compact-set approximation supplies no uniform lower-tail or derivative bound and does not equate exact bid-clearing events with common signal cutoffs. Existence, measurable selection, the corridor, and the stated margin bounds remain maintained conditions for the full disclosure argument. The independent two-dealer results in Sections A.2.14–A.2.17 establish a normalized zero-width equilibrium and finite-grid continuation results. They do not establish the maintained positive-width equilibrium class at every public history.

A.2.6 Talker rents and the invariant disclosure band

Let $\Pi^T(u)$ be the optimized exact talker payoff and $D(u)$ their assigned winning probability. Nonnegative optimized payoffs follow from the low-bid argument in the proof of Lemma A5, now using $\mathbb{E}[(b - v - u)_+]$. This argument also applies to silent dealers.

Lemma A6 (Talker envelope and payoff bound)**.** *Fix a mixed continuation configuration satisfying Assumption A1, including its margin comparison on the positive-winning range. The optimized talker payoff is nondecreasing, convex, and* 1*-Lipschitz on* $[0, \bar{u}]$*, with its continuous extension at zero. It is absolutely continuous,* $\Pi^{T\prime}(u) = D(u)$ *almost everywhere, and*

$$0 \le \Pi^T(u) = \int_0^u D(t)\, dt \le u. \tag{A.27}$$

Proof. For each fixed bid, the payoff is affine in u, with slope in $[0, 1]$. Taking suprema makes the optimized payoff convex and nondecreasing. Comparing every fixed bid at two idiosyncratic values and then taking suprema bounds the difference in optimized values by the difference in idiosyncratic values, so the value is 1-Lipschitz wherever finite. The margin bound below establishes finiteness. At a differentiability point, the slope of any assigned optimal bid equals the derivative of the value, giving the envelope identity almost everywhere.

If $D(u) = 0$, the assigned payoff is zero. If $D(u) > 0$ and $n_S \ge 2$, the bid clears a positive-probability set of silent types, so some zero-value silent bid is below it. The maintained eventual upper crossing supplies a zero-value bid above it. Continuity and strict signal monotonicity therefore give $x(u)$, including the upper boundary by continuity. The corridor implies $\rho(u) \le u$, and the margin comparison implies $\Pi^T(u) = D(u)\rho^{\mathrm{ex}}(u) \le u$. For the one-silent-dealer case, the same conclusion follows from the directly maintained margin bound. Together with nonnegative payoff, this gives $\lim_{u \downarrow 0} \Pi^T(u) = 0$. Absolute continuity completes the integral formula. □

Let $\bar{\Pi}^T(s; p)$ and $\bar{\Pi}^S(s; p)$ be the gross continuation values before idiosyncratic values are drawn, and let $\Delta(s; p) = \bar{\Pi}^T(s; p) - \bar{\Pi}^S(s; p)$. The disclosure best response is $\mathcal{B}(p)(s) = G_\kappa(\Delta(s; p))$.

Lemma A7 (Invariant disclosure band)**.** *Suppose the selected continuation games exist for every disclosure configuration reached under p, with the margin conditions of Assumption A1 whenever both groups are present. Then $\bar{\Pi}^T(s;p) \le \mu_u$, $\bar{\Pi}^S(s;p) \ge 0$, and $\Delta(s;p) \le \mu_u$ for every signal, where $\mu_u = \mathbb{E}[u]$. With $\kappa \sim U[0,K]$ and $\bar{u} < K$,*

$$0 \le \mathcal{B}(p)(s) \le \frac{\mu_u}{K} \le \frac{\bar{u}}{K} < 1. \tag{A.28}$$

Every symmetric disclosure equilibrium in this class obeys the same bound.

Proof. In mixed configurations, Lemma A6 bounds talker payoff by u. If all dealers talk, the standard symmetric private-value continuation auction, with public common-value mean added to every valuation, has zero lowest-type rent and winning probability at most one, giving the same bound. Silent dealers have nonnegative optimized gross payoffs in every configuration. Averaging first over the focal idiosyncratic value gives $\bar{\Pi}^T(s;p) \le \mathbb{E}[u] = \mu_u$ in every configuration. Averaging over public histories preserves that bound, while $\bar{\Pi}^S(s;p) \ge 0$. Apply the uniform cost cdf to $\Delta \le \mu_u \le \bar{u} < K$. □

Define $\underline{q} = 1 - \bar{u}/K > 0$.

Corollary A1 (Coexistence and a lower bound on silence)**.** *At a symmetric disclosure equilibrium under the stated continuation conditions and $\bar{u} < K$, every signal has positive probability of silence. In particular, any type that talks with positive probability also remains silent with positive probability, $w(s) \ge \underline{q}$, and*

$$q(v) = \int_{\mathbb{R}} w(t) f_\varepsilon(t - v)\, dt \ge \underline{q}. \tag{A.29}$$

Proof. Equation (A.28) gives $w(s) = 1 - p(s) \ge \underline{q}$. Integration against the signal-error density gives the bound on q. □

Lemma A8 (Full support of silent signals)**.** *For any policy satisfying $0 \le p \le \bar{u}/K < 1$, the conditional support of every silent signal is $\mathbb{R}$. For $n_S \ge 2$, $G(x \mid s, \mathcal{I}) > 0$ and $G_0(x \mid \mathcal{I}) > 0$ at every finite x. For $n_S = 1$, $G \equiv 1$ and $G_0(x \mid \mathcal{I}) > 0$.*

Proof. The density $w(t) f_\varepsilon(t - v)/q(v)$ is strictly positive everywhere. Conditional independence given v gives positive probability to each finite collection of lower-cutoff events. Integrating over v preserves positivity. □

A.2.7 Signal boundary posteriors

Fix disclosed signals $\mathcal{M}$ and let $f_M(v) = f(v \mid \mathcal{M})$ be the Gaussian posterior density with mean μ_T and variance $\sigma^2_{v|\mathcal{M}} > 0$. Fix $n = n_S \ge 1$ and a continuous disclosure policy p satisfying $0 \le p \le \bar{u}/K < 1$. Write $w = 1 - p$ and $\underline{q} = 1 - \bar{u}/K > 0$. The joint likelihood that a dealer is silent and has signal below x is $Q(x,v) = \int_{-\infty}^{x} w(t) f_\varepsilon(t - v)\, dt$, and the silence probability is $q(v) = Q(\infty, v)$.

Let Ω_s be the posterior with density proportional to $f_M(v) Q(s,v)^n$, and let Ω_s^s have density proportional to $f_M(v) f_\varepsilon(s - v) Q(s,v)^{n-1}$. Then $\psi(s;\mathcal{I}) = \mathbb{E}_{\Omega_s}[v]$ and $\varphi(s;\mathcal{I}) = \mathbb{E}_{\Omega_s^s}[v]$. Define the cutoff reverse hazard $W(v;s) = w(s) f_\varepsilon(s - v)/Q(s,v)$, its mean $\mu_1(s) = \mathbb{E}_{\Omega_s}[W]$, and the covariance $\kappa_1(s) = \mathrm{Cov}_{\Omega_s}(v, W)$.

For the lower-tail calculation, define $D = \sigma_\varepsilon^2 + n\sigma_{v|\mathcal{M}}^2$, $a = \sigma_{v|\mathcal{M}}^2/D$, $\tau^2 = \sigma_\varepsilon^2\sigma_{v|\mathcal{M}}^2/D$, $L_s = (\mu_T - s)/D$, and $m_s = (\sigma_\varepsilon^2\mu_T + n\sigma_{v|\mathcal{M}}^2 s)/D$. Constants in uniform bounds may depend on n, $\underline{q}$, and the two Gaussian variances. They do not depend on the policy or on μ_T when the tail is indexed by L_s.

Lemma A9 (Gaussian signal-boundary behavior)**.** *For every continuous policy in the stated band, $\psi'(s;\mathcal{I}) > 0$ at every finite s. For all sufficiently large L_s, uniformly over these policies, there are constants $c, C > 0$ such that $c \le \psi'(s;\mathcal{I}) \le C$ and*

$$\psi(s;\mathcal{I}) = m_s + O(L_s^{-1}), \qquad \varphi(s;\mathcal{I}) = m_s + O(L_s^{-1}), \qquad 0 < \varphi(s;\mathcal{I}) - \psi(s;\mathcal{I}) \le CL_s^{-1}. \tag{A.30}$$

For each fixed policy with $w(s) \to w_- > 0$ as $s \to -\infty$,

$$\psi'(s;\mathcal{I}) \longrightarrow na, \qquad \varphi(s;\mathcal{I}) - \psi(s;\mathcal{I}) \sim \frac{\sigma_{v|\mathcal{M}}^2}{\mu_T - s}. \tag{A.31}$$

Proof. Under Ω_s, let $R_i = s - S_i \ge 0$ be the gaps of the n silent signals below the cutoff. Conditional Gaussian regression gives

$$v = m_s - a\sum_{i=1}^{n} R_i + Z, \qquad Z \sim \mathcal{N}(0, \tau^2), \tag{A.32}$$

where Z is independent of the gaps. This follows by first conditioning on the entire signal vector. The Gaussian regression error is independent of that vector, so subsequent truncation and multiplication by the weights $w(S_i)$ preserve independence.

The marginal signal covariance is $\Sigma = \sigma_\varepsilon^2 I_n + \sigma_{v|\mathcal{M}}^2 \mathbf{1}\mathbf{1}^\top$, with $\Sigma^{-1}\mathbf{1} = D^{-1}\mathbf{1}$. Substituting $S_i = s - r_i$ into its density shows that the gap density is proportional to

$$\prod_{i=1}^{n} w(s - r_i) \exp\left(-L_s \sum_{i=1}^{n} r_i - \tfrac{1}{2} r^\top \Sigma^{-1} r\right), \qquad r \in \mathbb{R}_+^n. \tag{A.33}$$

Fix $L_0 > 0$. Since the largest eigenvalue of Σ^{-1} is at most σ_ε^{-2}, integration on $[0, L_s^{-1}]^n$ bounds the normalizing integral below by a positive constant times L_s^{-n} whenever $L_s \ge L_0$. Dropping the weights and quadratic term bounds the kth moment numerator of $\sum_i R_i$ above by $\Gamma(n+k)L_s^{-(n+k)}/\Gamma(n)$. Consequently, $\mathbb{E}[(\sum_i R_i)^k] = O(L_s^{-k})$ uniformly over the band for every fixed integer $k \ge 1$. Equation (A.32) implies $\mathbb{E}_{\Omega_s} v = m_s + O(L_s^{-1})$, uniformly bounded moments of $v - m_s$ of each fixed order, and

$$\mathrm{Var}_{\Omega_s}(v) = \tau^2 + a^2 \,\mathrm{Var}\left(\sum_i R_i\right) = \tau^2 + O(L_s^{-2}). \tag{A.34}$$

Next, differentiate the hazard in the common value. For $j = 0, 1$, let $I_j(s, v) = \int_0^\infty r^j w(s - r) \exp(-(v - s)r/\sigma_\varepsilon^2 - r^2/(2\sigma_\varepsilon^2))\, dr$. Substitution into Q gives

$$W(v; s) = \frac{w(s)}{I_0(s, v)}, \qquad \partial_v W(v; s) = \frac{w(s) I_1(s, v)}{\sigma_\varepsilon^2 I_0(s, v)^2} > 0. \tag{A.35}$$

Let $\lambda_0(v;s) = f_\varepsilon(s-v)/F_\varepsilon(s-v)$ be the unweighted Gaussian reverse hazard. Bounds on w give

$$\underline{q}\lambda_0 \le W \le \underline{q}^{-1}\lambda_0, \qquad \underline{q}^2\partial_v\lambda_0 \le \partial_v W \le \underline{q}^{-2}\partial_v\lambda_0 \le \frac{\underline{q}^{-2}}{\sigma_\varepsilon^2}. \tag{A.36}$$

For the last inequality, if $T \sim \mathcal{N}(v,\sigma_\varepsilon^2)$, then $\mathbb{E}[T \mid T \le s] = v - \sigma_\varepsilon^2\lambda_0(v;s)$. Differentiating in v gives $1 - \sigma_\varepsilon^2\partial_v\lambda_0 = \mathrm{Var}(T \mid T \le s)/\sigma_\varepsilon^2 \ge 0$. Positivity of $\partial_v\lambda_0$ follows from (A.35) with weight one. Rescaling the unweighted integrals by $(v-s)/\sigma_\varepsilon^2$ also gives $\partial_v\lambda_0 \to \sigma_\varepsilon^{-2}$ as $v - s \to \infty$.

For independent copies X, X' and any function g for which the covariance is finite,

$$\mathrm{Cov}(X, g(X)) = \tfrac{1}{2}\mathbb{E}[(X - X')(g(X) - g(X'))]. \tag{A.37}$$

Because $(m_s - s)/\sigma_\varepsilon^2 = L_s$, the derivative of $\lambda_0(v;s)$ converges uniformly to σ_ε^{-2} on compact sets of $v - m_s$. Use (A.37), first on such compact sets and then on their complements. The global derivative bound and uniform fourth moments obtained above make the complementary terms uniformly small. Together with (A.34), this gives $\mathrm{Cov}_{\Omega_s}(v,\lambda_0) \to a$ uniformly over policies. Equation (A.36) therefore bounds κ_1 above and away from zero for large L_s. Strict increase of W gives $\kappa_1 > 0$ at every finite s.

The same unweighted integral rescaling gives $\lambda_0(m_s;s)/L_s \to 1$. Its global Lipschitz bound and the bounded first moment of $v - m_s$ imply $\mathbb{E}_{\Omega_s}\lambda_0/L_s \to 1$ uniformly. The level comparison in (A.36) then gives two positive constant bounds for μ_1/L_s.

Differentiating the normalized density of Ω_s in the cutoff yields $\psi' = n\kappa_1$. Moreover, Bayes' rule gives $d\Omega_s^s/d\Omega_s = W/\mu_1$. Hence

$$\varphi(s;\mathcal{I}) - \psi(s;\mathcal{I}) = \frac{\kappa_1(s)}{\mu_1(s)}. \tag{A.38}$$

The bounds just proved establish (A.30), including the mean bound for φ. Continuity and boundedness of w and Gaussian integrability justify the cutoff differentiation.

Finally suppose $w(t) \to w_- > 0$. Set $\epsilon_s = \sup_{t\le s}|w(t)/w_- - 1| \to 0$. Both I_j and $w(s)$ differ from their constant-w_- counterparts by relative errors $O(\epsilon_s)$, uniformly in v. Equation (A.35) thus gives $W/\lambda_0 = 1 + O(\epsilon_s)$ and $(\partial_v W)/(\partial_v\lambda_0) = 1 + O(\epsilon_s)$. Apply the second comparison between two common values and use (A.37). It follows that $\kappa_1 = (1 + O(\epsilon_s))\,\mathrm{Cov}_{\Omega_s}(v,\lambda_0) \to a$. The level comparison gives $\mu_1/L_s \to 1$. Equation (A.38) and $a/L_s = \sigma^2_{v|\mathcal{M}}/(\mu_T - s)$ prove (A.31). □

A continuous nonincreasing policy in the band has a positive lower-tail silence limit. The exact limits therefore hold without an additional rate restriction on p. The positive derivative bounds and vanishing wedge are uniform over the whole continuous band. The exact limiting value of ψ' is a fixed-policy statement.

Lemma A10 (Lower-tail boundary valuation)**.** *Fix a public configuration with $n_S \ge 1$. For continuous policies satisfying $0 \le p \le \bar{u}/K < 1$, $\varphi(s;\mathcal{I}) \to -\infty$ as $s \to -\infty$, uniformly over the policy band.*

Proof. Equation (A.30) gives $\varphi(s) = m_s + O(L_s^{-1})$ uniformly, and m_s is affine in s with strictly positive slope. □

A.2.8 Exact lower support and the entry margin

Maintain a continuous policy in the band, $n_T \ge 1$, and $n = n_S \ge 2$. The following argument uses the exact clearing objects in Appendix A.2.4.

Let $\widetilde{\mu}_S(s) = \mathbb{E}[v \mid s, \mathcal{I}]$ be the silent dealer's posterior mean before observing the outcome of the auction. Let $\mu_S(s, \mathcal{M})$ and σ_S^2 be the mean and variance of $f(v \mid s, \mathcal{M})$. Gaussian conjugacy makes $\mu_S(s, \mathcal{M})$ affine in s with strictly positive slope and makes σ_S^2 independent of s. The full posterior multiplies this Gaussian density by $q(v)^{n-1}$. Since $\underline{q} \le q(v) \le 1$,

$$|\widetilde{\mu}_S(s) - \mu_S(s, \mathcal{M})| \le \underline{q}^{-(n-1)}\sigma_S. \tag{A.39}$$

To obtain this bound, center the numerator at μ_S, bound its absolute value by the Gaussian mean absolute deviation, and bound the normalizing denominator below by $\underline{q}^{n-1}$. Thus $\widetilde{\mu}_S(s) \to -\infty$ as $s \to -\infty$, uniformly over policies in the band for the fixed history.

Global optimality implies nonnegative optimized gross payoffs. Conditional on a bidder's private information, the expected negative part of the payoff from a bid b is at most $\mathbb{E}[\max\{b - v - u, 0\}]$. This bound tends to zero as $b \to -\infty$ because the conditional common value has a finite first moment. The supremum over bids is therefore nonnegative. This argument applies to talkers and silent dealers and uses no lower-support conclusion.

Clearing silent rivals multiplies this full posterior by $[C(b, v)/q(v)]^{n-1}$, a nonincreasing function of v. It follows that $V^{\mathrm{ex}}(s, b) \le \widetilde{\mu}_S(s)$ whenever the clearing probability is positive. At an equilibrium bid with positive winning probability, nonnegative optimized payoff therefore implies

$$b^S(s, u) \le V^{\mathrm{ex}}(s, b^S(s, u)) + u \le \widetilde{\mu}_S(s) + u. \tag{A.40}$$

We next compare a talker and a boundary silent dealer using the same actual rival bid distribution. Let b and x be finite with $b^S(x, 0) = b$ and $C(b, v) > 0$. Let Ω_T^b have density proportional to $f_M(v)C(b, v)^n$ and let $\Omega_S^{x,b}$ have density proportional to $f_M(v)f_\varepsilon(x - v)C(b, v)^{n-1}$. Their means are $\Psi^{\mathrm{ex}}(b)$ and $V^{\mathrm{ex}}(x, b)$.

Lemma A11 (Exact boundary information ordering)**.** *If $b^S(x, 0) = b$, the silent strategy is nondecreasing in u and strictly increasing in the signal, and $C(b, v) > 0$ for every v, then $\Omega_S^{x,b}$ strictly dominates Ω_T^b in the monotone likelihood-ratio order. In particular, $V^{\mathrm{ex}}(x, b) > \Psi^{\mathrm{ex}}(b)$.*

Proof. Monotonicity gives $\chi(b, u) \le x$ for every u. Dividing the densities shows that their likelihood ratio is proportional to $f_\varepsilon(x - v)/C(b, v)$. Its reciprocal can be written as

$$\frac{C(b, v)}{f_\varepsilon(x - v)} = \int_0^{\bar{u}} \int_{x - \chi(b,u)}^{\infty} w(x - r) \exp\left(-\frac{v - x}{\sigma_\varepsilon^2} r - \frac{r^2}{2\sigma_\varepsilon^2}\right) dr\, dF_u(u). \tag{A.41}$$

The lower integration limits are nonnegative. Positivity of C means that positive mass lies at $r > 0$, so the expression is strictly decreasing in v. The density ratio is therefore strictly increasing. Both posteriors have positive densities on the real line and finite first moments. Strict likelihood-ratio dominance gives the strict mean inequality. □

The ordering uses the zero-value bid only to locate the upper bound on rivals' signal cutoffs. All rival idiosyncratic values remain integrated in C.

Let $\underline{b} = \lim_{u \downarrow 0} \beta^T(u)$, allowing $\underline{b} = -\infty$.

Proposition A2 (No finite lower bid endpoint)**.** *Fix a continuous policy satisfying $0 \le p \le \bar{u}/K < 1$ and a public history with $n_T \ge 1$ and $n_S \ge 2$. Suppose talker bids are continuous and strictly increasing on $(0, \bar{u}]$ with finite $\bar{b}$, silent bids are continuous and strictly increasing in the signal and*

nondecreasing in u, and both sides optimize globally. If $b^S(s,0) > \bar{b}$ for all sufficiently large s, then $\underline{b} = -\infty$, H has support $(-\infty, \bar{b}]$, and

$$\sup_{u\in[0,\bar{u}]} b^S(s,u;\mathcal{I}) \longrightarrow -\infty \qquad \text{as } s \to -\infty. \tag{A.42}$$

Proof. Suppose first that $\underline{b}$ is finite. Atomlessness and continuity imply $H(\underline{b}) = 0$ and $H(b) > 0$ for $\underline{b} < b \leq \bar{b}$.

Some zero-value silent bids must lie below $\underline{b}$. Otherwise, strict signal monotonicity would give $b^S(s,0) > \underline{b}$ for every finite s, so every such type would clear the talker side with positive probability. They would also clear their silent rivals with positive probability. To see this, choose a rival signal $t < s$ so that $b^S(t,0) < b^S(s,0)$, and use continuity and full support to obtain a positive-probability set of signals near t and idiosyncratic values near zero bidding below $b^S(s,0)$. Conditional independence given v then gives a positive probability that all silent rivals are in such sets. Equation (A.40) would imply $b^S(s,0) \leq \widetilde{\mu}_S(s) \to -\infty$, contradicting the proposed lower bound.

Because the zero-value silent bid eventually exceeds $\bar{b}$, there is a unique finite x_0 with $b^S(x_0,0) = \underline{b}$. The same continuity and full-support argument implies $C(\underline{b}, v) > 0$ for every finite v. Every positive talker type bids above $\underline{b}$, clears silent rivals with positive probability, and clears rival talkers with positive probability. Nonnegative optimized payoff therefore gives $\Psi^{\text{ex}}(\beta^T(u)) + u - \beta^T(u) \geq 0$. Letting $u \downarrow 0$ yields

$$\Psi^{\text{ex}}(\underline{b}) \geq \underline{b}. \tag{A.43}$$

The limit is valid because the clearing probability at $\underline{b}$ is positive and the numerator and denominator defining the conditional mean are continuous in the bid. Gaussian integrability and bounded likelihoods justify dominated convergence.

Lemma A11 now gives $V^{\text{ex}}(x_0, \underline{b}) > \Psi^{\text{ex}}(\underline{b}) \geq \underline{b}$. By continuity in the signal and bid, a zero-value silent type with $s < x_0$ sufficiently close to x_0 obtains a strictly positive conditional margin by deviating to $\underline{b} + \epsilon$ for sufficiently small $\epsilon > 0$. This deviation clears talkers and silent rivals with positive probability. Their equilibrium bid is below $\underline{b}$ and earns zero because it never clears the talker side. The deviation contradicts global optimality. If optimality is imposed almost everywhere in u, the same strict profitable deviation persists for sufficiently small positive u by continuity. Thus $\underline{b} = -\infty$.

The continuous increasing talker schedule has image $(-\infty, \bar{b}]$, and full support of F_u gives the same support for H. Consequently, every finite silent bid clears talkers with positive probability. The argument above gives positive silent-rival clearance at a zero-value equilibrium bid, so (A.40) proves $b^S(s,0) \to -\infty$. At every finite bid of a general silent type, sufficiently low rival signals and sufficiently small rival idiosyncratic values then give positive rival-clearing probability. Apply (A.40) again and take the supremum over u to obtain $\sup_u b^S(s,u) \leq \widetilde{\mu}_S(s) + \bar{u} \to -\infty$. □

The proposition uses monotonicity, global optimality, and eventual upper crossing. It does not use the corridor, the comparison $\rho^{\text{ex}} \leq \rho$, or a scalar first-order condition.

Corollary A2 (Talker lower tail and silent compatibility)**.** *Fix a continuous policy in the band and a mixed configuration with $n_S \geq 2$ satisfying Assumption A1. Then $\beta^T(u;\mathcal{I}) \to -\infty$ as $u \downarrow 0$, and*

$$b^S(s;\mathcal{I}) - \psi(s;\mathcal{I}) \longrightarrow 0 \qquad \text{as } s \to -\infty. \tag{A.44}$$

Proof. Proposition A2 gives the talker limit and makes the zero-value silent bid lie in the talker

support for all sufficiently negative signals. The maintained corridor and Lemma A9 give

$$0 \le b^S(s,0) - \psi(s) \le \varphi(s) - \psi(s) = O(L_s^{-1}). \tag{A.45}$$

This proves compatibility along the full lower tail. □

Lemma A12 (Talker participation and rents)**.** *Fix a continuous policy in the band and a mixed configuration with $n_S \ge 2$ satisfying Assumption A1. Then:*

(i) $\Pi^T(0) = 0$ *by continuous extension and* $0 \le \Pi^T(u) \le u$.

(ii) Every $u > 0$ talker wins with positive probability and earns strictly positive expected surplus, with $0 < \rho^{\mathrm{ex}}(u) \le \rho(u) \le u$.

(iii) The talker support is $(-\infty, \bar{b}]$, and $\bar{b} \le \widetilde{\mu}_T + \bar{u}$.

Proof. Part (i) follows from the envelope bound. The uniform silent lower tail gives positive silent-clearing probability at every finite talker bid. A positive talker type also clears all rival talkers with positive probability because $F_u(u) > 0$. Hence $D(u) > 0$ for every $u > 0$, and $\Pi^T(u) = \int_0^u D(t)\,dt > 0$. Division by $D(u)$ gives the strict exact margin. The remaining margin bounds follow from the maintained comparison and corridor.

For part (iii), the support follows from Proposition A2. Relative to the full public posterior, clearing silence multiplies the density by $[C(b,v)/q(v)]^n$, a nonincreasing function of v. Hence $\Psi^{\mathrm{ex}}(b) \le \widetilde{\mu}_T$. Nonnegative winning margins imply $\beta^T(u) \le \widetilde{\mu}_T + u$, including the upper endpoint. □

The zero-value entry margin $s^\times(\mathcal{I})$ solves

$$b^S(s^\times(\mathcal{I}); \mathcal{I}) = \bar{b}(\mathcal{I}). \tag{A.46}$$

For each u, define $s^\times(u;\mathcal{I})$ by $b^S(s^\times(u;\mathcal{I}), u; \mathcal{I}) = \bar{b}(\mathcal{I})$, so $s^\times(0;\mathcal{I}) = s^\times(\mathcal{I})$.

Lemma A13 (Probabilistic screening at the entry margin)**.** *Fix a continuous policy in the band and a configuration with $n_T \ge 1$ and $n_S \ge 2$ satisfying Assumption A1. For every $u \in [0, \bar{u}]$, $s^\times(u;\mathcal{I})$ is finite and unique. For $s < s^\times(u;\mathcal{I})$, $H(b^S(s,u;\mathcal{I}) \mid \mathcal{I}) \in (0,1)$ and the probability of being outbid by the talker side is strictly decreasing in s. For $s \ge s^\times(u;\mathcal{I})$, that probability is zero. The entry margin is nonincreasing in u. Every positive talker type wins with positive probability and earns strictly positive expected surplus.*

Proof. Proposition A2 gives the lower-tail crossing for every u. The maintained eventual upper crossing for $u = 0$ and monotonicity in u give an upper crossing for every u. Continuity and strict signal monotonicity yield a finite unique solution. The talker cdf is strictly increasing on $(-\infty, \bar{b}]$ because β^T is continuous and strictly increasing and F_u has full support. This proves the outbid probabilities, including zero probability at equality because $H(\bar{b}) = 1$. Monotonicity in u gives the ordering of the entry margins. Participation and rents follow from Lemma A12. □

A.2.9 Scalar benchmark bidding by silent dealers

The calculations in this subsection concern the auxiliary auction in which all silent dealers have $u = 0$. Write $b^{S,\mathrm{bm}}$ for its symmetric silent strategy. Within this calculation, H, h, β^T, and $\bar{b}$ refer to the auxiliary talker side. The posterior objects G, g, ψ, and φ depend on the signal and

disclosure model and retain their earlier definitions. We require $n_S \geq 2$ for the rival-boundary density formulas.

If a silent dealer with true signal s bids as pseudo-type x, they clear silent rivals precisely when $S_{(1)} < x$. Thus

$$\Pr(\text{win} \mid s, \mathcal{I}, x) = H(b^{S,\text{bm}}(x) \mid \mathcal{I}) G(x \mid s, \mathcal{I}). \tag{A.47}$$

For the highest rival silent bid $B^-_{(1)} = b^{S,\text{bm}}(S_{(1)})$, differentiating its cdf at $x = (b^{S,\text{bm}})^{-1}(b)$ gives

$$f_{B^-_{(1)}}(b \mid s, \mathcal{I}) = \frac{g(x \mid s, \mathcal{I})}{(b^{S,\text{bm}})'(x; \mathcal{I})}. \tag{A.48}$$

These are exact formulas for the auxiliary auction. In the full game, even a focal dealer with $u = 0$ faces the integrated rival cutoffs in (A.22).

Define the value in the state with a rival at the signal boundary by $\vartheta(s; \mathcal{I}) = \mathbb{E}[v \mid s, \mathcal{I}, S_{(1)} = s]$. This differs from $\varphi(s; \mathcal{I})$, which conditions on all rival signals being below s.

Proposition A3 (Silent dealer first-order condition)**.** *In the auxiliary scalar auction with $n_S \geq 2$, suppose the silent strategy is continuously differentiable and strictly increasing. At an interior optimum above the talker support, if $G(s \mid s, \mathcal{I}) > 0$ and $g(s \mid s, \mathcal{I}) > 0$, then*

$$b^{S,\text{bm}}(s; \mathcal{I}) = \vartheta(s; \mathcal{I}) - \frac{(b^{S,\text{bm}})'(s; \mathcal{I})}{K(s \mid s, \mathcal{I})}. \tag{A.49}$$

Proof. At a bid above $\bar{b}$, the talker side is cleared with probability one. The payoff from pseudo-type x is $\int_{-\infty}^{x} [\mathbb{E}[v \mid s, \mathcal{I}, S_{(1)} = t] - b^{S,\text{bm}}(x)] g(t \mid s, \mathcal{I})\, dt$. Leibniz's rule gives derivative $[\mathbb{E}[v \mid s, \mathcal{I}, S_{(1)} = x] - b^{S,\text{bm}}(x)] g(x \mid s, \mathcal{I}) - (b^{S,\text{bm}})'(x) G(x \mid s, \mathcal{I})$. Set $x = s$, impose stationarity, and divide by $g > 0$. □

The bid equals the value at the rival boundary less first-price shading, with competitive intensity measured by the reverse hazard K. To compare the boundary value with the Gaussian talk-based posterior, define

$$\widetilde{\Lambda}_v(s; \mathcal{I}) \equiv \mathbb{E}[v \mid s, \mathcal{M}] - \vartheta(s; \mathcal{I}). \tag{A.50}$$

Writing $\widetilde{\mu}_S(s) = \mathbb{E}[v \mid s, \mathcal{I}]$, the decomposition is $\widetilde{\Lambda}_v = (\widetilde{\mu}_S - \vartheta) - (\widetilde{\mu}_S - \mu_S)$. The first term is the signed change associated with conditioning on the rival boundary. It need not be nonnegative. If p is nonincreasing, the second term is nonnegative because rivals' silence is favorable news. No fixed sign follows for their difference.

Since $\mu_S(s, \mathcal{M}) = \mu_S(0, \mathcal{M}) + \gamma' s$ by Lemma A1,

$$b^{S,\text{bm}}(s) = \mu_S(0, \mathcal{M}) + \gamma' s - \widetilde{\Lambda}_v(s) - \frac{(b^{S,\text{bm}})'(s)}{K(s \mid s)}. \tag{A.51}$$

We suppress $\mathcal{I}$ in the remaining scalar formulas. Fix a finite lower endpoint s_0 of an upper branch on which $K(s \mid s)$ is positive and continuous.

Lemma A14 (Monotonicity of the benchmark silent bid)**.** *Suppose $\widetilde{\Lambda}_v$ is continuously differentiable on this branch and*

$$\gamma' - \partial_s \widetilde{\Lambda}_v(s) > 0. \tag{A.52}$$

If the boundary value satisfies

$$b^{S,\text{bm}}(s_0) < \vartheta(s_0) = \mu_S(0, \mathcal{M}) + \gamma' s_0 - \widetilde{\Lambda}_v(s_0), \tag{A.53}$$

then every continuously differentiable solution of (A.51) *with that boundary value is strictly increasing on the branch.*

Proof. Let $D_S = \vartheta - b^{S,\mathrm{bm}}$. The first-order condition gives $(b^{S,\mathrm{bm}})' = KD_S$ and $D_S' = \vartheta' - KD_S$. If D_S had a first zero after s_0, its derivative there would be nonpositive because it approaches zero from above. The differential equation instead gives $D_S' = \vartheta' > 0$. Thus $D_S > 0$, and $K > 0$ implies a strictly positive bid slope. □

The linear form of the upper-branch equation is

$$(b^{S,\mathrm{bm}})'(s) + K(s \mid s)b^{S,\mathrm{bm}}(s) = K(s \mid s)\vartheta(s). \tag{A.54}$$

Define its integrating factor by

$$I(s \mid \mathcal{I}) = \exp\left\{\int_{s_0}^{s} K(t \mid t, \mathcal{I})\, dt\right\}. \tag{A.55}$$

Then $I(s_0) = 1$ and $I' = IK$. Multiplying the linear equation by I and integrating gives

$$b^{S,\mathrm{bm}}(s) = I(s)^{-1}\left[b^{S,\mathrm{bm}}(s_0) + \int_{s_0}^{s} I(x)K(x \mid x)\vartheta(x)\, dx\right]. \tag{A.56}$$

For $s > s_0$, the weights can be normalized as

$$h_s(x \mid \mathcal{I}) = \frac{I(x \mid \mathcal{I})K(x \mid x, \mathcal{I})}{I(s \mid \mathcal{I}) - 1}\mathbf{1}\{s_0 < x < s\}. \tag{A.57}$$

This is a density because $\int_{s_0}^{s} I(x)K(x \mid x)\, dx = I(s) - 1$. Hence

$$b^{S,\mathrm{bm}}(s) = I(s)^{-1}b^{S,\mathrm{bm}}(s_0) + (1 - I(s)^{-1})\mathbb{E}_X^{h_s}[\vartheta(X)]. \tag{A.58}$$

The solution averages the boundary bid and the boundary-state valuations encountered along the branch. The density h_s is generated by the first-order condition. It is not the literal conditional distribution of a rival signal.

A.2.10 The talker side

Return first to the exact game and fix the full silent strategy. If rival talkers use β^T, define $y(b) = \beta^{T,-1}(b)$ on its image, with the cdf extended as zero or one outside that image. The talker payoff is $D_T(b)[\Psi^{\mathrm{ex}}(b) + u - b]$, where $D_T(b) = F_u(y(b))^{n_T-1}Q_T(b)$. At a differentiable interior optimum with $D_T(b) > 0$,

$$[\Psi^{\mathrm{ex}}(b) + u - b]\,\partial_b \log D_T(b) = 1 - \partial_b \Psi^{\mathrm{ex}}(b). \tag{A.59}$$

For $n_T = 1$, the rival-talker factor is one. A talker faces no information change from defeating rival talkers, but changing the probability of clearing silence changes the common-value assessment.

In the auxiliary scalar auction only, let $x(b) = (b^{S,\mathrm{bm}})^{-1}(b)$. Then $Q_T(b) = G_0(x(b))$ and $\Psi^{\mathrm{ex}}(b) = \psi(x(b))$ for that auction. Write $D_{\mathrm{bm}}(b) = F_u(y(b))^{n_T-1}G_0(x(b))$. The exact talker condi-

tion specializes to

$$[\psi(x(b)) + u - b]\,\partial_b \log D_{\mathrm{bm}}(b) = 1 - \psi'(x(b))x'(b). \tag{A.60}$$

The log derivative has two terms. Rival talkers contribute $(n_T - 1)f_u(y(b))/(F_u(y(b))\beta^{T\prime}(y(b)))$, and silent rivals contribute $K_0(x(b))x'(b)$, where $x'(b) = 1/(b^{S,\mathrm{bm}})'(x(b))$. This specialization does not follow from setting only the focal silent dealer's u to zero.

In either game, strictly increasing talker bids give the highest-talker-bid cdf

$$H(b \mid \mathcal{I}) = F_u\big(\beta^{T,-1}(b;\mathcal{I})\big)^{n_T}. \tag{A.61}$$

In the exact weighted model with $n_T \geq 1$ and $n_S \geq 2$ under the hypotheses of Proposition A2, that image is $(-\infty, \bar{b}]$. The randomness in H comes from idiosyncratic values under pure bidding strategies.

A.2.11 Scalar bidding inside the talker support

Continue with the auxiliary auction and suppress public information. A zero-value silent dealer with true signal s choosing pseudo-type x has payoff

$$\Pi_{\mathrm{bm}}(x;s) = H(b^{S,\mathrm{bm}}(x)) \int_{-\infty}^{x} [\mathbb{E}[v \mid s, \mathcal{I}, S_{(1)} = t] - b^{S,\mathrm{bm}}(x)]g(t \mid s)\,dt. \tag{A.62}$$

At $x = s$, the integral equals $G(s \mid s)[\varphi(s) - b^{S,\mathrm{bm}}(s)]$. Differentiating by the product rule and Leibniz's rule, then dividing by $H(b^{S,\mathrm{bm}}(s))G(s \mid s) > 0$, gives

$$\frac{h(b^{S,\mathrm{bm}})}{H(b^{S,\mathrm{bm}})}(b^{S,\mathrm{bm}})'(\varphi - b^{S,\mathrm{bm}}) + K(\vartheta - b^{S,\mathrm{bm}}) - (b^{S,\mathrm{bm}})' = 0. \tag{A.63}$$

All scalar functions in this equation are evaluated at s, with $K = K(s \mid s)$. Solving for the slope gives

$$(b^{S,\mathrm{bm}})'(s) = \frac{K(s \mid s)[\vartheta(s) - b^{S,\mathrm{bm}}(s)]}{1 - \dfrac{h(b^{S,\mathrm{bm}}(s))}{H(b^{S,\mathrm{bm}}(s))}[\varphi(s) - b^{S,\mathrm{bm}}(s)]}. \tag{A.64}$$

The denominator is maintained positive on the branch where this equation is used. Above the talker support, $H = 1$ and $h = 0$, recovering the upper-branch equation. A downward slope kink at the crossing follows if the one-sided coefficients have finite limits, $h(\bar{b}^-) > 0$, and both $\varphi(s) - \bar{b}$ and $\vartheta(s) - \bar{b}$ are strictly positive there. Positivity of the density alone does not establish a strict kink.

The auxiliary upper branch starts at its own crossing of $\bar{b}$, with boundary bid $\bar{b} = \beta^T(\bar{u})$. The two scalar silent equations, the auxiliary talker condition, and appropriate boundary conditions characterize stationary auxiliary schedules. The exact continuation game instead uses (A.26) and (A.59), coupled through the full silent strategy and the clearing likelihood C. None of these first-order conditions alone establishes global optimality.

Remark A1 (Randomization in the pure common-value limit)**.** *When $\bar{u} = 0$, the scalar silent formulation applies directly. Lemma A5 excludes a deterministic common talker bid with positive winning probability. Atomless mixing by itself requires a common payoff over the mixing support, but it does not by itself establish that the common payoff is zero.*

Consider a candidate with zero optimized talker payoff, positive clearing probability, and a positive density of the highest talker bid on an overlap interval. Zero payoff implies $b^S(x) = \psi(x)$ there. Since $\psi' = K_0(\varphi - \psi)$, substitution in (A.63) gives the necessary inequality

$$K(x \mid x)[\vartheta(x) - \varphi(x)] < [K_0(x) - K(x \mid x)][\varphi(x) - \psi(x)]. \tag{A.65}$$

Indeed, the omitted term is $(h/H)\psi'(\varphi - \psi) > 0$. Without a positive density, the corresponding restriction is weak. The inequality concerns this zero-payoff candidate class, not all possible mixed strategies.

The preceding necessary inequality does not exclude every randomized talker equilibrium.

A.2.12 The auxiliary hard-truncation boundary

The preceding baseline support result uses a silence weight bounded away from zero. The hard-truncation benchmark instead sets $w(s) = \mathbf{1}\{s > \bar{s}\}$ and has a finite lower silent-signal boundary $\bar{s}$. In this subsection, $b^S(s,u)$, its zero-value slice $b^S(s)$, β^T, and $s^\times$ refer to that hard-truncation continuation game. It is distinct from both the weighted baseline and the scalar restriction on silent idiosyncratic values.

For this boundary calculation, set $d = \mu_T - \bar{s}$. Given a subset $\mathcal{S}'$ of the other silent dealers, write $B_{\mathcal{S}'}(b) = \{b^S(s_k, u_k) < b \text{ for every } k \in \mathcal{S}'\}$. Conditional on their idiosyncratic values, this event imposes upper signal cutoffs. Let $\underline{b}_H$ denote a proposed common finite lower bid limit. The following restrictions are used only in results invoking this hard-truncation boundary.

Assumption A2 (Overbid envelope and density regularity)**.** *There is $\delta_0 > 0$ such that for every compact set D of values of d, there are constants $0 < a_D \le A_D < \infty$ and $B_D < \infty$, uniform across the configurations under comparison, with the following properties.*

(i) For $s \in (\bar{s}, \bar{s} + \delta_0]$,

$$a_D(s - \bar{s}) \le b^S(s;\mathcal{I}) - \psi(s;\mathcal{I}) \le A_D(s - \bar{s}), \tag{A.66}$$

and $b^S(s;\mathcal{I}) - \psi(s;\mathcal{I})$ is nondecreasing on $(\bar{s}, s^\times(\mathcal{I})]$.

(ii) The conditional density of a silent signal is bounded above by B_D on $[\bar{s}, \bar{s} + \delta_0]$, and the probabilities of the required silence events are bounded away from zero. These bounds also hold after conditioning on any $B_{\mathcal{S}'}(b)$ whose induced rival cutoffs lie in this interval, whenever that event has positive probability.

(iii) The finite limits satisfy $\lim_{s \downarrow \bar{s}} b^S(s;\mathcal{I}) = \lim_{u \downarrow 0} \beta^T(u;\mathcal{I}) = \underline{b}_H(\mathcal{I})$, with $\bar{b}(\mathcal{I}) > \underline{b}_H(\mathcal{I})$ for $\bar{u} > 0$.

The first condition imposes linear vanishing of the overbid at a finite signal boundary. The weighted baseline instead has the asymptotic compatibility condition (A.44). The hard-truncation assumptions do not transfer that finite boundary to the weighted model.

Remark A2 (Boundary configurations)**.** *Return to the exact continuation game. If $n_S = 1$, there is no rival silent signal, so $G^{\text{ex}} \equiv 1$. Let $c = \mathbb{E}[v \mid s, \mathcal{I}] + u$. The single silent dealer maximizes $H(b)(c - b)$, with interior first-order condition $h(b)(c - b) = H(b)$. The objective has increasing differences in (b, c). On an interval where H is strictly increasing, adding the optimality inequalities at two values of c orders their maximizing bids. More generally, extremal selections are nondecreasing whenever the corresponding extrema exist. No optimal bid exceeds $\bar{b}$, because $H = 1$ there and a higher bid only raises payment.*

A finite signal threshold at which the bid reaches $\bar{b}$ requires an additional condition. For an atomless H with $H(\bar{b}) = 1$ and $H(b) < 1$ for $b < \bar{b}$, define

$$c_{\text{cap}} = \bar{b} + \sup_{b<\bar{b}} \frac{H(b)(\bar{b} - b)}{1 - H(b)}, \tag{A.67}$$

allowing $c_{\text{cap}} = \infty$. Comparison with every $b < \bar{b}$ shows that $\bar{b}$ is optimal exactly when $c \geq c_{\text{cap}}$. If this bound is finite and c exceeds it, the bid is uniquely $\bar{b}$. If it is infinite, the bid does not reach $\bar{b}$ at any finite valuation. Thus a finite plateau cannot be inferred from monotonicity alone.

If $n_S = 0$, the continuation game is a symmetric first-price private-value auction in the u_j, with the public posterior common-value mean added to each valuation. If $n_T = 0$, the game has silent types (s_j, u_j) with common and idiosyncratic value. Sections A.2.14–A.2.16 give separate results for $N = 2$. The mixed-history support formulas do not apply to these boundary cases.

A.2.13 Selection within the continuation auction

Fix public information and the full silent strategy. The exact objects Q_T and Ψ^{ex} are defined in (A.24). If rival talkers use a strictly increasing schedule β_-, a talker of type u obtains

$$U^T(u, b; \beta_-) = F_u(\beta_-^{-1}(b))^{n_T - 1} Q_T(b)[\Psi^{\text{ex}}(b) + u - b], \tag{A.68}$$

with the cdf extended outside the range of β_-. For $n_T = 1$, the rival-talker factor is one.

For $n_T \geq 2$, a stationary candidate is a continuous strictly increasing schedule on $(0, \bar{u}]$, continuously differentiable with positive slope on its interior, that tends to $-\infty$ as $u \downarrow 0$ and satisfies the exact talker first-order condition at every interior assigned bid. Admissibility here requires $Q_T(\beta(u)) > 0$ and $0 \leq \Psi^{\text{ex}}(\beta(u)) + u - \beta(u) \leq u$. In particular, the assigned payoff vanishes at the lower endpoint. These restrictions are part of the candidate class and do not follow from stationarity alone.

Solving the exact talker condition for the type derivative gives the slope field

$$\mathcal{F}(u, b) = \frac{(n_T - 1)\dfrac{f_u(u)}{F_u(u)} Q_T(b)[\Psi^{\text{ex}}(b) + u - b]}{Q_T(b)[1 - \partial_b \Psi^{\text{ex}}(b)] - Q_T'(b)[\Psi^{\text{ex}}(b) + u - b]}. \tag{A.69}$$

Lemma A15 (Ordering of talker stationary candidates)**.** *Let $n_T \geq 2$. Suppose Q_T and Ψ^{ex} are continuously differentiable on the relevant bid range. Suppose the denominator in (A.69) is positive around every candidate graph and $\mathcal{F}$ is continuous and locally Lipschitz in b, locally uniformly in u, including a one-sided neighborhood of $\bar{u}$. Then distinct stationary candidates are strictly ordered on $(0, \bar{u}]$.*

Proof. Every candidate satisfies $\beta'(u) = \mathcal{F}(u, \beta(u))$ in the interior. If two candidates meet, uniqueness of the initial-value problem would make them coincide locally and then throughout their common connected domain. The same backward uniqueness applies to contact at $\bar{u}$. Distinct candidates therefore never meet. Continuity gives one strict sign for their difference. □

Proposition A4 (Talker selection at fixed silent behavior)**.** *Fix $n_T \geq 2$ and silent behavior. Suppose Q_T is strictly increasing on the relevant bid range, Q_T and Ψ^{ex} are continuously differentiable, and the slope field (A.69) has a positive denominator and the continuity and local Lipschitz properties stated above. Then at most one admissible stationary candidate can be a globally optimal continuation-equilibrium talker schedule.*

Proof. By the ordering argument, two distinct candidates can be labeled $\beta_1 < \beta_2$. Their assigned winning probabilities are $D_i(u) = F_u(u)^{n_T-1} Q_T(\beta_i(u))$, so $D_1(u) < D_2(u)$ for every positive interior type. Let $\Pi_i(u) = U^T(u, \beta_i(u); \beta_i)$. The chain rule and stationarity give $\Pi_i' = D_i$ in the interior, with rival schedules held fixed. Admissibility gives $\Pi_i(u) \to 0$ as $u \downarrow 0$, so $\Pi_i(u) = \int_0^u D_i(t)\, dt$ and $\Pi_2(u) > \Pi_1(u)$ for every $u > 0$.

Now hold rival talkers at β_1 and let type u bid $\beta_2(u)$. This higher bid clears at least the fraction $F_u(u)^{n_T-1}$ of rival talkers that it clears under symmetric use of β_2. The silent-side payoff factor is unchanged, and it is nonnegative by admissibility. The deviation therefore gives at least $\Pi_2(u) > \Pi_1(u)$. Thus the lower candidate cannot be globally optimal, ruling out two equilibrium candidates. □

When $n_T = 1$, the talker maximizes $Q_T(b)[\Psi^{\text{ex}}(b) + u - b]$ pointwise. The argument below applies to positive types. At $u = 0$, the value has a continuous extension, but a finite maximizing bid need not exist.

Proposition A5 (One-talker selection)**.** *Fix $n_T = 1$ and silent behavior. Suppose Q_T is strictly increasing on an interval containing all maximizers and the pointwise argmax is nonempty and compact for every $u \in (0, \bar{u}]$. Then the argmax sets are ordered in u, multiple maximizers occur for at most countably many positive types, and any two selections agree F_u-almost everywhere. Among continuous schedules on $(0, \bar{u}]$, the talker response is pointwise unique.*

Proof. Take $u_1 < u_2$ and maximizing bids b_1, b_2. Adding their optimality inequalities gives $(u_2 - u_1)[Q_T(b_2) - Q_T(b_1)] \geq 0$, hence $b_2 \geq b_1$. Thus the intervals spanned by the argmax sets have disjoint interiors. Each nondegenerate interval contains a different rational number, so there are at most countably many types with multiple maximizers. Atomlessness gives almost-everywhere agreement. Continuous selections that agree off a countable set agree everywhere on $(0, \bar{u}]$. □

For a silent dealer with fixed signal s, let $\mathcal{A}_b$ be the event that bid b clears all rivals. Define $Q_S(b \mid s, \mathcal{I}) = \Pr(\mathcal{A}_b \mid s, \mathcal{I})$ and $V_S(b, s; \mathcal{I}) = \mathbb{E}[v \mid s, \mathcal{I}, \mathcal{A}_b]$ where $Q_S > 0$. Under the maintained atomlessness, $Q_S = HG^{\text{ex}}$ and $V_S = V^{\text{ex}}$. Their exact payoff is

$$U^S(b, s, u) = Q_S(b \mid s, \mathcal{I})[V_S(b, s; \mathcal{I}) + u - b], \tag{A.70}$$

with value zero for a bid that never wins. For fixed s, its dependence on u is affine, with slope Q_S.

Proposition A6 (Pointwise uniqueness of the continuous silent response)**.** *Fix a continuation environment and signal s. Suppose the global argmax of (A.70) is nonempty and compact for every $u \in [0, \bar{u}]$, and $Q_S(b \mid s, \mathcal{I})$ is strictly increasing on an interval containing those maximizers. Then the argmax sets are ordered in u, multiple maximizers occur for at most countably many u, and any two selections agree F_u-almost everywhere. Among continuous schedules, the response is pointwise unique.*

Proof. For $u_1 < u_2$ and maximizing bids b_1, b_2, the sum of the optimality inequalities is $(u_2 - u_1)[Q_S(b_2 \mid s, \mathcal{I}) - Q_S(b_1 \mid s, \mathcal{I})] \geq 0$. Strict increase of Q_S gives $b_2 \geq b_1$. The disjoint-interval argument in the preceding proof gives countability and almost-everywhere agreement. Continuity extends agreement to every u. □

Corollary A3 (Reduction of continuation selection)**.** *Fix a mixed configuration in the maintained continuous class. Suppose every fixed silent environment under consideration satisfies the talker-selection conditions appropriate to its number of talkers, and every fixed rival environment satisfies*

the compact-argmax and strictly increasing winning-probability conditions for the silent response. Then there is at most one continuous talker response conditional on silent behavior and at most one continuous silent response at each signal conditional on rival behavior. Any multiplicity of the joint continuation equilibrium must involve distinct self-consistent silent environments.

Proof. For $n_T \geq 2$, global optimality eliminates all but at most one admissible talker candidate. For $n_T = 1$, continuity gives a unique talker selection whenever it exists. The silent result gives a unique response at every fixed signal and environment. If two joint equilibria had identical silent behavior, their talker schedules would coincide as well. □

These results select responses within fixed continuation environments. Joint equilibrium existence and dependence on the disclosure policy remain separate requirements. The local contraction condition in Proposition A18 concerns selected exact continuation values and changes in rival disclosure configurations. The two-dealer candidate comparison below supplies one approximation result without assuming those requirements.

A.2.14 The two-dealer all-silent continuation

This subsection fixes $N = 2$ and $n_T = 0$. It uses the Gaussian primitives and real bid space of the baseline. It does not impose the mixed-history corridor or assume a positive-width continuation equilibrium. Write $V = \sigma_v^2$, $e = \sigma_\varepsilon^2$, $\widetilde{s} = s - \mu_0$, $\lambda = V/(V+e)$, $a = 1-\lambda$, $\gamma = V/(2V+e)$, $\sigma_*^2 = e(1+\lambda)$, and $\zeta = \lambda/\sigma_*^2$. Let ϕ_{σ_*} and Φ_{σ_*} be the density and cdf of $\mathcal{N}(0,\sigma_*^2)$. Fix $q_0 \in (0,1]$ and a continuous silence weight $w_p = 1-p$ satisfying $q_0 \leq w_p \leq 1$.

Conditional on the focal signal s and both dealers' silence, the rival signal t has density and cdf

$$\pi_p(t \mid s) = \frac{w_p(t)\phi_{\sigma_*}(t-\mu_0-\lambda\widetilde{s})}{\int_{\mathbb{R}} w_p(x)\phi_{\sigma_*}(x-\mu_0-\lambda\widetilde{s})\,dx}, \qquad \mathsf{F}_p(x \mid s) = \int_{-\infty}^{x} \pi_p(t \mid s)\,dt. \tag{A.71}$$

The focal dealer's own silence adds no information after conditioning on s. Conditional on both signals, the silence weights cancel and the common-value mean is $\theta_2(s,t) = \mu_0 + \gamma(\widetilde{s}+\widetilde{t})$, where $\widetilde{t} = t - \mu_0$. Define $\theta(s) = \theta_2(s,s)$, the diagonal reverse hazard $\rho_p(s) = \pi_p(s \mid s)/\mathsf{F}_p(s \mid s)$, and $\rho_0(s) = \phi_{\sigma_*}(a\widetilde{s})/\Phi_{\sigma_*}(a\widetilde{s})$.

The Gaussian kernel gives $q_0\rho_0 \leq \rho_p \leq q_0^{-1}\rho_0$. The function ρ_0 is decreasing and satisfies $\rho_0(s) \geq a(\mu_0 - s)/\sigma_*^2$ for $s \leq \mu_0$. Its upper-tail integrals obey $\int_{\mu_0}^{\infty} \rho_p \leq \log 2/(q_0 a)$ and $\int_{\mu_0}^{\infty}(t-\mu_0)\rho_p(t)\,dt \leq \sigma_*\sqrt{2/\pi}/(q_0 a^2)$. For the first bound, integrate the logarithmic derivative of Φ_{σ_*}. For the second, integrate by parts and use $-\log(1-y) \leq 2y$ for $0 \leq y \leq 1/2$.

Define the kernel, margin, and reference bid by

$$\mathcal{K}_p(x,s) = \exp\left(-\int_x^s \rho_p(t)\,dt\right), \qquad m_p(s) = 2\gamma\int_{-\infty}^{s} \mathcal{K}_p(x,s)\,dx, \qquad \beta_p^0(s) = \theta(s) - m_p(s). \tag{A.72}$$

The superscript zero denotes the auction in which both idiosyncratic values are zero, rather than the zero-value slice of a positive-width auction. Write $d_p(s) = \rho_p(s)m_p(s)$ and let κ_p^0 denote the reference bid's upper limit, whose existence and finiteness are established below.

Proposition A7 (Selected continuation at zero idiosyncratic-value width)**.** *For every continuous weight in the stated band, (A.72) is finite and continuously differentiable, with* $(\beta_p^0)'(s) = d_p(s) > 0$.

Its range is $(-\infty, \kappa_p^0)$ for a finite κ_p^0. It is a symmetric continuation equilibrium against all real-bid deviations, with strictly positive payoff at every finite signal and payoff tending to zero as $s \to -\infty$. It is the unique solution of $\beta' = \rho_p(\theta - \beta)$ satisfying $\theta(s) - \beta(s) \to 0$ as $s \to -\infty$.

Proof. For $x \le s \le \mu_0$, the lower hazard bound gives $\mathcal{K}_p(x,s) \le \exp\{-q_0 a[(\mu_0 - x)^2 - (\mu_0 - s)^2]/(2\sigma_*^2)\}$. For $x \le \mu_0 \le s$, it gives $\mathcal{K}_p(x,s) \le \exp[-q_0 a(\mu_0 - x)^2/(2\sigma_*^2)]$. Thus the integral is finite. Differentiating it gives $m_p' = 2\gamma - \rho_p m_p$ and $d_p = \rho_p m_p > 0$. The lower bound also gives $m_p(s) = O((\mu_0 - s)^{-1})$ in the lower tail. Hence $\beta_p^0(s) \to -\infty$ there. For $s \ge \mu_0$, $m_p(s) \le m_p(\mu_0) + 2\gamma\widetilde{s}$. The two upper-tail integrals above then make d_p integrable, giving a finite upper limit κ_p^0.

For a rival cutoff x, put $R_p(x \mid s) = \mathsf{F}_p(x \mid s)/\pi_p(x \mid s)$ and $D_p(x;s) = x - \mathbb{E}_p[T \mid T \le x, S = s] > 0$. The exponential-family representation of (A.71) yields $\partial_s \log R_p(x \mid s) = -\zeta D_p(x;s) < 0$. Define $\Lambda_p(x,s) = \rho_p(x) R_p(x \mid s)$, so $\Lambda_p(x,x) = 1$. If the focal dealer bids $\beta_p^0(x)$, their payoff is $P_0(x;s) = \int_{-\infty}^{x} [\theta_2(s,t) - \beta_p^0(x)]\pi_p(t \mid s)\, dt$. Direct differentiation gives

$$\partial_x P_0(x;s) = \pi_p(x \mid s)\{\gamma(s - x) + m_p(x)[1 - \Lambda_p(x,s)]\}. \tag{A.73}$$

When $x < s$, both terms in braces are positive. When $x > s$, both are negative. The payoff tends to zero as $x \to -\infty$, because the bid has a linear lower tail and the rival signal has Gaussian tails. The unique maximum within the bid range is therefore positive and occurs at $x = s$. A bid at κ_p^0 wins surely and has payoff $\lim_{x\to\infty} P_0(x;s)$, strictly below the maximum. A higher bid wins surely at greater cost. This covers every real bid.

Integration by parts gives the equilibrium payoff $m_p(s)\mathsf{F}_p(s \mid s) - \gamma \int_{-\infty}^{s} \mathsf{F}_p(t \mid s)\, dt$. Both terms tend to zero in the lower tail by the margin bound and Gaussian tail integrability. Finally, two solutions of the linear bid equation differ by $C \exp(-\int_{s_0}^{s} \rho_p)$ for a fixed finite s_0. This factor diverges as $s \to -\infty$, so the stated margin normalization forces $C = 0$. □

The normalization selects the reference branch. The proposition does not claim uniqueness among all possible continuation equilibria without a boundary restriction.

We next compare reference bids across policies. For two weights let $\delta = \operatorname{osc} \log(w_{p_2}/w_{p_1})$, where $\operatorname{osc} f = \sup f - \inf f$, and define

$$\mathcal{D}_{q_0} = \sigma_* \left[\frac{2}{\mathrm{e}} \sqrt{\frac{\pi}{q_0 a}} + \frac{\log 2}{q_0 a} \sqrt{\frac{\pi}{2 q_0 a}} + \frac{\sqrt{2/\pi}}{q_0 a^2} \right]. \tag{A.74}$$

Here e is the base of the exponential function.

Lemma A16 (Reference policy comparison)**.** *The reference schedules satisfy $\|\beta_{p_2}^0 - \beta_{p_1}^0\|_\infty \le 2\gamma \mathcal{D}_{q_0} \delta$, with $\delta \le 2\|p_2 - p_1\|_\infty / q_0$.*

Proof. Use the geometric interpolation $w_\tau = w_{p_1}^{1-\tau} w_{p_2}^{\tau}$ and put $\ell = \log(w_{p_2}/w_{p_1})$. Differentiating the diagonal hazard gives $\partial_\tau \log \rho_\tau(s) = \ell(s) - \mathbb{E}_\tau[\ell(T) \mid T \le s, S = s]$, whose absolute value is at most δ. With $A_\tau(x,s) = \int_x^s \rho_\tau$, differentiation of (A.72) therefore bounds $|\partial_\tau \beta_\tau^0(s)|$ by $2\gamma\delta J_\tau(s)$, where $J_\tau(s) = \int_{-\infty}^{s} A_\tau(x,s) e^{-A_\tau(x,s)}\, dx$.

If $s \le \mu_0$, write $x = s - z$. The hazard bound gives $A_\tau(x,s) \ge q_0 a z^2/(2\sigma_*^2)$. Using $y e^{-y} \le (2/\mathrm{e}) e^{-y/2}$ bounds J_τ by the first term of (A.74). If $s \ge \mu_0$, split the integral at μ_0. For $x \le \mu_0$, write $A_\tau(x,s) = A_\tau(x,\mu_0) + A_\tau(\mu_0,s)$ and use $A_\tau(\mu_0,s) \le \log 2/(q_0 a)$. Gaussian integration bounds this piece by the first two terms of (A.74). On $[\mu_0, s]$, drop e^{-A_τ} and apply Fubini to obtain the bound $\int_{\mu_0}^{s} (t - \mu_0)\rho_\tau(t)\, dt \le \sigma_* \sqrt{2/\pi}/(q_0 a^2)$. Thus $J_\tau \le \mathcal{D}_{q_0}$. Integrating over τ proves the bid bound. The last inequality follows from the q_0^{-1} Lipschitz bound for the logarithm on $[q_0, 1]$. □

A.2.15 A small-width all-silent candidate

Write $u = \epsilon z$, with $\epsilon = \bar{u}$, $z \sim F_0$ on $[0,1]$, and $\nu = \mathbb{E}z$. Hold the disclosure policy fixed when varying ϵ. For uniform comparisons, use the class of continuously differentiable weights satisfying $q_0 \le w_p \le 1$ and $|\partial_s \log w_p| \le L_0$, where q_0 and $L_0 < \infty$ are fixed. This class is closed under geometric interpolation.

Let $D_p(s) = D_p(s;s)$ and define the positive coefficient $\mathcal{H}_p(s) = \gamma + \zeta m_p(s) D_p(s)$ and the pass-through coefficient $\chi_p(s) = d_p(s)/\mathcal{H}_p(s)$. The candidate is

$$\widehat{\beta}^p_\epsilon(s,z) = \beta^0_p(s) + \epsilon[\nu + \chi_p(s)(z-\nu)]. \tag{A.75}$$

The common mean shift is $\epsilon\nu$. The response to a idiosyncratic value relative to its mean depends on the signal through χ_p. Define $c_0 = 1 + q_0^{-2}$, $d_* = 2\gamma q_0^{-2}$, $\chi_* = 2q_0^{-2}$, and

$$C_* = \sup_{p,s} \frac{|\chi_p'(s)|}{d_p(s)}, \qquad T_* = \sup_{p,s} \left| \chi_p(s) + m_p(s) \frac{\chi_p'(s)}{d_p(s)} \right|, \qquad G_* = \sup_{p,s} \frac{|(\log d_p)'(s)|}{\mathcal{H}_p(s)}. \tag{A.76}$$

The suprema are over the stated smooth weight class.

Lemma A17 (Bounds for the candidate coefficients)**.** *The constants in (A.76) are finite. Uniformly over the weight class, $0 < d_p \le d_*$, $0 < \chi_p \le \chi_*$, and $\chi_p(s) \to 2$ as $s \to -\infty$, while $\chi_p(s) \to 0$ as $s \to \infty$. Along geometric policy interpolation, $|\partial_\tau \log d_\tau| \le c_0\delta$ and $|\partial_\tau \log \chi_\tau| \le c_0\delta$. Consequently,*

$$\|\chi_{p_2} - \chi_{p_1}\|_\infty \le \chi_* c_0 \delta, \qquad \sup_{s,z} |\widehat{\beta}^{p_2}_\epsilon(s,z) - \widehat{\beta}^{p_1}_\epsilon(s,z)| \le (2\gamma \mathcal{D}_{q_0} + \epsilon \chi_* c_0)\delta. \tag{A.77}$$

Proof. Since ρ_0 is decreasing, $\int_x^s \rho_p \ge q_0 \rho_0(s)(s-x)$. Hence $m_p(s) \le 2\gamma/[q_0\rho_0(s)]$, giving $d_p \le 2\gamma q_0^{-2}$. The bound for χ_p follows from $\mathcal{H}_p \ge \gamma$.

For the tail estimates, set $h = a(\mu_0 - s)/\sigma_*^2$ when $s < \mu_0$. Conditional on $T \le s$ and $S = s$, the gap $y = s - T$ has density proportional to $w_p(s-y)\exp[-hy - y^2/(2\sigma_*^2)]$. Substitute $y = x/h$. The bound $|w_p(s-y) - w_p(s)|/w_p(s) \le L_0 y/q_0$ and the band bound give uniformly integrable exponential bounds for the first two moments. They yield $\rho_p = h + O(1)$, $D_p = h^{-1} + O(h^{-2})$, and $\mathrm{Var}_p(T \mid T \le s, S = s) = O(h^{-2})$. Applying the same rescaling to the integral for m_p, whose exponent is $hy + ay^2/(2\sigma_*^2) + O(y)$, gives $m_p = 2\gamma/h + O(h^{-2})$ and $d_p = 2\gamma + O(h^{-1})$.

Differentiating the diagonal conditional moments gives

$$\frac{\rho_p'}{\rho_p} = \frac{w_p'}{w_p} - \frac{a\widetilde{s}}{\sigma_*^2} - \rho_p + \zeta D_p, \qquad D_p' = 1 - \rho_p D_p - \zeta \,\mathrm{Var}_p(T \mid T \le s, S = s). \tag{A.78}$$

Together with $m_p' = 2\gamma - d_p$, the lower-tail estimates imply $(\log d_p)' = O(1)$ and $\mathcal{H}_p' = O(h^{-2})$. Since

$$\frac{\chi_p'}{d_p} = \frac{(\log d_p)'}{\mathcal{H}_p} - \frac{\mathcal{H}_p'}{\mathcal{H}_p^2}, \tag{A.79}$$

all three constants are bounded in the lower tail, and $\chi_p \to 2$.

In the upper tail, Lemma A16 bounds β^0_p uniformly by comparison with the constant weight. Thus $m_p = 2\gamma\widetilde{s} + O(1)$. Band bounds on the centered truncated Gaussian moments give $D_p = a\widetilde{s} + O(1)$ and bounded conditional variance. Equations (A.78) imply $(\log d_p)' = O(\widetilde{s})$, while $\mathcal{H}_p$ is of order $\widetilde{s}^2$ and $\mathcal{H}_p' = O(\widetilde{s})$. Equation (A.79) is therefore $O(1/\widetilde{s})$. Since $m_p = O(\widetilde{s})$, all three constants are bounded there. The hazard decays at a Gaussian rate, so $d_p \to 0$ and $\chi_p \to 0$. On

compact intervals, the band provides positive uniform lower bounds for ρ_p and m_p, and bounded conditional moments give the remaining bounds.

For the policy derivatives, the proof of Lemma A16 gives $|\partial_\tau \log \rho_\tau| \le \delta$. Also $|\partial_\tau \log m_\tau| \le \delta \int Ae^{-A}dx / \int e^{-A}dx$. Substitute $A = \int_x^s \rho_\tau$ in this ratio. The resulting density is proportional to $e^{-A}/\rho_\tau(x(A))$. The factor $1/\rho_0(x(A))$ is nonincreasing in A, so tilting an exponential density by this factor weakly reduces its mean from one. The ratio ρ_0/ρ_τ lies in $[q_0, q_0^{-1}]$, which bounds the displayed mean by q_0^{-2}. Thus $|\partial_\tau \log m_\tau| \le q_0^{-2}\delta$ and $|\partial_\tau \log d_\tau| \le c_0\delta$.

The gap D_τ is a ratio of two integrals of the weight against nonnegative kernels. Its logarithmic policy derivative is the difference of two averages of ℓ, hence has absolute value at most δ. Writing $\chi_\tau = \rho_\tau/[\gamma/m_\tau + \zeta D_\tau]$, the logarithmic derivative of its positive denominator is a weighted average of the two termwise logarithmic derivatives. This gives $|\partial_\tau \log \chi_\tau| \le c_0\delta$. Integrating $|\partial_\tau \chi_\tau| \le \chi_* c_0 \delta$ proves the first comparison, and Lemma A16 gives the second. □

Choose $\epsilon_0 > 0$ so that $\theta_* := \epsilon_0 C_* < 1$, and set $a_* = 1-\theta_*$, $R_* = \exp(\epsilon_0 G_*/a_*)$, $j_- = 1/[R_*(1+\theta_*)]$, and $j_+ = R_*/a_*$. The candidate has signal derivative at least $a_* d_p > 0$ and idiosyncratic-type derivative $\epsilon\chi_p > 0$ for $0 < \epsilon \le \epsilon_0$. Its range at every idiosyncratic type is $(-\infty, \kappa_p^0 + \epsilon\nu)$. It also satisfies a common lower bound $\widehat{\beta}_\epsilon^p(t,z) \ge -C(1+|t|)$, because the perturbation is bounded and the reference has a linear lower tail.

For a centered rival idiosyncratic type $r = z - \nu$, define the cutoff $c_p(x,r)$ at the bid $\epsilon\nu + \beta_p^0(x)$ by

$$\beta_p^0(c_p(x,r)) + \epsilon\chi_p(c_p(x,r))r = \beta_p^0(x). \tag{A.80}$$

Every finite x has a unique cutoff because each candidate schedule has the same range. Write $a_p(t) = \chi_p'(t)/d_p(t)$.

Lemma A18 (Candidate cutoff comparisons)**.** *For $0 \le \epsilon \le \epsilon_0$,*

$$|c_p(x,r) - x| \le \frac{\epsilon|r|}{\gamma a_*}, \qquad j_- \le \partial_x c_p(x,r) \le j_+, \qquad \sup_{x,r} |c_{p_2}(x,r) - c_{p_1}(x,r)| \le h_\epsilon \delta, \tag{A.81}$$

where $h_\epsilon = 2c_0\epsilon/(\gamma a_*)$.

Proof. Implicit differentiation gives $\partial_r c_p = -\epsilon/[\mathcal{H}_p(c_p)(1+\epsilon a_p(c_p)r)]$ and $\partial_x c_p = d_p(x)/[d_p(c_p)(1+\epsilon a_p(c_p)r)]$. Integrating the first formula from $r = 0$, where $c_p(x,0) = x$, gives the displacement bound. Along this path $|\partial_r \log d_p(c_p)| \le \epsilon G_*/a_*$. Consequently $R_*^{-1} \le d_p(c_p)/d_p(x) \le R_*$, which proves the derivative bounds.

Along geometric interpolation, write the cutoff equation as $\int_x^{c_\tau} d_\tau(t)\,dt + \epsilon\chi_\tau(c_\tau)r = 0$. Holding the endpoint fixed, the policy derivative of each term has absolute value at most $c_0\delta\epsilon|r|\chi_\tau(c_\tau)$. Dividing their sum by the endpoint derivative $d_\tau(c_\tau)(1+\epsilon a_\tau(c_\tau)r)$ gives $|\partial_\tau c_\tau| \le 2c_0\delta\epsilon|r|/(\gamma a_*)$. Integration proves the policy comparison. □

A.2.16 Global approximation and policy stability

For a fixed policy, let $\mathcal{R}_\epsilon^p(s,z)$ be the supremum payoff over all real bids against the rival candidate, minus the payoff from using the candidate at (s,z). This is the actual auction with idiosyncratic values $u = \epsilon z$. Put $A_* = 1/(\gamma a_*)$, $B_* = A_* + (1+T_*)/\gamma$, and $B_1 = B_* + A_*$.

Proposition A8 (Global approximate optimality)**.** *For every policy in the stated smooth class, every finite signal, every $z \in [0,1]$, and $0 < \epsilon \le \epsilon_0$, the candidate-rival problem has a positive finite*

optimized payoff attained at a real bid. There is a bounded nonnegative function $\mathcal{G}(s)$, common to the weight class and tending to zero at both signal tails, such that

$$0 \le \mathcal{R}_\epsilon^p(s,z) \le 2B_*\epsilon^2\mathcal{G}(s). \tag{A.82}$$

For each fixed policy,

$$\sup_{s\in\mathbb{R},\ z\in[0,1]} \mathcal{R}_\epsilon^p(s,z) = o(\epsilon^2) \qquad as\ \epsilon \downarrow 0. \tag{A.83}$$

Proof. Center the focal idiosyncratic type at $r_0 = z - \nu$ and subtract $\epsilon\nu$ from bids and valuations. A bid with reference index y has payoff $P_\epsilon(y;s,r_0) = \mathbb{E}_r \int_{-\infty}^{c_p(y,r)} [\theta_2(s,t) + \epsilon r_0 - \beta_p^0(y)]\pi_p(t \mid s)\,dt$. The own candidate index y_c satisfies $\beta_p^0(y_c) = \beta_p^0(s) + \epsilon\chi_p(s)r_0$ and hence $c_p(y_c, r_0) = s$.

At $t = c_p(y,r)$, differentiate the payoff in y. Its integrand has the positive factor $d_p(y)\pi_p(t \mid s)/[d_p(t)(1+\epsilon a_p(t)r)]$ multiplied by

$$(1+\epsilon a_p(t)r)m_p(t)[1-\Lambda_p(t,s)] + \gamma(s-t) + \epsilon[r_0 - rT_p(t)], \qquad T_p = \chi_p + m_p a_p. \tag{A.84}$$

Indeed, direct differentiation first gives the bracket $m_p(t) + \gamma(s-t) + \epsilon r_0 - \epsilon\chi_p(t)r - d_p(t)(1+\epsilon a_p(t)r)R_p(t \mid s)$. Substituting $d_pR_p = m_p\Lambda_p$ gives (A.84).

If $y < s - B_*\epsilon$, the cutoff bound puts every t below $s - (1+T_*)\epsilon/\gamma$. Then $1 - \Lambda_p(t,s) > 0$, and the second term exceeds the absolute value of the last term. The payoff derivative is positive. If $y > s + B_*\epsilon$, the inequalities reverse and the derivative is negative. The payoff tends to zero as $y \to -\infty$ by the linear bid tail and Gaussian integrability. It is therefore positive somewhere and attains its maximum with $|y^* - s| \le B_*\epsilon$. The endpoint payoff is the limit as $y \to \infty$, below the maximum because the payoff is strictly decreasing above this interval. Bids above the endpoint win surely at greater cost. Thus the argument covers every real bid. The cutoff bound also gives $|y_c - s| \le A_*\epsilon$.

We bound the derivative between y_c and a best response. Put $K_1 = \epsilon_0 B_1$ and define

$$\mathcal{D}(s) = \frac{a(\widetilde{s})^+ + (1+\lambda)K_1 + \sigma_*\sqrt{2/\pi}}{q_0}, \qquad \mathcal{M}(s) = \frac{2\gamma}{q_0\rho_0(\mu_0)} + 2\gamma[(\widetilde{s})^+ + K_1]. \tag{A.85}$$

Also set

$$\mathcal{P}(s) = \frac{\exp\{-[(a|\widetilde{s}| - K_1)^+]^2/(2\sigma_*^2)\}}{q_0\sigma_*\sqrt{2\pi}}. \tag{A.86}$$

For $|y-s| \le B_*\epsilon$, all cutoffs satisfy $|t-s| \le B_1\epsilon \le K_1$. A normal variable with cutoff displacement v_0 from its mean has expected truncated gap at most $v_0^+ + \sigma_*\sqrt{2/\pi}$. This follows from the monotonicity of the gap below zero and the decreasing reverse hazard above zero. Tilting by $w_p \in [q_0, 1]$ increases this bound by at most q_0^{-1}. It follows that $D_p(t;x) \le \mathcal{D}(s)$ whenever x lies between s and t. The inequalities $m_p(t) \le 2\gamma/[q_0\rho_0(t)]$ for $t \le \mu_0$ and $m_p' \le 2\gamma$ give $m_p(t) \le \mathcal{M}(s)$. Finally, the band and Gaussian density give $\pi_p(t \mid s) \le \mathcal{P}(s)$.

Integrating $\partial_s \log\Lambda_p(t,s) = -\zeta D_p(t;s)$ from t to s gives $|\log\Lambda_p(t,s)| \le \zeta\epsilon B_1\mathcal{D}(s)$. Hence $|1-\Lambda_p(t,s)| \le \zeta\epsilon B_1\mathcal{D}(s)e^{\zeta K_1\mathcal{D}(s)}$. The cutoff derivative proof gives $d_p(y)/[d_p(t)(1+\epsilon a_p(t)r)] \le R_*/a_*$. Thus $|\partial_y P_\epsilon| \le \epsilon\mathcal{G}(s)$ on the stated interval, where

$$\mathcal{G}(s) = \frac{R_*}{a_*}\mathcal{P}(s)\left[(1+\theta_*)\mathcal{M}(s)\zeta B_1\mathcal{D}(s)e^{\zeta K_1\mathcal{D}(s)} + \gamma B_1 + 1 + T_*\right]. \tag{A.87}$$

The bracket is bounded in the lower tail and grows at most polynomially times an exponential

linear in $\widetilde{s}$ in the upper tail. The Gaussian factor makes $\mathcal{G}$ bounded and vanishing at both tails. Integrating the derivative over a distance at most $2B_*\epsilon$ gives (A.82).

For the sharper conclusion, fix a policy and a compact signal interval. All candidate and optimal indices are within $O(\epsilon)$ of s. At the own candidate index, the rival cutoff satisfies $t = s + \epsilon(r_0 - r)/\mathcal{H}_p(s) + o(\epsilon)$ uniformly on this compact set and in both idiosyncratic types. Since $\Lambda_p(t,t) = 1$ and $\partial_s \log \Lambda_p(t,s) = -\zeta D_p(t;s)$, its derivative in the first argument at (s,s) is $\zeta D_p(s)$. Expanding the bracket in (A.84) therefore gives $\epsilon r[1 - T_p(s)] + o(\epsilon)$. Its positive prefactor equals $\pi_p(s \mid s) + o(1)$. Averaging over the centered rival type, whose mean is zero, yields $\partial_y P_\epsilon(y_c; s, r_0) = o(\epsilon)$ uniformly on the compact set.

At zero width, differentiating (A.73) at $y = s$ gives $\partial_{yy} P_0(s;s) = -\pi_p(s \mid s)\mathcal{H}_p(s) < 0$, uniformly bounded away from zero on the compact interval. The continuously differentiable weight and the candidate cutoff formula give a locally uniform expansion $\partial_y P_\epsilon = \partial_y P_0 + \epsilon G + o(\epsilon)$ with continuous G. At a best response the derivative is zero. Comparing this equation with the candidate derivative, using the $O(\epsilon)$ localization and the strictly negative limiting curvature, gives $y^* - y_c = o(\epsilon)$ uniformly. The derivative throughout that interval is $O(\epsilon)$, so its integral is $o(\epsilon^2)$.

Given any $\eta > 0$, choose a compact signal interval outside which $2B_*\mathcal{G}(s) < \eta$. Inside it, the regret divided by ϵ^2 tends uniformly to zero by the preceding argument. Outside it, (A.82) bounds the ratio by η. Letting $\epsilon \downarrow 0$ and then $\eta \downarrow 0$ proves (A.83). Uniform little-o across the whole policy class would require additional uniform continuity control on the weight derivatives. The established class-uniform conclusion is (A.82). □

We finally hold beliefs fixed and compare the payoffs induced by two rival candidates. Let p_0 be any measurable belief policy whose silence weight lies in $[q_0, 1]$, independently of the two smooth strategy policies. Write $V_i(s,u)$ for the optimized payoff against $\widehat{\beta}_\epsilon^{p_i}$ under this fixed belief. Let $c_\star = 1/\sqrt{2\pi \mathrm{e}}$ and define

$$L_{\mathrm{cand}}(\epsilon) = 2\gamma \mathcal{D}_{q_0} + h_\epsilon \left[\frac{d_*}{j_-} + \frac{\gamma c_\star}{q_0} \left(1 + \frac{j_+}{j_-} \right) \right]. \tag{A.88}$$

Proposition A9 (Global candidate strategy comparison)**.** *For $0 \le \epsilon \le \epsilon_0$, every pair of smooth strategy weights in the stated class, every band belief p_0, and every finite focal idiosyncratic value u, the optimized payoffs are finite and*

$$\sup_{s \in \mathbb{R}} |V_2(s,u) - V_1(s,u)| \le L_{\mathrm{cand}}(\epsilon) \operatorname{osc} \log \frac{w_{p_2}}{w_{p_1}} \le \frac{2L_{\mathrm{cand}}(\epsilon)}{q_0} \|p_2 - p_1\|_\infty. \tag{A.89}$$

The same constant applies uniformly over the belief weights.

Proof. Fix s and a clearing probability $\alpha \in (0,1)$. Under candidate i, a bid $\epsilon\nu + \beta_{p_i}^0(x)$ clears with probability $A_i(x) = \mathbb{E}_r[\mathsf{F}_{p_0}(c_{p_i}(x,r) \mid s)]$. This function is continuous and strictly increasing from zero to one, so there is a unique index x_i giving α. The cutoff comparisons imply $|x_2 - x_1| \le h_\epsilon \delta / j_-$. Indeed, moving x_2 that distance in either direction from x_1 brackets every first-profile cutoff and hence its average clearing probability. They also imply $\sup_r |c_{p_2}(x_2,r) - c_{p_1}(x_1,r)| \le h_\epsilon \delta(1 + j_+/j_-)$.

Let $Q_i(\alpha \mid s) = \epsilon\nu + \beta_{p_i}^0(x_i)$ be the corresponding bid quantile. The reference comparison and $d_p \le d_*$ give $|Q_2 - Q_1| \le [2\gamma\mathcal{D}_{q_0} + d_* h_\epsilon / j_-]\delta$. Put $Z = T - \mu_0 - \lambda\widetilde{s}$ and let $L_i(\alpha \mid s)$ be its unnormalized mean on the event of clearing candidate i at this quantile. Since $|Z| \pi_{p_0}(T \mid s) \le c_\star / q_0$, integration between the matched cutoffs gives $|L_2 - L_1| \le (c_\star/q_0) h_\epsilon \delta (1 + j_+/j_-)$.

The identity $\gamma(1+\lambda) = \lambda$ makes the fixed-quantile payoff equal to $\alpha[\mu_0 + \lambda\widetilde{s} + u - Q_i(\alpha \mid s)] + \gamma L_i(\alpha \mid s)$. At common α, the first term's common-value and own-value components cancel.

The remaining difference is bounded by $L_{\text{cand}}(\epsilon)\delta$. Taking suprema preserves this bound. Full-clearing bids are covered by $\alpha \uparrow 1$, and the supremum is at least zero by the lower-bid argument.

Finiteness follows directly from the linear lower bound on the rival bid B_i. At any bid, the payoff is bounded above by $\mathbb{E}_{p_0}[(v + u - B_i)^+ \mid s]$, which is finite under the Gaussian type law and the band bound. The logarithmic distance inequality from Lemma A16 completes the proof. □

The results establish global approximate optimality and policy stability of an explicit two-dealer candidate. Exact continuation existence at a fixed positive width and stability of that exact selection remain additional requirements for disclosure contraction. A uniform neighborhood of the candidate alone gives neither a Lipschitz remainder across policies nor a uniform lower bound on each idiosyncratic type's signal derivative.

A.2.17 Finite-grid continuation, measurable selection, and policy regret

This construction restricts bids to a finite set while retaining the unbounded Gaussian signals and common value. It concerns $N = 2$ and an all-silent history. It supplies existence and measurable selection in the restricted action space. Its moment estimates control expanding grids. The final policy comparison measures regret when the same prescription is used under a different disclosure policy.

Use $\lambda = \sigma_v^2/(\sigma_v^2+\sigma_\varepsilon^2)$ and $\theta_2(s,t)$ from Section A.2.14. Write the idiosyncratic value as ϵW, where W has any fixed probability law on $[0,1]$ and is independent across dealers and of their signals. In this subsection $\epsilon \geq 0$, so the zero-width case is included. Let $\mu^0(d\tau)$ be the unconditioned marginal law of $\tau = (S, W)$, and let g^0 be its signal marginal. The unconditioned joint type law is

$$Q^0(d\tau, d\tau') = k^0(s,t)\mu^0(d\tau)\mu^0(d\tau'), \qquad \tau' = (t,z). \tag{A.90}$$

The density k^0 retains the Gaussian signal correlation. Fix a constant $q \in (0,1]$. The policy space $\mathcal{P}_q$ consists of equivalence classes of measurable p with $0 \leq p \leq 1-q$, equipped with the $L^1(g^0)$ metric. The constant q is a lower bound on silence weights, distinct from the function $q(v)$ used elsewhere. No monotonicity or smoothness is imposed on p here. For these policy equivalence classes, $\|\cdot\|_\infty$ denotes the essential supremum under g^0.

Put $w_p = 1-p$, $z_p(s) = \mathbb{E}^0[w_p(T) \mid S = s]$, and $Z_p = \mathbb{E}^0[w_p(S)w_p(T)]$. The type marginal conditional on both silences, denoted by μ_p, and the corresponding correlation density satisfy

$$\frac{d\mu_p}{d\mu^0}(s,w) = \frac{w_p(s)z_p(s)}{Z_p}, \qquad k_p(s,t) = k^0(s,t)\frac{Z_p}{z_p(s)z_p(t)}. \tag{A.91}$$

Both density ratios $d\mu_p/d\mu^0$ and k_p/k^0 lie in $[q^2, q^{-2}]$.

Let $D = \{b_1, \ldots, b_J\} \subset \mathbb{R}$ be a nonempty finite bid set. A mixed prescription σ assigns probabilities $\sigma_j(\tau)$ to these bids. Equal sharing at a tie is represented by $a(b,b') = \mathbf{1}_{\{b>b'\}} + \frac{1}{2}\mathbf{1}_{\{b=b'\}}$. The actual payoff is $U_p = A_p/z_p(s)$, where the unnormalized payoff is

$$A_p(b \mid \tau, \sigma) = \mathbb{E}^0[w_p(T)[\theta_2(s,T) + \epsilon w - b]a(b,B') \mid \tau]\,. \tag{A.92}$$

The rival draws B' from $\sigma(\cdot \mid T, Z)$. A grid equilibrium is a symmetric measurable prescription maximizing this payoff over D almost everywhere in the own type.

Existence and measurable choice. Let $\mathcal{C}_D$ be the simplex-valued elements of $L^2(\mu^0)^J$ with the weak topology. This is a compact metrizable space. Write $\mathcal{E}_D(p)$ for its equilibrium subset.

Proposition A10 (Symmetric mixed equilibrium on a finite bid set)**.** *For every finite D, every policy in $\mathcal{P}_q$, and every finite idiosyncratic-value width, the two-dealer all-silent game restricted to D has a measurable symmetric mixed Bayesian equilibrium. All its interim values are finite. No smoothness of the policy or atomlessness of the idiosyncratic-value law is required.*

Proof. Let μ be the marginal type law conditional on both dealers being silent. Let $k(s,t)$ be the density of the joint signal law relative to the product of its marginals. Write $\tau = (s,w)$ and $\tau' = (t,z)$. The expected-payoff kernel for pure bids b_i, b_j is

$$K_{ij}(\tau,\tau') = k(s,t)\big[\theta_2(s,t) + \epsilon w - b_i\big]\left(\mathbf{1}_{\{b_i > b_j\}} + \tfrac{1}{2}\mathbf{1}_{\{b_i = b_j\}}\right). \tag{A.93}$$

This integrates out the common asset value using its exact conditional mean. Conditional on both signals, silence adds no further information about that value.

Each K_{ij} is square integrable under $\mu \otimes \mu$. To see this, first remove the silence weights. Standardize the two signals to marginal variance one. Their correlation is $\lambda \in (0,1)$. The square of their joint Gaussian density divided by the product of their marginal densities has a strictly negative quadratic exponent. Write Σ for their correlation matrix. The precision eigenvalues of $2\Sigma^{-1} - I$ are $(1-\lambda)/(1+\lambda)$ and $(1+\lambda)/(1-\lambda)$. Multiplication by the square of a linear function of the signals remains integrable. Restoring weights in $[q,1]$ changes the density ratios and marginal densities by bounded factors. The idiosyncratic-value and bid terms are bounded.

Let $\mathcal{C}$ be the set of vectors $\sigma = (\sigma_1, \ldots, \sigma_J)$ in $L^2(\mu)^J$ with $\sigma_i \geq 0$ and $\sum_i \sigma_i = 1$ almost everywhere. It is nonempty, convex, norm-closed, and bounded, hence weakly compact. The operators with kernels K_{ij} are compact on L^2. Thus, if two strategy sequences converge weakly, their expected payoff converges. One factor converges strongly after applying the compact operator and the other is bounded and converges weakly.

For completeness, obtain the fixed point using finite type partitions. Choose nested finite measurable partitions generating the type sigma field up to null sets. For each partition, the cell-constant strategies form a finite product of simplexes. A best response exists at each cell, and the response correspondence has nonempty convex compact values and a closed graph because its payoffs are continuous. The finite-dimensional Kakutani theorem gives a symmetric cell-constant fixed point σ_n. Extract a weakly convergent subsequence in $\mathcal{C}$, with limit σ. For any $\widehat{\sigma} \in \mathcal{C}$, its conditional expectations on the partitions converge strongly in L^2 to $\widehat{\sigma}$. The finite-partition best-response inequality, applied to these deviations, passes to the limit by the kernel compactness just proved. Hence σ maximizes expected payoff against itself over all measurable strategies.

If a set of own types of positive measure admitted a profitable deviation, choosing the best of the finitely many bids on that set would strictly increase expected payoff. This contradicts the preceding inequality. Thus the best-response condition holds almost everywhere, as required for Bayesian equilibrium. Measurable finite-action tie-breaking can specify responses at exceptional types. Gaussian conditional first moments and the finite bid set make every interim value finite. □

Lemma A19 (Closed graph of the finite-grid equilibrium correspondence)**.** *For every fixed finite grid and fixed finite width, $\mathcal{E}_D(p)$ is nonempty and compact. If $p_n \to p$ in $L^1(g^0)$ and $\sigma_n \in \mathcal{E}_D(p_n)$ converge weakly to $\sigma \in \mathcal{C}_D$, then $\sigma \in \mathcal{E}_D(p)$.*

Proof. For grid actions i, j, the unnormalized payoff operator on the fixed space has kernel $H^p_{ij}(\tau,\tau') = k^0(s,t)w_p(t)[\theta_2(s,t) + \epsilon w - b_i]a(b_i,b_j)$. Its square is integrable by the Gaussian kernel calculation in Proposition A10. Thus each operator is compact on $L^2(\mu^0)$. Nonemptiness is the finite-grid equilibrium proposition, whose conditional inequalities are unchanged by the positive normalization in (A.92).

If $p_n \to p$ in $L^1(g^0)$, the weights converge in measure and remain bounded. The square of the difference of these kernels is bounded by an integrable fixed function times $|p_n(t) - p(t)|^2$. Truncate that fixed function at height R. On the truncated part, boundedness of the policy differences and their L^1 convergence make the integral tend to zero. The remaining integral is uniformly small as $R \to \infty$. Hence the kernels converge in L^2, and the associated operators converge in operator norm.

For any fixed measurable mixed deviation $\widehat{\sigma} \in \mathcal{C}_D$, equilibrium gives the integrated inequality comparing σ_n with $\widehat{\sigma}$ against σ_n, using the reference own-type law and the unnormalized operators. Compactness sends the weakly convergent rival sequence to a strongly convergent payoff sequence. Operator-norm convergence permits replacing p_n by p. Pairing with the weakly convergent own prescription passes its payoff, and pairing with the fixed deviation passes the deviation payoff. The limiting inequality holds for every $\widehat{\sigma}$. Finite-action measurable maximization then gives the Bayesian inequalities almost everywhere. Thus the graph is closed. For fixed p, its section is a closed subset of compact $\mathcal{C}_D$ and is compact. □

Selection across disclosure policies. Fix any finite vector of bid costs $c_1, \ldots, c_J$. The cost of a prescription under the actual marginal law is $C_D(p, \sigma) = \sum_j c_j \int \sigma_j \, d\mu_p$. The squared negative part of a bid corresponds to $c_j = (b_j^-)^2$.

Theorem A1 (A measurable equilibrium choice minimizing bid cost)**.** *For every fixed finite grid and finite idiosyncratic-value width, there is a Borel map $p \mapsto \sigma_D(p)$ from $\mathcal{P}_q$ into $\mathcal{C}_D$ such that $\sigma_D(p) \in \mathcal{E}_D(p)$ and $C_D(p, \sigma_D(p)) = \min_{\sigma \in \mathcal{E}_D(p)} C_D(p, \sigma)$. In particular, an equilibrium choice is measurable without specifying a bid cost.*

Proof. First, the cost is jointly continuous in the stated topologies. Conditional expectation is an L^1 contraction, so $\|z_{p_n} - z_p\|_{L^1(g^0)} \le \|p_n - p\|_{L^1(g^0)}$. The numerator and denominator in (A.91) consequently converge in L^1 and in $\mathbb{R}$, respectively. Uniform boundedness of those densities upgrades their L^1 convergence to L^2 convergence. Pairing this convergent coefficient with a bounded weakly convergent strategy proves continuity of C_D.

The minimum $v(p)$ exists by compactness. For a fixed compact $F \subset \mathcal{C}_D$, define $v_F(p)$ to be the minimum cost on $\mathcal{E}_D(p) \cap F$, with value $+\infty$ for an empty intersection. Both v and v_F are lower semicontinuous. To verify this claim, take a convergent sequence of policies on which the liminf of the minimum is finite, choose minimizing strategies, extract a convergent subsequence in $\mathcal{C}_D$, and use the closed graph and continuity of the cost. The limit is feasible, and its cost is the limiting subsequential value, which is at least the minimum at the limiting policy. The same argument works with the additional closed restriction F.

Let $\mathcal{E}_D^*(p)$ be the compact nonempty set of cost-minimizing equilibria. It meets F exactly when $v_F(p) = v(p)$, a Borel condition. We now construct a selector rather than assume one. Choose a countable dense sequence of centers in the compact metric space $\mathcal{C}_D$. At stage one, take the first closed ball of radius $1/2$ meeting $\mathcal{E}_D^*(p)$. At stage n, intersect the previously chosen closed sets with the first closed ball of radius 2^{-n} whose intersection still meets $\mathcal{E}_D^*(p)$. Such a ball exists by density of the centers. For every fixed finite history of choices, each candidate intersection is a fixed compact set. The preceding hit test is Borel, so the first successful index at every stage is Borel in p.

The selected compact intersections are nested, meet $\mathcal{E}_D^*(p)$, and have diameters tending to zero. Their intersection with that set is a singleton. The chosen centers converge to it, making the resulting map Borel. Choosing all costs zero proves the final assertion. □

The selector is Borel measurable. Its continuity and a uniform modulus across grids require

additional arguments. Minimizing a bid cost also need not identify the normalized reference branch used in disclosure comparisons.

Uniform moments on expanding bid grids. For any symmetric prescription, let $\nu_p(d\tau, db) = \mu_p(d\tau)\sigma(db \mid \tau)$. Its actual two-dealer game law is $\Lambda_p = k_p(s,t)\nu_p \otimes \nu_p$. All game expectations and norms in the next estimates use this correlated law. Let $R = B \vee B'$ be the winning bid and let $C = R^-$, where $x^- = \max\{-x, 0\}$. Choose the smallest nonnegative grid bid h and a number $\delta > 0$ with $0 \le h \le \delta$. Assume every adjacent pair of grid bids whose upper member is at most h has a gap at most δ. No restriction is placed on gaps above h. Define $X = (h - R)^+$ and $Z_h = (h - \theta_2)^+$, so $C \le X$ and $\|Z_h\|_r \le \|\theta_2^-\|_r + \delta$.

Lemma A20 (Threshold inequality for negative payments)**.** *For every symmetric grid equilibrium under these gap conditions, and every real $r > 2$,*

$$\frac{r-2}{r(r-1)}\mathbb{E}X^r \le \frac{\mathbb{E}[(Z_h + 2\delta)X^{r-1}]}{r-1} + \frac{\delta}{2}\mathbb{E}[Z_h(X+\delta)^{r-2}]. \tag{A.94}$$

Proof. Shift all bids by h, so zero is available. For $x > 0$ outside the finitely many grid thresholds, let $c(x)$ be the smallest shifted bid strictly above $-x$. It is at most zero. Whenever a shifted bid is below $-x$, the lower-gap condition gives $c(x) \le -x + \delta$.

Consider the two separate unilateral deviations that replace the dealer's own shifted bid b by $\max\{b, c(x)\}$. Each has nonpositive expected gain. If both original shifted bids are below $c(x)$, equivalently $X > x$, each deviator would win outright. Summing their deviation gains and subtracting the two original payoffs gives at least $2x - X - 2\delta - Z_h$. The idiosyncratic-value contribution is nonnegative because the sum of the two idiosyncratic values is at least the winner's idiosyncratic value or their tie average. If only one bid is below $c(x)$, its upward deviation changes the allocation only when the rival bids exactly $c(x)$. Entering that tie gives a gain of at least $-Z_h/2$. In all remaining cases the summed gain is zero. Thus, for an event T_x recording entry into such a tie,

$$\mathbb{E}[(2x - X - 2\delta - Z_h)\mathbf{1}_{\{X > x\}}] \le \tfrac{1}{2}\mathbb{E}[Z_h \mathbf{1}_{T_x}]. \tag{A.95}$$

For a fixed realized pair, entry into a tie requires $c(x) = -X$, so its possible thresholds form a subset of $[X, X + \delta]$. Multiply the inequality by x^{r-2} and integrate. The main integral equals $(r-2)X^r/[r(r-1)]$. The negative terms integrate to $(Z_h + 2\delta)X^{r-1}/(r-1)$, and the tie integral is at most $\delta Z_h(X+\delta)^{r-2}/2$. On each finite grid, the threshold range and bids are bounded, while Z_h is integrable. Fubini's theorem therefore applies to the signed integrals and proves the result. □

For $r > 2$, put $A_r = r/(r-2)$ and $D_r = A_r + \exp(1)(r-1)/(2(r-2))$.

Theorem A2 (Higher winning-payment moments)**.** *Under the preceding lower-gap condition, every symmetric equilibrium satisfies*

$$\|C\|_r \le D_r\|\theta_2^-\|_r + \big[D_r + \max\{r, 2A_r\}\big]\delta, \qquad r > 2. \tag{A.96}$$

In particular, for $r \ge 3$,

$$\|C\|_r \le 6\|\theta_2^-\|_r + (r+9)\delta. \tag{A.97}$$

For a fixed policy, if $\delta_n \to 0$ along a sequence satisfying the lower-gap condition, then for each fixed $r > 2$,

$$\limsup_n \|C_n\|_r \le \frac{r}{r-2}\|\theta_2^-\|_r. \tag{A.98}$$

The bounds are independent of the nonnegative idiosyncratic-value width.

Proof. Lemma A20 gives (A.94). The finite grid bounds the ranges of X and the relevant thresholds, and Z_h is integrable. This justifies the integration. Put $a = \|X\|_r$ and $z = \|Z_h\|_r$. Hölder and Minkowski give

$$a^r \le A_r(z + 2\delta)a^{r-1} + \frac{r(r-1)}{2(r-2)}\delta z(a+\delta)^{r-2}. \tag{A.99}$$

For the second expectation, the Hölder exponents may include the constant function one because the total powers $1 + (r-2)$ sum to $r - 1 < r$.

If $a < r\delta$, there is already a bound. Otherwise divide (A.99) by a^{r-1}. Since $\delta/a \le 1/r$ and $(1 + 1/r)^{r-2} \le \exp(1)$, the result is $a \le D_r z + 2A_r\delta$. Combining the two cases gives $a \le D_r z + \max\{r, 2A_r\}\delta$. Since $C \le X$ and $z \le \|\theta_2^-\|_r + \delta$, this proves (A.96). For $r \ge 3$, $A_r \le 3$, $D_r \le 3 + \exp(1) < 6$, and $D_r + \max\{r, 2A_r\} \le r + 9$. This proves (A.97).

For the limiting estimate, the preceding bound makes a_n bounded at fixed r. Also $z_n \to \|\theta_2^-\|_r$ because $0 \le h_n \le \delta_n \to 0$. Along a subsequence attaining the limiting upper value of a_n, (A.99) gives $a^r \le A_r\|\theta_2^-\|_r a^{r-1}$. If $a = 0$, the conclusion is immediate, and otherwise division gives $a \le A_r\|\theta_2^-\|_r$. Finally $|X_n - C_n| \le h_n$, so their norms have the same limiting upper value. □

Choose $1 < \xi < 1/\lambda$ and put $\beta = (\xi - 1)/(2\xi) > 0$. For the next comparison, let $J_\xi = \int k_p^{1-\xi}\, d(\mu_p \otimes \mu_p)$ and $A_\xi = J_\xi^{1/(2\xi)}$. These constants are finite uniformly over $\mathcal{P}_q$. Indeed direct Gaussian integration after standardizing the signals gives

$$\int (k^0)^a\, d(\mu^0 \otimes \mu^0) = \frac{(1-\lambda^2)^{(1-a)/2}}{\sqrt{1-(a-1)^2\lambda^2}}, \qquad |a-1|\lambda < 1. \tag{A.100}$$

To verify the formula, the covariance matrix of the standardized pair has eigenvalues $1 + \lambda$ and $1 - \lambda$. The quadratic form in the integrand has precision $a\Sigma^{-1} + (1-a)I$. Its eigenvalues are positive exactly on the displayed domain. Multiplying their Gaussian integrals gives the formula. Applying it with $a = 1 - \xi$ and using (A.91) gives

$$J_\xi \le q^{-2\xi-2} \int (k^0)^{1-\xi}\, d(\mu^0 \otimes \mu^0) < \infty. \tag{A.101}$$

The marginal bid distribution under ν_p is denoted by $\nu_{p,B}$.

Lemma A21 (A reverse tail comparison)**.** *For any symmetric strategy law and any $x > 0$,*

$$\nu_{p,B}((-\infty, -x)) \le A_\xi \mathbb{P}_{\Lambda_p}(C > x)^\beta. \tag{A.102}$$

If $r > 0$, $m\beta > r$, and $\|C\|_m \le L_m$, then

$$\|B^-\|_r \le \left(\frac{m\beta}{m\beta - r}\right)^{1/r} A_\xi^{1/(m\beta)} L_m. \tag{A.103}$$

Proof. Let E_x be the event that both bids are below $-x$. Under the product law $\nu_p \otimes \nu_p$ its probability is $a(x)^2$, where $a(x) = \nu_{p,B}((-\infty, -x))$. Under the game law it is $\mathbb{P}(C > x)$. Because $k_p > 0$,

$$a(x)^2 = \int k_p^{-1}\mathbf{1}_{E_x}\, d\Lambda_p \le \left(\int k_p^{-\xi}\, d\Lambda_p\right)^{1/\xi} \mathbb{P}_{\Lambda_p}(E_x)^{1-1/\xi} = J_\xi^{1/\xi}\mathbb{P}(C > x)^{2\beta}. \tag{A.104}$$

The last equality uses the fixed product-type marginal before multiplication by k_p. Taking square roots proves (A.102).

Markov's inequality now gives $a(x) \le \min\{1, A_\xi L_m^{m\beta} x^{-m\beta}\}$. If $L_m = 0$, strict positivity of k_p gives $B^- = 0$ almost surely. Otherwise split the tail integral $\mathbb{E}(B^-)^r = r\int_0^\infty x^{r-1}a(x)\,dx$ at $x_0 = A_\xi^{1/(m\beta)} L_m$. The two integrals sum to $[m\beta/(m\beta - r)]x_0^r$. Taking the rth root proves the claim. □

For positive bids choose an available $b_0 \le \bar{b}_0$, where $\bar{b}_0$ is a fixed finite bound. Let $Q_p(s, b)$ be the conditional allocation probability including ties. Define

$$H_{\bar{b}_0}(s) = \mathbb{E}_p[\theta_2(s,T)^+ \mid s] + \epsilon + \mathbb{E}_p[(\bar{b}_0 - \theta_2(s,T))^+ \mid s]. \tag{A.105}$$

The expectations are conditioned on both silences. The Gaussian conditional law and the band give $H_{\bar{b}_0}(s) \le C_0(1 + |s|)$, with C_0 uniform over policies and bounded widths.

Lemma A22 (Positive bid moments)**.** *For every symmetric grid equilibrium with the stated available bid and every $r \ge 1$,*

$$\mathbb{E}_{\nu_p}(B^+)^r \le 2^r \mathbb{E}_{\mu_p} H_{\bar{b}_0}(S)^r. \tag{A.106}$$

Proof. The available bid gives payoff at least $-\mathbb{E}_p[(\bar{b}_0 - \theta_2)^+ \mid s]$. Optimality therefore gives $bQ_p(s,b) \le H_{\bar{b}_0}(s)$ for every used bid. Let $X_r = \mathbb{E}_{\nu_p}(B^+)^r$. Under the symmetric game law the allocation factors sum to one, so the expected positive winning payment to power r equals $2\int (b^+)^r Q_p\,d\nu_p$. Since $B \le R$, this implies

$$X_r \le 2\int (b^+)^{r-1} H_{\bar{b}_0}(s)\,d\nu_p. \tag{A.107}$$

For $r = 1$ the conclusion follows directly. For $r > 1$, Hölder bounds the right side by $2X_r^{(r-1)/r}(\mathbb{E}_{\mu_p} H_{\bar{b}_0}^r)^{1/r}$. The finite grid makes X_r finite. Divide when it is positive, and treat the zero case directly. □

Corollary A4 (Uniform bid and value moments)**.** *Fix $q > 0$ and a finite upper bound on ϵ. For grids satisfying the stated lower-gap condition with bounded δ, every finite moment of the bid and of the equilibrium interim value is bounded uniformly over policies in $\mathcal{P}_q$, grids, and symmetric equilibrium choices. In particular, their first-moment tails are uniformly integrable.*

Proof. For a desired negative bid moment of order r, choose $m \ge 3$ with $m\beta > r$. The negative winning-payment estimate and the reverse tail comparison give a bound on $\mathbb{E}(B^-)^r$. Every Gaussian moment in this bound is uniform in policy because the actual joint type law is at most q^{-2} times Q^0. The inverse-kernel constant is uniform by (A.101). The available bid $h \le \sup\delta$ permits the preceding positive-moment lemma to be applied uniformly. This bounds both tails at every order.

For the interim value, conditional Jensen and $|b|a(b,b') \le |R|$ give an L^r norm bound by $\|\theta_2\|_r + \epsilon + \|R\|_r$. The winning-bid norm is at most $2\|B\|_r$, so the value bound follows. Apply the moment bound at order $r+1$ to obtain uniform integrability at order r, including the first-moment assertion. □

These bounds imply tightness of the joint type-and-bid laws on grids that become dense on every bounded bid interval and have vanishing lower gaps. At fixed positive width, identifying a weak limit as an unrestricted equilibrium still requires control of allocation at limiting ties. Moment bounds alone do not give an atomless limit. For continuous policies at zero width, Proposition A7 already supplies an explicit normalized unrestricted equilibrium independently of this approximation argument.

Regret under a policy change. Fix a finite grid D containing a bid b_0 with $|b_0| \le d_0$. Let σ be any symmetric equilibrium on that grid under $p_1 \in \mathcal{P}_q$. Use the same conditional prescription for both dealers under another policy $p_2 \in \mathcal{P}_q$. Let $U_2(b \mid \tau, \sigma)$ be the actual normalized interim payoff under p_2, and let $U_2(\sigma \mid \tau, \sigma)$ average this payoff over the own prescription. Define its integrated regret by

$$\mathcal{R}_D(p_2, \sigma) = \int \left[\max_{b \in D} U_2(b \mid \tau, \sigma) - U_2(\sigma \mid \tau, \sigma) \right] \mu_{p_2}(d\tau). \tag{A.108}$$

Write $d_p = \|p_2 - p_1\|_\infty$, $M_\theta = \mathbb{E}^0|\theta_2(S,T)|$, and $L_- = \int b^- \mu^0(d\tau)\sigma(db \mid \tau)$. These are finite. Define

$$C_* = (1 + 2/q)M_\theta + (1 + 1/q)(\epsilon + L_-) + d_0/q. \tag{A.109}$$

Theorem A3 (A policy perturbation bound uniform along the grids)**.** *With the objects and finite-grid equilibrium just specified,*

$$0 \le \mathcal{R}_D(p_2, \sigma) \le \frac{2C_*}{q^2}\, d_p. \tag{A.110}$$

For bounded widths and grids with bounded lower gaps and a uniformly bounded near-zero bid, the constants can be chosen uniformly over those grids, their equilibrium selections, and policies in $\mathcal{P}_q$.

Proof. Choose a measurable pure best response $\widehat{b}(\tau)$ to σ under p_2 by taking the first maximizer among the finite bids. Set $Y = \theta_2(S,T) + \epsilon W$. Under the reference correlated law Q^0, including the independent bid randomizations, a prescribed or deviating payoff is $G = (Y - b)a(b, B')$.

First obtain an absolute payoff bound for either the equilibrium prescription σ under p_1 or the best response $\widehat{b}$ under p_2. In the corresponding case write $w_i(T)$ for its rival silence weight. Optimality against b_0 gives $\mathbb{E}^0[w_i(T)G] \ge -\mathbb{E}^0[w_i(T)(b_0 - \theta_2)^+] \ge -(M_\theta + d_0)$. This follows by integrating the unnormalized interim best-response inequalities under the reference own-type law. For a mixed prescription, integrate its own randomization as well.

Whenever a negative own bid wins or ties, its absolute payment is at most the rival's negative bid. Pointwise, $b^- a(b, B') \le (B')^-$. Thus $\mathbb{E}^0[w_i(T)b^- a] \le L_-$, since the reference rival type-and-bid marginal is $\mu^0\sigma$. Also $\mathbb{E}^0[w_i(T)Ya] \le M_\theta + \epsilon$. Writing the payoff as $Ya - b^+a + b^-a$ and applying its lower bound yields $\mathbb{E}^0[w_i(T)b^+a] \le 2M_\theta + \epsilon + L_- + d_0$. Because $w_i \ge q$, the same positive payment without the weight is at most q^{-1} times this bound. Finally $|G| \le |Y| + b^+a + b^-a$ gives $\mathbb{E}^0|G| \le C_*$. This argument applies separately to both compared prescriptions with the same C_* and does not involve the largest positive grid bid.

For almost every own type, equilibrium at p_1 gives $A_{p_1}(\widehat{b}) - A_{p_1}(\sigma) \le 0$. Only the rival silence weight changes in the unnormalized payoff. Therefore $A_{p_2}(\widehat{b}) - A_{p_2}(\sigma) \le d_p \mathbb{E}^0[|G_{\widehat{b}}| + |G_\sigma| \mid \tau]$. To integrate actual regret, use (A.91) and divide by the interim normalizer $z_{p_2}(s)$. The normalizers cancel exactly, giving $\mathcal{R}_D(p_2, \sigma) = Z_{p_2}^{-1} \int w_{p_2}(s)[A_{p_2}(\widehat{b}) - A_{p_2}(\sigma)]\mu^0(d\tau)$. Since $w_{p_2}(s) \le 1$, $Z_{p_2} \ge q^2$, and both absolute payoffs have expectation at most C_*, this proves (A.110).

For the uniform assertion, the widths and d_0 are bounded by hypothesis. Corollary A4 and the policy density bounds give a common bound on the negative first moment under the actual marginal law. Its density is at least q^2 times the reference marginal, so L_- has a common bound as well. No property of the selected equilibrium beyond optimality and those already proved moment bounds is used. □

The result compares the same strategy under two belief environments. It is valid for every finite-grid equilibrium selection, including a discontinuous selection. It does not bound the distance between an equilibrium selected at p_1 and one selected at p_2, and it does not establish the strategy channel of the disclosure contraction proof.

A.3 Analysis of the disclosure stage

We study the disclosure stage of the Perfect Bayesian Equilibrium defined in Section 3.2. All N dealers participate in the continuation auction. Disclosure precedes the independent draw of the idiosyncratic value u_j. A dealer therefore compares expected gross continuation values before paying their privately observed talking cost. Throughout the baseline analysis, continuation values use the measurable selection of finite-value, globally optimal continuation equilibria maintained in Assumption A1. This selection must be defined at every history used to evaluate either disclosure action, including a verifiable announcement that has zero probability under the conjectured policy.

Fix dealer j with signal $s_j = s$. Their rivals use a common talk propensity p, with silence weight $w = 1 - p$. Observed signal values enter the Gaussian posterior directly. The probabilities of disclosing those exact values cancel from Bayes' rule wherever the history is on path. At a forced announcement, we use the same Gaussian posterior and the remaining silence weights as the specified continuation belief. The dealer's own talking cost is sunk in the auction and does not enter their continuation value.

A.3.1 Continuation values and the disclosure decision

Let $\mathcal{T}_{-j}$ be the set of rival talkers and let $\mathcal{M}_{-j}$ be their disclosed signals, with identities retained. If dealer j talks, the public messages become $\mathcal{M}^{\text{in}}(s) = \mathcal{M}_{-j} \cup \{s\}$. If they remain silent, they are $\mathcal{M}^{\text{out}} = \mathcal{M}_{-j}$. The resulting public histories, including the disclosure identities, are denoted by $\mathcal{I}^{\text{in}}(s)$ and $\mathcal{I}^{\text{out}}$.

Write $R^T(\mathcal{I}; p)$ for a talker's expected gross rent in the selected continuation, integrating over all idiosyncratic values. Their disclosure payoff at the realized rival history is

$$\Pi^T(s; \mathcal{T}_{-j}, \mathcal{M}_{-j}; p) = R^T(\mathcal{I}^{\text{in}}(s); p). \tag{A.111}$$

Let $\Pi^S(s; \mathcal{T}_{-j}, \mathcal{M}_{-j}; p)$ be their optimized gross value from silence, again integrated over all idiosyncratic values. These are exact continuation values. An approximation with error $O(\bar{u})$ need not preserve the sign of a disclosure gain that is itself of order $\bar{u}$.

For comparison, in the auxiliary auction that sets every silent dealer's idiosyncratic value to zero, the payoff at the assigned benchmark bid is

$$\Pi^{S,\text{bm}}(s; \mathcal{I}^{\text{out}}) = H^{\text{bm}}(b^{S,\text{bm}}(s)) G(s \mid s, \mathcal{I}^{\text{out}})[\varphi(s; \mathcal{I}^{\text{out}}) - b^{S,\text{bm}}(s)]. \tag{A.112}$$

The factors are the probabilities of defeating talkers and other silent dealers, followed by the common-value margin conditional on winning. The exact game instead uses the bid-clearing objects in (A.25), which integrate over rivals' idiosyncratic values. The auxiliary strategy $b^{S,\text{bm}}$ need not equal the zero-value slice $b^S(s, 0)$ of the exact strategy. Under hard truncation, a signal outside the conjectured silent support is evaluated by an optimal deviation in the specified continuation, rather than by extrapolating an assigned strategy outside its domain.

Because disclosure is simultaneous, the dealer averages over rival signals and disclosure decisions. Define

$$\bar{\Pi}^a(s; p) = \mathbb{E}[\Pi^a(s; \mathcal{T}_{-j}, \mathcal{M}_{-j}; p) \mid s_j = s], \qquad a \in \{T, S\}. \tag{A.113}$$

The expectation also incorporates the idiosyncratic values already integrated into each continuation payoff. The gross gain from talking is

$$\Delta(s; p) = \bar{\Pi}^T(s; p) - \bar{\Pi}^S(s; p). \tag{A.114}$$

A dealer talks when $\kappa_j \leq \Delta(s;p)$. Equality can be assigned either action because the cost distribution is continuous. With costs uniform on $[0,K]$, the induced propensity is $\mathcal{B}(p)(s) = G_\kappa(\Delta(s;p))$.

Retaining a favorable signal can raise the value of silence because the dealer can use that information without revealing it to competitors. Disclosure also changes public beliefs, rival bids, and the distribution of continuation configurations. The sign of the total disclosure response must account for all of these changes.

A.3.2 Talker rents

At a fixed continuation history, let $D_T(u)$ be the winning probability of a talker with idiosyncratic value u, and let $Q^{\text{ex}}(u)$ be the probability that their assigned bid clears all silent rivals. Set $Q^{\text{ex}} \equiv 1$ when there are no silent rivals. Let $Q^0(u)$ be the clearing probability at the same talker bid and under the same exact rival strategies when silent idiosyncratic values are set to zero. This last comparison holds strategies fixed. Write $\mu_u = \mathbb{E}[u]$ for the mean idiosyncratic value.

Lemma A23 (Structure of talker rents)**.** *Suppose the selected talker value is finite, the lowest-type value is zero by continuous extension, and the symmetric talker strategy is strictly increasing on* $(0,\bar{u}]$*. Then*

$$R^T = \int_0^{\bar{u}} (1 - F_u(t)) D_T(t)\, dt, \qquad D_T(t) = F_u(t)^{n_T - 1} Q^{\text{ex}}(t). \tag{A.115}$$

If silent strategies are nondecreasing in their idiosyncratic values, $Q^{\text{ex}}(t) \leq Q^0(t)$*. In particular,* $0 \leq R^T \leq \mu_u$*.*

Proof. At every fixed bid the talker's payoff is affine in u, with slope equal to its winning probability. The optimized finite value is convex and 1-Lipschitz. Its derivative equals $D_T(u)$ almost everywhere, so zero lowest-type value gives $\Pi^T(u) = \int_0^u D_T(t)\, dt$. Tonelli's theorem then gives the rent integral. The talker defeats their $n_T - 1$ co-talkers exactly when their idiosyncratic values are below theirs. That event is independent of clearing silence, which proves the factorization. Setting silent idiosyncratic values to zero weakly lowers their bids and enlarges the clearing event. Finally, $D_T \leq 1$ gives $R^T \leq \int_0^{\bar{u}} (1 - F_u(t))\, dt = \mu_u$. □

The zero-value normalization follows from the maintained margin conditions in Lemma A6 in mixed configurations and from the selected symmetric private-value equilibrium when all dealers talk. The comparison with Q^0 does not compare two equilibria. In particular, it does not order exact clearing relative to the auxiliary scalar equilibrium.

Under hard truncation, write $w(s) = \mathbf{1}\{s > \bar{s}\}$ and $d = \mu_T - \bar{s}$. At fixed Gaussian primitives, idiosyncratic-value law, and counts (n_T, n_S), the public posterior and continuation game depend on disclosed signals through their Gaussian sufficient statistic. Translating μ_T, $\bar{s}$, values, signals, and bids by the same amount preserves payoffs and winning probabilities. Consequently, the equilibrium correspondence has this translation property. To write a selected rent as $R^T(n_T, n_S, d)$, we select as a function of the resulting Bayesian continuation game and require the selection to respect translations. We impose this convention only in the hard-truncation comparisons below. A general weighted policy does not reduce to the statistic d.

For those comparisons, suppose $u = \bar{u}Z$ with a fixed continuous distribution F_0 of full support on $[0,1]$. A swap moves one rival from talking to silence while retaining the focal talker. Denote by $Q^{\text{swap,old}}(u;d)$ the probability that the focal bid in the swapped equilibrium clears the original silent rivals.

Assumption A3 (Comparability of continuation clearing probabilities)**.** *Fix a compact set D of hard-truncation gaps. The compared selections depend on the public reports only through the resulting Bayesian continuation game, respect translations, and have zero lowest-type talker value. There are $\eta_D > 0$ and $C_D < \infty$, uniform over $d \in D$ and sufficiently small positive $\bar{u}$, such that:*

(i) $Q^{\text{ex}}(u; d) \geq \eta_D Q^{\text{ex}}(\bar{u}; d)$ *for* $u \in [\bar{u}/2, \bar{u}]$.

(ii) $Q^{\text{swap,old}}(u; d) \leq C_D Q^{\text{ex}}(u; d)$.

(iii) In each compared continuation with silent dealers, assigned positive-winning talker bids lie in the image of the exact zero-value silent slice and have nonnegative common-cutoff margin $\psi(x(u)) + u - \beta^T(u)$*, where* $b^S(x(u), 0) = \beta^T(u)$*. The top clearing probability is positive. Assumption A2 holds in the configurations with silence, including its conditional density bounds for every clearing event used in the swap.*

(iv) When $n_S \geq 1$*,* $Q^{\text{ex}}(u; d)$ *is nonincreasing in* d *for almost every* u*, after the selected strategies adjust.*

These are restrictions on selected continuation equilibria. The last condition is an equilibrium comparative static. The Gaussian argument below establishes its fixed-clearing-rule component only.

Proposition A11 (Talker rents and informed competition)**.** *Under Assumption A3, uniformly over* $d \in D$*, a swap with* $n_T \geq 2$ *satisfies*

$$R^T(n_T - 1, n_S + 1, d) \leq C'_D \bar{u}\, R^T(n_T, n_S, d) \tag{A.116}$$

for a finite constant C'_D *and sufficiently small positive* $\bar{u}$*. The inequality is strict once* $C'_D \bar{u} < 1$*. For* $n_S \geq 1$*, the rent is nonincreasing in* d*. The latter comparison is strict between two gaps if their clearing probabilities differ strictly on a positive-measure set of idiosyncratic types in* $(0, \bar{u})$*.*

Proof. In the swapped continuation, the common-cutoff margin gives $b^S(x(u), 0) - \psi(x(u)) \leq \bar{u}$. Monotonicity of this overbid and its local lower bound in Assumption A2 imply $x(u) - \bar{s} \leq \bar{u}/a_D$ for sufficiently small $\bar{u}$. Conditional on clearing the original silent rivals, the newly silent rival is therefore cleared with probability at most $\lambda_D \bar{u}$, where $\lambda_D = B_D/a_D$. These are conditional bounds because clearing other rivals changes the posterior of the common value.

The swapped winning probability is at most $F_u(u)^{n_T - 2} \lambda_D \bar{u}\, Q^{\text{swap,old}}(u; d)$. Assumption A3(ii), monotonicity of clearing in the talker type, and the rent envelope give an upper bound $C_D \lambda_D \bar{u}^2 Q^{\text{ex}}(\bar{u}; d)$ for the swapped rent. For the original rent, restrict the envelope integral to $[\bar{u}/2, 3\bar{u}/4]$. It is at least $c_D \bar{u} Q^{\text{ex}}(\bar{u}; d)$, where $c_D = \eta_D [1 - F_0(3/4)] F_0(1/2)^{n_T - 1}/4 > 0$. Dividing proves (A.116). Integrating the clearing order in part (iv) against the positive envelope weights proves the gap comparison and its strict version. □

Lemma A24 (Positive expected payoff from talking)**.** *Suppose the selected continuations satisfy the talker envelope and* $\bar{u} > 0$*. Then* $\bar{\Pi}^T(s; p) > 0$ *whenever a positive-probability set of rival histories gives positive talker winning probability on a positive-measure set of idiosyncratic types in* $(0, \bar{u})$*. This condition holds if all rivals talk with positive probability. It also holds for* $N \geq 3$ *under a continuous policy bounded strictly below one and the mixed-continuation conditions of Assumption A1.*

Proof. The rent envelope has a nonnegative integrand, with $1 - F_u(t) > 0$ for $t < \bar{u}$. The stated positive-winning condition therefore gives positive expected rent. When all rivals talk, the selected symmetric private-value auction gives positive rent to positive idiosyncratic types. For the second sufficient condition, all $N - 1$ rivals remain silent with positive probability. The resulting continuation after the focal disclosure has at least two silent rivals, so Lemma A12 gives positive rent at every positive idiosyncratic value. Averaging over this event proves the claim. □

The participation lemma for at least two silent rivals does not by itself cover the two-dealer mixed history. Positive talker rent in that history requires its own continuation argument. The upper bound $\bar{\Pi}^T \leq \mu_u$ uses the maintained margin conditions and does not require strict participation.

A.3.3 Signal comparisons and disclosure switches

The Gaussian structure provides exact comparisons at fixed bids and public histories. To obtain a disclosure comparison, we must also account for changes in the disclosed signal values and in the identities of rival talkers. We develop these components separately. All formulas in this subsection use the exact continuation game unless hard truncation is explicitly imposed.

Gaussian coupling. Define $\varrho = \sigma_v^2/(\sigma_v^2 + \sigma_\varepsilon^2) \in (0, 1)$.

Lemma A25 (Monotone Gaussian coupling)**.** *Conditional on $s_j = s$, there is a centered Gaussian vector $(\eta_v, \{\eta_k\}_{k \neq j})$ with distribution independent of s such that*

$$v(s) = \mu_0 + \varrho(s - \mu_0) + \eta_v, \qquad s_k(s) = \mu_0 + \varrho(s - \mu_0) + \eta_k \quad (k \neq j). \tag{A.117}$$

Increasing s by $h > 0$ raises every rival signal by ϱh and raises the focal signal relative to every rival by $(1 - \varrho)h$.

Proof. Gaussian conditioning gives $\mathbb{E}[v \mid s_j = s] = \mathbb{E}[s_k \mid s_j = s] = \mu_0 + \varrho(s - \mu_0)$. The conditional covariance matrix does not depend on s. Subtracting the conditional means gives the stated vector and coupling. □

For a nonincreasing policy, draw independent uniforms U_k and let rival k talk when $U_k \leq p(s_k(s))$. Along this coupling a rival can switch from talking to silence as s rises. A rival who continues to talk discloses a higher signal. Thus fixed disclosure identities do not mean fixed public information.

The own-signal derivative at a fixed public history. Fix a public history $\mathcal{I}$ at which the focal dealer is silent. Let n_S include that dealer. For a fixed bid b, let $\mathcal{C}_b$ be the event that every silent rival bids below b, using their full strategy $b^S(t, u; \mathcal{I})$. Define

$$V(s, b; \mathcal{I}) = \mathbb{E}[v \mid s, \mathcal{I}, \mathcal{C}_b], \qquad \bar{V}(s; \mathcal{I}) = \mathbb{E}[v \mid s, \mathcal{I}], \qquad G(b \mid s, \mathcal{I}) = \Pr(\mathcal{C}_b \mid s, \mathcal{I}). \tag{A.118}$$

Here G and V are the exact objects denoted by G^{ex} and V^{ex} in (A.25). If $n_S = 1$, set $G \equiv 1$ and $V = \bar{V}$. Let $H(b \mid \mathcal{I})$ be the cdf of the highest talker bid, with $H \equiv 1$ when there are no talkers. Conditional on the public history, talker idiosyncratic values are independent of the common value and silent types.

Write $f_s(v)$ for the Gaussian density given the observed signals $(s, \mathcal{M})$ before applying silence weights. Its variance is $\sigma_S^2 = 1/[\tau_v + (n_T + 1)\tau_\varepsilon]$, and the derivative of its mean is $\gamma' = \tau_\varepsilon \sigma_S^2$. All bid-conditioned quantities below are evaluated where clearing has positive probability.

Lemma A26 (Gaussian score identities)**.** *At a fixed public history, rival strategy profile, and finite bid,*

$$\partial_s V(s,b;\mathcal{I}) = \tau_\varepsilon \operatorname{Var}(v \mid s,\mathcal{I},\mathcal{C}_b) > 0, \tag{A.119}$$

and

$$\partial_s \log G(b \mid s,\mathcal{I}) = \tau_\varepsilon[V(s,b;\mathcal{I}) - \bar{V}(s;\mathcal{I})] \le 0. \tag{A.120}$$

The second inequality is strict when silent clearing is informative about v.

Proof. The Gaussian score is

$$\partial_s \log f_s(v) = \frac{\gamma'}{\sigma_S^2}[v - \mu_S(s)] = \tau_\varepsilon[v - \mu_S(s)]. \tag{A.121}$$

Let $C(b,v)$ and $q(v)$ be the likelihoods in (A.22). The numerator weight for clearing is $C(b,v)^{n_S-1}$, and the weight before clearing is $q(v)^{n_S-1}$. Both are independent of the focal signal when the bid, history, and rival strategies are fixed. Differentiating the normalized moment and the two normalizing constants gives the identities.

Conditional on silence, the rival signal density is proportional to $w(t)f_\varepsilon(t-v)$. Its likelihood ratio between two values of v is increasing in t because w cancels. Each signal cutoff is therefore cleared with probability nonincreasing in v. Averaging those probabilities over rival idiosyncratic values gives the same order for $C(b,v)/q(v)$. Thus clearing is adverse news, so $V \le \bar{V}$. This argument does not require the unnormalized likelihood $C(b,v)$ to be monotone in v. □

The fixed-bid payoff is $U^S(b,s;\mathcal{I},u) = H(b \mid \mathcal{I})G(b \mid s,\mathcal{I})[V(s,b;\mathcal{I}) + u - b]$. It retains the full distribution of rival idiosyncratic values.

Lemma A27 (Exact fixed-bid derivative)**.** *For fixed $(b,\mathcal{I},u)$ with positive winning probability,*

$$\partial_s U^S(b,s;\mathcal{I},u) = \tau_\varepsilon HG\{\operatorname{Var}(v \mid s,\mathcal{I},\mathcal{C}_b) - [\bar{V} - V][V + u - b]\}. \tag{A.122}$$

Proof. Differentiate $HG(V+u-b)$ with the bid, history, and rival strategies fixed, then apply (A.119) and (A.120). □

At a selected optimizer where the value is differentiable and the envelope theorem applies, the same expression gives the own-signal derivative of the optimized value at fixed $\mathcal{I}$. When $\bar{V} > V$, its sign is nonnegative exactly when

$$V(s,b^*;\mathcal{I}) + u - b^* \le \frac{\operatorname{Var}(v \mid s,\mathcal{I},\mathcal{C}_{b^*})}{\bar{V}(s;\mathcal{I}) - V(s,b^*;\mathcal{I})}. \tag{A.123}$$

If $\bar{V} = V$, the variance term gives a positive derivative whenever winning has positive probability. These are fixed-history comparisons. They do not sign the total disclosure derivative.

The derivative along the coupling. For the derivative calculations, let ω include the Gaussian vector, rival disclosure uniforms, and the focal idiosyncratic value. Write $\Pi^a(s,\omega)$ for the optimized continuation value after action a, retaining that focal idiosyncratic value and averaging over rivals' idiosyncratic values. Then $\mathbb{E}_\omega[\Pi^a(s,\omega)] = \bar{\Pi}^a(s;p)$. Let D_s^{fix} mean differentiation along (A.117) with the identities of rival talkers fixed. In particular, for the silent action,

$$D_s^{\text{fix}}\Pi^S = \partial_s \Pi^S\big|_{\mathcal{I}^{\text{out}}} + \varrho \sum_{k \in \mathcal{T}_{-j}} \partial_{m_k}\Pi^S. \tag{A.124}$$

The derivatives in the sum include changes in public beliefs and in the selected rival bids. They need not have a known sign.

Assumption A4 (Regularity for disclosure derivatives)**.** *For the derivative results, p is continuously differentiable and nonincreasing. At each fixed disclosure configuration, the selected continuation values admit the own-signal envelopes and the derivatives along the public-message path used below. They are locally absolutely continuous along that path. The derivatives and switch contributions have locally integrable dominating bounds that justify differentiation under expectation. At a zero where a pointwise derivative is invoked, that derivative and the displayed identities exist.*

This regularity is additional to finite values and measurability. It is used only for the differential sufficient conditions. The lattice existence result below allows discontinuous monotone policies and states its order conditions directly.

For a rival k, define $J_k(s,\omega)$ by forcing that rival to switch from talking to silence, while holding their signal and every other rival's disclosure decision fixed. It is the resulting change in the focal gain from talking. The forced values do not depend on U_k.

Lemma A28 (Local decomposition)**.** *Under Assumption A4, for almost every s,*

$$\partial_s \Delta(s;p) = \mathbb{E}_\omega[D_s^{\mathrm{fix}}\Pi^T - D_s^{\mathrm{fix}}\Pi^S] + \varrho\mathbb{E}_\omega\left[\sum_{k\neq j} |p'(s_k(s))| J_k(s,\omega)\right]. \tag{A.125}$$

Proof. Conditional on the Gaussian vector, integrate the independent disclosure uniforms first. The expected payoff is a finite multilinear function of the rival probabilities $p(s_k(s))$. Differentiating its continuation values gives the fixed-identity term. Differentiating the probability of rival k gives $\varrho p'(s_k)$ times the payoff difference between forcing that rival to talk and forcing their to be silent. Since $p' \leq 0$, this is $\varrho|p'(s_k)|J_k$. The assumed domination justifies integration over the remaining variables. □

Uniform upper-tail deactivation on a policy band. The high-signal comparison can be established without a lower bound on silent bids. Let $q_0 > 0$ be fixed and suppose $q_0 \leq w \leq 1$. This includes every baseline policy in the invariant band by taking $q_0 = 1 - \bar{u}/K$.

Lemma A29 (No overbidding relative to interim value)**.** *Fix a bidder's information. Suppose optimized payoff is nonnegative and winning is weakly adverse news about v. Every optimal bid with positive winning probability is at most the bidder's interim mean of v plus their idiosyncratic value.*

Proof. A higher bid would give strictly negative expected surplus conditional on winning and hence negative payoff. It could not attain a nonnegative optimized value. □

Proposition A12 (Eventual negativity away from full talk)**.** *Fix $q_0 > 0$ and the primitives. Suppose that, for every policy with $q_0 \leq 1 - p \leq 1$, the selected continuations have nonnegative finite values, finite ex ante action values, and talker values at most u. Suppose also that in the all-silent continuation winning is adverse news and the selected bids win with positive probability for almost every bidder type. Then there is a finite S^*, independent of the policy and continuation selection within this class, such that*

$$\Delta(s;p) < 0 \quad \textit{and} \quad \mathcal{B}(p)(s) = 0 \qquad (s \geq S^*). \tag{A.126}$$

Proof. Put $M = N - 1$ and $\sigma_1^2 = (\tau_v + \tau_\varepsilon)^{-1}$. In the all-silent history, the posterior given a dealer's signal t is the Gaussian posterior tilted by the M other silence probabilities. Their product lies in $[q_0^M, 1]$. Consequently its mean differs from $\mu_0 + \varrho(t - \mu_0)$ by at most $C_q = q_0^{-M}\sigma_1\sqrt{2/\pi}$. Lemma A29 bounds a rival's bid almost surely by $\mu_0 + \varrho(t - \mu_0) + C_q + \bar{u}$.

Conditional on the focal signal, the event A that all rivals are silent has probability at least q_0^M. The conditional law of the Gaussian coupling vector given A has density at most q_0^{-M} relative to its unconditioned law. Choose a with $\varrho^2 < a < \varrho$ and let the focal dealer bid $b_s = \mu_0 + a(s - \mu_0)$ after observing A. A rival's cap is $\mu_0 + \varrho^2(s - \mu_0) + \varrho\eta_k + C_q + \bar{u}$. Gaussian tail bounds therefore make the probability that any rival cap exceeds b_s tend to zero exponentially in $(s-\mu_0)^2$, uniformly over the band. The same bounds and Cauchy–Schwarz make the expected absolute surplus on this exceptional event tend to zero.

The mean of v conditional on (s, A) differs from $\mu_0 + \varrho(s - \mu_0)$ by at most C_q. The deviation payoff is therefore at least $(\varrho - a)(s - \mu_0) - C_q - o(1)$, uniformly over the policies and focal idiosyncratic values. Optimized payoffs in other configurations are nonnegative. For sufficiently large s, averaging gives $\bar{\Pi}^S(s; p) \geq q_0^M[(\varrho - a)(s - \mu_0) - C_q - o(1)]$. This diverges uniformly, whereas $\bar{\Pi}^T(s; p) \leq \mu_u$. The conclusion follows. □

In a symmetric all-silent continuation with bids strictly increasing in the signal and nondecreasing in the idiosyncratic value, a positive idiosyncratic type can beat rivals with lower signals and lower idiosyncratic values. Full-support silence then gives positive winning probability. The preceding proposition nevertheless states that property explicitly because continuation existence and regularity remain maintained.

Remark A3 (The baseline excludes full talk)**.** *Lemma A7 gives $p(s) \leq \bar{u}/K < 1$ at every best response. Silence therefore remains on path at every signal. Proposition A12 strengthens pointwise eventual negativity to a common upper region of zero disclosure on a fixed band, under its all-silent continuation conditions. It does not require p to be differentiable or monotone.*

A bound for the complete disclosure derivative. At the selected silent bid let $W = HG$, let $\Sigma_W = \mathrm{Var}(v \mid s, \mathcal{I}^{\mathrm{out}}, \mathcal{C}_{b^*})$, and use V and $\bar{V}$ from (A.118). When $\bar{\Pi}^S(s; p) > 0$, define $\kappa_V(s) = \mathbb{E}[(\bar{V} - V)\Pi^S]/\bar{\Pi}^S(s; p)$. Its exact fixed-history interpretation is

$$\tau_\varepsilon \kappa_V(s) = \mathbb{E}_{\Pi^S}\left[-\partial_s \log G(b^* \mid s, \mathcal{I}^{\mathrm{out}})\big|_{b^*, \mathcal{I}^{\mathrm{out}}} \right], \tag{A.127}$$

where $\mathbb{E}_{\Pi^S}$ weights observations by their silent payoff. If $\bar{\Pi}^S = 0$, define the product $\kappa_V \bar{\Pi}^S$ as zero. Values at events with zero winning probability are interpreted through their zero payoff weights.

Define the expected public-history contribution to silent payoff by $A_S(s) = \varrho\mathbb{E}[\sum_{k \in \mathcal{T}_{-j}} \partial_{m_k} \Pi^S]$, and the full fixed-identity talker derivative by $A_T(s) = \mathbb{E}[D_s^{\mathrm{fix}} \Pi^T]$. Whenever $\bar{\Pi}^T(s; p) > 0$, put $\mathcal{R}(s) = [A_S(s) - A_T(s)]/\bar{\Pi}^T(s; p)$. This is a signed quantity. It records the public-information contribution relative to talker payoff, including changes in the selected bidding strategies.

Lemma A30 (Payoff-level bound on silent drift)**.** *Under Assumption A4, at almost every signal where the displayed expectations are finite,*

$$\mathbb{E}[D_s^{\mathrm{fix}} \Pi^S] = A_S(s) + \tau_\varepsilon \mathbb{E}[W\Sigma_W] - \tau_\varepsilon \kappa_V(s)\bar{\Pi}^S(s; p) \geq A_S(s) - \tau_\varepsilon \kappa_V(s)\bar{\Pi}^S(s; p). \tag{A.128}$$

Proof. Apply the fixed-history envelope (A.122), add the public-message derivative in (A.124), and take expectations. The identity uses $W(V + u - b^*) = \Pi^S$. Dropping the nonnegative variance term gives the inequality. □

Write Π^a_{before} and Π^a_{after} for the values before and after a forced rival switch. Nonnegative values imply

$$J_k^+ \le \Pi^S_{\text{before}} + \Pi^T_{\text{after}}. \tag{A.129}$$

Let $f_{s_k|s}(t)$ be the Gaussian rival-signal density conditional on $s_j = s$. Define $C_J(s)$ as the maximum over rivals of the essential supremum over t of $f_{s_k|s}(t)\mathbb{E}[\Pi^S_{\text{before}} + \Pi^T_{\text{after}} \mid s_k = t, s_j = s]$. The expectation uses the other rivals' natural decisions and forces the specified rival's action. Finally, let $\mathrm{TV}(p) = p(-\infty) - p(+\infty)$ for a nonincreasing policy.

Lemma A31 (Total-variation bound on switch flow)**.** *For a continuously differentiable nonincreasing policy with finite $C_J(s)$,*

$$\varrho\mathbb{E}\left[\sum_{k\neq j} |p'(s_k)| J_k^+\right] \le \varrho(N-1)\,\mathrm{TV}(p) C_J(s). \tag{A.130}$$

Proof. Condition on $s_k = t$. The expectation for rival k is bounded by $C_J(s)\int |p'(t)|\,dt = C_J(s)\,\mathrm{TV}(p)$. Sum over the rivals and multiply by ϱ. □

Combining these results gives the complete upper bound

$$\partial_s \Delta(s;p) \le -\mathcal{R}(s)\bar{\Pi}^T(s;p) - \tau_\varepsilon \mathbb{E}[W\Sigma_W] + \tau_\varepsilon \kappa_V(s)\bar{\Pi}^S(s;p) + \varrho(N-1)\,\mathrm{TV}(p) C_J(s). \tag{A.131}$$

If talking payoff is zero, replace $-\mathcal{R}\bar{\Pi}^T$ by $A_T - A_S$. On the baseline band, $\mathrm{TV}(p) \le \bar{u}/K$. The formula requires no bound on $\sup|p'|$. The terms A_S and A_T remain essential because changing the focal signal moves public messages even without a disclosure switch.

The hard-truncation gap comparison. The next formulas concern the hard-truncation family, with cutoff $\bar{s}$, gap $d = \mu_T - \bar{s}$, and the preceding selection convention. After centering at the cutoff, the selection depends on the reports through (n_T, n_S, d). A general weighted disclosure policy retains its full silence weight.

Lemma A32 (Exact gap drift)**.** *On the talking branch, when the focal signal is disclosed and the identities of the $n_T - 1$ rival talkers are fixed, the Gaussian coupling gives*

$$\frac{d\mu_T}{ds} = \frac{\tau_\varepsilon[1 + \varrho(n_T - 1)]}{\tau_v + n_T\tau_\varepsilon} = \varrho. \tag{A.132}$$

Thus $dd/ds = \varrho$ when the hard cutoff is fixed.

Proof. Use $\varrho = \tau_\varepsilon/(\tau_v + \tau_\varepsilon)$ to obtain $1 + \varrho(n_T - 1) = (\tau_v + n_T\tau_\varepsilon)/(\tau_v + \tau_\varepsilon)$ and substitute. □

Let $m = n_S \ge 1$, let f_d be the Gaussian density with mean $\mu(d) = \bar{s} + d$ and variance $\sigma_M^2 = 1/(\tau_v + n_T\tau_\varepsilon)$, and let $q(v) = \Pr(s_k > \bar{s} \mid v)$. Conditional on the public silence events, the density of v is $g_d(v) = f_d(v)q(v)^m/Z_d$, where $Z_d = \int f_d(v)q(v)^m\,dv$. Denote expectation under this density by $\mathbb{E}_{\mathcal{I},d}$.

For a talker idiosyncratic type t, let $T(v;t,d)$ be the probability that their assigned bid clears all silent rivals, conditional on v and their silence. It includes all rival idiosyncratic values. Thus $Q^{\text{ex}}(t;d) = \int T(v;t,d)g_d(v)\,dv$. Let $\mathbb{E}_{T,d}$ denote expectation after additionally conditioning on this clearing event.

Lemma A33 (Exact decomposition of the gap derivative)**.** *Where $Q^{\mathrm{ex}} > 0$ and differentiation under the integrals is valid,*

$$-\partial_d \log Q^{\mathrm{ex}}(t;d) = \frac{\mathbb{E}_{\mathcal{I},d}[v] - \mathbb{E}_{T,d}[v]}{\sigma_M^2} - \mathbb{E}_{T,d}[\partial_d \log T(v;t,d)]. \tag{A.133}$$

The first term is nonnegative and is positive whenever the fixed clearing rule is informative about v.

Proof. The score of the normalized public density is $\partial_d \log g_d(v) = [v - \mathbb{E}_{\mathcal{I},d}[v]]/\sigma_M^2$. Differentiating the clearing integral therefore gives the identity. Conditional on silence, each signal distribution is MLR-increasing in v. At fixed bids and strategies, clearing is nonincreasing in v, also after averaging over rival idiosyncratic values. Hence $\mathbb{E}_{T,d}[v] \le \mathbb{E}_{\mathcal{I},d}[v]$, with strict inequality for a nonconstant clearing likelihood. □

The reference mean in (A.133) is the silence-conditioned public mean. Replacing it by the Gaussian mean $\mu(d)$ would omit the derivative of Z_d. The unnormalized mass $f_d(v)q(v)^m T(v;t,d)$ alone is insufficient for a clearing-probability comparison.

To express the response of the clearing rule, let $F_S(x \mid v)$ and $f_S(x \mid v)$ be the cdf and density of a signal conditional on v and $s_k > \bar{s}$. For rival idiosyncratic value z, let $\chi(t,z,d)$ solve $b^S(\chi(t,z,d),z;d) = \beta^T(t;d)$, on the interior range where this inverse is differentiable. Define $A(v;t,d) = \int F_S(\chi(t,z,d) \mid v)\, dF_u(z)$. Then $T = A^m$.

Lemma A34 (Gap response reduced to bid passthrough)**.** *Suppose the relevant cutoffs are interior and differentiation under the idiosyncratic-value integral is valid. Then*

$$\partial_d \log T(v;t,d) = m \frac{\int f_S(\chi(t,z,d) \mid v)\partial_d \chi(t,z,d)\, dF_u(z)}{\int F_S(\chi(t,z,d) \mid v)\, dF_u(z)}, \tag{A.134}$$

where

$$\partial_d \chi(t,z,d) = \frac{\partial_d \beta^T(t;d) - \partial_d b^S(\chi(t,z,d),z;d)}{\partial_x b^S(\chi(t,z,d),z;d)}. \tag{A.135}$$

If $\partial_x b^S > 0$, the pointwise restriction

$$\partial_d b^S(\chi(t,z,d),z;d) \ge \partial_d \beta^T(t;d) \quad \text{for } F_u\text{-almost every } z \tag{A.136}$$

is sufficient for $\partial_d T \le 0$.

Proof. Conditional independence gives $T = A^m$. Differentiate this expression and then the cutoff identity. Positive signal derivatives give the stated sign. □

In the auxiliary scalar auction the integral collapses to a single cutoff $x_c(t,d)$. Its response is $\partial_d \log T = mK_S(x_c \mid v)\partial_d x_c$, where $K_S = f_S/F_S$. In the exact game, the integral in (A.134) remains necessary. More generally, the sufficient gap-response inequality is

$$\mathbb{E}_{T,d}[\partial_d \log T(v;t,d)] < \frac{\mathbb{E}_{\mathcal{I},d}[v] - \mathbb{E}_{T,d}[v]}{\sigma_M^2}. \tag{A.137}$$

Assumption A5 (Gap-response condition)**.** *In the hard-truncation continuation family where a gap comparative static is invoked, the exact clearing probabilities are positive and differentiable on the specified type and gap domain, and (A.137) holds there.*

This condition gives strict pointwise decay. A uniform logarithmic decay rate additionally requires control as the talker type approaches zero and as public histories vary.

Lemma A35 (Uniform talker decay conditional on a uniform gap bound)**.** *Suppose, on the entire type and gap domain used in a rent integral, $-\partial_d \log Q^{\mathrm{ex}}(t;d) \geq \lambda_{\mathrm{gap}} > 0$ for almost every t. Suppose the positive rent is $R(d) = \int a(t)Q^{\mathrm{ex}}(t;d)\,dt$, with $a \geq 0$ independent of d, and differentiation under this integral is valid. Then*

$$-\partial_d \log R(d) \geq \lambda_{\mathrm{gap}}. \tag{A.138}$$

Proof. The negative logarithmic derivative of R is the average of $-\partial_d \log Q^{\mathrm{ex}}$ under the probability density proportional to $a(t)Q^{\mathrm{ex}}(t;d)$. It inherits the lower bound. □

The lemma covers typewise rents with $a(t) = F_u(t)^{n_T-1}\mathbf{1}\{t \leq u\}$ and expected rents with $a(t) = (1-F_u(t))F_u(t)^{n_T-1}$. A continuous positive logarithmic derivative has a positive minimum on a compact domain where clearing remains positive. This observation does not supply a uniform bound at the zero-clearing endpoint or over unbounded public histories.

Let $\pi_0(s) = \mathbb{E}[\Pi^T\mathbf{1}\{n_S = 0\}]/\bar{\Pi}^T(s;p)$ when talking payoff is positive. If the hard-truncation bounds apply uniformly to every mixed history entering the expectation, and the all-talk selected rent is unchanged by a translation of the public mean, then

$$-A_T(s) \geq \lambda_{\mathrm{gap}}\varrho[1-\pi_0(s)]\bar{\Pi}^T(s;p). \tag{A.139}$$

This follows from (A.132) and the uniform rent bound. It bounds the talker term only. The silent public-history term $A_S(s)$ is still needed in $\mathcal{R}(s)$. A general weighted policy also need not admit this hard-truncation reduction.

Returning to the exact weighted model, a sufficient condition for a negative disclosure derivative on the positive region is

$$\tau_\varepsilon \kappa_V(s)\bar{\Pi}^S(s;p) + \varrho(N-1)\,\mathrm{TV}(p)C_J(s) < \mathcal{R}(s)\bar{\Pi}^T(s;p). \tag{A.140}$$

It follows from (A.131) after dropping the favorable variance term. This condition controls the complete derivative rather than only the fixed-history signal comparison.

Remark A4 (Scope of the gap-response condition)**.** *The bid derivatives in* (A.135) *concern selected global best responses. Differentiating first-order conditions provides candidate equations for those derivatives, but does not establish selection, global optimality, or their uniform bounds. The exact cutoffs depend on every rival idiosyncratic value. Neither the scalar approximation nor pointwise gap decay establishes* (A.140) *for the weighted model.*

Integrability of the payoff-weighted objects. Fix a disclosure propensity p satisfying $0 \leq p(s) \leq \bar{u}/K < 1$. Write $w = 1-p$ and $\underline{q} = 1-\bar{u}/K > 0$, so that $\underline{q} \leq w(s) \leq 1$. This lower bound on silence probabilities controls posterior moments directly. A random-cutoff representation remains available, but its mixing distribution must be updated after conditioning on silence.

For the random-cutoff representation only, suppose w is nondecreasing and right-continuous. Let C be an auxiliary cutoff taking values in the extended real line, independent of the auction primitives, and let ν denote its law. Let $\mathbb{P}_0$ be the Gaussian reference law obtained by conditioning on observed signal values before applying the silence weights. If ℓ silent signals remain latent, write $A_c = \{s_i \geq c_i,\ i = 1,\dots,\ell\}$ for a cutoff vector c. A cutoff at $-\infty$ imposes no restriction, while a cutoff at $+\infty$ makes the event empty.

Lemma A36 (Random-cutoff representation)**.** *Every nondecreasing, right-continuous function $w : \mathbb{R} \to [0,1]$ admits a cutoff law ν such that $w(s) = \mathbb{P}(C \le s)$. The masses at $-\infty$ and $+\infty$ are $w(-\infty)$ and $1 - w(\infty)$, respectively. Consequently,*

$$q(v) = \mathbb{E}_C[\mathbb{P}(s \ge C \mid v)]. \tag{A.141}$$

For ℓ silent dealers, take independent cutoffs with prior law $\nu^{\otimes \ell}$. Whenever the silence event has positive probability, the posterior cutoff law is

$$\nu^*(dc) = \frac{\mathbb{P}_0(A_c)\nu^{\otimes \ell}(dc)}{\int \mathbb{P}_0(A_a)\nu^{\otimes \ell}(da)}. \tag{A.142}$$

The weighted posterior is the mixture of the truncated Gaussian laws $\mathbb{P}_0(\cdot \mid A_c)$ under ν^*. Components with $\mathbb{P}_0(A_c) = 0$ receive zero posterior weight.

Proof. Assign mass $w(-\infty)$ to $-\infty$, use the Stieltjes measure of w on $\mathbb{R}$, and assign the remaining mass $1 - w(\infty)$ to $+\infty$. This defines a probability law with $w(s) = \mathbb{E}_C[\mathbf{1}\{C \le s\}]$. Tonelli's theorem gives (A.141).

For any nonnegative measurable function h of the auction primitives, independence of the auxiliary cutoffs gives

$$\mathbb{E}_0\left[h \prod_{i=1}^{\ell} w(s_i)\right] = \int \mathbb{E}_0[h\mathbf{1}_{A_c}]\nu^{\otimes \ell}(dc). \tag{A.143}$$

Divide this identity by its counterpart with $h = 1$. The result is the asserted mixture with mixing law (A.142). □

The cutoffs are independent under their auxiliary prior. Conditioning on silence generally makes them dependent. Moreover, the disclosure-probability bound does not bound the support of their finite components. The posterior moment argument below requires no cutoff moment condition.

Let $Z \sim N(m, \Sigma)$ under the Gaussian reference law, where Z collects the common value and any remaining latent signals. Let $L(Z)$ be the product of the ℓ silence weights, and write $\mathbb{E}_w$ for expectation under the probability law with density $L/\mathbb{E}_0[L]$ relative to the reference law.

Lemma A37 (Gaussian moments under silence weights)**.** *Suppose $\underline{q}^{\ell} \le L(Z) \le 1$. For every nonnegative measurable function h,*

$$\mathbb{E}_w[h(Z)] \le \underline{q}^{-\ell}\mathbb{E}_0[h(Z)]. \tag{A.144}$$

For fixed Σ, fixed dimension, and $\ell \le N$, there are finite constants A_1 and A_2, independent of m and of the particular silence weights, such that

$$\mathbb{E}_w[\|Z\|^r] \le A_r(1 + \|m\|^r), \qquad r \in \{1, 2\}. \tag{A.145}$$

Proof. The numerator $\mathbb{E}_0[hL]$ is at most $\mathbb{E}_0[h]$, while the denominator $\mathbb{E}_0[L]$ is at least $\underline{q}^{\ell}$. This proves (A.144). Under the reference Gaussian law, $\mathbb{E}_0[\|Z\|] \le \|m\| + \sqrt{\operatorname{tr}\Sigma}$ and $\mathbb{E}_0[\|Z\|^2] = \|m\|^2 + \operatorname{tr}\Sigma$. Applying the first inequality proves (A.145). □

Conditional on observed signal values, disclosure probabilities associated with those exact observations cancel from Bayes' rule. The remaining likelihood factors are the silence weights. Thus the lemma applies to the continuation posteriors. At a forced announcement of an exact signal,

use the same Gaussian reference law and the remaining silence weights to specify the continuation belief.

A lower bound on selected silent bids provides one sufficient way to control continuation payoffs. Write η for the Gaussian component of the coupling variable ω, and write $x^- = \max\{-x, 0\}$. For each continuation configuration, let $\mathcal{I}(s, \omega)$ denote the public history generated by the focal signal s and the coupling variables.

Assumption A6 (Lower bound on silent bids)**.** *For every compact signal set $D \subset \mathbb{R}$, there is a finite constant A_D such that*

$$\left[b^S(s, u; \mathcal{I}(s, \omega))\right]^- \leq A_D(1 + \|\eta\|) \tag{A.146}$$

for every $s \in D$ and $u \in [0, \bar{u}]$, almost surely in the coupling variables. The bound holds uniformly over the ordinary continuation configurations and the forced disclosure configurations entering $C_J(s)$.

This is an additional condition on the selected exact continuations. It concerns only the negative part of silent bids. Talker payoffs are already bounded by $\bar{u}$. Their bids need not satisfy an affine bound because, when the talker support is unbounded below, $\beta^T(u; \mathcal{I})$ tends to $-\infty$ as u approaches zero even at a fixed public history.

Lemma A38 (Continuation payoff and selection-gap bounds)**.** *Maintain the band $\underline{q} \leq w \leq 1$, globally optimal continuation bids, nonnegative optimized payoffs, the talker-payoff bound $0 \leq \Pi^T \leq \bar{u}$, and Assumption A6. For every compact signal set D, there is a finite constant B_D such that $0 \leq \Pi^S(s, \omega) \leq B_D(1 + \|\eta\|)$ and*

$$|\bar{V} - V|\Pi^S(s, \omega) \leq B_D(1 + \|\eta\|^2) \tag{A.147}$$

for $s \in D$. The bounds hold in the ordinary and forced continuation configurations.

Proof. Fix a continuation history and a silent dealer's information. For $r \in \{1, 2\}$, let $M_r = \mathbb{E}[|v|^r \mid s, \mathcal{I}]$. The Gaussian reference mean is affine in the observed signals, and its covariance depends only on the identities of the observed signals. There are finitely many such configurations. Lemma A37 therefore gives, uniformly for $s \in D$, $M_1 \leq C_D(1 + \|\eta\|)$ and $M_2 \leq C_D(1 + \|\eta\|^2)$.

Let b be the selected silent bid, let Q be its winning probability, and let V be the common-value mean conditional on winning. The optimized payoff is $\Pi^S = Q(V + u - b)$. If $Q = 0$, set the payoff-weighted terms equal to zero. If $Q > 0$, conditional expectation and Jensen's inequality give $Q|V| \leq M_1$ and $QV^2 \leq M_2$. Hence $\Pi^S \leq M_1 + \bar{u} + b^-$.

Also, $|\bar{V}| \leq M_1$. Since the optimized margin is nonnegative whenever $Q > 0$, expanding $Q(|\bar{V}| + |V|)(|V| + \bar{u} + b^-)$ gives

$$|\bar{V} - V|\Pi^S \leq M_1^2 + M_2 + 2M_1(\bar{u} + b^-). \tag{A.148}$$

Substituting the posterior moment bounds and (A.146) proves both conclusions. This argument controls the payoff-weighted selection gap without requiring a separate affine bound on the mean conditional on winning.

The same calculation applies to the forced configurations because their posterior kernels satisfy the same silence-weight bounds and Assumption A6 covers their selected silent bids. □

Corollary A5 (Finiteness of the payoff-weighted objects)**.** *Maintain the band* $\underline{q} \le w \le 1$*, globally optimal continuation bids, nonnegative optimized payoffs, the talker-payoff bound* $0 \le \Pi^T \le \bar{u}$*, and Assumption A6. Then* $\bar{\Pi}^S(s;p) < \infty$*, and* $\kappa_V(s) < \infty$ *whenever* $\bar{\Pi}^S(s;p) > 0$*. Moreover,* $\sup_{s\in D} C_J(s) < \infty$ *for every compact signal set* D*.*

Proof. The Gaussian vector η has finite first and second moments. The bounds in Lemma A38 therefore make both $\bar{\Pi}^S(s;p)$ and the numerator defining $\kappa_V(s)$ finite. Division by a strictly positive $\bar{\Pi}^S(s;p)$ gives the assertion about $\kappa_V(s)$.

For $C_J(s)$, condition on $s_k = t$ and $s_j = s$. Under the Gaussian coupling, the conditional mean of η is affine in (s,t) and its conditional covariance is fixed. Consequently, $\mathbb{E}[\|\eta\| \mid s_k = t, s_j = s] \le C_D(1+|t|)$ uniformly for $s \in D$. The payoff bounds imply $\mathbb{E}[\Pi^S_{\text{before}} + \Pi^T_{\text{after}} \mid s_k = t, s_j = s] \le C_D(1+|t|)$.

The density $f_{s_k|s}(t)$ is Gaussian with positive variance and a mean bounded as s ranges over D. Thus there are constants $c_D > 0$ and $C_D < \infty$ such that

$$f_{s_k|s}(t)\,\mathbb{E}[\Pi^S_{\text{before}} + \Pi^T_{\text{after}} \mid s_k = t, s_j = s] \le C_D(1+|t|)e^{-c_D t^2}. \tag{A.149}$$

The right-hand side is bounded and tends to zero in both tails. Taking the essential supremum over t, the supremum over $s \in D$, and the maximum over the finitely many rivals proves the claim. □

To use these integrability conclusions, the selected exact continuation must satisfy Assumption A6. On the mixed overlap region, the maintained inequality $b^S(s,0;\mathcal{I}) \ge \psi(s;\mathcal{I})$ provides a starting point through bounds on Gaussian upper-truncation means. The auxiliary scalar integrating-factor formula suggests another possible bound above the talker support, but it does not by itself give a bound for the full idiosyncratic-value game. Its boundary and tail contributions would first need to be controlled and the resulting estimate transferred to exact strategies. The same lower bound must hold in all-silent and one-silent configurations and after forced rival decisions. An approximation valid only on compact signal sets does not supply these bounds across unbounded public histories. Until this continuation estimate is established, the finiteness conclusions for κ_V and C_J remain conditional on Assumption A6.

The random-cutoff identity concerns probability laws. Equilibrium strategies and values must still be solved under the resulting posterior. The moment bounds establish integrability of payoff levels and payoff-weighted means. Differentiability and domination retain the separate requirements in Assumption A4.

Conditions at a disclosure boundary. A unique upper boundary of the positive-talk region requires less than monotonicity of the entire best-response propensity. The following elementary reduction is useful for that distinction.

Proposition A13 (Zeros-only reduction)**.** *Suppose* $\Delta(\cdot;p)$ *is continuous, differentiable at every zero, negative for all sufficiently large signals, and positive somewhere. If every zero* s^* *satisfies*

$$\partial_s \Delta(s^*;p) < 0, \tag{A.150}$$

then there is a unique zero $\bar{s}(p)$ *and* $\{s : \Delta(s;p) > 0\} = (-\infty, \bar{s}(p))$*.*

Proof. A negative derivative at a zero makes the function positive immediately to its left and negative immediately to its right. Two such zeros would require a change from negative to positive between them. By continuity, the first zero after a negative value and before a positive value cannot

have positive values immediately to its left. This contradicts the derivative condition. Eventual negativity and positivity somewhere guarantee a zero. Its local sign and the absence of other zeros give the asserted positive set. □

At a zero with $\bar{\Pi}^T(s^*;p) = \bar{\Pi}^S(s^*;p) = \Pi^* > 0$, the complete derivative identity is

$$\partial_s \Delta(s^*;p) = -\mathcal{R}(s^*)\Pi^* - \tau_\varepsilon \mathbb{E}[W\Sigma_W] + \tau_\varepsilon \kappa_V(s^*)\Pi^* + \varrho \mathbb{E}\left[\sum_{k\neq j} |p'(s_k)| J_k\right]. \tag{A.151}$$

All expectations here are evaluated at s^*. The signed public-information term $\mathcal{R}$ includes the fixed-identity talker response and the movement of silent payoff with rival public messages.

Proposition A14 (Local condition at a candidate zero)**.** *Maintain Assumption A4. At a positive-payoff zero where the expectations in (A.151) are finite, a sufficient condition for a negative derivative is*

$$\tau_\varepsilon \kappa_V(s^*)\Pi^* + \varrho \mathbb{E}\left[\sum_{k\neq j} |p'(s_k)| J_k^+\right] < \mathcal{R}(s^*)\Pi^* + \tau_\varepsilon \mathbb{E}[W\Sigma_W]. \tag{A.152}$$

Proof. Replace each jump in (A.151) by its positive part and rearrange. □

The variance contribution has a lower bound that does not require a scalar cutoff or a lower bound on bids. Set $\underline{\sigma}^2 = 1/(\tau_v + N\tau_\varepsilon)$.

Lemma A39 (Variance floor)**.** *For every positive-probability exact clearing event at a fixed bid,*

$$\mathrm{Var}(v \mid s, \mathcal{I}, \mathcal{C}_b) \geq \underline{\sigma}^2. \tag{A.153}$$

Proof. Condition further on all N signals and all independent cost, idiosyncratic-value, and tie-breaking variables. Conditional on all signals, the common value has Gaussian variance $\underline{\sigma}^2$. Disclosure and clearing depend only on these signals and independent auxiliary variables. Conditioning on them does not reduce the remaining Gaussian variance. The law of total variance under the coarser information set proves the inequality. □

At a positive-payoff zero, define the average winning margin by

$$\bar{m}(s^*) = \frac{\Pi^*}{\mathbb{E}[W]}. \tag{A.154}$$

The denominator is positive because payoff is positive. This is the winning-probability-weighted average of the selected margin.

Proposition A15 (Normalized crossing condition)**.** *At a positive-payoff zero, suppose the exact derivative identity (A.151) holds with finite expectations and the total-variation switch bound applies. Then*

$$\frac{\mathbb{E}[W\Sigma_W]}{\Pi^*} \geq \frac{\underline{\sigma}^2}{\bar{m}(s^*)}. \tag{A.155}$$

A sufficient condition for a negative derivative is therefore

$$\tau_\varepsilon \kappa_V(s^*) + \frac{\varrho(N-1)\,\mathrm{TV}(p) C_J(s^*)}{\Pi^*} < \mathcal{R}(s^*) + \tau_\varepsilon \frac{\underline{\sigma}^2}{\bar{m}(s^*)}. \tag{A.156}$$

Proof. The variance floor gives $\mathbb{E}[W\Sigma_W] \geq \underline{\sigma}^2\mathbb{E}[W]$. Divide by Π^* and use (A.154). Then apply the total-variation switch bound to (A.152) and divide by Π^*. □

Proposition A16 (Threshold result at small idiosyncratic values)**.** *Consider a sequence with fixed Gaussian primitives and $\bar{u} \downarrow 0$. Maintain the exact derivative identity and finite expectations at every candidate zero. Suppose all such zeros have $\Pi^* > 0$. Uniformly over these zeros, suppose $\mathbb{E}[W] \geq \eta > 0$, κ_V and $\mathrm{TV}(p)C_J/\Pi^*$ are bounded, and $\mathcal{R}$ is bounded below. Then every zero has a negative derivative for sufficiently small $\bar{u}$.*

Proof. At a zero, $\Pi^* = \bar{\Pi}^T \leq \mu_u \leq \bar{u}$, so $\bar{m} \leq \bar{u}/\eta$. The favorable term $\tau_\varepsilon\underline{\sigma}^2/\bar{m}$ diverges at least as fast as $\tau_\varepsilon\underline{\sigma}^2\eta/\bar{u}$. The adverse terms on the left of (A.156) remain bounded, and its signed public-information term is bounded below. The sufficient inequality follows. □

The bound on $\mathcal{R}$ is substantive. A small level of talker rent does not bound its derivative, and it does not control the silent public-message response. Thus one cannot omit the talker and public-history terms solely because $\bar{u}$ is small. The positive lower bound on $\mathbb{E}[W]$ is also additional to indifference, since both payoff and winning probability could otherwise vanish.

Proposition A17 (Threshold disclosure from local continuation inequalities)**.** *Suppose $\Delta(\cdot; p)$ is continuous, differentiable at every zero, and positive somewhere. Maintain the exact derivative identity with finite expectations, and assume every zero has positive common continuation payoff. Suppose the all-silent conditions and fixed policy band in Proposition A12 hold. If every zero satisfies (A.152), or its sufficient version (A.156), then the positive-talk region is a lower signal interval with a unique finite upper boundary.*

Proof. The band argument gives eventual negativity. The local inequality gives a negative derivative at every zero. Proposition A13 completes the proof. □

Scope of the disclosure restriction. The Gaussian coupling, exact fixed-bid scores, variance floor, and posterior moment bounds are analytical consequences of the information structure. Upper-tail deactivation additionally uses the stated all-silent continuation properties. The remaining single-crossing conditions concern selected continuation values, their public-history response, and the gains at forced rival switches. Their integrability and regularity require the stated conditions. The hard-truncation gap comparison is one component of this analysis and does not establish the weighted-model inequality.

For smooth policies, (A.140) is sufficient for strict decline throughout the positive region when the value is locally absolutely continuous. A condition imposed only at zeros yields an upper boundary but need not make the full best-response propensity nonincreasing. The equilibrium result uses the stronger order condition on the entire policy domain.

Define $\mathcal{P}_{\bar{u}} = \{p : \mathbb{R} \to [0, \bar{u}/K] : p \text{ is nonincreasing}\}$, with the pointwise order. The domain includes discontinuous monotone policies. In addition to the continuation selection in Assumption A1, we maintain the following restrictions on exact disclosure values.

Assumption A7 (Continuation dominance and best responses)**.** *For every $p \in \mathcal{P}_{\bar{u}}$, the selected continuation values at natural and required forced histories are defined, and the ex ante disclosure values are finite. Moreover:*

(i) $\Delta(\cdot; p)$ is continuous and negative for all sufficiently high signals. If $s_1 < s_2$ and $\Delta(s_2; p) > 0$, then $\Delta(s_1; p) > \Delta(s_2; p)$.

(ii) The disclosure best response is order preserving: $p \leq \widetilde{p}$ implies $G_\kappa(\Delta(s;p)) \leq G_\kappa(\Delta(s;\widetilde{p}))$ for every s.

Part (i) states continuation dominance directly on the full lattice. The preceding differential inequalities provide sufficient tests on smooth subclasses, but their extension to discontinuous policies is not supplied by Gaussian affiliation or by a compact-set approximation. Part (ii) is the separate strategic-complementarity condition needed for the order argument. Neither restriction follows from uniqueness of a bidder's response to fixed rivals.

A.3.4 Selective disclosure equilibrium

We state the disclosure characterization with its selection and order conditions. The invariant-band estimate gives $\bar{\Pi}^T(s;p) \leq \mu_u$ and hence $\mathcal{B}(p)(s) \leq \mu_u/K$, obtained by averaging the typewise bound $\Pi^T(u) \leq u$. We retain the larger band $\mathcal{P}_{\bar{u}}$ to match the main statement.

Proof of Proposition 2. The payoff difference between talking and silence is $\Delta(s;p) - \kappa$. Integrating the cost threshold gives $\mathcal{B}(p)(s) = G_\kappa(\Delta(s;p))$. Lemma A7 implies $0 \leq \mathcal{B}(p)(s) \leq \bar{u}/K < 1$. Assumption A7(i) implies that the positive region of Δ is a lower interval and that Δ is strictly decreasing there. Its positive part is nonincreasing on the whole line. Thus $\mathcal{B}(p) \in \mathcal{P}_{\bar{u}}$.

The domain $\mathcal{P}_{\bar{u}}$ is a complete lattice. Pointwise suprema and infima of arbitrary nonempty families preserve monotonicity and the band. The constant policies zero and $\bar{u}/K$ are its bottom and top elements. Assumption A7(ii) therefore permits Tarski's theorem to be applied to this self-map. There are least and greatest fixed points. At any fixed point, the canonical cost threshold is $\bar{\kappa}(s) = \Delta(s;p^*)$. Continuity and eventual negativity imply that any nonempty positive region has a finite upper boundary $\bar{s}^*$ and equals $(-\infty, \bar{s}^*)$. Since $p^* < 1$ everywhere, every signal with positive talk probability also has positive silence probability.

For the signal ordering, let S' be an independent copy of the generic signal S. When p^* is nonconstant on a set of positive signal probability, $\mathrm{Cov}(S, p^*(S)) = \mathbb{E}[(S-S')(p^*(S)-p^*(S'))]/2 < 0$. Provided both disclosure events have positive probability, Bayes' rule gives

$$\mathbb{E}[S \mid \text{Talk}] = \mathbb{E}[S] + \frac{\mathrm{Cov}(S, p^*(S))}{\mathbb{E}[p^*(S)]}, \qquad \mathbb{E}[S \mid \text{Silent}] = \mathbb{E}[S] - \frac{\mathrm{Cov}(S, p^*(S))}{1 - \mathbb{E}[p^*(S)]}. \tag{A.157}$$

This proves the strict mean ordering. □

Tarski's theorem gives existence conditional on the selection and order assumptions. It does not guarantee that the greatest equilibrium is nontrivial. A sufficient condition for nontriviality is $\Delta(s;0) > 0$ at some signal. The greatest fixed point satisfies $p^* \geq \mathcal{B}(0)$ by order preservation, and continuity then gives talking on a set of positive signal probability. This payoff condition remains to be established if nontriviality is claimed for a particular parameterization.

Lemma A40 (Collective silence is not self-enforcing)**.** *Maintain the conditions of Proposition 2 and suppose the selected continuation games define $\Delta(s;0)$, with $\Delta(s;0) > 0$ on a set of signals with positive probability. Then the all-silent policy $p = 0$ is not an equilibrium. Any profile in which every dealer stays silent with probability one admits a strictly profitable unilateral deviation for a positive-probability set of signal and cost types.*

Proof. Fix a profile in which every rival stays silent with probability one. A dealer with signal s and talking cost κ who deviates to talking gains $\Delta(s;0) - \kappa$, since rival behavior coincides with $p = 0$. With $\kappa \sim U[0, K]$, the deviation is strictly profitable on the event $\kappa < \Delta(s;0)$, which has

probability $G_\kappa(\Delta(s;0)) = \Delta(s;0)/K > 0$ on the assumed signal set. Hence $\mathcal{B}(0) \neq 0$ and $p = 0$ is not a fixed point of $\mathcal{B}$. □

The lemma restates the nontriviality condition as an enforcement statement used in Section 5.3. Collective silence admits a profitable unilateral deviation exactly when equilibrium talk is nontrivial.

The continuous monotone policies do not in general form a complete lattice. They may nevertheless be used for contraction because they form a closed set under the sup norm. The distinction between these domains is substantive.

Local uniqueness and iteration. Let $C_b(\mathbb{R})$ be the Banach space of bounded continuous real functions with the sup norm, and let $\mathcal{P}^c_{\bar{u}} = \mathcal{P}_{\bar{u}} \cap C_b(\mathbb{R})$. For a fixed point $p^* \in \mathcal{P}^c_{\bar{u}}$ and $r > 0$, put $\mathcal{N}_r = \{p \in \mathcal{P}^c_{\bar{u}} : \|p - p^*\|_\infty \leq r\}$.

Proposition A18 (Sufficient local contraction of disclosure)**.** *Suppose $\mathcal{B}$ is defined on $\mathcal{N}_r$ and takes values in $\mathcal{P}^c_{\bar{u}}$. Suppose a finite S^* satisfies $\mathcal{B}(p)(s) = 0$ for every $p \in \mathcal{N}_r$ and $s \geq S^*$. If $L < K$ and*

$$\sup_{s \leq S^*} |\Delta(s;p_2) - \Delta(s;p_1)| \leq L\|p_2 - p_1\|_\infty \qquad (p_1, p_2 \in \mathcal{N}_r), \tag{A.158}$$

then $\mathcal{B}$ maps $\mathcal{N}_r$ into itself and contracts there with modulus at most L/K. The fixed point is unique in $\mathcal{N}_r$, and iteration from any $p_0 \in \mathcal{N}_r$ satisfies $\|\mathcal{B}^n(p_0) - p^*\|_\infty \leq (L/K)^n\|p_0 - p^*\|_\infty$.

Proof. The uniform cost cdf is globally $1/K$-Lipschitz, including its two kinks. Apply this bound and (A.158) below S^*, and use the common zero region above S^*. This gives $\|\mathcal{B}(p_2) - \mathcal{B}(p_1)\|_\infty \leq (L/K)\|p_2 - p_1\|_\infty$. Comparing with p^* makes the closed ball invariant. The monotonicity and band constraints are closed under uniform convergence, so $\mathcal{N}_r$ is complete. Banach's theorem gives uniqueness and geometric convergence. □

A bound on the whole signal line is also sufficient. The restricted version avoids requiring a comparison where both disclosure best responses are already zero. Proposition A12 supplies a common zero region when its uniform all-silent conditions hold. Preservation of the continuous policy class and the quantitative gain bound remain additional requirements.

What the Lipschitz estimate requires. For a precise decomposition, distinguish three policies. The composition policy c generates the rivals' disclosure decisions. The belief policy b determines silence-conditioned posteriors at a realized history. The strategy policy r selects the rival bidding environment, held fixed when beliefs are changed. Let $\Phi_s(c, b, r)$ be the expected optimized gain from talking under these choices, using the same Gaussian signal law conditional on s and the specified beliefs at forced announcements. These comparison values are defined only when both optimized action values and their expectations are finite. At a common policy, $\Phi_s(p, p, p) = \Delta(s; p)$.

For two policies p_1, p_2, define the belief, strategy, and composition differences, respectively, by

$$D_{\text{bel}}(s) = \Phi_s(p_1, p_2, p_1) - \Phi_s(p_1, p_1, p_1), \qquad D_{\text{str}}(s) = \Phi_s(p_1, p_2, p_2) - \Phi_s(p_1, p_2, p_1), \tag{A.159}$$

and $D_{\text{comp}}(s) = \Phi_s(p_2, p_2, p_2) - \Phi_s(p_1, p_2, p_2)$. Telescoping gives

$$\Delta(s;p_2) - \Delta(s;p_1) = D_{\text{bel}}(s) + D_{\text{str}}(s) + D_{\text{comp}}(s). \tag{A.160}$$

Both the direct belief and strategy comparisons are integrated under the single natural history law generated by p_1. The composition term changes that law while holding the continuation fixed.

The latter term has an exact finite-switch representation. Fix a continuation policy r and write $g_r(s; J, \mathbf{t})$ for its idiosyncratic-value-averaged gain when the set of rival talkers is J, with signals $\mathbf{t} = (t_1, \ldots, t_M)$ and $M = N - 1$. Only the disclosed coordinates enter public information. Let $\Delta_k g_r$ be the gain with rival k forced to talk minus the gain with their forced to be silent. For the kth comparison, let $\mathbb{E}^{\mathrm{hyb}}_{k,s}$ use policy p_2 for rivals $1, \ldots, k-1$, policy p_1 for rivals $k+1, \ldots, M$, and force the two actions of rival k. All signals retain their Gaussian law conditional on the focal signal. Define $S_{k,T}$ and $S_{k,S}$ as the focal silent continuation values in those two forced histories.

Lemma A41 (Composition by one-rival switches)**.** *At a fixed continuation r with finite comparison values,*

$$\Phi_s(p_2, r, r) - \Phi_s(p_1, r, r) = \sum_{k=1}^{M} \mathbb{E}^{\mathrm{hyb}}_{k,s}[(p_2(t_k) - p_1(t_k))\Delta_k g_r]. \tag{A.161}$$

Suppose its expected talker values lie in $[0, \mu_u]$ and its silent values are nonnegative. If

$$\mathbb{E}^{\mathrm{hyb}}_{k,s}[\max\{S_{k,T}, S_{k,S}\}] \le A_F$$

uniformly over the required policies, rivals, and $s \le S^*$, then the composition modulus can be taken as

$$L_{\mathrm{comp}} = M(\mu_u + A_F). \tag{A.162}$$

Proof. Conditional on the signal vector, disclosure decisions are independent Bernoulli variables. Replace their probabilities one at a time. Each difference is exactly the probability change for that rival multiplied by the forced gain difference, averaged over the remaining hybrid decisions. This proves (A.161). The difference between two numbers in $[0, \mu_u]$ is at most μ_u in absolute value. The difference between two nonnegative silent values is at most their maximum. Thus $|\Delta_k g_r| \le \mu_u + \max\{S_{k,T}, S_{k,S}\}$. Bound each probability difference by $\|p_2 - p_1\|_\infty$ and sum. □

Consequently, integrated bounds $|D_{\mathrm{bel}}| \le L_{\mathrm{bel}}\|p_2 - p_1\|_\infty$ and $|D_{\mathrm{str}}| \le L_{\mathrm{str}}\|p_2 - p_1\|_\infty$, uniform for $s \le S^*$ and arbitrary policy pairs in the chosen class, give (A.158) with $L = L_{\mathrm{bel}} + L_{\mathrm{str}} + M(\mu_u + A_F)$.

The required envelopes are integrated over the relevant history laws. A history-independent payoff bound need not hold when disclosed Gaussian signals are unbounded. Also, a forced talk in (A.161) carries no factor $p(t_k)$. A bound averaged only under naturally occurring disclosures does not control that term. Although $\Pr(|\mathcal{T}_{-j}| \ge 1 \mid s) \le MR$ on a policy band $0 \le p \le R$, this rarity bound applies to the natural history law and cannot be used to suppress a forced talk.

The two-dealer results in Sections A.2.16 and A.2.17 give an explicit approximate continuation, measurable finite-grid equilibrium selection, uniform bid moments, and linear policy-induced regret. The regret bound compares one prescription under two policies. The strategy channel compares two exact selected equilibrium prescriptions. Passing from the first comparison to the second requires an additional stability estimate and identification of the selected branch. A uniform approximation error independent of $\|p_2 - p_1\|_\infty$ also cannot supply the required pairwise Lipschitz bound as the policies approach one another.

For completeness, a zero-centered class permits a conditional contraction result without first finding a fixed point. For $R \in (0, 1)$, define $\mathcal{Q}_R = \{p \in C_b(\mathbb{R}) : 0 \le p \le R,\ p \text{ is nonincreasing}\}$.

Corollary A6 (A conditional existence certificate on a fixed band)**.** *Suppose selected exact continuation values are defined on $\mathcal{Q}_R$, satisfy $0 \le \bar{\Pi}^T \le \mu_u$ and $\bar{\Pi}^S \ge 0$, and generate continuous nonincreasing disclosure best responses. Suppose there is a common upper region of zero disclosure*

and a finite gain modulus L_R below it, valid for every pair in $\mathcal{Q}_R$. If

$$K > \max\{\bar{u}, \mu_u/R, L_R\}, \tag{A.163}$$

then $\mathcal{B}$ has a unique fixed point in $\mathcal{Q}_R$, lying in $\mathcal{Q}_{\mu_u/K}$, and iteration from any policy in $\mathcal{Q}_R$ converges geometrically with modulus at most L_R/K.

Proof. The payoff bound gives $0 \leq \mathcal{B}(p) \leq \mu_u/K < R$. Together with the assumed continuity and monotonicity, this makes $\mathcal{Q}_R$ invariant. It is a nonempty closed subset of $C_b(\mathbb{R})$. The cost-cdf argument gives a contraction with modulus $L_R/K < 1$. Banach's theorem and the same level bound give all conclusions. □

The continuation auction does not depend directly on the talking-cost scale K once the policy is fixed. Thus, if the continuation conditions and L_R are established on a fixed band independently of K, the displayed inequality is a one-point sufficient condition in the primitives. It does not establish those continuation conditions. This certificate isolates the outstanding quantitative requirements. Exact continuation existence and a compatible selection, integrated pairwise strategy estimates, forced-history payoff envelopes, and preservation of the policy class must all be established before checking the total against K. The paper's equilibrium existence result therefore remains conditional on Assumption A7, and local uniqueness remains conditional on (A.158). The greatest Tarski equilibrium remains the stated selection for the comparative statics.

A.3.5 Hard-truncation entry margin

The disclosure and bidding boundaries admit a direct comparison in the separate hard-truncation family. The comparison does not require that the conjectured deterministic disclosure cutoff be an equilibrium cutoff.

Proposition A19 (Disclosure cutoff and bidding entry margin)**.** *Fix a hard-truncation continuation with silent support $(\bar{s}^*, \infty)$, at least one talker and one silent dealer, and $\bar{u} > 0$. Suppose the exact zero-value silent slice is continuous and strictly increasing below the top talker bid $\bar{b}$, reaches $\bar{b}$ at a finite signal $s^\times$, and satisfies the common lower-limit condition in Assumption A2(iii). Then*

$$s^\times(\mathcal{I}^{\text{out}}) > \bar{s}^*. \tag{A.164}$$

Proof. The common lower limit satisfies $\lim_{s \downarrow \bar{s}^*} b^S(s, 0) = \underline{b}_H < \bar{b}$. Continuity therefore leaves a nonempty interval above $\bar{s}^*$ on which the silent bid remains below $\bar{b}$. The stated finite crossing must occur strictly above the lower support boundary. □

In the weighted baseline, silence has full signal support. Its upper disclosure boundary $\bar{s}^*$ is therefore not a lower boundary of silent signals, and (A.164) does not transfer to that model. The baseline entry margin is established separately in the bidding appendix under its stated continuation conditions.

A.3.6 Winner identity and symmetry among talkers

Proof of Proposition 4. Let B_T and B_S be the highest talker and silent bids. Conditional on $\mathcal{I}$, talker bids depend on independent idiosyncratic values, so the law of B_T does not depend on v. A silent signal given $(v, \mathcal{I})$ has density proportional to $w(s) f_\varepsilon(s - v)$. Between two common values, its

likelihood ratio is proportional to $\exp\{(v'-v)s/\sigma_\varepsilon^2\}$, since the silence weight cancels. It is therefore MLR-increasing in v.

At each idiosyncratic value, a signal-increasing silent strategy maps this signal order into first-order stochastic dominance of bids. Conditional independence of rival signals given v and independence of idiosyncratic values imply that the maximum silent bid also increases in this order. Hence $\ell_T(v) = \Pr(B_T > B_S \mid v, \mathcal{I})$ is nonincreasing in v. Under atomlessness, $\ell_S(v) = 1 - \ell_T(v)$ is nondecreasing.

For winner events $\mathcal{W}_T$ and $\mathcal{W}_S$ of positive probability, Bayes' rule gives

$$\frac{f(v \mid \mathcal{W}_a, \mathcal{I})}{f(v \mid \mathcal{I})} = \frac{\ell_a(v)}{\Pr(\mathcal{W}_a \mid \mathcal{I})}, \qquad a \in \{T, S\}. \tag{A.165}$$

The likelihood ratios give the claimed MLR ordering and its weak mean inequalities. Strict mean inequalities follow if the winning likelihood is nonconstant. This occurs when the talker distribution puts positive mass on bids that correspond to interior silent signal cutoffs for a positive-measure set of silent idiosyncratic values, with the silent strategy strictly increasing at those cutoffs. This is the overlap qualification used in the main text. □

These results concern the latent common value. They do not by themselves order transaction prices conditional on winner identity.

Let m_j be talker j's disclosed signal and let $\bar{S} = n_T^{-1} \sum_{j \in \mathcal{T}} m_j$. When signals are distinct, let R_J denote the ascending rank of the winning talker's signal among the n_T talks, and define its normalized rank by $U_J = (R_J - 1/2)/n_T$.

Proposition A20 (Symmetry within the talker side)**.** *Fix a public history with $n_T \geq 1$, an identical strictly increasing strategy for all talkers, independent identically distributed continuous idiosyncratic values, and $\Pr(\mathcal{W}_T \mid \mathcal{I}) > 0$. Conditional on a talker win, every talker wins with probability $1/n_T$. For distinct disclosed signals, R_J is discrete uniform on $\{1, \ldots, n_T\}$, and*

$$\Pr\left(U_J = \frac{r - 1/2}{n_T} \,\middle|\, \mathcal{W}_T, \mathcal{I}\right) = \frac{1}{n_T} \quad (r = 1, \ldots, n_T), \qquad \mathbb{E}[U_J \mid \mathcal{W}_T, \mathcal{I}] = \frac{1}{2}. \tag{A.166}$$

The winning talker's expected disclosed signal is $\mathbb{E}[m_J \mid \mathcal{W}_T, \mathcal{I}] = \bar{S}$.

Proof. For each talker, the probability of winning is $\int_0^{\bar{u}} F_u(u)^{n_T - 1} Q^{\mathrm{ex}}(u; \mathcal{I})\, dF_u(u)$. This expression does not depend on their identity or on which disclosed signal is theirs. The individual winning events are disjoint and exhaust $\mathcal{W}_T$, so their conditional probabilities equal $1/n_T$. With distinct messages, the same uniform probabilities apply to ranks. Averaging their normalized values gives $1/2$, and averaging the messages gives $\bar{S}$. □

The rank distribution is discrete at a fixed talker count. Pooling auctions with different counts does not turn it into a continuous uniform law.

Corollary A7 (Selection on post-disclosure idiosyncratic values)**.** *Under the symmetry conditions above, suppose F_u has density f_u. Conditional on $(\mathcal{W}_T, \mathcal{I})$, the winning talker's idiosyncratic value has density proportional to $f_u(u) F_u(u)^{n_T - 1} Q^{\mathrm{ex}}(u; \mathcal{I})$. This distribution weakly MLR-dominates F_u, with a nonconstant likelihood ratio whenever the individual winning probability is nonconstant in u on a set of positive F_u probability.*

Proof. For a given talker, the probability of winning is $D(u) = F_u(u)^{n_T - 1} Q^{\mathrm{ex}}(u; \mathcal{I})$. Bayes' rule weights their idiosyncratic-value density by $D(u)$. Symmetry gives the same conditional law for

the winning identity. The co-talker factor is nondecreasing, and the silent-clearing factor is nondecreasing because the talker's bid increases with u. Their product is the likelihood ratio up to normalization. □

A.3.7 Small-idiosyncratic-value limit under hard truncation

The following bounds concern a fixed hard-truncation continuation family. They do not give rates for the weighted baseline or for the probability of disclosure at its endogenous fixed point.

Proposition A21 (Small-$\bar{u}$ upper bounds)**.** *Fix a compact set D of hard-truncation gaps and counts $n_T \geq 1$, $n_S \geq 1$. Suppose the exact selected continuations have zero lowest-type talker value, strictly increasing talker bids, and silent bids nondecreasing in idiosyncratic value. Suppose the exact zero-value silent slice has a finite crossing $s^\times$ of the top talker bid and satisfies Assumption A2. At that crossing suppose $\psi(s^\times) + \bar{u} - \beta^T(\bar{u}) \geq 0$. Then, for sufficiently small $\bar{u}$, uniformly over $d \in D$ and with $\lambda_D = B_D/a_D$:*

(i) $0 < s^\times - \bar{s} \leq \bar{u}/a_D$.

(ii) For each silent dealer, $\Pr(\bar{s} < s_k < s^\times \mid \mathcal{I}) \leq \lambda_D \bar{u}$.

(iii) $\Pr(\mathcal{W}_T \mid \mathcal{I}) \leq (\lambda_D \bar{u})^{n_S}$.

(iv) $R^T(n_T, n_S, d) \leq \mu_u (\lambda_D \bar{u})^{n_S} \leq \lambda_D^{n_S} \bar{u}^{n_S+1}$.

Proof. At the crossing, the common-cutoff margin gives $b^S(s^\times, 0) - \psi(s^\times) \leq \bar{u}$. The common lower bid limit places the crossing strictly above $\bar{s}$. If it were at least δ_0 above $\bar{s}$, monotonicity of the overbid and its local lower bound would give an overbid of at least $a_D\delta_0$. For $\bar{u} < a_D\delta_0$ this is impossible. The local bound then gives $a_D(s^\times - \bar{s}) \leq \bar{u}$, proving part (i). Integrating the density bound over this interval proves part (ii).

For part (iii), a talker win requires every exact silent bid to be below the top talker bid. Because $b^S(s, u) \geq b^S(s, 0)$, each such clearing event requires $s < s^\times$. Apply the chain rule to these exact bid-clearing events. After conditioning on the preceding rivals' clearing events, Assumption A2(ii) bounds the next rival's probability of lying in $(\bar{s}, s^\times)$ by $\lambda_D \bar{u}$. Thus the joint exact clearing probability is at most $(\lambda_D \bar{u})^{n_S}$. This argument uses the conditional density assumption rather than multiplying marginal signal probabilities.

The same bound applies to $Q^{\text{ex}}(t)$ for every talker type because no assigned bid exceeds the top bid. In the rent envelope, $F_u(t)^{n_T-1} \leq 1$ and $\int_0^{\bar{u}} (1 - F_u(t))\, dt = \mu_u$. This proves part (iv). □

Remark A5 (Upper bounds only)**.** *A matching lower bound for $s^\times - \bar{s}$ would require the top common-cutoff margin to leave a bid increment of order $\bar{u}$, together with the upper overbid envelope. Matching lower bounds for winning probabilities would also require lower bounds on the relevant conditional densities and control of the rival idiosyncratic-value cutoffs. The stated assumptions provide upper bounds only.*

A.4 Testable predictions

This subsection connects the baseline model to the empirical tests. Predictions 1–4 in Section 4 state the tested implications in terms of observables. The statements here develop the formal mappings behind them and add restrictions that are not tested in the main text. The Gaussian information structure and symmetry among talkers give exact restrictions on posterior beliefs and

winner identity. Predictions for transaction prices, regression fit, bid dispersion, and responses to economic proxies require additional conditions. Throughout, equilibrium statements retain the continuation selection and disclosure assumptions in Appendices A.2 and A.3.

In the model, N is the total number of participating dealers and n_T is the number who talk. In the empirical notation, auction i has a set $\mathcal{T}_i$ of observed talkers, $N_i = |\mathcal{T}_i|$ talks, reported values S_{ij}, consensus $\bar{S}_i$, and transaction price P_i. Thus N_i is a talk count, not the model's total bidder count. For $N_i \geq 1$, let $S_i^{1:N_i} = \max_{j \in \mathcal{T}_i} S_{ij}$ denote the highest talk. Let $\mathcal{W}_{T,i}$ and $\mathcal{W}_{S,i}$ denote a talker win and a silent-dealer win. The distinction between talk counts and total participation matters when interpreting competition and the number of silent dealers.

The theoretical conditioning variable is the full public history $\mathcal{I}$, including the disclosure identities and values. Empirical controls approximate this information. Conditioning only on consensus and the talk count can leave differences in public information across the two winner groups. Conditional mean comparisons require integrable outcomes and positive probabilities for the groups being compared. The price tests also require a comparable observed transaction sample. Reserve acceptance and selective reporting introduce additional selection beyond the baseline auction.

A.4.1 Consensus predictiveness

For $n_T \geq 1$, write the Gaussian reference mean as $\mu_T = \alpha\mu_0 + \beta\bar{S}$, where $\alpha = \tau_v/(\tau_v + n_T\tau_\varepsilon)$ and $\beta = n_T\tau_\varepsilon/(\tau_v + n_T\tau_\varepsilon)$. This reference calculation conditions on the disclosed signal values before applying the silence likelihoods. The consensus weight lies in $(0, 1)$ and increases with the number and precision of disclosed signals. The full public posterior also uses silence, as in (9).

The positive response of the public posterior to consensus survives that additional conditioning. To see this, fix the disclosure identities, the policy p, and the counts (n_T, n_S). Write $q(v) = \int [1 - p(t)] f_\varepsilon(t - v)\, dt$, let $\sigma_M^2 = (\tau_v + n_T\tau_\varepsilon)^{-1}$, and let $\widetilde{\mu}_T = \mathbb{E}[v \mid \mathcal{I}]$. The prescribed public density is proportional to the Gaussian density with mean μ_T and variance σ_M^2, multiplied by $q(v)^{n_S}$.

Lemma A42 (Public posterior response)**.** *Suppose silence probabilities are bounded away from zero. At fixed policy, disclosure identities, and counts, each disclosed signal m_j satisfies*

$$\frac{\partial \widetilde{\mu}_T}{\partial m_j} = \tau_\varepsilon \operatorname{Var}(v \mid \mathcal{I}) > 0, \qquad \frac{\partial \widetilde{\mu}_T}{\partial \bar{S}} = n_T\tau_\varepsilon \operatorname{Var}(v \mid \mathcal{I}) > 0. \tag{A.167}$$

The consensus derivative varies consensus while holding the deviations of individual talks from consensus fixed.

Proof. With the policy and counts fixed, the silence likelihood is independent of μ_T. The score of the normalized public density with respect to μ_T is $(v - \widetilde{\mu}_T)/\sigma_M^2$. Differentiating its mean therefore gives $\partial\widetilde{\mu}_T/\partial\mu_T = \operatorname{Var}(v \mid \mathcal{I})/\sigma_M^2$. Each message changes the Gaussian mean by $\tau_\varepsilon\sigma_M^2$, and an equal change in all n_T messages changes it by $n_T\tau_\varepsilon\sigma_M^2$. The chain rule gives both identities. The bounded likelihood and Gaussian tails justify differentiation, and the posterior has positive variance. □

Thus a higher consensus raises the public common-value assessment at fixed policy and counts. This result concerns a posterior mean. A transaction price also reflects the response of bidding strategies, idiosyncratic values, and winner selection.

Prediction 5 (Consensus predictiveness)**.** *Conditional on the number of talks and other pre-bid covariates, transaction prices increase with consensus. Consensus explains a larger share of price variation in auctions won by talkers than in auctions won by silent dealers. Its predictive contribution increases with the number of talks.*

The all-talk continuation gives a useful exact benchmark for the first prediction. In the selected symmetric private-value auction, the public mean μ_T is added to every valuation and bid. The remaining component of the winning price depends only on the idiosyncratic values, whose distribution is independent of the disclosed signals. At a fixed total dealer count, $\mathbb{E}[P \mid \mathcal{I}, n_S = 0] = \mu_T + c_N$ for a constant c_N determined by that private-value auction. The consensus loading in this benchmark is β.

With silent rivals, the empirical loading need not equal either β or the derivative in (A.167). A talker's bid contains no undisclosed signal about v, whereas a silent dealer's bid reflects their retained signal. This distinction provides a reason for consensus to explain less of the price variation when a silent dealer wins. It does not order regression fit or coefficients without further restrictions. Although talker idiosyncratic values are independent of v before bidding conditional on $\mathcal{I}$, conditioning on a talker win can make them dependent through the event of defeating silent rivals.

Price loadings and regression fit should therefore be assessed separately. More talks increase the precision of the Gaussian reference information, but endogenous changes in the disclosed signals, the silent side, and bid shading can also change both statistics. At fixed signal precision, β falls when prior precision τ_v rises. A larger consensus loading for senior securities cannot consequently be attributed solely to lower fundamental variance. Relatively precise dealer signals, disclosure selection, and the contribution of retained private information can also matter.

A.4.2 Price decomposition

Fix $n_T \geq 2$ and distinguish one disclosed signal m_j from the average $\bar{S}_{-j}$ of the other $n_T - 1$ talks. In an auction won by a talker, j can be the winning dealer. Since $\bar{S} = [m_j + (n_T - 1)\bar{S}_{-j}]/n_T$, the Gaussian reference mean satisfies

$$\mu_T = \alpha\mu_0 + \frac{\beta}{n_T} m_j + \frac{\beta(n_T - 1)}{n_T} \bar{S}_{-j}. \tag{A.168}$$

For a silent dealer with private signal s, define $\tau_S = \tau_v + (n_T + 1)\tau_\varepsilon$, $\alpha' = \tau_v/\tau_S$, $\beta' = n_T\tau_\varepsilon/\tau_S$, and $\gamma' = \tau_\varepsilon/\tau_S$. Their reference mean is

$$\mu_S(s, \mathcal{M}) = \alpha'\mu_0 + \frac{\beta'}{n_T} m_j + \frac{\beta'(n_T - 1)}{n_T} \bar{S}_{-j} + \gamma' s. \tag{A.169}$$

Here j still indexes a talker, and s is the separate retained signal of the silent dealer. There is no observed own talk for a silent winner.

These identities give equal weight to each disclosed signal. If the other talks enter as an average, the coefficients on m_j and $\bar{S}_{-j}$ have ratio $1 : (n_T - 1)$. If they enter as a sum, their coefficients are equal. Lemma A42 also gives equal marginal effects of individual reports on the full public mean at fixed policy and counts. The latter response can be nonlinear in consensus.

Prediction 6 (Price decomposition)**.** *For auctions won by talkers, the transaction-price decomposition is consistent with a public-information component that assigns equal weight to each disclosed signal. With the remaining talks entered as an average, this component has coefficient ratio* $1 : (n_T - 1)$*. Prices in auctions won by silent dealers additionally reflect information retained by the winner.*

The exact ratio belongs to the Gaussian location parameters in (A.168) and (A.169). A regression of transaction prices need not reproduce it. Winner selection, the winner's-curse adjustment,

and shading can depend on public information, and a linear projection can approximate these nonlinear responses differently across regressors.

In the empirical decomposition, a convenient pair of regressors is the winning talker's signal divided by N_i and the sum of the remaining talks divided by N_i. Equal coefficients then correspond to a price component proportional to consensus. Because the consensus weight itself varies with N_i, pooling counts also requires interactions with the count or an explicit restriction on how the loading varies. Scaling the regressors alone does not make that loading constant. Auctions with one talk have no remaining-talk average, and auctions without talks do not enter this decomposition.

A.4.3 Belief convergence

The Gaussian reference posterior has variance

$$\sigma_M^2 = \frac{1}{\tau_v + n_T \tau_\varepsilon}. \tag{A.170}$$

This variance decreases with the number of disclosed signals. It describes learning from the signal values before the additional information in silence is applied.

A more general learning statement uses an actual refinement of information. Let $\mathcal{H} \subseteq \mathcal{H}'$ be nested information sets about the same common value. The law of total variance gives

$$\mathbb{E}[\mathrm{Var}(v \mid \mathcal{H}') \mid \mathcal{H}] = \mathrm{Var}(v \mid \mathcal{H}) - \mathrm{Var}(\mathbb{E}[v \mid \mathcal{H}'] \mid \mathcal{H}) \leq \mathrm{Var}(v \mid \mathcal{H}). \tag{A.171}$$

Additional information reduces posterior uncertainty on average within this information process. It need not reduce posterior variance at every realization. Nor are two auctions with different endogenous talk counts necessarily nested information sets. Their reported values and silence events can differ.

Prediction 7 (Belief convergence)**.** *When disclosure reduces the importance of retained common-value information in bidding, dispersion in observed bids is lower than dispersion in preceding talks, and the reduction is larger when more informative signals are disclosed.*

This is a prediction about bids motivated by common learning. At a fixed public history, talkers already share the same common-value posterior. Dispersion among their bids comes from idiosyncratic values under the maintained symmetric strategy. Silent dealers can have different posteriors because they retain different signals. Their bids combine that information with idiosyncratic values and shading.

Consequently, (A.170) does not prove an ordering of talk dispersion and bid dispersion. The two statistics can also use different dealer samples. A comparison based only on the highest or otherwise selected bids differs from one based on all bids. The empirical test should hold the security, units, timing, and observed sample as comparable as possible and distinguish variation in the common-value component from variation in idiosyncratic values.

A.4.4 Price crossing

Proposition 4 orders the conditional common value by winner identity. It does not give an unrestricted ordering of transaction prices. For a public history with both winner events of positive probability, define the price difference $C_P(\mathcal{I}) = \mathbb{E}[P \mid \mathcal{W}_S, \mathcal{I}] - \mathbb{E}[P \mid \mathcal{W}_T, \mathcal{I}]$. Write $S^{1:n_T} = \max_{j \in \mathcal{T}} m_j$ for the highest talk at that history.

Prediction 8 (Price crossing)**.** *In a parameter region where retained common-value information dominates the opposing effects of idiosyncratic values and shading, $C_P(\mathcal{I}) > 0$. At a fixed public history, the expected price premium and the expected price relative to the highest talk are then both larger for a silent-dealer win.*

A silent-dealer win selects favorable information absent from the disclosed signals. If that information contributes sufficiently strongly to bidding, the price rises further above consensus. A talker win instead selects on idiosyncratic value and on a relatively weak silent side. A sufficiently large idiosyncratic value can nevertheless generate a high talker price, so winner identity alone does not determine the price ordering.

At a fixed public history, consensus and the highest talk are constants. For either benchmark $c(\mathcal{I}) \in \{\bar{S}, S^{1:n_T}\}$,

$$\mathbb{E}[P - c(\mathcal{I}) \mid \mathcal{W}_S, \mathcal{I}] - \mathbb{E}[P - c(\mathcal{I}) \mid \mathcal{W}_T, \mathcal{I}] = C_P(\mathcal{I}). \tag{A.172}$$

The two conditional price-gap comparisons therefore have the same sign. In the data, the gap from the highest talk also differs from the gap from consensus by the talk spread $S_i^{1:N_i} - \bar{S}_i$. That spread can differ across winner groups. Unconditional comparisons of the two price gaps are consequently distinct empirical statistics. A positive difference between groups also does not require either group's average gap to have a particular sign relative to zero.

More disclosure can weaken the private-information channel by making more information public and leaving fewer silent dealers at a fixed total dealer count. Larger idiosyncratic-value dispersion can strengthen the ability of a talker to generate a high price. These forces motivate a weaker crossing when there are many talkers or large idiosyncratic values.

A.4.5 Winner identity

Under the monotonicity, overlap, and positive-probability conditions of Proposition 4, the exact common-value comparison is

$$\mathbb{E}[v \mid \mathcal{W}_S, \mathcal{I}] > \mathbb{E}[v \mid \mathcal{I}] > \mathbb{E}[v \mid \mathcal{W}_T, \mathcal{I}]. \tag{A.173}$$

Without informative overlap the corresponding inequalities can be weak. The result conditions on public information before bidding and compares the latent common value, irrespective of whether the price crossing holds.

To map this ordering into a future outcome, let Y be a value proxy measured so that a higher value indicates a better common-value outcome. Let $a_Y(\mathcal{I})$ be a public-history-dependent intercept and $b_Y(\mathcal{I}) > 0$ a slope. A sufficient measurement condition is

$$\mathbb{E}[Y \mid v, \mathcal{I}, \mathcal{W}_a] = a_Y(\mathcal{I}) + b_Y(\mathcal{I})v, \qquad a \in \{T, S\}, \tag{A.174}$$

where the intercept and slope are the same for both winner groups. This condition excludes a separate winner-group component in the proxy after conditioning on value and public information.

Prediction 9 (Winner identity)**.** *For future outcomes satisfying* (A.174)*, and at public histories with the strict common-value ordering in* (A.173)*, expected outcomes are higher following a silent-dealer win than following a talker win.*

Iterated expectations give the prediction directly. The difference in expected proxy outcomes equals $b_Y(\mathcal{I})$ times the difference in expected common values. Later benchmark-adjusted transaction prices, evaluated prices, or collateral outcomes can be informative proxies when their measurement

and subsequent shocks satisfy this comparison. A later price can also contain persistent dealer demand or liquidity effects, so its connection to the auction's common value requires empirical justification.

Controls should approximate the public information available before bidding and account for subsequent common shocks. Adding the auction price as a control asks whether winner identity contains information beyond what the price already reveals. Proposition 4 does not imply that stronger conditional statement. Likewise, replacing the full public history by a coarser set of controls need not preserve the theoretical ordering if omitted public information differs across winner groups.

A.4.6 Symmetry within the talker side

Fix a public history with identical strictly increasing talker strategies, independent identically distributed continuous idiosyncratic values, and a positive probability of a talker win. These are the conditions underlying Proposition A20. With distinct disclosed values, let R_i denote the ascending rank of the winning talker's signal among the N_i talks and let $U_i = (R_i - 1/2)/N_i$.

Prediction 10 (Symmetry among talkers)**.** *Conditional on the public history and a talker win, every talker wins with probability $1/n_T$, and the winning talker's expected signal equals consensus. In empirical notation,*

$$\Pr\left(U_i = \frac{r-1/2}{N_i} \,\middle|\, N_i, \mathcal{W}_{T,i}\right) = \frac{1}{N_i} \quad (r = 1, \ldots, N_i), \qquad \mathbb{E}[U_i \mid N_i, \mathcal{W}_{T,i}] = \frac{1}{2}. \tag{A.175}$$

Conditional on public information, each talker has the same common-value assessment and faces the same distribution of rival bids. Their probability of winning is determined by their independent idiosyncratic value, rather than by which disclosed signal is theirs. Proposition A20 first gives uniform probabilities over identities at a fixed history. Averaging over histories with the same talker count preserves the stated discrete rank distribution.

The null is discrete uniform on the rank midpoints for each observed N_i. Pooling different counts does not produce an exactly continuous uniform variable. If reported talks are rounded and tied, independent random ordering within each tied group preserves the rank test. Midranks retain the mean-one-half restriction under symmetry but generally have a different discrete distribution.

A complementary within-auction test uses the identities of all talkers, conditional on a talker win. The model implies that $S_{ij} - \bar{S}_i$ does not predict which talker wins after conditioning on the public history. Auction fixed effects absorb information shared by these talkers. For example, the conditional expectation of the sum of $(S_{ij} - \bar{S}_i)[\mathbb{1}\{j \text{ wins}\} - 1/N_i]$ over talkers is zero. This moment restriction tests one implication of uniform selection and need not detect every departure from it. It concerns selection among talkers, rather than the probability that the talker side defeats the silent side.

In the data, the null is rejected. Among 16,249 traded talker-won auctions with $N_i \geq 2$ and a winning talker identified among the talkers, ranks are assigned with independent random tie-breaking. The mean of U_i is 0.594 (security-clustered s.e. 0.002). A chi-square test of the discrete uniform law rejects within every count cell $N_i = 2, \ldots, 12$ ($p < 0.001$ in each). The mean of U_i is between 0.58 and 0.60 in every talk-count bin and between 0.59 and 0.61 in every seniority group. In the within-auction test, the coefficient of $\mathbb{1}\{j \text{ wins}\}$ on $S_{ij} - \bar{S}_i$ with auction fixed effects is 0.017 (s.e. 0.001), or 0.016 (0.002) when each auction receives equal weight. A systematic tendency for optimistic talkers to win of this kind rejects the joint timing, independence, and symmetry restrictions. Possible explanations include idiosyncratic values known before disclosure, correlation

between idiosyncratic values and signals, persistent differences across dealers, or private common-value information remaining after a talk. The rank test alone cannot distinguish these explanations.

Corollary A7 gives the complementary selection result. Conditional on a talker win, the winning idiosyncratic value is tilted upward relative to its unconditional distribution. Observable demand or inventory proxies can test that implication when they measure the relevant idiosyncratic value and their measurement errors are not selected by winning. A high silent bid is less informative about the decomposition of value. As Lemma A3 shows, different combinations of retained information and idiosyncratic value can generate the same silent bid, while a disclosed signal and a talker's bid separate those components under the known strategy. This distinction does not require a mass point in silent bids and does not make latent types directly observable in the data.

A.4.7 Comparative-static predictions

The fixed-point comparison requires an order on exact disclosure incentives. It does not assign that order to a primitive parameter without a continuation argument. Let θ and θ' denote two parameterizations with the same cost cdf. Let $\mathcal{P}$ be a common complete lattice of admissible policies, with the pointwise order, on which both selected continuation systems are defined. Write $\mathcal{B}_\theta(p)(s) = G_\kappa(\Delta(s; p, \theta))$ and let p^*_θ be its greatest fixed point. The canonical payoff cutoff is $\bar{\kappa}_\theta(s) = \Delta(s; p^*_\theta, \theta)$. Define $\bar{s}^*(\theta) = \sup\{s : p^*_\theta(s) > 0\}$, with $\sup \varnothing = -\infty$.

Corollary A8 (Monotone fixed-point comparison)**.** *Suppose $\mathcal{B}_\theta$ and $\mathcal{B}_{\theta'}$ are order-preserving self-maps of the common complete lattice $\mathcal{P}$. If $\mathcal{B}_{\theta'}(p) \le \mathcal{B}_\theta(p)$ for every $p \in \mathcal{P}$, then $p^*_{\theta'} \le p^*_\theta$ and $\bar{s}^*(\theta') \le \bar{s}^*(\theta)$. A sufficient condition for the map order is $\Delta(s; p, \theta') \le \Delta(s; p, \theta)$ for every s and $p \in \mathcal{P}$. If this sufficient condition holds and $\Delta(s; p, \theta)$ is additionally nondecreasing in p, then the canonical payoff cutoffs also satisfy $\bar{\kappa}_{\theta'}(s) \le \bar{\kappa}_\theta(s)$ for every signal.*

Proof. A policy satisfying $p \le \mathcal{B}_{\theta'}(p)$ also satisfies $p \le \mathcal{B}_\theta(p)$. The greatest fixed point of an order-preserving self-map of a complete lattice is the supremum of the policies satisfying this inequality. Inclusion of the two sets therefore gives $p^*_{\theta'} \le p^*_\theta$. Their positive sets are correspondingly nested, which orders the upper boundaries. With a common nondecreasing cost cdf, the stated pointwise gain order implies the map order.

For the last claim, the gain order at $p^*_{\theta'}$ gives $\Delta(s; p^*_{\theta'}, \theta') \le \Delta(s; p^*_{\theta'}, \theta)$. The additional monotonicity in the policy then gives $\Delta(s; p^*_{\theta'}, \theta) \le \Delta(s; p^*_\theta, \theta)$. These are the canonical payoff cutoffs in the two equilibria. □

The common domain is necessary for this comparison. The lattice of bounded monotone policies used in Appendix A.3.4 can serve this role if both maps are defined and invariant on it. A class containing only continuous monotone policies need not be a complete lattice. Together with the order between the maps, order preservation of each best response is enough to compare disclosure propensities and their positive regions. The additional gain monotonicity is used to order the raw payoff cutoffs. The best-response order alone does not rank negative gains because the cost cdf is flat there.

Prediction 11 (Comparative statics)**.** *Consider parameter regions in which both selected equilibrium systems meet the common-domain conditions above. Suppose adding a participating competitor lowers the gross gain from talking at every signal and rival policy. Suppose increasing σ_v^2 has the same effect, while increasing σ_ε^2 raises that gain. The greatest-equilibrium talk propensity, and its upper disclosure boundary then weakly decrease with competition and common-value uncertainty and weakly increase with signal noise.*

These statements concern finite parameter comparisons. A derivative of the disclosure boundary additionally requires its existence and differentiability. For competition, N is the total participating set in the model. An increase in the observed number of talks can instead reflect a change in disclosure propensities at a fixed N.

More competitors can lower talker rents by increasing the chance of facing informed silent rivals. They can also change the value of remaining silent and all selected bidding strategies. The hard-truncation rent comparison in Proposition A11 isolates one channel under its own assumptions. It does not establish the required uniform shift of the exact weighted-model gain when N changes.

A noisier private signal can reduce the value of retaining information and make disclosure less costly in informational terms. Greater prior uncertainty can increase the importance of private information and strengthen the incentive to retain a favorable signal. These forces motivate the stated directions. Their net effects also depend on disclosure selection, public learning, and equilibrium bidding, so the uniform gain orders remain substantive restrictions.

At fixed total participation and fixed dealer signals, the uncertainty comparison has a direct count implication. Independence of talking costs gives $\mathbb{E}_\theta[n_S \mid s_1, \dots, s_N] = \sum_{j=1}^{N}[1 - p_\theta^*(s_j)]$. A pointwise fall in p_θ^* therefore raises the expected number of silent dealers conditional on these signals.

An unconditional comparison must also allow the signal distribution to change. If f_θ is the marginal signal density, then

$$\mathbb{E}_\theta[n_S] = N\left[1 - \int_{\mathbb{R}} p_\theta^*(s) f_\theta(s)\, ds\right]. \tag{A.176}$$

For two parameterizations with the same N, the difference decomposes as

$$\mathbb{E}_{\theta'}[n_S] - \mathbb{E}_\theta[n_S] = N\int (p_\theta^* - p_{\theta'}^*) f_\theta\, ds + N\int p_{\theta'}^*(f_\theta - f_{\theta'})\, ds. \tag{A.177}$$

The first term is nonnegative under the pointwise policy order. The second term, which changes the signal distribution, need not have a general sign.

A sufficient condition for the unconditional ordering is a fall in disclosure at every fixed signal quantile. For Gaussian signals, write $s_\theta(z) = \mu_0 + \sqrt{\sigma_v^2 + \sigma_\varepsilon^2}\, z$ under parameterization θ, with z standard normal. If $p_{\theta'}^*(s_{\theta'}(z)) \le p_\theta^*(s_\theta(z))$ for almost every z, integrating under this common standard-normal law makes the unconditional talk probability weakly fall. This is an additional sufficient condition, not an equivalent restatement of the policy order at fixed raw signals.

Junior tranches, auctions with few talks, and periods of elevated aggregate uncertainty offer settings in which to examine these mechanisms. Tranche seniority and VIX are proxies rather than direct changes in a single primitive. They can also covary with signal precision, idiosyncratic values, liquidity, participation, and talking costs. The hypotheses of fewer talks and a larger silence premium therefore retain the gain-order and price-crossing qualifications. With total participation unobserved, fewer talks alone do not establish that more dealers are silent.

A.4.8 Reserve calibration

Let $\mathcal{F} \subseteq \mathcal{F}'$ be nested seller information sets, and suppose that v is square integrable. Write $\widehat{v} = \mathbb{E}[v \mid \mathcal{F}]$ and $\widehat{v}' = \mathbb{E}[v \mid \mathcal{F}']$. The seller's full-information reservation value is $v + \omega$.

Lemma A43 (Reserve calibration)**.** *Among all square-integrable $\mathcal{F}$-measurable reserves, $\widehat{v} + \omega$*

uniquely minimizes expected squared deviation from $v+\omega$, up to almost-sure equality. Moreover,

$$\mathbb{E}[(v-\widehat{v})^2]-\mathbb{E}[(v-\widehat{v}')^2]=\mathbb{E}[(\widehat{v}'-\widehat{v})^2]. \tag{A.178}$$

Proof. Let r be any square-integrable $\mathcal{F}$-measurable reserve. Conditional expectation gives

$$\mathbb{E}[(r-v-\omega)^2 \mid \mathcal{F}]=(r-\widehat{v}-\omega)^2+\mathrm{Var}(v \mid \mathcal{F}). \tag{A.179}$$

The conditional variance does not depend on r. The unique minimizer is therefore $r=\widehat{v}+\omega$ almost surely.

For the information comparison, write $v-\widehat{v}=(v-\widehat{v}')+(\widehat{v}'-\widehat{v})$. The second term is $\mathcal{F}'$-measurable, while the first has conditional mean zero given $\mathcal{F}'$. Their expected cross product is therefore zero. Squaring the identity and taking expectations proves (A.178). The reduction in mean squared error is strict exactly when $\widehat{v}'\neq\widehat{v}$ with positive probability. □

Equation (A.178) proves (31). The comparison applies to nested information sets in a common experiment. Auctions selected into different realized talk-count cells need not satisfy this nesting property.

A.4.9 Proof of Proposition 5

Proof. Let $v\sim\mathcal{N}(\mu_0,\sigma_v^2)$ and $s_i=v+\varepsilon_i$, where the errors are independent $\mathcal{N}(0,\sigma_\varepsilon^2)$ variables and are independent of v. Define $\tau_v=\sigma_v^{-2}$ and $\tau_\varepsilon=\sigma_\varepsilon^{-2}$. After observing n signals, the seller's posterior mean is

$$\mu_n=\frac{\tau_v\mu_0+\tau_\varepsilon\sum_{i=1}^n s_i}{\tau_v+n\tau_\varepsilon}, \tag{A.180}$$

and their posterior variance is $V_n=(\tau_v+n\tau_\varepsilon)^{-1}$. The estimation error $e_n=\mu_n-v$ consequently has distribution $\mathcal{N}(0,V_n)$.

By assumption, X is independent of v and of the signals used to construct μ_n. It is therefore independent of e_n. Failure occurs when $v+X<\mu_n+\omega$, or equivalently when $X-e_n<\omega$. The difference $X-e_n$ is Gaussian with mean m_X and variance $\sigma_X^2+V_n$. It follows that $q_n=\Phi((\omega-m_X)/\sqrt{\sigma_X^2+V_n})$, which proves (32).

To determine the sign, define $A=\omega-m_X$ and $q(V)=\Phi(A/\sqrt{\sigma_X^2+V})$. For $V>0$,

$$q'(V)=-\frac{A}{2(\sigma_X^2+V)^{3/2}}\phi\left(\frac{A}{\sqrt{\sigma_X^2+V}}\right), \tag{A.181}$$

where ϕ is the standard normal density. Since $V_{n+1}<V_n$, equation (A.181) gives the three cases in the proposition.

It remains to prove the diminishing effect when $\delta=m_X-\omega>0$. Extend the posterior variance to real $t\geq 0$ by setting $V(t)=(\tau_v+t\tau_\varepsilon)^{-1}$. Then $V'(t)=-\tau_\varepsilon V(t)^2$. Define

$$-\frac{dq(V(t))}{dt}=h(V(t)),\qquad h(V)=\frac{\delta\tau_\varepsilon V^2}{2(\sigma_X^2+V)^{3/2}}\phi\left(\frac{\delta}{\sqrt{\sigma_X^2+V}}\right). \tag{A.182}$$

The logarithmic derivative of h is

$$\frac{h'(V)}{h(V)}=\frac{2}{V}-\frac{3}{2(\sigma_X^2+V)}+\frac{\delta^2}{2(\sigma_X^2+V)^2}>0. \tag{A.183}$$

The sum of the first two terms is $(4\sigma_X^2+V)/(2V(\sigma_X^2+V)) > 0$, and the final term is positive. Thus h is strictly increasing in V. Since $V(t)$ is strictly decreasing, $h(V(t))$ is strictly decreasing in t. Finally,

$$q_n - q_{n+1} = \int_n^{n+1} h(V(t))\,dt. \tag{A.184}$$

Integrating a strictly decreasing function over successive intervals of equal length proves that $q_n - q_{n+1}$ is strictly decreasing in n. □

A.4.10 The seller's acceptance threshold

The calibrated reserve need not maximize the seller's expected payoff. To make this distinction explicit, hold the bid law fixed and allow the seller to choose an acceptance threshold. Conditional on a public history, suppose that $v \sim \mathcal{N}(\mu, V)$ and $B = v+X$, where $X \sim \mathcal{N}(m_X, \sigma_X^2)$ is independent of v and $\sigma_X^2 > 0$. The seller receives B upon sale and $v+\omega$ upon retention. Their expected payoff from a threshold r, chosen before observing B, is

$$W(r) = \mathbb{E}[B\mathbb{1}\{B \geq r\} + (v+\omega)\mathbb{1}\{B < r\}]. \tag{A.185}$$

Lemma A44 (Optimal acceptance threshold)**.** *Under the Gaussian bid specification and the seller objective in* (A.185)*, W has the unique maximizer*

$$r^{\text{opt}} = \mu + \omega + \frac{V}{\sigma_X^2}(\omega - m_X). \tag{A.186}$$

The corresponding failure probability is

$$q^{\text{opt}}(V) = \Phi\left(\frac{(\omega - m_X)\sqrt{V+\sigma_X^2}}{\sigma_X^2}\right). \tag{A.187}$$

If $m_X > \omega$, a reduction in V strictly increases this failure probability.

Proof. Let f_B be the density of B and define $\lambda = V/(V+\sigma_X^2)$. Gaussian regression gives $\mathbb{E}[v \mid B = b] = \mu + \lambda(b - \mu - m_X)$. Writing (A.185) as two integrals over b and differentiating at their common endpoint gives

$$W'(r) = f_B(r)\big(\mu + \omega + \lambda(r - \mu - m_X) - r\big). \tag{A.188}$$

The term in parentheses is affine in r with slope $-(1-\lambda) < 0$. Its unique zero is (A.186). Since $f_B(r) > 0$ everywhere, W increases before that zero and decreases after it. This proves global optimality and uniqueness.

The bid has mean $\mu + m_X$ and variance $V + \sigma_X^2$. Subtracting its mean from (A.186) gives $(\omega - m_X)(V+\sigma_X^2)/\sigma_X^2$. Standardizing proves (A.187). If $m_X > \omega$, the argument of Φ is negative and becomes less negative when V falls. The failure probability therefore increases. □

The reversal relative to Proposition 5 arises because the seller changes their threshold to maximize their own payoff. The calculation holds the bid distribution fixed and does not incorporate a bidding response to the new threshold.

A.4.11 Disclosure with one dealer

We first solve the one-dealer model. This benchmark isolates the private value of changing the seller's reserve from the effect of disclosure on competing bids.

Assumption A8 (One dealer)**.** *There is one potential dealer. The common value is $v \sim \mathcal{N}(\mu_0, \sigma_v^2)$, and the dealer observes $s = v + \varepsilon$, where $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ and both variances are positive. The disclosure cost κ is uniform on $[0, K]$. The idiosyncratic value has support $[0, \bar{u}]$, where $\bar{u} < \infty$. The variables v, ε, κ, u, and ξ_s are mutually independent. The dealer observes (s, κ) before choosing between exact disclosure of s and silence, and they observe u before bidding. Bids can take any real value, and nonparticipation pays zero. The seller observes ξ_s, which the dealer does not observe before bidding. Its distribution function F_ξ is continuous with full support on $\mathbb{R}$ and satisfies $\mathbb{E}[\xi_s^-] < \infty$. The seller uses $R = \widehat{v} + \omega + \xi_s$, and the reserve is not announced before bidding. Following an exact disclosure, the seller uses the Gaussian posterior conditional on the disclosed signal, including after an off-path disclosure. Following silence, their posterior satisfies Bayes' rule.*

A disclosure equilibrium consists of a disclosure rule, measurable optimal bids following disclosure and silence, and a Bayes-consistent seller posterior after silence.

Let $\varrho = \sigma_v^2/(\sigma_v^2 + \sigma_\varepsilon^2)$ and $\mu_s = \mu_0 + \varrho(s - \mu_0)$. If w is the dealer's valuation relative to the common-value estimate entering the reserve and z is their bid relative to that estimate, their optimized payoff is

$$\phi_\xi(w) = \sup_{z \in \mathbb{R}} F_\xi(z - \omega)(w - z)^+. \tag{A.189}$$

Including the positive part is equivalent to allowing the dealer to choose nonparticipation. Define $L_\xi(w) = \mathbb{E}[(w - \omega - \xi_s)^+]$.

Lemma A45 (Trading value and optimal bids)**.** *Under Assumption A8, ϕ_ξ is finite, strictly positive, strictly increasing, continuous, and one-Lipschitz. It satisfies*

$$0 < \phi_\xi(w) \le L_\xi(w) \le (w - \omega)^+ + \mathbb{E}[\xi_s^-]. \tag{A.190}$$

Moreover, $\phi_\xi(w) \to 0$ as $w \to -\infty$ and $\phi_\xi(w) \to \infty$ as $w \to \infty$. At every finite w, the supremum is attained at a finite $z < w$. The smallest maximizing bid defines a nondecreasing Borel function $z^*(w)$.

Proof. On the event $\{\xi_s \le z - \omega\}$, $(w - z)^+ \le (w - \omega - \xi_s)^+$. Multiplying by the event indicator and taking expectations gives $F_\xi(z - \omega)(w - z)^+ \le L_\xi(w)$. Taking the supremum over z proves the first upper bound in (A.190). The second follows from $(a + b)^+ \le a^+ + b^+$.

For each z, the objective in (A.189) is nondecreasing and has Lipschitz constant at most one as a function of w. Its finite supremum has the same properties and is continuous. Full support implies that every finite bid has a strictly positive acceptance probability. Choosing any $z < w$ therefore proves $\phi_\xi(w) > 0$.

For $w \le \omega$, the variable $(w - \omega - \xi_s)^+$ is bounded by ξ_s^- and converges pointwise to zero as $w \to -\infty$. Dominated convergence gives $L_\xi(w) \to 0$ and hence $\phi_\xi(w) \to 0$. Holding any finite z fixed shows that $\phi_\xi(w) \to \infty$ as $w \to \infty$.

We next prove attainment. The objective converges to zero as $z \uparrow w$. Integrability of ξ_s^- implies $|x| F_\xi(x) \to 0$ as $x \to -\infty$. Indeed, for $y > 0$, $y \Pr(\xi_s^- \ge y) \le \mathbb{E}[\xi_s^- \mathbb{1}\{\xi_s^- \ge y\}] \to 0$. Setting $x = z - \omega$ shows that $F_\xi(z - \omega)(w - z) \to 0$ as $z \to -\infty$. Since a strictly positive payoff is feasible, every maximizer lies in a compact subset of $(-\infty, w)$. Continuity gives attainment.

Let $w_2 > w_1$, and let z_1 maximize the objective at w_1. Then

$$\phi_\xi(w_2) \geq \phi_\xi(w_1) + F_\xi(z_1 - \omega)(w_2 - w_1) > \phi_\xi(w_1), \tag{A.191}$$

which proves strict monotonicity.

Finally, let z_i maximize the objective at w_i, where $w_2 > w_1$. Adding the two optimality inequalities gives $(w_2 - w_1)(F_\xi(z_2 - \omega) - F_\xi(z_1 - \omega)) \geq 0$. If $z_1 > z_2$, monotonicity of F_ξ forces equality of the two positive acceptance probabilities. The lower bid z_2 would then yield a strictly larger margin at w_1, contradicting optimality of z_1. Hence every optimizer at w_1 is weakly below every optimizer at w_2. Each maximizing set is nonempty and compact, so its smallest element exists. This smallest-element selection is nondecreasing and therefore Borel measurable. □

Following disclosure, the dealer and seller share the posterior mean μ_s. A bid $b = \mu_s + z$ gives expected profit $F_\xi(z - \omega)(u - z)$. The dealer's expected optimized talking payoff before observing u is therefore

$$\pi_T = \mathbb{E}_u[\phi_\xi(u)]. \tag{A.192}$$

Lemma A45 gives $0 < \phi_\xi(0) \leq \pi_T \leq \phi_\xi(\bar{u}) < \infty$.

Following silence, suppose that the seller's posterior mean is m. Conditional on the dealer's signal, their disclosure decision carries no information about v because the disclosure cost is independent of v given s. A bid $b = m + z$ gives expected profit $F_\xi(z - \omega)(\mu_s - m + u - z)$. Define

$$\pi_S(s, m) = \mathbb{E}_u[\phi_\xi(\mu_s - m + u)], \qquad \Delta(s, m) = \pi_T - \pi_S(s, m). \tag{A.193}$$

Suppose $K > \pi_T$ and write $b_0 = \pi_T / K < 1$. The disclosure best response to a proposed silent posterior mean m is $p_m(s) = \Delta(s, m)^+ / K$. The posterior mean generated by this policy is

$$\mathcal{H}(m) = \frac{\mathbb{E}[v(1 - p_m(s))]}{\mathbb{E}[1 - p_m(s)]}. \tag{A.194}$$

Set $C_1 = b_0 \sigma_v \sqrt{2/\pi} / (1 - b_0)$.

Proposition A22 (One-dealer disclosure equilibrium)**.** *Under Assumption A8, suppose $K > \pi_T$. The map $\mathcal{H}$ has a fixed point in $[\mu_0, \mu_0 + C_1]$, and every fixed point supports a disclosure equilibrium. Conditional on a fixed point m, the gross disclosure gain is continuous and strictly decreasing in s, with limits π_T as $s \to -\infty$ and minus infinity as $s \to \infty$. Its unique zero is*

$$s^*(m) = \mu_0 + \frac{m - \mu_0}{\varrho}. \tag{A.195}$$

Up to cost ties, disclosure occurs if and only if $\kappa < \Delta(s, m)$. Consequently, $p_m(s) > 0$ exactly when $s < s^*(m)$ and $p_m(s) \to \pi_T / K$ as $s \to -\infty$. Every equilibrium under the stated rules has this form, up to almost-sure equality and the selection of optimal bids.

Proof. Fix a finite m. Since μ_s is strictly increasing in s, Lemma A45 implies that $\pi_S(s, m)$ is continuous and strictly increasing in s. Boundedness of the idiosyncratic value gives $\phi_\xi(\mu_s - m) \leq \pi_S(s, m) \leq \phi_\xi(\mu_s - m + \bar{u})$. The limits of ϕ_ξ therefore imply that $\pi_S(s, m)$ converges to zero in the lower signal tail and to infinity in the upper signal tail. When $\mu_s = m$, its value is exactly π_T. These observations prove the stated properties of Δ and yield (A.195).

The dealer talks when the gross gain exceeds their disclosure cost. Since $\Delta(s, m) \leq \pi_T < K$, the upper endpoint of the cost distribution never binds. Thus the disclosure probability is exactly

$p_m(s) = \Delta(s,m)^+/K$. This function is continuous and nonincreasing and satisfies $0 \le p_m(s) \le b_0 < 1$. The denominator in (A.194) is therefore at least $1 - b_0$.

We next establish consistency of the seller's posterior mean. Conditional expectation gives $\mathrm{Cov}(v, p_m(s)) = \mathrm{Cov}(\mu_s, p_m(s)) \le 0$. To verify the inequality, take independent copies s and s' of the signal. Since μ_s is increasing and $p_m(s)$ is nonincreasing, $(\mu_s - \mu_{s'})(p_m(s) - p_m(s')) \le 0$. Its expectation is twice the covariance. It follows that

$$0 \le \mathcal{H}(m) - \mu_0 = \frac{-\mathbb{E}[(v-\mu_0)p_m(s)]}{\mathbb{E}[1-p_m(s)]} \le \frac{b_0 \mathbb{E}|v-\mu_0|}{1-b_0} = C_1. \tag{A.196}$$

Hence $\mathcal{H}$ maps the real line into $[\mu_0, \mu_0 + C_1]$.

For every s, continuity of ϕ_ξ makes $p_m(s)$ continuous in m. Dominated convergence applies to the numerator and denominator of (A.194). The denominator remains bounded away from zero, so $\mathcal{H}$ is continuous. The function $\mathcal{H}(m) - m$ is nonnegative at $m = \mu_0$ and nonpositive at $m = \mu_0 + C_1$. The intermediate value theorem supplies a fixed point.

Let m be a fixed point and define $q_m(v) = \mathbb{E}_\varepsilon[1 - p_m(v+\varepsilon)]$. The seller's posterior density following silence is

$$f_m^S(v) = \frac{f_v(v) q_m(v)}{\mathbb{E}[1-p_m(s)]}. \tag{A.197}$$

It is well defined because $q_m(v) \ge 1 - b_0$, and its mean is $\mathcal{H}(m) = m$. Following an on-path disclosure, the factor $p_m(s)$ is constant as a function of v and cancels from Bayes' rule. The posterior is therefore the Gaussian posterior conditional on s. Assumption A8 specifies the same posterior following an off-path exact disclosure.

Use the measurable optimizer from Lemma A45. Following disclosure, set $b_T(s,u) = \mu_s + z^*(u)$. Following silence, set $b_S(s,u) = m + z^*(\mu_s - m + u)$. These bids attain the corresponding continuation values. The disclosure rule compares those values net of the realized cost. Together with (A.197), they form an equilibrium.

Conversely, consider any equilibrium under the stated rules. Its talking payoff is π_T, while nonparticipation makes its silent payoff nonnegative. Its disclosure probability is therefore at most b_0. Silence has positive probability and generates a finite posterior mean, say m. Optimal bidding produces the values in (A.193), optimal disclosure produces p_m, and Bayes' rule requires $m = \mathcal{H}(m)$. This proves the characterization. □

The cutoff is unique conditional on the equilibrium posterior mean m. The fixed point itself need not be unique. The relevant disclosure bound is π_T/K, rather than $\bar{u}/K$, because the seller's private shock can give a dealer positive trading value even when their idiosyncratic value is zero.

A.4.12 Disclosure with competing dealers

We now prove Proposition 6. Maintain the Gaussian primitives, the independent uniform disclosure costs, the independent idiosyncratic values with common distribution F_u on $[0, \bar{u}]$, and the continuation beliefs stated in Section 6. Fix a finite population $N \ge 2$ and a disclosure equilibrium satisfying Assumption 1, with $K > \pi_T$. Independent random priorities resolve ties, and submitted bids may have atoms. The argument is conditional on this equilibrium.

Write $p_i(s)$ for dealer i's disclosure probability. For a focal dealer j with signal s, let $\Pi_j^{T,R}(s)$ and $\Pi_j^{S,R}(s)$ be their optimized continuation payoffs from disclosure and silence before their idiosyncratic value and the rival disclosure history are realized. Define $\Delta_j^R(s) = \Pi_j^{T,R}(s) - \Pi_j^{S,R}(s)$. All signal-indexed statements are understood almost everywhere.

Randomized monotonicity. The main-text assumption orders conditional action distributions rather than requiring deterministic bids. The following representation connects the two formulations.

Lemma A46 (Quantile representation of monotone action kernels)**.** *Order nonparticipation below every real bid. A Borel action kernel is nondecreasing in the first-order stochastic order in each coordinate of the dealer's type if and only if it can be generated by an independent uniform variable ζ through a Borel rule $a(t,\zeta)$ that is nondecreasing in each type coordinate for every fixed ζ.*

Proof. Map the ordered action space into $[0,1)$ by sending nonparticipation to zero and real bids through a strictly increasing transformation into $(0,1)$. If $a(t,\zeta)$ is pathwise nondecreasing, the indicator of every upper interval is nondecreasing in type for every ζ. Integration over ζ gives first-order stochastic monotonicity.

Conversely, let $F(t,x)$ be the conditional distribution function of the transformed action. Kernel measurability and right continuity make $(t,x) \mapsto F(t,x)$ jointly Borel. First-order stochastic monotonicity makes $F(t,x)$ nonincreasing when either coordinate of t rises. Define the left quantile $Q(t,z) = \inf\{x : F(t,x) \geq z\}$. The map Q is jointly Borel. If ζ is uniform on $[0,1]$, the inverse-transform theorem gives $Q(t,\zeta)$ the required conditional law. The ordering of the distribution functions makes $Q(t,\zeta)$ nondecreasing in each coordinate of t for every ζ. Transforming back to the original action space proves the result. □

The talking-payoff bound. For a submitted bid b, let $W_j(b)$ be the event that it beats all rivals, with nonparticipation treated as an action below every bid. Let $Q_j(b) = \Pr(W_j(b) \mid \mathcal{G}_j)$. When $Q_j(b) > 0$, define $M_j(b) = \mathbb{E}[v \mid \mathcal{G}_j, W_j(b)]$.

Lemma A47 (Winning and the talking-payoff bound)**.** *Under Assumption 1, $M_j(b) \leq \mathbb{E}[v \mid \mathcal{G}_j]$ whenever $Q_j(b) > 0$. A talker's optimized payoff conditional on their public history and idiosyncratic value u is at most $\phi_\xi(u)$. Consequently,*

$$\Pi_j^{T,R}(s) \leq \pi_T, \qquad \Delta_j^R(s) \leq \pi_T, \qquad 0 \leq p_j(s) \leq \frac{\pi_T}{K} < 1. \tag{A.198}$$

Proof. Fix the public history and the focal dealer's bidding information. Conditional on v, a silent rival i has signal density proportional to $(1 - p_i(t)) f_\varepsilon(t - v)$. For $v_2 > v_1$, the likelihood ratio of the two conditional densities is a positive constant times $\exp((v_2 - v_1)t/\sigma_\varepsilon^2)$. The silence weight cancels from the ratio. The rival's signal is therefore increasing in v in the likelihood-ratio order and hence in the first-order stochastic order.

Use a common-quantile coupling for each rival signal and the representation in Lemma A46 for each randomized action. Holding idiosyncratic values, private randomization devices, and the priority ordering fixed, an increase in v weakly raises every silent rival's action. A talker's action depends only on public information, their idiosyncratic value, and independent randomization, so its conditional distribution does not increase with v through an undisclosed signal.

For a fixed priority ordering, the focal dealer wins against a rival either when the rival's action is strictly below b or, if the focal dealer has priority, when it is weakly below b. The focal winning indicator is therefore nonincreasing in every rival action. Under the coupling, it is nonincreasing in v. Integrating over rivals' signals, idiosyncratic values, devices, and priorities gives a nonincreasing function $h(v) = \Pr(W_j(b) \mid v, \mathcal{G}_j)$.

Let V and V' be independent draws from the focal dealer's posterior. Then $2\operatorname{Cov}(V, h(V)) = \mathbb{E}[(V - V')(h(V) - h(V'))] \leq 0$. The posterior has a finite first moment because its Gaussian

density is multiplied by bounded silence weights and normalized by a positive constant. Hence $\mathbb{E}[v\mathbb{1}\{W_j(b)\} \mid \mathcal{G}_j] \le \mathbb{E}[v \mid \mathcal{G}_j]Q_j(b)$. Dividing by $Q_j(b) > 0$ proves the winner's-curse inequality.

Following exact disclosure, the talker shares the seller's conditional common-value mean $\widehat{v}$. For a fixed bid, seller acceptance is independent of the common value and rival actions. If $Q_j(b) > 0$, the bid's expected payoff satisfies

$$F_\xi(b - \widehat{v} - \omega)Q_j(b)(M_j(b) + u - b) \le F_\xi(b - \widehat{v} - \omega)(\widehat{v} + u - b)^+. \tag{A.199}$$

If the margin on the left is negative, the inequality remains valid. If $Q_j(b) = 0$, the bid earns zero. Setting $z = b - \widehat{v}$ and optimizing over bids and nonparticipation gives the bound $\phi_\xi(u)$. Averaging over idiosyncratic values and rival histories yields $\Pi_j^{T,R}(s) \le \pi_T$.

Nonparticipation gives $\Pi_j^{S,R}(s) \ge 0$, so $\Delta_j^R(s) \le \pi_T$. Uniform disclosure costs imply $p_j(s) = \min\{1, (\Delta_j^R(s))^+/K\}$. Since $K > \pi_T$, this proves (A.198). The argument does not use a pre-existing lower bound on silence probabilities. At a signal with positive disclosure probability, a cost above π_T cannot prefer disclosure to the zero outside option. At a signal with zero disclosure probability, the bound is immediate. The prescribed continuation beliefs extend the same comparison to an off-path exact disclosure. □

Bid atoms do not change the result. The priority ordering determines whether a tie uses a strict or weak comparison, and both comparisons are decreasing in a rival's action. The one-dealer bound is sharp over the model family. Lemma A45 also gives $\phi_\xi(u) \le \phi_\xi(0) + u$, so $\pi_T \le \phi_\xi(0) + \mathbb{E}[u] \le \phi_\xi(0) + \bar{u}$.

Posterior bounds under the equilibrium policy. Set $\underline{q} = 1 - \pi_T/K > 0$ and define $q_i(v) = \mathbb{E}_\varepsilon[1 - p_i(v + \varepsilon)]$. Equation (A.198) implies $1 - p_i(s) \ge \underline{q}$ and $q_i(v) \ge \underline{q}$. For a public history with talker set T and silent set S, write $n_T = |T|$. Let

$$\mu_T = \frac{\tau_v\mu_0 + \tau_\varepsilon \sum_{i\in T} s_i}{\tau_v + n_T\tau_\varepsilon} \tag{A.200}$$

be the Gaussian posterior mean based only on disclosed signal values. Define

$$C_0 = (\underline{q}^{-N} - 1)\sigma_v\sqrt{2/\pi}, \qquad C_0'' = (\underline{q}^{-(N-1)} - 1)\sqrt{2V_1/\pi}. \tag{A.201}$$

Lemma A48 (Bounds on posterior means)**.** *The seller's posterior density is*

$$f(v \mid \mathcal{I}) = \frac{\phi_{V_{n_T}}(v - \mu_T)\prod_{i\in S} q_i(v)}{\int_{\mathbb{R}} \phi_{V_{n_T}}(x - \mu_T)\prod_{i\in S} q_i(x)\,dx}. \tag{A.202}$$

It satisfies $|\widehat{v}(\mathcal{I}) - \mu_T| \le C_0$. At the all-silent history, a dealer with signal s_i has a conditional common-value mean within C_0'' of μ_{s_i}. At a one-talker history where the talker discloses s, the seller's posterior mean is within C_0'' of μ_s.

Proof. Conditional on v, signals and disclosure decisions are independent across dealers. A silent dealer contributes the likelihood $q_i(v)$. A disclosed signal contributes its Gaussian likelihood. Conditional on the disclosed signal, its disclosure probability is constant as a function of v and cancels from Bayes' rule. The specified off-path beliefs use the same expression. This proves (A.202).

Suppose a reference density is multiplied by a weight taking values in $[a, 1]$. The normalized

density ratio r then lies in $[a, a^{-1}]$. If Z is Gaussian with mean μ and variance V,

$$|\mathbb{E}[(Z-\mu)r(Z)]| = |\mathbb{E}[(Z-\mu)(r(Z)-1)]| \le (a^{-1}-1)\sqrt{2V/\pi}. \tag{A.203}$$

Apply this inequality with $a = \underline{q}^{|S|}$ and $V = V_{n_T} \le \sigma_v^2$ to obtain the bound C_0.

At the all-silent history, conditioning additionally on dealer i's signal cancels their own silence factor. The remaining $N-1$ silence factors weight a Gaussian posterior with mean μ_{s_i} and variance V_1. At a one-talker history, the $N-1$ silent-dealer factors weight the Gaussian posterior conditional on the disclosed signal. Taking $a = \underline{q}^{N-1}$ proves both C_0'' bounds. □

Silent payoffs in the signal tails. Define $c = \varrho\tau_v/(\tau_v + (N-1)\tau_\varepsilon) > 0$. Let $Z_v \sim \mathcal{N}(0, V_1)$ and $Z_i \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ for $i = 1, \dots, N-1$, with all these variables independent. Set $W = |Z_v| + \sum_{i=1}^{N-1} |Z_i| + C_0$, independently of ξ_s, and define

$$\psi(s) = \mathbb{E}[(W + \bar{u} - \omega - \xi_s - c(\mu_0 - s))^+], \qquad s \le \mu_0. \tag{A.204}$$

Lemma A49 (Silent payoffs in the signal tails)**.** *For every dealer j, $0 \le \Pi_j^{S,R}(s) \le \psi(s)$ for $s \le \mu_0$, and $\psi(s) \to 0$ as $s \to -\infty$. There are constants $d_1 > 0$ and $d_2 < \infty$, independent of the selected equilibrium, such that*

$$\Pi_j^{S,R}(s) \ge d_1(s - \mu_0) - d_2, \qquad s \ge \mu_0. \tag{A.205}$$

Proof. Conditional on the focal signal s, write $v = \mu_s + \eta_v$. The residual η_v is $\mathcal{N}(0, V_1)$ and is independent of s. Rival signal errors remain independent $\mathcal{N}(0, \sigma_\varepsilon^2)$ variables.

Let J be the realized set of rival talkers when the focal dealer remains silent, and write $n = |J| \le N-1$. Define $a_n = \tau_v/(\tau_v + n\tau_\varepsilon)$ and $d_n = \tau_\varepsilon/(\tau_v + n\tau_\varepsilon)$. Substitution of $s_i = v + \varepsilon_i$ into (A.200) gives

$$\mu_T - v = a_n\varrho(\mu_0 - s) - a_n\eta_v + d_n \sum_{i \in J} \varepsilon_i. \tag{A.206}$$

For $s \le \mu_0$, the deterministic term is at least $c(\mu_0 - s)$. Since $a_n \le 1$ and $d_n \le 1$, Lemma A48 yields

$$\widehat{v} - v \ge c(\mu_0 - s) - |\eta_v| - \sum_{i \ne j} |\varepsilon_i| - C_0. \tag{A.207}$$

This is a pathwise bound and does not require the set J to be independent of the signal errors.

Whenever a bid is accepted, $b \ge \widehat{v} + \omega + \xi_s$. The dealer's realized profit is therefore at most $(v + \bar{u} - \widehat{v} - \omega - \xi_s)^+$. Equations (A.207) and (A.204) bound the payoff from every silent strategy by $\psi(s)$. Nonparticipation gives the lower bound zero. The positive part defining $\psi(s)$ converges pointwise to zero as $s \to -\infty$ and is bounded by the integrable variable $W + \bar{u} + |\omega| + \xi_s^-$. Dominated convergence proves the lower-tail limit.

For the upper tail, fix $a > 0$ and let $\mathcal{I}_0$ be the all-silent public history. Consider the event A_s on which all rivals remain silent, $|\eta_v| \le a$, and $|\varepsilon_i| \le a$ for every rival. Conditional on any signal vector, each rival remains silent with probability at least $\underline{q}$. Hence

$$\Pr(A_s \mid s) \ge \underline{q}^{N-1} \Pr(|Z_v| \le a) \prod_{i=1}^{N-1} \Pr(|Z_i| \le a) \equiv \eta_a > 0. \tag{A.208}$$

On A_s, every rival signal is at most $\mu_s + 2a$. Lemma A48 and the maintained bid cap imply

that every rival bid is at most $\mu_0 + \varrho(\mu_s + 2a - \mu_0) + C_0'' + \bar{u}$. At $\mathcal{I}_0$, the focal dealer can submit

$$b_s = \mu_0 + \varrho(\mu_s + 2a - \mu_0) + C_0'' + \bar{u} + 1 \tag{A.209}$$

and decline to participate at other histories. This bid strictly clears every rival bid on A_s.

The seller's posterior mean at $\mathcal{I}_0$ is at most $\mu_0 + C_0$. The bid is therefore accepted with probability at least $F_\xi(b_s - \mu_0 - C_0 - \omega)$, which converges to one. On A_s, its surplus is at least $\mu_s - a - b_s$. Define $C_3 = (2\varrho + 1)a + C_0'' + \bar{u} + 1$. Then $\mu_s - a - b_s = (1 - \varrho)\varrho(s - \mu_0) - C_3$.

The possible loss outside the favorable event is at most $\ell(s) = \mathbb{E}[(b_s - \mu_s - Z_v)^+]$. Since $b_s - \mu_s \to -\infty$, Gaussian integrability gives $\ell(s) \to 0$. Global optimality therefore implies

$$\Pi_j^{S,R}(s) \geq \eta_a F_\xi(b_s - \mu_0 - C_0 - \omega)\big[(1 - \varrho)\varrho(s - \mu_0) - C_3\big] - \ell(s) \tag{A.210}$$

for all sufficiently large s. The right side has a strictly positive asymptotic linear slope. Choosing a smaller positive slope d_1 and increasing a finite intercept d_2 proves (A.205) for every $s \geq \mu_0$. □

A positive talking payoff with two dealers. Now set $N = 2$ and let dealer 1 be the focal talker. Write $\mathcal{I}_1^T(s)$ for the history at which they disclose s and dealer 2 remains silent. Define $G = s_2 - \mu_s$, $\sigma_r^2 = V_1 + \sigma_\varepsilon^2$, $\lambda_2 = \tau_\varepsilon/(\tau_v + 2\tau_\varepsilon)$, and $\sigma_c = \lambda_2\sigma_r = \sqrt{\lambda_2 V_1}$. Before conditioning on dealer 2's silence, $G \mid s \sim \mathcal{N}(0, \sigma_r^2)$. At $\mathcal{I}_1^T(s)$, its density is multiplied by $1 - p_2(\mu_s + G)$ and normalized. Its density ratio relative to the Gaussian reference law therefore lies in $[\underline{q}, \underline{q}^{-1}]$. The one-talker history itself has conditional probability at least $\underline{q}$.

Conditional on the rival's signal, their own silence factor cancels. Gaussian updating gives

$$v \mid \mathcal{I}_1^T(s), s_2, u_2 \sim \mathcal{N}(m_2, V_2), \qquad m_2 = \mu_s + \lambda_2 G. \tag{A.211}$$

The quantity m_2 is the rival's conditional common-value mean. It is also the focal dealer's conditional mean after additionally conditioning on the rival's signal. Let $e_s = \widehat{v}(\mathcal{I}_1^T(s)) - \mu_s$. Lemma A48 gives $|e_s| \leq C_0''$, where $C_0'' = (\underline{q}^{-1} - 1)\sqrt{2V_1/\pi}$.

Nondegeneracy of F_u permits values $0 \leq u_0 < u_* \leq \bar{u}$ such that $F_0 = F_u(u_0) > 0$ and $m_* = \Pr(u \geq u_*) > 0$. Fix such a pair and set $\delta = (u_* - u_0)/2 > 0$. For $x > 0$, write $\ln^+ x = \max\{0, \ln x\}$. Define

$$T_0 = C_0'' + \frac{\sigma_c^2}{\delta} \ln^+\left(\frac{2\sigma_c^2}{\underline{q}^2\delta^2 F_0}\right). \tag{A.212}$$

Under full lower support of the seller shock, define

$$\pi_0 = \frac{\underline{q}^2 m_* F_0 \delta^2}{2\sigma_c} \phi\left(\frac{T_0 + \bar{u} - \delta + C_0''}{\sigma_c}\right) F_\xi(-T_0 - \omega) > 0. \tag{A.213}$$

Lemma A50 (Uniform talking payoff with two dealers)**.** *Suppose $N = 2$ and F_u is nondegenerate. Under the maintained continuation conditions, $\Pi_j^{T,R}(s) \geq \pi_0$ for every dealer j and almost every signal s.*

Proof. Consider dealer 1 at $\mathcal{I}_1^T(s)$ with idiosyncratic value $u \geq u_*$. For $t \geq C_0''$, let their submit $b = \widehat{v} - t$. Its acceptance probability is $F_\xi(-t - \omega)$. Define $X = m_2 - \widehat{v} = \lambda_2 G - e_s$. Conditional on (s_2, u_2), the rival's action and private randomization are independent of the common value. Hence the expected surplus from the focal bid, conditional on winning, is $X + t + u$.

The rival's bid cap guarantees that the focal bid wins strictly whenever $X + u_2 < -t$. In

particular, define

$$E = \{u_2 \le u_0,\ \delta - t - u \le X < -t - u_0\}. \tag{A.214}$$

On E, the focal dealer wins and their conditional margin is at least δ. The interval for X has length $u - \delta - u_0 \ge \delta$. Counting this positive contribution and charging every negative-margin state as a loss gives

$$\mathbb{E}[(X + t + u)\mathbb{1}\{W_1(b)\}] \ge \delta \Pr(E) - \mathbb{E}[(X + t + u)^-]. \tag{A.215}$$

Under the Gaussian reference law, $X = \sigma_c Z - e_s$, where $Z \sim \mathcal{N}(0,1)$. Set $x_p = (t+u-\delta-e_s)/\sigma_c \ge 0$. The standardized interval in (A.214) lies to the left of zero. Its density is minimized at its left endpoint. The lower density-ratio bound, the interval length, and independence of u_2 imply $\delta \Pr(E) \ge \underline{q} F_0 \delta^2 \phi(x_p)/\sigma_c \equiv P$.

For $x \ge 0$, define $\Upsilon(x) = \mathbb{E}[(-Z - x)^+] = \phi(x) - x\Phi(-x) \le \phi(x)$. The upper density-ratio bound gives $\mathbb{E}[(X + t + u)^-] \le \underline{q}^{-1}\sigma_c \Upsilon(x_p + \delta/\sigma_c) \equiv L$. At the same posterior shift e_s,

$$\frac{L}{P} \le \frac{\sigma_c^2}{\underline{q}^2 F_0 \delta^2} \exp\left(-\frac{\delta x_p}{\sigma_c}\right). \tag{A.216}$$

To obtain the final inequality, use $\Upsilon(y) \le \phi(y)$ and the Gaussian density ratio, then discard the additional factor $\exp(-\delta^2/(2\sigma_c^2)) \le 1$.

Since $\sigma_c x_p \ge t + u_* - \delta - C_0'' \ge t - C_0''$, the choice $t = T_0$ makes the right side of (A.216) at most $1/2$. Equation (A.215) therefore gives a conditional payoff of at least $F_\xi(-T_0 - \omega)\underline{q} F_0 \delta^2 \phi(x_p)/(2\sigma_c)$.

Only now take the uniform bound over histories and focal idiosyncratic values. Since $u \le \bar{u}$ and $|e_s| \le C_0''$, $0 \le x_p \le (T_0 + \bar{u} - \delta + C_0'')/\sigma_c$. The standard normal density is decreasing on the positive half-line. Replacing $\phi(x_p)$ by the density in (A.213) gives a uniform conditional payoff bound. The one-talker history has probability at least $\underline{q}$, and the event $u \ge u_*$ has probability m_*. The focal dealer can make the deviation on this history and idiosyncratic-value event and choose nonparticipation elsewhere. Global optimality then yields (A.213). The same argument applies after reversing dealer identities. □

The positive and negative contributions in the proof are compared at the same posterior shift before uniform bounds are imposed. This order is essential. The strict cap-clearance event also makes the proof insensitive to bid atoms. The constant π_0 establishes a strict payoff sign but can be economically small.

Corollary A9 (A lower-endpoint atom)**.** *Replace full lower support of the seller shock by a finite essential infimum $\xi_{\min}$ with a positive atom at $\xi_{\min}$. Suppose $\mathbb{E}[\xi_s^-] < \infty$ and*

$$T_0 + \omega + \xi_{\min} \le 0. \tag{A.217}$$

Then $F_\xi(-T_0 - \omega) > 0$, the constant in (A.213) remains strictly positive, and the lower-tail conclusion of Proposition 6 continues to hold.

Proof. Condition (A.217) is equivalent to $-T_0 - \omega \ge \xi_{\min}$. If the inequality is strict, the definition of essential infimum gives $F_\xi(-T_0 - \omega) > 0$. If equality holds, the atom at $\xi_{\min}$ gives the same conclusion. The proof of Lemma A50 is otherwise unchanged. The posterior and upper-tail arguments require only an integrable negative part and convergence of the distribution function to one in its upper tail. □

Remark A6 (A nonempty endpoint branch)**.** *The endpoint condition is not vacuous. Take $N = 2$, $\mu_0 = 0$, $\tau_v = \tau_\varepsilon = 1$, $\bar{u} = 1$, and let F_u be uniform on $[0,1]$. Suppose $K \ge 100\pi_T$, so $\underline{q} \ge 0.99$.*

Choose $u_0 = 1/4$ and $u_* = 3/4$. Then $\delta = 1/4$, $F_0 = 1/4$, $V_1 = 1/2$, and $\sigma_c^2 = 1/6$. At $\underline{q} = 0.99$, $C_0'' < 0.00570$ and $T_0 < 2.0593$.

Set $\omega = 0$. Let $\xi_s = -3$ with probability $1/2$, and let $\xi_s = -3 + E$ with probability $1/2$, where E is an independent unit-rate exponential variable. The shock has essential infimum -3, a positive atom there, an integrable negative part, and a distribution function continuous on $(-3, \infty)$. Moreover, $T_0 + \omega + \xi_{\min} < -0.9407$ and $F_\xi(1) = 1 - \frac{1}{2}e^{-4} < 1$. Thus the lower-endpoint condition holds with slack. This example verifies the primitive support conditions. It does not establish existence of the all-silent continuation or of a complete disclosure equilibrium.

Proof of Proposition 6. Lemma A47 proves (34). Lemma A49 implies that the silent payoff eventually exceeds π_T in the upper signal tail. Hence $\Delta_j^R(s) < 0$ for all sufficiently large s. Optimal disclosure gives a finite s_+ such that $p_j(s) = 0$ for almost every $s \geq s_+$. Since the dealer population is finite, s_+ can be chosen uniformly across dealers.

Now suppose $N = 2$ and F_u is nondegenerate. Lemma A50 gives $\Pi_j^{T,R}(s) \geq \pi_0 > 0$. Since $\psi(s) \to 0$ as $s \to -\infty$, choose $s_- \leq \mu_0$ such that $s_- < s_+$ and $\psi(s_-) \leq \pi_0/2$. The function ψ is nondecreasing in s, so $\Pi_j^{S,R}(s) \leq \psi(s) \leq \pi_0/2$ for every $s \leq s_-$. Therefore $\Delta_j^R(s) \geq \pi_0/2$. Since $\Delta_j^R(s) \leq \pi_T < K$, the uniform-cost best response gives $p_j(s) = (\Delta_j^R(s))^+/K \geq \pi_0/(2K)$. This proves (35). In a symmetric equilibrium, the dealer-specific policies coincide with p^*. The same bounds apply dealer by dealer in an asymmetric equilibrium. □

Remark A7 (Scope and remaining existence problem)**.** *Proposition 6 proves incentive bounds conditional on a disclosure equilibrium satisfying Assumption 1. It does not construct that equilibrium. The positive lower-tail conclusion is proved for two potential dealers. Extending it to $N \geq 3$ requires a joint event on which a single bid clears every remaining rival while retaining positive expected surplus. Separate marginal rent bounds for individual rivals do not establish such an event.*

The available continuation analysis constructs a two-talker equilibrium under separate local smoothness and log-concavity conditions. It also constructs a true-payoff one-talker equilibrium with a bid floor and removes that floor on the lower-endpoint-atom branch. General unfloored one-talker existence under full lower support remains open. All-silent existence also remains open. At the all-silent history, product-order single crossing fails for an admissible monotone rival rule, and the full monotone strategy class is not invariant under best replies at that rule. These results rule out a broad fixed-point route but do not prove nonexistence of an all-silent equilibrium or of a smaller invariant class.

A complete disclosure-equilibrium proof must establish all-silent existence for every policy in the endogenous band, place the continuation profiles in an explicit standard Borel strategy space, prove that the equilibrium correspondence has an analytic graph, and obtain a continuation selection jointly measurable in histories and policies. It must then show that the induced disclosure gain is stable under weak-star convergence of policies. Until these steps are completed, measurable continuation selection and the disclosure fixed point remain open.

A.4.13 Failure conditional on the talk count

For any finite number of dealers, suppose the seller uses $R = \widehat{v} + \omega + \xi_s$, where ξ_s is independent of dealer information and randomization and is not observed by dealers before bidding. Let n_T be the realized number of talkers, and let $\mathcal{X}$ denote controls measurable with respect to dealer-side primitives. Define $L = B - \widehat{v} - \omega$. If no dealer bids, set $B = -\infty$ and $F_\xi(-\infty) = 0$.

Conditional on dealer information, bids, and the public history, failure occurs exactly when

$\xi_s > L$. Independence of the seller shock and iterated expectations give

$$\Pr(\text{failure} \mid n_T = k, \mathcal{X}) = \mathbb{E}[1 - F_\xi(L) \mid n_T = k, \mathcal{X}] \tag{A.218}$$

whenever the conditioning event has positive probability.

Equation (A.218) permits arbitrary dependence between the highest bid and the seller's estimate. Because $1 - F_\xi$ is bounded and nonincreasing, a sufficient condition for failures to decrease from count k to count $k+1$ is that $L \mid (n_T = k+1, \mathcal{X})$ first-order stochastically dominate $L \mid (n_T = k, \mathcal{X})$. The Gaussian information experiment establishes no such ordering. Endogenous talk counts select auctions by the information that induced disclosure and by the continuation bids generated by that information. The posterior-MSE result and the disclosure-tail theorem therefore do not by themselves imply that auctions with more talk have lower equilibrium failure probabilities.

Appendix B The Leave-One-Out Within-List Exercise

This appendix reports the leave-one-out exercise summarized in Section 6.2. For CUSIP i on list L_i, define the leave-one-out average number of talks on the other CUSIPs in the same list,

$$Z_i = \frac{1}{|L_i| - 1} \left(\sum_{j \in L_i} N_j - N_i \right), \tag{B.219}$$

and consider the two-stage procedure that projects N_i on Z_i (and the controls) and estimates

$$Y_i = \beta_0 + \beta_1 \widehat{N}_i + \beta_2' \mathbb{X}_i + \alpha_{\ell(i)} + \varepsilon_i, \tag{B.220}$$

where $\alpha_{\ell(i)}$ denotes list fixed effects and Y_i is an outcome such as trade failure or cover disclosure.

This is not an instrumental-variables design, and we do not present it as one. Within a list ℓ, the leave-one-out measure can be written as $Z_i = (T_\ell - N_i)/(|L_\ell| - 1)$, where $T_\ell = \sum_{j \in L_\ell} N_j$ is the list total and $|L_\ell|$ the list size. Once list fixed effects absorb T_ℓ and $|L_\ell|$, the remaining within-list variation in the measure is an exact linear transformation of the endogenous regressor itself, $\widetilde{Z}_i = -\widetilde{N}_i/(|L_\ell| - 1)$. The projection of N_i on Z_i is therefore mechanical, any endogeneity in N_i is inherited by Z_i with the opposite sign, and the resulting estimator is a weighted within-list regression of the outcome on N_i, with weights driven by heterogeneity in list size $|L_\ell|$ across lists. Accordingly, Table B.1 reports no first-stage F-statistics, and the leave-one-out columns should be read as reweighted within-list associations. Panel A displays the (mechanical) projection coefficient for completeness; its magnitude reflects the list-size weighting, not attention substitution.

Panel B reports the trade-failure estimates. The leave-one-out estimate (-0.008) is substantially smaller in magnitude than the OLS estimate (-0.060), indicating that the OLS association likely overstates the relationship due to reverse causality or omitted variables: CUSIPs that are ex ante easier to trade may both attract more talks and exhibit lower failure rates. Taken at face value, the leave-one-out estimate associates an additional talker with a 0.8 percentage point lower probability of trade failure, with larger magnitudes for junior tranches (e.g., -1.6 percentage points for BB/B). Panel C shows the corresponding cover-disclosure estimates; the leave-one-out estimate associates each additional talker with a 2.8 percentage point higher likelihood that the cover is disclosed.

A design that could support causal statements would require variation not algebraically constructed from N_i, for example predetermined dealer coverage, external dealer workload, broker routing, operational outages, or historical talk propensities measured outside the current list. We

leave such designs to future work. Standard errors clustered by list would additionally be appropriate throughout this table.

Table 6: Winner's Talk vs. Others' Talk

This table reports the winner's-talk regressions for talker-won auctions. All variables are measured relative to the ML benchmark. Panel A estimates equation (24), the decomposition into the winner's own talk and the others' average. Panel B estimates equation (25), where the equal-weight null implies $\gamma = 0$. Panel C estimates equation (26), the restriction $\beta + \gamma = 1$ of Panel B with the price premium $P_i - \bar{S}_i$ as the outcome, whose ρ feeds the winner-selection adjustment of the crossing test. Standard errors clustered by security in parentheses. ***, **, * denote significance at 1%, 5%, 10%.

	Dep. var. Price Surprise $\tilde{P}_i$ (Talker Wins Only)				
	(1) All	(2) AAA/AA	(3) A/BBB	(4) BB/B	(5) Equity
Panel A. Own Talk versus Others' Average					
Winner's talk $\tilde{S}_i$	0.273***	0.319***	0.174***	0.174***	0.386***
	(0.022)	(0.034)	(0.028)	(0.031)	(0.075)
Others' average $\tilde{\bar{S}}_{-i}$	0.754***	0.673***	0.832***	0.852***	0.658***
	(0.021)	(0.034)	(0.027)	(0.029)	(0.075)
Implied contribution					
Own talk	27%	32%	17%	17%	37%
Others' talks	73%	68%	83%	83%	63%
R^2	0.973	0.994	0.970	0.973	0.984
Obs	16,342	4,872	3,960	4,942	1,484
Panel B. Consensus and Own Talk					
Consensus $\tilde{\bar{S}}_i$ (β)	0.749***	0.822***	0.978***	0.881***	0.637***
	(0.024)	(0.035)	(0.051)	(0.033)	(0.066)
Winner's talk $\tilde{S}_i$ (γ)	0.231***	0.169***	0.016	0.112***	0.335***
	(0.026)	(0.034)	(0.058)	(0.035)	(0.075)
R^2	0.973	0.994	0.964	0.971	0.986
Consensus only, $\gamma = 0$ (β)	0.965***	0.990***	0.994***	0.989***	0.933***
	(0.004)	(0.003)	(0.011)	(0.009)	(0.009)
R^2	0.971	0.994	0.964	0.971	0.983
Panel C. Unit Total Pass-Through					
$S_i - \bar{S}_i$ (ρ)	0.164***	0.160***	0.055**	0.077***	0.243***
	(0.023)	(0.034)	(0.027)	(0.026)	(0.051)
R^2	0.278	0.490	0.470	0.383	0.508

Table 7: Nontalker Information and Consensus Predictiveness

This table reports regressions of price surprises on consensus surprises and nontalker winner indicators. All variables are measured relative to the ML benchmark. Panel A reports the consensus coefficient from the predictiveness regression with day fixed effects, estimated separately by winner type within each seniority group. The formal test of the difference between winner types is the interaction in Panel B, which pools the sample with common day effects. Panel B reports the interaction regressions of equation (27). Standard errors clustered by security in parentheses. ***, **, * denote significance at 1%, 5%, 10%.

Panel A. Consensus Predictiveness by Winner Type

	Dep. var. Price Surprise $\tilde{P}_i$ (Talker/Nontalker Wins Only)				
	(1) All	(2) AAA/AA	(3) A/BBB	(4) BB/B	(5) Equity
Consensus surprise $\tilde{\tilde{S}}_i$, talker-won	0.962***	0.990***	0.994***	0.986***	0.932***
	(0.004)	(0.003)	(0.010)	(0.008)	(0.008)
Consensus surprise $\tilde{\tilde{S}}_i$, nontalker-won	0.935***	0.980***	0.978***	0.958***	0.895***
	(0.004)	(0.006)	(0.011)	(0.008)	(0.010)
R^2, talker-won	0.971	0.994	0.964	0.971	0.983
R^2, nontalker-won	0.952	0.984	0.921	0.969	0.964

Panel B. Interaction Regressions

	Dep. var. Price Surprise $\tilde{P}_i$				
	(1) All	(2) AAA/AA	(3) A/BBB	(4) BB/B	(5) Equity
Consensus surprise $\tilde{\tilde{S}}_i$	0.961***	0.985***	0.992***	0.981***	0.932***
	(0.004)	(0.003)	(0.011)	(0.008)	(0.007)
Nontalker$_i$	0.127***	0.037***	0.129**	0.124***	0.152**
	(0.020)	(0.011)	(0.052)	(0.041)	(0.073)
$\tilde{\tilde{S}}_i \times$ Nontalker$_i$	−0.025***	−0.002	−0.012	−0.015*	−0.035***
	(0.006)	(0.005)	(0.015)	(0.009)	(0.010)
R^2	0.958	0.987	0.936	0.967	0.969
Obs	37,026	12,906	9,361	9,742	4,061

Table 8: Winning Margin

This table reports the winning margin (price minus cover, in percentage points of par) by winner type. Panel A reports summary statistics. Panel B reports regressions with day fixed effects and standard errors clustered by security. ***, **, * denote significance at 1%, 5%, 10%.

Panel A. Winning Margin by Winner Type

	(1) All	(2) Talker Wins	(3) Nontalker Wins	(4) Difference
Winning margin (pp of par)	0.240	0.247	0.233	−0.014***
				(0.005)
By seniority				
AAA/AA	0.090	0.081	0.095	0.014***
A/BBB	0.237	0.226	0.246	0.020**
BB/B	0.371	0.366	0.376	0.009
Equity	0.412	0.431	0.398	−0.033*
By talk intensity				
Few talkers ($N_i \leq 5$)	0.253	0.277	0.247	−0.029**
Many talkers ($N_i > 5$)	0.236	0.244	0.226	−0.018***

Panel B. Regression Evidence

	Dep. var. Winning Margin (Price – Cover)						
	(1) All	(2) All	(3) All	(4) AAA/AA	(5) A/BBB	(6) BB/B	(7) Equity
Nontalker$_i$	−0.025***	−0.017***	−0.005	0.007*	0.005	−0.030**	−0.043*
	(0.005)	(0.005)	(0.015)	(0.004)	(0.009)	(0.013)	(0.023)
N_i		0.006***	0.007***	−0.000	−0.000	−0.012***	−0.006
		(0.001)	(0.001)	(0.001)	(0.002)	(0.004)	(0.005)
Nontalker$_i \times N_i$			−0.001				
			(0.001)				
Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.198	0.199	0.199	0.392	0.335	0.285	0.341
Obs	30,322	30,322	30,322	10,590	7,974	7,559	3,216

Table 9: Disclosure Regimes Compared

This table compares disclosure regimes in the numerical counterfactual of Section 7.1 ($\mu_0 = 0$, $\sigma_v = 1$, $\sigma_\varepsilon = 0.8$, $N = 8$, $\bar{u} = 0.02$, 40,000 simulated auctions on common random numbers). Revenue, dealer surplus, total surplus, and allocative loss are per-auction means in units of the common value. Pr(efficient) is the probability that the dealer with the highest idiosyncratic value wins. Corr(price, v) is the correlation between the transaction price and the common value. Revenue and dealer surplus are not reported for the voluntary regime, whose silent-side bid levels are not pinned down.

Regime	Revenue	Dealer surplus	Total surplus	Allocative loss	Pr(efficient)	Corr(price, v)
Banned (no talk)	−0.056	0.066	0.0111	0.0078	0.13	0.87
Voluntary (selective talk)	—	—	0.0120	0.0069	0.27	0.44
Mandated (all talk)	0.018	0.001	0.0189	0.0000	1.00	0.96
Mandated − Banned	0.0735	−0.0657	0.0078			

Table B.1: Price Talk and Market Quality: Leave-One-Out Weighted Within-List Associations

This table reports within-list associations between talk intensity and market-quality outcomes. The OLS columns regress the outcome on the number of talks N_i. The LOO columns replace N_i with its projection on the leave-one-out measure (B.219). Because list fixed effects absorb the list total and list size, this projection is an exact within-list transformation of N_i, so the LOO columns are weighted within-list associations rather than instrumental-variables estimates. Panel A reports the projection coefficient. All specifications include controls, day fixed effects, and list fixed effects.

Panel A. Projection of N_i on the Leave-One-Out Measure (Mechanical)

	Dep. var. Number of Talks N_i					
	(1) All	(2) All	(3) AAA/AA	(4) A/BBB	(5) BB/B	(6) Equity
Z_i (leave-one-out mean)	−5.253***	−5.249***	−5.712***	−5.440***	−5.537***	−4.145***
	(0.115)	(0.115)	(0.249)	(0.288)	(0.200)	(0.482)

Panel B. Trade Failures

	Dep. var. Trade Failure$_i$					
	OLS	LOO	LOO. AAA/AA	LOO. A/BBB	LOO. BB/B	LOO. Eq
N_i	−0.060***	−0.008***	−0.002*	−0.013***	−0.016***	−0.020
	(0.002)	(0.001)	(0.001)	(0.002)	(0.003)	(0.018)
Obs	48,277	47,079	16,784	13,072	15,179	2,044

Panel C. Cover Disclosure

	Dep. var. Cover Disclosed$_i$					
	OLS	LOO	LOO. AAA/AA	LOO. A/BBB	LOO. BB/B	LOO. Eq
N_i	0.033***	0.028***	0.024***	0.026***	0.030***	−0.020
	(0.004)	(0.001)	(0.001)	(0.002)	(0.003)	(0.025)
Obs	35,581	35,653	14,501	9,905	10,145	1,102

Controls, Day FE, List FE in all specifications

Table B.2: The Within-Dealer Silence Premium

This table re-estimates the crossing regressions from Table 3 adding dealer fixed effects. The sample is restricted to auctions won by dealers who also win at least one auction in the opposite talk status (95% of auctions). All specifications include dealer and day fixed effects. Standard errors clustered by security in parentheses. ***, **, * denote significance at 1%, 5%, 10%.

Panel A. Price Premium ($P_i - \bar{S}_i$)					
	(1) All	(2) AAA/AA	(3) A/BBB	(4) BB/B	(5) Equity
Nontalker$_i$	0.099***	0.026***	0.071***	0.150***	0.188***
	(0.014)	(0.009)	(0.023)	(0.034)	(0.068)
R^2	0.198	0.435	0.371	0.334	0.398
Obs	34,922	12,065	8,859	9,213	3,786
Panel B. Max-Talk Crossing ($P_i - S_i^{1:N_i}$)					
	(1) All	(2) AAA/AA	(3) A/BBB	(4) BB/B	(5) Equity
Nontalker$_i$	0.176***	0.068***	0.146***	0.230***	0.128
	(0.024)	(0.012)	(0.037)	(0.056)	(0.107)
R^2	0.228	0.413	0.419	0.394	0.425
Obs	34,922	12,065	8,859	9,213	3,786
Dealer FE, Day FE in all specifications					

Table B.3: Evidence of Learning from Price Talk

Panel A reports average winning-bid revisions ($P_i - S_i$, in percentage points of par) for winning dealers, conditional on whether the winner's talk is below or above consensus. Because the sample conditions on winning, revisions partly reflect winner selection. Panel B reports within-auction dispersion measured as the gap between the two highest observations, averaged across auctions with a disclosed cover, in percentage points of par. Talk dispersion is the highest minus the second-highest talk. Bid dispersion is the price minus the cover. Both are winsorized at the 1st and 99th percentiles.

Panel A. Bid Revisions Toward Consensus								
	(1) All	(2) AAA	(3) AA	(4) A	(5) BBB	(6) BB	(7) B	(8) Equity
All revisions	−0.18	−0.03	−0.05	−0.04	−0.19	−0.19	−0.62	−0.84
Talk below consensus	0.47	0.14	0.27	0.49	0.54	0.74	0.48	0.75
Talk above consensus	−0.51	−0.12	−0.23	−0.32	−0.55	−0.69	−1.25	−1.53
Panel B. Dispersion Reduction								
Talk dispersion	0.482	0.122	0.258	0.428	0.545	0.679	0.970	1.690
Bid dispersion	0.226	0.066	0.148	0.185	0.270	0.354	0.429	0.563
Reduction (%)	53	46	43	57	50	48	56	67