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.

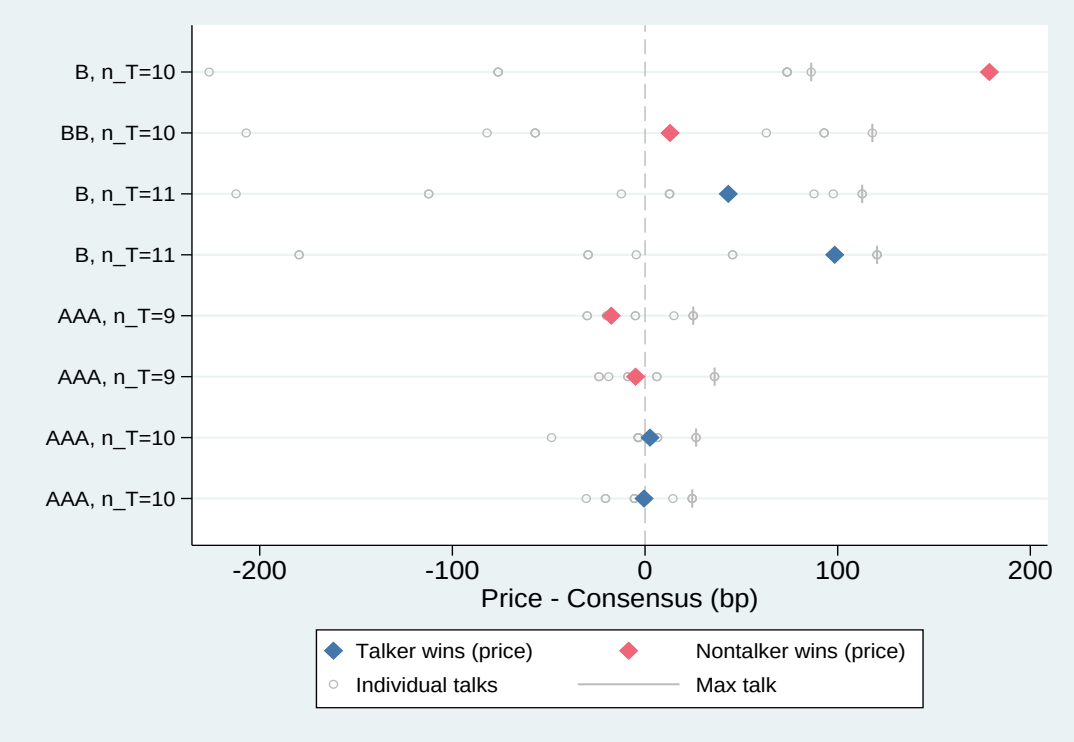

Figure 1: Anatomy of BWIC auctions. Each row is one auction, centered on its consensus $\bar{S}_i$ (dashed line at zero). Open circles are individual talks, and diamonds are transaction prices for talker-won (blue) and nontalker-won (red) auctions. Rows are labeled by seniority and number of talks.

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			

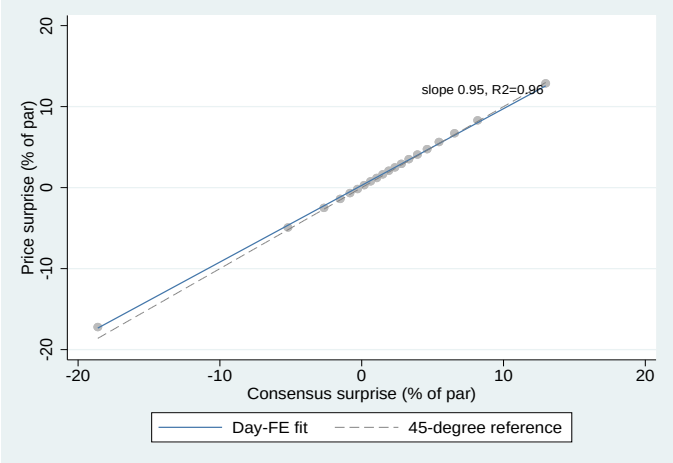

Figure 2: Consensus surprise and price surprise. Binscatter of the price surprise $\tilde{P}_i$ against the consensus surprise $\tilde{\tilde{S}}_i$, in percent of par, in twenty equal-count bins. The solid line is the day-fixed-effects fit from Table 2, Panel B, and the dashed line is the 45-degree 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.

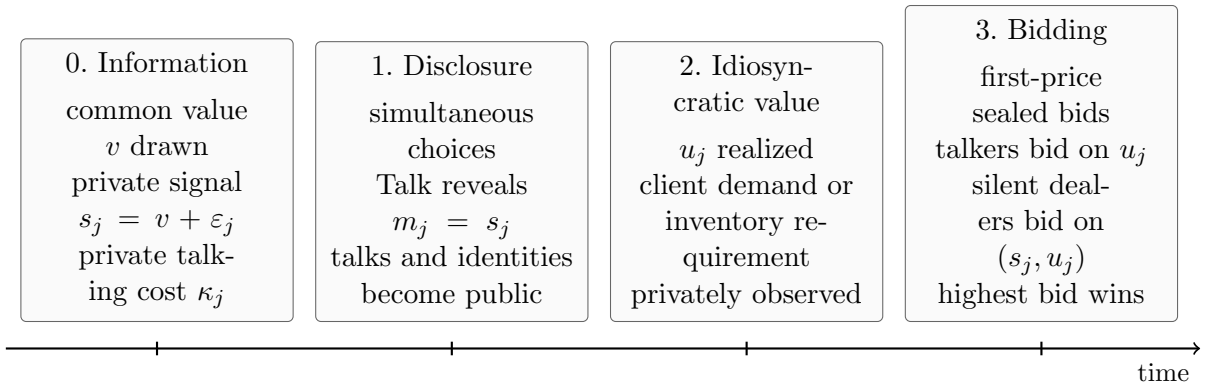

Figure 3: Model timing. Dealers observe their signals s_j and talking costs κ_j and decide whether to talk. Idiosyncratic values u_j are then realized, and dealers submit sealed bids. The common value v is never observed.

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

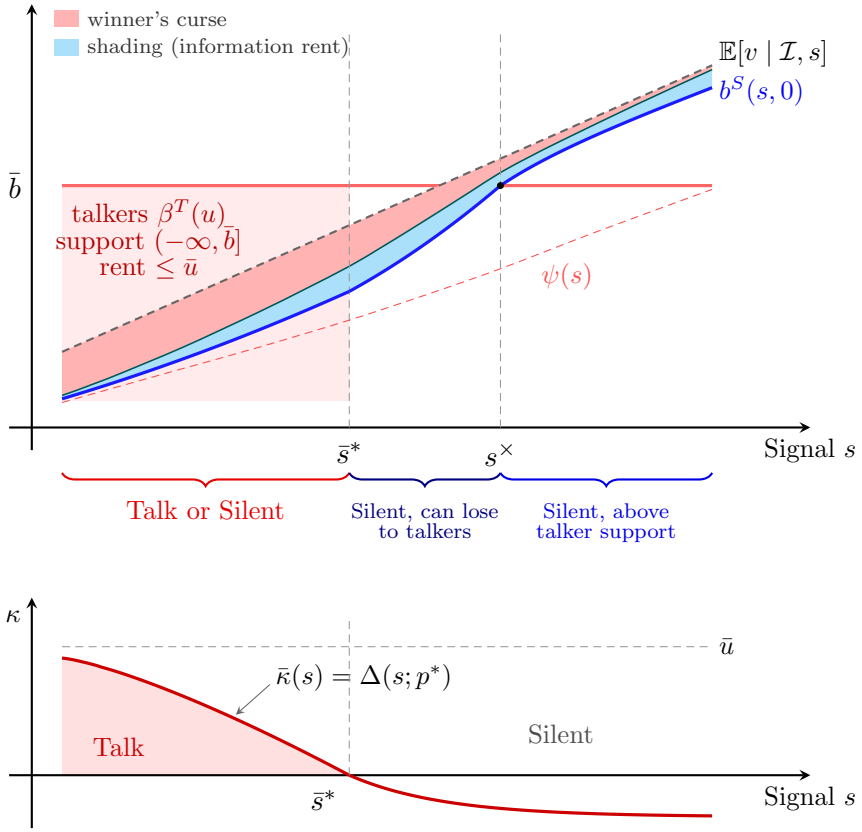

Figure 4: Disclosure and bidding in equilibrium. Top panel: the zero-value silent bid $b^S(s,0)$ (blue), the silent dealer's posterior mean $\mathbb{E}[v \mid \mathcal{I}, s]$ (dashed), and the win-conditional mean (teal). The upper shaded wedge is the winner's-curse discount and the lower wedge is bid 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

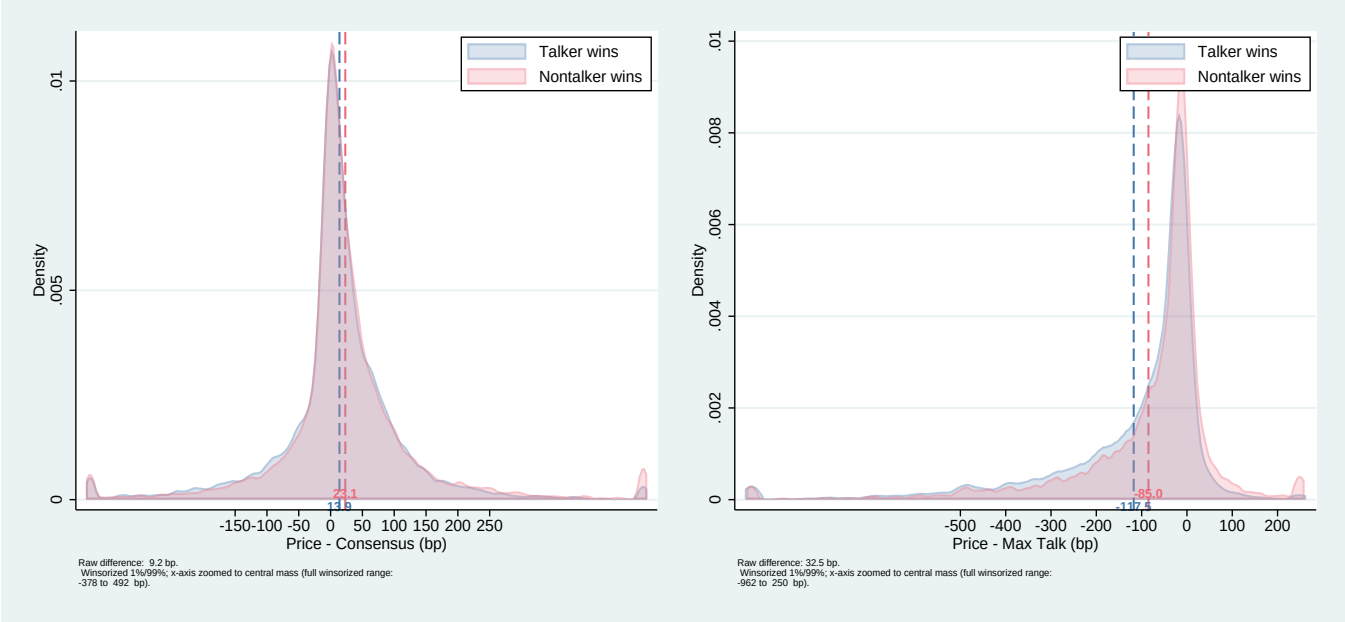

Figure 5: Crossing distributions by winner type. Kernel densities of the price premium $P_i - \bar{S}_i$ (left) and the max-talk crossing $P_i - S_i^{1:N_i}$ (right), in bp, for talker-won (blue) and nontalker-won (red) auctions. Dashed lines mark group means, winsorized at 1% and 99%, so they differ 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

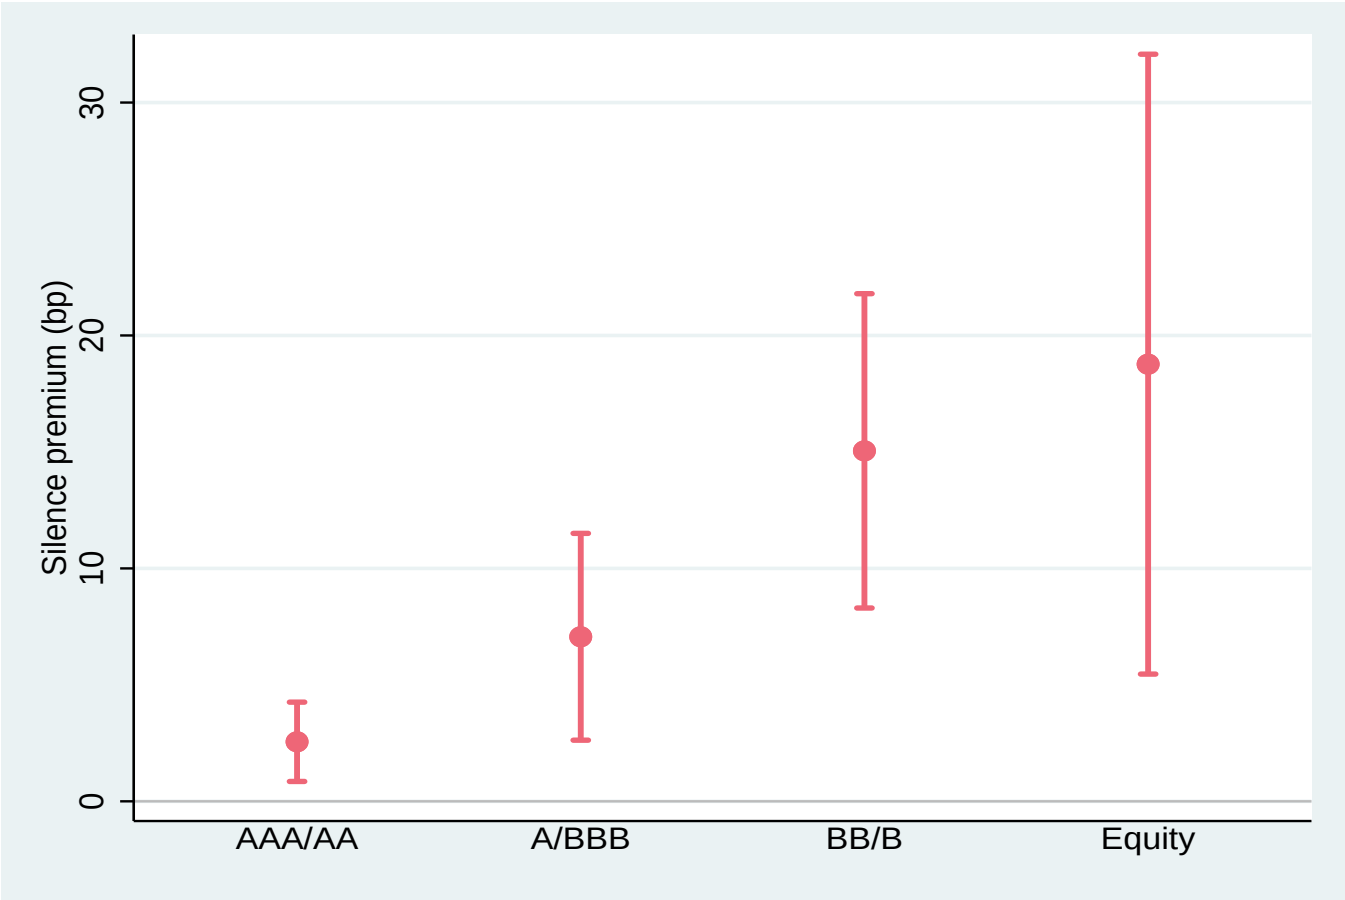

Figure 6: Within-dealer silence premium by seniority. Each point is the Nontalker$_i$ coefficient for the price premium $P_i - \bar{S}_i$ from the within-dealer regressions in Table B.2, Panel A, columns (2) to (5), in bp. Whiskers are 95% confidence intervals based on security-clustered standard errors.

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.

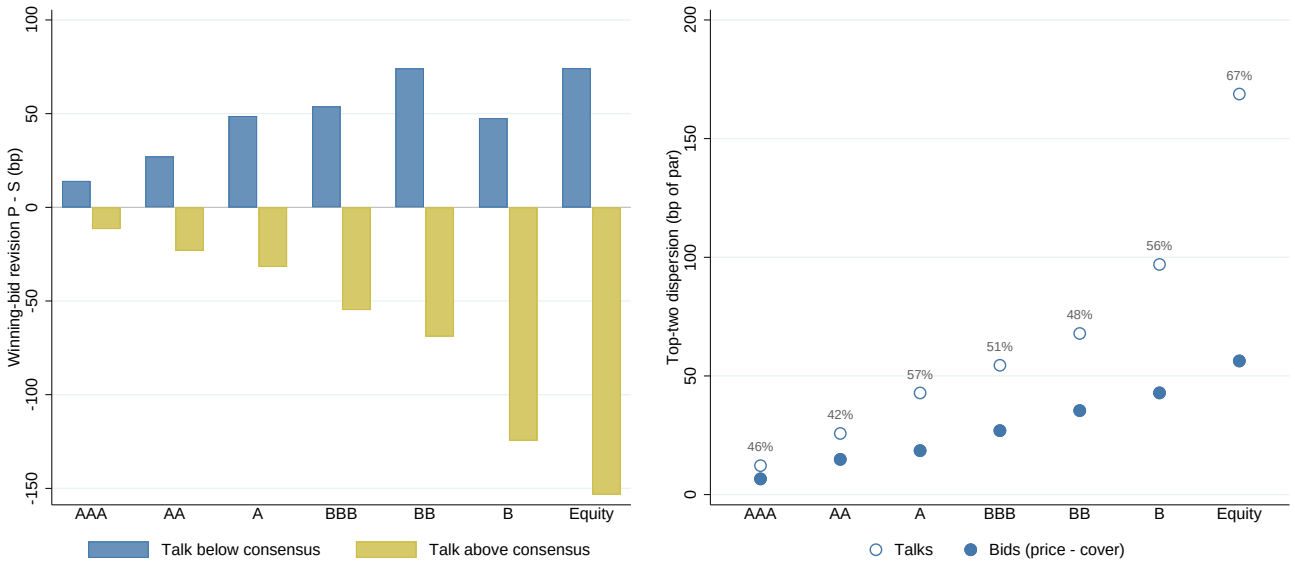

Figure 7: Winning-bid revisions and dispersion. Left: average winning-bid revision $P_i - S_i$ for winners whose talk lies below (blue) or above (yellow) the consensus. Right: top-two dispersion of talks (open) and bids (filled), with the percentage reduction labeled. Underlying values are in Table B.3.

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