

Taking Orders and Taking Notes: Dealer Information Sharing in Treasury Markets*

Nina Boyarchenko, David O. Lucca, and Laura Veldkamp

Federal Reserve Bank of New York and
New York University, Stern School of Business

June 28, 2016

Abstract

The use of order flow information by financial firms has come to the forefront of the regulatory debate. A central question is: Should a dealer who acquires information by taking client orders be allowed to use or share that information? We explore how information sharing affects dealers, clients and issuer revenues in U.S. Treasury auctions. Because one cannot observe alternative information regimes, we build a model, calibrate it to auction results data, and use it to quantify counter-factuals. We estimate that yearly auction revenues with full-information sharing (with clients and between dealers) would be \$5 billion higher than in a “Chinese Wall” regime in which no information is shared. When information sharing enables collusion, the collusion costs revenue, but prohibiting information sharing costs more. For investors, the welfare effects of information sharing depend on how information is shared. Surprisingly, investors benefit when dealers share information with each other, not when they share more with clients. For the market, when investors can bid directly, information sharing creates a new financial accelerator: Only investors with bad news bid through intermediaries, who then share that information with others. Thus, sharing amplifies the effect of negative news. Tests of two model predictions support its key features.

*The views expressed here are the authors’ and are not representative of the views of the Federal Reserve Bank of New York or of the Federal Reserve System. We thank Bruno Biais, Giovanni Cespa, John Cochrane (through his blog), Darrell Duffie, Ken Garbade, Luis Gonzalez, Robin Greenwood, Zhiguo He, Pablo Kurlat, Richard Tang, Hongjun Yan, Haoxiang Zhu and participants at the NBER AP meetings, Chicago Booth Asset Pricing Conference, Duke, Miami, U.S. Treasury Roundtable on Treasury Markets, NYU, the Spring 2015 Macro Finance Society Meetings, SED 2015, Gerzensee Summer Institute 2015, EFA 2015, EEA 2015, FIRS 2016, Bank of Canada and SFI Lausanne for comments. Nic Kozeniauskas, Karen Shen, Arnav Sood and Peifan Wu provided excellent research assistance. Emails: nina.boyarchenko@ny.frb.org; david.lucca@ny.frb.org; lveldkam@stern.nyu.edu.

“[B]efore the Treasury holds an auction, salespeople at 22 primary dealers field billions of dollars in bids for government debt. Traders working at some of these financial institutions have the opportunity to learn specifics of those bids hours ahead of the auctions [and] also have talked with counterparts at other banks via online chatrooms [...]. Such conversations, both inside banks and among them, could give traders information useful for making bets on one of the most powerful drivers of global markets [...].” — Bloomberg (2015), “As U.S. Probes \$12.7 Trillion Treasury Market, Trader Talk Is a Good Place to Start.”

Recent financial market misconduct, involving misuse of information about clients’ orders, cost the firms involved record fines and lost reputation. It also prompted investigations and calls for curbing dissemination of order flow information, between and within dealers. Recent investigations reportedly involve U.S. Treasury auctions (Bloomberg, 2015 above). But the use of order flow information has been central to our understanding of Treasury auctions (Hortaçsu and Kastl, 2012), to market making theory generally (Kyle, 1985) and to market practice for decades. In describing Treasury market pre-auction activities in the 1950s, Robert Roosa (1956) noted that “Dealers sometime talk to each other; and they all talk to their banks and customers; the banks talk to each other.” Furthermore, sharing order-flow information—or, colloquially, “market color”—with issuers is even mandatory for primary dealers both in the U.S. and abroad. Of course, if information sharing leads to collusion, that has well-known welfare costs. But if collusion could be prevented with separate remedies, is information sharing by itself problematic? The strong conflicting views on a seemingly long-established practice raise the question of who gains or loses when order-flow information is shared.¹

Measuring the revenue and welfare effects of information sharing directly would require data with and without sharing. In the absence of such data, we use a calibrated model. Our setting is an institutionally-detailed model of U.S. Treasury auctions, which we select because of the available data, the absence of other dealer functions,² and their enormous economic importance. In the model, dealers observe client orders and may use that in-

¹Thus far actions for misconduct have been successfully brought against participants in the interbank lending (Libor) and foreign exchange markets. Regulations on information sharing in sovereign auctions vary and are evolving. As of 2011, the UK Debt Management Office sanctioned that UK primary dealers, or Gilt-edged Market Makers, “whilst not permitted to charge a fee for this service, may use the information content of that bid to its own benefit” (GEMM Guidebook, 2011). The 2015 GEMM Guidebook, instead, states that “information about trading interests, bids/offers or transactions may be subject to confidentiality obligations or other legal restrictions on disclosure (including pursuant to competition law). Improper disclosure or collusive behaviour will fall below the standards expected of GEMMs, and evidence or allegations of such behaviour will be escalated to the appropriate authority(ies).” We are not aware of analogous rules in the context of U.S. Treasury auctions. In practice, a financial intermediary’s use of client information, including sharing such information with other clients or using the information for other benefit to such intermediary, may violate legal requirements, be they statutory, regulatory or contractual, market best practices or standards. This paper does not take a view as to whether the described use of client information with respect to Treasury auction activity is legal or proper. The objective of the paper is to study the economic effects of alternative information sharing arrangements.

²Dealers in Treasury auctions do not diversify or transform risks, do not locate trading counterparties and cannot monitor issuers because they cannot influence fiscal policy.

formation to inform their own strategy, share some of the information with clients, or exchange information with other dealers. Then all agents submit menu bids to a uniform-price, common-value auction. To quantify the effects of order flow information sharing and sign welfare results, we calibrate the model to auction results and allotment data as well as market pricing information and on-the-run premia. Then we compare the model’s revenue and utility predictions with varying degrees and types of information sharing. Finally, we derive two testable predictions from our model and show that both are supported by auction data.

The model teaches us that the primary beneficiary of information sharing is the U.S. Treasury, who benefits from the higher bids of better-informed buyers. We estimate that moving from full information sharing benchmark to a “Chinese wall” policy of no information-sharing between or within dealers would lower Treasury auction revenues by \$4.8 billion annually. While the idea that better-informed investors bid more is not a new finding, the issue is rarely raised in policy debates, presumably because the magnitude of the effect is not known.

Our second finding is that dealer information sharing with other dealers and sharing with clients have opposite effects on investor utility. When all dealers share information with their clients, it typically makes the clients worse off. This is a form of the well-known [Hirshleifer \(1971\)](#) effect, which arises here because better-informed clients have more heterogeneous beliefs and therefore share risk less efficiently. But surprisingly, when dealers share information with each other and then transmit the same amount of information to their clients, investor welfare improves. Our model shows how inter-dealer information sharing makes beliefs more common, and thereby improves risk-sharing and welfare. In essence, information sharing with clients is similar to providing more private information, while inter-dealer sharing functions effectively makes information more public.

Third, since information sharing has been associated with coordinated trades in foreign-exchange misconduct (for example, to manipulate benchmark rates), we consider a setting in which dealers who share information also collude. In a collusive equilibrium, dealers who share information also bid as a group, or coalition, that considers price impact of the coalition as a whole. We find that dealer information sharing and collusion jointly suppress auction prices and reduce Treasury revenue. However, if the dealers share enough information with clients, the revenue costs may disappear.

Fourth, we uncover a new financial accelerator: Only investors with bad news employ intermediaries, who then share that information with others. Thus the combination of information sharing and intermediation choice can amplify the effect of negative news and raise the probability of auction failure.

These findings are not meant to imply that dealers should have *carte blanche* in using information in any way they choose. The model assumes that clients know how dealers use their information, and that order flow information is aggregated. While we consider the case of collusive bids by dealers, our setting does not clearly span the range of malpractices that may have been undertaken. In effect, we ask: If dealers disclose how information is used, what are the costs and benefits of limiting information sharing?

Treasury auctions are unique in their importance and their complexity. Our model balances a detailed description with a tractable and transparent model which highlights insights that are broadly applicable. The basis for the model is a standard, common-value, uniform-price auction with heterogeneous information, limit orders and market orders. On top of this foundation, we add five features that distinguish Treasury auctions from other settings.

Feature 1: Dealers learn and share order flow information. The assumption that bidders have private signals about future Treasury values and that dealers learn from observing their order flow is supported by [Hortaçsu and Kastl \(2012\)](#). Using data from Canadian Treasury auctions, they find that order flow is informative about demand and asset values. They further show that information about order flow accounts for a significant fraction of dealers’ surplus. In our setting, dealers not only collect this information but also share it.

Feature 2: Strategic bidding [Bikhchandani and Huang \(1993\)](#) present evidence that Treasury auctions are imperfectly competitive (see also [Song and Zhu, 2016](#)). Because of their bidding volume (40 percent of the total), primary dealers, which currently are only 22, bid strategically by taking into account the price impact of their bids. As in [Kyle \(1989\)](#), the strategic aspect of primary dealers’ bids is a central feature of our model. Without it, for example, the choice of intermediated or direct bidding that we discuss below would be trivial. By including the various types of bidders, our model predicts not only average post-auction appreciation (as in [Lou, Yan, and Zhang, 2013](#)), but also a relationship between post-auction price appreciation and auction allocations by investor types.

Feature 3: Non-competitive bidding In every auction, a group of bidders, labelled “non-competitive,” place market orders that are not conditional on the market-clearing price. Such bidders are an important feature of the model, because they prevent the price from perfectly aggregating all private information. We assume that competitive bidders are different in that they submit limit orders (quantity schedules that are conditional on realized prices). We show in the data that competitive bids incorporate information in realized prices, just as they do in our model.

Feature 4: Minimum bidding requirements Primary dealers are expected to bid at

all auctions an amount equal to the pro-rata share of the offered amount, with bids that are “reasonable” compared to the market. A dealer that consistently fails to bid for a large enough quantity at a high enough price could lose his primary dealer status.³ We model this requirement as a shadow cost for low bids.

Feature 5: Direct and indirect bidding U.S. Treasury auctions are mixed auctions, meaning that investors can choose to bid indirectly, through a dealer, or directly, without any intermediary. Section 5 examines a large, strategic bidder’s choice of how to bid.

Each of these model features contributes to our understanding of the symbiotic relationship between investors and intermediaries: it is the process of intermediating trades that reveals information to dealers and thus empowers them. Information sharing is what induces clients to use intermediaries and induces large investors to intermediate.

Contribution to the existing literature. Our paper contributes to several strands of literature. First, it is connected to work in the microstructure literature that studies how order flow information contributes to price formation. For example, dealers learn from sequential order flow in [Easley, Kiefer, O’Hara, and Paperman \(1996\)](#) and leverage asymmetric information and market power in [Kyle \(1985\)](#) and [Medrano and Vives \(2004\)](#). In [Babus and Parlato \(2015\)](#), dealers fragment a market. They consider how fragmentation inhibits risk-sharing, while we consider the effect on information-sharing. What distinguishes our model most from previous work is its analysis of information-sharing and its attention to the institutional features of Treasury auctions. While, as in most models with exponential utility, the equilibrium price is a linear function of the signals in the economy, the information sharing changes the linear weights placed on each signal. The institutional detail we add is not just window-dressing on a standard model. Intermediation choice and minimum bidding requirements are what make information sharing necessary. When bidders can choose to bid directly, dealers must share some of their information in order to attract clients. Conversely, allowing dealers to extract information from order flow is what induces them to be primary dealers and subject themselves to costly minimum bidding (underwriting) requirements. Without this both types of information sharing, the primary dealer system as we know it could not operate.

Primary dealers perform underwriting services, as they do in the initial public offering (IPO) literature. This literature typically finds that intermediaries lower issuers’ revenues but also revenue variance ([Ritter and Welch, 2002](#)). We show, instead, that when deal-

³Prior to 1992, an active primary dealer had to be a “consistent and meaningful participant” in Treasury auctions by submitting bids roughly commensurate with the dealer’s capacity. See appendix E in [Brady, Breedon, and Greenspan \(1992\)](#). In 1997, the New York Fed instituted an explicit counterparty performance scorecard and dealers were evaluated based on the volume of allotted securities. In 2010 the NY Fed clarified their primary dealer operating policies and strengthened the requirements. See New York Fed website for the most recent rules.

ers share information, the conventional wisdom of underwriting is reversed: information intermediaries raise expected revenue but also revenue variance.

Work by [Hortaçsu and McAdams \(2010\)](#) and others, who study how auction design affects revenues,⁴ complements our project. Similarly, [He, Krishnamurthy, and Milbradt \(2016\)](#) explore why US Treasuries are safe. We fix the auction format to a uniform-price menu auction, fix the distribution of future Treasury values, and instead focus on how intermediation and information sharing affect revenue and surplus.

The emergent literature on intermediary asset pricing also explores the idea that intermediaries are central to determining the equilibrium price of an asset. In [He and Krishnamurthy \(2013\)](#) and [Brunnermeier and Sannikov \(2014\)](#), capital-constrained intermediaries provide households with access to risky asset markets and thus improve economy-wide risk sharing. In contrast, agents in our model choose optimally whether to invest through an intermediary, whose role is to improve information sharing among investors. This new information-sharing role for financial intermediaries helps explain why agents that are not precluded from accessing markets directly might choose intermediation nonetheless.


1 Baseline Auction Model with Primary Dealers

The auction setting is a simple and familiar one, with a structure similar to [Kyle \(1989\)](#). The novel feature of the model lies in its rich information structure as determined by how information is shared between the agents. [Figure 1](#) summarizes the alternative sharing arrangements that we consider, for a simplified setting with no signal noise and only a few market participants. Dealers are denoted with the letter “D,” while investors with the letter “I.” Panel a) shows the case of no information sharing (“Chinese walls”), where each auction participant only observes his private information s_i . Competitive bidders can submit a menu of price-contingent quantities. Auction theory teaches us that each bidder should avoid the winner’s curse by choosing a quantity for each price that would be optimal if he observed that market-clearing price and included it in his information set. Thus, the information set of investor i is effectively $\{s_i, p\}$.

When information is shared between dealers and customers (panel b), an investor’s information set now not only includes her private signal but also the dealer’s, and the dealer

⁴Theoretical work by [Chari and Weber \(1992\)](#), [Bikhchandani and Huang \(1989\)](#), [Back and Zender \(1993\)](#), and [Wilson \(1979\)](#) considers the merits of uniform-price auctions versus other possible alternatives. Empirical work by [Nyborg and Sundaresan \(1996\)](#), [Malvey, Archibald, and Flynn \(1995\)](#) and [Malvey and Archibald \(1998\)](#) compares revenues from 1992-1998 when the U.S. Treasury used both uniform and discriminatory price auctions. [Armantier and Sbaï \(2006\)](#) use French Treasury auction bids to structurally estimate the benefits of uniform price auctions.

Figure 1: Information sets with alternative sharing assumptions. Letter D denotes dealers; letters I denotes investors (either large or small) bidding through a dealer or not (direct bidding); p is the equilibrium price. Dashed lines indicate common information sets.



also observes an extended information set. This information set is further increased in the case of cross-dealer information sharing (panel c). In this case, each investor observes the information set pertaining to its dealer, and also that of the other dealer. Common to case b) and c) is the fact that information sharing with dealers and customers leads to more signal pooling. Investors who bid independently from the intermediary keep their signal private (panel d) resulting in a more dispersed information set both for the direct bidder and other bidders.

While this simplified setting conveys the essence of information sharing, our model is richer along many dimensions. We consider four type of bidders to match key features of Treasury auction participation: small and large limit-order bidders, intermediaries (or dealers) and non-price contingent bidders. Limit-order bidders and intermediaries place price-contingent bids, which specify for each clearing price p , a price-quantity pair. Limit-order bidders can be small (price takers) or large (strategic bidders). We refer to large and

small limit-order bidders as the investors. Dealers are just like large limit-order bidders but they also intermediate bids from other limit-order bidders and face minimum bidding requirements. Dealers place bids directly in the auction while small and other large limit-order bidders bid indirectly through the dealers. (Section 5 allows direct bidding as well.) Non-price-contingent bidders are the fourth type of agent that places bids. In contrast to other investors and dealers, these bidders place market orders that only specify a quantity but not a price (called noncompetitive bids). In practice, noncompetitive bidders are small retail investors or foreign central banks who participate at auctions to invest dollar-denominated foreign reserve balances, for example by rolling expiring securities into newly issued ones. As opposed to other investors, this foreign official auction demand is not driven by the security fundamentals but by exchange rate policies, simply placing bids for a given amount of securities, and injecting noise in auction prices.

Agents, assets and preferences The model economy lasts for one period and agents can invest in a risky asset (the newly issued Treasury security) and a riskless storage technology with zero net return. The risky asset is auctioned by Treasury in a fixed number of shares (normalized to 1) using a uniform-price auction with a market-clearing price p . The fundamental value of the newly issued asset is unknown to the agents and normally distributed: $f \sim N(\mu, \tau_f^{-1})$. By assuming that there is one final value of the asset, we are describing a common value auction. But this is purely for convenience. Appendix C describes a private value setup that delivers the same results.

We index investors (small and large) and dealers with $i = 1, \dots, N$, where $N \equiv N_L + N_I + N_D$, N_L , N_I and N_D denote respectively the total number of large investors, small investors and dealers. Each small investor has initial wealth W_i , and chooses the quantity of the asset to hold, q_i (which can in principle be negative) at price p per share, in order to maximize his expected utility,⁵

$$\mathbb{E}[-\exp(-\rho_i W_i)], \tag{1}$$

where ρ_i denotes agent i 's coefficient of absolute risk aversion. In our data, investors with larger balance sheets hold larger positions of risky assets. To allow our model to match these wealth effects, we assign large investors and dealers a smaller absolute risk aversion. The budget constraint for small and large investors dictates that final wealth is $W_i = W_0 + q_i(f - p)$.

Dealer and large investors solve the same problem of small investors but they also internalize

⁵Technically, the price of each Treasury is fixed at par and auction participants bid coupon payments. We interpret p as the price per present value dollar of coupons.

the effect they have on market prices. They maximize their final utility with risk aversion ρ_i subject to the same constraints as well as the market clearing condition. We assume that all large investors and dealers share the same risk aversion and signal precision, or $\rho_i = \rho_l$ and $\tau_{\varepsilon,i} = \tau_{\varepsilon,l}$ for $i \in \{\mathcal{N}_L, \mathcal{N}_D\}$, where \mathcal{N}_j is the set of agents of type j . Similarly, all small investors are symmetric: $\rho_i = \rho_I$ and $\tau_{\varepsilon,i} = \tau_{\varepsilon,s}$ for $i \in \mathcal{N}_I$. The net quantity of market orders, x , is unknown to other investors and normally distributed $x \sim N(\bar{x}, \tau_x^{-1})$.

Minimum Bidding Requirements In the current design of the primary dealer system, dealers are expected to bid for a pro-rata share of the auction at “reasonable” prices compared to the market. A dealer may violate the minimum bidding requirement in any given auction. But over time, if a dealer is consistently allotted an insufficient share, his primary dealer status could be revoked. To capture the essence of this dynamic requirement in a static model, we model the bidding requirement as a cost levied on a dealer who purchases too little. This cost is a stand-in for the shadow cost of a dynamic constraint. Conversely, a dealer who purchases a large dollar amount of Treasuries faces a relaxed bidding constraint in the future. We model this benefit as a current transfer. Thus, for a dealer who purchases a dollar amount qp of Treasuries through the auction, we assume a low-bid penalty of $\chi_0 - \chi qp$. Thus, a dealer d 's budget constraint is

$$W_d = W_{0,d} + q_d(f - p) - (\chi_0 - \chi q_d p). \quad (2)$$

Describing Information Sets and Updating Beliefs with Correlated Signals Depending on the information structure, investors and dealers observe three pieces of information: 1) their own private signal, 2) signals from others who may share information with them and 3) the equilibrium price of the asset. We explain each in turn. Before trading, each investor and dealer gets a signal about the payoff of the asset. These signals are unbiased, normally distributed and have private noise:

$$s_i = f + \varepsilon_i,$$

where $\varepsilon_i \sim N(0, \tau_{\varepsilon,i}^{-1})$. In practice, the fundamental value of a newly-issued Treasury security depends on the term structure of interest rates, as implied by market prices of other Treasury securities, and the specific liquidity value of the newly auctioned security relative to older securities. In the model calibration, we focus on how information sharing affects the specific liquidity value taking as given the term structure, which is largely driven by monetary policy as opposed to information sharing at the auction.

Second, by placing orders through dealers, customers reveal their order flow to their dealer,

which in the model is equivalent to sharing their private signal. Each dealer d receives orders from an equal number of investors. He observes the orders of $\nu_I \equiv N_I/N_D$ small and $\nu_l \equiv N_L/N_D$ large investors. Since bids will turn out to be linear functions of beliefs, observing average bids and observing average signals is equivalent. The dealer can construct \bar{s}_d , which is an optimal signal-precision-weighted average of his and his clients' private signals as:

$$\bar{s}_d = \frac{\sum_{i \in \mathcal{I}_d} \tau_i s_i}{\sum_{i \in \mathcal{I}_d} \tau_i} = \frac{\tau_{\varepsilon,s} \sum_{k \in \mathcal{I}_d^s} s_k + \tau_{\varepsilon,l} \left(\sum_{j \in \mathcal{I}_d^l} s_j + s_d \right)}{\nu_I \tau_{\varepsilon,s} + (1 + \nu_l) \tau_{\varepsilon,l}}, \quad (3)$$

where the second equality follows from the fact that signal precision is common within each bidder type. Dealers, in turn, can share some of this order flow information with their clients. Dealer information sharing takes the form of a noisy signal about \bar{s}_d , which is the summary statistic for everything the dealer knows about the asset fundamental f . That noisy signal is $\bar{s}_{\xi d} = \bar{s}_d + \xi_d$ where $\xi_d \sim N(0, \tau_{\xi}^{-1})$ is the noise in the dealers' advice, which is iid across dealers d . Our model captures noisy dealer advice, as well as two extreme cases: perfect information-sharing between dealers and clients ($\tau_{\xi} = \infty$) and no information-sharing ($\tau_{\xi} = 0$). For now, we assume that each dealer discloses $\bar{s}_{\xi d}$ to each of his clients, but not to other dealers (Figure 1, case b). We return to inter-dealer sharing later.

The final piece of information that all agents observe is the auction-clearing price p . Of course, the agent does not know this price at the time he bids. However, the agent conditions his bid $q(p)$ on the realized auction price p . Thus, each quantity q demanded at each price p conditions on the information that would be conveyed if p were the realized price. Since p contains information about the signals that other investors received, an investor uses a signal derived from p to form his posterior beliefs about the asset payoff. Let $s(p)$ denote the unbiased signal constructed from auction-clearing price. We guess and verify that $(s(p) - f) \sim N(0, \tau_p^{-1})$, where τ_p is a measure of the informativeness of the auction-clearing settle price.

Thus the vector of signals observed by an investor j assigned to dealer $d(j)$ is $S_j = [s_j, \bar{s}_{\xi d(j)}, s(p)]$. This is the same vector for large and small investors. The only difference for the two investors is that large investors' private signals s_j are more precise and have a different (higher) covariance with price information. A dealer observes the same signals, except that he sees the exact order flows, instead of a noisy signal of them. For dealer d , $S_d = [s_d, \bar{s}_d, s(p)]$. For every agent, we use Bayes' law to update beliefs about f . Bayesian updating is complicated by correlation in the signal errors. To adjust for this

correlation, we use the following optimal linear projection formulas:

$$\mathbb{E}[f|S_j] = (1 - \beta' \mathbf{1}_m)\mu + \beta' S_j \quad \text{where} \quad (4)$$

$$\beta_j \equiv \mathbb{V}(X_j S)^{-1} \mathbf{Cov}(f, S_j) \quad (5)$$

$$\mathbb{V}[f|S_j] = \mathbb{V}(f) - \mathbf{Cov}(f, S_j)' \mathbb{V}(S_j)^{-1} \mathbf{Cov}(f, S_j) \equiv \hat{\tau}_j^{-1}, \quad (6)$$

where m is the number of signals in the vector S_j , the covariance vector is $\mathbf{Cov}(f, S_j) = \mathbf{1}_m \tau_f^{-1}$ and the signal variance-covariance $\mathbb{V}(S_j)$, is worked out in the appendix. The vector $\beta_j = [\beta_{s_j}, \beta_{\xi_j}, \beta_{p_j}]$ dictates how much weight an agents puts on his signals $[s_j, s_{\xi d(j)}, s(p)]$ in his posterior expectation. In a Kalman filtering problem, β is like the Kalman gain. In an econometric forecasting problem, β s are the OLS coefficients that multiply the independent variables to forecast the dependent variable – in this case, the payoff f .⁶ We define an equilibrium in the auction as

Equilibrium. A Nash equilibrium is

1. A menu of price-quantity pairs bid by each small investor i that solves

$$\begin{aligned} \max_{q_i(p)} \mathbb{E}[-\exp(-\rho W_i)|S_i] \\ \text{s.t. } W_i = W_{0i} + q_i(f - p). \end{aligned}$$

The optimal bid function is the inverse function: $p(q_i)$.

2. A menu of price-quantity pairs bid by each large investor that maximizes

$$\max_{q_j, p} \mathbb{E}[-\exp(-\rho_l W_j)|S_j] \quad (7)$$

$$\text{s.t. } W_j = W_{0,j} + q_j(f - p), \quad (8)$$

$$x + \sum_{i=1}^{N_I} q_i + \sum_{j=1}^{N_D} q_j + \sum_{k=1}^{N_L} q_k = 1. \quad (9)$$

The second constraint is the market clearing condition and reflects that the strategic players must choose their quantity taking into account the effect their demand has on market clearing, and, hence, the realized price.

3. A menu of price-quantity pairs bid by each dealer (dealer and large investor) that

⁶More precisely, the OLS formulas with known means, variances and covariances. The OLS additive constant α is $(\mathbf{1}_N - \beta)' \mathbf{1}_N \mu$. β is the infinite sample version of $(X'X)^{-1} X'y$. The conditional mean here is analogous to the optimal linear estimate in the OLS problem. This equivalence holds because in linear-normal systems, both OLS and Bayesian estimators are consistent.

maximizes

$$\max_{q_d, p} \mathbb{E}[-\exp(-\rho_l W_d) | S_d] \quad (10)$$

$$\text{s.t. } W_d = W_{0,d} + q_d(f - p) - \chi_0 + \chi q_d p, \quad (11)$$

$$x + \sum_{i=1}^{N_I} q_i + \sum_{j=1}^{N_D} q_j + \sum_{k=1}^{N_L} q_k = 1. \quad (12)$$

The dealer's budget constraint reflects the minimum bidding requirement faced by the dealers.

4. An auction-clearing (settle) price that equates demand and supply: $x + \sum_{i=1}^{N_I} q_i + \sum_{j=1}^{N_D} q_j + \sum_{k=1}^{N_L} q_k = 1$.

2 Solving the Model

We first solve for optimal bid schedules of large, small investors and dealers. Then, we work out the auction equilibrium and vary the amount of information being shared either between investors and dealers or between dealers. We consider three cases as illustrated in Figure 1: 1) dealers and customers share information; 2) dealers also share information with other dealers; and 3) no information is shared either with customers or between dealers.

Since all investors' posterior beliefs about f are normally distributed, we can use the properties of a log-normal random variable to evaluate the expectation of each agent's objective function. It then follows that the FOC of the small investors' problem is to bid

$$q_i(p) = \frac{\mathbb{E}[f|S_i] - p}{\rho \mathbb{V}[f|S_i]}, \quad (13)$$

which is a standard portfolio expression in an exponential-normal portfolio problem. The fact that it is an auction setting rather than a financial market doesn't change how choices are made. The novelty of the model is in how dealers' information sharing affects the conditional mean and variance of the asset payoff.

For large strategic bidders, we substitute the budget constraint in the objective function, evaluate the expectation and take the log. The strategic investor maximization problem then simplifies to $\max_{q_j, p} q_j(\mathbb{E}[f|S_j] - p) - \frac{1}{2}\rho_l q_j^2 \mathbb{V}[f|S_j]$ subject to the market clearing condition (9), where the price is *not taken as given*.⁷ The first order condition with respect

⁷In the baseline model we rule out collusion, and relax this assumption in Section 4.3.

to q_j reveals that dealers and large investors bid

$$q_j(p) = \frac{\mathbb{E}[f|S_j] - p}{\rho_l \nabla[f|S_j] + dp/dq_j}. \quad (14)$$

Importantly, this expression differs from equation (13) by the term dp/dq_j , which measures the price impact of a strategic investor bid. As the price impact increases, the dealer’s demand becomes less sensitive to his beliefs about the value of the security.

Dealers are just like large investors, except that in addition to strategic price impact, they also face minimum bidding requirements. We substitute the dealer’s budget constraint (2) and the market -clearing expression for equilibrium price (9) in the objective (1), take the expectation and the first order condition with respect to q_D , to obtain

$$q_D(p) = \frac{\mathbb{E}[f|X_D S] - p(1 - \chi)}{\rho_l \nabla[f|X_D S] + (1 - \chi)dp/dq_D}. \quad (15)$$

Note that the bidding requirement shows up like a dealer price subsidy, encouraging the dealer to purchase more of the asset. It also mitigates the effect of dealer market power by multiplying the dp/dq_D term by a number less than one.

2.1 Equilibrium auction-clearing price: 3 cases

In order to understand the effect of client information sharing, dealer information sharing and no information sharing, we solve for the equilibrium auction outcomes in each of these three cases.

The no-information-sharing world we consider is one with “Chinese walls,” where dealers cannot use client information to inform neither their own nor their clients’ purchases. In recent years, a number of financial firms have implemented such a separation of brokerage activities and transactions for their own account. Regulators have also recommended that banks establish and enforce such internal controls to address potential conflicts of interest.⁸ In our Chinese wall model, each agent sees only their own private signal s_i and the price information $s(p)$ which they can condition their bid on, but not any signal from the dealer: $S_i = [s_i, s(p)]$.

In every version of the model, adding up all investors’ and dealers’ asset demands as well as the volume of market orders x and equating them with total supply delivers the equilibrium auction price. As in most models with exponential utility (e.g. Kyle (1989)), the price turns out to be a linear function of each signal. The innovation in this model is

⁸For example, the Financial Stability Board (FSB) 2014 report on “Foreign Exchange Benchmarks”

that information sharing changes the linear weights. To the extent that signals are shared with more investors, that signal will influence the demand of more investors, and the weight on those signals in the price function will be greater. In this model, the signals are the investor's private signal s_i or s_j (for large investors), in each dealer's average order flow signal s_d , the volume of market orders x , and the signal noise in each dealer's signal to his clients ξ_d :

Result 1. *Suppose all investors bid through dealers. Consider the following three information-sharing regimes.*

1. *Dealers share information imperfectly with clients, but not with other dealers.*
2. *Dealers share information with clients and ψ other dealers.*
3. *There is no information sharing at all. Dealers cannot use client trades as information on which to condition their own bid (Chinese walls).*

In all three cases, the auction revenue is a linear function of signals s_i , market orders x , and dealer signal noise ξ_d :

$$p = A + B_I \bar{s}_I + B_L \bar{s}_{NL} + B_D \bar{s}_{ND} + Cx + D \bar{\xi}_d \quad (16)$$

where $\bar{s}_{Nz} \equiv N_z^{-1} \sum_{i=1}^{N_z} s_i$ are the average signals of individuals ($z = I$), large investors ($z = L$) and dealers ($z = D$); $\bar{\xi}_d \equiv \sum_{d=1}^{N_d} \xi_d$ is the average dealer's signal noise. The equilibrium pricing coefficients A, B_I, B_L, B_D, C and D that solve each model differ by model and are reported in appendix A.

Standard competitive market models often have simple solutions for the price coefficients. The complication here is two-fold: 1) there are large strategic agents whose demand is not linear in the coefficients of the price function and 2) shared signals are correlated with price information. Both sources of extra complexity are essential to understand how the number of dealers and their information sharing affects auction revenue.

A primary effect of information sharing in this model is higher auction revenue. The reason is that sharing information leaves all investors better informed. Investors who perceive an asset to be less risky will hold it at a lower risk premium, or at a higher price. A lower risk premium is a less negative A . We see in the solution that this risk premium ($-A$) decreases when information is shared and uncertainty is lower. While this type of effect shows up in many imperfect information asset pricing models, it offers a new perspective on how restricting sharing of information affects auction revenues. This effect is largely neglected in the policy discourse on information regulation.

With Chinese walls, when dealers can no longer use the information in their clients' orders,

the functional difference between dealers and large investors disappears. The difference between direct bidding and indirect bidding is similarly eviscerated. In other words, eliminating all information sharing effectively eliminates intermediation as well. The finding that there is no longer any meaningful distinction between a dealer and a non-dealer large investor is reflected in the fact that in the price formula, if the number of dealers and large investors is equal and the dealers do not face a minimum bidding requirement, then the coefficients on the signals of dealers s_d and the signal of large investors s_j are equal as well.

Auction Revenue Since we normalized Treasury asset supply to one, price and auction revenue in this model are the same. Our objective is to determine what the expected revenue is, what the variance of that revenue is, and how this mean and variance vary with information sharing. The unconditional expected revenue will be a linear function of the unconditional mean of the asset payoffs μ and the unconditional mean quantity of market orders \bar{x} : $A + B_{total}\mu + C\bar{x}$, where $B_{total} = B_I = B_L + B_D$. Unconditional revenue variance will be $B_I^2\tau_{\epsilon I}^{-1}/N_I + B_L^2\tau_{\epsilon L}^{-1}/N_L + B_D^2\tau_{\epsilon D}^{-1}/N_D + C^2\tau_x^{-1} + D^2\tau_\xi^{-1}/N_D$.

3 Mapping the Model to the Data

To measure the impact of information sharing on auction revenue and bidders' welfare, we calibrate the model parameters using data from two main data sources: Treasury auction results and market prices. In 2013 alone, Treasury issued nearly \$8 trillion direct obligations in the form of marketable debt as bills, notes, bonds and inflation protected securities (TIPS), in about 270 separate auctions.⁹ Our sample starts in September 2004 and ends in June 2014. To study a comparable sample and estimate yield curves, we restrict attention to 2-, 3-, 5-, 7- and 10-year notes and exclude bills, bonds and TIPS.

In each auction, competitive bids specify a quantity and a rate, or the nominal yield for note securities. Non-competitive bids specify a total amount to purchase at the market-clearing rate. Each bidder can only place a single non-competitive bid with a maximum size of \$5 million. Competitive bids can be direct or indirect. To place a direct bid, investors submit electronic bids to Treasury's Department of the Public Debt or the Federal Reserve Bank of New York. Indirect bids are placed on behalf of their clients by depository institutions (banks that accept demand deposits), or brokers and dealers, which include all institutions

⁹Treasury bills are auctioned at a discount from par, do not carry a coupon and have terms of not more than one year. Bonds and notes, instead, pay interest in the form of semi-annual coupons. The maturity of notes range between 1 and 10 years, while the term of bonds is above 10 year. For TIPS, the coupon is applied to an inflation-adjusted principal, which also determines the maturity redeemable principal. TIPS maturities range between 1 and 30 years.

registered according to Section 15C(a)(1) of the Securities Exchange Act. Though most indirect bids are placed through primary dealers, indirects bids also include those placed by the New York Fed on behalf foreign and international monetary authorities (FIMA) that hold securities in custody at the Fed. We return to these types of indirect bids below.

On the auction day, bids are received prior to the auction close. The auction clears at a uniform price, which is determined by first accepting all non-competitive bids, and then competitive bids in ascending yield or discount rate order. The rate at the auction (or stop-out rate) is then equal to the interest rate that produces the price closest to, but not above, par when evaluated at the highest accepted discount rate or yield at which bids were accepted.


We first discuss the calibration of auction participation by types of bidders using auction results data, which are made publicly available by the U.S. Treasury. For each maturity, we compute the mean share of securities allotted to primary dealers, direct and indirect bidders, after excluding amounts allotted to the Fed’s own portfolio through roll-overs of maturing securities, which are an add-on to the auction. As discussed above, the definition of indirect bidders from official auction results includes competitive bids placed by foreign officials through the NY Fed.¹⁰ Bids by foreign official investors are driven by exogenous factors such as foreign exchange strategies and their need to roll-over large quantities of reserve balances, as opposed to fluctuations in short-term market values. As a result we apply a simple calculation to reclassify these bids as part of the noise trader group (non-competitive). Although auction result data do not detail indirect bids from FIMA customers, we use information on foreign security holdings and on foreign bids from investment allotment data to reconstruct the amount of these bids at each auction.¹¹

Figure 2 plots Treasury auction participation by type of bids over time and the corresponding target moments (first and second) are reported in Table 1. As shown by the dark grey area, primary dealers bidding for their own account, are the largest bidder category at auctions accounting for about half of all security allotments. Indirect bidders, excluding estimated FIMA bids, are the second largest at 32 percent (light gray). Direct bids (medium gray) are about 8 percent and non-competitive bids, as computed above, are about 11 percent (red areas).

¹⁰FIMA customers can place *non-competitive* bids for up to \$100 million per account and \$1 billion in total. Additional bids need to be placed competitively.

¹¹For example, from Treasury International Capital (TIC) data, as of August 2014, about 6 trillion of securities are held by foreign investors, while from the Fed Board’s H4.1 release, FIMA holdings at the New York Fed are about \$3.4 trillion as of that time. Assuming that the portfolio composition and bidding strategy of FIMA and non-FIMA are similar then an estimate of FIMA’s share of competitive bids reported as indirect ones on that date is: $3.3/6 \times$ all foreign bids (from investment allotment) less FIMA non-competitive bids that are reported separately.

Figure 2: Allotted shares by bidders across all auctions. Market orders (FIMA other) are constructed from indirect FIMA bids as discussed in Section 3. Source: Treasury auction results.



We turn next to the calibration of the security fundamental value. We first note that the type of uncertainty faced by Treasury bidders is different from the risks faced by corporate bond investors. Because sovereign secondary markets are deep and liquid, Treasury investors can hedge issuer-specific risks by shorting already-issued securities. Newly issued government securities do, however, carry a liquidity premium relative to already-issued securities. Investors’ demand for specific issues is the key determinant of these liquidity differences. As a results, key underwriting risks for bidders are issue-specific rather than issuer-specific. In our model, we assume that each investor observes a signal of the issue-specific value of the newly issued security, and uses this signal to form an expectation of the value of the security after the auction.

To calibrate the first and second moments of p and f , it is important to first note that, up to rounding, the auction price clears at par. The stop-out coupon rate is, instead, uncertain and will be a function of issue-specific value as well as the term structure of interest rates at the time of the auction, which depends on factors unrelated to the auction, such as monetary policy and inflation expectations. We focus on issue-specific fundamentals, or the “on-the-run” value of the issue, for two reasons. First, an investor can easily hedge interest rate risk into the auction by shorting a portfolio of currently outstanding securities. Second, from the issuer perspective, the stop-out rate could be very low because of low interest rates, but an issue could still be “expensive” relative to the rate environment due to auction features, which is what we are after. To strip out the aggregate interest-rate

effects, we assume that the bidder enters the auction with an interest-rate-neutral portfolio, which holds one unit of the auctioned security and shorts a replicating portfolio of bonds trading in the secondary market. This portfolio is equivalent to the excess revenue on the current issue, relative to outstanding securities. Thus, price p in our model corresponds to *the auction price, minus the present value of the security’s cash-flows*, where future cash flows are discounted using a yield curve. To compute this measure, we estimate a Svensson yield curve following the implementation details of [Gürkaynak, Sack, and Wright \(2007\)](#) but using intraday price data as of 1pm, which is when the auction closes (data from Thomson Reuters TickHistory). The fundamental value f in the model corresponds to *the value of the interest-rate neutral portfolio on the date when the security is delivered to the winning bidders* (close of issue date). The issue date in our sample lags the auction date by an average of 5.5 days with a standard deviation of about 2.3 days. For example, in [Table 1](#), the average revenue from selling a new coupon-bearing security is 37.18 basis points higher than the replicating portfolio formed using outstanding securities. Thus, we calibrate the model to have this average asset payoff. This excess revenue is positive across all maturities. This is the well-known “on-the-run” premium ([Lou, Yan, and Zhang, 2013](#); [Amihud and Mendelson, 1991](#); [Krishnamurthy, 2002](#)). [Appendix D](#) details exactly how we calculate payoffs, explores other possible ways of hedging the interest rate risk, and discusses the role of the when-issued (WI) market.

Table 1: Calibration targets and model-implied values. Prices and excess revenues are all expressed in basis points.

	Data	Model
A	-17.01	-7.55
Price sensitivity to fundamental	0.97	0.91
C	124.38	73.88
Error Std. Dev.	29.72	23.12
Expected excess revenue	37.18	38.72
Volatility of excess revenue	72.64	70.81
Indirect share	0.25	0.51
Volatility of indirect share	0.09	0.73
Dealer share	0.53	0.43
Volatility of dealer share	0.14	0.19
Direct share	0.10	0.00
Volatility of direct share	0.09	0.01

We fit the parameters of the full model to aggregate moments. The full model differs from the one presented in the previous section for the inclusion of a direct bidder in the model, which, as we discuss in [Section 5](#), is a key distinguishing feature of Treasury auctions. The objective function matches a few moments from the model to their empirical counterparts:

the pricing coefficients A , B and C in equation (16), the mean and variance of the price of the interest-rate-risk-neutral portfolio p at auction, the mean and variance of the price of the portfolio f at issuance, the mean allotted share and variance of non-competitive bids x (including the FIMA trades, or market orders), the mean allotted share to primary dealers, $\sum_{d=1}^D q_d$, the mean allotted share to indirect bidders ($\sum_{i=1}^N q_i$), and the mean allotted share to direct bidders, q_L .¹² We obtain sample estimates of A , B and C by regressing the stop-out-price at each auction on a constant, the end-of-day secondary price on the issue date of the auction security (data from Bloomberg LP) and the non-competitive bids. As shown in Table 1, consistent with the model, excess revenues are positively correlated to the fundamental value on issue date (positive $B = B_I + B_L + B_D$), and it also increases with the share of securities allocated to market orders (positive C). The model moments are computed by making 100000 draws of realization of the fundamental f , all the signals in the economy s_i and the non-competitive demand x from the model, and calculating the equilibrium outcomes.

Table 2: Calibrated parameters μ , χ_0 , $\tau_f^{-\frac{1}{2}}$, $\tau_{\varepsilon,s}^{-\frac{1}{2}}$ and $\tau_{\varepsilon,l}^{-\frac{1}{2}}$ are all expressed in basis points.

μ	$\tau_f^{-\frac{1}{2}}$	$\tau_{\varepsilon,s}^{-\frac{1}{2}}$	$\tau_{\varepsilon,l}^{-\frac{1}{2}}$	$\tau_x^{-\frac{1}{2}}$	\bar{x}	ρ	ρ_L	χ_0	χ	N_S	N_L	N_D
40.8	73.5	529.9	272.3	0.06	0.12	47825	505	0.06	0.05	240	40	20

We set the level of minimum bids χ_0 to be equal to the pro-rata share of the issuance at the expected price in the baseline model, with perfect information sharing with clients and no information sharing with other dealers. This reflects the spirit of the minimum bidding requirement: dealers have an effective price concession when they bid for a larger fraction of the auction or at a higher price.

The final parameter, τ_ξ regulates how much information dealers share with clients. Without micro data, we cannot infer this value, and assume that $\tau_\xi = \infty$ (perfect sharing) in calibrating the other parameters. We explore the stability of the calibrated parameters and of the equilibrium outcomes under different assumptions on information sharing between clients and dealers and between dealers in appendix E. In reality, different dealers probably engage in different degrees of information-sharing. Instead, we show results from the whole spectrum of potential values, from $\tau_\xi = 0$ (no information sharing) to $\tau_\xi = \infty$ (perfect sharing).

Given these parameters, we solve the model by solving for the equilibrium pricing coefficients in Result 1. This amounts to solving for a fixed point in a set of up to seven

¹²We use the pricing coefficients A , B and C for calibration, but not D . The reason is that D multiplies the dealer signal noise ξ , which is not observed. Thus, D is part of the estimation residual.

non-linear equations (five for pricing coefficients and two for demand elasticities of dealers and large investors). We iterate to convergence, using the average violation of the market clearing condition (9) to ensure that we find the equilibrium pricing coefficients. The average violation of the market clearing condition at the solution does not exceed 4 basis points for the models in Result 1. For some models, we have to use multiple starting points to ensure that the maximum is a global one.

4 Results: Effects of Information Sharing

We examine two forms of information sharing: First, the case in which dealers vary the degrees of information with their clients but do not communicate with each other. Then, we hold the precision of client communication fixed and vary the number of dealers that dealers share information with. In both cases, we find that information sharing increases auction revenues as well as revenue volatility. The surprising finding is that small investors dislike, as a group, when dealers share more precise information with them, but benefit when dealers share information with each other. The intuition for this puzzling finding is that client information sharing increases information asymmetry and inhibits risk sharing, as in Hirshleifer (1971), while inter-dealer talk reduces information asymmetry and improves risk sharing.

Since the quantity of Treasury securities sold is normalized to 1, the auction price and auction revenues are the same. Therefore, in the plots that follow, we report the expected price, varying one exogenous parameter at a time. In each exercise, all parameters other than the one being varied are held at their calibrated values. The one exception is χ , the minimum bidding penalty. For simplicity, we turn that off ($\chi = \chi_0 = 0$) to start, and return to examine its effect in section 4.2.


4.1 Information Sharing and Auction Revenue

The top-left panel of Figure 3 plots expected auction revenues as a function of different levels of dealer information sharing with clients. The horizontal axis shows the precision of the dealer signal τ_ξ from zero (no information sharing) to infinity (perfect information sharing). More information sharing means that dealers reveal their information \bar{s}_d with less noise to their clients. The figure shows that moving from no sharing to perfect information sharing increases expected revenue by 1.5 basis points.


Information sharing makes investors better informed which in turn makes Treasuries less risky to investors, eliciting stronger bids and increasing auction revenues. To put this yield

Figure 3: Dealer Information Sharing. Top row: dealer information sharing with clients; bottom row: dealer information sharing with other dealers. In the top row, the horizontal axis shows the precision of the dealer signal τ_ξ from zero (no information sharing) to infinity (perfect information sharing). Expected utility is plotted as a fraction of the utility each type gets in the Chinese wall equilibrium.


(a) Sharing with clients: Expected Revenue




(b) Sharing with clients: Expected Utility



(c) Sharing with dealers: Expected Revenue



(d) Sharing with dealers: Expected Utility



effect in perspective, applying this estimated effect to annual Treasury issuance (or about \$8 trillion), the model implies that total auction revenues would increase about \$1.2 billion when going from no dealer sharing with customers to perfect sharing with them.

In unreported results, we also find that sharing information with clients also increases the variance of auction revenue by about 1.6 basis points. This higher variance arises because dealers make investors better informed. Absent any information about the future value of a security, bidders would always bid the same amount and revenue would be constant. With more precise information, bidders condition their bids on this information. When the fundamental value of the securities fluctuate, investors learn this information with a high degree of accuracy, and use this information in their bids leading to more volatile auction revenues. One parameter that is important for these quantitative results is the variance of non-competitive bids. When these bids are less predictable, auction clearing prices are less clear signals about the true value of the asset. The value of information aggregation increases, which makes dealers more valuable in expected auction revenue terms.

In terms of bidders' welfare, the top-right panel of Figure 3 shows that dealers' utility declines when they share more information. That's not surprising since they are giving up some of their information advantage. But it also shows that small and large investors' utility declines with information sharing. Information acquisition is like a prisoners' dilemma in this setting. Each investor would like more of it. But when they all get more, all are worse off. One reason is that better-informed investors bid more for the asset. By raising the price, they transfer more welfare to the issuer (Treasury).

When dealers share information among themselves, as opposed to with customers, auction revenues also increase (Figure 3, bottom-left panel). As we increase the number of dealers with which each dealer shares information with auction revenues increase by about 1.5 basis points.¹³ In doing this exercise, we hold dealer information-sharing with clients fixed by assuming that all dealers share all information with their clients. In additional analysis, we find that when prior uncertainty about the future value of the asset is high (precision τ_f is low), or the variance of non-competitive bids grow, information sharing raises revenue by more. The increased auction revenue effect is similar to that resulting from information sharing with clients. In both cases, additional information makes the average bidder for the asset less uncertain. Since dealers disclose some of their information to their clients, all investors have more precise information sets. All else equal, a reduction in risk prompts bidders to bid more for the asset.

¹³Since we assume dealers are symmetric, we need the number of dealers in an information-sharing collective to be a factor of 20, the calibrated number of dealers. Thus, we stop at 9, which implies that two groups of 10 dealers each are sharing information with each other. Any more information sharing beyond this level would be perfect inter-dealer sharing.

4.2 Client vs. Dealer Information Sharing: Utility Effects

A key insight of our model is how client and dealer information sharing differ. While both types of information sharing reduce uncertainty and increase auction revenue, client and dealer information sharing have opposite effects on investor utility. The reason for this opposite effect lies in how each type of information sharing affects information asymmetry and risk-sharing.


One might expect that when dealers share information with each other, investors are harmed. In fact, the opposite is true here. When dealers share information with each other, their information sets become more similar. That is the essence of information sharing. Since dealers' beliefs are more similar, the signals that dealers share with their clients also become more similar. With similar signals, investors' beliefs become more similar. As a result, their bids and auction allocations become more symmetric as well. When allocations are more similar, they are closer to the full-information optimal asset allocation. Because investor preferences are concave, this reduction in information and investment asymmetry improves average investor utility. The welfare effects of dealer inter- and within-dealer information is different.

In contrast, sharing information with clients increases information asymmetry. When dealers share little information with clients, the clients beliefs are not very different. They all average their priors with a heterogeneous, but imprecise, private signal. Because the private information is imprecise, beliefs mostly reflect prior information, which is common to all investors. But different dealers transmit different signals. When investors get the more precise dealer's signal, they weigh it more heavily in their beliefs, which makes investors' beliefs quite different from each other. This increase in information asymmetry makes ex-ante similar investors hold different amounts of securities ex-post. Asymmetric information pushes the asset allocation further away from the efficient diversified benchmark. The consequent reduction in risk sharing reduces utility.

Figure 4 shows how the two types of information sharing affect information asymmetry as measured by the cross-sectional dispersion of investments – the average squared deviation of each investor's auction allocation from the average for that investor's type. The fact that this dispersion increases with client information sharing and decreases with dealer information sharing illustrates how client information sharing increases information asymmetry and dealer information sharing reduces it. This is why the two types of information sharing affect risk-sharing and welfare in opposite ways.

The result that more informative signals can increase information asymmetry and thereby reduce utility is the same force that is at work in [Hirshleifer \(1971\)](#). Instead, sharing

Figure 4: Client information sharing makes allocations more heterogeneous. Dealer information sharing reduces asymmetry this dispersion. Average squared deviation is a cross-sectional measure of dispersion of Treasury holdings: $(1/N_j) \sum_{i \in \mathcal{N}_j} (q_i - \bar{q}_j)^2$, where \bar{q}_j is the average Treasury allocation of investors of type j : $\bar{q}_j = (1/N_j) \sum_{i \in \mathcal{N}_j} q_i$.



information between dealers makes agents better informed, but in a way that makes information more symmetric. This has the opposite effect on utility. Dealer information sharing is more like giving investors more public information. The results here are the opposite of the [Morris and Shin \(2002\)](#) result that in a coordination game with negative coordination externalities, public information is welfare-reducing. In our setting, the Morris-Shin assumptions are reversed: actions are substitutes instead of complements and there are positive instead of negative externalities of correlated actions because correlated investments share risk more efficiently. Thus in our setting, inter-dealer sharing making information more correlated (more public), which is welfare-improving.

Non-competitive bidder profits Since most of our non-competitive bids come for foreign monetary authorities, we do not focus our analysis on their profits. Presumably, the U.S. Treasury will not change policy to ensure that the Chinese central bank faces low prices. However, we do note that whenever information is shared, resulting in more informed competitive bidders, the profits of non-competitive bidders declines.

The role of minimum bidding requirements Primary dealers are required to be consistent, active participants in Treasury auctions. Today, primary dealers are expected to bid at all auctions an amount equal to the pro-rata share of the offered amount, with bids that are “reasonable” compared to the market. The inclusion of minimum bidding penalties in the model is therefore realistic, but also helps to calibrate the model in a sensible way. Without them, it would be hard to explain why dealers bid for so much of the auction.

However, removing low bid penalties does not change our main findings. In most cases, the penalties raise revenue, since dealers are incentivized to bid more aggressively. But bidding requirements leave the effect of client information-sharing on revenue and utility unchanged. One difference is that, with bidding requirements, dealer information sharing has a non-monotonic effect on revenue. The reason is that bidding requirements make dealers less responsive to changes in price. Therefore prices have to move more to induce dealers to bid more or less to clear the market. These large swings in price make prices more informative about dealer’s signals. That leaves less scope for precision improvements through dealer information-sharing.

4.3 What if Information Sharing Enabled Collusion?

One of the reasons that information-sharing has raised concerns is that dealers who share information can also collude. Many textbook analyses show the economic losses associated with collusion. We do not repeat those arguments here. Instead, we look at how information sharing interacts with the costs of collusion.


Suppose that every time dealers shared information with each other, that group of dealers colluded, meaning that they bid as one dealer in order to amplify their price impact. How would this collusion and information sharing jointly affect auction revenue? It turns out that the answer depends on how much information is shared with clients. Figure 5 shows that when dealers pass most of their information on to their clients, sharing information and colluding with other dealers increases revenue. Collusion, by itself, is of course revenue reducing. But the joint effect of better informed bidders and colluding dealers is a net positive for revenue. The problem arises when dealers talk, collude and don’t inform their clients: that reduces revenue. Notice, however, that the proposed “remedy” of imposing Chinese walls reduces revenue by more than collusion.

5 Mixed Auctions: Choosing Direct or Indirect Bidding

A key distinguishing feature of U.S. Treasury auctions is that they are mixed auctions: Any investor can either place an intermediated bid through a primary dealer, or bid directly.¹⁴ The option to bid directly is important because it amplifies the effect of low signals on auction revenue.

¹⁴While direct bidding has been historically allowed since 1992, electronic bidding systems and the elimination of deposit requirements for all bidders have facilitated direct bids. Direct bidding has grown from 2 percent of all bids in 2003 to about 10 percent in 2014. While auction results do not disclose the number of direct bidders, public remarks of Treasury officials suggest there were about 1200 direct bidders in 2001, and 825 in 2004.

Figure 5: Collusion reduces revenue when client information sharing is too low. Figure plots average equilibrium auction revenue against the number of other dealers that share information. We assume here that when dealers share information, they also bid as one. These results differ from previous ones because here, varying information-sharing along the x-axis also varies the extent of collusion.



5.1 A Model with Intermediation Choice

Consider one investor choosing between bidding directly or indirectly through an intermediary. Without loss of generality, we assume that this choice is made by one of the large investors bidding through dealer 1. The large investor’s choice to bid directly or indirectly affects the information structure of that investor, its dealer, other investors bidding with that same dealer, and the information content of the price $s(p)$.

When the large investor chooses to bid directly on his own behalf, he observes only his own signal and the price information: $X_L S = [s_L, s(p)]$. His dealer’s signal is the average of the first ν_I investors’, the first $\nu_I - 1$ large investors’ and the dealer’s signal:

$$\bar{s}_1 = \frac{\tau_{\epsilon,s} \sum_{k \in \mathcal{I}_1^s} s_k + \tau_{\epsilon,l} \left(\sum_{j \in \mathcal{I}_1^l} s_j + s_d \right)}{\nu_I \tau_{\epsilon,s} + \nu_I \tau_{\epsilon,l}}$$

As in the previous model, investor i who bids through intermediary d observes signals $X_i S = [s_i, \bar{s}_d, s(p)]$.

Solution: Auction Outcomes Solving this model introduces a technical challenge. The decision to bid directly or indirectly becomes in itself a signal. We assume that the dealer who would intermediate this trade observes the large investor’s bidding decision and transmits this information to clients, with noise. If the large investor bids through the dealer, the dealer can infer exactly what the large investor’s signal is. But if the large investor bids directly, the dealer only knows that the investor’s signal is above a threshold.

The information that has been revealed is that a normal variable (the large investor’s signal) lies in two disjointed truncated regions of the distribution. This is problematic because doing Bayesian updating of beliefs with truncated normals would require involved simulation methods. Embedding that updating problem in the non-linear fixed point problem we already have would render the model intractable.

We circumvent this problem by constructing an approximating normal signal. Through simulation, we first estimate the mean and variance of the large investor’s signal, conditional on choosing direct bidding. Then, whenever the large investor chooses to bid directly, we allow the dealer who would have intermediated that trade to make inference from the direct bidding decision, by observing a *normally distributed* signal with the same mean and variance as the true information. This signal is included in the precision-weighted average signal of dealer 1, \bar{s}_1 . (See appendix B for details.)

If the large investor bids through the dealer, the problem and the solution are the same as in the previous section. With direct bidding, the auction price (revenue) is a linear function of the dealer-level average individual investor signals, \bar{s}_d (where \bar{s}_1 includes the information inferred from the large investor’s decision to bid directly), the signal of the large direct bidder, s_L , and of market orders x .

Result 2. *With N_D dealers and 1 large investor who bids directly, the auction revenue is $p = A + B_{d1}\bar{s}_1 + B_{d\neq 1}/(N_D - 1) \sum_{d=2}^{N_D} \bar{s}_d + Cx + B_L s_L$, where the coefficient formulas are in reported appendix B.*

5.2 Understanding Intermediation Choice: Why Is Bad News Revealed?

When an investor bids directly, no one observes their order flow and their signal remains private. When they bid through an intermediary, they reveal their signal realization to the dealer, but also learn from the dealer’s signal. An investor whose signal indicates a high future value for the security expects to take a large position, which will make his utility more sensitive to the auction-clearing price. Sharing his good news with others will increase the clearing price and negatively affect his expected utility. Thus, an investor with good news prefers not to share his information and bids directly. Conversely, when the news is bad, the investor expects to take a small position in the auction making his utility not as sensitive to the clearing price. With a low signal, the investor is less concerned about sharing his signal but still benefits from learning new information from other investors. Thus, low-signal investors are more likely to bid indirectly through the dealer. When negative signals are shared, they affect bids of many investors and their price impact is amplified. Dealer information sharing makes the accelerator stronger by making intermediation more costly

and direct bidding more likely.

To solve for the large investor's choice of whether to bid directly or indirectly $l \in \{Ld, Li\}$, we compute expected utility conditional on signals and a realized price. When the investor chooses whether to invest through a dealer the only signal that he has seen is his private signal s_i . Thus the intermediation choice maximizes expected utility with an additional expectation over the information that the large investor has not yet observed. Computing the expectation over possible price realizations and dealer signals, we find that expected utility is

$$EU(l) = -\exp(\rho_L W_L)(1 + 2\theta_l \Delta \mathbb{V}_l)^{-\frac{1}{2}} \exp\left(-\frac{\mu_{rl}^2}{\theta_l^{-1} + 2\Delta \mathbb{V}_l}\right). \quad (17)$$

The intermediation decision affects utility in three ways: through the expected profit per unit allotted μ_{rl} , the sensitivity of demand to expected profit θ_l , and through the ex-ante variance of expected profit $\Delta \mathbb{V}_l$. These three terms are:¹⁵

$$\mu_{rl} \equiv \mathbb{E}\{\mathbb{E}[f|X_l S] - p|s_i\}, \quad (18)$$

$$\theta_l \equiv \rho_L[\rho_L \mathbb{V}[f|S_L] + dp/dq_L]^{-1} \left(1 - \frac{1}{2}\rho_L[\rho_L \mathbb{V}[f|S_L] + dp/dq_L]^{-1} \mathbb{V}[f|S_L]\right), \quad (19)$$

$$\Delta \mathbb{V}_l \equiv \mathbb{V}\{\mathbb{E}[f|X_l S] - p|s_i\} = \mathbb{V}[f - p|s_i] - \mathbb{V}[f|X_l S]. \quad (20)$$


The first term μ_{rl} embodies the main cost of intermediation: It reveals one's private information s_i to others. This effect shows up as a reduction in μ_{rl} , the ex-ante expectation profit per share, after all signals are observed. Information sharing reduces μ_{rl} for two reasons. First, since many investors all condition their bids on the shared information $X_{d1} S$, the expectation conditional on that information, $E[f|X_{d1} S]$, has a large effect (closer to 1) on the auction-clearing price. Thus the difference $E[f|X_l S] - p$ is closer to zero with intermediation. Second, improving the precision of other investors' information lowers their risk, raises the expected price p , which in turn, lowers μ_{rl} (see eq (18)).

Why does the asymmetry appear? In equation (17), a decrease in μ_{rl}^2 decreases expected utility because $\theta_l > 0$ (see appendix). In principle, a very negative signal could make μ_{rl} a large negative number, which would also trigger direct bidding. However, on average μ_{rl} is positive. That positive mean reflects the positive risk premium. So while it is possible that a very negative signal triggers direct bidding, it is highly unlikely.

The second term θ_l captures the main advantage of intermediation: Dealers give their clients an extra signal, which makes them better informed. Better information allows the large investor to make better bids, increasing expected utility. In appendix A, we

¹⁵See the appendix for derivations of the following three equations and support for the analysis that follows.

Figure 6: Mixed auctions: Probability of bidding directly, conditional on signal realization above (solid line) or below (dashed line) its mean.



show that θ_l is positive and strictly decreasing in the posterior variance of the asset payoff $\mathbb{V}[f|X_l S]$. Thus, intermediation improves the investor’s information, which decreases variance $\mathbb{V}[f|X_l S]$, increases θ , and (holding all other terms equal) increases expected utility.

The third effect of intermediation, which operates through ex-ante variance $\Delta\mathbb{V}_l$, is ambiguous and turns out to be quantitatively unimportant.¹⁶

A zero-profit signal (bad news) is always shared. Note from equation (17) that as $\mu_{rl} \rightarrow 0$, the first effect disappears and an increase in the ex-ante variance of the profit will unambiguously increase expected utility. The reason for this is that the strength of the second effect depends on the mean of the expected profit, μ_{rl} . When $\mu_{rl} \neq 0$ the increase in the ex-ante variance $\Delta\mathbb{V}_l$ increases the probability that the expected profit $\mathbb{E}[f|X_l S] - p$ will be close to zero as well as increasing the probability of large observations. So intuitively the gains from the increase in ex-ante variance are larger when μ_{rl} is closer to zero. We use these three effects to understand the intermediation choice results below.

5.3 Intermediation Choice and the Financial Accelerator

This asymmetry, whereby bad news is shared and good news is not, is a new channel through which intermediation can amplify shocks. Figure 6 shows that when signals about

¹⁶When the large investor trades through a dealer, his uncertainty $\mathbb{V}[f|X_l S]$ declines. From equation (20) we can see that this increases the ex-ante variance of the expected profit $\Delta\mathbb{V}_l$. This is because more information makes the investor’s beliefs change more, which means a higher ex-ante variance. This change in $\Delta\mathbb{V}_l$ has two opposing effects on expected utility. First, an increase in $\Delta\mathbb{V}_l$ increases the exponential term in equation (17), which decreases $EU(l)$. This effect arises because the large investor is risk averse and higher $\Delta\mathbb{V}_l$ corresponds to more risk in continuation utility. The second effect is that an increase in $\Delta\mathbb{V}_l$ reduces $(1 + 2\theta_L \Delta\mathbb{V}_l)^{-\frac{1}{2}}$, which increases $EU(l)$. The intuition for this is that when the variance of the expected profit is larger, there are more realizations with large magnitude (more weight in the tails of the distribution). Since these are the states that generate high profit, this effect increases expected utility.

the value of a financial asset are negative (dashed line), this information is more likely to be shared with a dealer and his clients (lower probability of direct bidding). Positive signals (solid line) are less likely to be shared because an investor who receives a positive signal then expects to take a large portfolio position in the asset and faces a high expected cost from sharing his information. But sharing a bad signal places that signal in the information set of many more investors and leads to a large number of investors to demand less of the asset. Thus, bad signals may affect the demand of more investors than good signals do and have a larger effect on asset prices. When bad news is amplified and good news is not, revenue is negatively skewed, a prediction we test in the next section.

Moving right along the horizontal axis of Figure 6 represents an increase in the precision of the information dealers share with their clients. Increasing this precision strengthens the asymmetry. An increase of 250 bps, which is equivalent to doubling the precision of a large investor’s signal, loses a dealer 20% of his clients with positive information and only 2% of clients with negative information, for an average loss of just over 10% of the dealer’s clientele.

6 Testing Two Model Predictions

So far, we used the model to make qualitative and quantitative predictions about the effects of various information-sharing regimes on treasury auction revenues and investor utilities. But we have not yet shown that the data supports our interpretation of how and what information is shared. Therefore, our final section compares two central predictions of the model to the data.

Testable Prediction 1: Informed Traders’ Demand Forecasts Profits. An essential feature of the model is that a subset of agents have private information about the future resale value of Treasury securities. If they did not, then observing order flow would not be useful for dealers. The hallmark of informed trading is that such trades, as opposed to uninformed ones, forecast profits. An uninformed agent cannot systematically buy more securities when profits ($f - p$) are high and sell when they are low. Absent information about the difference between the fundamental and the auction price $f - p$, such an investment strategy would not be a measurable one. In other words, $\mathbf{Cov}(q_i, f - p) > 0$ is evidence of informed trading.

The data counterpart to profit $f - p$ is post-auction appreciation, which is the difference between the resale price of the asset in the secondary market, minus the price paid at auction. Since we assumed that all competitive bidders have private information and all

Table 3: Regression of $f - p$ on competitive bidders' auction share. The dependent variable is difference between the value of the interest-rate neutral portfolio on issue date (f) and on auction date (p) expressed in dollar units. Robust standard errors reported in square brackets. *** significant at 1%, ** significant at 5%, *significant at 10%.


	(1)	(2)	(3)
Competv. Share		0.0115*** [0.0031]	0.0147*** [0.0043]
Const	0.0003** [0.0001]	0.0017*** [0.0005]	0.0018*** [0.0005]
Adj. R2	0.00	0.05	0.23
Obs.	494	494	494
Month FEs?	No	No	Yes
Tenor FEs?	No	No	Yes

non-competitive bidders are uninformed, the data counterpart to q_i of an informed trader is the share of the auction awarded to competitive bidders. We use share of the auction, rather than face value because the size of auctions varies and this introduces noise in our regression.

Corollary 1. *High competitive share predicts high post-auction appreciation: $\partial E[f - p]/\partial \bar{q} > 0$, where $\bar{q} \equiv \int q_i di$ is the share of the auction awarded to competitive bidders.*

We test this prediction, using the auction data. Table 3 reports estimates of a regression of the price appreciation of the hedged portfolio from the time of the auction close to the issue date, or $f - p$. As shown in the first column of the table, the value of the newly issued security appreciates on average by about 3 basis points between the auction date and the issue date (column 1), consistent with the findings of [Lou, Yan, and Zhang \(2013\)](#). As shown in the second column, this appreciation is higher the higher the share of competitive bids into the auction (column 2) with a highly statistically significant coefficient (t-stat = 3.5, column 2). This empirical result is robust to the inclusion of month and tenor fixed effects (column 3). This effect is consistent with non-competitive being noise traders and competitive bidders being informed. It is not a mechanical result from high demand. When informed traders demand is high, the price is lower on average, relative to the payoff. It is that low price relative to fundamental value that induces informed investors to buy more. We run analogous regressions using data generated from the model and find similar results. This result does not prove that primary dealers aggregate this information. But it does support the notion that client order flow contains some private information to be aggregated.

Figure 7: Distribution of post-auction price appreciation $f - p$: Model and Data. Post-auction appreciation is the difference between the value of the interest-rate neutral portfolio on issue date (f) and on auction date (p)



Testable Prediction 2: Positive Skewness in Post-Auction Appreciation The financial accelerator effect from the intermediation choice shows up as a distribution of auction-clearing prices that has unconditional negative skewness. Negative skewness in the price p translates into positive skewness in the post-auction appreciation $f - p$. Table 4 reveals that the unconditional skewness in p is $-.75\%$. This translates into positive skewness in post-auction appreciation of 5.7% . When shocks are good, they have a moderate effect on the asset price and the auction revenue. But with a bad realization of the asset’s value, large investors observe negative signals. These investors choose to share their low signals with primary dealers, which in turn lowers the demands of other investors and has a significantly negative effect on auction revenues. Depressed auction revenue corresponds to high post-auction appreciation.

The sign of the skewness prediction is consistent with the empirical distribution of post-auction appreciation $f - p$, (see histogram in Figure 7). However the magnitude of skewness in the data is much stronger. Our model’s potential to generate skewness is limited by the fact that we only allow one bidder the choice of bidding directly or indirectly. One agent alone can only generate limited skewness in aggregate revenue. If we could compute a model with many agents making an intermediation choice, this skewness would likely be amplified.

7 Conclusions

Recent instances of market abuse involving sharing of confidential client information has led to calls to restrict the use of order flow information by financial intermediaries. While the need for regulation and sanctions may be evident in the case of collusive behaviour, in a setting in which all agents are informed about how information is shared, gains and losses of information sharing are not as apparent. Using data from U.S. Treasury auctions, we estimate a structural auction model to quantify the costs and benefits of information sharing both between dealers and between dealers and customers.

We find that, information sharing raises auction revenues, as bidders are better informed. Investors' welfare depends on how information is shared. Surprisingly, we find that investors are worse off when dealers share more information with them, but are better off when the dealers share information among themselves. The model analysis shows that client information sharing is like private information, which makes beliefs more different from each other, while inter-dealer talk is like public information, which makes beliefs more similar. Once we understand that analogy, the first finding that sharing with clients reduces welfare looks similar to a [Hirshleifer \(1971\)](#) effect. The second finding shows how Hirschleifer's effect can be reversed when information-sharing makes information sets more common. We also study the choice of investors to bid directly or through dealers, as well as the effect of minimum bidding requirements on primary dealers, which are an essential part of Treasury auctions.

While the paper uses the model to study the role of information in Treasury auctions, an alternative interpretation of the model sheds light on a related policy question: What is the optimal number of primary dealers? The number of primary dealers has varied over time. In 1960, there were 18 primary dealers. Amid the rapid rise in federal debt and interest rate volatility of the 1970's, the number of primary dealers rose to 46 in the mid-1980s. Subsequently, the population of primary dealers dwindled to about 22 today. The experiment in which dealers collude is equivalent to combining pairs of dealers, with half the resulting number of dealers. When we reinterpret the collusion results as reducing the number of dealers, we find that restricting the number of dealers improves revenue, but only if the information sharing with clients is sufficiently high.

The common theme throughout the paper is a reversal of the common wisdom about dealers as underwriters. The prevailing thinking about underwriters is that they lower auction revenue, but also revenue risk (for example, [Ritter and Welch, 2002](#)). In the information model that we present, we find the exact opposite: when investors bid through dealers, both mean and variance of auction revenue increase. The stark difference in these

predictions highlights how policy prescriptions may be heavily dependent on the exact role of intermediation in a given market. While many intermediaries perform roles other than information aggregation, this role is a key one in Treasury auctions and is likely to be present in some form in other markets as well. The unique features of Treasury auctions makes them a useful laboratory to isolate, investigate and quantify this new facet of intermediation.

References

- AMIHUD, Y., AND H. MENDELSON (1991): “Liquidity, maturity, and the yields on US Treasury securities,” *The Journal of Finance*, 46(4), 1411–1425.
- ARMANTIER, O., AND E. SBAÏ (2006): “Estimation and comparison of Treasury auction formats when bidders are asymmetric,” *Journal of Applied Econometrics*, 21(6), 745–779.
- BABUS, A., AND C. PARLATORE (2015): “Strategic Fragmented Markets,” Working paper, New York University.
- BACK, K., AND J. F. ZENDER (1993): “Auctions of divisible goods: on the rationale for the treasury experiment,” *Review of Financial Studies*, 6(4), 733–764.
- BIKHCHANDANI, S., AND C.-F. HUANG (1989): “Auctions with resale markets: An exploratory model of Treasury bill markets,” *Review of Financial Studies*, 2(3), 311–339.
- (1993): “The economics of Treasury securities markets,” *The Journal of Economic Perspectives*, pp. 117–134.
- BRADY, N., R. BREEDEN, AND A. GREENSPAN (1992): “Joint report on the government securities market,” *US Government Printing Office*.
- BRUNNERMEIER, M. K., AND Y. SANNIKOV (2014): “A macroeconomic model with a financial sector,” *The American Economic Review*, 104(2), 379–421.
- CHARI, V. V., AND R. J. WEBER (1992): “How the U.S. Treasury Should Auction Its Debt,” *Federal Reserve Bank of Minneapolis*, 16(4), 1–8.
- DUFFIE, D. (1996): “Special repo rates,” *The Journal of Finance*, 51(2), 493–526.
- EASLEY, D., N. KIEFER, M. O’HARA, AND J. PAPERMAN (1996): “Liquidity, Information and Infrequently Traded Stocks,” *The Journal of Finance*, 51(4), 1405–1436.
- FLECKENSTEIN, M., F. A. LONGSTAFF, AND H. LUSTIG (2014): “The TIPS-Treasury Bond Puzzle,” *the Journal of Finance*, 69(5), 2151–2197.
- FLEMING, M. J., AND K. D. GARBADE (2007): “Dealer behavior in the specials market for US Treasury securities,” *Journal of Financial Intermediation*, 16(2), 204–228.
- GÜRKAYNAK, R. S., B. SACK, AND J. H. WRIGHT (2007): “The US Treasury yield curve: 1961 to the present,” *Journal of Monetary Economics*, 54(8), 2291–2304.
- HE, Z., AND A. KRISHNAMURTHY (2013): “Intermediary Asset Pricing,” *American Economic Review*, 103(2), 732–770.
- HE, Z., A. KRISHNAMURTHY, AND K. MILBRADT (2016): “What Makes US Government Bonds Safe Assets?,” *American Economic Review*, 106(5), 519–523.
- HIRSHLEIFER, D. (1971): “The private and social value of information and the reward of inventive activity,” *American Economic Review*, 61, 561–574.
- HORTAÇSU, A., AND J. KASTL (2012): “Valuing Dealers’ Informational Advantage: A Study of Canadian Treasury Auctions,” *Econometrica*, 80(6), 2511–2542.

- HORTAÇSU, A., AND D. MCADAMS (2010): “Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market,” *Journal of Political Economy*, 118(5), 833–865.
- JORDAN, B. D., AND S. D. JORDAN (1997): “Special repo rates: An empirical analysis,” *The Journal of Finance*, 52(5), 2051–2072.
- KRISHNAMURTHY, A. (2002): “The bond/old-bond spread,” *Journal of Financial Economics*, 66(2), 463–506.
- KYLE, A. S. (1985): “Continuous auctions and insider trading,” *Econometrica: Journal of the Econometric Society*, pp. 1315–1335.
- (1989): “Informed Speculation with Imperfect Competition,” *Review of Economic Studies*, 56(3), 317–355.
- LOU, D., H. YAN, AND J. ZHANG (2013): “Anticipated and Repeated Shocks in Liquid Markets,” *Review of Financial Studies*, 26(8), 1891–1912.
- MALVEY, P. F., AND C. M. ARCHIBALD (1998): “Uniform-price auctions: Update of the Treasury experience,” *US Treasury*.
- MALVEY, P. F., C. M. ARCHIBALD, AND S. T. FLYNN (1995): “Uniform-Price Auctions: Evaluation of the Treasury Experience,” *Office of Market Finance, US Treasury*.
- MEDRANO, L., AND X. VIVES (2004): “Regulating Insider Trading when Investment Matters,” *Review of Finance*, 8, 199–277.
- MORRIS, S., AND H. S. SHIN (2002): “Social value of public information,” *The American Economic Review*, 92(5), 1521–1534.
- NYBORG, K. G., AND S. M. SUNDARESAN (1996): “Discriminatory versus uniform Treasury auctions: Evidence from when-issued transactions,” *Journal of Financial Economics*, 42, 63–104.
- RITTER, J. R., AND I. WELCH (2002): “A Review of IPO Activity, Pricing, and Allocations,” *The Journal of Finance*, 57(4), 1795–1828.
- ROOSA, R. V. (1956): *Federal Reserve Operations in the Money and Government Securities Markets*, vol. 332. Federal Reserve Bank of New York.
- SONG, Z., AND H. ZHU (2016): “QE Auctions of Treasury Bonds,” Discussion paper, MIT working paper.
- SVENSSON, L. E. (1994): “Estimating and interpreting forward interest rates: Sweden 1992-1994,” Discussion paper, National Bureau of Economic Research.
- WILSON, R. (1979): “Auctions of shares,” *Quarterly Journal of Economics*, 93(4), 675–689.

A Appendix: Derivations and Proofs

A.1 Result 1

This result has three cases. We prove each separately, reasoning through case 1 in the most detail and then pointing out similar steps that arise in cases 2 and 3. Each proof is divided into two parts: The first part takes the linear equilibrium price equation as given and produces beliefs and asset demands for the case in question. The second part imposes market clearing, uses coefficient-matching to solve for the price equation coefficients, and in doing so shows that the linear price hypothesis is an equilibrium outcome.

A.1.1 Case 1: Dealer-Client Sharing, No Dealer-Dealer Sharing

Information Structure In this case, dealers share information about average order flow with their clients by means of a noisy signal. First, recall the following definitions:

- $N_I, N_L, N_D, N \equiv$ The number of small, large, dealer, and total investors, respectively.
- $\bar{s}_d \equiv$ The dealer's unbiased estimate of average order flow, constructed as a precision-weighted average of her client's orders.

First, recall that the information set for non-dealer investor i is a three-dimensional object: $S_i = [s_i, \bar{s}_d + \xi_i, s(p)]$, where the first entry is i 's private signal, the second is that which i receives from his dealer, and the third is the information i derives from the price.

These three signals load on the following vector of orthogonal shocks:

$$Z = [\epsilon_1, \dots, \epsilon_N, \xi_i, \dots, \xi_{N_D}, x]' \quad (21)$$

The variance matrix for this vector is:

$$\mathbb{V}(Z) = \text{diag}([\tau_{e,s}^{-1} \mathbf{1}_{N_I}, \tau_{e,l}^{-1} \mathbf{1}_{N_L+N_D}, \tau_{\xi}^{-1} \mathbf{1}_{N_D}, \tau_x^{-1}]) \quad (22)$$

Next, we construct the loading matrix Π_i for investor i , which maps shocks Z to signals $S_i = f + \Pi_i Z$. This requires the construction of some auxiliary objects. Let ϕ_i denote the $1 \times N$ vector of zeros with a 1 in the i -th position (i.e., an investor identifier), and $\psi_d(i)$ the $1 \times N_D$ vector of zeros with a 1 in the $d(i)$ -th position (i.e., a dealer identifier).

It follows that the first row of the loading matrix is simply $[\phi_i, 0 \cdot \mathbf{1}_{N_D}, 0]$, as this row simply identifies the investor's private shock. The next row requires the construction of the precision-weighted average the dealer uses to construct \bar{s}_d . Let $t(i)$ be a *type operator for precision*, which returns τ_s if the investor indexed by i is small, and τ_l otherwise (recall that precisions are identical within types). Let $\mathbb{1}_i(j)$ be an indicator variable for whether i and j bid through the same dealer, and $\|\mathbf{v}\|$ be the l_1 -norm of a vector. Then, ω_i is given as follows, written as a list comprehension over $1, 2, \dots, N$.

$$\omega_i = [t(j) \mathbb{1}_i(j) \text{ for } j \text{ in } 1, 2, \dots, N] \quad (23)$$

The second row then requires identification of the signal and noise that i receives from his dealer: i.e., $[\frac{\omega_i}{|\omega_i|}, \psi_d(i), 0]$.

The third row requires us to extract relevant information from the supply shock. This first requires rewriting our price assumption such that it represents our unbiased signal about asset fundamentals:

$$p = A + B_I \bar{s}_I + B_L \bar{s}_L + B_D \bar{s}_d + Cx + D\bar{\xi}_d \implies \quad (24)$$

$$s(p) = \frac{p - A - Cx}{B_I + B_L + B_D} \quad (25)$$

Writing $\tilde{B} = B_I + B_L + B_D$, we can express this as $s(p) = \frac{B_I}{\tilde{B}} \cdot \frac{\sum s_i}{N_I} + \frac{B_L}{\tilde{B}} \cdot \frac{\sum s_j}{N_L} + \frac{B_D}{\tilde{B}} \cdot \frac{\sum x_d}{N_D} + \frac{C}{\tilde{B}}(x - \bar{x})$, which means that the third row here is $[\frac{B_I}{\tilde{B}} \cdot \mathbf{1}_{N_I}, \frac{B_L}{\tilde{B}} \cdot \mathbf{1}_{N_L}, \frac{B_D}{\tilde{B}} \cdot \mathbf{1}_{N_D}, 0 \cdot \mathbf{1}_{N_D}, \frac{C}{\tilde{B}}]$, assuming that the indexes run first through small investors, than to large ones, than to dealers.

Therefore, we can write:

$$\Pi_i = \begin{bmatrix} \phi_i & 0 \cdot \mathbf{1}_{N_D} & 0 \\ \frac{\omega_i}{|\omega_i|} & \psi_d(i) & 0 \\ \frac{B_I}{\tilde{B}} \cdot \mathbf{1}_{N_I}, \frac{B_L}{\tilde{B}} \cdot \mathbf{1}_{N_L}, \frac{B_D}{\tilde{B}} \cdot \mathbf{1}_{N_D} & 0 \cdot \mathbf{1}_{N_D} & \frac{C}{\tilde{B}} \end{bmatrix} \quad (26)$$

With this, we can compute the Bayesian updating weights mentioned earlier in the paper, which came from the projection formulas we used to excise shock covariance from the problem. These are the entries $\beta_{si}, \beta_{\bar{s}i}, \beta_{pi}$ of the vector given by:

$$\beta_i = \mathbb{V}(S_i)^{-1} \mathbf{1}_3 \tau_f^{-1} \quad (27)$$

Using these weights, and equations (4) and (6), we can express the expectation and variance of the posterior beliefs $f|S_i$:

$$\mathbb{E}(f|S_i) = (1 - \beta_{\bar{s}i} - \beta_{pi} - \beta_{si})\mu + \beta_{si}s_i + \beta_{\bar{s}i}\bar{s}_{\xi d(i)} + \beta_{pi} \frac{p - A - Cx}{\tilde{B}} \quad (28)$$

$$\mathbb{V}(f|S_i) = \tau_f^{-1} - \tau_f^{-1}(\beta_{si} + \beta_{\bar{s}i} + \beta_{pi}) \quad (29)$$

Recall that $\bar{s}_{\xi d(i)}$ is the realized signal that a dealer shares with her clients.

All that remains is to do the same thing for dealers, and the information structure will be explicitly solved. Here, dealers are essentially similar to large investors, except that they observe \bar{s}_d noiselessly (i.e., $\xi = 0$).

We can then produce an signal vector and shock weighting matrix:

$$S_d = [\bar{s}_d, s(p)] \quad (30)$$

$$\Pi_d = \begin{bmatrix} \frac{B_I}{B} \cdot \mathbf{1}_{N_I}, \frac{B_L}{B} \cdot \frac{\omega_i}{|\omega_i|} \cdot \mathbf{1}_{N_L}, \frac{B_D}{B} \cdot \mathbf{1}_{N_D} & 0 & 0 \\ 0 \cdot \mathbf{1}_{N_D} & \frac{C}{B} \end{bmatrix} \quad (31)$$

From these loadings, we can produce Bayesian updating weights exactly as we did above, noting that here we'd have a two-dimensional object as opposed to a three-dimensional one.¹⁷

These determine the posterior expectations and variances of all agents in the model whose behavior is contingent on information, we've produced the applicable information structure.

Coefficient Matching For large investors, let the sensitivity of asset demand to changes in the expected per-share profit be denoted $M_L = \rho_Z \nabla(f|S_Z) + \frac{dp}{dq_L}$.

The asset demand for each small investor is given by the first-order condition solved in (13), i.e.:

$$q_i(p) = \frac{\mathbb{E}(f|S_i) - p}{\rho_I \nabla(f|S_i)} \quad (32)$$

The total demand for each large investor is written analogously in (39):

$$q_j(p) = \frac{\mathbb{E}(f|S_j) - p}{M_L} \quad (33)$$

The dealers' demand function looks similar, except that dealers are subject to a sharper budget constraint, in virtue of the shadow costs associated with the minimum bid penalty. Therefore, we can write the demand function as given in (15):

$$q_d(p) = \frac{\mathbb{E}(f|S_D) - p(1 - \chi)}{\rho_I \nabla(f|S_D) + (1 - \chi) \frac{dp}{dq_L}} \quad (34)$$

The next step is to determine $\frac{dp}{dq_L}$, which we do by applying the market-clearing condition to the demands of all agents save one large investor:

$$x + \sum_{i=1}^{N_I} q_i(p) + \sum_{j=1}^{N_L-1} q_j(p) + \sum_{d=1}^{N_D} q_d(p) + q_L = 1 \quad (35)$$

Or:

¹⁷This is partly a matter of convenience: separating the dealer's own signals from her clients' would restore a three-dimensional structure.

$$x + \rho_I^{-1} \hat{\tau}_I \sum_{i=1}^{N_I} (\mathbb{E}(f|S_i) - p) + M_L^{-1} \sum_{j=1}^{N_L-1} (\mathbb{E}(f|S_j) - p) + M_D^{-1} \sum_{d=1}^{N_D} (\mathbb{E}(f|S_d) - p(1 - \chi)) + q_L = 1 \quad (36)$$

Next, we substitute in for the conditional expectations in terms of our Bayesian updating weights, and use the implicit function theorem to calculate $\frac{dp}{dq_L}$, which we need to determine M_L .

$$\begin{aligned} 1 &= x + \frac{\hat{\tau}_I}{\rho_I} \sum_{i=1}^{N_I} \left((1 - \beta_{si} - \beta_{\bar{s}i} - \beta_{pi})\mu + \beta_{si}s_i + \beta_{\bar{s}i}\bar{s}_{\xi d(i)} + \beta_{pi} \frac{p - A - C\bar{x}}{\tilde{B}} - p \right) \\ &+ M_L^{-1} \sum_{j=1}^{N_L-1} \left((1 - \beta_{sj} - \beta_{\bar{s}j} - \beta_{pj})\mu + \beta_{sj}s_j + \beta_{\bar{s}j}\bar{s}_{\xi j} + \beta_{pj} \frac{p - A - C\bar{x}}{\tilde{B}} - p \right) \\ &+ M_D^{-1} \sum_{d=1}^{N_D} \left((1 - \beta_{\bar{s}d} - \beta_{pd})\mu + \beta_{\bar{s}d}\bar{s}_d + \beta_{pd} \frac{p - A - C\bar{x}}{\tilde{B}} - p(1 - \chi) \right) \\ &+ q_L \end{aligned} \quad (37)$$

Taking the derivative with respect to q_L , we obtain:

$$0 = \left(N_I \hat{\tau}_I \rho_I^{-1} \frac{\beta_{pi} - \tilde{B}}{\tilde{B}} + (N_L - 1) M_L^{-1} \frac{\beta_{pL} - \tilde{B}}{\tilde{B}} + N_D M_D^{-1} \frac{\beta_{pd} - \tilde{B}(1 - \chi)}{\tilde{B}} \right) \frac{dp}{dq_L} + 1$$

At this point, we can solve for $\frac{dp}{dq_L}$, and then express M_L^{-1} implicitly, as follows:

$$M_L = \rho_L \hat{\tau}_L^{-1} - \left(N_I \hat{\tau}_I \rho_I^{-1} \frac{\beta_{pi} - \tilde{B}}{\tilde{B}} + (N_L - 1) M_L^{-1} \frac{\beta_{pL} - \tilde{B}}{\tilde{B}} + N_D M_D^{-1} \frac{\beta_{pd} - \tilde{B}(1 - \chi)}{\tilde{B}} \right)^{-1}$$

To determine the dealers' demand $q_d = \frac{\mathbb{E}(f|S_d) - p(1 - \chi)}{M_D}$, the solution is identical, except that the dealers' signal precision is higher (as she receives a noiseless signal about average order flow), and therefore her posterior variance $\hat{\tau}^{-1}$ is lower.

Following the same steps, we find that we can express M_D^{-1} as follows:

$$M_D = \rho_D \hat{\tau}_D^{-1} - \left(N_I \hat{\tau}_I \rho_I^{-1} \frac{\beta_{pi} - \tilde{B}}{\tilde{B}} + N_L M_L^{-1} \frac{\beta_{pL} - \tilde{B}}{\tilde{B}} + (N_D - 1) M_D^{-1} \frac{\beta_{pd} - \tilde{B}(1 - \chi)}{\tilde{B}} \right)^{-1}$$

Finally, we can express the price coefficients as functions of these M 's. To do this, we substitute all the investors' demands into the market-clearing condition we used before

(i.e., we sum over all investors, instead of all but one), and then match coefficients to yield the following results:

$$A = C \left[N_I \hat{\tau}_I \rho^{-1} (1 - \beta_{sI} - \beta_{\bar{s}I} - \beta_{pI}) \mu + N_L M_L^{-1} (1 - \beta_{sL} - \beta_{\bar{s}L} - \beta_{pL}) \mu + N_D M_D^{-1} (1 - \beta_{\bar{s}D} - \beta_{pD}) \mu - \tilde{B}^{-1} (A + C \bar{x}) (\beta_{pI} + \beta_{pL} + \beta_{pD}) - 1 \right] \quad (38)$$

$$B_I = C N_I (\hat{\tau}_I \rho^{-1} \beta_{sI} + M_D^{-1} \beta_{\bar{s}d} \omega_I) + N_I D \omega_I \quad (39)$$

$$B_L = C N_L (M_L^{-1} \beta_{sL} + M_D^{-1} \beta_{\bar{s}d} \omega_L) + N_L D \omega_L \quad (40)$$

$$B_D = C N_D M_D^{-1} \beta_{\bar{s}d} \omega_D + N_D D \omega_D \quad (41)$$

$$C = - \left(N_I \hat{\tau}_I \rho^{-1} \frac{\beta_{pI} - \tilde{B}}{\tilde{B}} + N_L M_L^{-1} \frac{\beta_{pL} - \tilde{B}}{\tilde{B}} + N_D M_D^{-1} \frac{\beta_{pd} - \tilde{B}(1 - \chi)}{\tilde{B}} \right)^{-1} \quad (42)$$

$$D = C \left(\hat{\tau}_I \rho^{-1} \frac{N_I}{N_D} \beta_{\bar{s}I} + M_L^{-1} \frac{N_L}{N_D} \beta_{\bar{s}L} \right) \quad (43)$$

The solution to this model is the set of β 's for each type of agent, price coefficients, and M coefficients which jointly solve the above, the implicit definitions of M_L and M_D , and the definitions of the Bayesian updating weights.

A.1.2 Case 2: Dealer-Client and Dealer-Dealer Information Sharing

In this setup, dealers share information with clients using the same noisy signal as before, but they also share information with ψ other dealers. Dealer-dealer sharing is symmetric, which requires that the number of dealers in an information-sharing collective be a factor of 20. We also require that $\psi \neq 19$, as that would imply perfect inter-dealer sharing. Thus, we only consider $\psi \in \{0, 1, 3, 4, 9\}$.

It would be repetitive to re-derive each part of the preceding analysis when most of it can be preserved. So, rather than do that, we will wherever possible argue by analogy; that is, by pointing out that such-and-such term maps to such-and-such combination of terms.

Perfect Dealer-Dealer Sharing First, observe that if dealers transmit their signals \bar{s}_d to one another noiselessly, then this result follows from the above. This is so because the effects of dealer-dealer sharing would simply be as if each dealer observed a greater proportion of the overall order flow. Each dealer's signal \bar{s}'_d , which we define to be their signal about order flow after dealer-dealer sharing is complete, would be more precise, and consequently each investor's signal $\bar{s}'_d + \xi_d$ would be more precise, but no part of the above machinery would be rendered invalid.

Imperfect Dealer-Dealer Sharing, Information Structure Define a *transmission signal* $\gamma_D \equiv \bar{s}_D + \delta_D$, where δ_D is dealer-dealer transmission noise which is i.i.d. across dealers. We say that $\delta_d \sim N(0, \sigma_{\gamma_D}^2)$.

First, note that from the point of view of non-dealer investor i , all objects are nearly exactly the same. We need to add another set of orthogonal shocks $\delta_1, \dots, \delta_D$ to identify those dealers who share with i 's dealer:

$$Z = [\epsilon_1, \dots, \epsilon_N, \xi_1, \dots, \xi_D, \delta_1, \dots, \delta_D, x]' \quad (44)$$

Agents' signals load on Z according to the signal loading matrix:

$$\Pi_i = \begin{bmatrix} \phi_i & \mathbf{0}_{\mathbf{N}_D} & \mathbf{0}_{\mathbf{N}_D} & 0 \\ \frac{\omega_i}{|\omega_i|} & \psi_{d(i)} & \zeta_{d(i)} & 0 \\ \frac{B_L}{B} \cdot \mathbf{1}_{\mathbf{N}_I}, \frac{B_L}{B} \cdot \mathbf{1}_{\mathbf{N}_L}, \frac{B_D}{B} \cdot \mathbf{1}_{\mathbf{N}_D} & 0 \cdot \mathbf{1}_{\mathbf{N}_D} & 0 \cdot \mathbf{1}_{\mathbf{N}_D} & \frac{C}{B} \end{bmatrix} \quad (45)$$

In the above, $\omega_i(j)$ returns the type-specific precision for investor j if i and j bid through the same dealer, *or* if j bids through a dealer who talks with i 's dealer. $\psi_{d(i)}$ is the dealer identifier as before, and we define a new $1 \times N_D$ vector $\zeta_{d(i)}$, which is 0 everywhere, except for those dealers who *talk to* i 's dealer.

Which in turn entails the following posterior expectations and variances:

$$\mathbb{E}(f|S_i) = (1 - \beta_{\bar{s}'i} - \beta_{pi} - \beta_{si})\mu + \beta_{si}s_i + \beta_{\bar{s}'i}\bar{s}\xi_{d(i)} + \beta_{pi}\frac{p - A - Cx}{\bar{B}} \quad (46)$$

$$\mathbb{V}(f|S_i) = \tau_f^{-1} - \tau_f^{-1}(\beta_{si} + \beta_{\bar{s}'i} + \beta_{pi}) \quad (47)$$

Here, the weights must be re-derived using the formulas introduced in the preceding case. The main difference will be to the term $\beta_{\bar{s}'i}$, as dealer-dealer sharing induces agents to weight the signal more heavily, and dealer-dealer noise induces agents to weight it less heavily.

The setup for dealer j is slightly different, in that j observes the aggregated dealer signal \bar{s}'_d without ξ noise, and that she observes her clients' order flow without ζ noise.

So, define the dealer's signal vector as the 3-dimensional object $[\bar{s}, s', s(p)]$, where the first term denotes information received from clients, the second term denotes information received from other dealers, and the third term denotes information received from the unobserved price.¹⁸

This enables us to write the following loading matrix for dealer j , which in turn entails a specific set of Bayesian updating weights by the formulas above. Below, define λ to be the original ω (i.e., the signal aggregator over our dealer's clients), and ω similarly to the above

¹⁸In practice, s' will be a ψ -dimensional object, as each dealer speaks with ψ others. However, as all dealer-dealer noise is i.i.d., and all dealers are symmetric, this reduces to a simple average. In our notation, s' refers to this averaged signal of all dealer's messages.

(i.e., the signal aggregator over clients of *other* dealers, who talk to our dealer).

$$\Pi_i = \begin{bmatrix} \frac{\lambda_i}{|\lambda_i|} & 0 \cdot \mathbf{1}_{\mathbf{N}_D} & 0 \cdot \mathbf{1}_{\mathbf{N}_D} & 0 \\ \frac{\omega_i}{|\omega_i|} & 0 \cdot \mathbf{1}_{\mathbf{N}_D} & \zeta_{d(i)} & 0 \\ \frac{B_L}{\bar{B}} \cdot \mathbf{1}_{\mathbf{N}_I}, \frac{B_L}{\bar{B}} \cdot \mathbf{1}_{\mathbf{N}_L}, \frac{B_D}{\bar{B}} \cdot \mathbf{1}_{\mathbf{N}_D} & 0 \cdot \mathbf{1}_{\mathbf{N}_D} & 0 \cdot \mathbf{1}_{\mathbf{N}_D} & \frac{C}{\bar{B}} \end{bmatrix} \quad (48)$$

Imperfect Dealer-Dealer Sharing, Coefficient Matching We inherit the demand functions from the previous section. For non-dealer investors, we also inherit the Bayesian updating weights, with the exception that $\beta_{\bar{s}}$ is replaced with $\beta_{\bar{s}'}$, reflecting changes in the signal the dealer shares with her clients.

For dealers, we have a new set of beliefs:

$$\begin{aligned} \mathbb{E}(f|S_j) &= (1 - \beta_{\bar{s}d} - \beta_{pd} - \beta_{s'\zeta d})\mu + \beta_{\bar{s}d}\bar{s}_d + \beta_{s'\zeta d}s'_{\zeta d} + \beta_{pd}\frac{p - A - C\bar{x}}{\bar{B}} \\ \mathbb{V}(f|S_j) &= \tau_f^{-1} = \tau_f^{-1}(\beta_{\bar{s}d} + \beta_{s'\zeta d} + \beta_{pd}) \end{aligned}$$

In the above, the signal $s'_{\zeta d}$ is the realized (i.e., noisy) signal that d shares with her clients after dealer-dealer sharing is complete. The term $\beta_{s'\zeta d}$ is the corresponding Bayesian weight, which is computed using the formula above.

The only difference in the market-clearing condition is that M_D now modifies a term with instances of $\beta_{s'd}s'_{\zeta d}$, i.e.:

$$1 = x + \dots + M_D^{-1} \left((1 - \beta_{\bar{s}d} - \beta_{pd} - \beta_{s'\zeta d}) + \beta_{\bar{s}d}\bar{s}_d + \beta_{s'\zeta d}s'_{\zeta d} + \beta_p\frac{p - A - C\bar{x}}{\bar{B}} - p \right)$$

It's clear that adding in that constant doesn't affect our ability to match coefficients — any terms which aren't common, like that representing dealer-dealer noise, are simply present in the constant A . Therefore, we still have linearity in prices under noisy (yet symmetric) dealer-dealer signal transmission, and the result holds.

A.1.3 Case 3: No Information Sharing (“Chinese Wall”)

In this model, dealers do not use or share any information derived from client order flow. Practically speaking, it is as if each type of investor submits bids on their own behalf, rather than through an intermediary. Each investor's information set is therefore a 2×1 vector $S_i = [s_i, s(p)]$, comprised of their private signal s_i and the counterfactual price signal $s(p)$.

Since there is no longer a meaningful informational distinction between dealers and large investors, we refer to both as large investors, i.e.: $N'_L \equiv N_L + N_D$.¹⁹

¹⁹In this regime, the only difference between dealers and non-dealer large investors is that dealers are subject to a minimum bidding penalty.

Information Structure As above, the linear form we assume for prices entails a specific form for $s(p)$:

$$s(p) = \frac{p - A - Cx}{\tilde{B}} \quad (49)$$

The above signal is unbiased, but as we noted, the price signal and the private signals have correlated errors. The Bayesian updating weights β come from optimal linear projection formulas that correct for this covariance.

The shock vector is now simply:

$$Z = [\epsilon_1, \dots, \epsilon_N, x]', \quad (50)$$

where individuals are ordered by putting small investors first, then large investors, then dealers. The variance vector of this matrix is:

$$\mathbb{V}(Z) = \text{diag}(\tau_{\epsilon,I}^{-1} \mathbf{1}_{\mathbf{N}_I}, \tau_{\epsilon,L}^{-1} \mathbf{1}_{\mathbf{N}'_L}, \tau_x^{-1}) \quad (51)$$

The loading matrix is given by:

$$\Pi_i = \begin{bmatrix} \phi_i & 0 \\ \frac{B_I}{\tilde{B}} \mathbf{1}_{\mathbf{N}_I}, \frac{B_L}{\tilde{B}} \mathbf{1}_{\mathbf{N}'_L} & \frac{C}{\tilde{B}} \end{bmatrix} \quad (52)$$

We can then determine the posterior beliefs exactly as previously. Because dealers are indistinguishable from large investors here, these formulas are true for all agents:

$$\mathbb{E}(f|S_z) = (1 - \beta_{sz} - \beta_{pz})\mu + \beta_{sz}s_z + \beta_{pz} \frac{p - A - Cx}{\tilde{B}} \quad (53)$$

$$\mathbb{V}(f|S_z) = \tau_f^{-1} - \tau_f^{-1}(\beta_{sz} + \beta_{pz}), \quad (54)$$

where z is a type variable $z \in \{I, L\}$.

Coefficient Matching We can follow exactly the same process as in the above two proofs to yield the final set of coefficients. For clarity, only those coefficients are reported here:²⁰

$$\begin{aligned} A &= -C \left(1 + (N_I \hat{\tau}_I \beta_{pI} + N_L M_L^{-1} \beta_{pL} + N_D M_D^{-1} \beta_{pd}) \frac{A + C\bar{x}}{\tilde{B}} \right) \\ &\quad + C \left(N_I \frac{\hat{\tau}_I}{\rho} (1 - \beta_{sI} - \beta_{pI}) + N_L M_L^{-1} (1 - \beta_{sL} - \beta_{pL}) + N_D M_D^{-1} (1 - \beta_{sd} - \beta_{pd}) \right) \mu \\ B_I &= CN_I \rho^{-1} \beta_{sI} \\ B_L &= CN_L M_D^{-1} \beta_{sL} \\ B_D &= CN_D M_D^{-1} \beta_{sD} \\ C &= -\tilde{B} \left[N_I \rho^{-1} (\beta_{pI} - \tilde{B}) + N_L M_L^{-1} (\beta_{sL} - \tilde{B}) + N_D M_D^{-1} (\beta_{sD} - \tilde{B}(1 - \chi)) \right]^{-1} \\ D &= 0 \end{aligned}$$

²⁰Recall that once again $N_L \equiv$ the number of non-dealer large investors, and $N_D \equiv$ the number of dealers.

The existence of a set of coefficients verifies the price conjecture. Since the supply of the asset is one, auction revenue is the price of the asset. The solution to this model is a joint solution to (7)-(11) and (13)-(18)

A.2 Corollary 1

We argue by contradiction here, by showing that noncompetitive shares decreased the expected profit $\mathbb{E}(\pi) = \mathbb{E}(f - p)$. Since the total supply of the asset is 1, which is the sum of competitive and non-competitive demand, the competitive share must increase $\mathbb{E}(\pi)$. We can break the partial derivative $\frac{\partial \pi}{\partial x}$ into the two terms, the first of which is zero by construction, as f and x are exogenous, independent variables.

The second term can be expressed as $\frac{\partial A+Bf+Cf}{\partial x}$. Note that $\frac{\partial p}{\partial x} > 0$ if C is negative.

First, we assume that $N_I \gg N_L, N_D$. This reduces the expression for C to:

$$C = -N_I \frac{\beta_{sI} - \tilde{B}}{\tilde{B}} \quad (55)$$

C is positive if $\tilde{B} > \beta_{sI}$, or that the price puts more weight on the true payoff f than individuals do.

A.3 Result 2

Recall that this result states that in an auction with N_D dealers, and one large investor who bids directly, prices follow a particular linear form. We solve the model generally with a minimum bidding penalty, modeled by a per-period cost χ , after which $\chi = 0$ lets us evaluate results for the no-penalty model.

This proof is divided into three parts. The first two are the standard sections we have seen before, and the third considers the case where the dealer incorporates the information gleaned from the intermediation decision of the large investor into the signal about order flow that she shares with her clients.

Information Structure

When the first large investor chooses to bid directly on his own behalf, the aggregated signal for the first dealer d is the precision-weighted average of her proportional share of small and large investors, less one large investor, i.e.:

$$s_1 = \frac{\tau_{\epsilon,I} \sum_{k \in I^I} s_k + \tau_{\epsilon,l} \left(\sum_{j \in I^L} s_j + s_d \right)}{\tau_{\epsilon,I} N_I / N_D + \tau_{\epsilon,l} (N_L / N_D - 1) + \tau_{\epsilon,l}} \quad (56)$$

As in the previous model, investor i who bids through dealer d observes signals $S_i = [s_i, \bar{s}_d + \xi_d, s(p)]$. The large investor bidding directly observes only his own signal s_j , and the counterfactual price signal $s(p)$, i.e.: $S_j = [s_j, s(p)]$.

Variance-Covariance Matrix of Signals The only difference in the signal construction is that the first large bidder does not trade through the large dealer. Thus, the large bidder's signals are only his own private signal and the price. The first dealer's information (and, therefore, her clients') are less precise because they miss the one large investor.²¹

We can construct loading matrices Π that map shocks into signals as before. Let

$$\bar{Z} = [\epsilon_{L1}, \epsilon_1, \dots, \epsilon_{N_D}, x] \quad (57)$$

Then:

$$\Pi_L = \begin{bmatrix} 1 & 0 & \dots & 0 \\ \frac{B_L}{\tilde{B}} & \frac{B_{d1}}{\tilde{B}} & \frac{B_{d \neq 1}}{\tilde{B}} \mathbf{1}_{N_D-1} & \frac{C}{\tilde{B}} \end{bmatrix} \quad (58)$$

$$\Pi_{d=1} = \begin{bmatrix} 0 & 1 & 0 \dots & 0 \\ \frac{B_L}{\tilde{B}} & \frac{B_{d1}}{\tilde{B}} & \frac{B_{d \neq 1}}{\tilde{B}} \mathbf{1}_{N_D-1} & \frac{C}{\tilde{B}} \end{bmatrix} \quad (59)$$

$$\Pi_{d \neq 1} = \begin{bmatrix} 1 & & \psi_i & 0 \\ \frac{B_L}{\tilde{B}} & \frac{B_{d1}}{\tilde{B}} & \frac{B_{d \neq 1}}{\tilde{B}} \mathbf{1}_{N_D-1} & \frac{C}{\tilde{B}} \end{bmatrix} \quad (60)$$

where $\tilde{B} \equiv B_{d1} + B_{d \neq 1} + B_L$ and $d \neq 1$ stands for all dealers, besides dealer 1.

Then, for all agents j who are clients of dealer $d(j)$, or for the large investor $d(j) = L$, the signal variance-covariance is

$$\mathbb{V}(S_j) = \tau_f^{-1} + \bar{\Pi}_{d(j)} \mathbb{V}(\bar{Z}) \bar{\Pi}'_{d(j)} \quad \forall j \quad (61)$$

Bayesian Updating Weights For each agent who bids through a dealer, there are three signals — the dealer's signal ξ_s , the private signal s_i , and the counterfactual price signal $s(p)$. We can compute corresponding Bayesian weights using the machinery introduced previously. These have precision given by:

$$\hat{\tau}_l = \tau_f^{-1} (1 - \beta_s - \beta_p) \quad (62)$$

Coefficient Matching We work through the same process as above (i.e., substitute all but one agent into the market clearing constraint to solve implicitly for M_L and M_D , and then go back and match coefficients). This yields the following set of price sensitivities:

²¹Note that in the following, the first dealer's information set — and therefore ϵ_1 — include the information encoded in the first large investor's intermediation decision.

$$\begin{aligned}
\left(\frac{dp}{dq_L}\right)^{-1} &= -(\nu_I \hat{\tau}_I \rho^{-1} + (\nu_L - 1)M_{L(d=1)}^{-1} + M_{d=1}^{-1})\frac{\beta_{p,d1} - \tilde{B}}{\tilde{B}} \\
&- \chi M_{d=1}^{-1} - \chi(N_D - 1)M_{d \neq 1}^{-1} - (N_D - 1)(\nu_I \hat{\tau}_L \rho_L^{-1} + \nu_L M_{L(d \neq 1)}^{-1} + M_{d \neq 1}^{-1})\frac{\beta_{p,d \neq 1} - \tilde{B}}{\tilde{B}} \\
\left(\frac{dp}{dq_{L(d=1)}}\right)^{-1} &= \left(\frac{dp}{dq_L}\right)^{-1} - M_L^{-1}\frac{\beta_{pL} - \tilde{B}}{\tilde{B}} + M_{L(d=1)}^{-1}\frac{\beta - \tilde{B}}{\tilde{B}} \\
\left(\frac{dp}{dq_{L(d \neq 1)}}\right)^{-1} &= \left(\frac{dp}{dq_L}\right)^{-1} - M_L^{-1}\frac{\beta_{pL} - \tilde{B}}{\tilde{B}} + M_{L(d \neq 1)}^{-1}\frac{\beta - \tilde{B}}{\tilde{B}} \\
\left(\frac{dp}{dq_{(d=1)}}\right)^{-1} &= \left(\frac{dp}{dq_L}\right)^{-1} - M_L^{-1}\frac{\beta_{pL} - \tilde{B}}{\tilde{B}} + M_{d=1}^{-1}\left(\chi + \frac{\beta_{p,d1}\tilde{B}}{\tilde{B}}\right)
\end{aligned}$$

This holds true for dealers, recalling their altered demand function:

$$M_D \equiv \rho_L \hat{\tau}_L + (1 - \chi)\frac{dp}{dq_L}$$

$$\begin{aligned}
\left(\frac{dp}{dq_{(d=1)}}\right)^{-1} &= \left(\frac{dp}{dq_L}\right)^{-1} - M_L^{-1}\frac{\beta_{pL} - \tilde{B}}{\tilde{B}} + M_{d=1}^{-1}\left(\chi + \frac{\beta_{p,d1}\tilde{B}}{\tilde{B}}\right) \\
\left(\frac{dp}{dq_{(d \neq 1)}}\right)^{-1} &= \left(\frac{dp}{dq_L}\right)^{-1} - M_L^{-1}\frac{\beta_{pL} - \tilde{B}}{\tilde{B}} + M_{d \neq 1}^{-1}\left(\chi + \frac{\beta_{p,d1}\tilde{B}}{\tilde{B}}\right)
\end{aligned}$$

And the following set of pricing coefficients:

$$\begin{aligned}
A &= C \left(-1 + \left(\nu_I \hat{\tau}_1 \rho^{-1} + (\nu_l - 1) M_{L(d=1)}^{-1} + M_{d=1}^{-1} \right) \left((1 - \beta_{\bar{s},d1} - \beta_{p,d1}) \mu - \frac{\beta_{p,d1}}{\tilde{B}} (A + C\bar{x}) \right) \right) \\
&+ C (N_D - 1) \left(\nu_I \hat{\tau}_2 \rho^{-1} + \nu_l M_{L(d \neq 1)}^{-1} + M_{d \neq 1}^{-1} \right) \left((1 - \beta_{\bar{s},d \neq 1} - \beta_{p,d \neq 1}) \mu - \frac{\beta_{p,d \neq 1}}{\tilde{B}} (A + C\bar{x}) \right) \\
&+ C M_L \left((1 - \beta_{sL} - \beta_{pL}) \mu - \frac{\beta_{pL}}{\tilde{B}} (A + C\bar{x}) \right) \tag{63}
\end{aligned}$$

$$B_{d=1} = C \left(\nu_I \hat{\tau}_1 \rho^{-1} + (\nu_l - 1) M_{L(d=1)}^{-1} + M_{d=1}^{-1} \right) \beta_{\bar{s},d1} \tag{64}$$

$$B_{d \neq 1} = C (N_D - 1) \left(\nu_I \hat{\tau}_2 \rho^{-1} + \nu_l M_{L(d \neq 1)}^{-1} + M_{d \neq 1}^{-1} \right) \beta_{\bar{s},d \neq 1} \tag{65}$$

$$B_L = C M_L^{-1} \beta_L(1) \tag{66}$$

$$\begin{aligned}
C^{-1} = & -M_L \left(\frac{\beta p L - \tilde{B}}{\tilde{B}} \right) - \left(\nu_I \hat{\tau}_1 \rho^{-1} + (\nu_l - 1) M_{L(d=1)}^{-1} + M_{d=1}^{-1} \right) \left(\frac{\beta_{p,d1} - \tilde{B}}{\tilde{B}} \right) - \chi M_{d=1}^{-1} \\
& - (N_D - 1) \left(\nu_I \hat{\tau}_2 \rho^{-1} + \nu_l M_{L(d \neq 1)}^{-1} + M_{d \neq 1}^{-1} \right) \left(\frac{\beta_{p,d \neq 1} - \tilde{B}}{\tilde{B}} \right) - (N_D - 1) \chi M_{d \neq 1}^{-1}. \quad (67)
\end{aligned}$$

B Learning from the intermediation decision of the large investor

In this appendix, we consider the model where dealer 1 incorporates the information that can be learned from the intermediation choice of the large investor in the average signal shared with the rest of his clients. When the large investor chooses to bid through the dealer, the dealer observes the investor's signal directly and the equilibrium outcomes are the same as before. We solve this model generally with a minimum bidding penalty. Then we can set $\chi = 0$ to get the results for the no-penalty model.

Consider now the case when the large investor bids directly in the auction. Recall that the large investor chooses to bid directly whenever $\mathbb{E}U(Ld) \geq \mathbb{E}U(Li)$. Using (17), we can rewrite this as

$$(1 + 2\theta_{Ld} \Delta \mathbb{V}_{Ld})^{-\frac{1}{2}} \exp \left(-\frac{\theta_{Ld} \mu_{r,Ld}^2}{1 + 2\theta_{Ld} \Delta \mathbb{V}_{Ld}} \right) \leq (1 + 2\theta_{Li} \Delta \mathbb{V}_{Li})^{-\frac{1}{2}} \exp \left(-\frac{\theta_{Li} \mu_{r,Li}^2}{1 + 2\theta_{Li} \Delta \mathbb{V}_{Li}} \right).$$

Equivalently, the large investor chooses to bid directly when

$$\exp \left(-\frac{\theta_{Ld} \mu_{r,Ld}^2}{1 + 2\theta_{Ld} \Delta \mathbb{V}_{Ld}} + \frac{\theta_{Li} \mu_{r,Li}^2}{1 + 2\theta_{Li} \Delta \mathbb{V}_{Li}} \right) \leq \sqrt{\frac{1 + 2\theta_{Ld} \Delta \mathbb{V}_{Ld}}{1 + 2\theta_{Li} \Delta \mathbb{V}_{Li}}},$$

or

$$\frac{\theta_{Ld} \mu_{r,Ld}^2}{1 + 2\theta_{Ld} \Delta \mathbb{V}_{Ld}} - \frac{\theta_{Li} \mu_{r,Li}^2}{1 + 2\theta_{Li} \Delta \mathbb{V}_{Li}} \geq -\frac{1}{2} \log(1 + 2\theta_{Ld} \Delta \mathbb{V}_{Ld}) + \frac{1}{2} \log(1 + 2\theta_{Li} \Delta \mathbb{V}_{Li}).$$

The left hand side of the above is a quadratic function of the signal of the large investor, s_{N_I+1} . Let $\Upsilon_{l,u}$ be the two solutions to the quadratic equation for the boundaries of the indirect bidding region, so that the large investor bids indirectly whenever $\Upsilon_l \leq s_{N_I+1} \leq \Upsilon_u$; equivalently, the large investor bids directly if $s_{N_I+1} \geq \Upsilon_u$ or $s_{N_I+1} \leq \Upsilon_l$.

To solve the model where all agents make rational inferences from the intermediation decision, we first solve our model without this information and determine the cutoffs ν_l and ν_h . We simulate the model to determine the probability of direct bidding $Pr(v)$. We also compute, conditional on choosing to bid directly, what the direct bidder's average signal is (s_v) and the variance of that signal (τ_v^{-1}). Next, we construct a normally-distributed, conditionally independent approximating signal with the same mean and variance: $s_v = f + \epsilon_v$ where $\epsilon_v \sim N(0, \tau_v^{-1})$. In cases where the large investor bids directly, we allow dealer 1 to observe s_v and incorporate it in his advice to his clients. When the large investor bids indirectly, his dealer observes his signal exactly, as in the indirect bidding model we solved before. The other dealers and their clients do not observe the intermediation

decision and instead use $Pr(v)$ to weight the price signal that would be extracted in both scenarios.

Variance-covariance matrix of signals When the large investor bids directly, the first dealer's signal is the average of the first ν_l investors', the first $\nu_l - 1$ large investors', the first dealer's signal and the new signal s_v that arises from observing the direct bidding decision. By Bayes' law, the posterior expectation, conditional on all these signals is:

$$\bar{s}_1 = \frac{\tau_{\varepsilon,s} \sum_{k \in \mathcal{I}_1^s} s_k + \tau_{\varepsilon,l} \left(\sum_{j \in \mathcal{I}_1^l} s_j + s_d \right) + \tau_v s_v}{\nu_l \tau_{\varepsilon,s} + \nu_l \tau_{\varepsilon,l} + \tau_v}$$

Bayes law also tells us that the precision is $\hat{\tau}_{d1} = \nu_l \tau_{\varepsilon,s} + \nu_l \tau_{\varepsilon,l} + \tau_v$. The signals of all other dealers are the same as before because those dealers are not aware of the large investor's intermediation decision.

As in the previous model, investor i who bids through intermediary d observes signals $\tilde{S}_i = [s_i, \bar{s}_d, s(p)]$. Since dealer share information perfectly here, the investor's own signal is redundant. Thus, for notational convenience, we continue as if $\tilde{S}_i = [\bar{s}_{d(i)}, s(p)]$, $\forall i \neq L$. The large investor bidding directly observes only his own signal and the price information: $\tilde{S}_L = [s_L, s(p)]$.

We can construct the signals as the true payoff f plus weights Π on orthogonal shocks Z . Let

$$\bar{Z}_v = [\epsilon_L \quad \bar{\epsilon}_1 \quad \bar{\epsilon}_2 \quad \dots \quad \bar{\epsilon}_{N_d} \quad (x - \bar{x})]'$$

where $\epsilon_l \equiv S_L - f$ and $\bar{\epsilon}_d \equiv \bar{s}_d - f$ are the noise the the signals of the large investors and the dealer and x is the non-competitive bids. Note that the information from the intermediation decision of the direct bidder is incorporated in the first dealer's signal, and thus in $\bar{\epsilon}_1$. The variances of these shocks are:

$$\mathbb{V} [\bar{Z}^{direct}] = \text{diag}([\tau_L^{-1}, \hat{\tau}_{d1}^{-1}, \tau_d^{-1} \mathbf{1}_{N_d-1}, \tau_x^{-1}]). \quad (68)$$

The signals of the large investor L , dealer 1 and his clients $d(1)$ and all others $d \neq 1$ have loadings on the orthogonal shocks given by (58), (59) and (60). Then, for all agents j who are clients of dealer $d(j)$, or for the large investor $d(j) = L$, the signal variance-covariance is $\mathbb{V}(S_j) = \tau_f^{-1} + \Pi_d(j) V[\bar{Z}_v] \Pi_d(j)'$ $\forall j$.

Bayesian updating weights β and posterior precisions τ The posterior expectation of the asset value is a linear combination of the prior and two signals: (1) the signal s_L or $s_{d(i)}$ provided by the dealer; and (2) the signal conveyed by the price $(p - A - C\bar{x})/B$. β_s and β_p are the weights the agent places on the two signals. The weights are given by (5), the resulting conditional expectation is given by (4), with precision (62).

Equilibrium Price Once we've adjusted the Bayesian updating weights β , the rest of the solution of the direct bidding model follows just as before. The equilibrium price coefficients are a joint solution to the coefficient equations (63) - (67), the M equation (A.1.1), in conjunction with the price sensitivities (63) - (A.3) and the β 's (5), which in turn depend

on the Π 's (58), (59) and (60) and the variances of the orthogonal shocks (61). Setting $\chi = 0$ yields the solution to the model without the minimum bidding requirement.

C A Private Value Auction Model

The model economy lasts for two periods and agents can invest in a risky asset (the newly issued Treasury security) and a riskless storage technology with zero net return. The risky asset is auctioned by Treasury in a fixed number of shares (normalized to 1) using a uniform-price auction with a market-clearing price p . The value of the newly issued asset to small investor i is \tilde{f}_i , which is known only to investor i . The small investors' valuations are correlated. They have an aggregate and an idiosyncratic component that are not separately observed: $\tilde{f}_i = f + v_i$, where $f \sim N(\mu, \tau_f^{-1})$ and $v_i \sim N(0, \tau_v^{-1})$. Each investor knows the private component of their valuation (perhaps a hedging motive), but not the common valuation f or the sum \tilde{f}_i . Large investors' valuation for the treasury is $\tilde{f}_i = f$, which is also unknown to them.

As before, preferences are exponential,

$$\begin{aligned} \max_{q_i(p)} \mathbb{E}[-\exp(-\rho W_i) | S_i] \\ \text{s.t. } W_i = W_{0i} + q_i(\tilde{f}_i - p). \end{aligned}$$

The timing, information sets, intermediation choices are all as before, except for one twist. An investor's order flow no longer perfectly reveals their signal s_i . The dealer cannot perfectly disentangle private information s_i from private valuation v_i . So, we assume that dealers see their clients bids q_i , but not that dealers know s_i directly.

The definition of equilibrium is the same as before, with the adjusted preferences and the one change to the information set.

Updating Investors who bid through dealers have access to four pieces of information. They know their priors and their private signal. They get information from their dealer and they observe the market price. The first two pieces of information are the same as before. The second two change.

By placing orders through dealers, customers reveal their order flow to their dealer. From the first order condition for portfolio choice q_i , that order is

$$q_i = \frac{E[f|S_i] + v_i - p}{\rho \text{Var}[f|S_i]}$$

Since risk aversion and conditional variances are known, and price p is observable to all, observing q_i lets a dealer observe $E[f|S_i] + v_i$. We can then decompose $E[f|S_i]$ into the private signal, common prior and the price:

$$E[f|s_i, p] = \frac{\tau_f \mu + \tau_s s_i + \tau_p p}{\tau_f + \tau_s + \tau_p}$$

Table 4: Descriptive statistics for the calibrated, simulated model with direct and indirect bidding and low-bid penalty. Revenue is in basis points, and allocations are in percent.

	Revenue	Dealer allocation	Direct allocation	Indirect allocation
Mean	38.7236	43.3994	0.4698	50.5758
Std. Dev.	70.8100	19.4523	0.7284	73.1065
Skew	-0.0075	-0.0091	1.2365	0.1008
Kurtosis	3.0060	3.0326	3.2103	4.2614

Thus, the dealer can form a signal $s_{qi} = [(\tau_f + \tau_s + \tau_p)(E[f|S_i] + v_i) - (\tau_f\mu + \tau_p\eta_p)]/\tau_s$. This yields a signal $s_{qi} = s_i + e_{qi}$ where $e_{qi} = (\tau_f + \tau_s + \tau_p)/\tau_s v_i$, which is distributed $N(0, \tau_q^{-1})$, where the signal variance is $\tau_v^{-1}(\tau_f + \tau_s + \tau_p)^2/\tau_s^2$. But s_{qi} is a noisy signal about the underlying noisy signal s_i . Thus s_{qi} as a signal of f has a conditional variance equal to the sum of the variance of each of the signal components. Let $\tau_q^{-1} \equiv \tau_i^{-1} + \tau_v^{-1}(\tau_f + \tau_s + \tau_p)^2/\tau_s^2$. Then τ_q is the precision of s_{qi} as a signal about f .

Thus, if we re-write the dealer's Bayesian signal-precision-weighted average of his and his clients' private signals, it is

$$\bar{s}_d = \frac{\sum_{i \in \mathcal{I}_d} \tau_q s_{qi}}{\sum_{i \in \mathcal{I}_d} \tau_i} = \frac{\tau_q \sum_{k \in \mathcal{I}_d^s} s_{qk} + \tau_{\varepsilon,l} \left(\sum_{j \in \mathcal{I}_d^l} s_j + s_d \right)}{\nu_I \tau_q + (1 + \nu_l) \tau_{\varepsilon,l}}, \quad (69)$$

Once we redefine this signal \bar{s}_d and adjust its precision to be $\tau_\xi = \nu_I \tau_q + (1 + \nu_l) \tau_{\varepsilon,l}$, we can then proceed to solve the model, exactly the same as before.

The last piece of the model that could be affected by the presence of private values is the information extracted from prices. But, it turns out that once we adjust the dealers' signal precision, there is not further change needed to price informativeness. The reason is that the average of private values is known. It's zero. That's common knowledge. So, when we aggregate demand to clear the market, we get

$$\begin{aligned} \int \frac{E[f|S_i] + v_i - p}{\rho \text{Var}[f|S_i]} di &= 1 \\ &= \int \frac{E[f|S_i] - p}{\rho \text{Var}[f|S_i]} di \end{aligned}$$

because $\int v_i di = 0$ and v_i is uncorrelated with any other variables in the model.

So, while private values affect individuals' demands, they average out and have no direct effect on prices. The only aggregate effect of private values is through their effect on the dealers' information \bar{s}_d and its precision τ_ξ .

D Measuring Treasury Payoffs

This appendix provides additional detail about how payoffs are calculated. Because of lags between trade and settlement dates, the appendix also provides detail on funding costs. The first subsection describes what those terms are and argues that they are small and stable. The second subsection discusses an alternative hedging strategy, known as a coupon roll. The third explains why information from the when-issued-market (or WIs) is not relevant in our setting.

Funding position. In the model, winning bids pay p and the common fundamental value is f . In Treasury auctions bidders bid a coupon rate rather than a price. The price is always set to \$100 up to rounding, which we rescale to \$1 for the purposes of this discussion. To assess auction results from the issuer perspective we discount future interest payments using a yield curve estimated on outstanding Treasury securities. Economically this means that we measure issuance cost relative to other debt outstanding at the time of the auction. Newly issued Treasury securities are typically valued more than older securities because of their better liquidity, a phenomenon known as the on-the-run premium (see e.g., Vayanos and Weill 2008). As a result of the on-the-run premium, the discounted value of Treasury's future interest and principal payments is smaller the price at which the security sells (\$1), and we define net auction revenue as the gap between the two:

$$\hat{R}_{\text{auction}} = 1 - \left(\sum_{t=0}^T Z_{\text{auction}}(t) C + Z_{\text{auction}}(T) \right), \quad (70)$$

where C is the coupon determined at the auction, T is the maturity, $Z_{\text{auction}}(i) = \exp(-iy_{\text{auction}}(i))$ is the price at the time of the auction of a zero-coupon bond maturing at i , $y(i)$ is the i th maturity yield from the yield curve estimated on outstanding securities at the time of the auction.

Trades in the secondary Treasury market settle on the business day following a trade, meaning that securities are delivered and cash is paid a day after a transaction is agreed upon. In Treasury auctions, instead, investors pay bids to Treasury and receive securities on the issuance date, which occurs one to 14 days following the date of the auction. This different settlement rule is the source of extra funding cost/income in our setting.

We measure f as the market price of the security on the issuance date, which is when the security is first available to investors. The value of f depends on the general level of interest rates and the on-the-run premium. While fluctuations in interest rates between auction and issuance date create risk for investors, this risk can be hedged with other outstanding Treasuries. We assume that investors hedge interest rate risk optimally by selling a replicating portfolio of other Treasury securities. On the auction date, the investor buys the new security and shorts the replicating portfolio of off-the-run issues. On the issuance date, the investor reverses by selling the new security and covering the short in older securities. The per-unit value of the hedged portfolio at auction is equal to $-\hat{R}_{\text{auction}}$, and to:

$$\hat{f}_{issue} = \left(\sum_{t=0}^T Z_{issue}(t) C + Z_{issue}(T) \right) - P_{issue}, \quad (71)$$

on the issuance date, where P_{issue} is the market price of the new security on that date. Detailed steps in the investment strategy are:

1. Auction date:
 - (a) Place bid
 - (b) For each unit of successful bid allotted, sell T zero coupon bond each priced at $Z_{auction}(t)$ and in amounts equal to C for $t < T$ and $1 + C$ for $t = T$. The zero coupon bonds could either be stripped Treasuries (as in [Fleckenstein, Longstaff, and Lustig, 2014](#)) or proxied with a combination of coupon securities.
2. Post-auction date:
 - (a) Borrow (to post-issuance date) the amount $Z_{auction} = \sum_{t=0}^T Z_{auction}(t) C + Z_{auction}(T)$ paying the per-diem unsecured rate r_b .
 - (b) Borrow zero-coupon bonds with reverse repos (to post-issuance date) and receive the per-diem repo rate r_{repo} . Deliver the T zero coupons to the auction-date buyer.
3. Issuance date:
 - (a) Borrow \$1 at rate r_b . Receive new issue from, and pay \$1, to Treasury; sell issue in the secondary market
 - (b) Buy portfolio of T zero-coupon bonds at $Z_{issuance}$
4. Post-issuance date:
 - (a) Receive payment of p_{issue} and repay the issuance-date loan
 - (b) Receive T zero-coupon bonds and deliver into the reverse repo;
 - (c) Receive payment of $Z_{auction}$ from reverse-repo and pay Z_{issue} to settle the issue-date purchase; Repay post-auction date loan

The cash flows from this position at the post-issue date are:

$$(P_{issue} - 1) + (Z_{auction} - Z_{issue}) + \frac{(\text{issue date} - \text{auction date})}{360} \times (r_{repo} - r_{borrow}) \times Z_{auction} - \frac{r_{borrow}}{360} \times 1 \quad (72)$$

In our calculations we disregard the two funding terms because they are small and don't vary much when $r_{repo} \approx r_{borrow}$. The repo rate for old issues, which are being funded between the post-auction and post-issue date, typically trades within a few basis points to the unsecured rate r_b , so the funding terms are small. Repo rates for new (or first-off-the-run) securities instead can trade far off from uncollateralized rates and be volatile because they funding rates balance the supply and demand of new securities, which can be in high demand to take short position in interest rates (see e.g. [Duffie, 1996](#); [Jordan and Jordan, 1997](#)). As per the detailed steps above the new issue is never shorted or funded, as it is

sold as soon as it is received by the investor. Thus fluctuations in the special-repo rate do not affect the returns in our position.

Coupon roll An investor could achieve approximately the same hedged position by shorting only the previously on-the-run (same maturity) security. This strategy is fairly common around Treasury auctions as discussed by [Fleming and Garbade \(2007\)](#). While this would be a preferred approach in practice, the paper focuses on a OTR strategy for two reasons. First, interest hedging with the former on-the-run is imperfect because maturities are not matched and additional accrued interest calculations would need to be accounted for. Second, the repo rate for recently issued securities can trade “special”, that is at a significant gap to the r_{borrow} so that the funding terms would become more important. At the same time historical special repo rates are not readily available, so we focus on OTR for which these terms are not important.

Information from the When-Issued Market Days ahead of the auction date, when details such as amounts and maturities are set, dealers and other market participants begin to trade the not-yet-available security. Trades on when-issued securities (or WIs), are quoted in coupon rates that investors are willing to pledge, in advance of the auction, to pay to receive the treasury on the issuance date. The price at which WIs trade is indeed a good predictor of what the auction-clearing coupon rate will be. But in the model, investors can condition their bids on the realized auction price p . Thus all information that is in the auction is already in their information sets. Pre-auction information from the WI market is a noisy signal about something they effectively know perfectly already. As such, it is redundant information and does not affect investors’ bids in our model.

E Calibration Robustness

In Section 3, we calibrated the full model to Treasury auction outcomes assuming that dealers share fully their average signal with their clients, so that $\tau_\xi = +\infty$, and that the 20 dealers do not share information with each other, so that $\psi = 0$. In this Appendix, we investigate alternative information sharing assumptions for the calibration. It is important to note that the parameters of the distribution of the fundamental value of the security, μ and τ_f , and the parameters of the non-competitive agent demand, \bar{x} and τ_x , remain fixed over the alternative calibrations as we can directly observe these moments from the data.

We consider three alternative assumptions for τ_ξ and N_P :

1. Full information sharing with clients ($\tau_\xi = +\infty$) and information sharing with one other dealer ($\psi = 1$).
2. No information sharing with clients ($\tau_\xi = 0$) and no information sharing with the other dealers ($\psi = 0$).
3. No information sharing with clients ($\tau_\xi = 0$) and information sharing with one other dealer ($\psi = 1$).

Though this is by no means an exhaustive list of possible assumptions for the true information sharing structure that occurs in practice, we find the differences in calibration outcomes to be representative of the changes to the calibrated parameters – and corresponding quantitative implications of the comparative statics conducted in the paper – under alternative assumptions.

We evaluate the alternative calibration assumptions along several dimensions. First, Table 5 reports the calibrated parameters under the baseline assumption used in the main body of the paper and the three alternatives. Comparing first the baseline calibration with alternative 1 ($\tau_\xi = +\infty$, $\psi = 1$), we see that the calibrated parameters are nearly identical, with slightly higher risk aversion for the large investors and dealers, lower signal precision of private signals, lower risk aversion for the small investors and higher minimum bidding requirement penalty. Comparing next the baseline calibration to alternative 2 ($\tau_\xi = 0$, $\psi = 0$), we see that the risk aversion of large investors and dealers decreases, precision of the private signals observed by large investors and dealers increases, while the calibration of the small investors moves in the opposite direction (high risk aversion and lower signal precision). Finally, alternative 3 ($\tau_\xi = 0$, $\psi = 1$) moves the calibrated parameters back toward the baseline calibration.

Table 5: Calibrated parameters μ , χ_0 , $\tau_f^{-\frac{1}{2}}$, $\tau_{\varepsilon,s}^{-\frac{1}{2}}$ and $\tau_{\varepsilon,l}^{-\frac{1}{2}}$ are all expressed in basis points.

	μ	$\tau_f^{-\frac{1}{2}}$	$\tau_{\varepsilon,s}^{-\frac{1}{2}}$	$\tau_{\varepsilon,l}^{-\frac{1}{2}}$	$\tau_x^{-\frac{1}{2}}$	\bar{x}	ρ	ρ_L	χ	N_S	N_L	N_D
Baseline	40.8	73.5	494.9	295.1	0.06	0.12	49402	1068	0.06	240	40	20
Info share with 1 dealer	40.8	73.5	491.1	295.4	0.06	0.12	49355	1079	0.08	240	40	20
No info sharing with clients	40.8	73.5	5386.7	107.5	0.06	0.12	99927	20	0.04	240	40	20
Info share with 1 dealer; no client info sharing	40.8	73.5	414.8	236.4	0.06	0.12	17534	874	0.06	240	40	20

Table 6 evaluates the goodness-of-fit of the four different assumptions on the information sharing structure. Comparing the four calibrations, we see that, while they all do fairly well in matching the moments of price (A , price sensitivity to fundamentals, C , and error standard deviation) and the expected excess revenue and volatility of excess revenue, the alternatives are less able to match average shares allotted. Indeed, if we compare the minimized objective of the calibration (the sum of squared deviations), we see that the baseline calibration achieves the lowest minimized objective.

Finally, in Table 7, we compare the full information sharing outcome and the Chinese wall outcome for the four calibrations, in terms of expected revenue, revenue volatility, expected utility of a small investor and expected utility of a dealer (or large investor). The full information outcome looks very similar for all four calibration alternatives, with the only exception the revenue volatility and agents' utility for alternative 2 ($\tau_\xi = 0$, $\psi = 0$), corresponding to the much lower risk aversion and much higher signal precision of the large investors and dealers in this calibration. Comparing the Chinese wall outcomes to the full information sharing outcomes, we see that alternative 2 has the lowest expected excess revenue gain from moving from no information sharing to full information sharing but the largest expected utility gain for both large and small investors. Even in this case, however, the Treasury would lose 2.7 bps per dollar of security auctioned (or \$2.1 billion annually)

Table 6: Calibration targets and model-implied values. Prices and excess revenues are all expressed in basis points.

	Data	Baseline	Info share with 1 dealer	No info sharing with clients	Info share with 1 dealer; No info sharing with clients	
Price sensitivity to fundamental	A 0.97	-7.51 0.91	-6.38 0.91	-28.54 0.92	-25.38 0.90	
Error Std. Dev.	C 29.72	73.57 23.12	64.81 22.73	240.03 17.10	223.46 19.29	
Expected excess revenue		37.18	38.73	38.65	38.64	39.10
Volatility of excess revenue		72.64	70.81	70.46	70.82	70.10
Volatility of indirect share	Indirect share	0.25	0.51	0.70	0.52	0.55
Volatility of dealer share	Dealer share	0.09	0.73	0.65	0.26	0.27
Volatility of dealer share	Dealer share	0.53	0.43	0.24	0.35	0.32
Volatility of direct share	Direct share	0.14	0.19	0.14	0.25	0.27
Volatility of direct share	Direct share	0.10	0.00	0.01	0.01	0.01
Volatility of direct share	Direct share	0.09	0.01	0.01	0.01	0.00
Sum of squared errors		-	19758.66	90853.26	58752.48	70498.48

in moving from full information sharing to no information sharing.

Table 7: Full information and Chinese Wall outcomes under different calibration assumptions. Prices and excess revenues are all expressed in basis points.

<i>A. Full information outcomes</i>				
	Expected revenue	Revenue volatility	Small client utility	Large agent utility
Baseline	39.987	79.881	-0.545	-0.545
Info share with 1 dealer	39.986	79.885	-0.544	-0.544
No info sharing with clients	40.767	343.439	-0.049	-0.049
Info share with 1 dealer; no client info sharing	40.367	78.115	-0.545	-0.546
<i>B. Chinese Wall outcomes</i>				
	Expected revenue	Revenue volatility	Small client utility	Large agent utility
Baseline	32.808	58.237	-1.041	-1.041
Info share with 1 dealer	32.781	58.208	-1.041	-1.041
No info sharing with clients	38.044	71.020	-1.097	-1.036
Info share with 1 dealer; no client info sharing	36.976	64.302	-1.085	-1.082