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

Nina Boyarchenko, David O. Lucca, and Laura Veldkamp

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

March 25, 2016

Abstract

The use of order flow information by financial firms has come to the forefront of the regulatory debate. Central to this discussion is whether a dealer who acquires information by taking client orders can share it with traders within the same firm, with other clients or other firms. We explore how information sharing affects dealers, clients and issuer revenues in U.S. Treasury auctions. Because one cannot observe outcomes of alternative information regimes, we build a model and use it construct counter-factuals. Calibrating the model to auction results data, we estimate that yearly auction revenues with full-information (sharing with clients, within and between dealers) would be about $5 billion higher than in a "Chinese Wall" regime in which no information is shared. When information sharing is associated with collusion, or group bidding, auction revenues decline but are higher than when no information is shared. Small investors are not necessarily hurt by increased information sharing, and their welfare depends on how information is shared. Surprisingly, investors benefit when dealers share information with each other, not when they share more with clients. In a mixed auction, where investors can opt out of intermediated trade and bid directly, information sharing creates a new financial accelerator: Only investors with bad news bid through intermediaries, who then share that information with others. Sharing amplifies the effect of negative news and generates revenue skewness, or a higher probability of "auction failure." Tests of two model predictions support two 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, Ken Garbade, Luis Gonzalez, Robin Greenwood, Pablo Kurlat, Richard Tang, Haoxiang Zhu and participants at the Chicago Booth Asset Pricing Conference, Duke, the 2014 U.S. Treasury Roundtable on Treasury Markets, NYU, the Spring 2015 Macro Finance Society Meetings, SED 2015, Gerzensee Summer Institute, EFA 2015, EEA 2015, Bank of Canada and SFI Lausanne for comments. Nic Kozeniauskas, Karen Shen 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. As noted in the Bloomberg (2015) article quoted above, recent investigations reportedly involve U.S. Treasury auctions. But the use of order flow information has been central to market making theory (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 information to inform their own strategy, share some of the information with clients, or

[1]Thus 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 only to study the economic effects of alternative information sharing arrangements.

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

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 was associated with coordinated trades (benchmark fixing) in foreign-exchange misconduct, 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' information is also shared 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 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 Following the discussion in Bikhchandani and Huang (1993), we assume an imperfect competitive setting. 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 (1985), 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, called "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] Section 4.4 models such a 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 how intermediation and information sharing are intertwined: it is the process of intermediating trades that reveals information to dealers and that empowers them. Information sharing is what induces clients to use intermediaries and induces large investors to intermediate.

Contribution to the existing literature. Our paper is related to four main 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 Parlatore (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. In our model, bidders can choose whether to bid through a dealer or to bid directly, circumventing the information pools. Dealers know that, in order to attract clients, they must share some of their information with them. Without this dealer-to-client information sharing, the primary dealer system could not operate in our model. Conversely, allowing dealers to extract information from order flow is equally important to induce them being subject to underwriting (minimum bidding) requirements, which are costly. These institutional details of our model are essential to understanding information sharing.

Primary dealers perform underwriting services and as such our paper is also related to 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 dealers share information, the conventional wisdom of underwriting is reversed: information intermediaries raise expected revenue but also revenue variance.

[3] See Appendix E in Brady et al. (1992) for pre-1992 policies. Prior 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. 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.

Finally, our paper is related to a literature studying how auction design affects revenues.[4] This literature complements our project, which fixes the auction format to a uniform-price menu auction and focuses on the effect of intermediation and information sharing.

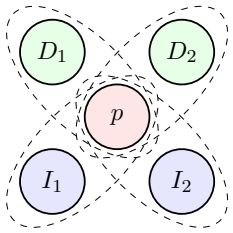
1 Baseline Auction Model with Primary Dealers

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. Dealers are denoted with the letter "D," while investors with the letter "I." Panel a) shows the case of no information sharing either between customers and dealers or between dealers (or, "Chinese walls"). In this case each auction participant only observes his private information y_i and conditions demand on the realized market-clearing price, which is a noisy signal about all agents' information. Optimal bids use this additional price information (represented by p in the figure), making the information set of investor i effectively $\{y_i, p\}$. 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 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 not only observes the information set pertaining to its dealer, but 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. In the model we consider a larger number of dealer and investors than depicted in Figure 1. In this more general setting, when a dealer shares information with its customers, investors not only learn from the dealer specific signal but also from the average signal of other customers that a dealer can infer from the bidding functions. In addition, we also consider intermediate cases of partial information sharing, in which case the order flow signal shared by a dealer with its customers or other dealers is only a noisy realization of the actual information a dealer observes.

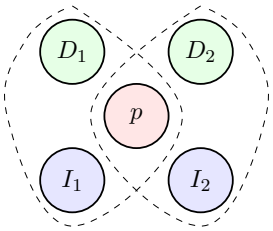
We consider four type of bidders to match key features of Treasury auction participation:

[4] Theoretical 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 et al. (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) and Hortaçsu and McAdams (2010) use French and Turkish Treasury auction bids to structurally estimate the benefits of uniform price auctions.

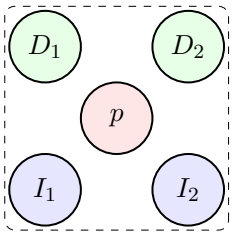
Figure 1: Information sets with alternative sharing assumptions. Letter D denotes dealers; letters I denotes investors (either large or small) bidding through a dealer or not (direct bidding); p is the equilibrium price. Dashed lines indicate common information sets.



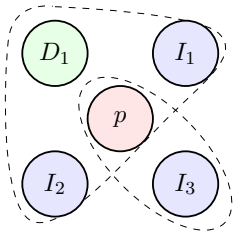
(a) No sharing with customers or dealers (Chinese Walls)



(b) Sharing with customers not with other dealers



(c) Sharing with customers and dealers



(d) Sharing with customers only; one direct bidder

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. 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, 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 acting as noise traders for other investors.

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: $f \sim N(\mu, \tau_f^{-1})$.

We index investors (small and large) and dealers with $i = 1, \ldots, N$, 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,

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

where ρ_i denotes agent i's coefficient of absolute risk aversion.[5] The budget constraint for small and large investors dictates that final wealth is $W_i = W_0 + q_i(f - p)$.

[5] We use exponential (CARA) utility here for all agents to keep the problem tractable. Of course, this rules out wealth effects on portfolio choices. At the same time, we want to capture the idea that investors with larger balance sheets naturally hold larger positions of risky assets. Therefore, we assign large investors and dealers a smaller absolute risk aversion. In other words, we capture wealth effects with differences in risk aversion.

Dealer and large investors solve the same problem of small investors but they also internalize 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 large investors and dealers share the same risk aversion and signal precision within each agent type: $\rho_i = \rho_l$ and $\tau_{\varepsilon,i} = \tau_{\varepsilon,l}$ for $i \in \{\mathcal{N}_L, \mathcal{N}_D\}$, where with $\mathcal{N}_j$ we denote 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 any given auction, no set minimum bid. 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 $\chi_0 - \chi qp$, or equivalently a benefit of $\chi qp - \chi_0$. Thus, a dealer d's budget constraint is

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

Describing Information Sets and Updating Beliefs with Correlated Signals Depending on the information structure, investors and dealers may 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 can be decomposed into two components. A component that depends on the value of outstanding (previously issued) Treasury securities, as implied by the shape of the yield curve, and the specific liquidity (or relative) value of the newly auctioned asset relative to that value. In the model calibration, we focus on how information sharing affects the specific liquidity value taking as given the shape of the yield curve, which is largely driven by monetary policy.

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 $N_I/N_D \equiv \nu_I$ small and $N_L/N_D \equiv \nu_l$ 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:

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

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 ones' 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. But because the signals agents observe have correlated signal errors, we need to use a procedure that adjusts for this correlation. To update beliefs, we use the following optimal linear

projection formulas[6]

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

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

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

where m is the number of signals in the vector S_j, the covariance vector is $\mathbf{Cov}\left(f, S_j\right) = 1_m \tau_f^{-1}$ and the signal variance-covariance $\mathbb{V}\left(S_j\right)$, is worked out in the appendix. The vector $\beta_j = [\beta_{sj},\ \beta_{\xi j},\ \beta_{pj}]$ 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, these are the OLS coefficients that multiply the independent variables to forecast the dependent variable – in this case, the payoff f. 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

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

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

2. A menu of price-quantity pairs bid by each strategic player (dealer and large investor) that maximizes

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

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

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

   The second constraint is the market clearing condition and reflects that the strategic players must choose their quantity and the price so that the market clears.

[6]Note that the following formulas are just like the OLS formulas in a context where all the means, variances and covariances of variables are known. The OLS additive constant α is $(1_N - \beta)' 1_N \mu$. 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 systems, both OLS and Bayesian estimators are consistent. It also has the Bayesian interpretation as a weighted average of normal priors and signals, where each is weighted by their relative precision. Here that precision is adjusted to account for correlation.

3. 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

To solve 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 the following set of price-quantity pairs:

$$q_i(p) = \frac{1}{\rho}\mathbb{V}[f|S_i]^{-1}\left(\mathbb{E}[f|S_i] - p\right). \tag{10}$$

This 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*. The first order condition with respect to q_j reveals that dealers and large investors bid

$$q_j\left(p\right) = \frac{\mathbb{E}[f|S_j] - p}{\rho_l \mathbb{V}[f|S_j] + dp/dq_j}. \tag{11}$$

Importantly, this expression differs from equation (10) 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 fact 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

$$q_D(p) = \frac{\mathbb{E}[f|X_D S] - p(1-\chi)}{\rho_l \mathbb{V}[f|X_D S] + (1-\chi)dp/dq_D} \tag{12}$$

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 the bank's own purchases or anyone else's. 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.[7] In our model with Chinese walls, each agent sees only their own private signal. Formally, the investors' information set contains only their 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 (or noise trader demand) 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 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.*

[7] For example, the Financial Stability Board (FSB) 2014 report on "Foreign Exchange Benchmarks"

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, non-competitive trades x, and dealer signal noise ξ_d:

$$p = A + B_I\,\bar{s}_I + B_L\,\bar{s}_{NL} + B_D\,\bar{s}_{ND} + Cx + D\,\bar{\xi}_d \tag{13}$$

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 model typically have simple closed form 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) information sharing makes an investor's signal more correlated with the signal from the realized auction price that bidders condition on. 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 that is raises 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 result should not surprise anyone familiar with imperfect information asset pricing models, it forms the basis for understanding the more subtle results that arise with different types of information sharing. It is also a point that 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, 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 and compare to other primary dealer arrangements. 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 $B_I^2\tau_{\epsilon I}^{-1}/N_I + B_L^2\tau_{\epsilon L}^{-1}/N_L + B_D^2\tau_{\epsilon D}^{-1}/N_D + C^2\tau_x^{-1}$.

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.[8] 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 (value) to purchase at the market-clearing rate. Each bidder can only place a single non-competitive bid with a maximum size of $5 million. 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 registered according to Section 15C(a)(1) of the Securities Exchange Act. In practice, though, most indirect bids are placed through primary dealers, with the exception of bids 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

[8] Treasury 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.

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 FIMA competitive bids placed through the NY Fed.[9] We apply a simple calculation to reclassify these bids as part of the noise trader group (non-competitive), as they do not reflect the short-term issue-specific value (more on this in the next paragraph) because foreign official investors have long investment horizons invest reflecting foreign exchange strategies and these bids cannot provide order flow information to primary dealers or other participants because these bids are placed through the NY Fed. 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.[10]

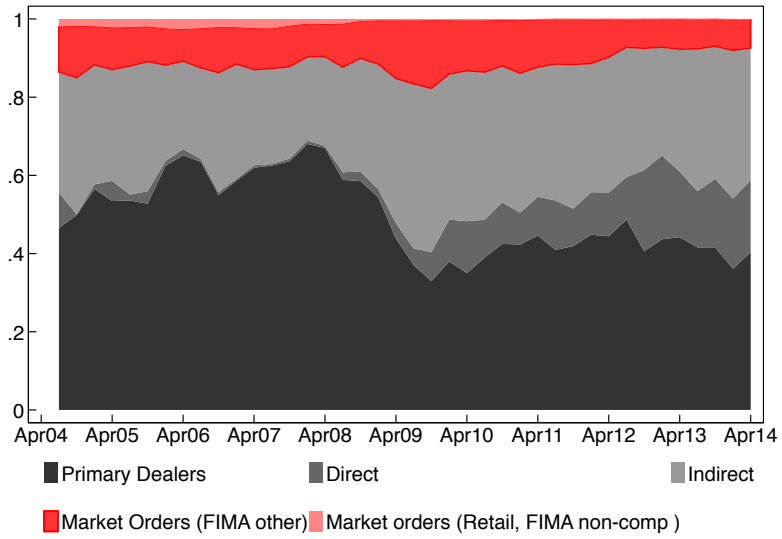
Figure 2 reports Treasury auction participation by type of bids over time and the actual target moments (first and second) are shown in Table 1. As shown by the dark grey area, primary dealers bidding for their own account, are the largest bidder category at auctions accounting for about half of all 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).

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

[9]FIMA 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.

[10]For 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.



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).

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 (the model of Section 4.4) to aggregate moments. The full model differs from the one presented in the previous section for the inclusion of minimum bidding requirements for dealers, which, as we discuss in the next sections, are key features 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 (13), 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.[11] 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

The final parameter, τ_ξ regulates how much information dealers share with clients. Without micro data, we cannot infer this value. 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 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

[11] We 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.

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 and the variance of that 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 ($\chi = 0$) to start and the return to examine its effect in section 4.4.

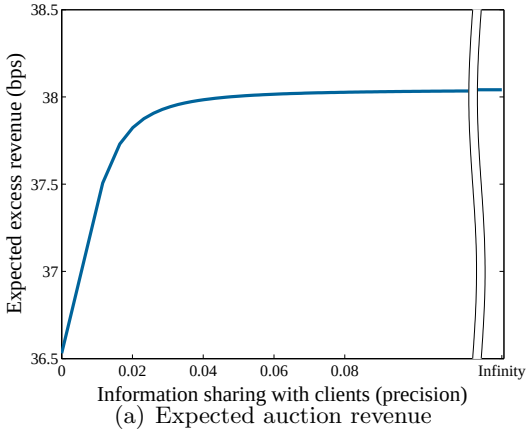
4.1 Information Sharing and Auction Revenue

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 revenue by 1.2 basis points.

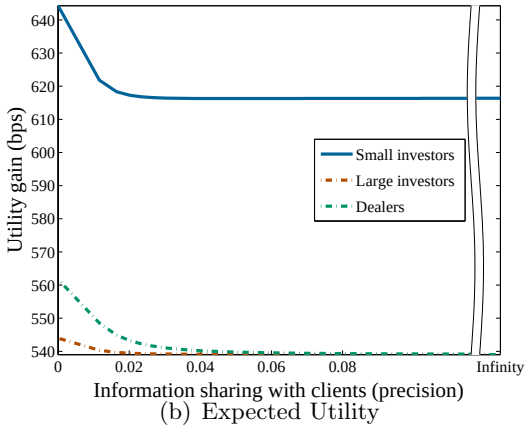
Information sharing makes investors better informed which in turn makes Treasuries less risky to investors, eliciting stronger bids and increasing auction revenues. Applying this estimated effect to annual Treasury issuance (or about $ 8 trillion), the model implies that total auction revenues would increase about $ 960 million when going from no dealer sharing with customers to perfect sharing with them.

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

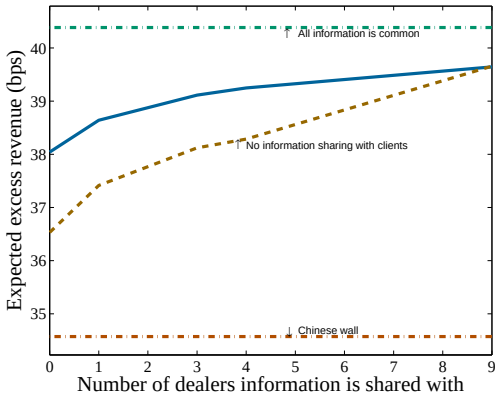
Figure 3: Dealer Information Sharing.Top row: dealer information sharing with clients; bottom row: dealer information sharing with other dealers. In the top row, the horizontal axis shows the precision of the dealer signal τ_ξ from zero (no information sharing) to infinity (perfect information sharing).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



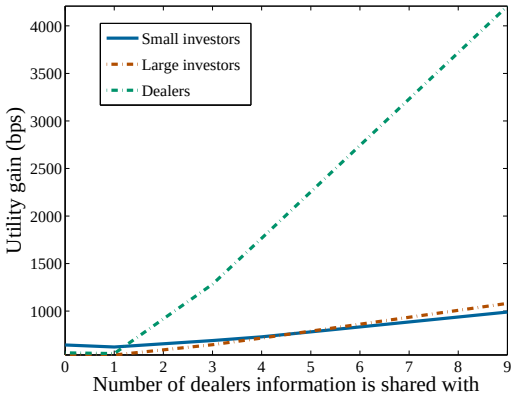
(a) Expected auction revenue



(b) Expected Utility



(c) Expected auction revenue



(d) Expected Utility

more valuable in expected auction revenue terms.

Finally, 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, it also raises auction revenues (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.[12] 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.

The variability of auction revenue also rises with inter-dealer information sharing, just as with client information sharing. The magnitude of this effect is quite small. Sharing with one versus 5 other dealers increases the standard deviation of revenue by a couple of basis points.

4.2 Client and dealer information sharing and their opposite effects on risk-sharing and utility

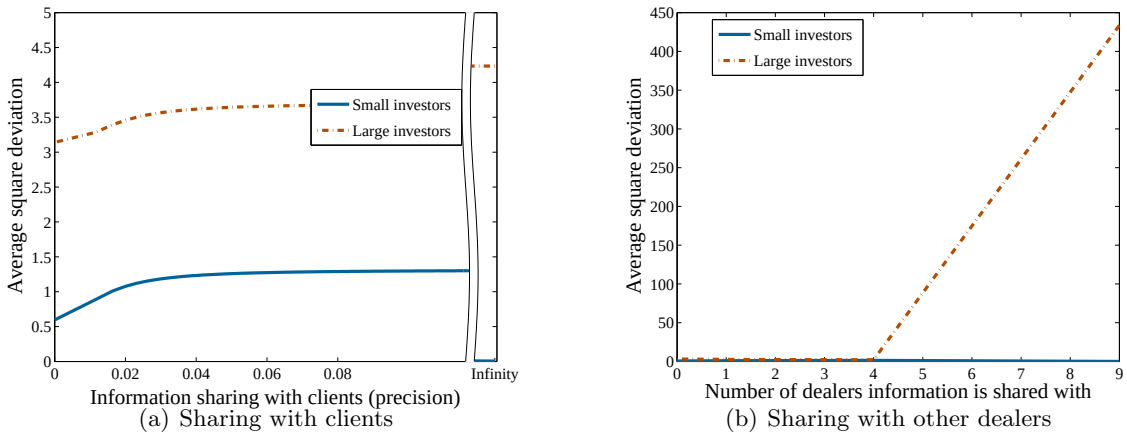
A key insight of our model is how client and dealer information sharing differ. While both types of information sharing reduce uncertainty and increase auction revenue, client and dealer information sharing have opposite effects on investor utility. The reason for this opposite effect lies in how each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

[12] Since 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.

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. When investors get the more precise dealer's signal, they weigh it more heavily in their beliefs. Since different dealers transmit different signals, this makes investors' beliefs quite different from each other. This increase in information asymmetry will make ex-ante similar investors hold different amounts of securities ex-post. This asymmetric information effect pushes the asset allocation further away from the efficient diversified benchmark. The consequent reduction in risk sharing reduces utility.

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: $\bar{q}_j=(1/N_j)\sum_{i\in\mathcal{N}_j}q_i$.



(a) Sharing with clients

(b) Sharing with other dealers

Figure 4 illustrates this key force. It shows how the two types of information sharing affect

information asymmetry in different ways. Differences in beliefs show up as differences in treasury allocations. What the plots shows is 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.

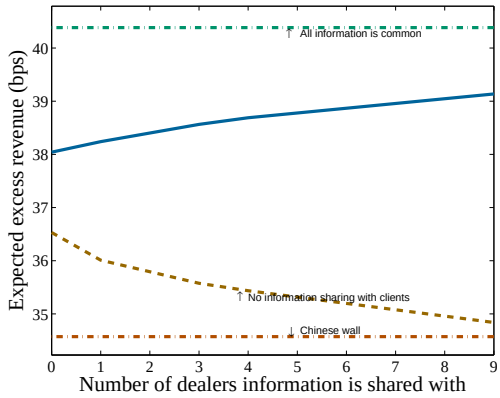
The result that client information increases information asymmetry and thereby reduces utility is the same force that is at work in Hirshleifer (1971). What is new is that sharing information between dealers makes agents better informed, but in a way that makes information more symmetric. This has the opposite effect on utility. Dealer information sharing is more like giving investors more public information. The results here are the converse 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) is welfare-improving.

4.3 What if Information Sharing Enabled Collusion?

One of the reasons that information-sharing raises 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 as well. That reduces revenue. But notice that the remedy of imposing Chinese walls reduces revenue by far more than collusion.

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.



4.4 The Effect of Minimum Bidding Requirements

Primary dealers are required to be consistent, active participants in Treasury auctions. We include minimum bidding penalties because they are realistic, but also because they help us to calibrate the model in a sensible way. Without them, it is hard to explain why dealers bid for so much of the auction. However, our results reveal that low bid penalties do not change our main findings. (See Appendix C.1 for more detailed results.) In most cases, they raise revenue, since dealers are incentivized to bid more aggressively. Bidding requirements leave the effect of client information-sharing on revenue and utility unchanged. The one prediction that does change is that, with bidding requirements, dealer information sharing has a non-monotonic effect on revenue, even without collusion. The reason is that bidding requirements make dealers less responsive to changes in price. That means that prices have to move a lot 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. Therefore, there is less scope for precision improvements through dealer information-sharing.

5 Mixed Auctions: Choosing Direct or Indirect Bidding

A key distinguishing feature of U.S. Treasury auctions is that they are mixed auctions in the sense that non-dealer investors can either place intermediated bids through primary

dealers, or bid directly.[13] The option to bid directly is important in our setting because it limits how much information dealers can share. Dealers who share too much of the clients' information with others or share too little information with their clients will find themselves without clients to observe. Intermediation choice also affects auction revenue and amplifies the effect of low signals on auction revenue.

5.1 A Model with Intermediation Choice

To explore these effects in a simple setting, we now allow one investor to choose 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 (through an intermediary) only affects the information structure of that investor, its dealer, and other investors bidding with that same dealer. By affecting the information sets of these investors, the choice to bid directly affects the information content of the price $s(p)$ as well.

When the large investor chooses to bid directly on his own behalf, the first dealer's signal is the average of the first ν_I investors', the first $\nu_l - 1$ large investors' and the first dealer's signals:

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

As in the previous model, investor i who bids through intermediary d observes signals $X_i S = [s_i, \bar{s}_d, s(p)]$. The large investor bidding directly observes only his own signal and the price information: $X_L S = [s_L, s(p)]$.

Solution: Auction Outcomes Solving this model introduces a technical challenge. If the large investor always bids directly or indirectly, then the nature of the solution would be as before. But more generally 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 shares this information (as well as all other information) with his clients. The technical difficulty is that the inferred signal about auction participation decision is not normal, but instead binomial. If the large investor bids indirectly, this is not a problem because the dealer can infer exactly what the large investor's signal

[13] While 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.

is and report that to his clients. But if the large investor bids directly, the dealer does not find out the large investor's signal and only knows that it is above the threshold that would have made the investor report in equilibrium. The information that has been revealed is that a normal variable (the large investor's signal) lies in some truncated region 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 an equilibrium model like this that is already complex to solve would make it intractable.

We solve this problem by constructing a tractable approximate signal that agents can use to update their beliefs as follows. 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, using 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 C 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 $p = A + B_{d1}\bar{s}_1 + B_{d\neq 1}/(N_D - 1)\sum_{d=2}^{N_D}\bar{s}_d + C(N_D, 1)x + B_L s_L$, where the coefficient formulas are in reported Appendix.*

5.2 Understanding Intermediation Choice: Why Is Bad News Revealed?

When an investor bids directly, no one observes their order flow and their information remains private. When they bid through an intermediary, they share their signal realization with the intermediary through their bidding behavior but also learn information about the overall order flow observed by the intermediary. 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 also benefits from learning new information from other investors. Thus, low-signal investors are more likely to bid indirectly through the dealer. But 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

$$EU(l) = -\exp(\rho_L W_L)(1 + 2\theta_l \Delta\mathbb{V}_l)^{-\frac{1}{2}} \exp\left(-\frac{\mu_{rl}^2}{\theta_l^{-1} + 2\Delta\mathbb{V}_l}\right). \tag{14}$$

The intermediation decision affects utility in three ways: through the expected profit per unit 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:[14]

$$\mu_{rl} \equiv \mathbb{E}\{\mathbb{E}[f|X_l S] - p|s_i\}, \tag{15}$$

$$\theta_l \equiv \rho_L[\rho_L \mathbb{V}[f|S_L] + dp/dq_L]^{-1}\Big(1 - \frac{1}{2}\rho_L[\rho_L \mathbb{V}[f|S_L] + dp/dq_L]^{-1}\mathbb{V}[f|S_L]\Big), \tag{16}$$

$$\Delta\mathbb{V}_l \equiv \mathbb{V}\{\mathbb{E}[f|X_l S] - p|s_i\} = \mathbb{V}[f - p|s_i] - \mathbb{V}[f|X_l S]. \tag{17}$$

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. When the large investor bids through a dealer, their information set and their expectations will be the same as for every other agent who bids through that dealer $d1$: $E[f|X_l S] = E[f|X_{d1} S]$. Information sharing reduces μ_{rl} for two reasons. First, since many investors all condition their bids on this 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 raises the expected price p, which in turn, lowers μ_{rl} (see eq (15)). In equation (14), a decrease in μ_{rl} decreases expected utility because $\theta_l > 0$ (see Appendix).

The second term θ_l captures the main advantage of intermediation: Dealers give their

[14] See the Appendix for derivations of the following three equations and support for the analysis that follows.

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 show that θ_l is positive and strictly decreasing in the posterior variance of the asset payoff $\mathbb{V}[f|X_lS]$. Thus, intermediation improves the investor's information, which decreases variance $\mathbb{V}[f|X_lS]$, 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.[15]

A zero-profit signal (bad news) is always shared. Note from equation (14) that as $\mu_{rl} \to 0$, 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 $\mathbb{E}[f|X_lS] - p$ 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 larger when μ_{rl} is closer to zero. We use these three effects to understand the intermediation choice results below.

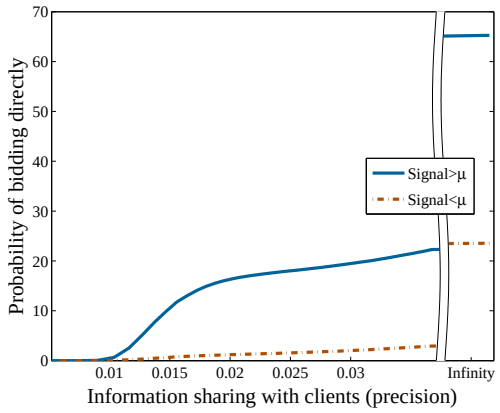
5.3 Quantitative Results: Intermediation Choice

Figure 6 confirms the result from the previous section that a large investor with a high signal always bids directly, while an investor with a low signal may or may not choose to bid through a dealer. It also teaches us that increasing the precision of client information sharing by 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.

The middle panel of Figure 6 also illustrates that intermediation choice changes the relationship between information sharing and auction revenues. Whereas before revenues increased monotonically with information precision, now they first increase and then decrease. The increase happens for the same reason as before: More information makes the asset less risky and encourages strong bidding, which pushes up the price. Revenue de-

[15] When the large investor trades through a dealer, his uncertainty $\mathbb{V}[f|X_lS]$ declines. From equation (17) 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 (14), 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.

Figure 6: Mixed auctions: Probability of bidding directly, conditional on signal realization.



creases in signal precision in the region where additional information sharing with clients makes the large investor less likely to bid through the dealer. When the large investor bids directly, he is more uncertain and the other clients who would have seen his signal are more uncertain. This uncertainty restrains demand for the asset and results in a lower average price. As in the previous models, expected auction revenue and its standard deviation move in lock step.

However, this time, most agents' utility increases with information sharing. This is in contrast to the case with no intermediation choice where information sharing accentuated asymmetric information and reduced risk sharing. In this model, more information sharing reduces the probability that the large investor, "the decider", will choose to bid through the dealer. For the large bidder, the cost of bidding through the dealer is that his information is then shared with others. When more information is shared, this cost rises and the probability of indirect bidding falls (panel a). If the large bidder bids directly, dealer 1 and his clients get less precise information. This information loss of dealer 1 makes all the other dealers and their clients better off because they bid against less informed competitors.

A New Financial Accelerator. The result about direct versus indirect bidding decision, reveals a new channel through which intermediation can amplify shocks. When signals about the value of a financial asset are negative, this information is more likely to be shared with a dealer and his clients. Positive signals 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.

6 Testing the Model Predictions

In this section we test two implications of the model. We first verify our assumption on the noise trader group. We then show that post-auction appreciation displays positive skewness as implied by intermediation choice model specification.

6.1 Test 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. The hallmark of informed trading is that such trades, as opposed to uninformed ones, forecast profits. An uninformed agent only having information about what is revealed in equilibrium prices, 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. That's 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 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 and, if auction size is correlated with asset characteristics, could potentially bias our estimates.

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 ?? 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 et al. (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

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***	0.0147***
		[0.0031]	[0.0043]
Const	0.0003**	0.0017***	0.0018***
	[0.0001]	[0.0005]	[0.0005]
Adj. R2	0.00	0.05	0.23
Obs.	494	494	494
Month FEs?	No	No	Yes
Tenor FEs?	No	No	Yes

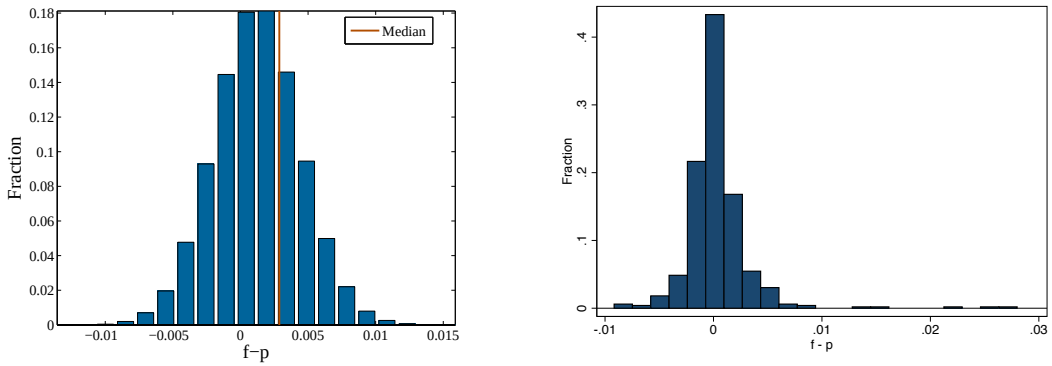
traders demand is high, the price is lower on average, relative to the payoff. It is that low price relative to fundamental value that induces informed investors to buy more. We run analogous regressions using data generated from the model and find similar results. This result does not prove that primary dealers aggregate this information. But it does support the notion that the dealers' clients have some private information to be aggregated.

6.2 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 $f-p$. The first panel of Table 4 reveals that the unconditional skewness in p is -1.5%. 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. That one agent alone can only generate limited skewness is aggregate revenue If we could compute

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)



(a) Model (b) Data

a model with many agents making an intermediation choice, this effect would likely be amplified.

7 Conclusions

Recent instances of market abuse involving access 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. The model allows us to run counter-factuals and policy evaluation. We find that, with more information sharing auction revenues increase, as bidders are better informed, but also revenue volatility increases, as investors are more sensitive to the information that they observe. Investors' welfare depends on how information is shared. Surprisingly, we find that investors, as a class, are worse off when dealers share more information with them, but are better off when the dealers share information among themselves. The model helps explain why 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. One 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 Hirschliefer'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 financial markets, 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. 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 what exactly the role of intermediation in a given market is. 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, 1411–1425.

Armantier, O. and E. Sbaï (2006): "Estimation and comparison of Treasury auction formats when bidders are asymmetric," *Journal of Applied Econometrics*, 21, 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, 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, 311–339.

——— (1993): "The economics of Treasury securities markets," *The Journal of Economic Perspectives*, 117–134.

Brady, N., R. Breeden, and A. Greenspan (1992): "Joint report on the government securities market," *US Government Printing Office*.

Chari, V. V. and R. J. Weber (1992): "How the U.S. Treasury Should Auction Its Debt," *Federal Reserve Bank of Minneapolis*, 16, 1–8.

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.

Gürkaynak, R. S., B. Sack, and J. H. Wright (2007): "The US Treasury yield curve: 1961 to the present," *Journal of Monetary Economics*, 54, 2291–2304.

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, 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, 833–865.

Krishnamurthy, A. (2002): "The bond/old-bond spread," *Journal of Financial Economics*, 66, 463–506.

Kyle, A. S. (1985): "Continuous auctions and insider trading," *Econometrica: Journal of the Econometric Society*, 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, 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, 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, 1795–1828.

Roosa, R. V. (1956): *Federal Reserve Operations in the Money and Government Securities Markets*, vol. 332, Federal Reserve Bank of New York.

Svensson, L. E. (1994): "Estimating and interpreting forward interest rates: Sweden 1992-1994," Tech. rep., National Bureau of Economic Research.

Wilson, R. (1979): "Auctions of shares," *Quarterly Journal of Economics*, 93, 675–689.

A Appendix: Derivations and Proofs

A.1 Auction with Dealers and Imperfect Information-Sharing

When dealers reveal a noisy signal of average order flow $\bar{s}_d + \xi_d$, their clients' private signal becomes relevant because it is not contaminated with the noise ξ_d. Therefore, for investor i, $S_i = [s_i, \bar{s}_d + \xi_i, s(p)]$.

These three signals for investor j load on the following orthogonal shocks

$$Z = \left[\begin{array}{cccccccc} \varepsilon_1 & \dots & \varepsilon_N & \xi_1 & \dots & \xi_N & x \end{array} \right]'. \tag{18}$$

with loadings Π_i for any non-dealer investor i:

$$\Pi_i = \left[\begin{array}{ccc} \iota_i & 0 & 0 \\ \iota'_{d(i)} \otimes 1_{N/N_D} N_D/N & \iota_i & 0 \\ B/(B+F)1_{N_I} & F/(B+F)1_{N_L+N_D} & C/(B+F) \end{array} \right] \tag{19}$$

where ι_i is a $1 \times N$ vector of zeros with a 1 in the ith position, $\iota_{d(i)}$ is a $1 \times N_D$ vector of zeros with a 1 in the $d(i)$th position, and $d(i)$ is the dealer of investor i.

The variance matrix of the Z shocks is

$$\mathbb{V}[Z] = diag([\tau_{\epsilon,s}^{-1} 1_{N_I},\ \tau_{\epsilon,l}^{-1} 1_{N_L+N_D},\ \tau_{\xi}^{-1} 1_N,\ \tau_x^{-1}]). \tag{20}$$

Given this Π matrix and $V[Z]$, we can construct $V(S_i) = \tau_f^{-1} + \Pi_i V[Z] \Pi_i'$.

Given the signal variance-covariance matrix for dealers and investors, the Bayesian updating weights are given by (35). For non-dealers, there are three Bayesian updating weights $\beta = [\beta_s, \beta_{\bar{s}}, \beta_p]'$. Using (4) and (6), we find posterior beliefs

$$\mathbb{E}[f|S_j] = (1 - \beta_{\bar{s}j} - \beta_{pj})\,\mu + \beta_{\bar{s}j} \bar{s}_{d(j)} + \beta_{pj} \frac{p - A - C\bar{x}}{B+F}, \tag{21}$$

with conditional variance $\mathbb{V}[f|S_j] = \tau_f^{-1} - \tau_f^{-1}(\beta_{sj} + \beta_{\bar{s}j} + \beta_{pj}) \equiv \hat{\tau}_j^{-1}$.

The dealers in this setting are slightly different because they are the one set of agents who observe their order flow noiselessly. For dealers, there are only two relevant signals $S_d[\bar{s}_d, s(p)]$ where

$$\Pi_d = \left[\begin{array}{ccc} \iota_d \otimes 1_{N/N_D} N_D/N & 0 & 0 \\ B/(B+F)1_{N_I} & F/(B+F)1_{N_L+N_D} & C/(B+F) \end{array} \right] \tag{22}$$

are the loadings on the same Z shock vector as for investors. Thus, for dealers, $V(S_d) = \tau_f^{-1} + \Pi_d V[Z] \Pi_d'$, which allows us to compute updating weights β_d and posterior means and variances. For dealers, since there are only two relevant signals, there are also only two Bayesian weights: $\beta_D = [\beta_{\bar{s}D}, \beta_{pD}]'$. Given these weights, conditional expectations and variances are computed with the same formula as above.

Proof of Result 1 For large investors and dealers, let the sensitivity of asset demand to changes in the expected profit per share be denoted $M_L = V[f|S_L] + dp/dq_L$ and $M_D = V[f|S_D] + dp/dq_D$.

The asset demand for each small investors is (10) and each large investor is (36). We equate total demand to supply (1) $\int q_i di + x = 1$ and substitute in for all demands, except for that of one large investor q_L to get

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

Then, we substitute in for the conditional expectations and use the implicit function theorem to calculate dp/dq_L, which we need to determine M_L. Substituting for the expectations yields

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

Taking the derivative with respect to q_L, we obtain

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

so that we can solve for dp/dq_L and express M_L implicitly as

$$M_L^{-1} = \rho_L\hat{\tau}_L^{-1} - \left(N_I\hat{\tau}_I\rho^{-1}\frac{\beta_{pi} - \tilde{B}}{\tilde{B}} + (N_L - 1)\,M_L\frac{\beta_{pL} - \tilde{B}}{\tilde{B}} + N_D M_D\frac{\beta_{pd} - \tilde{B}}{\tilde{B}}\right)^{-1}. \quad (23)$$

To determine dealers' demand $q_d = M_D(\mathbb{E}[f|S_d] - p)$, we need to solve for M_D in the same fashion. The solution is identical, except that the dealers' signal precision is higher and therefore, his posterior variance $\hat{\tau}^{-1}$ is lower. Following the same steps, we find that

$$M_D^{-1} = \rho_L\hat{\tau}_D^{-1} - \left(N_I\hat{\tau}_I\rho^{-1}\frac{\beta_{pi} - \tilde{B}}{\tilde{B}} + N_L M_L\frac{\beta_{pL} - \tilde{B}}{\tilde{B}} + (N_D - 1)M_D\frac{\beta_{pd} - \tilde{B}}{\tilde{B}}\right)^{-1}. \quad (24)$$

Finally, we can express the price coefficients as functions of these M's. Do to this, we substitute all the investors' demands in the market-clearing condition (same expression as before, without the q_L term and summing from $j = 1$ to N_L, instead of $N_L - 1$). Coefficient matching yields the following price coefficients:

$$A = C\left[N_I\hat{\tau}_I\rho^{-1}(1-\beta_{sI}-\beta_{\bar{s}I}-\beta_{pI})\mu + N_L M_L(1-\beta_{sL}-\beta_{\bar{s}L}-\beta_{pL})\mu \right.$$
$$\left. + N_D M_D(1-\beta_{sD}-\beta_{\bar{s}D}-\beta_{pD})\mu - \tilde{B}^{-1}(A+C\bar{x})(\beta_{pI}+\beta_{pL}+\beta_{pD}) - 1\right] \quad (25)$$

$$B_I = CN_I\left(\hat{\tau}_I\rho^{-1}\beta_{sI} + M_D\beta_{\bar{s}d}\omega_I\right) + N_I D\omega_I \quad (26)$$

$$B_L = CN_L\left(M_L\beta_{sL} + M_D\beta_{\bar{s}d}\omega_L\right) + N_L D\omega_L \quad (27)$$

$$B_D = CN_D M_D\beta_{\bar{s}d}\omega_D + N_D D\omega_D \quad (28)$$

$$C = -\left(N_I\hat{\tau}_I\rho^{-1}\frac{\beta_{pI}-\tilde{B}}{\tilde{B}} + N_L M_L\frac{\beta_{pL}-\tilde{B}}{\tilde{B}} + N_D M_D\frac{\beta_{pd}-\tilde{B}}{\tilde{B}}\right)^{-1} \quad (29)$$

$$D = C\left(\hat{\tau}_I\rho^{-1}\frac{N_I}{N_D}\beta_{\bar{s}I} + M_L\frac{N_L}{N_D}\beta_{\bar{s}L}\right) \quad (30)$$

The solution to this model is the set of β's for each type of agent, price coefficients, and M coefficients that jointly solve (35), (25)-(30), (23) and (24).

Proof of Corollary 1: Noncompetitive share reduces expected profit. We can break $\partial E[f-p]/\partial x$ into the difference $\partial E[f]/\partial x - \partial E[p]/\partial x$. The first term, $\partial E[f]/\partial x = 0$, by assumption. Both f and x are assumed to be exogenous, independent variables. Using the equilibrium price fomula, the second term is $\partial E[A + Bf + Cx]/\partial x = C$. Finally, because $C > 0$, $\partial E[f-p]/\partial x > 0$ and the result holds.

A.2 Price with Dealer Information Sharing

Proof of result 3:

A.3 Auction with Chinese Walls

In this model, intermediaries do not use or disclose any aggregated information. It is as if each type of limit order investor submits bids on their own behalf rather than through intermediaries. Each investor's information set is a 2×1 vector $S_i = [s_i, s(p)]$ containing their private signal s_i and the price information $s(p)$, but not an average signal provided by an intermediary.

Given the price formula in Result 3, an investor's unbiased price signal is

$$s(p) = \frac{p - A - C\bar{x}}{\tilde{B}} = \frac{B_I}{\tilde{B}}N_I^{-1}\sum_{i=1}^{N_I} s_i + \frac{B_L}{\tilde{B}}N_L^{-1}\sum_{j=1}^{N_L} s_j + \frac{C}{\tilde{B}}(x-\bar{x})$$

where $\tilde{B} \equiv B_I + B_L$. Since $(x-\bar{x})$ is mean zero, $s(p)$ is an unbiased signal about f. However, the price signal and the private signals have correlated signal errors. The Bayesian updating weights β come from optimal linear projection formulas that correct for this covariance.

Let Z be the vector of orthogonal shocks.

$$Z = \left[\begin{array}{cccc} \varepsilon_1 & \dots & \varepsilon_N & x \end{array} \right]'. \tag{31}$$

where we order agents by putting small, individual investors first, then large investors, then dealers. The variance matrix of this vector is

$$\mathbb{V}[Z] = diag([\tau_{\epsilon,I}^{-1} 1_{N_I},\ \tau_{\epsilon,l}^{-1} 1_{N_L},\ \tau_x^{-1}]). \tag{32}$$

Then, we can represent the vector of all m signals observed by investor i as

$$S_i = 1_m f + \Pi_i Z,$$

where 1_m is an $m \times 1$ vector of ones. The loading matrix Π_i of signals on orthogonal shocks is

$$\Pi_i = \left[\begin{array}{ccc} & \iota_i & \\ (B_I/\tilde{B})\, 1_{N_I} & (B_L/\tilde{B})\, 1_{N_L} & C/\tilde{B} \end{array} \right] \tag{33}$$

where ι_i is a vector of zeros with a 1 in the ith position, 1_N is a $1 \times N$ vector of 1's, N is the total number of agents ($N = N_I + N_L + N_D$, all exogenous), and B_I, B_L and C are equilibrium pricing coefficients (endogenous).

With this Π_i matrix, we can express the Bayesian updating weights β that small and large investors s, l put on their private signal and the price signal. With the β's, we can then easily compute posterior precision $\hat{\tau}$. The variance-covariance matrix of signals is

$$\mathbb{V}(S_i) = \tau_f^{-1} + \Pi_i V[Z] \Pi_i' \tag{34}$$

The 2×1 vector of Bayesian updating weights $\beta_i = [\beta_{si}, \beta_{pi}]'$ is

$$\beta_i = \mathbb{V}(S_i)^{-1} 1_m \tau_f^{-1}. \tag{35}$$

Let $\beta_I = [\beta_{sI},\ \beta_{pI}]$ be the vector of weights used by the small, individual investors and $\beta_L = [\beta_{sL},\ \beta_{pL}]$ be the vector of weights used by the large investors. Then, using (4) and (6), the posterior beliefs of a small investor are

$$\mathbb{E}[f|S_i] = (1 - \beta_{sI} - \beta_{pI})\, \mu + \beta_{sI} s_i + \beta_{pI} \frac{p - A - C\bar{x}}{\tilde{B}}$$

with conditional variance $\mathbb{V}[f|S_i] = \tau_f^{-1} - \tau_f^{-1} (\beta_{sI} + \beta_{pI}) \equiv \hat{\tau}_I^{-1}$. The beliefs of a large investor are

$$\mathbb{E}[f|S_i] = (1 - \beta_{sL} - \beta_{pL})\, \mu + \beta_{sL} s_i + \beta_{pL} \frac{p - A - C\bar{x}}{\tilde{B}}.$$

with conditional variance $\mathbb{V}[f|S_i] = \tau_f^{-1} - \tau_f^{-1} (\beta_{sL} + \beta_{pL}) \equiv \hat{\tau}_L^{-1}$.

Proof of Result 3: Price in auction with Chinese walls. The small investors' demand functions (bids) in the competitive market model are given by equation (10) and the large investors' bids by equation (36). In this and future results, it is convenient to

re-express the large investors' and dealers' demand as

$$q_j(p) = M_L\left(\mathbb{E}[f|X_j S] - p\right) \qquad \text{where} \tag{36}$$

$$M_l \equiv \left[\rho_l \mathbb{V}[f|S_l] + dp/dq_l\right]^{-1} \qquad \text{for any large player } l \tag{37}$$

where M_l describes the sensitivity of l's treasury demand to changes in expected profit.

Substituting these bids into the market clearing condition $x + \sum_{i=1}^{N} q_i = 1$ yields

$$x + \rho^{-1}\hat{\tau}_I \sum_{i=1}^{N_I}(\mathbb{E}[f|X_i, S] - p) + M_L \sum_{j=1}^{N_L - 1}\left(\mathbb{E}[f|X_j, S] - p\right) + q_L = 1.$$

Substituting for $\mathbb{E}[f|X_i, S]$, we obtain

$$\begin{aligned} 1 = x &+ \rho^{-1}\hat{\tau}_I \sum_{i=1}^{N_I}\left(\left(1 - \beta_{sI} - \beta_{pI}\right)\mu + \beta_{sI} s_i + \beta_{pI}\frac{p - A - C\bar{x}}{\tilde{B}} - p\right) \\ &+ M_L \sum_{j=1}^{N_L - 1}\left(\left(1 - \beta_{sL} - \beta_{pL}\right)\mu + \beta_{sL} s_i + \beta_{pL}\frac{p - A - C\bar{x}}{\tilde{B}} - p\right) + q_L. \end{aligned}$$

Taking the derivative with respect to q_L, we obtain

$$0 = \left(N_I \hat{\tau}_I \rho^{-1}\frac{\beta_{pI} - \tilde{B}}{\tilde{B}} + \left(N_L - 1\right) M_L \frac{\beta_{pL} - \tilde{B}}{\tilde{B}}\right)\frac{dp}{dq_L} + 1,$$

so that M_L (defined as in (37)) is given implicitly by

$$M_L^{-1} = \rho_L \hat{\tau}_l^{-1} - \left(N_I \hat{\tau}_I \rho^{-1}\frac{\beta_{pI} - \tilde{B}}{\tilde{B}} + \left(N_L - 1\right) M_L \frac{\beta_{pL} - \tilde{B}}{\tilde{B}}\right)^{-1}. \tag{38}$$

Substituting the conditional expectations into the market clearing conditions gives us an equation that, together with M_L and β's pins down an implicit solution for the price p:

$$\begin{aligned} 1 = x &+ \rho^{-1}\hat{\tau}_I \sum_{i=1}^{N_I}\left(\left(1 - \beta_{sI} - \beta_{pI}\right)\mu + \beta_{sI} s_i + \beta_{pI}\frac{p - A - C\bar{x}}{\tilde{B}} - p\right) \\ &+ M_L \sum_{j=1}^{N_L}\left(\left(1 - \beta_{sL} - \beta_{pL}\right)\mu + \beta_{sL} s_i + \beta_{pL}\frac{p - A - C\bar{x}}{\tilde{B}} - p\right). \end{aligned}$$

Collecting terms in p and matching coefficients yields the pricing coefficients

$$A = -C\left(1 + \left(N_I \frac{\hat{\tau}_I}{\rho}\beta_{pI} + N_L M_L \beta_{pL}\right)\frac{A + C\bar{x}}{\tilde{B}}\right) \tag{39}$$

$$+ C\left(N_I \frac{\hat{\tau}_I}{\rho}\left(1 - \beta_{sI} - \beta_{pI}\right) + N_L M_L \left(1 - \beta_{sL} - \beta_{pL}\right)\right)\mu \tag{40}$$

$$B_I = C N_I \rho^{-1} \hat{\tau}_I \beta_{sI} \tag{41}$$

$$B_L = C N_L M_L \beta_{sL} \tag{42}$$

$$C = -\tilde{B}\left[N_I \rho^{-1}\hat{\tau}_I\left(\beta_{pI} - \tilde{B}\right) + N_L M_L\left(\beta_{pL} - \tilde{B}\right)\right]^{-1}. \tag{43}$$

The elasticity of the large investor's demand to information is given by (38). 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 problem is a joint solution to (7) - (11) and (13)- (18).

A.4 Direct Bidding and Minimum Bidding Penalties

We solve the model generally with a minimum bidding penalty. Then we can set $\chi = 0$ to get the results for the no-penalty model.

When the large investor chooses to bid directly on his own behalf, the first dealer's signal is the average of the first N_I/N_D investors', the first $N_L/N_D - 1$ large investors' and the first dealer's signals:

$$\bar{s}_1 = \frac{\tau_{\varepsilon,I}\sum_{k\in\mathcal{I}^I} s_k + \tau_{\varepsilon,l}\left(\sum_{j\in\mathcal{I}^L} s_j + s_d\right)}{\tau_{\varepsilon,I} N_I/N_D + \tau_{\varepsilon,l}(N_L/N_D - 1)}$$

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 we are considering the case with perfect dealer information sharing, the investor's own signal is redundant because it is included in the dealer's average signal $\bar{s}_d$. Thus, for notational convenience, we drop the coefficient on s_i and 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: $\tilde{S}_L = [s_L, s(p)]$.

Variance-covariance matrix of signals The only difference in the signal construction is that we assume 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 his clients') are less precise because they miss the one large investor. And finally, the price information has a slightly different composition.

We can construct a condensed Π matrix as before. Let

$$\bar{Z}^{direct} = \left[\begin{array}{ccccccc} \epsilon_{L1} & \bar{\epsilon}_1 & \bar{\epsilon}_2 & \ldots & \bar{\epsilon}_{Nd} & (x - \bar{x}) \end{array}\right]'$$

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 intermdiation decision of the direct bidder in incorporated in the first dealer's signal, and thus in $\bar{\epsilon}_1$. The variances of these shocks are:

$$\mathbb{V}\left[\bar{Z}^{direct}\right] = diag([\tau_L^{-1},\ (\tau_d - \tau_L)^{-1},\ \tau_d^{-1}1_{N_d-1},\ \tau_x^{-1}]). \tag{44}$$

Then, we can express signals as the true payoff f plus the following loadings on each of the orthogonal Z shocks:

$$\bar{\Pi}_L = \begin{bmatrix} 1 & 0 & \dots & 0 \\ B_L/\tilde{B} & B_{d1}/\tilde{B} & (B_{d\neq 1}/\tilde{B})\, 1'_{(N_D-1)} & C/\tilde{B} \end{bmatrix} \tag{45}$$

$$\bar{\Pi}_{d1} = \begin{bmatrix} 0 & 1 & 0\dots & 0 \\ B_L/\tilde{B} & B_{d1}/\tilde{B} & (B_{d\neq 1}/\tilde{B})\, 1'_{(N_D-1)} & C/\tilde{B} \end{bmatrix} \tag{46}$$

$$\bar{\Pi}_{d\neq 1} = \begin{bmatrix} 1 & \iota_d & & 0 \\ B_L/\tilde{B} & B_{d1}/\tilde{B} & (B_{d\neq 1}/\tilde{B})\, 1'_{(N_D-1)} & C/\tilde{B} \end{bmatrix} \tag{47}$$

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

$$\mathbb{V}(S_j) = \tau_f^{-1} + \bar{\Pi}_{d(j)} V[\bar{Z}^{direct}] \bar{\Pi}'_{d(j)} \qquad \forall j. \tag{48}$$

Bayesian updating weights β and posterior precisions τ For each agent, there are two signals: (1) the signal s_i provided by the dealer or the large investor's own signal, in the case of direct bidding; and (2) the signal conveyed by the price $(p - A - C\bar{x})/B)$. The posterior expectation of the asset value is a linear combination of the prior and each of these two signals. β_{si} and β_{pi} are the weights the agent places on the two signals. According to Bayes' law, the β's are the precision of the signal, divided by the sum of the precisions of the prior and signals. But in this case, the signals are correlated. We need to use optimal linear projection formulas to adjust for this correlation. The 2×1 vector $\beta = [\beta_s, \beta_p]'$ is given by (35). The conditional expectation is then (4). The posterior precision is (6), which can always be expressed as a simple function of the Bayesian updating weights β:

$$\hat{\tau}_l^{-1} = \tau_f^{-1}(1 - \beta_s - \beta_p) \tag{49}$$

where τ_f is the prior variance of the asset payoff, and is exogenous.

Proof of Results 2 and ??: Price when the large investor bids directly (with and without minimum bidding. In order to solve for equilibrium price, we need to determine asset demands of the large investors. That in turn depends on their price impact. To determine price impact dp/dq_j for an agent j, we need to know what the equilibrium price function looks like. The following steps derive a set of equations that jointly solve this fixed point problem.

Substitute demand functions of all agents, except one (dealer 1 in this case) into the market clearing constraint:

$$
\begin{aligned}
1 = x + q_{d1} + \left(\frac{N_I}{N_D}\hat{\tau}_{d1}\rho^{-1} + \frac{N_L}{N_D}M_{L(d=1)}\right)\left(\left(1-\beta_{\bar{s},d1}-\beta_{p,d1}\right)\mu + \beta_{\bar{s},d1}\bar{s}_1 + \beta_{p,d1}s(p) - p\right) \\
+ \sum_{d=2}^{N_D}\left(\frac{N_I}{N_D}\hat{\tau}_{d\neq 1}\rho^{-1} + \left(\frac{N_L}{N_D}-1\right)M_{L(d\neq 1)} + M_{d\neq 1}\right)\left(\left(1-\beta_{\bar{s},d\neq 1}-\beta_{p,d\neq 1}\right)\mu + \beta_{\bar{s},d\neq 1}\bar{s}_d + \beta_{p,d\neq 1}s(p) - p\right) \\
+ \chi\left(N_D-1\right)M_{d\neq 1}p + M_L\left(\left(1-\beta_{sL}-\beta_{pL}\right)\mu + \beta_{sL}s_{N_I+1} + \beta_{pL}s(p) - p\right),
\end{aligned}
$$

where M's are defined as in (37) and $\hat{\tau}_{d1}$ is the posterior precision of the first dealer. The first line equates the supply (1) to the demands of non-competitive investors (x), dealer 1 (q_{d1}) and the demand of all the investors who bid through dealer 1, the second line and first term of the third line is the demand of other dealers and their investors, and the rest of the third line is the large direct bidder.

For each type of agent j, the variable M_j is the sensitivity of j's treasury demand to j's expected profit. For small investor i, this sensitivity is $\hat{\tau}_i\rho^{-1}$ and is simply expressed as such in the formula above. For large investors and dealers, this object complicated by their market power. For each large, non-dealer investor l, the first-order condition tells us that this sensitivity is given by (37): $M_l \equiv \left[\rho_l\hat{\tau}_l + dp/dq_l\right]^{-1}$. The price sensitivity to demand dp/dq_l depends on whether the large investor is a client of dealer 1 $dp/dq_{L(d=1)}$, a client of some other dealer $dp/dq_{L(d\neq 1)}$, or the direct bidder dp/dq_L. Taking the market-clearing condition with all demand substituted in, except q_L, and differentiating with respect to q_L and rearranging, we obtain the sensitivity of price to the demand of the large direct bidder:

$$
\begin{aligned}
\left(\frac{dp}{dq_L}\right)^{-1} = &-\left(\nu_I\hat{\tau}_1\rho^{-1} + \left(\nu_l-1\right)M_{L(d=1)} + M_{d=1}\right)\left(\frac{\beta_{p,d1}-\tilde{B}}{\tilde{B}}\right) - \chi M_{d=1} - \chi\left(N_D-1\right)M_{d\neq 1} \\
&- \left(N_D-1\right)\left(\nu_I\hat{\tau}_2\rho^{-1} + \nu_l M_{L(d\neq 1)} + M_{d\neq 1}\right)\left(\frac{\beta_{p,d\neq 1}-\tilde{B}}{\tilde{B}}\right). \qquad (50)
\end{aligned}
$$

Following the same steps, we obtain the sensitivity of price to the demand of the large investors who bid through dealer 1 (the dealer who lost the business of the large bidder)

$$
\left(\frac{dp}{dq_{L(d=1)}}\right)^{-1} = \left(\frac{dp}{dq_L}\right)^{-1} - M_L\left(\frac{\beta_{pL}-\tilde{B}}{\tilde{B}}\right) + M_{L(d=1)}\left(\frac{\beta_{p,d1}-\tilde{B}}{\tilde{B}}\right), \qquad (51)
$$

and the sensitivity of price to the demand of the large investors who bid through other dealers

$$
\left(\frac{dp}{dq_{L(d\neq 1)}}\right)^{-1} = \left(\frac{dp}{dq_L}\right)^{-1} - M_L\left(\frac{\beta_{pL}-\tilde{B}}{\tilde{B}}\right) + M_{L(d\neq 1)}\left(\frac{\beta_{p,d\neq 1}-\tilde{B}}{\tilde{B}}\right). \qquad (52)
$$

For each dealer d, the first-order condition tells us that the sensitivity of demand to expected profit differs from that of the other large investors because of the minimum bidding

penalty χ:

$$M_d \equiv \left[\rho_l \hat{\tau}_l + (1-\chi) dp/dq_d\right]^{-1}. \tag{53}$$

The price sensitivity to demand dp/dq_d depends on whether this is dealer 1 (the one who lost the business of 1 large investor) or another dealer. Differentiating the market-clearing condition, we find that the price impact dp/dq_{d1} of dealer 1 is

$$\left(\frac{dp}{dq_{d_1}}\right)^{-1} = \left(\frac{dp}{dq_L}\right)^{-1} - M_L\left(\frac{\beta_{pL}-\tilde{B}}{\tilde{B}}\right) + M_{d=1}\left(\chi + \left(\frac{\beta_{p,d1}-\tilde{B}}{\tilde{B}}\right)\right) \tag{54}$$

and the sensitivity of the price to the demand of any other dealer $d \neq 1$ is:

$$\left(\frac{dp}{dq_{d\neq 1}}\right)^{-1} = \left(\frac{dp}{dq_L}\right)^{-1} - M_L\left(\frac{\beta_{pL}-\tilde{B}}{\tilde{B}}\right) + M_{d\neq 1}\left(\chi + \left(\frac{\beta_{p,d\neq 1}-\tilde{B}}{\tilde{B}}\right)\right) \tag{55}$$

Finally, we substitute the demands of all agents back into the market clearing condition, use our price guess to write the price signal as $s(p) = (p - A - C\bar{x})/\tilde{B}$, to obtain

$$\begin{aligned}
1 = {} & x + \left(\nu_I \hat{\tau}_1 \rho^{-1} + (\nu_l - 1) M_{L(d=1)} + M_{d=1}\right)\left((1-\beta_{\bar{s},d1}-\beta_{p,d1})\mu + \beta_{\bar{s},d1}\bar{s}_1 + \beta_{p,d1}\frac{p-A-C\bar{x}}{\tilde{B}} - p\right) \\
& + \sum_{d=2}^{N_D}\left(\nu_I \hat{\tau}_2 \rho^{-1} + \nu_l M_{L(d\neq 1)} + M_{d\neq 1}\right)\left((1-\beta_{\bar{s},d\neq 1}-\beta_{p,d\neq 1})\mu + \beta_{\bar{s},d\neq 1}\bar{s}_d + \beta_{p,d\neq 1}\frac{p-A-C\bar{x}}{\tilde{B}} - p\right) \\
& + \chi M_{d=1} p + \chi\left(N_D - 1\right) M_{d\neq 1} p + M_L\left((1-\beta_{sL}-\beta_{pL})\mu + \beta_{sL} s_{N_I+1} + \beta_{pL}\frac{p-A-C\bar{x}}{\tilde{B}} - p\right).
\end{aligned}$$

Equating coefficients, we obtain the system of equations for the price coefficients

$$\begin{aligned}
A = {} & C\left(-1 + \left(\nu_I \hat{\tau}_1 \rho^{-1} + (\nu_l - 1) M_{L(d=1)} + M_{d=1}\right)\left((1-\beta_{\bar{s},d1}-\beta_{p,d1})\mu - \frac{\beta_{p,d1}}{\tilde{B}}(A + C\bar{x})\right)\right) \\
& + C\left(N_D - 1\right)\left(\nu_I \hat{\tau}_2 \rho^{-1} + \nu_l M_{L(d\neq 1)} + M_{d\neq 1}\right)\left((1-\beta_{\bar{s},d\neq 1}-\beta_{p,d\neq 1})\mu - \frac{\beta_{p,d\neq 1}}{\tilde{B}}(A + C\bar{x})\right) \\
& + CM_L\left((1-\beta_{sL}-\beta_{pL})\mu - \frac{\beta pL}{\tilde{B}}(A + C\bar{x})\right) \qquad (56)
\end{aligned}$$

$$B_1 = C\left(\nu_I \hat{\tau}_1 \rho^{-1} + (\nu_l - 1) M_{L(d=1)} + M_{d=1}\right)\beta_{\bar{s},d1} \tag{57}$$

$$B_2 = C\left(N_D - 1\right)\left(\nu_I \hat{\tau}_2 \rho^{-1} + \nu_l M_{L(d\neq 1)} + M_{d\neq 1}\right)\beta_{\bar{s},d\neq 1} \tag{58}$$

$$F = CM_L \beta_L(1) \tag{59}$$

$$\begin{aligned}
C^{-1} = {} & -M_L\left(\frac{\beta pL - \tilde{B}}{\tilde{B}}\right) - \left(\nu_I \hat{\tau}_1 \rho^{-1} + (\nu_l - 1) M_{L(d=1)} + M_{d=1}\right)\left(\frac{\beta_{p,d1}-\tilde{B}}{\tilde{B}}\right) - \chi M_{d=1} \\
& - \left(N_D - 1\right)\left(\nu_I \hat{\tau}_2 \rho^{-1} + \nu_l M_{L(d\neq 1)} + M_{d\neq 1}\right)\left(\frac{\beta_{p,d\neq 1}-\tilde{B}}{\tilde{B}}\right) - \left(N_D - 1\right)\chi M_{d\neq 1}. \qquad (60)
\end{aligned}$$

This verifies the price conjecture. Since supply of the asset is one, price and revenue are equal.

The solution to this model is a joint solution to the price coefficient equations (56) - (60), the M equation (37), in conjunction with the price sensitivities (50) - (55) and the β's (),

which in turn depend on the Π's (45), (46) and (47) and the variances of the orthogonal shocks (44). Setting $\chi = 0$ yields the solution to the model without the minimum bidding requirement.

A.5 Indirect Bidding with Perfect Information Sharing and Bidding Penalties.

In this case, all the dealers in the market are symmetric, and the beliefs of any investor bidding through dealer d are the same as his dealer's beliefs. We begin with the preliminaries of belief updating.

Define the weight the dealer's signal puts on each small investor's signal as $\omega_I \equiv \tau_{\epsilon,s}/(\nu_I\tau_{\epsilon,s}+(1+\nu_l)\tau_{\epsilon,l})$, where τ's are exogenous signal precisions. $\nu_I \equiv N_I/N_D$ is small investors per dealer and $\nu_l \equiv N_L/N_D$ is large investors per dealer (both exogenous). Then s_d is a signal-precision-weighted average of all of the dealer's customer signals as well as the dealer's own signal:

$$\bar{s}_d = \frac{\tau_{\varepsilon,s}\sum_{k\in\mathcal{I}_d^s} s_k + \tau_{\varepsilon,l}\left(\sum_{j\in\mathcal{I}_d^l} s_j + s_d\right)}{\nu_I\tau_{\varepsilon,s} + (1+\nu_l)\tau_{\varepsilon,l}} \equiv \omega_I \sum_{k\in\mathcal{I}_d^s} s_k + \omega_l \sum_{k\in\mathcal{I}_d^l} s_k + \omega_l s_d.$$

For a given agent who trades through dealer d, there are only 2 relevant signals: the dealer's signal $s_d = f + \bar{\epsilon}_d$ where $\bar{\epsilon}_d \sim N(0, \tau_d^{-1})$ and $\tau_d = \nu_I\tau_s + (\nu_l+1)\tau_l$. The ν's, τ_s and τ_l are all primitives. We can express each dealer's signals and the price signal as the following Π loading matrix times a vector Z of orthogonal shocks: $S_d = f + \bar{\Pi}_d\bar{Z}$:

$$\bar{\Pi}_d = \begin{bmatrix} & \iota_d & & 0 \\ 1/N_d & \dots & 1/N_d & C/B \end{bmatrix} \tag{61}$$

$$\bar{Z} = \begin{bmatrix} \bar{\epsilon}_1 & \dots & \bar{\epsilon}_{Nd} & x \end{bmatrix}' \qquad \text{s.t.}$$

$$\mathbb{V}\left[\bar{Z}\right] = diag([\tau_d^{-1}1_{N_d},\ \tau_x^{-1}]). \tag{62}$$

where ι_d is a vector of zeros with a 1 in entry d and B and C are equilibrium pricing coefficients

Then, for all agents j who are clients of dealer d, signal variance-covariance is

$$\mathbb{V}(S_j) = \tau_f^{-1} + \bar{\Pi}_d V[\bar{Z}]\bar{\Pi}_d' \qquad \forall j \in \mathcal{I}_d. \tag{63}$$

Then, the Bayesian updating weights and posterior precision are the same formulas as before (35) and (**??**). Using (4) and (6), we find that the posterior belief is

$$\mathbb{E}\left[f|X_dS\right] = (1-\beta_{\bar{s}}-\beta_p)\,\mu + \beta_{\bar{s}}\bar{s}_d + \beta_p\frac{p-A-C\bar{x}}{B}, \tag{64}$$

with conditional variance $\mathbb{V}\left[f|X_dS\right] = \tau_f^{-1} - \tau_f^{-1}\left(\beta_{\bar{s}}+\beta_p\right) \equiv \hat{\tau}^{-1}$.

Proof of Result ??: Price in the N dealer model with all investors bidding through a dealer. The small investors' demand functions (bids) in the competitive market model are given by equation (10) and the large investors' bids by equation (36). Dealers' bids are (12). Substituting these bids into the market clearing condition $x + \sum_{i=1}^{N} q_i = 1$ yields

$$1 = x + \sum_{d=1}^{N_D} \left(\nu_I \hat{\tau} \rho^{-1} + (\nu_l + 1) M_L\right) \left(\mathbb{E}\left[f | X_d S\right] - p\right) + \chi N_D M_D p.$$

Substituting in (64) for the conditional expectation and then taking the derivative with respect to q_L, we obtain the price impact of a large investor

$$0 = 1 + \left(N_I \hat{\tau} \rho^{-1} + (N_L + N_D - 1) M_L\right) \left(\frac{\beta(2) - B}{B} + \chi N_D M_D\right) \frac{dp}{dq_L},$$

so that M_L (defined as in (37)) is given implicitly by

$$M_L^{-1} = \rho_L \hat{\tau}^{-1} - B \left(\left(N_I \hat{\tau} \rho^{-1} + (N_L + N_D - 1) M_L\right) (\beta(2) - B)\right)^{-1}. \tag{65}$$

Similarly, taking the derivative with respect to q_d, we obtain the price impact of a dealer dp/dq_d. Substituting into (37) yields the dealer demand coefficient

$$M_D^{-1} = \rho_l \hat{\tau}^{-1} - (1 - \chi) \left[\left(N_D \left(\nu_I \hat{\tau} \rho^{-1} + \nu_l M_L + M_D\right) - M_D\right) \frac{\beta(2) - B}{B} + \chi \left(N_D - 1\right) M_D\right]^{-1}. \tag{66}$$

Finally, solving the market clearing condition for the settle price and equating the coefficients, we obtain the pricing coefficients

$$A = C \left(-1 + N_D \left(\nu_I \hat{\tau} \rho^{-1} + \nu_l M_L + M_D\right) \left(\left(1 - \beta\left(1\right) - \beta\left(2\right)\right) \mu - \frac{\beta(2)}{B} \left(A + C \bar{x}\right)\right)\right) \tag{67}$$

$$B = C N_D \left(\nu_I \hat{\tau} \rho^{-1} + \nu_l M_L + M_D\right) \beta(1) \tag{68}$$

$$C = -\left[N_D \left(\left(\nu_I \hat{\tau} \rho^{-1} + \nu_l M_L + M_D\right) \left(\frac{\beta(2) - B}{B}\right) + \chi M_D\right)\right]^{-1}. \tag{69}$$

The solution to this model is a solution to (35), (65)-(66) and (67)-(69).

B Extension: Learning from the intermediation decision of the large investor

In this appendix, we consider an extension of 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 (14), 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

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

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 υ_l and υ_h. We simulate the model to determine the probability of direct bidding $Pr(\upsilon)$. Then, we also compute, conditional on choosing to bid directly, what the direct bidder's average signal is (s_υ) and the variance of that signal (τ_υ^{-1}). Next, we construct a normally-distributed, conditionally independent approximating signal with the same mean and variance: $s_\upsilon = f + \epsilon_\upsilon$ where $\epsilon_\upsilon \sim N(0, \tau_\upsilon^{-1})$. In cases where the large investor bids directly, we allow dealer 1 to observe s_υ 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 $Pr(\upsilon)$ 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_I investors', the first $\nu_l - 1$ large investors', the first dealer's signal and the new signal s_υ that arises from observing the direct bidding decision. By Bayes' law, the posterior expectation, conditional on all these signals is:

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

Bayes law also tells us that the precision is $\hat{\tau}_{d1} = \nu_I\tau_{\varepsilon,s} + \nu_l\tau_{\varepsilon,l} + \tau_\upsilon$. 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

$$\bar{Z}_v = \left[\begin{array}{ccccc} \epsilon_L & \bar{\epsilon}_1 & \bar{\epsilon}_2 & \dots & \bar{\epsilon}_{Nd} \quad (x - \bar{x}) \end{array} \right]'$$

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 in incorporated in the first dealer's signal, and thus in $\bar{\epsilon}_1$. The variances of these shocks are:

$$\mathbb{V}\left[\bar{Z}^{direct}\right] = diag([\tau_L^{-1},\ \hat{\tau}_{d1}^{-1},\ \tau_d^{-1} 1_{N_d - 1},\ \tau_x^{-1}]). \tag{70}$$

The signals of the large investor L, dealer 1 and his clients $D1$ and all others $d \neq 1$ have loadings on the orthogonal shocks given by (45), (46) and (47). 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 $\mathbb{V}(S_j) = \tau_f^{-1} + \bar{\Pi_d}(j) V[\bar{Z}_v] \bar{\Pi_d}(j)' \qquad \forall j$.

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 (35), the resulting conditional expectation is given by (4), with precision (49).

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 (56) - (60), the M equation (37), in conjunction with the price sensitivities (50) - (55) and the β's (35), which in turn depend on the Π's (45), (46) and (47) and the variances of the orthogonal shocks (70). Setting $\chi = 0$ yields the solution to the model without the minimum bidding requirement.

C Additional Results

C.1 Minimum bidding results

Figure 8 reports the effect of information sharing on auction revenue and utility in a world where dealers face minimum bidding constraints. The revenue is higher than the model without the constraint. The reason for this is that the constraint is modelled as a penalty for bidding too low. When dealers bid want to avoid the penalty, they bid higher and that pushes auction revenue up. The other salient difference from the results without penalties

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.

Panel A: Full Sample

	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

Panel B: Direct Bidding

	Revenue	Dealer allocation	Direct allocation	Indirect allocation
Mean	57.6447	40.8394	1.3716	44.6400
Std. Dev.	67.6537	20.1439	0.5585	21.9002
Skew	0.0028	0.0022	0.0956	-0.0031
Kurtosis	3.0656	3.0607	3.0917	3.0744

Panel C: Indirect Bidding

	Revenue	Dealer allocation	Direct allocation	Indirect allocation
Mean	28.8670	44.7330	0.0000	53.6680
Std. Dev.	70.4135	18.9455	0.0000	88.6057
Skew	0.0120	0.0120	-	-0.0120
Kurtosis	2.9911	2.9911	-	2.9911

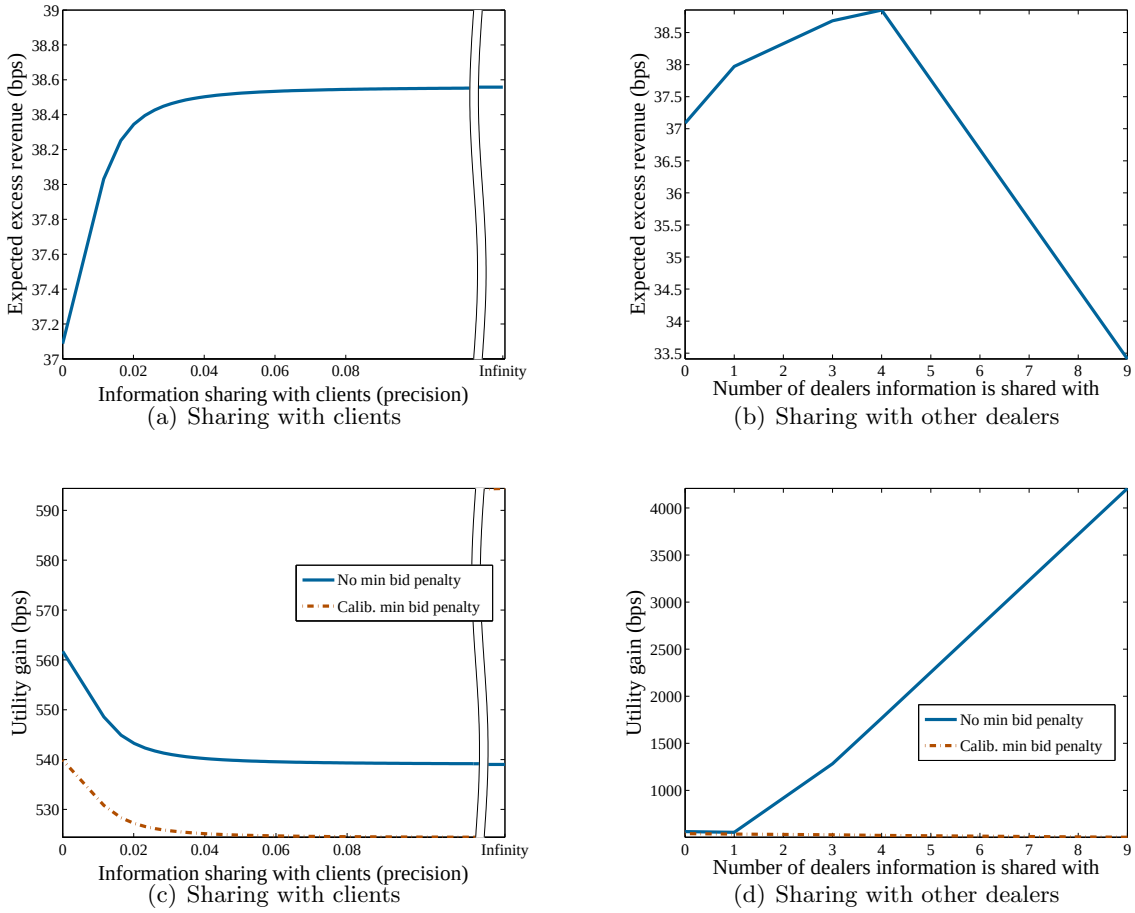
Panel D: High signal

	Revenue	Dealer allocation	Direct allocation	Indirect allocation
Mean	59.0374	41.2442	1.4022	44.1724
Std. Dev.	67.6489	20.1445	0.5565	21.8918
Skew	0.0064	0.0058	0.0865	-0.0064
Kurtosis	3.0754	3.0687	3.1068	3.0857

Panel E: Low signal

	Revenue	Dealer allocation	Direct allocation	Indirect allocation
Mean	17.1202	41.5724	0.0000	68.4497
Std. Dev.	69.9178	18.8121	0.0000	87.9818
Skew	0.0277	0.0277	-	-0.0277
Kurtosis	2.9739	2.9739	-	2.9739

Figure 8: Minimum Bidding Results Top row shows auction revenue. Bottom row plots utility.



(a) Sharing with clients

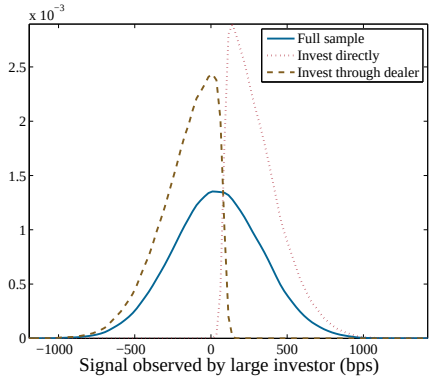
(b) Sharing with other dealers

(c) Sharing with clients

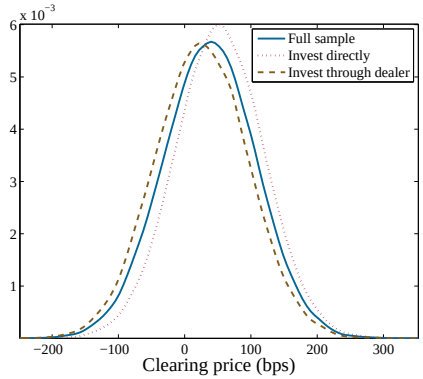
(d) Sharing with other dealers

is that revenue is non-monotonic in dealer information sharing (panel b). This is because of how penalties interact with price informativeness.

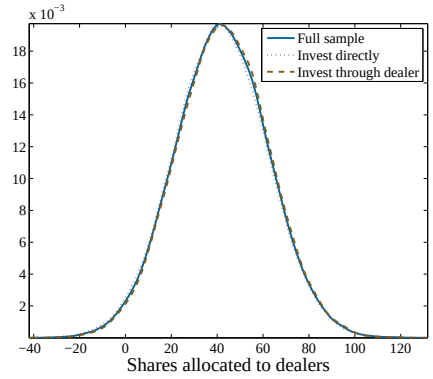
Figure 9: Distribution of simulated outcomes



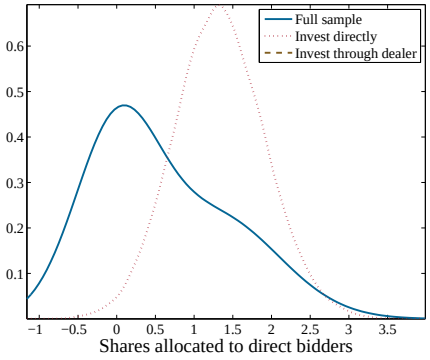
(a) Signals observed by large investor



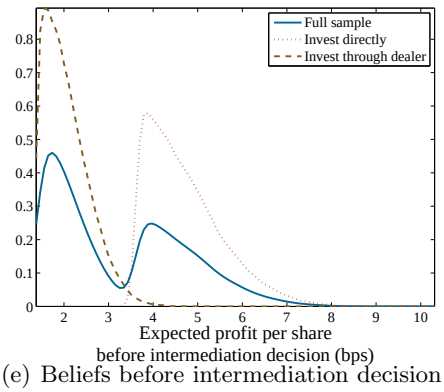
(b) Auction clearing price



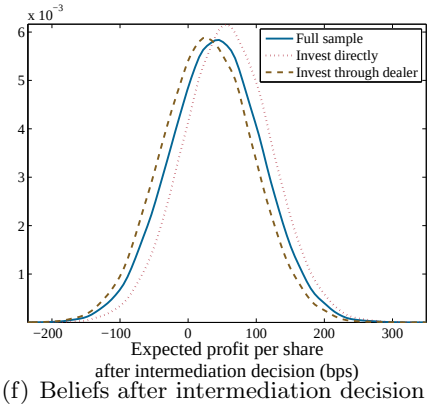
(c) Shares allocated to dealers



(d) Shares allocated to direct bidders



(e) Beliefs before intermediation decision



(f) Beliefs after intermediation decision