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

Nina Boyarchenko, David O. Lucca, and Laura Veldkamp

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

January 2, 2018

Abstract

The use of order flow information by financial firms has come to the forefront of the regulatory debate. A central question is: Should a dealer who acquires information by taking client orders be allowed to use or share that information? We explore how information sharing affects dealers, clients and issuer revenues in U.S. Treasury auctions. Because one cannot observe alternative information regimes, we build a model, calibrate it to auction results data, and use it to quantify counter-factuals. We estimate that yearly auction revenues would be $2.4 billion higher with full-information sharing with clients and between dealers. When information sharing enables collusion, the collusion costs revenue; but if dealers share information with clients, prohibiting information sharing may cost more. For investors, the welfare effects of information sharing depend on how information is shared. The model shows that investors can benefit when dealers share information with each other, not when they share more with clients.

*Revision in progress. 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 our editor, Ali Hortacsu and our two anonymous referees, as well as Bruno Biais, Giovanni Cespa, John Cochrane (through his blog), Darrell Duffie, Ken Garbade, Luis Gonzalez, Robin Greenwood, Zhiguo He, Pablo Kurlat, Richard Tang, Hongjun Yan, Haoxiang Zhu and participants at the NBER AP meetings, Chicago Booth Asset Pricing Conference, Duke, Miami, U.S. Treasury Roundtable on Treasury Markets, NYU, the Spring 2015 Macro Finance Society Meetings, Gerzensee Summer Institute 2015, EFA 2015, EEA 2015, FIRS 2016, Bank of Canada and SFI Lausanne for comments. Nic Kozeniauskas, Matias Covarrubias, Karen Shen, Arnav Sood and Peifan Wu provided excellent research assistance. Emails: nina.boyarchenko@ny.frb.org; david.lucca@ny.frb.org; lveldkam@stern.nyu.edu.

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

Recent financial market misconduct, involving misuse of information about clients' orders, cost the firms involved record fines and lost reputation. It also prompted investigations and calls for curbing dissemination of order flow information, between and within dealers. Recent investigations reportedly involve U.S. Treasury auctions (Bloomberg, 2015 above). But the use of order flow information has been central to our understanding of Treasury auctions (Hortaçsu and Kastl, 2012), to market making theory generally (Kyle, 1985) and to market practice for decades. In describing Treasury market pre-auction activities in the 1950s, Robert Roosa (1956) noted that "Dealers sometime talk to each other; and they all talk to their banks and customers; the banks talk to each other." Furthermore, sharing order-flow information–or, colloquially, "market color"–with issuers is even mandatory for primary dealers both in the U.S. and abroad. Of course, if information sharing leads to collusion, that has well-known welfare costs. But if collusion could be prevented with separate remedies, is information sharing in 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 quantitative model. Our setting is an institutionally-detailed model of U.S. Treasury auctions, which we select because of the available data, the absence of other dealer functions,[2] and their enormous economic importance. In the model, dealers observe client orders and may use that in-

[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 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.

formation to inform their own strategy, share some of the information with clients, or exchange information with other dealers. Then all agents submit continuous bid functions to a uniform-price auction that features both common and private values. To quantify the effects of order flow information sharing and sign welfare results, we calibrate the model to auction results, allotment data as well as information about post-auction returns using market prices on the co-called on-the-run premium, or the differential value of a new versus an old Treasury security. In this setting, bids reflect risk premia associated with reselling Treasuries, at an unknown price, in the secondary market. This risk premia informs the model about how much uncertainty bidders face, and thus, how much information they have, on average. After calibrating the model, we study the model-implied revenue and bidders' utilities with varying degrees and types of information sharing. We then extend the model to think about how information sharing affects bidders and dealers incentives to participate in the auction. Finally, we provide some empirical support for key model assumptions.

The model teaches us that the primary beneficiary of information sharing is the U.S. Treasury because better-informed buyers bid more. Based on the model parameters, moving from the calibrated status-quo of a partial information sharing arrangements to full information sharing would raise Treasury auction revenues by $2.4 billion annually. If instead, all information sharing were prohibited, revenue would fall by $80 million. 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 Hirshleifer (1971) effect, which arises here because better-informed clients have more heterogeneous beliefs and therefore share risk less efficiently. But surprisingly, when dealers share information with each other and then transmit the same amount of information to their clients, investor welfare improves. Our model shows how inter-dealer information sharing makes beliefs more common, and thereby improves risk-sharing and welfare. In essence, information sharing with clients is similar to providing more private information, while inter-dealer sharing functions effectively makes information more public.

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

Fourth, our findings contribute to our understanding of a symbiotic relationship between investors and intermediaries: it is the process of intermediating trades that reveals information to dealers. Information sharing is what induces clients to use intermediaries and induces large investors to intermediate.

These findings are not meant to imply that dealers should have carte blanche in using information in any way they choose. The model assumes that clients know how dealers use their information, and that order flow information is aggregated. 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 foundation for the model is Kyle (1989). We adapted that framework to a correlated-values, uniform-price auction with heterogeneous information, limit orders and market orders. Our bidders are fully strategic. They exploit their price impact (bid shading) and bid in such as way as to optimally manipulate the beliefs of others (signal jamming). On top of this foundation, we assume that dealers observe informative signals from order flows and that bidders share correlated values.

The assumption that bidders have private signals about future Treasury valuations 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.

Our model differs from previous treasury auction models (Hortaçsu and Kastl, 2012) by assuming that bidders have both a private and a common component to their valuation. The common component is the secondary market resale value. Many auction participants speculate on post-auction appreciation and sell within a week of the auction. Even non-speculators consider the fact they could obtain the same securities at a market price after the auction. This new assumption of correlated values is central to understanding the costs and benefits of information sharing. The risk that another bidder will bid on my information, because it is also informative about their own return, influence the price, and reduce my return, is what makes information sharing potentially costly.

In order to match additional institutional features we also model the auction as "mixed auctions", meaning that investors can bid indirectly (through a dealer) or directly (without any intermediary). Finally, we account for minimum bidding requirements of primary dealers, who have historically been expected to bid "consistently" at all auctions for amounts, which today, are equal to the pro-rata share of the offered amount.

Contribution to the existing literature. Previous papers have estimated the extent of client information shared with Treasury dealers.[3] We consider that dealers may send signals back to their clients and to other dealers. Also, we explore the incentive of a bidder to bid directly, or through a dealer. More broadly, our work offers a different framework for measurement. We use risk premia and the covariance of prices and payoffs to infer how much investors know. Our risk-based estimation approach predicts different revenue, market power and utility effects of information. Relative to Hortaçsu and Kastl (2012) and Hortaçsu, Kastl, and Zhang (2016), our model misses the realism of bids that are step functions. Instead, we assume demand is continuous and linear. This simplification allows for risk-averse utility, correlated values and asymmetric information. Our model captures bidders speculating on post-auction appreciation, whereas private values models describe buy-and-hold bidders, who receive a known payout at maturity. The costs and benefits of information sharing depend on this difference. For speculators with correlated values and asymmetric information, observing others' information helps the speculator determine their own value of the asset more accurately. Speculators worry that sharing their information with others will induce others to bid more aggressively because information reduces their risk. Without risk aversion, this effect disappears. So, while our model compromises realism in bidding, it enables us to examine new effects of information sharing.

Our main ideas are connected to a microstructure literature that studies how pre-trade order flow information contributes to price formation (O'Hara, 1995, Chapter 9), raises bid-ask spreads (Bloomfield and O'Hara, 1999) and affects utility of informed and uninformed traders (Fishman and Longstaff, 1992; Röell, 1990). 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).

[3] Hortaçsu and McAdams (2010) and others studies how auction design affects revenues, or how it affects information acquisition (Cole, Neuhann, and Ordonez, 2017) complements our project.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 alternatives. Empirical work by Nyborg and Sundaresan (1996), Malvey, Archibald, and Flynn (1995) and Malvey and Archibald (1998) compares revenues from 1992-1998 when the U.S. Treasury used both uniform and discriminatory price auctions. Armantier and Sbaï (2006) use French Treasury auction bids to structurally estimate the benefits of uniform price auctions. Similarly, He, Krishnamurthy, and Milbradt (2016) explore why US Treasuries are safe. We fix the auction format to a uniform-price multi-unit auction, fix the distribution of future Treasury values, and instead focus on how information sharing affects revenue and surplus.

Treasury primary dealers serve similar purposes to equity and initial public offering (IPO) dealers. Di Maggio, Franzoni, Kermani, and Sommavilla (2017) document that equity dealers share clients' bid information with other clients, just like the dealers in our model do. The IPO literature finds that intermediaries act as underwriters who stabilize prices, in return for expected profit (Ritter and Welch, 2002). We show, instead, that when dealers share information, the conventional wisdom of underwriting is reversed: information intermediaries raise expected revenue. Finally the paper is also related to work on dual-capacity trading, where brokers can submit customer orders and trade on their own accounts (see Fishman and Longstaff, 1992; Röell, 1990,).

The idea that intermediary behavior determines the equilibrium price of an asset arises in He and Krishnamurthy (2013) and Brunnermeier and Sannikov (2014). Their capital-constrained intermediaries provide households with access to risky asset markets and thus improve risk sharing. In Babus and Parlatore (2015), dealers fragment a market, which inhibits risk-sharing. In contrast, we explore how information-sharing induces some agents, who could access markets directly, to choose intermediation.

1 A Treasury auction Model with information sharing

The auction setting is similar to Kyle (1989), where strategic bidders submit continuous bid functions for an asset with a common value, or in our case a correlated value. The novel feature of the model lies in its rich information structure. We allow bidders to share information. Figure 1 summarizes the alternative sharing arrangements that we consider, for a simplified setting with only a few market participants. Dealers are denoted with the letter "D," while investors with the letter "I." Panel a) shows the case of no information sharing ("Chinese walls"), where each auction participant only observes his private information s_i.

When information is shared between dealers and customers (panel b), an investor's information set includes both her private signal and the dealer's; the dealer also observes this extended information set. With cross-dealer information sharing (panel c), each investor observes his dealer's and the other dealer's information. 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. Since all bidders can condition on every possible price, each bidder can use the information that would be conveyed by that realized price. Thus bids are formed as if the realized price were in every bidder's information set.

While this simplified setting conveys the essence of information sharing, our model is richer

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 sets in which information is shared.

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

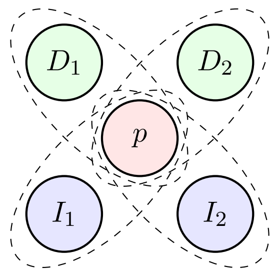

(b) Sharing with customers, not with other dealers

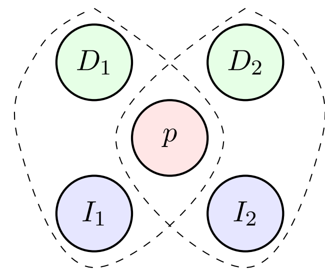

(c) Sharing with customers and dealers

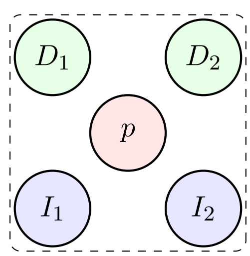

(d) Sharing with customers only; one direct bidder

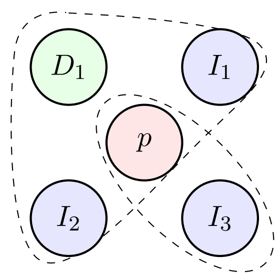

along many dimensions, as discussed next.

Assets 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 secondary market common value of the newly issued asset is unknown to the agents and normally distributed: $f \sim N(\mu, \tau_f^{-1})$.

Bidders To match key features of Treasury auctions, we consider four types of auction participants: dealers, as well as direct, indirect and non-price contingent bidders. The first three types we call speculative because they have a common value component to their payoff, that depends on the asset's future value. Each speculative bidder/dealer can submit a continuous function that specifies a quantity demanded, for every possible clearing price p. All dealers and direct bidders place bids directly in the auction. Indirect bidders are speculative bidders who bid through a dealer, instead of bidding directly. For now, the number of each type of bidders is fixed. Later, we examine the choice to bid (in)directly. There are N_I indirect bidders, which we index by $i = \{1, \ldots, N_I\}$ and N_J direct investors, which we index by $j = \{1, \ldots, N_J\}$.

In addition to the common value f of the Treasury security, each speculative bidder has a private value. For direct and indirect investors $v_i \sim$ *i.i.d.* $N(0, \tau_{vI}^{-1})$ and $v_j \sim$ *i.i.d.* $N(0, \tau_{vJ}^{-1})$ per share. There are many reason why investors may value Treasury issues differently (see Hortaçsu and Kastl, 2012). For example, a depository institution might address a duration mismatch, a foreign official may be investing dollar-denominated reserves, or an investor might cover a short position in the forward Treasury market (known as the when-issued market).

Each bidder has initial wealth W_{i0}, and chooses the quantity of the asset to hold, $q_i \in \mathbb{R}$ at price p per share, to maximize his expected utility,[4]

$$\mathbb{E}[-\exp(-\rho(W_i + q_i v_i))], \tag{1}$$

where ρ denotes absolute risk aversion. The budget constraint dictates that final wealth is initial wealth, plus earnings from post-auction appreciation:

$$W_i = W_{i0} + q_i(f - p). \tag{2}$$

All speculative bidders internalize the effect they have on market prices. Because they strategically consider their price impact, they are not perfectly competitive. They maximize their utility subject to the budget constraint as well as the market clearing condition.

While all other participants submit price-contingent (limit) orders, the non-price contingent bidders submit market orders. (Treasury parlance calls these "non-competes," which are in practice relatively small). These bidders have only a private value for the asset and do not condition their bids on price. Non-price contingent orders are exogenous and random. The aggregate non-price contingent demand is $\delta \sim N(0, \tau_\delta^{-1})$.

Dealers There are N_D dealers, which we index by $d = \{1, \ldots, N_D\}$. Then $N = N_I + N_J + N_D$ is the total number of speculative auction participants. Like investors, dealers' utility also depends on the common resale value of the asset, as well as a private value component.

[4]Technically, the price of each Treasury is fixed at par and auction participants bid coupon payments. Here p is the present discounted value of coupons computed from other outstanding Treasury securities.

But for dealers, the private value may arise, in part, from regulations known as "minimum bidding requirements."[5] In any given auction, a dealer may violate the requirement. But if over time, 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 give dealers private values for each share that are typically positive, but decreasing in asset shares $v_d = \chi + \frac{\chi_0}{q_d}$, where q_d represents the number of shares awarded to dealer d at the market price. The decreasing private value represents the idea that when the dealer's bid q_d is too low, raising that bid reduces the risk of penalties for the dealer. When the bid is already high and the requirement is satisfied, additional shares might relax future bidding constraints, but provide diminishing value. This cost is a stand-in for the shadow cost of a dynamic constraint.

Dealers choose asset demand functions $q_d(p)$ to maximize

$$\mathbb{E}[-\exp(-\rho(W_d + q_d v_d))] \qquad \text{where } W_d \text{ is given by (2).} \tag{3}$$

Describing Information Sets and Updating Beliefs with Correlated Signals Bidders can potentially observe five possible pieces of information: 1) their own private signal, 2) signals from others who may share information with them, 3) the equilibrium price of the asset, and 4) their private value v_i. For dealers, the private value is common knowledge as it derives from a common and publicly known requirement while investors' private values are private information. Finally, it is a feature of Treasury auctions that all non-price contingent bids are publicly revealed before bidding closes. Therefore, we assume 5) that δ is common knowledge. We explain each in turn.

Before trading, each bidder 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 $\varepsilon_i \sim N(0, \tau_\varepsilon^{-1})$.

By placing orders through dealers, customers reveal their order flow $q_i(p)$ to their dealer,

[5] In the current design of the primary dealer system, dealers are expected to bid for a pro-rata share of the auction at "reasonably competitive" prices Federal Reserve Bank of New York (2016). 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. See appendix E in Brady, Breeden, and Greenspan (1992) In 1997, the New York Fed instituted an explicit counterparty performance scorecard and dealers were evaluated based on the volume of allotted securities (FOMC, November 2007). In 2010 the NY Fed clarified their primary dealer operating policies and strengthened the requirements (Federal Reserve Bank of New York, 2016).

which in the model is equivalent to sharing their expected value of the security $\mathbb{E}_i[f] + v_i$. Each dealer d receives orders from N_I/N_D clients.[6] The dealer constructs $\tilde{s}_d$, which is an average of his clients' expected valuations:

$$\tilde{s}_d = \frac{N_D}{N_I} \left(\sum_{i \in \mathcal{I}_d} \mathbb{E}_i[f] + v_i \right), \tag{4}$$

where $\mathcal{I}_d$ is the set of investors bidding through dealer d.

Dealers, in turn, can share some of this order flow information with their clients. Dealer-client information sharing takes the form of a noisy signal about $\tilde{s}_d$, which is the summary statistic for everything the dealer learned from client order flow. That noisy signal is $s_{\xi i} = \tilde{s}_d + \xi_i$ where $\xi_i \sim N(0, \tau_\xi^{-1})$ is the noise in the dealers' advice. The noise ξ_i varies by dealer and by client, but sums to zero for each dealer, meaning that dealers do not systematically mislead clients. Section 4.4 considers dealers who do mislead clients, by not truthfully revealing their information sharing. 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 ($\tau_\xi = 0$).

In addition, dealers may share information with other dealers. Let ψ be the size of the group of dealers who share their information with each other. In other words, each dealer reveals all of his or her signals to $\psi - 1$ other dealers. No sharing between dealers is the case where $\psi = 1$. All information sharing is mutual.

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 $s_i(p)$ denote the unbiased signal i constructs from auction-clearing price. We guess and verify that $(s_i(p) - f) \sim N(0, \tau_{pi}^{-1})$, where τ_{pi} is a measure of the informativeness of the auction-clearing settle price. Recall that direct and indirect bidders have private values that are private information. Adjusting their price inference for their own valuation, they infer $s(p|v_i)$ or $s(p|v_j)$ from a realized price p.

Signal vectors for the three types of agents are as follows: An investor who bids directly

[6] It is quite plausible that a dealer might also include his own private signal in the information he transmits to clients. However, the paper focuses on the effect of dealers' sharing of client order flow information and we therefore exclude dealers' private information from $\tilde{s}$, in order to make clear that our results reflect the sharing of order flow information. Note also that dealers' signals to clients covary with clients' private and public information. Our solution method accounts for this covariance.

observes a vector of signals $S_j = [s_j, s(p|v_j)]$. Investors who bid through dealers observe the larger signal vector $S_i = [s_i, s_{\xi i}, s(p|v_i)]$. While these investors observe an extra signal, they also will end up having signals and thus making bids that covary more with price information. A dealer observes the same signals as an indirect investor, except that he sees the exact order flows, instead of a noisy signal of them. For dealer d, $S_d = [s_d, \tilde{s}_d, s(p)]$. Since non-price contingent bids δ and dealer valuations v_d are common knowledge, we don't include them in S. But every speculative bidder accounts for them.

For every agent, we use Bayes' law to update beliefs about f. Bayesian updating is complicated by correlation in the signal errors. To adjust for this correlation, we use the following optimal linear projection formulas:

$$\mathbb{E}\left[f|S_j\right] = (1-\beta' 1_m)\mu + \beta' S_j \qquad \text{where} \tag{5}$$

$$\beta_j \equiv \mathbb{V}\left(S_j\right)^{-1} \mathbf{Cov}\left(f, S_j\right) \tag{6}$$

$$\mathbb{V}\left[f|S_j\right] = \mathbb{V}\left(f\right) - \mathbf{Cov}\left(f, S_j\right)' \mathbb{V}\left(S_j\right)^{-1} \mathbf{Cov}\left(f, S_j\right) \equiv \hat{\tau}_j^{-1}, \tag{7}$$

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.

Equilibrium. A Nash equilibrium is

1. A bid function by each direct or indirect bidder that maximizes

$$\max_{q_i, p} \mathbb{E}[-\exp(-\rho(W_i + q_i v_i))|S_i] \tag{8}$$

$$\text{s.t. } W_i = W_{0,i} + q_i(f-p) \qquad \text{and} \qquad (12) \tag{9}$$

   The second constraint (12) is the auction clearing condition and reflects that the speculative bidders choose their quantity, taking into account the effect their demand has on the equilibrium price.

2. A bid function for each dealer that maximizes

$$\max_{q_d, p} \mathbb{E}[-\exp(-\rho(W_d + q_d v_d))|S_d] \tag{10}$$

$$\text{s.t. } W_d = W_{0,d} + q_d(f-p) \qquad \text{and} \qquad (12) \tag{11}$$

3. An auction-clearing (settle) price that equates demand and supply:

$$\sum_{i=1}^{N_I} q_i + \sum_{j=1}^{N_J} q_j + \sum_{d=1}^{N_D} q_d + \delta = 1. \tag{12}$$

Of course, in practice, customers submit orders, then dealers observe and share some information about these orders. Then, bids get revised, and so forth. Our equilbrium concept is a Nash or rational expectations equilibrium. It represents the point where this Tatonnement process converges.

Similarly, in practice, auction prices cannot possibly be observed while bids are still being formed. However, auction theory teaches us that each bidder should avoid the winner's curse by choosing a quantity for each price that would be optimal, if he observed that market-clearing price and included it in his information set. Since bidders can set a different demand for every possible price, and thus condition on the information contained in every possible price, the information set of investor i is effectively $\{s_i, p\}$.[7]

2 Solving the Model

We first solve for optimal bid schedules of investors and dealers. Then, we work out the auction equilibrium with different information being shared. 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.

We then substitute the budget constraint in the objective function, evaluate the expectation and take the log. The investor maximization problem simplifies to $\max_{q_j,p} \quad q_j(\mathbb{E}[f|S_j] + v_j - p) - \frac{1}{2}\rho q_j^2 \mathbb{V}[f|S_j]$ subject to the market clearing condition (12), where the price is *not taken as given*.[8] The first order condition with respect to q_j reveals that investors bid

$$q_j\,(p) = \frac{\mathbb{E}[f|S_j] + v_j - p}{\rho \mathbb{V}[f|S_j] + dp/dq_j}. \tag{13}$$

For dealers, the expression is almost identical. The only difference arises from the gap in signal vector S that the dealer conditions on and the form of the private value v_d. Since $v_d q_d = \chi q_d + \chi_0$, the constant χ_0 drops out when taking the first order condition and the

[7] There is a long history of including market-clearing prices in information sets, including the literature building on Grossman and Stiglitz (1980) and Kyle (1989).

[8] In the baseline model we rule out collusion, and relax this assumption in Section 4.4.

optimal dealer bid is

$$q_j(p) = \frac{\mathbb{E}[f|S_j] + \chi - p}{\rho \mathbb{V}[f|S_j] + dp/dq_j}. \tag{14}$$

2.1 Equilibrium auction-clearing price: Three cases

In order to understand the implications of different information sharing arrangements, we solve for auction outcomes in the three cases 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.

The no-information-sharing world is one with "Chinese walls," where dealers cannot use client information to inform their own or their clients' purchases. In recent years, a number of financial firms have reportedly 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.[9] In our Chinese wall model specification, each agent sees only their own private signal s_i and the price information $s_i(p)$ which they can condition their bid on, but not any signal from the dealer: $S_i = [s_i, s_i(p)]$.

In the information sharing cases, investors observe the larger signal vector $S_i = [s_i, s_{\xi i}, s_i(p)]$. The signal $s_{\xi i}$ includes information from clients and/or information shared across dealers. In these cases, the investors' own information will also be shared with others.

The equilibrium auction price is obtained by adding up all investors' and dealers' asset demands as well as the volume of market orders x and equating them with total supply. 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 price weights, which affects utility. 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.

Result 1. *Under each of the following three information-sharing regimes:*

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

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

auction revenues are always a linear function of signals s_i and investors' average private values $\bar{v}$:

$$p = A + B_I \bar{s}_I + B_J \bar{s}_J + B_D \bar{s}_D + C_I \bar{v}_I + C_J \bar{v}_J + D\delta \quad (15)$$

where $\bar{s}_I \equiv N_I^{-1} \sum_{i=1}^{N_I} s_i$, $\bar{s}_J \equiv N_J^{-1} \sum_{i=1}^{N_J} s_j$ and $\bar{s}_D \equiv N_D^{-1} \sum_{i=1}^{N_D} s_d$ are the average signals of indirect bidders (I), direct bidders (J) and dealers (D), and δ is the non-price contingent demand. The equilibrium pricing coefficients A, B_I, B_J, B_D, C_I, and C_J that solve each model differ by model and are reported in Appendix A.

Standard competitive market models often have simple solutions for the price coefficients, but this is not true in our setting. The complication is two-fold: 1) there are strategic agents whose demands are not linear in the coefficients of the price function and 2) shared signals are correlated with price information. Both sources of complexity are essential to understand information sharing affects auction revenue. Appendix A proves that, in the two extremes of Chinese walls and perfect information sharing, an equilibrium exists.

Through the lens of the pricing equation (15), the primary effect of information sharing in this model is higher auction revenue as sharing information leaves all investors better informed. Investors who perceive an asset to be less risky will hold it at a lower risk premium, or at a higher price. A lower risk premium is a less negative A. We see in the solution that this risk premium ($-A$) decreases when information is shared and uncertainty is lower. While this type of effect shows up in many imperfect information asset pricing models, it offers a new perspective on how restricting sharing of information affects auction revenues. From a policy perspective, this effect is equally important and has largely been 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 indirect and direct bidders and dealers disappears. In other words, eliminating all information sharing effectively eliminates intermediation as well.[10]

Auction Revenue Because the supply of the Treasury asset is normalized to one, price and auction revenue are the same. Our objective is to determine how expected revenue varies with information sharing. Since private values and non-price contingent bids are both mean zero, expected revenue is linear in the unconditional mean of the asset payoffs μ: $A + B_{total}\mu$, where $B_{total} = B_I + B_J + B_D$.

[10] The finding that there is no longer any meaningful distinction between a dealer and a non-dealer large investor is reflected in the fact that in the price formula, if the number of dealers and large investors is equal and the dealers do not face a minimum bidding requirement, then the coefficients on the signals of dealers s_d and the signal of large investors s_j are equal as well.

3 Mapping the Model to the Data

The model has twelve parameters. We map the model to the data by fixing the number of agents (three parameters) and then calibrating the remaining nine parameters to twelve moments from Treasury auction allotments and market pricing data. The rest of this section provides detail on the calibration. 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 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.[11]

In each auction, price-contingent (called "competitive") bids specify a quantity and a rate, or the nominal yield for note securities. "Non-competes," is treasury parlance for non-price contingent bidders; they specify a total amount to purchase at the market-clearing rate. As discussed in more detail below, we are largely focusing on the price-contingent bidding, which accounts for more about 99% of allotted amounts. Price-contingent 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.

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-price contingent bids, and then price-contingent bids in ascending yield 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 yield, at which bids were accepted.

We first discuss auction participation, as measured by market share of quantity won by type of bidders based on auction results data published by the U.S. Treasury. For each maturity, we compute the mean share of securities allotted to primary dealers, direct and indirect bidders as a fraction of all price-contingent bids. Price-contingent bids, the sum of direct, indirect and primary dealer bids, account for 99% of all bids in the auction.[12] More

[11] 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.

[12] We thus exclude, amounts allotted to the Fed's own portfolio through roll-overs of maturing securities, which are an add-on to the auction. Starting in April 2008, Treasury begun releasing noncompetitive results ahead of the auction close, meaning that just like the Fed's own bids, non-competitive bids are known to other investors.

precisely, primary dealers bidding for their own account, are the largest bidder category at auctions accounting for 53% on average of all price-contingent bids. Indirect bidders are the second largest at about 37% and direct bids account for about 10%.

The next set of calibration moments are the mean and variance of auction prices and secondary market values. Measuring payoff risk (speculative risk) is central to our calibration strategy. Importantly, risks faced by speculative Treasury bidders is different from those faced by investors underwriting corporate bonds. Because the U.S. sovereign secondary market is deep and liquid, Treasury investors can hedge issuer-specific risks by shorting already-issued securities. Newly issued government securities do, however, carry a liquidity premium relative to already-issued securities. Investors' demand for specific issues is the key determinant of these liquidity differences. As a result, key underwriting risks for bidders are issue-specific rather than issuer-specific.

When calculating the first and second moments of auction prices and secondary market payoffs (p and f), 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, namely the expected path of short-term interest rates and term premia. In our calibration, 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 also 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 speculative payoff 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 2, the average revenue from selling a new coupon-bearing security is about 37 basis points higher than the replicating portfolio formed using outstanding securities. Thus, we calibrate the model to have this average asset payoff. This excess revenue is positive across all maturities. This is the well-known "on-the-run" premium (Lou, Yan, and Zhang, 2013; Amihud and Mendelson, 1991; Krishnamurthy, 2002). Appendix B details exactly how we calculate payoffs and explores other possible ways of hedging the interest rate risk.

Table 1: Calibrated parameters: The valuation-related parameters in the Table (μ,$\tau_f^{-\frac{1}{2}}$,$\tau_\varepsilon^{-\frac{1}{2}}$,$\tau_{vI}^{-\frac{1}{2}}$,$\tau_{vJ}^{-\frac{1}{2}}$,$\tau_\xi^{-\frac{1}{2}}$ and χ) are expressed in basis points. The standard deviation of the non-price contingent parameter $\tau_\delta^{-\frac{1}{2}}$ is measured as a share.

Directly Measured					Jointly Estimated						
μ	$\tau_f^{-\frac{1}{2}}$	N_I	N_J	N_D	$\tau_\varepsilon^{-\frac{1}{2}}$	$\tau_{vI}^{-\frac{1}{2}}$	$\tau_{vJ}^{-\frac{1}{2}}$	$\tau_\xi^{-\frac{1}{2}}$	$\tau_\delta^{-\frac{1}{2}}$	ρ	χ
40.36	72.77	200	50	20	6.46	4.61	0.40	0.93	0.01	448.66	55.37

Calibration Table 1 lists the thirteen parameters in the model with the exception of the constant term χ_0 in the dealer's private value that plays no difference in the demand functions.[13] The three parameters that govern the number of market participants (N_D, N_I and N_J) are chosen directly to approximate the observed number of dealers (about 20) and produce 10 clients per dealer. Since indirect bidders take down almost four times as much of the auction as direct bidders do, we set the number of direct bidders $N_J = N_I/4 = 50$.

Of the remaining 9 parameters, two can be matched directly to data. The mean and standard deviation of the common value (μ and $\tau_f^{-1/2}$) correspond directly to the first and second moments of the secondary market payoffs, described above. Thus, two moments are used to match two parameters.

Risk aversion closely governs the price of the Treasury and thus auction revenues. For a given set of other parameters, we can calculate analytically the value of risk aversion that matches revenues. Thus, for each set of candidate parameters, risk aversion ρ is chosen such that the expected revenue (p) of the auction matches exactly observed data.

We estimate the remaining six parameters jointly to provide the best fit to the aggregate moments in Table 2. Note that there are more moments than parameters. Since the model is not an exact representation of reality, it cannot match all the moments. The over-identifying moments provide extra information to guide parameter estimation and gauge

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

Table 2: Calibration targets and model-implied values. Prices and excess revenues are all expressed in basis points. N_I, N_J and N_d are set directly, resulting in three over-identifying moments.

	Data	Model
Expected speculative payoff (μ)	40.36	40.36
Stdev speculative payoff ($\tau_f^{-1/2}$)	72.77	72.77
Expected revenue (p)	36.74	36.74
Stdev revenue	71.97	73.08
Price constant (A)	-1.56	-3.55
Price sensitivity to fundamental (B)	0.97	1.00
Pricing Error Stdev. (σ_ϵ)	30.19	8.03
Indirect share (%)	36.89	34.83
Dealer share (%)	53.31	56.53
Volatility of dealer share	14.50	7.39
Direct share (%)	9.80	8.64
Volatility of direct share	8.56	18.41

the fit of the model to the data.

The estimation objective function for these remaining six parameters includes the variance of the auction revenue (or the on-the-run premium at the auction), the mean allotted share to primary dealers ($\sum_{d=1}^{N_D} q_d$), indirect bidders ($\sum_{i=1}^{N_I} q_i$), and to direct bidders as well as the variance of the direct and dealer share. In addition, we estimate the empirical counterpart of the equilibrium pricing equation (15):

$$p_t = \gamma_0 + \gamma_1 f_t + \epsilon_t, \tag{16}$$

where, p_t and f_t are "on-the-run premiums" at the auction and issuance dates, respectively. From (15) and (16), we see that excess revenues are positively correlated to the fundamental value on issue date ($\gamma_1 = B_I + B_J + B_D > 0u$). The loading of about one suggests that the auction price reflects expectations for secondary market value, nearly one-for -one. The estimate of σ_ϵ is the variance of the residual from that regression.

The model moments are computed by drawing 100,000 realizations of the fundamental f, all the signals S_i, and non-price contingent demands δ, and calculating the average equilibrium outcomes. 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 (12) to ensure that we find the equilibrium pricing coefficients. At our solution, the average violation of the market clearing condition is about 10^{-12}. We use multiple starting points to ensure that the maximum is a global one.

Estimated parameters Looking at the values of the estimated parameters in Table 1 we note the standard deviation of the signal s_i ($\tau_\varepsilon^{-\frac{1}{2}}$) observed by investors and dealers is estimated to be about a tenth as large as the unconditional standard deviation of the fundamental f ($\tau_f^{-\frac{1}{2}}$). The estimated volatility of private values of direct and indirect bidders increase the standard deviation of their respective shares, or $\tau_{vI}^{-\frac{1}{2}} > \tau_{vJ}^{-\frac{1}{2}}$. Finally, the size of the dealer's private value, which is the same for all dealers, is ten times as large as τ_{vI} and of the same order of magnitude of the unconditional variation in f. This suggests that the model estimates a sizable private value component for dealers.

4 Results: Effects of Information Sharing

We examine two forms of information sharing. We first study information sharing between dealers and clients by varying the precision of the dealer signal to their clients, without allowing dealers to communicate amongst each other. Then, we hold the precision of client communication fixed and vary the number of other dealers that each dealer shares information with. In both cases, we find that information sharing increases auction revenues as well as revenue volatility. The surprising finding is that small investors dislike, as a group, when dealers share more precise information with them, but sometimes 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 can reduce information asymmetry and improve risk sharing.

Since the quantity of auctioned securities is fixed and normalized to 1, the auction price and auction revenues are the same. In the plots that follow, we study expected excess revenues varying one exogenous parameter at a time. In each exercise, all parameters other than the one being varied, are held at their calibrated values from Table 1.

4.1 Information Sharing Raises Auction Revenue

Information sharing – of either kind – raises auction revenue. Indeed more information sharing makes the average bidder better informed, which in turn makes Treasuries less risky to the average investor, eliciting stronger bids, resulting in higher auction revenues. However, the quantitative revenue effects of client-sharing and dealer-sharing are quite different. The left panel of Figure 2 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

Figure 2: Dealer Information Sharing. Left: dealer information sharing with clients; right: dealer information sharing with other dealers. In the left panel, the horizontal axis shows the precision of the dealer signal τ_ξ from zero (no information sharing) to infinity (perfect information sharing).

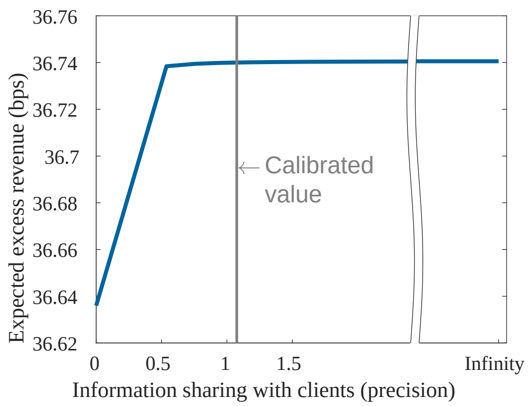

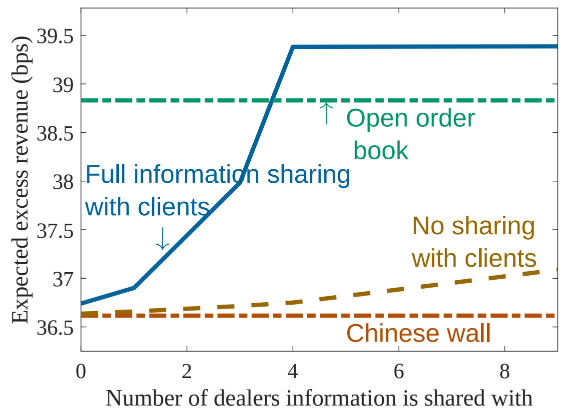

information sharing). More information sharing means that dealers reveal their information $\bar{s}_d$ with less noise to their clients. In the absence of inter-dealer talk ($\psi = 1$ in left panel), moving from no sharing to perfect information sharing with clients results in a very small increase in expected revenue of a tenth of a basis points. The vertical line on the plot represents the amount of client information sharing implied by the model calibration. This calibrated status quo corresponds to revenues of 36.74 bps. How much can client information sharing raise revenue? Given an annual Treasury issuance of about \$8 trillion, the model implies that going from no sharing to perfect sharing with clients would increase total auction revenues by a modest \$80 million. Furthermore, the model suggests no revenue gain from encouraging further client information disclosure compared to the calibrated status-quo. Indeed the revenue curve to the right of the current level of information sharing is flat. Importantly this result assumes that dealers do not share information with each other. The model suggests that without dealer sharing, the benefits of client information sharing are limited.

The biggest revenue gains arise when both types of information sharing take place as shown in the right panel of Figure 2. As shown by the brown dashed line, if dealers do not share information with their clients, the revenue benefits of dealer talk are small, less than 0.5 basis points. In contrast, the combination of inter-dealer sharing and sharing information with clients is a powerful revenue generator. In the presence of full client information sharing (solid blue line), increasing the number of dealers with which each dealer shares information, auction revenues increase by almost 3 basis points.[14] Given the current level

[14]Since we assume dealers are symmetric, we need the number of dealers in an information-sharing

of client information sharing, the additional revenue from allowing all dealers to share with four (or more) other dealers amounts to \$ 2.4 billion.

In unreported results, we also find that both types of information sharing reduce the variance of auction revenue, but that this effect is quite small (order of less than a basis point). The reason for the small change in variance is because two countervailing effects nearly offset each other. In a model with only common values, an increase in information sharing would increase revenue volatility: As investors put more weight on their more informative signals, the auction clearing price becomes more sensitive to changes in the fundamental value f. With correlated values, when information sharing makes the auction price more responsive to the speculative return, it also becomes less responsive to private values. The result is small changes in revenue variance.

In additional analysis, we find that when prior uncertainty about the future value of the asset is high (precision τ_f is low), or if the variance of non-price contingent bids grows, information sharing raises revenue by more. The reason is that both make bidders more uncertain ex-ante. When bidders are more uncertain, there is more scope for information sharing to reduce risk. All else equal, a reduction in risk prompts bidders to bid more for the asset.

One proposed policy is an open order book. An open book allows all bids to be observed by all market participants. The expected revenue of an open order book is just like the revenue of full dealer information sharing. This takes into account the fact that when all bids are observed, bidders behave strategically and try to manipulate their bids to affect others' beliefs. If instead of a strategic bidder market, this were a large, competitive market, the benefits of an open order book would rise further to 40.36 basis points. That's a 3.6 basis point increase over the status quo, corresponding to an additional \$2.9 billion in Treasury revenue.

The role of minimum bidding requirements Primary dealers are required to be consistent, active participants in Treasury auctions. While rules have evolved over time, today, primary dealers are expected to bid at all auctions an amount equal to the pro-rata share of the offered amount, with bids that are "reasonable" compared to the market. The inclusion of minimum bidding penalties in dealers' private values is realistic but also helps to calibrate the model in a sensible way. Absent a reason for a high private value and given common risk aversion for all bidders, it would be hard our model to explain why dealers bid for so much of the auction. One way to see this is by looking at the estimated level of

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.

$\chi = 55$ in Table 1, which is about as large as the standard deviation of the fundamental in the model ($\tau_f^{-\frac{1}{2}} = 73$). In words, the model needs a private value component for dealers which is of the same order the common component to rationalize the observed shares. In unreported results we show that the main effect of minimum bidding requirements is that a higher penalty (χ) raises expected revenue by boosting demand by primary dealers, since dealers are incentivized to bid more aggressively. But bidding requirements leave the effect of client information-sharing on revenue and utility unchanged.

4.2 Bid Shading and Signal Jamming

Since our bidders have price impact and are strategic, they optimally use their bids to influence the auction-clearing price ("bid shading"), which is the central focus Hortaçsu, Kastl, and Zhang (2016). Our bidders can also influence others' beliefs, so as to impact others' bids ("signal jamming"). In this section we provide additional details on bid shading and jamming, their impact on expected revenues and how they interact with information sharing.

Each speculator's and primary dealer's bid depends on her expected private and common value (numerator of (13)), and on the sensitivity to that expected value (denominator of (13)). The sensitivity (denominator) has a risk aversion $\rho V[f|S_i]$ term. If investor i is more risk averse, then she bids for a smaller position in the asset. The second term is dp/dq_i. This is a strategic effect that captures her ability to influence the auction price. We break that strategic effect into two parts. One is bid shading (BS).

If the bidder reduces their demand by one unit, and others' bidding best responses stayed fixed, this effect captures the change in the auction-clearing price. Bid shading is the part of dp/dq_i that would remain, even if, when an investor reduces his bid, others do not make inference from the slightly lower price: $BS = dp/dq_i|_{\beta_p=0}$, where $\beta_p = 0$ means that other bidders are not drawing inference from the price signal. The other part of the strategic effect is the ability to influence others' beliefs through prices. This is signal jamming: $SJ = (dp/dq_i)^{-1} - BS^{-1}$. We define these two terms formally in terms of model primitives in equations (103) and (101) in the Appendix.

The extent of bid shading and signal jamming depends greatly on how and how much information is shared. Information sharing, of either type, reduces both bid shading and signal jamming, for each investor. When dealers share more information, either with their clients or with each other, everyone is better informed. When all bidders are better informed, small changes in price have little additional information value. Thus, the price impact of a trade is reduced. Instead, highly informed traders infer that if the price is

Figure 3: Bid Shading and Signal Jamming Revenue Effects. This Figure shows the effect of removing either bid shading ($dp/dq_i|_{\beta_p=0}$) or signal jamming ($(dp/dq_i)^{-1} - BS^{-1}$) from bidders' demands, one at a time. Each line represents an average price (revenue). Left: dealer information sharing with clients; right: dealer information sharing with other dealers. Formulas for bids without jamming and shading are (101) and (103) in the Appendix

(a) With Client Information Sharing **(b)** With Dealer Talk

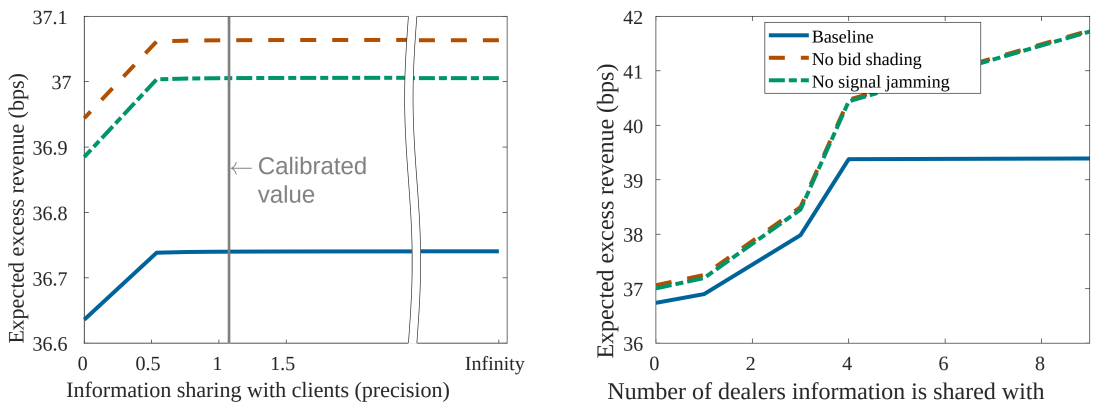

high, someone bid high for non-informational reasons and they will buy less. Because of this informed bidder response, price impact is reduced. The maximal effect of client information sharing on dp/dq is only equal to a tenth of a percent. For dealer talk, the effect is on the order of one percent.[15]

What is perhaps most interesting is the effect that each component has on Treasury auction revenue. Figure 3 shows that, not surprisingly, removing bid shading and signal jamming causes bidders to bid more. But as shown in panel a, the size of the revenue increase hardly depends on how much information dealers share with clients. In contrast, when information is shared between dealers, bid shading and signal jamming have a much larger effect on expected revenues. The reason is that highly-informed traders bid for larger positions. Shading a large bid has a larger revenue effect than shading a smaller one. As more dealers talk, we see that gap in revenue widens to about 2.5 bps, between auctions with and without strategic bids (Figure 3, panel b).

In sum, while revenues increase with dealer talk, this increase is partially offset by strategic bidding. By making bidders' beliefs more precise and more correlated, the effect of market power on auction revenue is strenghthened.

[15]The absolute magnitude of these effects is difficult to interpret because they depend on the meaning of one auction unit. Here, we define normalized offering amounts to one unit, so only relative comparisons within the model are meaningful. As we discuss next, revenue effects have interpretable units.

4.3 Client vs. Dealer Information Sharing: Utility Effects

So far, we studied effects of client and dealer information sharing on expected auction revenues. In this section, we study effects on bidders' welfare. A key insight of the model is that client and dealer information sharing are quite different for bidders' welfare. The reason for this opposite effect lies in how each type of information sharing affects information asymmetry and risk-sharing.

It would be logical to think that if dealers are passing along more of their information to their clients, that clients would be happy about that. That turns out not to be the case. Figure 4 plots investors' utility levels relative to the "Chinese wall" benchmark of no information sharing between dealers and customers. Panel (a) shows that bidders' utility declines when dealers share more information with them. 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 prices, they transfer more surplus to the issuer (Treasury).

Figure 4: Clients Lose from Client Information Sharing, Can Gain from Dealer Talk. Both panels plot the change in clients' expected utility from information sharing, as a fraction of the utility each type gets in the Chinese wall (no sharing) equilibrium. Client information sharing makes allocations more heterogeneous. This reduces client expected utility. Dealer and client information sharing reduces this asymmetry and can improve utility.

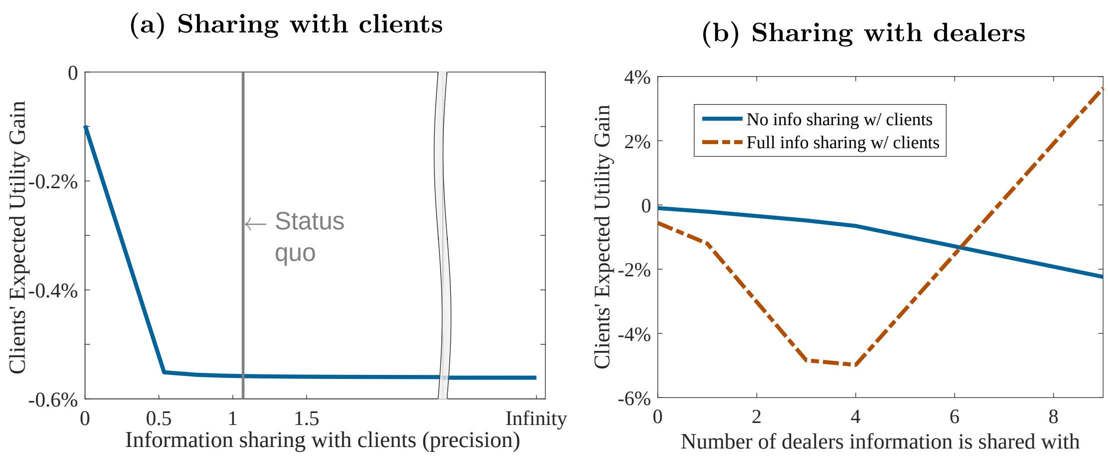

The other mechanism at work is that sharing information with clients increases information asymmetry. When dealers share little information with clients, clients' beliefs are not very different. They all average their priors with a heterogeneous, but imprecise, private signal. Because private information is imprecise, beliefs mostly reflect prior information, which is common to all investors. When different dealers transmit different signals, and investors get a more precise dealer's signal, they weigh it more heavily in their beliefs; this makes investors' beliefs differ. This increase in information asymmetry makes ex-ante similar investors hold different amounts of securities ex-post. Asymmetric information pushes the asset allocation further away from the efficient diversified benchmark. Because investor preferences are concave, investment asymmetry hurts average investor utility. In short, information reduces risk sharing, which reduces utility.

One might also expect that when dealers share information with each other, investors are harmed. When dealers share information with each other and do not pass this better information on to the clients, clients do suffer as shown by the blue line in panel (b) of Figure 4. However, if dealers share what they know with their clients, clients can benefit from inter-dealer talk as shown by the right side of the dashed line in panel (b). When only a few dealers talk, the limited dealer talk increases belief and investment dispersion, just like client information sharing. But when many dealers exchange information, 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. The similarity of these signals offset the dispersion increase arising from more precise information. When dealers share information with four other dealers, belief and investment dispersion stabilizes. When clients get more precise (shared) information from their dealers, but do not face the downside of more asymmetric auction outcomes, their utility rises.

When information is shared, two features of the information environment change simultaneously. First, the agents involved have more precise forecasts of post-auction appreciation. This creates the increase in auction revenues, which is common to both types of information sharing. Second, there may be more or less market-wide forecast disagreement, depending on how information is shared. Client (dealer) sharing is like observing a private (public) signal in a strategic game. Client information sharing is like more private information because it pushes beliefs further apart. The result that more informative private signals can increase information asymmetry and thereby reduce utility is the same force that is at work in Hirshleifer (1971). In contrast, dealer sharing makes information sets more similar or more public. A large literature examines the different effects of public and private information in strategic environments. In some of those environments, coordinated actions are socially costly (Morris and Shin, 2002; Angeletos and Pavan, 2007); therefore, public signals are bad because they enable this costly coordination. In other environments, like Lucas (1972) island models, coordination is socially beneficial (Woodford, 2011). In the island models, like in this model, public information is good: In a way, dealer talk is equivalent to merging some of Lucas' islands together.

Profits of non-price contingent bidders Whenever information is shared, speculative bidders become better informed, surplus is transfered from bidders to the issuer and profits of non-price contingent bidders decline. Who are these bidders that lose out? Some non-price contingent bidders are small retail investors. Many are bids placed by the New York Fed on behalf foreign and international monetary authorities (FIMA) that hold securities in custody at the Fed.[16]

4.4 What if Information Sharing Enabled Collusion?

One reason why some call for curbing information sharing is that dealers who share information may also collude. Many textbook analyses show 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 also 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? Without collusion, dealer information sharing increases expected revenue because of the reduction in investors' risk (Figure 2, panel b). With collusion, the effect depends on client information sharing. When no information is shared with clients, investors don't perceive a risk reduction, and revenue declines as collusion increases. With client information sharing, there is a small region in which the joint effect of information sharing and collusion is to increase revenue slightly, before dropping below the "Chinese walls" benchmark (Figure 5).

When expected revenues decline, investors utilities are higher, while taxpayers are worse off. In other words, if information enables collusion, the issuer is the main loser. From an auction revenue perspective, this does not necessarily mean that prohibiting information-sharing would be optimal. If anti-collusion laws could be effectively enforced, without prohibiting the sharing of information, that would be the best possible outcome for Treasury revenue.

Lying about dealer talk Perhaps not all investors know that dealers swap order flow information with other dealers. Of course, one could enforce laws about disclosure of information practices, without prohibiting information sharing. But agents understanding of others' strategies do matter in the results that we have discussed so far. When a set of dealers share information and collude but others are not aware, auction revenue falls

[16] 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.

Figure 5: Collusion reduces revenue. Average equilibrium auction revenue, assuming that when ψ dealers share information, they also bid as one. These results differ from previous figures because here, varying information-sharing along the x-axis also varies the extent of collusion.

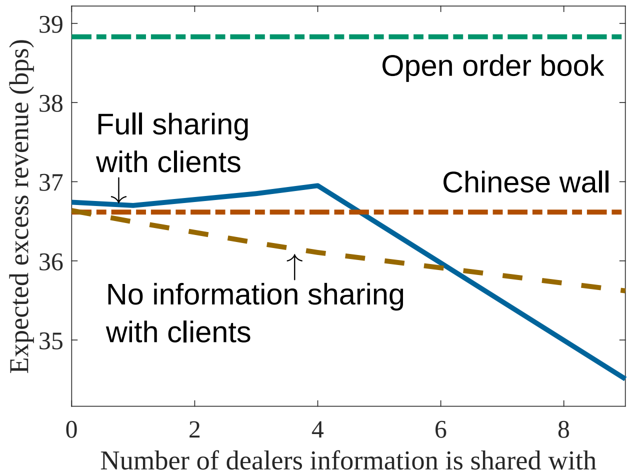

by more than in Figure 5.[17] When inter-dealer information sharing is undisclosed, even if information is subsequently shared with clients, revenue declines. This is because if clients are not aware that their information is very precise, they do not bid aggressively. Thus hidden information sharing fails to raise auction revenue.

5 How and When to Bid?

So far, the paper takes as given that all bidders participate in the auction, some are dealers, and some bid directly, while others bid indirectly. In reality, all of these margins involve choices. While a full analysis of every relevant margin would require more than one paper, an extension of the model provides insights about each choice, that we hope might invite further analysis.

A distinguishing feature of U.S. Treasury auctions is that they are mixed auctions: Any investor can either place an intermediated bid through a primary dealer, or bid directly. When choosing how to bid, information sharing arrangements matter. In order to understand these effects of information sharing, explore the choice of a single bidder, deciding whether to bid directly or indirectly (through a dealer). Once clients have a direct/indirect bidding choice, it becomes clear why dealers may share some of their information with clients. Absent any information sharing, clients would have no incentive to bid with them. Similarly, dealer's ability to use client information is what incentivizes them to be dealers.

[17] Details of this model variation and its results are reported in Appendix A.5.

Finally, we discuss an investor's choice to bid into the auction, or to purchase securities before the auction in the when-issued (WI) market.

Model: Choosing to Bid Directly or Indirectly Consider one investor choosing between bidding directly or indirectly through an intermediary (the dealer). The investor's choice of how to bid affects the information structure of that investor, its dealer, other investors bidding with that same dealer, and the information content of the price $s(p)$.[18] If investor i bids indirectly, through dealer d, the model and the signals are as discussed in the baseline specification. But when the investor chooses to bid directly on his own behalf, he observes only his own signal, private value and the price information: $S_i = [s_i, v_i, s_i(p)]$. The order flow signal of the dealer that investor j refused to bid through, now has a signal based on $N_I/N_D - 1$ clients' order flow signals. This dealer also knows one more piece of important information: that one investor, an investor who typically bids through him, decided to bid directly.

Solving the model with an endogenous direct versus indirect bidding decision introduces a technical challenge. The decision to bid directly or indirectly itself becomes a signal. We assume that the dealer who would intermediate this trade observes the investor's bidding decision and transmits this information to clients, with noise. If the investor bids through the dealer, the dealer observes his bid, as before. If the investor bids directly, the dealer learns that the investor's signal must lie in one of two disjoint regions of the distribution. This is problematic because doing Bayesian updating of beliefs with truncated normals would require involved simulation methods. Embedding that updating problem in our solution would render it intractable.

We circumvent this problem by constructing an approximating normal signal. Through simulation, we first estimate the mean and variance of the investor's signal, conditional on choosing direct bidding. Then, whenever the investor chooses to bid directly, the dealer, who would have intermediated that trade, makes inference from the direct bidding decision. That dealer observes a *normally distributed* signal, $s_q = f + e_q$ (q for quit) with the same mean and variance as the signals of the simulated direct bidders. This normal signal is included in the precision-weighted average signal of dealer d'. In appendix B we show how this signal can be used to construct a precision-weighted average signal of dealer d' and derive the updated equilibrium pricing condition.

[18] There may well be fixed costs associated with bidding directly for large investors, such as registering with the online direct bidding system known as TAAPSLink for large investors or setting up one's own trading desk. We abstract from such costs because they are difficult to quantify and do not change the main price asymmetry result.

Client Information Sharing and the Direct/Indirect Bidding Decision. Examining expected utility in our model clarifies the incentives to bid directly and indirectly. Our results reveal how information sharing affects bidders' utility of bidding through dealers. This, in turn, helps to explain why information sharing with clients takes place.

When the investor chooses whether to invest through a dealer, he has seen his private signal s_i. In addition, the bidder knows his private value v_i. Thus the intermediation choice maximizes expected utility, with an additional expectation over the information that the investor has not yet observed. Computing the expectation over possible price realizations and dealer signals, but conditioning on an investor's private signal, we find that expected utility of any type of bidder i is[19]

$$EU_i = -\exp(\rho W_{0i})(1 + 2\theta_i \Delta\mathbb{V}_i)^{-\frac{1}{2}} \exp\left(-\frac{\mu_{ri}^2}{\theta_i^{-1} + 2\Delta\mathbb{V}_i}\right). \tag{17}$$

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

$$\mu_{ri} \equiv \mathbb{E}\{\mathbb{E}[f|S_i] + v_i - p|s_i\}, \tag{18}$$

$$\theta_i \equiv \rho[\rho\mathbb{V}[f|S_i] + dp/dq_i]^{-1}\Big(1 - \frac{1}{2}\rho[\rho\mathbb{V}[f|S_i] + dp/dq_i]^{-1}\mathbb{V}[f|S_i]\Big), \tag{19}$$

$$\Delta\mathbb{V}_l \equiv \mathbb{V}\{\mathbb{E}[f|S_i] + v_i - p|s_i\} = \mathbb{V}[f - p|s_i] - \mathbb{V}[f|X_l S]. \tag{20}$$

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

The second term θ_i captures the main advantage of intermediation: Dealers give their clients an extra signal, which makes them better informed. Better information allows the large investor to make better bids, increasing expected utility. In Appendix A, we show that θ_i is positive and strictly increasing in the posterior precision of the asset payoff $\mathbb{V}[f|S_i]^{-1}$. Thus, intermediation improves the investor's information, which decreases variance $\mathbb{V}[f|S_i]$, increases θ_i, and (holding all other terms equal) increases expected utility.

[19] See Appendix A.6 for derivation.

The third effect of intermediation, which operates through ex-ante variance $\Delta\mathbb{V}_i$, is ambiguous and turns out to be quantitatively unimportant.[20]

Client information sharing and the existence of dealers. This utility expression reveals why information sharing is so integral to the primary dealer system. If dealers shared no information with clients, the cost of using a dealer, a reduction in μ_{ri}, would still be present because the dealer observes the client's order and may trades on this information. That is costly because it reduces the client's expected speculative profit. However, the benefit of using a dealer disappears. If clients get no information from dealers, the signal set S_i and thus the speculative risk $\mathbb{V}_i$ do not change. In such an environment with only costs and no benefits, to indirect bidding, primary dealers would cease to exist.

Of course, one could prohibit primary dealers from trading on clients' information, which is the "Chinese wall" solution we examined before. But this would also not support a dealer system. First, investors would now be indifferent between bidding directly and indirectly in the sense that both μ_{ri} and θ_i would be unaffected. Second, there are costs to being a dealer, in the form of regulatory or minimum bidding constraints. In sum, while we do not explicitly model dealers' and bidders participation decision, these results suggest that information sharing is what induces bidders to bid through dealers, and dealers to participate in the primary dealer system. In this sense, information sharing is at the core of the primary dealer system.

When-issued markets. Before each Treasury auction, investors can bet on the auction-clearing price by transacting in the when-issued (often called "WI") market. WI is an over-the-counter forward market. In a way, purchasing WI contracts is like direct bidding: An investor who bids in the WI market does not learn from a dealer's signal, implying a lower θ_i. At the same time, the investor does not reveal her order flow to a dealer who shares that information with others (implying a higher μ_{ri}). Of course, the person with whom the investor transacts will know the order and the market price will reflect it.

The decision of an investor to bid in the WI market, as opposed to the actual auction, depends on risk preferences and on information sharing. Investors who purchase securities

[20]When the large investor trades through a dealer, his uncertainty $\mathbb{V}[f|S_i]$ declines. From equation (20) we can see that this increases the ex-ante variance of the expected profit $\Delta\mathbb{V}_i$. This is because more information makes the investor's beliefs change more, which means a higher ex-ante variance. This change in $\Delta\mathbb{V}_i$ has two opposing effects on expected utility. First, an increase in $\Delta\mathbb{V}_i$ increases the exponential term in equation (17), which decreases EU_i. This effect arises because the large investor is risk averse and higher $\Delta\mathbb{V}_i$ corresponds to more risk in continuation utility. The second effect is that an increase in $\Delta\mathbb{V}_i$ reduces $(1+2\theta_i\Delta\mathbb{V}_i)^{-\frac{1}{2}}$, which increases EU_i. 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.

in the WI market limit auction uncertainty by purchasing newly-issued securities at a predetermined price. But this price will reflect in equilibrium risk compensation on the part of the sellers. Indeed, securities outstanding in the WI market are in zero net supply, meaning that whenever an investor is long a WI, another will be short. The other feature that differentiates WI from auction bidding is that the opportunity to bid through a dealer allows an investor to benefit from information sharing.

WI activity may also affect the benefits of information sharing. WI market commitments affect investors' private values. When private values are more important, shared information about common values is less important. One might think that, because the WI market is often an accurate forecast of the auction-clearing price, this would matter as well. However, since our model allows bidders to condition on every possible price, they have no use for a price forecast. They simply form bids by asking the question, if p were the auction clearing price, what would I learn and how would I want to bid? Continuous price-contingent bids make price forecasts redundant.

6 Supporting Evidence: Correlated Values and Informed Bidders

This section's evidence supports three key features of the model: the common value assumption, the private value assumption and the information content of signals about the common value. Although we do not have bidder-level data, our modeling strategy allows us to infer values and information from publicly available data on secondary market outcomes and risk premia.

Much of the previous literature on Treasury auctions assumes that valuations for auctioned securities are private, an assumption that is supported by the findings of (Hortaçsu and Kastl, 2012). There are a number of reasons why bidders may have private values for newly-issued Treasury securities. First, buy-and-hold bidders do not resell in the secondary market and their private valuation will reflect a number of bidder specific factors, including their duration-hedging needs, their investment horizon and regulatory requirements, such as liquidity coverage ratios, which tilt preferences towards Treasury securities. Some investors may also be entering the auction having sold the newly-issued security in the WI market, in which case their valuation would be influenced by the the magnitude of the short that needs to be covered. Finally, primary dealers may be required to bid for the auctioned securities over and beyond their secondary market value in order to maintain their primary dealer status.

Table 3: Regression of $f - p$ on speculative bidders' auction share $\bar{q}$. The dependent variable is the post-auction appreciation measured as the difference between the value of the interest-rate neutral portfolio on issue date (f) and on auction date (p) expressed in basis points units. speculative bidders are bids by primary dealers and direct/indirect bids of domestic investors. Non-PD speculative shares is domestic bidders except primary dealers. All shares are measured in percent. Robust standard errors reported in square brackets. *** significant at 1%, ** significant at 5%, *significant at 10%.

	(1)	(2)	(3)	(4)
Spec Share		0.77***	0.92***	
		[0.21]	[0.26]	
PD Share				0.97***
				[0.26]
non-PD Spec Share				0.70**
				[0.30]
Const	2.71**	-57.63***	-73.59***	-2.01
	[1.36]	[16.05]	[21.56]	[13.38]
Adj. R2	0.00	0.05	0.24	0.24
Obs.	494	494	494	494
Month FEs?	No	No	Yes	Yes
Tenor FEs?	No	No	Yes	Yes

On the other hand, a common value component is also likely present for two reasons. First, many speculators bid in the auction, in order to resell in the secondary market, at a profit. These speculators share a common profit objective, and thus a common valuation, given by the secondary market price. Data on inventory holdings of primary dealers, which are the largest participants in Treasury auctions, suggest that they typically resell all of their holdings by the first week of trading (Fleming and Jones, 2015). Similarly, hedge funds, who are active bidders in auctions (see Wall Street Journal, 2015, for example) hold securities only briefly because of their limited capital. A second reason for correlated values is that any investor who bids in the auction has the option to wait. Any investor can purchase the newly-issued security in the secondary market, on issuance date at the price f. If expectations about f are very low, even a buy-and-hold investor should wait to purchase in the secondary market, lowering demand at the auction.

Model predictions We turn next to three model predictions.

Prediction 1. *If values are correlated and signals s_i are informative about the fundamental f, The auction clearing price p is positively correlated with the fundamental f.*

The common value assumption is what causes the auction clearing price to be correlated with the fundamental value. This correlation arises because the price coefficients on signals,

B_I, B_J and B_D in equation (15), are all positive.

The next result is that private values lower the correlation of bidders' bids with the post-auction appreciation payoff f.

Prediction 2. *Consider an investor k for whom the private value is less important than others, or $\tau_{vk}^{-1} < \tau_{vJ}^{-1}, \tau_{vI}^{-1}$. Then a higher allotted share q_k, at any given price $\bar{p}$ has a higher correlation with post-auction appreciation $f - p$.*

From the first order condition (13) the quantity $q_j(p)$ demanded by each bidder is proportional to the sum of the expected return $\mathbb{E}[f|S_j] - p$ and the private value v_j. If agent k has a less volatile private value v_k component, she is more of a speculator. Then k's first order condition implies a lower correlation (or regression loading) of post-auction returns $\mathbb{E}[f|S_k] - p$ and allotted shares q_k.

Prediction 3. *Consider an investor k that has a more precise signal S_k about the fundamental f, then a higher allotted share q_k predicts higher post-auction appreciation $f - p$.*

When the signal s is informative about the common value f, the correlation between realized $f - p$ and the expected $E[f - p|S_i]$ rises. Information aligns beliefs with outcomes. The quantity q_i, demanded by any bidder is a linear function of $E[f - p|S_i]$. Thus, the covariance of q_i and $f - p$ is higher for a better informed investor.

Supporting evidence Evidence supporting Prediction 1 was provided in Section 3. Indeed to calibrate the A and B parameters in the model-implied equilibrium equation of Proposition 1, we estimated parameters of the regression (equation 16) of p on f. Consistent with the common value component we find a highly statistically significant coefficient, which is nearly equal to one. In words, the auction price p moves (on average) nearly one-for-one with the value of the fundamental f.

Predictions 2 and 3 imply a positive correlation between post-auction returns and allotted shares of investors that either weigh their common values more or that have more precise signals. Thus observing a higher correlation between an investor's allotted share and post-auction returns does not identify its type. We thus bring additional information from prior literature to identify investors that are more likely to possess stronger common value components or more precise signals. We then study how allotted shares of these investors' types are related to post-auction appreciation.

In reality, the importance of private values can differ across investors. While Treasury auction results provide information on types of bids (direct, indirect, primary dealers, non-price contingent), Treasury also releases data on allotted shares by investor types (dealers, investment funds, retail, depository institutions, pension funds and foreign investors). Prior literature contributions assume that foreign holders of U.S. Treasuries, which are largely official investors, hold Treasuries for non-speculative motives, for example, related to their foreign exchange policies and their respective domestic economic conditions. For example work by Bernanke, Reinhart, and Sack (2004) and Warnock and Warnock (2009) on quantitative easing policies and the "savings glut hypothesis" make the (stronger) assumption that foreign official purchases of Treasuries are completely exogenous to U.S. economic and market conditions. Furthermore Krishnamurthy and Vissing-Jorgensen (2012) provide supporting evidence for this hypothesis in the context of the convenience yield. Based on these findings, we define a speculative share by subtracting the share allotted to foreign investors.

We therefore estimate the relationship between this speculative share and post-auction 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 Table 3, the value of the newly issued security appreciates on average 2.7 basis points between the auction date and the issue date (column 1).[21] As shown in the second column, this appreciation is higher when the share of speculative bids is higher (column 2), with a highly statistically significant coefficient (t-stat = 3.5, column 2). The point estimate of 0.77 implies that a 10% increase in the speculative share (Std. Dev. = 9.1%) is associated with a sizable positive effect of 7.7% in $f - p$ (Std. Dev = 30%) This empirical result is stronger after including month and tenor fixed effects (column 3), which supports Prediction 2.

With respect to Prediction 3, prior literature (for example Hortaçsu and Kastl, 2012, in the context of Canadian auctions) find that dealers are at an informational advantage. This is also true in our model in which dealers share a noisier signal with clients meaning that dealers are better informed about f. Consistent with Prediction 3, post-auction appreciation is increasing with the allotted share to primary dealers, which display a larger coefficient vis-a-vis other speculative bidders. The coefficient on the primary dealer share is 0.97 versus a coefficient of 0.70 on the speculative share excluding primary dealers (p-val=.11).

Importantly neither the speculative bidders or the primary dealer results in Table 3 are the mechanical result from higher demand. When speculative 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 and speculators to buy more securities.

[21] This average post-auction appreciation estimate is consistent with the findings of Lou, Yan, and Zhang (2013). See Data Appendix for a full set of summary statistics on these variables.

7 Conclusion

Recent news about dealers sharing clients' order flow information with other clients or dealers has prompted calls to restrict financial intermediaries' use of order flow information. The need to prohibit collusion and misleading clients about information sharing are quite clear. But when collusion does not occur and information sharing is common knowledge, the gains and losses of information sharing are not as apparent. Using data from U.S. Treasury auctions, we calibrate a structural auction model, with correlated values, to quantify the costs and benefits of information sharing both between dealers as well as between dealers and customers.

To analyze the impact information sharing, this paper brings new features to empirical auction models (see Kastl, 2017, for a review): an uncertain common value component, risk aversion and of course information sharing. These new components offer a new way to estimate the role of information in auctions. Instead of relying on bid-level dispersion, whose measurement relies on confidential data, we back out bidders' uncertainty and average bidders' information from measured risk premia. This new approach is essential to understanding information sharing. If valuations are private, then by observing order flow one only learns about other bidders' demand. When valuations are common and uncertain, order flow is informative about everyone's valuation of the asset. With this perspective in mind, one can understand our main results. We find that the way in which information is shared matters. Dealer sharing information with other dealers about the common value makes bidders' beliefs more correlated. On the opposite, dealer sharing information with clients makes beliefs about the common value more dispersed. Beliefs about the common value are the heart of our utility results. Without the uncertain, common value, this effect, as we describe it, would not be present.

The model also shows that information sharing raises auction revenues by making bidders better informed. Dealer talk benefits issuers by raising auction revenues, but also can improve risk sharing by lowering asymmetric information. These results assume full disclosure about how information is used and no collusion; model extensions show that these can overturn the welfare effect. While our model does not detect whether collusion and misrepresentation occurs, it suggests different remedies to enforce disclosure and anti-collusion laws, perhaps without prohibiting information sharing.

Against a backdrop of policy initiatives aimed at curbing infomation sharing, other novel features of our model – investors' choice to bid directly or through dealers and dealers' minimum bidding requirements – highlight that information sharing is an integral part of Treasury auctions and the primary dealer system. Without client information sharing, clients would not want to bid through dealers. At the same time, prohibiting dealer use of client order flow data, while imposing bidding requirements on dealers, creates a dealer system with large costs and very limited benefits. Information sharing may not be optimal, depending on feasible alternatives and social welfare criterion. But it has upsides, as well as downsides.

References

Amihud, Y., and H. Mendelson (1991): "Liquidity, maturity, and the yields on US Treasury securities," *The Journal of Finance*, 46(4), 1411–1425.

Angeletos, G.-M., and A. Pavan (2007): "Efficient use of information and social value of information," *Econometrica*, 75(4), 1103–1142.

Armantier, O., and E. Sbaï (2006): "Estimation and comparison of Treasury auction formats when bidders are asymmetric," *Journal of Applied Econometrics*, 21(6), 745–779.

Babus, A., and C. Parlatore (2015): "Strategic Fragmented Markets," Working paper, New York University.

Back, K., and J. F. Zender (1993): "Auctions of divisible goods: on the rationale for the treasury experiment," *Review of Financial Studies*, 6(4), 733–764.

Bernanke, B., V. Reinhart, and B. Sack (2004): "Monetary policy alternatives at the zero bound: An empirical assessment," *Brookings papers on economic activity*, 2004(2), 1–100.

Bikhchandani, S., and C.-f. Huang (1989): "Auctions with resale markets: An exploratory model of Treasury bill markets," *Review of Financial Studies*, 2(3), 311–339.

Bloomfield, R., and M. O'Hara (1999): "Market transparency: Who wins and who loses?," *The Review of Financial Studies*, 12(1), 5–35.

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

Brunnermeier, M. K., and Y. Sannikov (2014): "A macroeconomic model with a financial sector," *The American Economic Review*, 104(2), 379–421.

Chari, V. V., and R. J. Weber (1992): "How the U.S. Treasury Should Auction Its Debt," *Federal Reserve Bank of Minneapolis*, 16(4), 1–8.

Cole, H., D. Neuhann, and G. Ordonez (2017): "A Walrasian Theory of Sovereign Debt Auctions with Asymmetric Information," Working paper.

Di Maggio, M., F. Franzoni, A. Kermani, and C. Sommavilla (2017): "The Relevance of Broker Networks for Information Diffusion in the Stock Market," Working paper, University of California at Berkeley.

Duffie, D. (1996): "Special repo rates," *The Journal of Finance*, 51(2), 493–526.

Easley, D., N. Kiefer, M. O'Hara, and J. Paperman (1996): "Liquidity, Information and Infrequently Traded Stocks," *The Journal of Finance*, 51(4), 1405–1436.

Federal Reserve Bank of New York (2016): "Policy on Counterparties for Market Operations," .

Fishman, M. J., and F. A. Longstaff (1992): "Dual trading in futures markets," *The Journal of Finance*, 47(2), 643–671.

Fleckenstein, M., F. A. Longstaff, and H. Lustig (2014): "The TIPS-Treasury Bond Puzzle," *the Journal of Finance*, 69(5), 2151–2197.

Fleming, M. J., and K. D. Garbade (2007a): "Dealer behavior in the specials market for US Treasury securities," *Journal of Financial Intermediation*, 16(2), 204–228.

——— (2007b): "Dealer behavior in the specials market for US Treasury securities," *Journal of Financial Intermediation*, 16(2), 204–228.

Fleming, M. P., and C. Jones (2015): "Dealers Positions and the Auction Cycle," Discussion paper, Liberty Street Economics.

FOMC (November 2007): "Minutes of the Federal Open Market Committee," *Federal Reserve Bulletin*.

Grossman, S., and J. Stiglitz (1980): "On the impossibility of informationally efficient markets," *American Economic Review*, 70(3), 393–408.

Gürkaynak, R. S., B. Sack, and J. H. Wright (2007): "The US Treasury yield curve: 1961 to the present," *Journal of Monetary Economics*, 54(8), 2291–2304.

He, Z., and A. Krishnamurthy (2013): "Intermediary Asset Pricing," *American Economic Review*, 103(2), 732–770.

He, Z., A. Krishnamurthy, and K. Milbradt (2016): "What Makes US Government Bonds Safe Assets?," *American Economic Review*, 106(5), 519–523.

Hirshleifer, D. (1971): "The private and social value of information and the reward of inventive activity," *American Economic Review*, 61, 561–574.

Hortaçsu, A., and J. Kastl (2012): "Valuing Dealers' Informational Advantage: A Study of Canadian Treasury Auctions," *Econometrica*, 80(6), 2511–2542.

Hortaçsu, A., J. Kastl, and A. Zhang (2016): " Bid Shading and Bidder Surplus in the U.S. Treasury Auction System ," *American Economic Review*, forthcoming.

Hortaçsu, A., and D. McAdams (2010): "Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market," *Journal of Political Economy*, 118(5), 833–865.

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

Kastl, J. (2017): "Recent advances in empirical analysis of financial markets: industrial organization meets finance," in *Advances in Economics and Econometrics: Volume 2: Eleventh World Congress*, vol. 59, p. 231. Cambridge University Press.

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

Krishnamurthy, A., and A. Vissing-Jorgensen (2012): "The aggregate demand for treasury debt," *Journal of Political Economy*, 120(2), 233–267.

Kyle, A. S. (1985): "Continuous auctions and insider trading," *Econometrica: Journal of the Econometric Society*, pp. 1315–1335.

——— (1989): "Informed Speculation with Imperfect Competition," *Review of Economic Studies*, 56(3), 317–355.

Lou, D., H. Yan, and J. Zhang (2013): "Anticipated and Repeated Shocks in Liquid Markets," *Review of Financial Studies*, 26(8), 1891–1912.

Lucas, R. E. (1972): "Expectations and the Neutrality of Money," *Journal of economic theory*, 4(2), 103–124.

Malvey, P. F., and C. M. Archibald (1998): "Uniform-price auctions: Update of the Treasury experience," *US Treasury.*

Malvey, P. F., C. M. Archibald, and S. T. Flynn (1995): "Uniform-Price Auctions: Evaluation of the Treasury Experience," *Office of Market Finance, US Treasury.*

Medrano, L., and X. Vives (2004): "Regulating Insider Trading when Investment Matters," *Review of Finance*, 8, 199–277.

Morris, S., and H. S. Shin (2002): "Social value of public information," *The American Economic Review*, 92(5), 1521–1534.

Nyborg, K. G., and S. M. Sundaresan (1996): "Discriminatory versus uniform Treasury auctions: Evidence from when-issued transactions," *Journal of Financial Economics*, 42, 63–104.

O'Hara, M. (1995): *Market microstructure theory*, vol. 108. Blackwell Cambridge, MA.

Ritter, J. R., and I. Welch (2002): "A Review of IPO Activity, Pricing, and Allocations," *The Journal of Finance*, 57(4), 1795–1828.

Röell, A. (1990): "Dual-capacity trading and the quality of the market," *Journal of Financial intermediation*, 1(2), 105–124.

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," Discussion paper, National Bureau of Economic Research.

Wall Street Journal (2015): "An Obscure Hedge Fund Is Buying Tens of Billions of Dollars of U.S. Treasurys," .

Warnock, F. E., and V. C. Warnock (2009): "International capital flows and US interest rates," *Journal of International Money and Finance*, 28(6), 903–919.

Wilson, R. (1979): "Auctions of shares," *Quarterly Journal of Economics*, 93(4), 675–689.

Woodford, M. (2011): *Interest and prices: Foundations of a theory of monetary policy.* princeton university press.

Internet Appendix: Not for Publication

A Solution and Proofs

A.1 Model solution details

Step 1. Describe information from conditioning on market-clearing prices: Define M_I, M_J, and M_D as follows.

$$M_I = \left(\rho\mathbb{V}[f|S_i] + \frac{dp}{dq_i}\right)^{-1} \qquad \texttt{for indirect investors} \tag{21}$$

$$M_J = \left(\rho\mathbb{V}[f|S_j] + \frac{dp}{dq_j}\right)^{-1} \qquad \texttt{for direct investors} \tag{22}$$

$$M_D = \left(\rho\mathbb{V}[f|S_d] + \frac{dp}{dq_d}\right)^{-1} \qquad \texttt{for dealers} \tag{23}$$

Now we define $\bar{v}_I$ and $\bar{v}_J$ as follows:

$$\frac{1}{N_I}\sum_{i=1}^{N_I} v_i \equiv \bar{v}_I \tag{24}$$

$$\frac{1}{N_J}\sum_{j=1}^{N_J} v_j \equiv \bar{v}_J \tag{25}$$

If $v_i \sim iid\, N(0, \tau_{vI}^{-1})$ and $v_j \sim iid\, N(0, \tau_{vJ}^{-1})$, then $var(\bar{v}_I) = 1/N_D\tau_{vI}^{-1} \equiv \tau_{\bar{v}I}^{-1}$ and $var(\bar{v}_J) = 1/N_J\tau_{vJ}^{-1} \equiv \tau_{\bar{v}J}^{-1}$. These $\tau_{\bar{v}}$ terms are the noise in prices.

We conjecture a linear price function p:

$$p = A + B_I\bar{s}_I + B_J\bar{s}_J + B_D\bar{s}_D + C_I\bar{v}_I + C_J\bar{v}_J + D\delta \tag{26}$$

where $\bar{s}_I = \frac{1}{N_I}\sum_i s_i$, $\bar{s}_J = \frac{1}{N_J}\sum_j s_j$, and $\bar{s}_D = \frac{1}{N_D}\sum_{d=1}^{N_D} s_d$.[22]

The price signal for dealers is:

$$s(p) = \frac{p - A - D\delta}{B_I + B_J + B_D} \tag{27}$$

Note that since non-price contingent bids are common knowledge, the dealer and all other bidders simpliy subtract this from this price signal. It is not a source of price noise.

Investors have private information about price noise $\bar{v}$. Investors can subtract the effect of their own private value from $\bar{v}$ and obtain a slightly more precise estimate of the noise, and thus extract a slightly more precise signal from prices. Indirect investor i faces precise noise from indirect investors of $\bar{v}_I - \mathbb{E}[\bar{v}_I|v_i]$. For a direct investor, the residual direct investor price noise is $\bar{v}_I - \mathbb{E}[\bar{v}_I|v_i]$. Note that dealers have no private value information. Their v_d values are known to all and get incorporated

[22] Why isn't dealer signal noise ξ_i in the price? Because each dealer is assumed to transmit unbiased signals with zero signal noise, on average, to his clients, it implies that dealer signal noise realizations ξ_d, have no effect on the market-clearing price. Note that each ξ is equally-weighted in demand by each client i. For the purpose of solving the price, we can drop the ξ_d in demand (but not it's precision τ_ξ).

in the price constant term A. Thus, the price signals for indirect and direct investors are:

$$s(p|v_i) = \frac{p - A - C_I v_i/N_I - D\delta}{B_I + B_J + B_D} \tag{28}$$

$$s(p|v_j) = \frac{p - A - C_J v_j/N_J - D\delta}{B_I + B_J + B_D} \tag{29}$$

Next steps are aim to construct conditional expectations of f for indirect, direct investors, and dealers: $\mathbb{E}[f|s_i, s_{\xi i}, s(p|v_i)]$, $\mathbb{E}[f|s_j, s(p|v_j)]$, and $\mathbb{E}[f|\tilde{s}_d, s(p)]$.

Step 2. Give all signals a state-space representation: The vector of orthogonal shocks Z is a column vector of size $N_Z = N + 2N_I + N_J + 1$, where

$$Z = [\epsilon_1, \dots, \epsilon_N, v_1, \dots, v_{NI+NJ}, \xi_1, \dots, \xi_{NI}, \delta]' \tag{30}$$

and the variance matrix of Z is

$$var(Z) = diag([\tau_\epsilon^{-1}\mathbf{1_N}, \tau_{vI}^{-1}\mathbf{1}_{N_I}, \tau_{vJ}^{-1}\mathbf{1}_{N_J}, \tau_\xi^{-1}\mathbf{1}_{N_I}, \tau_\delta^{-1}]) \tag{31}$$

where $\mathbf{1_N}$ is a vector of 1s of size N. Let ϕ_i be a vector of size N_Z of zeros with one 1 in ith position. For example, $\phi_3 = [0, 0, 1, 0, \dots, 0]$. Then $s_i = f + \phi_i Z$ and $\tilde{s}_d = f + \frac{N_D}{N_I+N_D}\sum_{i\in d(i)} \phi_i Z$. Here, $\frac{N_D}{N_I+N_D}$ means the number of clients per dealer.

Dealers' price signal. How do we represent price signals in state space? For a dealer, the price signal is

$$s(p) = \frac{p - A - D\delta}{\tilde{B}} = \frac{B_I\bar{s}_I + B_J\bar{s}_J + B_D\bar{s}_D}{\tilde{B}} + \frac{C_I}{\tilde{B}}\bar{v}_I + \frac{C_J}{\tilde{B}}\bar{v}_J. \tag{32}$$

To represent $\bar{v}$, we need a $1 \times N_Z$ vector ϕ_{vI} that is N zeros, followed by N_I ones, followed by $N_J + N_I$ zeros:

$$\phi_{vI} = [\mathbf{0_N}, \mathbf{1_{N_I}}, \mathbf{0_{N_J}}, \mathbf{0_{N_I+1}}] \tag{33}$$

$$\phi_{vJ} = [\mathbf{0_N}, \mathbf{0_{N_I}}, \mathbf{1_{N_J}}, \mathbf{0_{N_I+1}}] \tag{34}$$

$$\tag{35}$$

These vectors select out the private values of I or J investors from Z. Then $\bar{v}_I = (1/N_I)\phi_{vI} \cdot Z$ and $\bar{v}_J = (1/N_J)\phi_{vJ} \cdot Z$.

Next, we define binary vectors that select from Z the signal noise of all indirect investors, all direct investors, or all dealers:

$$\phi_{\epsilon I} = [\mathbf{1_{N_I}}, \mathbf{0_{N_Z - N_I}}] \tag{36}$$

$$\phi_{\epsilon J} = [\mathbf{0_{N_I}}, \mathbf{1_{N_J}}, \mathbf{0_{N_Z - N_I - N_J}}] \tag{37}$$

$$\phi_{\epsilon D} = [\mathbf{0_{N_I+N_J}}, \mathbf{1_{N_D}}, \mathbf{0_{N_Z - N}}]. \tag{38}$$

Finally, we define binary vectors that select from Z the signal noise ϵ or the private value v_i of all clients of a given dealer d:

$$\tilde{\phi}_{\epsilon d'}(i) = 1 \quad \text{if} \quad d' = d(i) \;\text{ and 0 otherwise.} \tag{39}$$

$$\tilde{\phi}_{vd'}(N+i) = 1 \quad \text{if} \quad d' = d(i) \;\text{ and 0 otherwise.} \tag{40}$$

Then dealers price signal is

$$s(p) = \frac{p - A - D\delta}{\tilde{B}} = f + \frac{B_I}{\tilde{B}N_I}\phi_{\epsilon I}\cdot Z + \frac{B_J}{\tilde{B}N_J}\phi_{\epsilon J}\cdot Z + \frac{B_D}{\tilde{B}N_D}\phi_{\epsilon D}\cdot Z + \frac{C_I}{N_I\tilde{B}}\phi_{vI}\cdot Z + \frac{C_J}{N_J\tilde{B}}\phi_{vJ}\cdot Z \equiv f + \pi_p Z \tag{41}$$

Dealers' information from indirect bids. The dealer observes the bid function (13) of each indirect investor i that bids through the dealer: $i : d(i) = d$. The dealer can multiply the bid by $\rho V[f|S_i] + dp/dq_i$, which all depend on known parameters, to infer $E[f|s_i] + v_i - p$, for each value p. The dealer can add p back in to infer $E[f|s_i] + v_i$. Recall from (5) that the conditional expectation is $(1 - \beta' 1_m)\mu + \beta' S_i$. The dealer knows the prior belief μ and can thus subtract that to infer $\beta' S_i + v_i$. Breaking out the signal vector in its component parts yields $\beta_{Is}s_i + \beta_{I\xi}s_{\xi i} + \beta_{Ip}(s(p) - C_I/(\tilde{B}N_I)v_i) + v_i$. Note that the indirect bidder's private value affects his bid in two ways, once directly affecting demand, and once by changing the way he interprets the price. The dealer can again take out the known terms $\beta_{Ip}s(p)$ and $\beta_{I\xi}s_{\xi i}$, which the dealer knows since he sent that signal to his client. Thus, the known component of beliefs that each dealer subtracts from average client valuations is $s_{public} = (1 - \beta' 1_m)\mu + \beta_{Ip}s(p) + \beta_{I\xi}s_{\xi i}$. That leaves $\beta_{Is}s_i + (1 - \beta_{Ip}C_I/(\tilde{B}N_I))v_i$. Dividing by β_{Is}, we obtain the unbiased signal that a dealer can infer from each of his N_I/N_D clients indirect bids. Since each of these signals is equally precise, the dealer optimally averages them to yield:

$$\tilde{s}_d = f + \frac{N_D}{N_I}\sum_{i\in d(i)} \epsilon_i + \beta_{Is}^{-1}\left(1 - \frac{\beta_{Ip}C_I}{\tilde{B}N_I}\right)\frac{N_D}{N_I}\sum_{i\in d(i)} v_i \tag{42}$$

So a dealer's 3 signals can be represented as

$$\begin{bmatrix} s_d \\ \tilde{s}_d \\ s(p) \end{bmatrix} = \begin{bmatrix} f \\ f \\ f \end{bmatrix} + \begin{bmatrix} \phi_i \\ \pi_d \\ \pi_p \end{bmatrix} \cdot Z \tag{43}$$

where

$$\pi_d = \frac{N_D}{N_I}\tilde{\phi}_{\epsilon d} + \beta_{Is}^{-1}\left(1 - \frac{\beta_{Ip}C_I}{\tilde{B}N_I}\right)\frac{N_D}{N_I}\tilde{\phi}_{vd}.$$

Indirect and direct bidders' price signals: Indirect and direct investors remove the effect of their own private valuations from the price when they condition on it. Their signals are the same as the dealers' price signal $s(p)$, minus all the terms that load on v_i or v_j for that investor:

$$s(p|v_i) = s(p) - \frac{C_I}{N_I\tilde{B}}\phi_{N+i} \cdot Z \tag{44}$$

$$s(p|v_j) = s(p) - \frac{C_J}{N_J\tilde{B}}\phi_{N+N_I+j} \cdot Z \tag{45}$$

Note that ϕ_{N+i} has a 1 in the position that corresponds to v_i in Z and ϕ_{N+N_J+j} has a 1 in the position that corresponds to v_j for direct investor j.

Indirect bidders' information from dealers: The dealer takes all the information collected from all his clients $\tilde{s}_d$, adds noise ξ_i to it for each bidder i and transmits the resulting $s_{\xi i}$ signal to his client. This signal has the same state space representation as the dealer's signal $\tilde{s}_d$, with one additional term $\phi_{N_Z - N_D + i}$ that adds the signal noise ξ_i:

$$s_{\xi i} = f + (\pi_d + \phi_{N_Z - N_D + i})Z. \tag{46}$$

Signals for indirect investors are:

$$\begin{bmatrix} s_i \\ s_{\xi i} \\ s(p|v_i) \end{bmatrix} = \begin{bmatrix} f \\ f \\ f \end{bmatrix} + \begin{bmatrix} \phi_i \\ \pi_d + \phi_{N_Z - N_D + i} \\ \pi_p - \frac{C_I}{N_I \tilde{B}} \phi_{N+i} \end{bmatrix} \cdot Z \tag{47}$$

Direct bidders' signals: The signal vector for direct investors is:

$$\begin{bmatrix} s_j \\ s(p|v_j) \end{bmatrix} = \begin{bmatrix} f \\ f \end{bmatrix} + \begin{bmatrix} \phi_i \\ \pi_p - \frac{C_J}{N_J \tilde{B}} \phi_{N+N_I+j} \end{bmatrix} \cdot Z \tag{48}$$

In sum the 3×1 signal loading matrix of a dealer's signals, an indirect bidder's signals and a direct bidder's signals can be written as:

$$\Pi_{D,d} \equiv \begin{bmatrix} \phi_i & \pi_d & \pi_p \end{bmatrix}' \tag{49}$$

$$\Pi_{I,i} \equiv \begin{bmatrix} \phi_i & \pi_{d(i)} + \phi_{N_Z - N_D + i} & \pi_p - \frac{C_I}{N_I \tilde{B}} \phi_{N+i} \end{bmatrix}' \tag{50}$$

$$\Pi_{J,j} \equiv \begin{bmatrix} \phi_i & \pi_p - \frac{C_J}{N_J \tilde{B}} \phi_{N+N_I+j} \end{bmatrix}' \tag{51}$$

Given this state space representation, we just apply (5),(6), and (7) to form conditional expectations and variance. $v(S_j) = \tau_f^{-1} \mathbf{1}\mathbf{1}' + \Pi_j v(z) \Pi_j'$ delivers $\beta, \mathbb{E}[f|S]$, and $\mathbb{V}[f|S]$. Let $\beta_I, \beta_J, \beta_D$ be the weights on signals given by (6) for each type of market participant. Then we can express the conditional expectation of the payoff as $\mathbb{E}[f|S_i] = (1 - \beta'\mathbf{1})\mu + \beta' S_i$. Let $\beta_I = [\beta_{Is}, \beta_{I\xi}, \beta_{Ip}]'$ and define β_J and β_D components analogously.

Step 3. Solve for the equilibrium price:

Using the first-order conditions (13) and (14) and the definition of M from (21), (22) and (23), we rewrite the market clearing condition as:

$$\sum_{i=1}^{N_I} \left(\mathbb{E}(f|S_i) + v_i - p\right) M_I + \sum_{j=1}^{N_J} \left(\mathbb{E}(f|S_j) + v_j - p\right) M_J + \sum_{d=1}^{N_D} \left(\mathbb{E}(f|S_d) + \chi - p\right) M_D + \delta = 1 \tag{52}$$

Substituting in the state space representation of conditional expectations:

$$\begin{aligned} 1 \quad = \quad & M_I[\sum_{i=2}^{N_I} (1 - \beta_I' \mathbf{1}')\mu + \beta_I' S_i - p] + N_I M_I \bar{v}_I \\ + \quad & M_J[\sum_{j=1}^{N_J} (1 - \beta_J' \mathbf{1}')\mu + \beta_J' S_j - p] + N_J M_J \bar{v}_J \\ + \quad & M_D[\sum_{d=1}^{N_D} (1 - \beta_D' \mathbf{1}')\mu + \beta_D' S_d - p] + N_D M_D \chi + \delta \end{aligned} \tag{53}$$

Define $\tilde{M} = N_I M_I + N_J M_J + N_D M_D$ and break out the signal vectors into their constituent parts.

$$\begin{aligned}
1 &= \tilde{A} + N_I M_I(\beta_{Is}\bar{s}_I + \beta_{I\xi}\bar{s}_\xi + \beta_{Ip}\bar{s}(p|v_i)) \\
&+ N_J M_J(\beta_{Js}\bar{s}_J + \beta_{Jp}\bar{s}(p|v_j)) + N_D M_D(\beta_{Ds}\bar{s}_D + \beta_{D\xi}\bar{s}_\xi + \beta_{Dp}s(p)) \\
&- \tilde{M}p + N_I M_I\bar{v}_I + N_J M_J\bar{v}_J + N_D M_D\chi + \delta
\end{aligned} \tag{54}$$

where $\tilde{A}$ is a collection of all μ terms $\tilde{A} = M_I N_I(1-\beta_I'\mathbf{1}')\mu + M_J N_J(1-\beta_J'\mathbf{1}')\mu + M_D N_D(1-\beta_D'\mathbf{1}')\mu$.

Define the average signal transmitted by a dealer to clients as: $\bar{s}_\xi = (1/N_I)\sum_{i=1}^{N_I} s_{\xi i}$. Note that this is the same as the average order flow information observed by dealers because we have assumed that dealer signal noise averages to zero. From the analysis of dealer's information from indirect bids, we showed that this order flow information is a weighted sum of indirect bidders' signals and their private values. Averaging this signal yields $\bar{s}_\xi = \bar{s}_I + \beta_{Is}^{-1}(1-\beta_{Ip}C_I/(\tilde{B}N_I))\bar{v}_I$.

Substitute in the conditional price signals to get:

$$\begin{aligned}
1 &= \tilde{A} + M_I N_I(\beta_{Is}\bar{s}_I + \beta_{I\xi}\bar{s}_\xi + \beta_{Ip}(s(p) - \frac{C_I}{N_I\tilde{B}}\bar{v}_I)) + M_J N_J(\beta_{Js}\bar{s}_J + \beta_{Jp}(s(p) - \frac{C_J}{N_J\tilde{B}}\bar{v}_J)) \\
&+ M_D N_D(\beta_{Ds}\bar{s}_D + \beta_{D\xi}\bar{s}_\xi + \beta_{Dp}s(p)) - \tilde{M}p + M_I N_I\bar{v}_I + M_J N_J\bar{v}_J + N_D M_D\chi + \delta
\end{aligned} \tag{55}$$

Using $s(p) = (p - A - D\delta)/\tilde{B}$, we can gather coefficients of p. Then let $\tilde{Q} \equiv \tilde{M} - (N_I M_I\beta_{Ip} + N_J M_J\beta_{Jp} + N_D M_D\beta_{Dp})\frac{1}{\tilde{B}}$.

Gathering terms in p and then matching A to all constants:

$$A = \tfrac{1}{\tilde{Q}}\left[\tilde{A} - 1 + A(\tilde{Q} - \tilde{M}) + N_D M_D\chi\right] \tag{56}$$

$$= \tfrac{1}{\tilde{M}}\left[\tilde{A} - 1 + N_D M_D\chi\right] \tag{57}$$

where the second line comes from collecting terms in A. Note that $\tilde{A}$ does not contain A terms.

Substituting in A, substituting in for $\bar{s}_\xi$, and rearranging the market clearing equation yields,

$$\begin{aligned}
\tilde{Q}p &= A\tilde{Q} + M_I N_I\bar{v}_I + M_J N_J\bar{v}_J + M_I N_I\beta_{Is}\bar{s}_I + M_J N_J\beta_{Js}\bar{s}_J \\
&+ (M_I N_I\beta_{I\xi} + M_D N_D\beta_{D\xi})(\bar{s}_I + \beta_{Is}^{-1}(1 - \beta_{Ip}\frac{C_I}{N_I\tilde{B}})\bar{v}_I) + M_D N_D\beta_{Ds}\bar{s}_D \\
&- M_I\beta_{Ip}\frac{C_I}{\tilde{B}}\bar{v}_I - M_J\beta_{Jp}\frac{C_J}{\tilde{B}}\bar{v}_J + (\tilde{Q} - \tilde{M})D\delta + \delta
\end{aligned} \tag{58}$$

Matching coefficients, gives us the solution for equilibrium prices, in terms of pricing impact M:

$$B_I = \tfrac{1}{\tilde{Q}}[M_I N_I\beta_{Is} + (M_I N_I\beta_{I\xi} + M_D N_D\beta_{D\xi})] \tag{59}$$

$$B_J = \tfrac{1}{\tilde{Q}}M_J N_J\beta_{Js} \tag{60}$$

$$B_D = \tfrac{1}{\tilde{Q}}M_D N_D\beta_{Ds} \tag{61}$$

$$C_I = \tfrac{1}{\tilde{Q}}M_I N_I\left(1 - \beta_{Ip}\tfrac{C_I}{N_I\tilde{B}}\right) + \tfrac{1}{\tilde{Q}}(M_I N_I\beta_{I\xi} + M_D N_D\beta_{D\xi})\beta_{Is}^{-1}\left(1 - \beta_{Ip}\tfrac{C_I}{N_I\tilde{B}}\right) \tag{62}$$

$$C_J = \tfrac{1}{\tilde{Q}}M_J N_J\left(1 - \beta_{Jp}\tfrac{C_J}{N_J\tilde{B}}\right). \tag{63}$$

$$D = \tfrac{1}{\tilde{M}} \tag{64}$$

Step 4. Determine price impact and demand: To solve for dp/dq_j, start with market-clearing in (55) but write one indirect investor's demand as an exogenous amount q_1:

$$\begin{aligned}
1 &= \tilde{A} + M_I(N_I-1)(\beta_{Is}\bar{s}_I + \beta_{I\xi}\bar{s}_{DI} + \beta_{Ip}(s(p) - \frac{C_I}{(N_I-1)\tilde{B}}\bar{v}_I)) \\
&+ M_J N_J(\beta_{Js}\bar{s}_J + \beta_{Jp}(s(p) - \frac{C_J}{N_J\tilde{B}}\bar{v}_J)) + M_D N_D(\beta_{Ds}\bar{s}_{DI} + \beta_{Dp}s(p)) \\
&- (\tilde{M} - M_I)p + M_I(N_I-1)\bar{v}_I + M_J N_J \bar{v}_J + \delta + q_1
\end{aligned} \tag{65}$$

Recall that $ds(p)/dp = 1/\tilde{B}$. Then use the implicit function theorem to solve for

$$\frac{dp}{dq_I} = \left[\tilde{M} - M_I - \frac{1}{\tilde{B}}\left(M_I(N_I-1)\beta_{Ip} + M_J N_J \beta_{Jp} + M_D N_D \beta_{Dp}\right)\right]^{-1} \tag{66}$$

Similarly, for direct bidders,

$$\frac{dp}{dq_J} = \left[\tilde{M} - M_J - \frac{1}{\tilde{B}}\left(M_I N_I \beta_{Ip} + M_J(N_J-1)\beta_{Jp} + M_D N_D \beta_{Dp}\right)\right]^{-1} \tag{67}$$

and for dealers,

$$\frac{dp}{dq_D} = \left[\tilde{M} - M_D - \frac{1}{\tilde{B}}\left(M_I N_I \beta_{Ip} + M_J N_J \beta_{Jp} + M_D(N_D-1)\beta_{Dp}\right)\right]^{-1} \tag{68}$$

Then the model solution is characterized jointly by the M's, the price coefficients and the updating formulas.

A.2 Result 1

This result has three cases. We consider each separately.

Case 1: Only Dealer-Client Information Sharing The set of equations (59)-(63), along with (66)-(68) substituted into M_D, M_I and M_J constitute a set of 8 equations in 8 unknowns. The fact that we could impose optimality, budget and market clearing conditions and then write price as a linear function of $\bar{s}_I$, $\bar{s}_J$, $\bar{s}_D$, $\bar{v}_I$ and $\bar{v}_J$, proves the linear price conjecture.

Case 2: Dealer-Dealer information sharing In this setup, dealers share information with clients using the same noisy signal as before, but they also share information with $\psi - 1$ other dealers. Then ψ is the size of the dealer chat room. Dealer-dealer sharing is symmetric, which requires that the number of dealers in an information-sharing collective be a factor of 20. We also require that $\psi \neq 20$, as that would imply perfect inter-dealer sharing. Thus, we only consider $\psi \in \{1, 2, 4, 5, 10\}$.

It would be repetitive to re-derive each part of the preceding analysis when most of it can be preserved. So, rather than do that, we simply point out the piece of the solution that is different.

First, define a set of dealers who share information with any given dealer d. Let $chat(d) = \{d', d'', \ldots\}$ such that dealers d', d'' and d share information. Information sharing means observing the order flow of the clients of all the dealers in $chat(d)$. Since all investor order flow signals are equally informative, all dealers in a chat group average all the order flow signals they see in the same way. Each dealer now sees a new, more precise composite order flow signal which is:

$$\tilde{s}_d = f + \frac{N_D}{\psi N_I} \sum_{d' \in chat(d(i))} \sum_{i \in d'} s_i + \beta_{Is}^{-1} \left(1 - \frac{\beta_{Ip} C_I}{\tilde{B} N_I}\right) \frac{N_D}{\psi N_I} \sum_{d' \in chat(d(i))} \sum_{i \in d'} v_i \tag{69}$$

Thus the dealer's signals have the same state space representation as in (43), except that we redefine π_d, the weight the order flow information puts on the underlying shocks, as

$$\pi_d = \frac{N_D}{\psi N_I} \sum_{d' \in chat(d(i))} \tilde{\phi}_{\epsilon d'} + \beta_{Is}^{-1} \left(1 - \frac{\beta_{Ip} C_I}{\tilde{B} N_I}\right) \frac{N_D}{\psi N_I} \sum_{d' \in chat(d(i))} \tilde{\phi}_{vd'}.$$

For indirect bidders, the signals are the same, except that π_d is redefined, as above. For direct bidders, there is no change in the signal vector. Of course, information charing will change bids and thus change the variance and covariance of auction-clearing price. But all these effects will show up through the change in the signal vector, represented by the change in π_d.

This change in the model does not change the linearity of the price in $\bar{s}_I$, $\bar{s}_J$, $\bar{s}_D$, $\bar{v}_I$ and $\bar{v}_J$. The equations (59)-(63), along with (66)-(68) substituted into M_D, M_I and M_J still characterize the solution to the model.

Case 3: No Information Sharing ("Chinese Wall") In this model, dealers do not use or share any information derived from client order flow. Practically speaking, it is as if each type of investor submits bids on their own behalf, rather than through an intermediary. Each investor's information set is therefore a 2×1 vector $S_i = [s_i, s_i(p)]$, comprised of their private signal s_i and the counterfactual price signal $s_i(p)$. In this regime, the only difference between dealers and non-dealer large investors is that dealers are subject to a minimum bidding penalty, while investors have private values that are private information.

We can solve this model by simply changing the signal vector weights on each of the underlying shocks. In sum the 2×1 signal loading matrix of a dealer's signals, an indirect bidder's signals and a direct bidder's signals now becomes:

$$\Pi_{D,d} \equiv \begin{bmatrix} \phi_i & \pi_p \end{bmatrix}' \tag{70}$$

$$\Pi_{I,i} \equiv \begin{bmatrix} \phi_i & \pi_p - \frac{CM_I}{\tilde{B}} \phi_{N+i} \end{bmatrix}' \tag{71}$$

$$\Pi_{J,j} \equiv \begin{bmatrix} \phi_i & \pi_p - \frac{CM_J}{\tilde{B}} \phi_{N+j} \end{bmatrix}' \tag{72}$$

Once we adjust the signal structure, the rest of the solution method goes through unchanged. In the price coefficient solutions (59)-(63), this change implies that the weight put on dealers' order flow signals (now no longer existent) $\beta_{I\xi}$ and $\beta_{D\xi}$ are both 0. Pricing simplifies to:

$$B_I = \frac{1}{\tilde{Q}} M_I N_I \beta_{Is} \tag{73}$$

$$B_J = \frac{1}{\tilde{Q}} M_J N_J \beta_{Js} \tag{74}$$

$$B_D = \frac{1}{\tilde{Q}} M_D N_D \beta_{Ds} \tag{75}$$

$$C_I = \frac{1}{\tilde{Q}} M_I N_I \left(1 - \beta_{Ip} \frac{C_I}{N_I \tilde{B}}\right) \tag{76}$$

$$C_J = \frac{1}{\tilde{Q}} M_J N_J \left(1 - \beta_{Jp} \frac{C_J}{N_J \tilde{B}}\right). \tag{77}$$

The existence of a set of coefficients verifies the price conjecture. Since the supply of the asset is one, auction revenue is the price of the asset. The solution to this model is a joint solution to

(66)-(68) substituted into M_D, M_I and M_J and the price coefficient equations above.

A.3 Equilibrium Existence

We prove equilibrium existence for two versions of the model, each of which represents one extreme of the information sharing spectrum. The first model has no information sharing (Chinese walls) and the second has complete sharing between dealers and their clients. The proof strategy is to write the solution as one equation in the ratio of price coefficients (B/C) and then use the intermediate value theorem to show a solution to that equation exists.

We employ the following simplifying assumptions throughout.

1. We assume that $N_I = 2 * N_D = N/3$ and $N_J = 0$.
2. The minimum bidding requirement cost χ is heterogeneous, private information, and has the same distribution as v_i. Thus, we call it v_d.

The role of these assumptions is to preserve as much symmetry as possible. Without symmetry, the number of terms in the pricing equation multiplies and it becomes impossible to characterize the model's solution with one equation. These simplifications imply that all investors have the same information set. Thus, they all put the same weight on each signal when they form their beliefs. Also, symmetry between investors implies that $M_J = M_D = M$. Therefore, the price conjecture and price signal can be written as:

$$p = A + B\bar{s} + C\bar{v} \tag{78}$$

$$s_i(p) = \frac{p - A - Cv_i}{B} = f + \bar{\epsilon} + \frac{N-1}{N}\frac{C}{B}\bar{v}$$

Case 1: Chinese Wall

Under no information sharing, we have that $S_i = S_d = [s_i, s_i(p)]$ and the market clearing condition is:

$$\sum_{i=1}^{N/3}(E[f|S_i] + v_i - p)M + \sum_{d=1}^{2N/3}(E[f|S_d] + v_d - p)M = 1 \tag{79}$$

Replacing $E[f|S]$ and dividing by MN we have:

$$\mu(1 - \beta_s - \beta_p) + \beta_s\bar{s} + \beta_p\bar{s}(p) + \bar{v} - p = \frac{1}{MN}$$

where $\bar{s}(p) = \frac{p-A}{B} - \frac{C}{BN}\bar{v}$. Then:

$$p(1 - \beta_p B^{-1}) = \mu(1 - \beta_s - \beta_p) - A\beta_p B^{-1} - (MN)^{-1} + \beta_s\bar{s} + \left(1 - \beta_p\frac{C}{BN}\right)\bar{v}$$

Matching coefficients we get:

$$A = \frac{\mu(1 - \beta_s - \beta_p) - A\beta_p B^{-1} - (MN)^{-1}}{1 - \beta_p B^{-1}} \tag{80}$$

$$B = \frac{\beta_s}{1 - \beta_p B^{-1}} \tag{81}$$

$$C = \frac{1 - \beta_p B^{-1} C/N}{1 - \beta_p B^{-1}} \tag{82}$$

The signals' weight in the optimal linear predictor are:

$$\beta = V[S]^{-1} 1_m \tau_f^{-1}.$$

In this case, $m = 2$ and

$$V[S] = \begin{bmatrix} \tau_f^{-1} + \tau_s^{-1} & \tau_f^{-1} + \tau_s^{-1}/N \\ \tau_f^{-1} + \tau_s^{-1}/N & \tau_f^{-1} + \tau_p^{-1} \end{bmatrix}$$

. Then:

$$\begin{aligned} V[S]^{-1} &= \frac{1}{(\tau_f^{-1} + \tau_s^{-1})(\tau_f^{-1} + \tau_p^{-1}) - (\tau_f^{-1} + \tau_s^{-1}/N)^2} \begin{bmatrix} \tau_f^{-1} + \tau_p^{-1} & -(\tau_f^{-1} + \tau_s^{-1}/N) \\ -(\tau_f^{-1} + \tau_s^{-1}/N) & \tau_f^{-1} + \tau_s^{-1} \end{bmatrix} \\ &= \frac{1}{\tau_f^{-1}[\tau_p^{-1} - \tau_s^{-1}/N] + \tau_f^{-1}[\tau_s^{-1} - \tau_s^{-1}/N] + \tau_s^{-1}[\tau_p^{-1} - \tau_s^{-1}/N^2]} \begin{bmatrix} \tau_f^{-1} + \tau_p^{-1} & -(\tau_f^{-1} + \tau_s^{-1}/N) \\ -(\tau_f^{-1} + \tau_s^{-1}/N) & \tau_f^{-1} + \tau_s^{-1} \end{bmatrix} \end{aligned}$$

Multiplying by $1_m \tau_f^{-1}$, we obtain:

$$\beta_s = \frac{\tau_f^{-1}[\tau_p^{-1} - \tau_s^{-1}/N]}{\tau_f^{-1}[\tau_p^{-1} - \tau_s^{-1}/N] + \tau_f^{-1}\tau_s^{-1}(N-1)/N + \tau_s^{-1}[\tau_p^{-1} - \tau_s^{-1}/N^2]}$$

We multiply the numerator and denominator by $\tau_s N/(N-1)\tau_f[\tau_p^{-1} - \tau_s^{-1}/N]^{-1}$ to get:

$$\beta_s = \frac{\tau_s N/(N-1)}{\tau_s N/(N-1) + [\tau_p^{-1} - \tau_s^{-1}/N]^{-1} + \tau_f^{-1} N/(N-1)[\tau_p^{-1} - \tau_s^{-1}/N^2][\tau_p^{-1} - \tau_s^{-1}/N]^{-1}}$$

Finally, we replace $\tau_p^{-1} = \frac{\tau_s^{-1}}{N} + \left(\frac{C}{B}\right)^2 \frac{N-1}{N} \frac{\tau_v^{-1}}{N}$:

$$\beta_s = \frac{\tau_s}{\tau_s + (B/C)^2 N\tau_v + \tau_f[1 + (B/C)^2 \tau_v/\tau_s]} \tag{83}$$

Following exactly the same procedure for the second row, we arrive to an analogous expression for β_p:

$$\beta_p = \frac{(B/C)^2 N\tau_v}{\tau_s + (B/C)^2 N\tau_v + \tau_f[1 + (B/C)^2 \tau_v/\tau_s]} \tag{84}$$

Now that we have the vector β, we can solve for the posterior variance:

$$\begin{aligned} \hat{\tau}^{-1} &= \tau_f^{-1} - 1_m' \tau_f^{-1} \beta && (85) \\ &= \frac{1 + (B/C)^2 \tau_v/\tau_s}{\tau_s + (B/C)^2 N\tau_v + \tau_f[1 + (B/C)^2 \tau_v/\tau_s]} && (86) \end{aligned}$$

In order to complete the characterization of the equilibrium, we need to compute the price impact to obtain M. From the market clearing condition we have:

$$\sum_{i=1}^{N/3-1} (E[f|S_i] + v_i - p)M + \sum_{d=1}^{2N/3} (E[f|S_d] + v_d - p)M + q = 1$$

Replacing $E[f|S]$ and dividing by $M(N-1)$ we have:

$$p(1-\beta_p B^{-1}) = \mu(1-\beta_s-\beta_p) - A\beta_p B^{-1} - (M(N-1))^{-1} + \beta_s \bar{s} + \left(1-\beta_p \frac{C}{BN}\right)\bar{v} + q(M(N-1))^{-1}$$

$$dp(1-\beta_p B^{-1}) = dq(M(N-1))^{-1}$$

and we obtain

$$\frac{dp}{dq} = (M(N-1))^{-1}(1-\beta_p B^{-1})^{-1}$$

Using the definition of M, we have that:

$$M^{-1} = \rho\hat{\tau} + M^{-1}(N-1)^{-1}(1-\beta_p B^{-1})^{-1}$$

Thus, the full characterization of the equilibrium is:

$$A = \mu(1-\beta_s-\beta_p) - (MN)^{-1} \tag{87}$$

$$B = \beta_s + \beta_p \tag{88}$$

$$C = \frac{1}{1-\frac{N-1}{N}\beta_p B^{-1}} \tag{89}$$

$$\beta_s = \frac{\tau_s}{\tau_s + (B/C)^2 N\tau_v + \tau_f[1+(B/C)^2\tau_v/\tau_s]} \tag{90}$$

$$\beta_p = \frac{(B/C)^2 N\tau_v}{\tau_s + (B/C)^2 N\tau_v + \tau_f[1+(B/C)^2\tau_v/\tau_s]} \tag{91}$$

$$\hat{\tau}^{-1} = \frac{1+(B/C)^2\tau_v/\tau_s}{\tau_s + (B/C)^2 N\tau_v + \tau_f[1+(B/C)^2\tau_v/\tau_s]} \tag{92}$$

$$M = \frac{1-(N-1)^{-1}(1-\beta_p B^{-1})^{-1}}{\rho\hat{\tau}} \tag{93}$$

Note that β_s and β_p only depend on C/B, so C and B do not depend on A or M. Also, since β_s and β_p are positive, it is easy to see that B and C are positive as well.

Starting from this observations, we are going to prove that the close system given by the equations for B and C plus the definitions of β_s and β_p has a solution. First, we divide the expressions for B and C and then plug in the definitions of β_s and β_p :

$$\begin{aligned}\frac{B}{C} &= (\beta_s+\beta_p)\left[1-\frac{N-1}{N}\frac{\beta_p}{\beta_s+\beta_p}\right] \\ &= \beta_s + \frac{\beta_p}{N} \\ &= \underbrace{\frac{\tau_s+(B/C)^2\tau_v}{\tau_s + (B/C)^2 N\tau_v + \tau_f[1+(B/C)^2\tau_v/\tau_s]}}_{\text{RHS}}\end{aligned}$$

The main idea of the proof is that while the left hand side is increasing in (B/C), the right hand side (RHS) is decreasing (since B and C are positive, we can use the derivative with respect to $(B/C)^2$). To see this, let $x = (B/C)^2$, then:

$$\text{RHS} = \frac{\tau_s + x\tau_v}{\tau_s + xN\tau_v + \tau_f[1 + x\tau_v/\tau_s]}$$

$$\begin{aligned}\frac{\partial \text{RHS}}{\partial x} =& \frac{\tau_v[\tau_s + xN\tau_v + \tau_f + x\tau_f\tau_v/\tau_s] - [x\tau_v + \tau_s][N\tau_v + \tau_f\tau_v/\tau_s]}{(\tau_s + xN\tau_v + \tau_f[1 + x\tau_v/\tau_s])^2} \\ =& \frac{\tau_v\tau_s - N\tau_v\tau_s}{(\tau_s + xN\tau_v + \tau_f[1 + x\tau_v/\tau_s])^2} < 0\end{aligned}$$

Finally, we use the intermediate value theorem to show existence. To do this, we need to show that there exist a value for (B/C) such that $(B/C) <$ RHS and another value such that $(B/C) >$ RHS. For $(B/C) = 0$, we have that $RHS = \tau_s/(\tau_s + \tau_f) > 0$. For $(B/C) = 1$, we have that RHS $= (\tau_s + \tau_v)/(\tau_s + N\tau_v + \tau_f[1 + \tau_v/\tau_s]) < 1$. Then, since the function is continuous, there exists an equilibrium value for (B/C) between 0 and 1.

Given the equilibrium value for (B/C), equations (90) ,(91) give us the values for β_s and β_p and then we can plug in and obtain the equilibrium values for the rest of the variables.

Case 2: Perfect information sharing

Under perfect information sharing a dealer observe the signal of each client and share that information with the other clients. With $N_I = 2N_D$, every agent observe 3 signals and every signal is observed by 3 agents. Let $\bar{s}_3 = \sum_{i \in d(i)} s_i$, then we have that all agents have the same information set $S_i = \{\bar{s}_3, s(p)\}$.

Since agents give the same weight to each of the three private signals they observe, the market clearing condition is the same as before. Thus, the equilibrium conditions we found for the chinese wall still holds, but for different β's. In particular, the unconditional variance now is:

$$V[S] = \begin{bmatrix} \tau_f^{-1} + \tau_s^{-1}/3 & \tau_f^{-1} + \tau_s^{-1}/N \\ \tau_f^{-1} + \tau_s^{-1}/N & \tau_f^{-1} + \tau_p^{-1} \end{bmatrix}$$

. Then, using the same steps than in the chinese wall, we get:

$$\beta_s = \frac{\tau_s}{\tau_s + \frac{N-1}{N-3}(B/C)^2 N\tau_v + \tau_f[1 + (B/C)^2\tau_v/\tau_s]} \tag{94}$$

$$\beta_p = \frac{\frac{N-1}{N-3}(B/C)^2 N\tau_v}{\tau_s + \frac{N-1}{N-3}(B/C)^2 N\tau_v + \tau_f[1 + (B/C)^2\tau_v/\tau_s]} \tag{95}$$

This modifications do not affect any part of the proof. Now, the equation for (B/C) is:

$$\frac{B}{C} = \underbrace{\frac{\tau_s + \frac{N-1}{N-3}(B/C)^2\tau_v}{\tau_s + \frac{N-1}{N-3}(B/C)^2 N\tau_v + \tau_f[1 + (B/C)^2\tau_v/\tau_s]}}_{\text{RHS}}$$

$$\text{where} \quad \frac{\partial \text{RHS}}{\partial x} = \frac{\frac{N-1}{N-3}\tau_v\tau_s(1 - N) + \tau_v\tau_f(\frac{N-1}{N-3} - 1)}{(\tau_s + x\frac{N-1}{N-3}N\tau_v + \tau_f[1 + x\tau_v/\tau_s])^2} < 0$$

For $(B/C) = 0$, we have that $RHS = \tau_s/(\tau_s + \tau_f) > 0$. For $(B/C) = 1$, we have that RHS $=$ $(\tau_s + \frac{N-1}{N-3}\tau_v)/(\tau_s + \frac{N-1}{N-3}N\tau_v + \tau_f[1 + \tau_v/\tau_s]) < 1$. Then, there exists an equilibrium with perfect information sharing.

A.4 Bid Shading and Signal Jamming

The term M for each investor type measures the bid sensitivity to changes in expected returns. Since expected returns are typically positive (this is a compensation for the risk of the uncertain common value), a larger value of M denotes a smaller average bid and a lower average equilibrium price.

From (21), (22) and (23), we know that the sensitivity M for ech of the three types of bidders is the inverse of a sum of $\rho V[f|S]$ term that measures risk aversion and risk, plus a dp/dq term that arises because strategic bidders internalize the impact they have on price. Bid shading and signal jamming are about this strategic dp/dq term.

The inverse of this price impact term is the sum of a direct effect, which is bid shading, and an indirect effect, which works through its effect on the beliefs of others:

$$\left(\frac{dp}{dq_I}\right)^{-1} = \tilde{M} - M_I - \left(M_I(N_I - 1)\beta_{p,I} + M_J N_J \beta_{p,J} + M_D N_D \beta_{p,3}\right)/\tilde{B}$$
$$\left(\frac{dp}{dq_J}\right)^{-1} = \tilde{M} - M_J - \left(M_I N_I \beta_{p,I} + M_J \left(N_J - 1\right) \beta_{p,J} + M_D N_D \beta_{p,3}\right)/\tilde{B}$$
$$\left(\frac{dp}{dq_D}\right)^{-1} = \tilde{M} - M_D - \left(M_I N_I \beta_{p,I} + M_J N_J \beta_{p,J} + M_D \left(N_D - 1\right) \beta_{p,3}\right)/\tilde{B}$$

The first two terms of each expression capture the direct effect of one bidder's demand on the price. $\tilde{M}$ is the sum of every bidder's demand sensitivity to return. The more sensitive demand is to return, the less the return needs to change to clear the market. This sum represents the inverse price elasticity to a change in demand. The higher it is, the larger the price impact of a change in bids. But if one bidder changes their demand, they do not absorb their own change in demand. Thus, bid shading is the price impact when all bidders, except one bidder of type I (or J or D), adjust their bids to the new equilibrium price:

$$S_I = \tilde{M} - M_I; \qquad S_J = \tilde{M} - M_J; \qquad S_D = \tilde{M} - M_D \tag{96}$$

The large term on the right describes how the price change affects others' demands through their beliefs. $\beta_{p,I}$ is the sensitivity of bidder type I's beliefs to a one-unit change in the price p. The sum in the parentheses is the sum of all the effects on beliefs of every bidder, except the one changing their demand (they do not fool themself). The final term $\tilde{B}$ maps these changes in beliefs to a change in price. Thus, signal jamming is defined as

$$SJ_I = \frac{dp}{dq_I} - S_I; \qquad SJ_J = \frac{dp}{dq_J} - S_J; \qquad SJ_D = \frac{dp}{dq_D} - S_D \tag{97}$$

HKZ measure of bid shading Hortacsu, Kastl and Zhang (2017) define bid shading as the quantity-weighted expected difference between the bidder's marginal valuation for the last unit awarded and the price paid. In our notation:

$$B(v_i, S_i) = \frac{\mathbb{E}[q_i \left(\frac{\partial EU}{\partial q_i} - p\right)]}{\mathbb{E}[q_i]} \tag{98}$$

Our model has two key differences. First, in HKZ, the uncertainty is about the realized price. Valuations are known. In our setting, bidders' uncertainty is about the payoff. So, we use marginal *expected* utility, in place of HKZ's marginal utility.

The second key difference is that our utility is not quasi-linear in bid payments. Instead, there is a risk-averse utility function over the value of the asset (itself a financial value) net of the payment. THere are two possible ways to deal with this

1. Instead of taking marginal utility of the payment, then subtracting the price ($MU(f) - p$), a natural adaption would be to compute the expected marginal value of the asset, net of payment($MEU(f - p)$). In other words, we bring the price paid inside the utility function because that's internally consistent with our model. Log expected utility (from text just before eqn (14)) is $q_j(\mathbb{E}[f|S_j] + v_j - p) - \frac{1}{2}\rho q_j^2 \mathbb{V}[f|S_j]$. The associated marginal utility is:

$$\frac{\partial EU}{\partial q_i} = \mathbb{E}[f|S_j] + v_j - p - q_j \frac{dp}{dq_j} - \rho q_j \mathbb{V}[f|S_j]. \tag{99}$$

   Marginal utility (without the log) is just a rescaling: $\frac{\partial EU}{\partial q_i} \times E[U]$. The problem is that this quantity is always zero. Why? Because it's our first order condition. Bid for more q_i until the marginal additional unit yields zero marginal utility.

2. However, we could instead be more true to the HKZ definition by keeping the price out of the marginal utility. We compute marginal utility, as if price were zero, and then subtract the price. We get

$$\left.\frac{\partial EU}{\partial q_i}\right|_{p=0} - p = q_i \frac{\partial p}{\partial q_i} \tag{100}$$

   In our model, the price impact term $\frac{\partial p}{\partial q_i}$ is a function of parameters, not of random variables. So, we can pull it out of the expectation. When we substitute this into 98, the $E[q_i]$ terms cancel and we get

$$B(v_i, S_i) = \frac{\partial p}{\partial q_i}. \tag{101}$$

   This measure is plotted below in Figure 6. It shows that information sharing, of either kind, reduces bid shading. The fact that only one line is visible indicates that, in this model, the price impact of a dealer trade, a client trade or a direct bidder trade are indistinguishable.

Since the main question of the paper is about how information sharing affects auction revenue, our primary measure of bid sharing is how much revenue is lost to bid shading (and signal jamming), and how that interacts with information sharing. To compute this revenue loss, we simply turn off some or all of the dp/dq term in the first-order condition. This term is the only piece of demand that differs from the demand of a fully competitive, measure-zero bidder. When we set $dp/dq = 0$ and re-solve the model (agents are aware that others are not strategic and correctly infer different information from the auction-clearing price), we capture all lost revenue due to strategic bidding. Then, we break up that lost revenue into two pieces: bid shading and signal jamming as follows.

Signal jamming Signal jamming is the revenue lost because bidders try to influence each others' beliefs. We compute optimal signal jamming now, both analytically and quantitatively. We compare its magnitude to price impact more generally and to the magnitude of bid shading. Equation (67)

Figure 6: A Price Impact Measure of Bid Shading. This plots bid shading, as defined in (101).

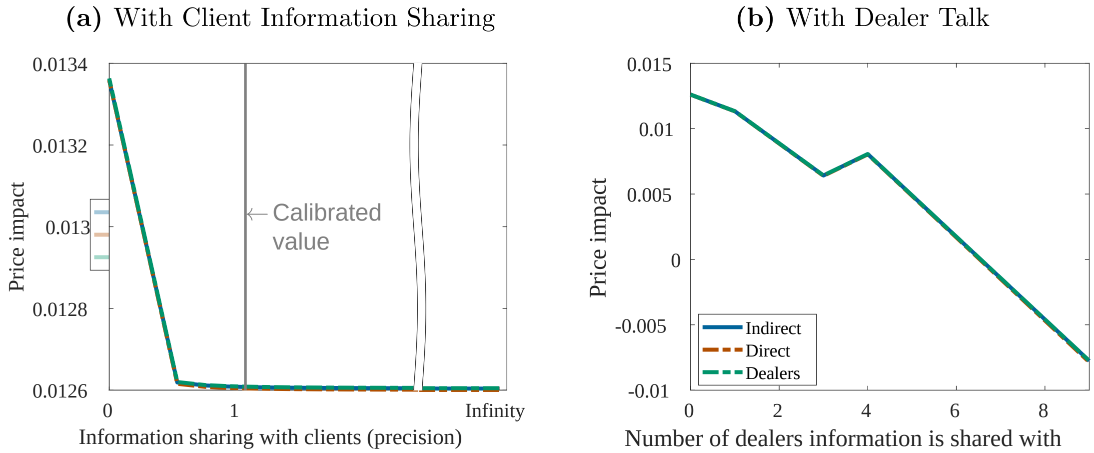

shows that the price impact of an indirect investor (for example, other classes of agents have analogous expressions) is:

$$\frac{dp}{dq_I} = \left[\tilde{M} - M_I - \frac{1}{\tilde{B}}\left(M_I(N_I-1)\beta_{Ip} + M_J N_J \beta_{Jp} + M_D N_D \beta_{Dp}\right)\right]^{-1} \tag{102}$$

where $\tilde{M} - M_I$ represents the price elasticity of all other market participants collectively. This is the direct effect of one unit of additional demand from one I investor on the market price. The long last term is signal jamming:

$$\text{Signal Jamming} = \frac{1}{\tilde{B}}\left(M_I(N_I-1)\beta_{Ip} + M_J N_J \beta_{Jp} + M_D N_D \beta_{Dp}\right) \tag{103}$$

The β_{Ip}, β_{Jp}, and β_{Dp} terms measure how much a change in the price affects other indirect, direct investor' and dealers' beliefs. Investors in our model consider how their bids affect the information transmitted by the price and they are optimally adjusting their bid to distort that price signal.

A.5 Auction price with dealer collusion

When dealers collude, they share information and then bid in order to maximize their joint utility. From an information point of view, if collusion takes place in pairs (each dealer shares information and bids jointly with 1 other dealer), it is as if there are $N_D/2$ dealers, each with twice as many orders as before. If collusion takes place in groups of size ψ, the information structure is as if there are N_D/ψ dealers. The only difference between the collusion model and the reduced-number of dealers model is that the demand of each collusive group is larger than it would be if there were only 1 dealer. Two colluding bidders bid have a larger appetite for risk. One can think of collusion as a contractual arrangement whereby each dealer commits to give half his profits to the other dealer and thereby internalizes his effect on the other dealer.

The portfolio optimization problem of colluding dealers is

$$\max_{q_d, q_{d'}, p} \mathbb{E}\left[-\exp\left(-\rho\frac{1}{2}((W_d + q_d v_d) + (W_{d'} + q_{d'} v_{d'}))\right) | S_d\right] \tag{104}$$

$$\text{s.t. } W_d = W_{0,d} + q_d(f - p), \text{ and } W_{d'} = W_{0,d'} + q_{d'}(f - p) \tag{105}$$

$$\sum_{i=1}^{N_I} q_i + \sum_{j=1}^{N_J} q_j + \sum_{d=1}^{N_D} q_d = 1. \tag{106}$$

Taking the expected value of the lognormal yields

$$-\exp\left(\text{const} - \rho\frac{1}{2}((q_d + q_{d'})(\mathbb{E}[f|S_j] - p + \chi)) + \frac{\rho^2}{8}(q_d^2 + q_{d'}^2 + 2q_d q_{d'})\mathbb{V}[f|S_j]\right)$$

where const is the constant that depends on initial wealth. Then computing the first order condition with respect to q_d reveals that

$$q_d(p) + q_{d'}(p) = 2 \cdot \frac{\mathbb{E}[f|S_j] + \chi - p}{\rho\mathbb{V}[f|S_j] + 2dp/dq_j}. \tag{107}$$

So, the two colluding dealers jointly bid for twice as much of the asset, but adjusted for twice the price impact. This is the same formula as that which would hold for one dealer who has 1/2 the risk aversion. Therefore, we numerically solve the collusion model by reducing the number of dealers from N_D to N_D/ψ and reducing each dealer's risk aversion from ρ to ρ/ψ for $\psi \geq 1$.

Figure 7: Lying about Dealer Talk Reduces Revenue. Figure plots average equilibrium auction revenue, against the number of other dealers that share information. We assume here that when dealers share information, no one else knows.

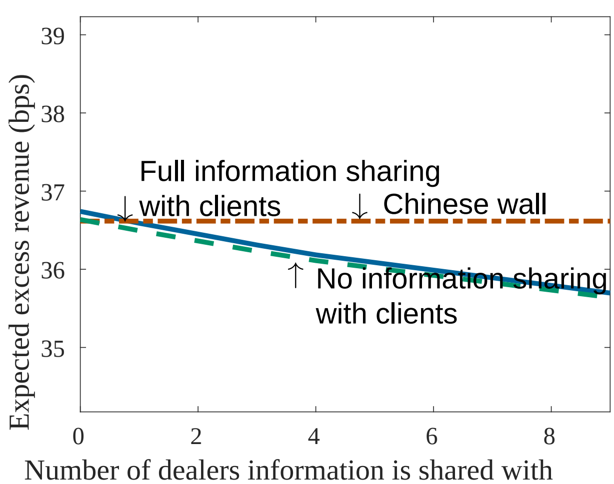

Lying about Dealer Talk A related issue is that in practice, not all market participants may know that dealers are swapping order flow information. Of course, this also has a separate legal remedy. One can enforce laws about disclosure of information practices, without prohibiting the information sharing. But our results on what happens when others are not aware highlights the importance of the assumption that agents understand others' strategies.

When a set of dealers share information and others are not aware, auction revenue falls. This is true even if the information is shared with clients. If the clients are not aware that their information is very precise, they do not bid as if they are better informed. By not bidding aggressively, these clients fail to push up auction revenue as they do in the baseline case. Just as with collusion, when revenue declines, bidder utilities rise. All bidders are better off because prices are lower. But taxpayers are left to foot the bill.

To compute the revenue in Figure 7, we simulated a version of our model where a set of ψ dealers share information and bid collusively on that more precise information. We vary the size of the set of dealers. But every other bidder and dealer bids using the no-dealer-sharing bid functions. The idea is that if they are unaware of the information sharing, then their strategy should be unchanged by it. For each ψ, we resolved for the equilibrium pricing coefficients and then computed the average auction revenue.

A.6 Intermediation Choice: Solution with one bidder who switches

Our objective is to illustrate the properties of the intermediation decision of a client. To do that, we simplify the model by assuming that all participants in the auction (including dealers) have private values drawn from $v_i \sim N\left(0, \tau_v^{-1}\right)$ and focusing on the model without demand shocks. We focus on the case where dealers share information perferctly with their clients and study, instead, how the intermediation decision changes when dealers share information with each other. Without loss of generality, we assume that client 1 of dealer 1 is the agent making the intermediation decision.

If one bidder switches from being an indirect bidder through a dealer to a direct bidder through treasury direct, how does the signal structure for bidders and dealers change? If the dealer did not make any inference from the direct bidding choice of the client, the the solution would be the same as before, only adjusting the number of indirect and direct bidders. But a rational dealer who observes a regular client not showing up infers that the client's signal must be in a particular range. We propose a solution method that includes that inferred information.

Define a conglomerate to be the set of dealers that share information with each other, as well as all their clients. Without loss, let conglomerate 1 be the conglomerate that the marginal bidder would be bid through, if he decided to bid through a dealer. This is the group of agents that learn from seeing bidder 1 bid directly or indirectly. The intermediation decision of the client depends on both the client's signal and private value. The key to our solution method is that we approximate the truncated normal signal that can be extracted from the intermediation choice with a normal signal $s_q = f + m_q + e_q$, with the same mean m_q and variance τ_q^{-1} as the true signal. Denote also by p_q the probability of the client choosing to bid directly.

If bidder 1 chooses to bid through the dealer, the dealer sees the intermediation decision, which reveals that bidder 1's order flow must be in a range. But the intermediating dealer also sees exactly what bidder 1's order flow is. The additional information from seeing the choice to bid indirectly is redundant. Thus, in this case, we do not need to construct an approximated dealer signal from the intermediation decision. Just seeing the order flow contains all the relevant information.

In cases where the bidder bids indirectly, we solve the model using an approximating normal signal. The normal signal is included in the precision-weighted average signal of dealer d'.

$$\tilde{s}_{d'} = \iota \left(\frac{1}{N_I/N_D - 1} \left(\sum_{i \in \mathcal{I}_d} \mathbb{E}_i[f] + v_i \right) - s_{public} \right) + (1-\iota) s_q, \tag{108}$$

where $\mathcal{I}_d$ is the reduced set of investors bidding through dealer d – excluding the direct bidder – and s_{public} is solved for in Appendix B. The dealer is constructing $\tilde{s}_d$ from an average of his clients' expected valuations plus private values, minus a term s_{public} that includes all public information in $E_i[f]$, and from the information s_q inferred from the direct bidding decision. If investor j bids through the dealer, the problem and the solution are the same as in the baseline model.

The equilibrium price if bidder 1 chooses to bid indirectly (through the dealer) can be expressed as

$$\begin{aligned} p = A + B_I \frac{\nu_I - 1}{N_I + N_D} \bar{s}_{I1} + B_I \frac{N_D + N_I - \nu_I}{N_D + N_I} \bar{s}_{I2} + \frac{B_I}{N_D + N_I} s_1 + \frac{B_J}{N_J} \bar{s}_J \\ + C_I \frac{\nu_I - 1}{N_I + N_D} \bar{v}_{I1} + C_I \frac{N_D + N_I - \nu_I}{N_D + N_I} \bar{v}_{I2} + \frac{C_I}{N_I + N_D} v_1 + C_J \bar{v}_J, \end{aligned} \tag{109}$$

where

$$\bar{s}_{I1} = \frac{1}{\nu_I - 1} \left(\sum_{i=2}^{\psi N_I / N_D} s_i + \sum_{d=1}^{\psi} s_d \right); \qquad \bar{v}_{I1} = \frac{1}{\nu_I - 1} \left(\sum_{i=2}^{\psi N_I / N_D} v_i + \sum_{d=1}^{\psi} v_d \right);$$

$$\bar{s}_{I2} = \frac{1}{N_D + N_I - \nu_I} \left(\sum_{i=\psi N_I / N_D + 1}^{N_I} s_i + \sum_{d=\psi+1}^{N} s_d \right); \quad \bar{v}_{I2} = \frac{1}{N_D + N_I - \nu_I} \left(\sum_{i=\psi N_I / N_D + 1}^{N_I} v_i + \sum_{d=\psi+1}^{N} v_d \right);$$

$$\bar{s}_J = \frac{1}{N_J + 1} \sum_{j=1}^{N_J + 1} s_j; \qquad \bar{v}_J = \frac{1}{N_J + 1} \sum_{j=1}^{N_J + 1} v_j,$$

and $\nu_I = N_D / (\psi (N_I + N_D))$.

If bidder 1 chooses to bid directly, conglomerate 1 learns from that decisions, and all the other agents observe that decision but do not learn from the truncated normal signal, we can express the equilibrium price as

$$p = A_d + B_{I1} \bar{s}_{I1} + B_{I2} \bar{s}_{I2} + \frac{B_{J,d}}{N_J + 1} s_1 + \frac{B_{J,d} N_J}{N_J + 1} \bar{s}_J \tag{110}$$

$$+ C_{I1} \bar{v}_{I1} + C_{I2} \bar{v}_{I2} + \frac{C_{J,d}}{N_J + 1} v_1 + \frac{C_{J,d} N_J}{N_J + 1} \bar{v}_J + F s_q. \tag{111}$$

Notice that, in this case, we have one more direct bidder, and dealer conglomerate 1 has one less client than all the other conglomerates.

In this model, the price signal for dealers and their clients in conglomerate 1 is now

$$s\left(p|\, v_i\right) = \frac{p - A - C_{I1} v_i / (\nu_I - 1) - F m_q}{\tilde{B}},$$

for dealers and their clients in all other conglomerates is

$$s\left(p|\, v_i\right) = \frac{p - A - C_{I2} v_i / (N_D + N_I - \nu_I)}{\tilde{B}},$$

and for direct bidders

$$s\left(p|\, v_j\right) = \frac{p - A - C_J v_j / (N_J + 1)}{\tilde{B}},$$

where $\tilde{B} = B_{I1} + B_{I2} + B_J + F$.

The vector of orthogonal shocks Z is a a column vector of size $N_Z = 2 * N + 1$, where

$$Z = \left[\epsilon_1, \ldots, \epsilon_N, v_1, \ldots, v_N, e_q\right],$$

and the variance matrix of Z is

$$var\left(Z\right)=diag\left(\left[\tau_{\epsilon}^{-1}\mathbf{1}_{N},\tau_{v}^{-1}\mathbf{1}_{N},\tau_{q}^{-1}\right]\right).$$

Consider now representing the price signals. Let

$$\begin{aligned}\phi_{vI,1}&=\left[\mathbf{0}_{N},0,\mathbf{1}_{\psi N_{I}/N_{D}-1},\mathbf{0}_{N_{I}-\psi N_{I}/N_{D}},\mathbf{0}_{N_{J}},\mathbf{1}_{\psi},\mathbf{0}_{N_{D}-\psi},0\right]\\ \phi_{vI,2}&=\left[\mathbf{0}_{N},0,\mathbf{0}_{\psi N_{I}/N_{D}-1},\mathbf{1}_{N_{I}-\psi N_{I}/N_{D}},\mathbf{0}_{N_{J}},\mathbf{0}_{\psi},\mathbf{1}_{N_{D}-\psi},0\right]\\ \phi_{vJ}&=\left[\mathbf{0}_{N},1,\mathbf{0}_{N_{I}-1},\mathbf{1}_{N_{J}},\mathbf{0}_{N_{D}},0\right]\end{aligned}$$

be the vectors that select the private values of dealers and their clients and of direct bidders, respectively. Then, $\bar{v}_{I1}=\left(1/\left(\nu_{I}-1\right)\right)\phi_{vI,1}\cdot Z$, $\bar{v}_{I2}=\left(1/\left(N_{I}+N_{D}-\nu_{I}\right)\right)\phi_{vI,2}\cdot Z$ and $\bar{v}_{I}=\left(1/(N_{J}+1)\right)\phi_{vJ}\cdot Z$. Similarly, the vectors that select the signal noise are given by

$$\begin{aligned}\phi_{\epsilon I,1}&=\left[0,\mathbf{1}_{\psi N_{I}/N_{D}-1},\mathbf{0}_{N_{I}-\psi N_{I}/N_{D}},\mathbf{0}_{N_{J}},\mathbf{1}_{\psi},\mathbf{0}_{N_{D}-\psi},\mathbf{0}_{N},0\right]\\ \phi_{\epsilon I,2}&=\left[0,\mathbf{0}_{\psi N_{I}/N_{D}-1},\mathbf{1}_{N_{I}-\psi N_{I}/N_{D}},\mathbf{0}_{N_{J}},\mathbf{0}_{\psi},\mathbf{1}_{N_{D}-\psi},\mathbf{0}_{N},0\right]\\ \phi_{\epsilon J}&=\left[1,\mathbf{0}_{N_{I}-1},\mathbf{1}_{N_{J}},\mathbf{0}_{N_{D}},\mathbf{0}_{N},0\right].\end{aligned}$$

Thus, the price can be represented as

$$\begin{aligned}p&=A+\tilde{B}f+\frac{B_{I1}}{\nu_{I}-1}\phi_{\epsilon I,1}\cdot Z+\frac{B_{I2}}{N_{D}+N_{I}-\nu_{I}}\phi_{\epsilon I,2}\cdot Z+\frac{B_{J}}{N_{J}+1}\phi_{\epsilon J}\cdot Z\\ &+\frac{C_{I1}}{\nu_{I}-1}\phi_{vI,1}\cdot Z+\frac{C_{I2}}{N_{D}+N_{I}-\nu_{I}}\phi_{vI,2}\cdot Z+\frac{C_{J}}{N_{J}+1}\phi_{vJ}\cdot Z+F\phi_{q}\cdot Z\\ &\equiv A+\tilde{B}f+\tilde{B}\pi_{p}Z.\end{aligned}$$

With this representation of the equilibrium price, the information that a dealer or one of its clients in conglomerate 1 extracts from the price is

$$s\left(p|\,v_{i}\right)=\frac{p-A-Fm_{q}}{\tilde{B}}-\frac{C_{I1}}{\left(\nu_{I}-1\right)\tilde{B}}\phi_{N+i}\cdot Z\equiv s\left(p\right)-\frac{F}{\tilde{B}}m_{q}-\frac{C_{I1}}{\left(\nu_{I}-1\right)\tilde{B}}\phi_{N+i}\cdot Z,$$

a dealer or one of its clients in any other conglomerate extracts from the price is

$$s\left(p|\,v_{i}\right)=s\left(p\right)-\frac{C_{I2}}{\left(N_{I}+N_{D}-\nu_{I}\right)\tilde{B}}\phi_{N+i}\cdot Z,$$

and the signal that a direct investor extracts from the price is

$$s\left(p|\,v_{j}\right)=s\left(p\right)-\frac{C_{J}}{\left(N_{J}+1\right)\tilde{B}}\phi_{N+N_{I}+j}\cdot Z.$$

It now remains to determine how dealers aggregate their own signals together with the signals they get from their clients (and other dealers). Similarly to the belief weighting in the no intermediation choice model, a dealer in conglomerate 1 optimally averages all the signals to yeild

$$\tilde{s}_{d1}=f+\frac{1}{\nu_{I}-1}\sum_{i\in d_{\psi}(i)}\epsilon_{i}+\beta_{Is,1}^{-1}\left(1-\frac{\beta_{Ip,1}C_{I1}}{\tilde{B}\left(\nu_{I}-1\right)}\right)\frac{1}{\nu_{I}-1}\sum_{i\in d_{\psi}(i)}v_{i},$$

and a dealer belonging to any other conglomerate optimally averages all the signals to yeild

$$\tilde{s}_{d2} = f + \frac{1}{\nu_I} \sum_{i \in d_\psi(i)} \epsilon_i + \beta_{Is}^{-1} \left(1 - \frac{\beta_{Ip} C_{I2}}{\tilde{B}\left(N_I + N_D - \nu_I\right)}\right) \frac{1}{\nu_I} \sum_{i \in d_\psi(i)} v_i.$$

Thus, the signals for the indirect investors and dealers in conglomerate 1 are given by

$$\begin{bmatrix} s_i \\ s_{\xi i} \\ s\left(p|v_i\right) \\ s_q \end{bmatrix} = \begin{bmatrix} f \\ f \\ f \\ f \end{bmatrix} + \begin{bmatrix} \phi_i \\ \pi_{d1} \\ \pi_p - \frac{C_{I1}}{(\nu_I - 1)\tilde{B}} \phi_{N+i} \\ \phi_q \end{bmatrix} \cdot Z,$$

where

$$\pi_{d1} = \frac{1}{\nu_I - 1} \tilde{\phi}_{\epsilon d,1} + \beta_{Is,1}^{-1} \left(1 - \frac{\beta_{Ip,1} C_{I1}}{\tilde{B}\left(\nu_I - 1\right)}\right) \frac{1}{\nu_I} \tilde{\phi}_{vd,1}.$$

Thus, the signals for the indirect investors and dealers in any other conglomerate are given by

$$\begin{bmatrix} s_i \\ s_{\xi i} \\ s\left(p|v_i\right) \end{bmatrix} = \begin{bmatrix} f \\ f \\ f \end{bmatrix} + \begin{bmatrix} \phi_i \\ \pi_{d2} \\ \pi_p - \frac{C_{I2}}{(N_I + N_D - \nu_I)\tilde{B}} \phi_{N+i} \end{bmatrix} \cdot Z,$$

where

$$\pi_{d2} = \frac{1}{\nu_I} \tilde{\phi}_{\epsilon d,2} + \beta_{Is,2}^{-1} \left(1 - \frac{\beta_{Ip,2} C_{I2}}{\tilde{B}\left(N_I + N_D - \nu_I\right)}\right) \frac{1}{\nu_I} \tilde{\phi}_{vd,2}.$$

The signals for the direct dealers are given by

$$\begin{bmatrix} s_j \\ s\left(p|v_j\right) \end{bmatrix} = \begin{bmatrix} f \\ f \end{bmatrix} + \begin{bmatrix} \phi_j \\ \pi_p - \frac{C_J}{(N_J + 1)\tilde{B}} \phi_{N+N_I+j} \end{bmatrix} \cdot Z.$$

Using the first-order conditions and this belief representation, we can now rewrite the market clearing condition as:

$$
\begin{aligned}
1 &= M_{I,1} \sum_{i=2}^{\psi N_I/N_D} \left(\mathbb{E}\left[f|S_i\right] + v_i - p\right) + M_{I,2} \sum_{i=\psi N_I/N_D+1}^{N_I} \left(\mathbb{E}\left[f|S_i\right] + v_i - p\right) \\
&+ M_J \sum_{j=1}^{N_J+1} \left(\mathbb{E}\left[f|\, S_j\right] + v_j - p\right) + M_{I,1} \sum_{d=1}^{\psi} \left(\mathbb{E}\left[f|\, S_d\right] + v_d - p\right) \\
&+ M_{I,2} \sum_{d=\psi}^{N_D} \left(\mathbb{E}\left[f|\, S_d\right] + v_d - p\right) \\
&= M_{I,1} \sum_{i=2}^{\psi N_I/N_D} \left(\left(1 - \beta_{I1}'\mathbf{1}'\right)\mu + \beta_{I1}' S_i - p\right) + M_{I,2} \sum_{i=\psi N_I/N_D+1}^{N_I} \left(\left(1 - \beta_{I2}'\mathbf{1}'\right)\mu + \beta_{I2}' S_i - p\right) \\
&+ M_J \sum_{j=1}^{N_J+1} \left(\left(1 - \beta_J'\mathbf{1}'\right)\mu + \beta_J' S_i - p\right) + M_J \left(N_J + 1\right)\bar{v}_J \\
&+ M_{I,1} \sum_{d=1}^{\psi} \left(\left(1 - \beta_{I1}'\mathbf{1}'\right)\mu + \beta_{I1}' S_d - p\right) + M_{I,2} \sum_{d=\psi}^{N_D} \left(\left(1 - \beta_{I2}'\mathbf{1}'\right)\mu + \beta_{I2}' S_d - p\right) \\
&+ \left(\nu_I - 1\right) M_{I,1}\bar{v}_{I1} + \left(N_D + N_I - \nu_I\right) M_{I,2}\bar{v}_{I2}.
\end{aligned}
$$

Define $\tilde{M} = \left(\nu_I - 1\right) M_{I1} + \left(N_I + N_D - \nu_I\right) M_{I2} + \left(N_J + 1\right) M_J$. Breaking out the signal vectors into the individual components, we obtain

$$
\begin{aligned}
1 = \tilde{A} &+ \left(\nu_I - 1\right) M_{I1} \left(\beta_{Is,1}\bar{s}_{I1} + \beta_{I\xi,1}\left(\nu_I - 1\right)\tilde{s}_{d1} + \beta_{Ip,1}\left(s\left(p\right) - \frac{C_{I1}}{\left(\nu_I - 1\right)\tilde{B}}\bar{v}_{I1} - \frac{F}{\tilde{B}} m_q\right) + \beta_q s_q\right) \\
&+ \left(N_I + N_D - \nu_I\right) M_{I2} \left(\beta_{Is,2}\bar{s}_{I2} + \beta_{I\xi,2}\left(N_I + N_D - \nu_I\right)\tilde{s}_{d2} + \beta_{Ip,2}\left(s\left(p\right) - \frac{C_{I2}}{\left(N_I + N_D - \nu_I\right)\tilde{B}}\bar{v}_{I2}\right)\right) \\
&+ \left(N_J + 1\right) M_J \left(\beta_{Js}\bar{s}_J + \beta_{Jp}\left(s\left(p\right) - \frac{C_J}{\left(N_J + 1\right)\tilde{B}}\bar{v}_J\right)\right) - \tilde{M}p,
\end{aligned}
$$

where

$$
\tilde{A} = \left(\nu_I - 1\right) M_{I1}\left(1 - \beta_{I1}'\mathbf{1}'\right)\mu + \left(N_I + N_D - \nu_I\right) M_{I2}\left(1 - \beta_{I2}'\mathbf{1}'\right)\mu + M_J\left(N_J + 1\right)\left(1 - \beta_J'\mathbf{1}'\right)\mu.
$$

Using $s(p) = (p - A)/\tilde{B}$, we can collect terms in p to obtain

$$
\begin{aligned}
\tilde{Q}p = \tilde{A} - 1 - A\left(\tilde{Q} - \tilde{M}\right) &- M_{I1}\left(\nu_I - 1\right)\beta_{Ip,1}\frac{F}{\tilde{B}} m_q + M_{I1}\left(\nu_I - 1\right)\left(\beta_{Is,1} + \beta_{I\xi,1}\right)\bar{s}_{I1} \\
&+ M_{I2}\left(N_I + N_D - \nu_I\right)\left(\beta_{Is,2} + \beta_{I\xi,2}\right)\bar{s}_{I2} + M_J\left(N_J + 1\right)\beta_{Js}\bar{s}_J \\
&+ M_{I1}\left(\nu_I - 1\right)\left(1 + \frac{\beta_{I\xi,1}}{\beta_{Is,1}}\right)\left(1 - \frac{C_{I1}\beta_{Ip,1}}{\tilde{B}\left(\nu_I - 1\right)}\right)\bar{v}_{I1} \\
&+ M_{I2}\left(N_I + N_D - \nu_I\right)\left(1 + \frac{\beta_{I\xi,2}}{\beta_{Is,2}}\right)\left(1 - \frac{C_{I2}\beta_{Ip,2}}{\tilde{B}\left(N_I + N_D - \nu_I\right)}\right)\bar{v}_{I2} \\
&+ M_J\left(N_J + 1\right)\left(1 - \frac{C_J\beta_{Jp}}{\tilde{B}\left(N_J + 1\right)}\right)\bar{v}_J + \left(\nu_1 - 1\right) M_{I1}\beta_q s_q
\end{aligned}
$$

where $\tilde{Q} = \tilde{B}^{-1}\left(M_{I1}\left(\nu_I - 1\right)\left(\tilde{B} - \beta_{Ip,1}\right) + M_{I2}\left(N_I + N_D - \nu_I\right)\left(\tilde{B} - \beta_{Ip,2}\right) + M_J\left(N_J + 1\right)\left(\tilde{B} - \beta_{Jp}\right)\right)$.

Matching coefficients to the price equation, we obtain

$$\begin{aligned}
A &= \frac{1}{\tilde{Q}}\left(\tilde{A}-1+A\left(\tilde{Q}-\tilde{M}\right)\right) = \frac{1}{\tilde{M}}\left(\tilde{A}-1-M_{I1}\left(\nu_I-1\right)\beta_{Ip,1}\frac{F}{\tilde{B}}m_q\right)\\
B_{I1} &= \frac{1}{\tilde{Q}}M_{I1}\left(\nu_I-1\right)\left(\beta_{Is,1}+\beta_{I\xi,1}\right)\\
B_{I2} &= \frac{1}{\tilde{Q}}M_{I2}\left(N_I+N_D-\nu_I\right)\left(\beta_{Is,2}+\beta_{I\xi,2}\right)\\
B_J &= \frac{1}{\tilde{Q}}M_J\left(N_J+1\right)\beta_{Js}\\
C_{I1} &= \frac{1}{\tilde{Q}}M_{I1}\left(\nu_I-1\right)\left(1+\frac{\beta_{I\xi,1}}{\beta_{Is,1}}\right)\left(1-\frac{C_{I1}\beta_{Ip,1}}{\tilde{B}\left(\nu_I-1\right)}\right)\\
C_{I2} &= \frac{1}{\tilde{Q}}M_{I2}\left(N_I+N_D-\nu_I\right)\left(1+\frac{\beta_{I\xi,2}}{\beta_{Is,2}}\right)\left(1-\frac{C_{I2}\beta_{Ip,2}}{\tilde{B}\left(N_I+N_D-\nu_I\right)}\right)\\
C_J &= \frac{1}{\tilde{Q}}M_J\left(N_J+1\right)\left(1-\frac{C_J\beta_{Jp}}{\tilde{B}\left(N_J+1\right)}\right)\\
F &= \frac{1}{\tilde{Q}}M_{I1}\left(\nu_1-1\right)\beta_q.
\end{aligned}$$

Finally, analogously to the no intermediation choice model, the price impact of indirect bidders and dealers in conglomerate 1 is given by

$$\frac{dp}{dq_{I1}} = \tilde{B}\left[M_{I1}\left(\nu_I-2\right)\left(\tilde{B}-\beta_{Ip,1}\right)+M_{I2}\left(N_D+N_I-\nu_I\right)\left(\tilde{B}-\beta_{Ip,2}\right)+M_J\left(N_J+1\right)\left(\tilde{B}-\beta_{Jp}\right)\right]^{-1},$$

the price impact of indirect bidders and dealers in any other conglomerate by

$$\frac{dp}{dq_{I2}} = \tilde{B}\left[M_{I1}\left(\nu_I-1\right)\left(\tilde{B}-\beta_{Ip,1}\right)+M_{I2}\left(N_D+N_I-\nu_I-1\right)\left(\tilde{B}-\beta_{Ip,2}\right)+M_J\left(N_J+1\right)\left(\tilde{B}-\beta_{Jp}\right)\right]^{-1},$$

and the price impact of direct bidders by

$$\frac{dp}{dq_J} = \tilde{B}\left[M_{I1}\left(\nu_I-1\right)\left(\tilde{B}-\beta_{Ip,1}\right)+M_{I2}\left(N_D+N_I-\nu_I\right)\left(\tilde{B}-\beta_{Ip,2}\right)+M_JN_J\left(\tilde{B}-\beta_{Jp}\right)\right]^{-1}.$$

Agents outside of conglomerate 1 do not observe the intermediation decision. That is, they perceive the price to be a probability-weighted average of the pricing coefficients in (109) and (110):

$$p = \bar{A}+\bar{B}_{I1}\bar{s}_{I1}+\bar{B}_{I2}\bar{s}_{I2}+\bar{B}_{J1}s_1+\bar{B}_{J2}\bar{s}_J+\bar{C}_{I1}\bar{v}_{I1}+\bar{C}_{I2}\bar{v}_{I2}+\bar{C}_{J1}v_1+\bar{C}_{J2}\bar{v}_J+\bar{F}s_q, \tag{112}$$

where

$$\begin{aligned}
\bar{A} &= \left(1-p_q\right)A+p_qA_d; & \bar{B}_{I1} &= \left(1-p_q\right)B_I\frac{\nu_I-1}{N_I+N_D}+p_qB_{I1};\\
\bar{B}_{I2} &= \left(1-p_q\right)B_I\frac{N_D+N_I-\nu_I}{N_D+N_I}+p_qB_{I2}; & \bar{B}_{J1} &= \left(1-p_q\right)\frac{B_I}{N_D+N_I}+p_q\frac{B_{J,d}}{N_J+1};\\
\bar{B}_{J2} &= \left(1-p_q\right)B_J+p_q\frac{B_{J,d}N_J}{N_J+1}; & \bar{C}_{I1} &= \left(1-p_q\right)C_I\frac{\nu_I-1}{N_I+N_D}+p_qC_{I1};\\
\bar{C}_{I2} &= \left(1-p_q\right)C_I\frac{N_D+N_I-\nu_I}{N_D+N_I}+p_qC_{I2}; & \bar{C}_{J1} &= \left(1-p_q\right)\frac{C_I}{N_D+N_I}+p_q\frac{C_{J,d}}{N_J+1};\\
\bar{C}_{J2} &= \left(1-p_q\right)C_J+p_q\frac{C_{J,d}N_J}{N_J+1}; & \bar{F} &= p_qF.
\end{aligned}$$

Dealers and clients of conglomerate 1, on the other hand, know the intermediation choice made by client 1, and perceive the price to be different conditional on the intermediation choice.

A.7 The Role of Risk Aversion

Since risk aversion is always a difficult parameter to identify with aggregate data, we show results with risk aversion that is 50% higher and 50% lower than our baseline value of 448. Table 4 shows that while the exact revenue and utility numbers change, the ordering and magnitudes are quite stable.

Table 4: The Role of Risk Aversion.

	Baseline τ_ξ	Chinese wall	Open order book
		Auction Revenue	
Baseline	36.740	32.884	39.311
1.5ρ	38.831	36.014	40.708
$\rho/1.5$	36.616	32.699	39.227
		Bidder Utility	
Baseline	0.810	1.915	0.975
1.5ρ	1.457	1.709	0.923
$\rho/1.5$	1.436	1.925	0.983

B Measuring Treasury Payoffs

This appendix provides additional detail about how payoffs are calculated. Because of lags between trade and settlement dates, the appendix also provides detail on funding costs. We begin by reporting summary statistics of post-auction appreciation and the speculative (competitive) share (Table 5). Then, we go into detail about how these variables are constructed and what alternative methods yield. The first subsection describes what those terms are and argues that they are small and stable. The second subsection discusses an alternative hedging strategy, known as a coupon roll. The third explains why information from the when-issued-market (or WIs) is not relevant in our setting.

Table 5: Summary statistics for Table 3. Shares are in percent; post-auction appreciation $(f - p)$ is measured in basis points.

	Post-auction appreciation $(f - p)$	Spec. share	PD share	non-PD Spec. share
Mean	2.7	78.8	53.3	25.5
P25	-8.6	74.1	42.2	14.2
P50	-1.0	79.9	51.2	25.8
P75	10.8	84.5	63.7	34.8
Stdev	30.2	9.1	14.4	13.4
Obs	494	494	494	494

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

$$\hat{R}_a = 1 - \left(\sum_{t=0}^{T} Z_a(t)\, C + Z_a(T) \right), \tag{113}$$

where C is the coupon determined at the auction, T is the maturity, $Z_a(j) = exp(-i \times y_a(j))$ is the price at the time of the auction of a zero-coupon bond maturing at i, $y(j)$ is the jth maturity yield from the yield curve estimated on outstanding securities at the time of the auction.

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

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

$$\hat{f}_i = \left(\sum_{t=0}^{T} Z_i(t)\, C + Z_i(T) \right) - P_i, \tag{114}$$

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

1. Auction date:
   (a) Place bid
   (b) For each unit of successful bid alloted, sell T zero coupon bond each priced at $Z_a(t)$ and in amounts equal to C for $t < T$ and $1 + C$ for $t = T$. The zero coupon bonds could either be stripped Treasuries (as in Fleckenstein, Longstaff, and Lustig, 2014) or proxied with a combination of coupon securities.
2. Post-auction date:
   (a) Borrow (to post-issuance date) the amount $\mathcal{Z}_a = \sum_{t=0}^{T} Z_a(t)\, C + Z_a(T)$ paying the per-diem unsecured rate r_b.
   (b) Borrow zero-coupon bonds with reverse repos (to post-issuance date) and receive the per-diem repo rate r_r. Deliver the T zero coupons to the auction-date buyer.
3. Issuance date:

   (a) Borrow \$1 at rate r_b. Receive new issue from, and pay \$1, to Treasury; sell issue in the secondary market
   (b) Buy portfolio of T zero-coupon bonds at $\mathcal{Z}_{issuance}$

4. Post-issuance date:
   (a) Receive payment of p_i and repay the issuance-date loan
   (b) Receive T zero-coupon bonds and deliver into the reverse repo;
   (c) Receive payment of $\mathcal{Z}_a$ from reverse-repo and pay $\mathcal{Z}_i$ to settle the issue-date purchase; Repay post-auction date loan

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

$$(P_i - 1) + (\mathcal{Z}_a - \mathcal{Z}_i) + \frac{(\text{date}_i - \text{date}_a)}{360} \times (r_r - r_b) \times \mathcal{Z}_a - \frac{r_b}{360} \times 1 \tag{115}$$

In our calculations we disregard the two funding terms because they are small and don't vary much when $r_r \approx r_b$. The repo rate for old issues, which are being funded between the post-auction and post-issue date, typically trades within a few basis points to the unsecured rate r_b, so the funding terms are small. For example (see e.g. Duffie, 1996), reports that first off-the-runs repo rates around about 25 basis points below the (general collateral) repo rate. This difference has only a minimal impact on the payoff as the position is only held between the auction and issue dates. Furthermore, off-the-run securities rarely go "on special" as indirectly observed in the Federal Reserve's securities loan auctions (Fleming and Garbade, 2007a). Instead, repo rates for new (or first-off-the-run) securities can trade far off from uncollateralized rates and be volatile because funding rates balance the supply and demand of new securities, which can be in high demand to take short position in interest rates (see e.g. Duffie, 1996; Jordan and Jordan, 1997). As per the detailed steps above the new issue is never shorted or funded, as it is sold as soon as it is received by the investor. Thus fluctuations in the special-repo rate do not affect the returns in our position.

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