Competition and Information Leakage

by

Joshua Mollner
Markus Baldauf

Working Paper No. 00096-01

Finance Theory Group

www.financetheory.com

Competition and Information Leakage*

Markus Baldauf† Joshua Mollner‡

July 28, 2023

Journal of Political Economy, forthcoming

Abstract

When seeking to trade in over-the-counter markets, institutional investors typically restrict both the number of potential counterparties they contact and the information they disclose (e.g., by requesting two-sided rather than one-sided quotes). We rationalize these important facts in a model featuring endogenous front-running. Although an additional contact intensifies competition and aids in finding a natural counterparty, it also intensifies information leakage—which can be costly if it helps a losing dealer to front-run. We also address information design: the client optimally provides no information about her trading direction when requesting quotes. We conclude with implications for market design and regulation.

Keywords: counterparty search, over-the-counter (OTC) markets, request for quotes (RFQ), trading platform design, information design, price impact, front-running

JEL Codes: D82, D83, G14, G23

* *First version:* December 7, 2020. We thank Anat Admati, Sabrina Buti (discussant), Eric Budish, Douglas Diamond, Piotr Dworczak, Timothy Feddersen, Lorenzo Garlappi, Lewen Guo (discussant), Lawrence Harris, Joel Hasbrouck, Zhiguo He, Maureen O'Hara, Terrence Hendershott, Esen Onur, Alexei Orlov, Carlos Ramírez, Marzena Rostek, Alvaro Sandroni, James Schummer, Andrzej Skrzypacz, Chester Spatt, Alireza Tahbaz-Salehi, Basil Williams (discussant), Haoxiang Zhu, Marius Zoican (discussant), and three anonymous referees for helpful comments. We thank Kamila Giesbrecht and various other industry participants for useful conversations. Baldauf gratefully acknowledges support by the Social Sciences and Humanities Research Council of Canada. Baldauf thanks the University of Chicago, Booth School of Business, where he was visiting when large parts of this paper were written. We also thank the editor, Emir Kamenica.

†University of British Columbia, Sauder School of Business: baldauf@mail.ubc.ca, https://sites.google.com/site/mbaldauf/.

‡Northwestern University, Kellogg School of Management: joshua.mollner@kellogg.northwestern.edu, https://sites.google.com/site/joshuamollner/.

1 Introduction

Most financial trading occurs over-the-counter—that is, away from centralized exchanges. Trading in this way requires searching for a counterparty, which, due to regulatory reform and technological innovation, is now much easier than ever before. Nevertheless, when seeking to trade, institutional investors typically request quotes from only a handful of dealers. Given the magnitude of institutional transaction costs, as well as the clear benefits of competition and liquidity, these facts are difficult to reconcile.[1] We rationalize this tension in a model that features information leakage—the risk that dealers who are contacted but not selected might front-run on the market—which, indeed, market participants themselves recognize as an important consideration. Even if the explicit cost of contacting an additional dealer is small, the implicit cost of information leakage may loom large. Of course, institutions must balance this cost against the potential benefits of more competitive quotes and better odds of locating a natural counterparty.

We formalize this tradeoff in a theoretical model. A client wishes to either buy or sell a security and seeks fulfillment through a dealer. This client could represent either an institutional investor (such as a pension fund or an endowment) or a traditional firm accessing the financial system. Our main results answer how she optimally orchestrates her counterparty search. One lever is how many dealers to contact. We show that the client may find it optimal to contact fewer than all available dealers. In that sense, our analysis suggests that search frictions in financial markets are often best understood through the lens of information leakage: search leaks information, which exacerbates front-running. The client's other lever is what information to provide while soliciting quotes. Optimizing over all information structures, we show that the client is always better off when providing no information about her trading direction. This rationalizes the common practice of requesting two-sided quotes.

These results further our understanding of the complex tradeoffs faced by participants in over-the-counter markets, some of which are among the largest markets in the world. As we discuss below, our results also have implications for the design and regulation of trading platforms employed in many of these markets. Stakeholders in these settings should be cognizant of front-running.

Model and equilibrium. In the model, the client issues a request for quotes (RFQ), contacting either one or two dealers and providing them with a potentially informative signal of her trading direction. Each dealer is either long or short the security, where these positions are unknown to the client (although, for simplicity, we assume they are common knowledge among the dealers). The contacted dealers respond with two-sided quotes: one quote at which to buy and another at which to sell. The client conducts a sealed-bid, second-price auction (with reserve). The winning dealer then learns the client's trading direction, while the losing dealer can only make inferences based on the RFQ itself, together with the auction's outcome. The winning dealer might fulfill the client's order in either of two ways. First, he might internalize the order against his inventory. However,

[1] For example, Nasdaq (2022) and SIFMA (2021) estimate institutional transaction costs for U.S. equities of around \$70 billion per year.

we assume position limits for the dealers, so that internalization is possible only if the dealer is long (short) and the client's order is to buy (sell). Second, the winning dealer might fulfill the client's order by trading on the market, where our model permits two periods of on-market trading. The other dealer might also trade on the market, whether he lost the RFQ or was not contacted at all.

In equilibrium, a losing dealer trades on the market in roughly two distinct ways. If he is long (short) while the client seeks to buy (sell), then he might provide liquidity to the winning dealer by selling (buying), which reduces the winning dealer's trading costs. On the other hand, he might front-run the winning dealer by buying (selling) initially and subsequently reversing those trades, which increases the winning dealer's trading costs. Both types of trading are profitable for the losing dealer. Anticipating this, dealers' bids account for both (*i*) the trading costs they would incur if they won, and (*ii*) the opportunity cost consisting of the profits they would obtain if they lost.

Our main results pertain to *RFQ policies*, which specify how the number of dealers she contacts, the signal she provides, and her reserve prices will depend on the client's realized trading direction. One question is what information the client ought to provide while issuing the RFQ. Simply contacting a dealer already indicates *that* the client is seeking to trade, which may be informative, even absent any information about *how* she is seeking to trade. Furthermore, the number of dealers contacted might itself signal how she is seeking to trade (for instance, if her policy is to contact two dealers more frequently when buying than when selling). But she might provide other information beyond this. In the extreme of *minimal disclosure*, she provides no such additional information. In our model, the prior distribution of client trades has two-point support: the client will seek either to buy or to sell a fixed amount. Hence, minimal disclosure is equivalent to asking contacted dealers to "make a two-sided market" (and saying nothing else about her trading direction).

Can the client do better by revealing additional information? In general, she provides the dealers with a signal about her trading direction, as in Bayesian persuasion (Kamenica and Gentzkow, 2011). For example, one possibility is the opposite extreme: fully revealing her trading direction (in other words, asking for a "one-sided market"). We obtain a strong result: minimal disclosure is unambiguously optimal for the client among all information structures. Notably, this result is in line with common industry practice, where additional information is rarely volunteered at the RFQ stage (e.g., Risk.net, 2018).

To understand this optimality of minimal disclosure, notice that any such additional information is relevant only through its effect on what the losing dealer knows. (The winning dealer learns the direction of the client's trade after the auction, regardless of what is revealed in conjunction with the RFQ.) Thus, additional information facilitates a losing dealer's ability to trade on the market, potentially increasing the scope for both front-running and liquidity provision. This creates three effects: (*i*) the losing dealer's profits (i.e., the winning dealer's opportunity cost) increase, (*ii*) the winning dealer's trading costs increase whenever the losing dealer uses the additional information to front-run, and (*iii*) the winning dealer's trading costs decrease whenever the losing dealer uses the information to provide liquidity. The optimality of minimal disclosure indicates that the dominant effects are *(i)* and *(ii)*, both of which lead dealers to bid less aggressively, to the detriment of the client.

Finally, we answer the related question about how many dealers the client should contact. A broad search has benefits. It intensifies competition among the dealers. It also improves the odds of locating a natural counterparty—a dealer who might, by virtue of an existing position, be able to internalize the client's order rather than expose it to the market. Nevertheless, the client might optimally contact fewer than all available dealers. This contrasts with conventional auction models, where the presence of additional bidders tends to unambiguously benefit the auctioneer (Bulow and Klemperer, 1996). The reason is that these dealers interact not only in the auction but also after the auction, on the market. In particular, dealers who lose in the auction after being contacted may front-run on the market, raising the winning dealer's trading costs; a dealer may therefore bid *less* aggressively when facing more competitors in the auction. Hence, information leakage raises the client's procurement cost not through its effect on first-order beliefs (i.e., that more dealers know the client's trade) but rather through second-order beliefs (i.e., that dealers know that more dealers know). The client optimally contacts only a single dealer when this risk of front-running is highest: when she needs to buy (or, respectively, sell) and when there is a sufficiently large prior probability that the dealers are initially short (or, respectively, long).

Thus, our model rationalizes two important facts about over-the-counter markets: when seeking to trade, institutional investors typically (*i*) limit the number of dealers whom they contact and (*ii*) limit information disclosure (e.g., by requesting two-sided rather than one-sided quotes). To emphasize that front-running is the key to rationalizing these facts, we contrast our results against those that would be obtained if front-running were exogenously shut down. Formally, the appendix considers a version of the model in which the losing dealer cannot trade in the first period. What we see is that, without front-running, our model cannot rationalize these facts, reverting instead to traditional intuitions about auction design: more bidders are better, and there is no need to conceal information.

Applications. Historically, over-the-counter trading was primarily conducted via telephone. Since then, new technologies have reduced search costs. Bloomberg chat and email have largely replaced the traditional phone call. Moreover, a variety of electronic *RFQ platforms* have emerged, where a trader enters information about a desired trade, selects a set of dealers to simultaneously request quotes from, and from the responses may then select one quote to trade against.

Our model of the client's procurement process most closely reflects the operation of these RFQ platforms. Indeed, shedding light on trading incentives on RFQ platforms is valuable, given their economic importance. These platforms play a key role in fixed income (e.g., treasuries, corporate bonds, municipal bonds, mortgage-backed securities), in derivatives (e.g., interest rate swaps, credit default swaps, options), as well as in foreign exchange, equities, energy markets, and crypto.[2,3] And consistent with their centrality to our analysis, information leakage and front-running are often cited as important considerations by regulators and market participants in these settings (CFTC, 2013; BIS, 2016; CFTC, 2018; EDMA Europe, 2018; FX Markets, 2021; SEC, 2022).

These RFQ platforms can differ in various specifics, including whether they constrain a trader in choosing how many potential counterparties to contact and whether to reveal her trading direction. One interpretation of our results is as a positive description of the behavior to be expected on platforms that do not constrain traders' choices, where our results rationalize some behaviors observed in practice. Other platforms do constrain traders' choices, and so another interpretation (explored further in Section 3.7) is that our results have normative implications for platform design and regulation. Regarding regulation, our analysis supports arguments against a CFTC rule that mandates contacting a minimum number of potential counterparties when trading certain swaps. Regarding platform design, our analysis highlights the welfare implications of details concerning order exposure, information policy, and anonymity. We also argue that considerations of information leakage can help to explain the puzzle of why electronic platforms have not seized a larger role in over-the-counter markets.

Not all over-the-counter trading occurs on electronic platforms. Beyond the application to electronic RFQ platforms, our model of the client's procurement process might also represent certain forms of off-platform over-the-counter trading. In these settings, the client orchestrates the RFQ herself, e.g., using email or Bloomberg chat.

The on-market trading of our model could, depending on asset class, represent a centralized exchange (as exists in, e.g., equities trading). Alternatively, it could represent an inter-dealer market, in which case our analysis is most applicable to settings where the inter-dealer market is active, as for index swaps, foreign exchange, or thickly-traded bonds (and less applicable to settings with inactive inter-dealer markets, such as single-name swaps or thinly-traded bonds).

Related literature. One connection is to the literature on over-the-counter markets pioneered by Duffie, Gârleanu and Pedersen (2005). In these models, exogenous technological constraints explain why a client contacts only a small number of potential counterparties.[4] This explanation

[2] Several such platforms have been studied in the academic literature. For example, Hendershott and Madhavan (2015), Hendershott, Livdan and Schürhoff (2021), and O'Hara and Zhou (2021) study the trading of U.S. corporate bonds on MarketAxess; Collin-Dufresne, Junge and Trolle (2020) and Riggs, Onur, Reiffen and Zhu (2020) study index credit default swaps on swap execution facilities such as those of Bloomberg and Tradeweb; and Allen and Wittwer (forthcoming) study Canadian government bonds on CanDeal. Other platforms include CME Group's directed request for quote (DRFQ) for block trades in certain futures and options. In energy markets, RFQ platforms include those of Enmacc and ShapeQ, and in crypto markets, Paradigm, Dyfn, and PowerTrade.

[3] Based on a variety of sources, we conservatively compute the fraction of trading that takes place via RFQ platforms to be 31% in U.S. treasuries, 22.5% in U.S. corporate bonds, 14.3% in swaps, 60% in options, 21.6% in foreign exchange, and 50% in European exchange-traded products. Many of these asset classes feature separate dealer-to-client and dealer-to-dealer segments; in such cases, RFQ platforms typically account for even larger fractions of the dealer-to-client trading. See Online Appendix H for details.

[4] In a similar spirit, a sequential searcher might stop before exploring every option if search entails explicit costs (Stigler, 1961; Weitzman, 1979, or various IO search models). Likewise, an auctioneer might limit participation if it is costly to solicit bidders (Lauermann and Wolinsky, 2017; Riggs, Onur, Reiffen and Zhu, 2020).

has, however, become increasingly less compelling as new technologies—electronic platforms, email, Bloomberg chat—have largely eroded explicit search costs. Indeed, explicit search costs in many modern electronic markets are not just small, but effectively zero (Budish, Lee and Shim, forthcoming). Instead, our model suggests that the more pertinent search friction is the implicit cost of information leakage.[5]

Although we are not the first to recognize that information leakage may constitute a search friction, we contribute by proposing a model of trading in which it does so endogenously. In the paper most closely related to ours, Hendershott and Madhavan (2015) explore a similar tradeoff between competition and information leakage. Although their primary emphasis is an empirical analysis of corporate bond trading, they offer a simple model to serve as a theoretical framework. In their model, competition and information leakage are captured by exogenous objects: d parametrizes the markup that a dealer sets in bilateral trade, π the competitiveness of a multi-dealer auction, q the probability of auction failure, and s the adverse price movement due to information leakage if the auction fails, where each can flexibly depend on a vector of characteristics x. The benefits of competition are therefore captured by $[1 - q(x)][d(x) - \pi(x)]$, and the costs of information leakage by $q(x)s(x)$.[6] We instead endogenize this tradeoff: we derive the costs of information leakage via a continuation trading game, and the benefits of competition by solving for equilibrium bids. Endogenizing the cost of information leakage not only puts sound theoretical foundations underneath the intuitive tradeoffs but also allows us to analyze when this cost is likely to loom large or small, as well as how it interacts with the information provided to potential counterparties at the RFQ stage.

Another novel contribution is to explore the role of information design in this setting. Our result on the optimality of minimal disclosure provides a microfoundation for the common practice of requesting two-sided quotes. Methodologically, our analysis uses tools from Bayesian persuasion (Kamenica and Gentzkow, 2011). Information design plays a role in our model because bidders' values in the client's procurement auction are endogenous, determined by expectations about subsequent trading—which information provision affects. This distinguishes our setting from auction models in which information provision serves instead to provide bidders with better information about their own exogenous values (Milgrom and Weber, 1982; Bergemann and Pesendorfer, 2007; Eső and Szentes, 2007; Ganuza and Penalva, 2010; Li and Shi, 2017; Bergemann et al., 2022). The subsequent trading also connects our model to the literature on auctions with resale, partic-

[5] Besides information leakage, several other frictions could similarly inhibit competition. For example, an auctioneer might limit participation if it is costly for bidders to learn their values (Levin and Smith, 1994) or to prepare their bids (Menezes and Monteiro, 2000; Wang, 2023; Yueshen and Zou, 2023); in common-value or almost-common-value auction settings, due to a winner's curse (Bulow and Klemperer, 2002; Riggs, Onur, Reiffen and Zhu, 2020); as well as in tournaments and contests (Taylor, 1995; Fullerton and McAfee, 1999; Che and Gale, 2003). In sequential search, a trader might stop before reaching all potential counterparties if quotes expire immediately, because a repeated contact signals that the trader has a low outside option (Zhu, 2012). And a trader might opt for bilateral bargaining over simultaneous multilateral search if the former provides her with more bargaining power (Glebkin, Yueshen and Shen, 2023).

[6] The cost of information leakage is also exogenous in Burdett and O'Hara (1987), who analyze a related problem in which a block trader—a dealer who has already been approached by a large client and is now assembling a syndicate to take the other side of the trade—determines how many potential counterparties to contact.

ularly to papers where the resale motive stems from bidders' inventory constraints (Bikhchandani and Huang, 1989).[7] Information provision has an effect in our setting because it promotes front-running. Similar front-running occurs also in the literature on predatory trading (Brunnermeier and Pedersen, 2005; Carlin, Lobo and Viswanathan, 2007; Sannikov and Skrzypacz, 2016). In our model the extent of front-running is endogenous, influenced, for example, by the client's information design choices.

2 Model

We begin with a formal description of the model. We then discuss modeling choices and introduce parametric assumptions.

2.1 Setup

Players. There are three players: a client and two dealers, labelled A and B. The client realizes a need to either buy or sell $\bar{s} > 0$ shares. We denote her trading need by a random variable s, where a realization $s = \bar{s}$ means that she seeks to buy, and $s = -\bar{s}$ that she seeks to sell. One thing the client must decide is how many dealers to contact. Because they are ex ante symmetric, it is without loss to assume that if only one dealer is contacted, it is dealer A. We therefore define $\mathcal{I}_1 = \{A\}$ and $\mathcal{I}_2 = \{A, B\}$ as the sets of dealers whom the client might contact.

Timing and information sets. First, the client observably commits to an *RFQ policy*, which consists of a finite realization space Σ together with a profile of distributions $(\pi_{s'})_{s' \in \{-\bar{s}, \bar{s}\}}$ over $\Sigma \times \{1, 2\} \times \mathbb{R}^2$. The interpretation is that, given the realized s, the distribution π_s will be used to generate a realization $(\sigma, M, \bar{b}) \in \Sigma \times \{1, 2\} \times \mathbb{R}^2$. At that point, the dealers $i \in \mathcal{I}_M$ will be contacted, informed of σ, and invited to participate in a second-price auction with M bidders and with reserve prices $\bar{b} = (\bar{b}_{-\bar{s}}, \bar{b}_{\bar{s}})$.[8]

Second, Nature draws the client's trading need $s \in \{-\bar{s}, \bar{s}\}$, where $\phi_0 \in [0, 1]$ is the prior probability of $\bar{s}$. Nature also draws a vector of initial dealer inventories $(e^A, e^B) \in \{-\bar{e}, \bar{e}\} \times \{-\bar{e}, \bar{e}\}$, with $\bar{e} \geq 0$. Both dealers commonly observe the entire vector of realized inventories. We

[7] In other models, the resale motive originates from ex-ante asymmetries among the bidders (Gupta and Lebrun, 1999; Zheng, 2002; Calzolari and Pavan, 2006; Garratt and Tröger, 2006; Hafalir and Krishna, 2008; Carroll and Segal, 2019); from new information learned after the auction (Haile, 2003); or from collusion (Garratt, Tröger and Zheng, 2009). Also related is the literature on auctions with externalities, which is motivated by settings with downstream interaction among the bidders (Jehiel, Moldovanu and Stacchetti, 1996; Jehiel and Moldovanu, 1996, 2000). A difference is that, in our model, the auction affects the downstream interaction not only through the allocation it generates but also through the number of bidders invited to participate and the information revealed.

[8] As in Kamenica and Gentzkow (2011), RFQ policies can implement any Bayes-plausible distribution over posterior beliefs about s. For example, let $\Sigma = \{\texttt{Sell}, \texttt{Buy}\}$ and for $\alpha \in [0, 1]$

$$\pi_{-\bar{s}} = \begin{cases} (\texttt{Sell}, 2, (\bar{s}^2, \bar{s}^2)) & \text{w.p. } \frac{1+\alpha}{2}, \\ (\texttt{Buy}, 2, (\bar{s}^2, \bar{s}^2)) & \text{w.p. } \frac{1-\alpha}{2}, \end{cases} \qquad \pi_{\bar{s}} = \begin{cases} (\texttt{Sell}, 2, (\bar{s}^2, \bar{s}^2)) & \text{w.p. } \frac{1-\alpha}{2}, \\ (\texttt{Buy}, 2, (\bar{s}^2, \bar{s}^2)) & \text{w.p. } \frac{1+\alpha}{2}. \end{cases}$$

At one extreme, the policy with $\alpha = 0$ reveals nothing about s. At the other, $\alpha = 1$ fully reveals s.

parameterize the joint distribution of (e^A, e^B) so that (*i*) $\rho \in [-1, 1]$ is the correlation of e^A and e^B, and (*ii*) for each dealer i, $\psi \in \left(\frac{\rho^-}{1+\rho^-}, \frac{1}{1+\rho^-}\right)$ is the marginal probability that $e^i = \bar{e}$, where $\rho^- := -\min(\rho, 0)$. To that end:

$$(e^A, e^B) = \begin{cases} (\bar{e}, \bar{e}) & \text{w.p. } \psi[1 - (1-\psi)(1-\rho)], \\ (\bar{e}, -\bar{e}) & \text{w.p. } \psi(1-\psi)(1-\rho), \\ (-\bar{e}, \bar{e}) & \text{w.p. } \psi(1-\psi)(1-\rho), \\ (-\bar{e}, -\bar{e}) & \text{w.p. } (1-\psi)[1 - \psi(1-\rho)]. \end{cases}$$

Third, the client follows through on the RFQ policy to which she previously committed. That is, $(\sigma, M, \bar{b}) \sim \pi_s$ is drawn and observed by each dealer $i \in \mathcal{I}_M$. Such a realization is an *RFQ*.

Fourth, each dealer $i \in \mathcal{I}_M$ submits a bid: a vector $b^i = (b^i_{-\bar{s}}, b^i_{\bar{s}}) \in \mathbb{R}^2$, where each component represents the smallest commission that dealer i will accept to facilitate the corresponding trade.

Fifth, a second-price auction with reserve is held. Let $b_s^{(m)}$ denote the mth order statistic among $(b^i_s)_{i \in \mathcal{I}_M}$. As this is a procurement auction, low bids are preferred. The winning dealer is chosen uniformly at random from $\left\{i \in \mathcal{I}_M : b^i_s = \min(\bar{b}_s, b_s^{(1)})\right\}$. If there is a winner, then the auction also determines a procurement cost $c := \min(\bar{b}_s, b_s^{(2)})$. The winning dealer then observes s, while any losing dealers observe only the identity of the winner. If there is no winner, the game ends.[9]

Sixth, two periods of on-market trading occur. For the first such period, each dealer $i \in \{A, B\}$ simultaneously submits a market order to buy $x^i_1 \in \mathbb{R}$ shares. These first-period trades occur at the price $p_1 = p_0 + \theta(x^A_1 + x^B_1)$, where $\theta > 0$ and $p_0 \in \mathbb{R}$. Having observed p_1, each dealer $i \in \{A, B\}$ then simultaneously submits a market order to buy $x^i_2 \in \mathbb{R}$ shares. These second-period trades occur at the price $p_2 = p_1 + \theta(x^A_2 + x^B_2)$. In making trading decisions, each dealer i is constrained to keep his final inventory within $\bar{e}$ shares, long or short: if he wins, then $e^i + x^i_1 + x^i_2 - s \in [-\bar{e}, \bar{e}]$, and if he does not win, then $e^i + x^i_1 + x^i_2 \in [-\bar{e}, \bar{e}]$. In addition, if a dealer i is not contacted, then he cannot trade in the first period: $x^i_1 = 0$. And to simplify the proofs, we also assume the constraint $x^i_1 \in [-\bar{s}, \bar{s}]$, which never binds and is made only to reduce the number of cases that must be formally checked.[10]

Seventh, the client pays the winning dealer $c + sp_0$ in exchange for s shares. Outstanding positions are then liquidated for a dividend of p_0.

Payoffs. If there is a winner, then the client's procurement cost is c. We say that an RFQ policy is *optimal* for the client if it minimizes her expected procurement cost within the class of policies that

[9] Ending the game here is equivalent to requiring that all subsequent on-market trades by the dealers be for zero shares (i.e., $x^i_1 = x^i_2 = 0$ for $i \in \{A, B\}$). We could allow for nonzero on-market trades when there is no winner; however, no such trading would take place in the equilibrium of any such subgame. Thus, ending the game is a harmless assumption made only for simplicity.

[10] Here is the intuition for why this constraint on first-period trade sizes does not bind in equilibrium. The most the winning dealer ever needs to trade on the market (to satisfy constraints on his final inventory) is $\bar{s}$. Due to price impact, he would not wish to trade more than that amount in total—and certainly not more than that amount in the first period. Given that, the other dealer also would not wish to trade more than $\bar{s}$ in the first period.

lead to execution with probability one.[11] If dealer i wins, his payoff is $c+(e^i+x_1^i+x_2^i)p_0-p_1x_1^i-p_2x_2^i$. If dealer i does not win, his payoff is $(e^i+x_1^i+x_2^i)p_0-p_1x_1^i-p_2x_2^i$. All players are risk-neutral expected utility maximizers.

2.2 Remarks

Modeling the client. We have assumed that the client can obtain execution only by trading with a dealer (and cannot trade on the market directly).[12] This assumption is realistic for many institutional investors who commonly lack either the infrastructure, expertise, membership rights, or regulatory clearance to access markets directly.

Our formulation of the client's RFQ policy generalizes how communication is modeled in Bayesian persuasion (Kamenica and Gentzkow, 2011). In that literature, a sender, who will privately observe her type $s \in \mathcal{S}$, commits to a signal realization space Σ together with a profile of distributions $(\pi_{s'})_{s' \in \mathcal{S}}$ over Σ. In our formulation, the client likewise commits to a policy that determines the signal she will send—but also the number of dealers she will contact and the reserve prices she will set. A standard justification for this commitment assumption is that commitment power may come informally through reputation and repeated interaction. This standard justification seems particularly well-suited to our setting, as clients in practice do indeed tend to interact repeatedly with a limited number of dealers. Trading firms might also commit to such policies through the business protocols they formulate and require their trader employees to follow. Moreover, Online Appendix B shows that some of our main results in fact *do not depend* on the client's ability to commit.

The client-dealer interactions that we model often occur repeatedly in reality. (Indeed, as just discussed, the standard commitment assumption might be motivated by repeated interaction.) Although this paper focuses on a one-shot version of this interaction, dynamics could introduce several forces. For example, the client could threaten to punish a dealer who was detected to have front-run in the past by avoiding him in the future. Such threats could deter front-running, making the client better off.[13] In the appendix we solve the model under the assumption that the losing dealer does not front-run, which we show does indeed make the client better off. However, this also leads the client to behave differently in important ways: she ceases to restrict her search, and she discloses weakly more information. Such behavior would be inconsistent with various stylized facts about counterparty search: clients often restrict their searches, and they exhibit a strict preference for secrecy. We conclude from this analysis that, in practice, front-running is not eliminated by

[11] Our analysis can therefore be thought of as characterizing one extreme point on the frontier of the tradeoff between execution probability and expected procurement cost. Moreover, this is the point on the frontier that the client would prefer if she had lexicographic preferences (or, equivalently, infinite disutility in the event that execution fails). See Online Appendix G for a further discussion of this frontier.

[12] For models of endogenous choice between trading on the market and trading with a dealer, see, e.g., Seppi (1990); Dugast, Üslü and Weill (2022); Lee and Wang (2022).

[13] On the other hand, dynamics could also allow the dealers to soften their competition through tacit collusion, making the client worse off. A natural way of dealing with this equilibrium multiplicity is to focus, as we do, on equilibria of the one-shot version of the game.

implicit threats to punish it, perhaps because of various barriers to carrying out such threats.[14]

We assume the client contacts dealers simultaneously. Although counterparty search was once done primarily over the phone (and hence sequentially), it is today more commonly done via email, Bloomberg chat, or through RFQ platforms, where search is more often simultaneous.[15] We also assume the contacted dealers know how many other dealers have been contacted, which is consistent with the operation of many RFQ platforms.[16] Thus, in our model, dealers are uncertain about neither the number of dealers that have been contacted nor the sequence in which they have been contacted. Instead, dealers' uncertainty concerns the direction of the client's trade. Though our model contains no asymmetric information about the security's fundamental value, uncertainty about the client's trading direction is relevant because trades have price impact.

We assume the client uses a second-price auction with reserve. Identical results would be obtained if the auction mechanism were a first-price auction with reserve, essentially by the revenue equivalence theorem (Myerson, 1981). However, as in Bertrand competition, we would then need to either (*i*) specify an appropriate tie-breaking rule, or (*ii*) use a model with discrete bidding increments. A second-price auction circumvents these issues. Nevertheless, this assumption is not completely without loss: given that the dealers' initial inventories may be correlated, more complex auction formats permit full surplus extraction (e.g., Crémer and McLean, 1988). Our analysis rules out such possibilities because we think standard auction formats are more realistic.

We also treat the client's order as an indivisible unit. Partly, this is for simplicity: it limits the potential outcomes of the auction. Partly, this is to reflect reality: our discussions with industry participants reveal that splitting an order among competing dealers would typically violate industry norms. Moreover, treating the order as indivisible might in fact be optimal: share auctions often produce worse outcomes for the auctioneer than corresponding unit auctions, due to the classic demand/supply reduction problem (Wilson, 1979; Viswanathan and Wang, 2004).[17]

In assuming that the client's desired trade is $s \in \{-\bar{s}, \bar{s}\}$, we permit uncertainty about only the direction of her trade—not about its size. This assumption reflects the operation of many RFQ platforms, where a trader can request two-sided quotes after specifying a trade size. Even for off-platform trading, clients commonly disclose their trade size when requesting quotes via phone, email, or Bloomberg chat. Nevertheless, in Online Appendix E we consider a version of the model with uncertainty about trade size (but no longer uncertainty about trade side) and obtain

[14] Barriers include: (*i*) data on dealer trades and/or market prices might be unavailable; (*ii*) even if available, such data would permit only statistical detection of front-running; and (*iii*) clients have only a limited number of relationships, so an implicit threat to destroy one might lack credibility.

[15] For example, at least 40% of customer-sell-to-dealer trades in U.S. collateralized loan obligations are conducted via an auction-like mechanism in which bids are simultaneously solicited from dealers via email (Hendershott, Li, Livdan and Schürhoff, 2022).

[16] For example, this assumption is consistent with CanDeal (Allen and Wittwer, forthcoming), as well as the Bloomberg and Tradeweb swap execution facilities (Riggs, Onur, Reiffen and Zhu, 2020). We also assume that all contacted dealers observe the same signal, which likewise reflects the operation of RFQ platforms.

[17] Share auctions are not the only way to permit divisibility. For example, the client could permit each dealer to win exactly half of her total order. Online Appendix D shows this cannot be optimal for the client, essentially because a dealer who receives part of the order fails to internalize the externalities its trading creates for dealers who receive other parts.

qualitatively similar results.

Modeling the dealers. It is primarily for tractability that we assume there are only two dealers. Nevertheless, in Online Appendix F we consider an extension of the model that permits an arbitrary number of dealers, and we find that our main results generalize. In particular, the client might optimally contact fewer than all available dealers—not only when two dealers are available (as in the baseline analysis) but also for any number of available dealers.

We assume that dealers have position limits of $\bar{e}$ shares, in that each dealer's final inventory (after both periods of on-market trading and after any trade with the client) must not exceed $\bar{e}$ shares, either long or short. This position limit can be interpreted as a risk limit or a capital constraint. We also assume that dealers are risk neutral (and experience no holding costs) for inventories within the position limits. Because risk aversion would create a separate trading motive for the dealers, this is to isolate trading dynamics that stem from the client's presence.[18]

The primary interpretation of a dealer's inventory in our model is his proprietary position, which could originate from speculation or from a previous (and unmodeled) interaction with a different client. Alternatively, his inventory could represent the trading desires of his other (unmodeled) clients. Indeed, if these other clients are inclined to buy (sell), then that would, all else equal, induce the dealer to bid more aggressively for a selling (buying) order from the focal client—just as a long (short) propriety position would.[19]

If not contacted, dealer i is constrained to set $x_1^i = 0$. Without this constraint, such a dealer might indeed want to trade in the first period—but only as an artifact of how we have structured the model. To explain, a more realistic (albeit unnecessarily complicated) model would feature many trading days, each with two trading periods (representing, perhaps, the morning and the afternoon), where on only some days would a client arrive seeking to trade. As the probability of a client arriving on any given day diminishes, so too would the first-period trading motive of a dealer who was not contacted that day. On the other hand, simply being contacted tells a dealer *that* a client has arrived, which may be informative in such a model.[20] To simplify, we instead model a single client arrival at a fixed time. Then, to align equilibrium behavior with what would be expected in a more realistic model such as the one described above, that a dealer trades in the first period only if contacted is something that must be imposed directly.

We assume that each dealer knows the other's realized initial inventory. This assumption of perfect knowledge—although extreme—captures in a stylized way that clients are likely to have worse

[18] Another consideration is that it would be efficient for the client to insure a risk-averse dealer against price shocks. For an investigation of optimal contracts with a risk averse dealer (in a model with price shocks), see Baldauf, Frei and Mollner (2023, 2022).

[19] A *riskless principal trade* is said to occur when a dealer conducts opposing trades with two separate clients.

[20] This elaboration of the model highlights a potential puzzle: why does a client not *always* request quotes from dealers, even on days that she is not seeking to trade? With such a signal-jamming strategy, being contacted no longer reveals to a dealer that the client is seeking to trade, potentially ameliorating information leakage. The answer may be that such behavior—making RFQs that are not in good faith—violates industry norms. Indeed, Riggs, Onur, Reiffen and Zhu (2020) find that, in their dataset of RFQs made on swap execution facilities, 92% result in trades, which is inconsistent with significant use of such signal-jamming strategies in practice.

information about dealers than dealers have about each other. Owing to frequent interactions, dealers might in practice be quite familiar with their competitors. They might also learn about competitors' positions by, for example, inspecting published indications of interest, by inspecting trade data, or via communication (e.g., informal backroom chats or so-called "talking your book"). This assumption is also useful for tractability, as without such perfect knowledge, a dealer's uncertainty would be two-dimensional (about both the client's trading direction and his competitor's inventory), leading to a complex Bayesian updating problem.

A natural case of interest is that in which the client has symmetric buying and selling needs (i.e., $\phi_0 = \frac{1}{2}$) and in which the dealers have equal probabilities of being long and short (i.e., $\psi = \frac{1}{2}$). There is, however, value in allowing for asymmetric parametrizations. For example, settings where dealers are more likely to be long than short are captured by $\psi > \frac{1}{2}$. Likewise, settings where the client is more likely to be buying than selling (e.g., because the client is a pension fund with young members) are captured by $\phi_0 > \frac{1}{2}$.

Dealers in practice do not always respond to RFQs. In their dataset, Riggs, Onur, Reiffen and Zhu (2020) find an overall response rate of nearly 90%. In our model, bids that do not meet the client's reserve price can be interpreted as non-responses.

Modeling the market. We assume an exogenous price process, with θ as the coefficient of permanent price impact. This process resembles, for example, the baseline model of Bertsimas and Lo (1998). Moreover, it could be micro-founded by adding a competitive fringe of long-term investors to the model. Specifically, assume that at each trading period the aggregate demand of these investors is $Y(p) = \frac{1}{\theta}(p_0 - p)$. Such a downward-sloping demand could stem from several sources, including risk aversion, institutional frictions, or concerns about adverse selection (although, to be clear, our model does not feature asymmetric information about fundamentals). Note also that this aggregate demand depends only on the current price p, so that these traders do not attempt to profit from short-term price swings. Brunnermeier and Pedersen (2005) make the same set of assumptions, which they argue could be appropriate if these long-term traders lacked the information, skills, or time necessary for predicting price changes.

The assumed price process is a deterministic function of the previous period's price and the current period's trades. This is mainly for tractability: adding price shocks would complicate the belief updating of losing dealers.

In modeling price impact, we aim to capture transitory price pressures away from efficient prices, which revert over the long run and do not affect the security's fundamental value. For this reason, we assume that dealer positions are liquidated at p_0, the initial price, rather than at p_2, the price prevailing in the last trading period.

2.3 Parametric assumptions

Henceforth, we set—without loss of generality—the initial price level to $p_0 = 0$, the coefficient of permanent price impact to $\theta = 1$, and dealers' position limits to $\bar{e} = 1$.

We also assume the client's trade size is $\bar{s} \leq 2$. This is not without loss of generality, but yields conclusions qualitatively similar to what obtains with larger $\bar{s}$.[21]

3 Equilibrium

We begin by describing our solution concept. With that in hand, we solve backward, first deriving outcomes following any RFQ policy, then deriving an optimal policy.

3.1 Solution concept

The solution concept is a refinement of weak perfect Bayesian equilibrium (WPBE). To begin, observe that a WPBE consists of the following elements: (*i*) an RFQ policy for the client, (*ii*) actions for the dealers (i.e., bids and trades), and (*iii*) beliefs for the dealers.

Bids. Three aspects of the refinement concern bids. In a classic second-price auction with private values, each bidder has two payoff-relevant outcomes (winning and losing) and a unique dominant strategy—bidding the difference in its values across those outcomes. Our setting, however, comprises not only an auction but also a continuation trading game that affects the relative values of winning and losing in the auction. Thinking of the entire game, the set of undominated strategies is quite large and in particular does little to pin down bids. Nevertheless, the auction can be analyzed in isolation if the equilibrium of the continuation trading game is computed and used to fix the dealers' values for the auction (as by backward induction). Conceiving of the auction in isolation in this way, a dealer has a unique dominant bid if he is the only one contacted: to bid the difference in his values between winning and losing. Our refinement requires this bid to be played.

The situation is somewhat more complicated when both dealers are contacted, because the auction is then one with externalities. That is, each dealer has *three* payoff-relevant outcomes: either he wins, his competitor wins, or neither wins. Even if values in the auction are fixed via backward induction, there is in general no dominant strategy. But if such an auction's reserve price were non-binding (so that the outcome in which neither wins does not occur in equilibrium), then there is an analogue: each dealer should bid the difference in its values between the outcomes in which he wins and in which his competitor wins. Our refinement requires these bids to be played when the reserves are non-binding.

We also require symmetric bidding strategies.

Beliefs. The remaining aspects of the refinement concern dealer beliefs. Although the winning dealer observes whether the client seeks to buy or sell, the other dealer does not observe this directly and must instead infer from what he does observe. Our refinement applies to how this other dealer infers from the winning dealer's first-period trades.

[21] Larger values of $\bar{s}$ would need to be handled as mathematically separate cases because additional constraints come into play. For example, with $\bar{s} > 2$, the winning dealer can never fully internalize the client's order, and so must trade at least some amount on the market regardless of his initial position.

We focus on equilibria in which these beliefs have a threshold structure. To state this formally, suppose (without loss of generality) that dealer A is the winner, and let $\mu_2^B(x_1^A)$ denote the probability that $s = \bar{s}$ under B's beliefs as a function of A's first-period trade. We require this to be a function that jumps from zero to one: (i) for all $x \in [-\bar{s}, \bar{s}]$, $\mu_2^B(x) \in \{0, 1\}$, and (ii) for all $x' < x''$, $\mu_2^B(x') = 1 \implies \mu_2^B(x'') = 1$ and $\mu_2^B(x'') = 0 \implies \mu_2^B(x') = 0$. One implication of this requirement is that we focus on separating equilibria in which the winning dealer's first-period trade fully reveals the realized s to the other dealer.

We also require these beliefs to satisfy an analogue of the intuitive criterion (Cho and Kreps, 1987). To explain, again suppose (without loss of generality) that dealer A is the winner. Dealer A's objective in the continuation trading game is to minimize his trading costs $p_1 x_1^A + p_2 x_2^A$, subject to the constraints on his final inventory. Given knowledge of A's objective, what can dealer B believe about s following an out-of-equilibrium choice for x_1^A? Fix an equilibrium. For $s' \in \{-\bar{s}, \bar{s}\}$, let $C_*^A(s')$ represent A's equilibrium trading costs given s', and let $C^A(s', x_1^A, \mu_2^B)$ represent A's trading costs from x_1^A and the equilibrium x_1^B, together with the second-period trades in the equilibrium of the game continuing from (x_1^A, x_1^B) given s' and given that B believes $s = \bar{s}$ with probability μ_2^B.[22] We require that for all out-of-equilibrium x_1^A, neither of the following pairs of conditions holds:

$$C_*^A(\bar{s}) < \min_{\mu_2^B \in [0,1]} C^A(\bar{s}, x_1^A, \mu_2^B) \quad \text{and} \quad C_*^A(-\bar{s}) > C^A(-\bar{s}, x_1^A, 0)$$

$$C_*^A(-\bar{s}) < \min_{\mu_2^B \in [0,1]} C^A(-\bar{s}, x_1^A, \mu_2^B) \quad \text{and} \quad C_*^A(\bar{s}) > C^A(\bar{s}, x_1^A, 1)$$

After stating Lemma 1, we provide an example to illustrate this criterion and to review the motivation that Cho and Kreps (1987) provide for it.

Henceforth, we use *equilibrium* to refer to this refinement of WPBE.

3.2 Contacting one dealer

We begin by analyzing continuation equilibrium in subgames following RFQs that contact one dealer, who without loss of generality we assume to be dealer A.

Recall that an RFQ policy is defined by a signal realization space Σ and distributions $(\pi_{s'})_{s' \in \{-\bar{s}, \bar{s}\}}$ over $\Sigma \times \{1, 2\} \times \mathbb{R}^2$. Fix such an RFQ policy. Fix, moreover, some RFQ $(\sigma, M, \bar{b})$ with $M = 1$. It suffices to focus on RFQs that occur on path, so we assume $\pi_{-\bar{s}}(\sigma, M, \bar{b}) + \pi_{\bar{s}}(\sigma, M, \bar{b}) > 0$. We also define

$$\phi = \frac{\phi_0 \pi_{\bar{s}}(\sigma, M, \bar{b})}{(1 - \phi_0)\pi_{-\bar{s}}(\sigma, M, \bar{b}) + \phi_0 \pi_{\bar{s}}(\sigma, M, \bar{b})}$$

as the posterior probability of $s = \bar{s}$ induced by this RFQ. Note that $\phi \in [0, 1]$ and that it may differ from the prior ϕ_0.

[22]The main subtlety is in determining $C^A(s, x_1^A, \mu_2^B)$. Technically, the equilibrium that we have fixed does not specify what second-period trades would be under out-of-equilibrium beliefs μ_2^B. Nevertheless, the structure of the model allows these to be uniquely computed.

Lemma 1. *In the class of RFQs that contact $M = 1$ dealer and lead to execution with probability one, the minimum expected procurement cost is $\hat{c}_1 = \frac{3\bar{s}^2}{4}$. It is achieved by RFQs featuring reserve prices $\bar{b} = (\frac{3\bar{s}^2}{4}, \frac{3\bar{s}^2}{4})$.*

To prove Lemma 1, we begin by deriving the unique equilibrium actions in the subgame following any RFQ that contacts one dealer. After sketching that derivation, we explain how the lemma's claims follow. Formal proofs are in Online Appendix A.

Continuation equilibrium. Because both dealers observe the entire vector (e^A, e^B), the four possible realizations of that vector can be analyzed separately. Let us focus on the case of $(e^A, e^B) = (1, -1)$. Suppose that dealer B's beliefs in this case are that $s = \bar{s}$ with probability

$$\mu_2^B = \begin{cases} 1 & \text{if } -\frac{\bar{s}}{6} \leq x_1^A \leq \bar{s}, \\ 0 & \text{if } -\bar{s} \leq x_1^A < -\frac{\bar{s}}{6}. \end{cases} \tag{1}$$

In words, if dealer A's first-period trade is above (below) the threshold $-\frac{\bar{s}}{6}$, then dealer B believes with certainty that the client is a buyer (seller). For an informal derivation of the unique on-path actions, consider two subcases:

- First, suppose the client's order is to buy (i.e., $s = \bar{s}$). Dealer A, being initially long (i.e., $e^A = 1$), can simply internalize the order. As a result, no on-market trading happens in equilibrium. Mathematically, if we use $p_1 = x_1^A$ and $p_2 = x_1^A + x_2^A + x_2^B$, and if we ignore constraints on final inventory (which are slack in the equilibrium), then dealer A minimizes his trading cost $x_1^A x_1^A + (x_1^A + x_2^A + x_2^B)x_2^A$, and dealer B minimizes $(x_1^A + x_2^A + x_2^B)x_2^B$, leading to $x_2^A = x_2^B = -\frac{x_1^A}{3}$. Inducting backward, we obtain $x_1^A = 0$, so that $x_2^A = x_2^B = 0$ on path. Panel (d) of Figure 1 illustrates these equilibrium trades.

  Plugging in these trades, dealer A incurs no trading costs if he wins, so that bidding $b_{\bar{s}}^A = 0$ is optimal—and in fact required by the refinement described in Section 3.1.

- Second, suppose the client's order is instead to sell (i.e., $s = -\bar{s}$). Dealer A now cannot internalize and must sell on the market. In equilibrium, dealer B provides liquidity to dealer A by buying on the market. Mathematically, if we assume that $x_2^A = -\bar{s} - x_1^A$ (which ensures that dealer A's final inventory just meets the constraint $e^A + x_1^A + x_2^A - s \leq 1$) and if we ignore all other constraints on final inventory, then dealer A minimizes $x_1^A x_1^A + (-\bar{s} + x_2^B)(-\bar{s} - x_1^A)$, and dealer B minimizes $(-\bar{s} + x_2^B)x_2^B$, leading to $x_2^B = \frac{\bar{s}}{2}$. Inducting backward, we obtain $x_1^A = -\frac{\bar{s}}{4}$. Panel (c) of Figure 1 illustrates these equilibrium trades.

  Plugging in these trades, dealer A incurs trading costs of $\frac{7\bar{s}^2}{16}$ if he wins, implying a bid of $b_{-\bar{s}}^A = \frac{7\bar{s}^2}{16}$.

Note that the beliefs (1) comport with the above derivation of equilibrium actions. Implicit in the derivation is that dealer B can condition x_2^B on the realized s; although he cannot do so directly, he can do so *indirectly* because x_1^A reveals s under these beliefs. Moreover, these beliefs are indeed consistent with dealer A's equilibrium first-period trades x_1^A whenever Bayes' rule can be applied.

Alternatively, if $(e^A, e^B) = (1, 1)$, then equilibrium trades are as illustrated in Panels (a) and (b) of Figure 1. The remaining inventory configurations (in which dealer A is initially short) are handled symmetrically.

Figure 1: Trading dynamics when contacting one dealer

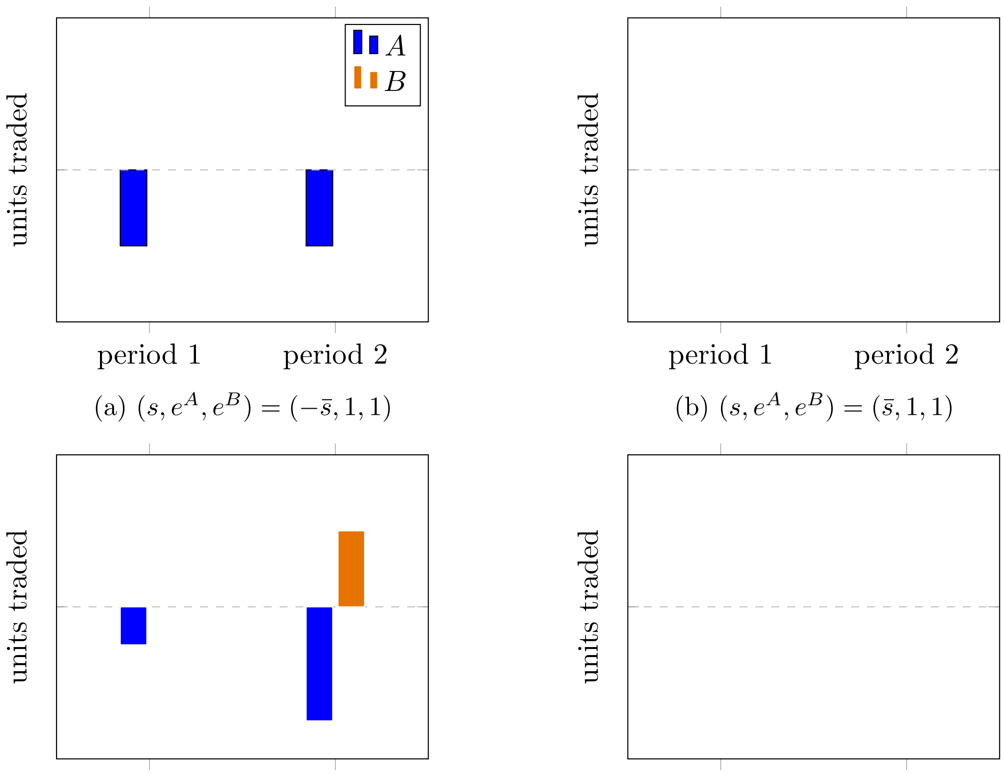

(a) $(s, e^A, e^B) = (-\bar{s}, 1, 1)$

(b) $(s, e^A, e^B) = (\bar{s}, 1, 1)$

(c) $(s, e^A, e^B) = (-\bar{s}, 1, -1)$

(d) $(s, e^A, e^B) = (\bar{s}, 1, -1)$

Suppose dealer A is the only one contacted, and suppose he wins the client's order. The four panels depict trades on the continuation equilibrium path for four cases of the client's trading need s and the dealer inventories (e^A, e^B). The remaining four cases are symmetric. Positive (negative) bars represent orders to buy (sell). No trading occurs in panels (b) and (d). Vertical axes suppress units because all quantities scale linearly with $\bar{s}$. Mathematical expressions are provided by Lemma A1.

The client's cost of procurement. The client's procurement cost is determined by a second-price procurement auction. In this case where only one dealer is contacted, the reserve sets the price, with execution failing to occur when this reserve is below the dealer's bid. To obtain execution with probability one, the client's reserve must be at least the dealer's bid in the worst case. If the client sells, the worst case is $(e^A, e^B) = (1, 1)$. As we specify in the proof of Lemma 1, $b^A_{-\bar{s}} = \frac{3\bar{s}^2}{4}$ in this case. (By symmetry, this is also the worst-case value for $b^A_{\bar{s}}$.) Therefore, the expected procurement cost resulting from an RFQ that contacts one dealer, induces a belief that ϕ is the probability of $s = \bar{s}$, and leads to execution with probability one is at least

$$\hat{c}_1 = \phi\frac{3\bar{s}^2}{4} + (1-\phi)\frac{3\bar{s}^2}{4} = \frac{3\bar{s}^2}{4},$$

as claimed by Lemma 1. Moreover, an RFQ with reserve prices $\bar{b} = (\frac{3\bar{s}^2}{4}, \frac{3\bar{s}^2}{4})$ achieves this bound.

Note that $\hat{c}_1$ does not depend on the probability of $s = \bar{s}$ induced by the RFQ (which we have labelled ϕ). This is for two reasons. First, the reserve prices specified by Lemma 1 are constant in ϕ. Intuitively, this is because there is no role for information design when $M = 1$: the contacted dealer submits a separate quote for each state $s \in \{-\bar{s}, \bar{s}\}$, and he will moreover learn the state before he trades. Second, these reserve prices are equal, so that although $\hat{c}_1$ can be thought of as a weighted average of the two (with weight ϕ on $\bar{b}_{\bar{s}}$), the weight does not matter. Intuitively, this is due to the model's symmetric structure. As we will see, this changes when two dealers are contacted: the analogous quantity $\hat{c}_2$ will be a non-constant (and in fact non-linear) function of ϕ, which creates a role for information design.

Likewise, $\hat{c}_1$ also does not depend on ρ and ψ, which govern the distribution of (e^A, e^B). This is because the client's cost is set by her reserve prices, and because the reserve prices specified by Lemma 1 are driven by the worst case. Although (ρ, ψ) affect the distribution over the various cases for (e^A, e^B), they do not affect equilibrium outcomes in those cases—and in particular do not affect the worst case. As we will see, this also changes when two dealers are contacted: the client's cost may then be set by the losing dealer's bid, and thus no longer driven solely by the worst case.

The intuitive criterion. Finally, we provide an example to illustrate what role the intuitive criterion (Cho and Kreps, 1987) plays in our equilibrium selection. As above, focus on the case of $(e^A, e^B) = (1, -1)$. There is another WPBE in which dealer B's beliefs entail a threshold lower than that in (1)

$$\mu_2^B = \begin{cases} 1 & \text{if } -\frac{\bar{s}}{2} < x_1^A \leq \bar{s}, \\ 0 & \text{if } -\bar{s} \leq x_1^A \leq -\frac{\bar{s}}{2}, \end{cases} \tag{2}$$

and in which dealer A sells more in the first period when $s = -\bar{s}$, in particular, selling just enough to meet the threshold: $x_1^A = -\frac{\bar{s}}{2}$. The intuition is that when $s = -\bar{s}$, A's inventory constraint binds, and he must sell on the market. If B believes that $s = -\bar{s}$, then in the second trading period, he provides liquidity to A by buying on the market. And because A would like B to provide this liquidity to him, $x_1^A = -\frac{\bar{s}}{2}$ is set at the threshold required for B to believe $s = -\bar{s}$. If a less extreme choice for x_1^A (e.g., $-\frac{\bar{s}}{4}$, as in Lemma 1) would induce B to provide liquidity in the second period, then A would prefer that. Intuitively, this outcome would allow A to sell only a minority of the necessary amount in the first period when B is not providing liquidity, while waiting to sell the majority until the second period when B is providing liquidity. This outcome is, however, incompatible with the beliefs (2). Under (2), A must instead settle for selling the smallest amount in the first period that does induce B to provide liquidity (i.e., $x_1^A = -\frac{\bar{s}}{2}$).

This alternative WPBE fails the version of the intuitive criterion described in Section 3.1. Indeed, in terms of the notation introduced there, A's equilibrium trading costs are $C_*^A(\bar{s}) = 0$ and $C_*^A(-\bar{s}) = \frac{\bar{s}^2}{2}$. Then consider the out-of-equilibrium choice $x_1^A = -\frac{\bar{s}}{4}$. We can compute $\min_{\tilde{\mu}_2^B \in [0,1]} C^A(\bar{s}, -\frac{\bar{s}}{4}, \tilde{\mu}_2^B) = C^A(\bar{s}, -\frac{\bar{s}}{4}, 0) = \frac{3\bar{s}^2}{64} > C_*^A(\bar{s}) = 0$ and $C^A(-\bar{s}, -\frac{\bar{s}}{4}, 0) = \frac{7\bar{s}^2}{16} < \frac{\bar{s}^2}{2}$. Intuitively, if $s = \bar{s}$, then the best that A could expect to achieve from a choice of $x_1^A = -\frac{\bar{s}}{4}$ is a trading cost of $\frac{3\bar{s}^2}{64}$ (achieved if B believes $s = \bar{s}$ with probability zero), which is strictly greater than his equilibrium trading cost of zero.[23] Thus, it would be difficult to make sense of this deviation if $s = \bar{s}$. Could the deviation make sense if $s = -\bar{s}$? Suppose $s = -\bar{s}$ and A deviates to $x_1^A = -\frac{\bar{s}}{4}$, hoping that—since it would not make sense if $s = \bar{s}$—the deviation will induce B to believe $s = -\bar{s}$. If the deviation does indeed induce B to believe this, then the deviation would make sense: A's trading cost would be $\frac{7\bar{s}^2}{16}$ (as in Lemma 1's equilibrium), which is strictly less than his trading cost in this equilibrium of $\frac{\bar{s}^2}{2}$. Thus, we have an argument that—contrary to the beliefs (2)—dealer B should infer from a deviation to $x_1^A = -\frac{\bar{s}}{4}$ that $s = -\bar{s}$.

3.3 Contacting two dealers

We proceed analogously to analyze continuation equilibrium in subgames following RFQs that contact two dealers. Fix an RFQ policy. Fix also some RFQ $(\sigma, M, \bar{b})$ with $M = 2$. As before, assume without loss of generality that $\pi_{-\bar{s}}(\sigma, M, \bar{b}) + \pi_{\bar{s}}(\sigma, M, \bar{b}) > 0$, and define

$$\phi = \frac{\phi_0 \pi_{\bar{s}}(\sigma, M, \bar{b})}{(1-\phi_0)\pi_{-\bar{s}}(\sigma, M, \bar{b}) + \phi_0 \pi_{\bar{s}}(\sigma, M, \bar{b})}$$

as the posterior probability of $s = \bar{s}$ induced by this RFQ.

Lemma 2. *In the class of RFQs that contact $M = 2$ dealers, induce beliefs ϕ, and lead to execution with probability one, the minimum expected procurement cost is a strictly convex and differentiable function $\hat{c}_2(\phi)$ that satisfies $\hat{c}_2(\frac{1}{2}) < \hat{c}_1$. It is achieved by RFQs featuring reserve prices $\bar{b} = (\bar{s}^2, \bar{s}^2)$.*

The client's procurement cost is determined by a second-price procurement auction. In this case where two dealers are contacted, it could in principle be set by either the reserve price or the losing dealer's bid. What would transpire if the reserves were low enough to occasionally set the price? Given our restriction to symmetric bidding strategies, it can be shown that the continuation equilibrium would entail mixed bidding strategies, in which with positive probability, neither dealer would meet the reserve. Thus, to obtain execution with probability one, the reserves must be high enough to never set the price.[24] To prove Lemma 2, we therefore begin by analyzing equilibria of subgames following RFQs that contact two dealers and in which the reserves are sufficiently high in this sense. After sketching that derivation, we explain how the lemma's claim follows.

[23]When B believes that $s = \bar{s}$ with probability $\tilde{\mu}_2^B$, he chooses $x_2^B = \frac{(8-7\tilde{\mu}_2^B)\bar{s}}{4(4-\tilde{\mu}_2^B)}$. If we actually have $s = \bar{s}$ and $x_1^A = -\frac{\bar{s}}{4}$, then A best responds with $x_2^A = \frac{(3\tilde{\mu}_2^B - 2)\bar{s}}{4(4-\tilde{\mu}_2^B)}$, leading to a total trading cost of $\frac{(3-2\tilde{\mu}_2^B)(1+\tilde{\mu}_2^B)\bar{s}^2}{4(4-\tilde{\mu}_2^B)}$. This cost is minimized at $\tilde{\mu}_2^B = 0$, where it evaluates to $\frac{3\bar{s}^2}{64}$.

[24]Equivalently, the client could run an auction without reserves. Reserve prices play a role only when the client contacts just a single dealer (*cf.* Section 3.2).

Continuation equilibrium. In the previous section, only one dealer was contacted, and the other dealer was unable to trade in the first period. Here, both dealers are contacted, and not only the winner but also the loser may trade in the first period. In some cases the losing dealer uses this ability to front-run the winner's trades. To see this front-running, focus on the case of $(e^A, e^B) = (1, 1)$, where both dealers begin long. The losing dealer does not know the direction of the client's trade, but does know that the winning dealer, being initially long, will have to sell on the market if the client's order is to sell. To exploit this, the losing dealer sells in the first period (as panels (a) and (b) of Figure 2 depict), which affords him flexibility to buy back in the second period if the client's order is indeed to sell. In this case, he will have sold high and bought low, netting a profit. Hence, the amount that he sells in the first period depends on his beliefs. If $\phi = 1$, so that he is sure the client is buying, then there is no scope to front-run, and he in fact does not trade. But as ϕ decreases, he sells progressively more in the first period (ultimately selling $\frac{\bar{s}}{3}$ if $\phi = 0$).

For the case of $(e^A, e^B) = (1, -1)$, if dealer B loses the auction, he infers that the client's order is to buy. (This belief will be correct in equilibrium, given equilibrium bids that we subsequently derive.) If the client's order is indeed to buy, then A can internalize and need not trade on the market. Anticipating that A will not trade, B does not trade either (as panels (c) and (d) depict). Of course, if the client's order is instead to sell, then A would need to sell on the market (as panel (c) depicts), but this scenario is off-path, so B assigns no weight to it in deciding his first-period trade.

To explain how bids are pinned down, focus on a particular realization for (e^A, e^B) and for s. Having characterized trading behavior, we can compute dealer A's trading costs in each of the three potential auction outcomes: (*i*) he wins, (*ii*) B wins, and (*iii*) neither wins. If, however, the reserve price $\bar{b}_s$ is sufficiently high, then outcome (*iii*) becomes irrelevant. Given that, the refinement described in Section 3.1 requires A's bid to be determined by the difference in his value across outcomes (*i*) and (*ii*). Having derived this bid, we can then fill in what it means for the reserve price to be sufficiently high in this sense: $\bar{b}_s$ must be at least this difference. Choosing $\bar{b} = (\bar{s}^2, \bar{s}^2)$, as in Lemma 2, satisfies all such constraints.

The client's cost of procurement. To compute $\hat{c}_2$, the expected procurement cost under $\bar{b} = (\bar{s}^2, \bar{s}^2)$—or any other reserves high enough that they never set the price—we simply calculate a weighted average of the losing bids. The weights are dictated by the parameters ψ and ρ, which govern the distribution of (e^A, e^B), as well as ϕ, which captures the distribution of s conditional on the RFQ. In addition to affecting these weights, ϕ also affects the bids themselves because the losing dealer, unable to observe the realized s, relies on these beliefs to select his first-period trade.

The comparison between $\hat{c}_2(\phi)$ and $\hat{c}_1$ is driven by the following tradeoff. On the one hand, contacting an additional dealer may intensify competition among the dealers for the client's business (the *competition effect*). The additional dealer might also be able to internalize the client's order and thereby provide fulfillment more efficiently (the *sampling effect*). On the other hand, dealers

Figure 2: Trading dynamics when contacting two dealers

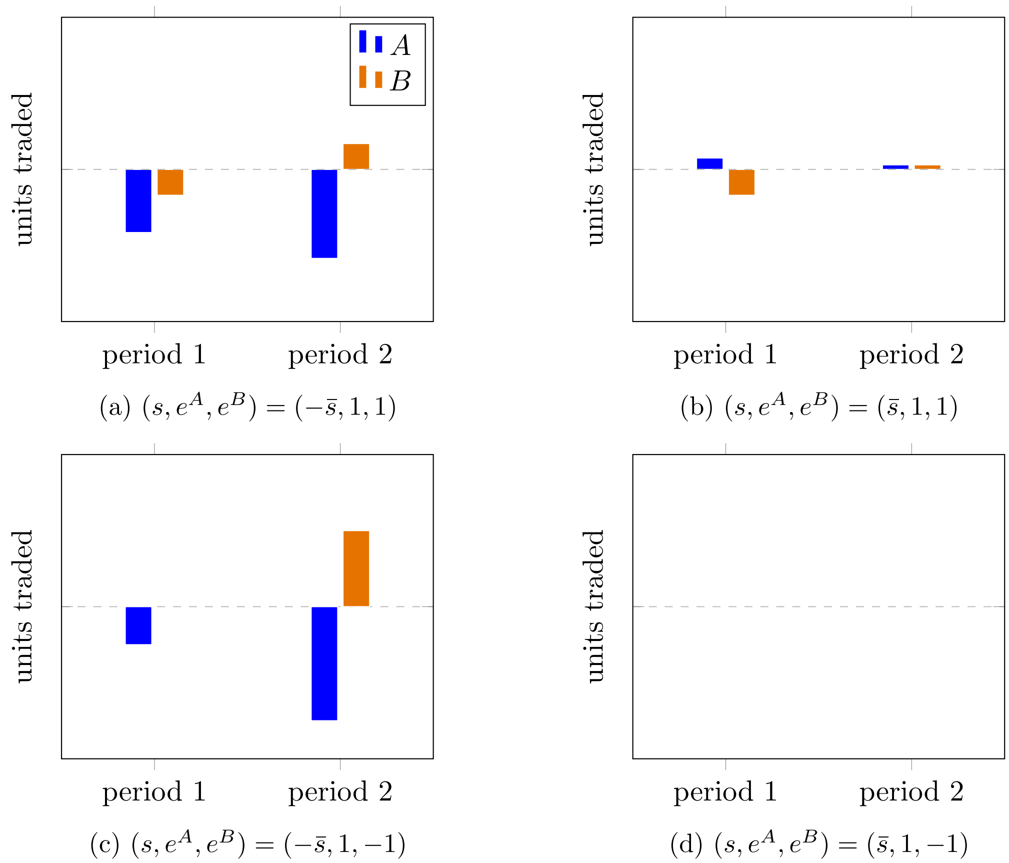

Suppose both dealers are contacted, suppose dealer A wins the client's order, and suppose the client is believed to be a buyer with probability $\phi = 1/2$. The four panels depict trades on the continuation equilibrium path for four cases of the client's trading need s and the dealer inventories (e^A, e^B). The remaining four cases are symmetric. Likewise, symmetric analysis covers the case in which dealer B wins. Positive (negative) bars represent orders to buy (sell). No trading occurs in panel (d). Vertical axes suppress units because all quantities scale linearly with $\bar{s}$. Mathematical expressions are provided by Lemma A2.

who are contacted but not selected might front-run on the market (the *front-running effect*).[25] (See Online Appendix C for a formalization of these three effects.)

If the front-running effect dominates, then $\hat{c}_2(\phi) > \hat{c}_1$. For example, consider the case of $(e^A, e^B) = (1, 1)$, in which both dealers begin long. We show in the proof of Lemma 2 that if $\phi = 0$, so that the dealers are certain the client is a seller, then each dealer i bids $b^i_{-\bar{s}} = \bar{s}^2$ when both are contacted. In contrast, as previously mentioned, $\frac{3\bar{s}^2}{4}$ is the analogous bid when only one dealer is contacted. The bid is larger here for two reasons. First, the losing dealer's front-running raises the winning dealer's trading costs.[26] Second, winning now entails an additional opportunity

[25] Thus, information leakage is costly in our setting because it informs losing dealers *that* the client is seeking to trade, which leads to front-running. This distinguishes our paper from others in which the information that is leaked pertains to fundamental information (e.g., Brunnermeier, 2005; Liu, Vogel and Zhang, 2018; Kondor and Pintér, 2022).

[26] This is because price impact from the losing dealer's first-period trade moves p_1 against the winning dealer.

cost—namely, the potential for profitable front-running that would exist if the other dealer were to win. Both effects lead to less aggressive bidding. It follows that if the dealers are likely to begin long (i.e., $\psi \approx 1$), then $\hat{c}_2(0) > \hat{c}_1$, meaning that for RFQs revealing the client to be a seller, she is better off contacting only a single dealer.

If the competition and sampling effects dominate, then $\hat{c}_2(\phi) < \hat{c}_1$. For example, Lemma 2 says that $\hat{c}_2(\frac{1}{2}) < \hat{c}_1$ (regardless of ψ and ρ). This means that for RFQs revealing the client to be equally likely a seller or a buyer, she is better off contacting both dealers.

3.4 Optimal RFQ policies

The previous sections derived continuation outcomes following any RFQ. We now induct backward so as to obtain the optimal RFQ policy. (Recall that an RFQ policy is optimal if it minimizes expected procurement cost subject to obtaining execution with probability one.) This problem is potentially complex because the set of policies to optimize over is large and rich. How many dealers should the client contact—always one, always two, or should she randomize in some way? What signals should she provide about her desired trade—full disclosure, minimal disclosure, or something intermediate? How should she treat buying and selling—symmetrically or asymmetrically? Perhaps surprisingly, the model yields sharp answers to these questions.

3.4.1 Optimality of minimal disclosure

Depending on the RFQ policy, the number of dealers contacted may be an informative signal about the client's trading direction. We say that an RFQ policy is *minimally disclosing* if it never reveals anything beyond that—neither via the arbitrary signal σ nor via the reserve prices $\bar{b}$. For example, all policies of the following form are minimally disclosing: let $\Sigma = \{\sigma_0\}$, fix $\bar{b}_1, \bar{b}_2 \in \mathbb{R}^2$, and let $(\pi_{s'})_{s' \in \{-\bar{s}, \bar{s}\}}$ put positive weight on at most the two RFQs $(\sigma_0, 1, \bar{b}_1)$ and $\big(\sigma_0, 2, \bar{b}_2\big)$. An implication of the following result is that one such policy is optimal.

Proposition 3. *A minimally-disclosing RFQ policy is optimal. In particular, there exists an optimal RFQ policy with the following properties:* (*i*) Σ *is a singleton* $\{\sigma_0\}$*; and* (*ii*) *the distributions* $(\pi_{s'})_{s' \in \{-\bar{s}, \bar{s}\}}$ *put positive weight on at most the two RFQs* $\big(\sigma_0, 1, (\frac{3\bar{s}^2}{4}, \frac{3\bar{s}^2}{4})\big)$ *and* $\big(\sigma_0, 2, (\bar{s}^2, \bar{s}^2)\big)$.

RFQ policies of the form described in the proposition use reserve prices $\bar{b} = (\frac{3\bar{s}^2}{4}, \frac{3\bar{s}^2}{4})$ when $M = 1$ dealers are contacted and $\bar{b} = (\bar{s}^2, \bar{s}^2)$ when $M = 2$. That this is consistent with optimality follows from Lemmas 1 and 2. So to establish the result, it suffices to show that beginning from an RFQ policy of the form described in the proposition, the client cannot benefit by deviating to a more informative signal structure. Given that $\hat{c}_1$ is a constant, a more informative signal has no effect when only one dealer is contacted. Furthermore, $\hat{c}_2(\phi)$ is convex, so that by standard arguments from Bayesian persuasion (Kamenica and Gentzkow, 2011), the client is worse off under a more informative signal when two dealers are contacted.

(On the other hand, because the model features permanent price impact, the losing dealer's round-trip trade has no influence on p_2.)

To build some intuition for this result, note that the disclosure policy is relevant only through its effect on the information available to the losing dealer in trading period 1 (as the winning dealer always learns s, and the loser always infers it after period 1.) So, the question is how the loser's period-1 trade depends on his information about s. In principle, better information could facilitate either front-running or liquidity provision. But in equilibrium, he uses better information to front-run.[27] Front-running raises both the winner's trading costs and the loser's profits (hence, the winner's opportunity cost). Both effects lead dealers to bid less aggressively, increasing the client's cost of procurement. It follows that the client optimally reveals nothing about s (beyond what may be revealed through the number of dealers she contacts).

Notably, this result is in line with common industry practice, where additional information is rarely volunteered at the RFQ stage. For example, clients typically attempt to disguise their trading direction by requesting two-sided quotes instead of one-sided quotes (e.g., Risk.net, 2018). Our model rationalizes that behavior as optimal.

3.4.2 Optimal policy for the number of dealers to contact

We can therefore restrict attention to RFQ policies of the form described by Proposition 3. For such RFQ policies, and for all $s \in \{-\bar{s}, \bar{s}\}$, let q_s denote the probability with which only one dealer is contacted (so that with complementary probability $1 - q_s$ two dealers are contacted). The problem then reduces to choosing a policy for the number of dealers to contact by optimizing over $(q_{-\bar{s}}, q_{\bar{s}}) \in [0,1]^2$. The policy choice has an effect in two ways: (i) given a fixed belief ϕ, it can affect the client's cost, doing so if $\hat{c}_1 \neq \hat{c}_2(\phi)$, and (ii) it can manipulate the beliefs themselves.

To characterize the optimal policy, we define $C(\phi)$ as the convex closure of $\min\{\hat{c}_1, \hat{c}_2(\phi)\}$:

$$C(\phi) = \inf \Big\{ z \;\Big|\; (\phi, z) \in co\big(\min\{\hat{c}_1, \hat{c}_2\}\big) \Big\},$$

where $co\big(\min\{\hat{c}_1, \hat{c}_2\}\big)$ denotes the convex hull of the graph of $\min\{\hat{c}_1, \hat{c}_2\}$. By construction, C is the largest convex function that is everywhere weakly less than both $\hat{c}_1$ and $\hat{c}_2(\phi)$. As we will see, $C(\phi_0)$ is a lower bound on the client's procurement cost. What is more, this lower bound is achievable. To describe an RFQ policy that achieves this bound, it is useful to define two cutoffs: $\underline{\phi}$ and $\overline{\phi}$, defined precisely to ensure that $C(\phi) = \hat{c}_2(\phi)$ if and only if $\phi \in [\underline{\phi}, \overline{\phi}]$.

Definition 1. Define $\underline{\phi}, \overline{\phi} \in [0,1]$ as follows. If $\hat{c}_2(0) \leq \hat{c}_1$, $\underline{\phi} = 0$; if $\hat{c}_2(1) - \hat{c}_2'(1) \geq \hat{c}_1$, $\underline{\phi} = 1$; otherwise, define it implicitly as the unique $\underline{\phi} \in (0,1)$ that solves $\hat{c}_2(\phi) - \phi \hat{c}_2'(\phi) = \hat{c}_1$. If $\hat{c}_2(1) \leq \hat{c}_1$, $\overline{\phi} = 1$; if $\hat{c}_2'(0) + \hat{c}_2(0) \geq \hat{c}_1$, $\overline{\phi} = 0$; otherwise, define it implicitly as the unique $\overline{\phi} \in (0,1)$ that solves $(1-\phi)\hat{c}_2'(\phi) + \hat{c}_2(\phi) = \hat{c}_1$.[28]

[27] For example, consider the full-disclosure regime. In that case, the losing dealer can condition his first-period trade on s. Applying Lemma A2 in the appendix, he uses this ability to trade with the winner in period 1 (while planning to trade against him in period 2). Although he may provide liquidity on net, his period 1 trades reflect front-running. Garblings of the full-disclosure policy mitigate this front-running by interfering with the loser's ability to condition his trades on s (and hence the direction of the winner's trading).

[28] To see that this provides a unique definition for $\underline{\phi}$, it suffices to show that $\hat{c}_2(\phi) - \phi\hat{c}_2'(\phi)$ is strictly decreasing on

The next result reports some properties of these cutoffs.

Proposition 4. *The cutoffs $\underline{\phi}$ and $\overline{\phi}$ (i) are each weakly increasing in ψ and (ii) satisfy $\underline{\phi} \leq \overline{\phi}$, where the inequality is strict if both $\underline{\phi} > 0$ and $\overline{\phi} < 1$.*

So, as the notation suggests, we indeed have $\underline{\phi} \leq \overline{\phi}$. Moreover, these cutoffs respond monotonically to changes in the parameter ψ. (The same need not be true for ρ.) We can now characterize the optimal RFQ policy.

Proposition 5. *$C(\phi_0)$ is the expected cost of procurement under any optimal RFQ policy. Moreover, one optimal policy is in the class described by Proposition 3, where the probabilities of contacting only a single dealer are*

$$(q_{-\bar{s}}, q_{\bar{s}}) = \begin{cases} \left(\frac{\underline{\phi}-\phi_0}{\underline{\phi}(1-\phi_0)}, 0\right) & \text{if } \phi_0 \in [0, \underline{\phi}), \\ (0,0) & \text{if } \phi_0 \in [\underline{\phi}, \overline{\phi}], \\ \left(0, \frac{\phi_0-\overline{\phi}}{\phi_0(1-\overline{\phi})}\right) & \text{if } \phi_0 \in (\overline{\phi}, 1]. \end{cases}$$

In words, the three cases depend on how the prior ϕ_0 relates to the cutoffs $\underline{\phi}$ and $\overline{\phi}$. According to the proposition, if the client is ex ante sufficiently likely to sell (in the sense that $\phi_0 < \underline{\phi}$), then her optimal policy entails contacting one dealer only if her realized trading need is in fact to sell. It follows that contacting two dealers must induce a belief that she is a buyer with a probability greater than ϕ_0, and in fact this belief is precisely $\underline{\phi}$ under the optimal policy. Everything is reversed if she is instead sufficiently likely to buy (in the sense that $\phi_0 > \overline{\phi}$). Otherwise, her optimal policy is to always contact two dealers.

One implication of this result is that the client might optimally restrict the number of dealers she contacts. This is in line with common industry practice. In swaps trading, for example, Riggs, Onur, Reiffen and Zhu (2020) document that the modal number of dealers contacted is the legal minimum of three (their Fig. 3). Similarly, in Canadian fixed income and derivatives, CanDeal (2020) reports that 51.6% of RFQs contact fewer than four dealers, the maximum on the platform.[29] Our model rationalizes this behavior without exogenous search frictions—as assumed by Duffie, Gârleanu and Pedersen (2005) and much of the follow-on literature. Rather, the reason is that information leakage is an endogenous search friction: contacting a dealer makes him aware that the client is seeking to trade, which may lead to front-running.

the unit interval. Because its derivative is $-\phi\hat{c}_2''(\phi)$, the conclusion follows from the convexity of $\hat{c}_2(\cdot)$. Analogously, $\overline{\phi}$ is well-defined because $\hat{c}_2(\phi) + (1-\phi)\hat{c}_2'(\phi)$ is strictly increasing on the unit interval.

[29] Other papers report that a given client's trades tend to be concentrated within a small subset of dealers. This is explained by relationships playing an important role in many asset classes (Bernhardt, Dvoracek, Hughson and Werner, 2004; Di Maggio, Kermani and Song, 2017; O'Hara, Wang and Zhou, 2018; Hendershott, Li, Livdan and Schürhoff, 2020; Riggs, Onur, Reiffen and Zhu, 2020; Hau, Hoffmann, Langfield and Timmer, 2021; Allen and Wittwer, forthcoming). Being costly to develop, a client will in practice have relationships with only a limited number of dealers. Relationship formation costs include, for example, the need to establish credit or clearing relationships, various administrative considerations, or the promise of future rents necessary to delay holdup (as modeled by Board, 2011), as well as the opportunity cost of not exploiting existing relationships. Yet, these considerations alone cannot explain why a client would not always contact the entire set of dealers with whom she has a relationship. The endogenous front-running, which we model, can explain this.

When the risk of front-running looms large, the client may contact only a single dealer. Front-running looms large, for example, if the dealers are likely to be initially long (so that $\underline{\phi}$ is large, by Proposition 4*(i)*), while the client is ex ante likely to sell (so that ϕ_0 is small), and if the client's realized trading need is indeed to sell. In such circumstances, the client optimally contacts only a single dealer with positive probability; this is why Proposition 5 says that $q_{-\bar{s}} > 0$ is optimal when $\phi_0 \in [0, \underline{\phi})$. Symmetric intuition explains why $q_{\bar{s}} > 0$ is optimal when $\phi_0 \in (\overline{\phi}, 1]$. Otherwise, the risk of front-running is small enough to endure in exchange for the countervailing benefits of an additional dealer, so that the client contacts only a single dealer with probability zero. This is why $q_{-\bar{s}} = q_{\bar{s}} = 0$ when $\phi_0 \in [\underline{\phi}, \overline{\phi}]$, why $q_{\bar{s}} = 0$ when $\phi_0 \in [0, \underline{\phi})$, and why $q_{-\bar{s}} = 0$ when $\phi_0 \in (\overline{\phi}, 1]$ in Proposition 5.

3.4.3 Examples

We illustrate Proposition 5 with two examples.

Asymmetric inventories. Suppose $\psi = 0.85$ and $\rho = 1$. Then $\hat{c}_1$ and $\hat{c}_2(\phi)$ are as depicted in the first panel of Figure 3. With this parametrization, $\hat{c}_2(\phi)$ is a decreasing function, and it in fact crosses $\hat{c}_1$ from above. The intuition is that when $\phi \approx 0$, the client is likely to be selling; at the same time, because $\psi = 0.85$, the dealers are likely to be long. Hence, if two dealers are contacted, the likely outcome is that the winning dealer will have to sell on the market, while the losing dealer will front-run, which raises the client's ultimate cost of procurement (relative to what it would have been if only one dealer had been contacted). On the other hand, when $\phi \approx 1$, the client is likely to be buying (and as before, the dealers are likely to be long). Hence, if two dealers are contacted, the likely outcome is that both could internalize the client's order, both will bid aggressively for the order, and the client's cost of procurement will be small (relative to what it would have been if only one dealer had been contacted).

The second panel of Figure 3 depicts $C(\phi)$, which is the convexification of the lower envelope of $\{\hat{c}_1, \hat{c}_2(\phi)\}$. This second panel also depicts $\underline{\phi}$; geometrically, this is the value of ϕ for which the line connecting $(0, \hat{c}_1)$ to $\big(\phi, \hat{c}_2(\phi)\big)$ is tangent to $c_2(\phi)$. Alternatively, this is the minimum value for which $C(\phi) = \hat{c}_2(\phi)$. We also have $\overline{\phi} = 1$ in this example, but suppress it in the figure.

The third panel depicts the case of $\phi_0 \in [\underline{\phi}, \overline{\phi}]$. According to Proposition 5, the optimal RFQ policy always contacts two dealers and discloses no information about the client's order. Under this policy, dealers' beliefs therefore always coincide with the prior, and the client's expected cost is $\hat{c}_2(\phi_0)$, which because $\phi_0 \in [\underline{\phi}, \overline{\phi}]$ is equal to $C(\phi_0)$.

Finally, the fourth panel depicts the case of $\phi_0 \in [0, \underline{\phi})$. According to Proposition 5, the optimal policy always contacts two dealers when $s = \bar{s}$, and it mixes between one and two dealers when $s = -\bar{s}$. Hence, if one dealer is contacted, dealers believe $s = \bar{s}$ with probability $\phi = 0$. Moreover, the mixing that occurs when $s = -\bar{s}$ is designed to ensure that, conditional on two dealers being contacted, they are induced to believe that $s = \bar{s}$ with probability $\phi = \underline{\phi}$. No further information is disclosed beyond this. Under this policy, the client's expected procurement cost is therefore an appropriate convex combination of $\hat{c}_1$ and $\hat{c}_2(\underline{\phi})$, which is precisely what $C(\phi_0)$ captures.

Figure 3: Optimal RFQ policies

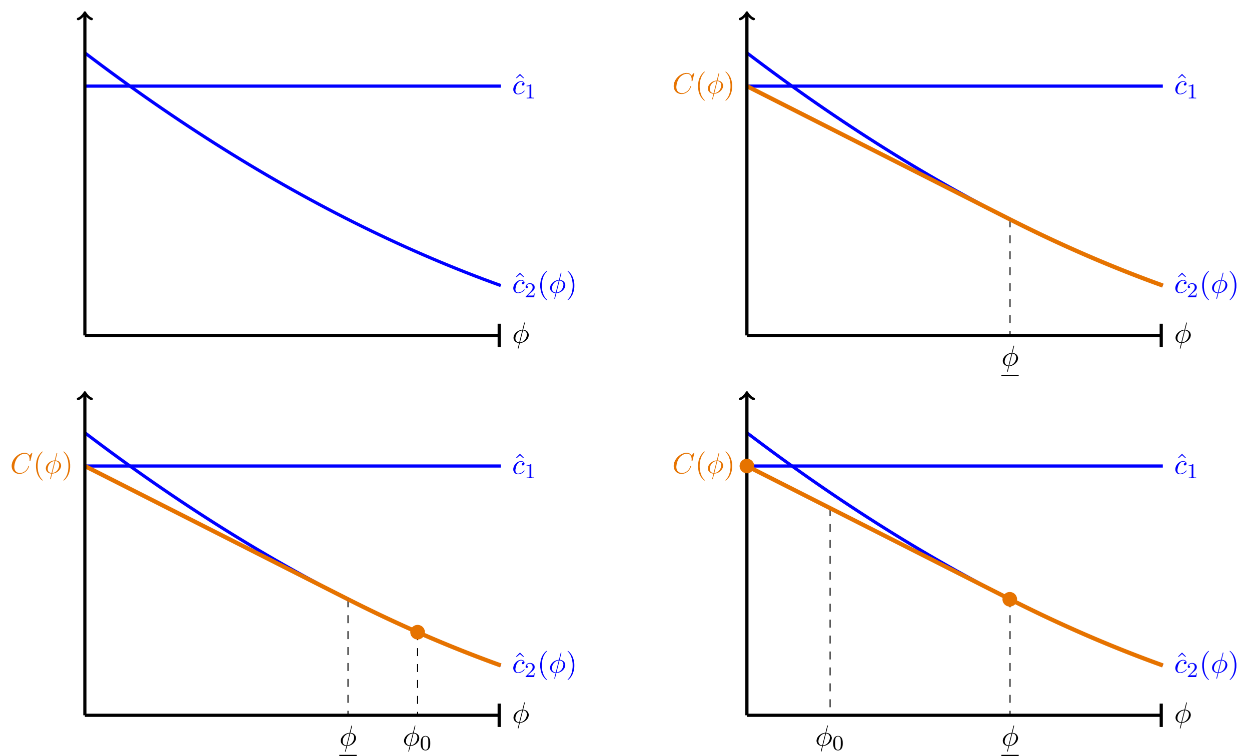

Dealer inventories are distributed according to the parameters $\psi = 0.85$ and $\rho = 1$. In all panels, $\hat{c}_1$ and $\hat{c}_2(\phi)$ (in blue) respectively depict expected procurement costs when contacting one or two dealers. Starting from the second panel, $C(\phi)$ (in orange) is, according to Proposition 5, the client's optimal procurement cost, and $\underline{\phi}$ is as in Definition 1. There are two cases that depend on ϕ_0, the prior probability of the client buying. If $\phi_0 \geq \underline{\phi}$ (as in the third panel), then the optimal RFQ policy always contacts two dealers, inducing a posterior belief that the client is a buyer which coincides with the prior ϕ_0. If $\phi_0 < \underline{\phi}$ (as in the fourth panel), then the optimal RFQ policy is such that contacting one dealer (two dealers) induces a posterior belief that the client is a buyer with probability zero (with probability $\underline{\phi}$). Vertical axes suppress units because all costs scale linearly with $\bar{s}^2$. Horizontal axes measure probability and depict the unit interval.

Symmetric inventories. If dealers are equally likely to be long and short (i.e., $\psi = \frac{1}{2}$), then $\hat{c}_2(\cdot)$ is minimized at $\phi = \frac{1}{2}$ and grows symmetrically in both directions to reach a maximum of $\hat{c}_2(0) = \hat{c}_2(1) = \frac{(15+\rho)\bar{s}^2}{32}$. It follows that regardless of ρ, $\hat{c}_2(\cdot)$ is everywhere less than $\hat{c}_1$, implying that $\underline{\phi} = 0$ and $\overline{\phi} = 1$, so that by Proposition 5, in these cases it is optimal for the client to always contact both dealers.

The intuition is that when dealers are equally likely to be long and short, there is a relatively high probability (between $\frac{1}{2}$ and $\frac{3}{4}$, depending on ρ) that one or more of the dealers will be able to internalize the client's order. The risk of front-running therefore does not loom particularly large. As a result, the client unambiguously benefits from inducing additional competition for her order.

3.5 Robustness

Online Appendix B investigates the robustness of our results to the assumption that the client can commit to an RFQ policy. We find that the result of Proposition 3 (on the optimality of minimal disclosure) remains fully intact: the client would not want to deviate from a commitment to minimal disclosure *even if she could*. The result of Proposition 5 (on the optimal policy for determining the number of dealers to contact) remains qualitatively intact: although the precise characterization does change without commitment, it remains the case that the client might restrict the number of dealers she contacts.

3.6 Testable implications

Our model makes a number of predictions that are testable empirically. One set of predictions concerns information disclosure: (*i*) if clients disclose more information about their order, then their procurement costs will increase on average; (*ii*) moreover, this effect will be stronger when more dealers are contacted; (*iii*) consequently, clients will tend to avoid disclosing information.

Another set concerns the number of dealers contacted: (*iv*) contacting more dealers tends to raise (lower) a client's procurement costs when dealers have much (little) difficulty internalizing; (*v*) consequently, in situations where dealers tend to be long, clients will tend to contact more dealers when buying than when selling; (*vi*) moreover, this difference is magnified when the client is ex ante more likely to sell; (*vii*) in particular, clients will sometimes contact fewer than all available dealers. Under the assumption that clients are pursuing their optimal RFQ policies, another prediction is: (*viii*) procurement costs are convex in the prior probability that the client is a buyer.

Other implications concern on-market trading: (*ix*) other dealers will initially tend to trade "with the wind" (i.e., in the same direction as the winning dealer), before reversing their direction to go "against the wind;" (*x*) moreover, the amount of initial "with the wind" trading increases when more dealers are contacted and when more information is disclosed.

Some of these predictions are already confirmed by existing empirical evidence. As mentioned, prediction (*iii*) is supported by the typical industry practice of asking for two-sided quotes when possible. Likewise, a great deal of evidence supports prediction (*vii*). The other predictions provide fertile ground for future empirical work.

RFQ platforms provide a particularly clean setting for testing these empirical implications. See Online Appendix I for a discussion of how to construct the variables that would be used in these tests.

3.7 Policy implications

Our analysis so far has solved for how an unconstrained client optimally behaves. But in practice, clients are sometimes constrained by existing institutions and regulations; for these cases, our analysis speaks to when such constraints might harm the client.

3.7.1 Regulation

One application of our results is to regulations that require a minimum number of potential counterparties to be contacted in certain situations. For example, the CFTC had once proposed to mandate that, for certain swaps, RFQs be sent to no fewer than five dealers.[30] (The proposal was since adopted with the requirement reduced to three.) Comment letters from Bloomberg, BlackRock, MetLife, Barclays, Morgan Stanley, and others objected, claiming that it is sometimes advantageous to contact fewer dealers—precisely because doing so limits the front-running and information leakage that we model:

> Several commenters specifically noted that the five market participant requirement may result in increased spreads for participants because non-executing market participants in the RFQ could "front run" the transaction in anticipation of the executing market participant's forthcoming and offsetting transactions. (CFTC, 2013)

Our analysis resonates with this concern about front-running. Indeed, our results suggest that there do exist circumstances in which concerns about front-running loom so large that a client optimally contacts only a single dealer. In those cases, mandates such as this one constitute a binding constraint and are likely to exacerbate transaction costs.

3.7.2 Platform design

In recent years, electronic trading platforms have been introduced with the goal of reducing search costs for many asset classes that had traditionally traded over the counter. One might have expected these platforms to be very popular. Indeed, in many classic models of over-the-counter markets (e.g., Duffie, Gârleanu and Pedersen, 2005), traders benefit from interventions that reduce or eliminate search costs. Yet, even with regulatory pushes (e.g., the Dodd-Frank Act), traders have been slow to adopt these platforms. As O'Hara and Zhou (2021, p. 369) write, "the puzzle remains why electronic trading has not taken on a larger role." Our analysis speaks to this puzzle by highlighting the importance of information leakage and the front-running it can lead to. Accordingly, it is not enough to simply reduce explicit search costs: platforms must also mitigate the implicit search cost of information leakage. The following design aspects are relevant:

Order exposure. Where information leakage is a concern, traders might rationally avoid platforms that expose their orders to a large number of potential counterparties. Instead, traders might be more receptive of platforms that permit them to fine-tune how broadly their orders are exposed. Not every RFQ platform offers this functionality; yet many of the more successful ones do (e.g., Bloomberg, CanDeal, CME's DRFQ, MarketAxess, Tradeweb).

[30] The originally stated rationale for this requirement was "to ensure that multiple participants have the ability to reach multiple counterparties" (CFTC, 2011).

Information policy. Many RFQ platforms require the client to reveal both the size and side of her desired trade, effectively mandating an information policy of full disclosure.[31] Our analysis highlights that clients might prefer more flexible information policies—in fact, our results imply that full disclosure is the *worst* information policy for the client.

Trader anonymity. Another variable in RFQ platform design is whether quotes can be requested anonymously or not. To capture a population of heterogenous clients, consider a version of the model in which the parameter ϕ_0 is first drawn from a distribution F. Suppose further that the client is exogenously required to use the RFQ policy that always contacts two dealers and discloses no information; this simplifies the analysis by making the client non-strategic, in particular preventing her from signaling her realized ϕ_0 through her RFQ policy. To capture a platform with name-disclosed trading, assume that ϕ_0 is observed by the dealers in conjunction with the RFQ. In this regime, the average client procurement cost would be $\int \hat{c}_2(\phi)\,dF(\phi)$. To capture a platform with anonymous trading, assume that ϕ_0 is never revealed. Letting $\phi_{avg} = \int \phi\,dF(\phi)$, the average client procurement cost would be then $\hat{c}_2(\phi_{avg})$, which is less, by the convexity of $\hat{c}_2(\cdot)$. Thus, clients are better off on average under anonymity in our model.[32]

Periodic auctions. Certain aspects of the RFQs that we model resemble the periodic auctions held by some equities exchanges (e.g., Cboe BYX, Cboe Europe) and alternative trading systems (e.g., CODA). Indeed, practitioners worry about information leakage and front-running in this context also (e.g., ESMA, 2019; SEC, 2021). Typically, these venues allow investors to initiate a uniform-price batch auction in a particular symbol. During the auction period, other investors, seeing that an auction is in progress, may submit orders also. Although these auctions typically conceal the direction of the initiating investor's order, the simple fact that an auction has been initiated could nevertheless be informative, potentially leading to front-running and suboptimal execution for reasons similar to those we model. This type of information leakage would be mitigated if—rather than being held only when initiated by an investor—batch auctions occurred at scheduled intervals, as in proposals made by academic literature (Budish, Cramton and Shim, 2014, 2015; Budish, Lee and Shim, forthcoming).

4 Conclusion

When seeking to trade in over-the-counter markets, institutional investors typically contact only a small number of potential counterparties and limit information disclosure (e.g., by requesting two-

[31] For example, this is the case for MarketAxess (O'Hara and Zhou, 2021), the Bloomberg and Tradeweb swap execution facilities (Riggs, Onur, Reiffen and Zhu, 2020), and CanDeal (Allen and Wittwer, forthcoming).

[32] Of course, we are unlikely to have $\phi_{avg} = \arg\min \hat{c}_2(\phi)$, and so a subset of client types may be worse off under anonymity. Note that if such client types could signal or disclose who they are, then we might expect the anonymous regime to unravel (as in, e.g., Grossman, 1981; Milgrom, 1981). Accordingly, it might be important for a platform to prevent such signaling or disclosure. Further outside our model, issues of unraveling would be accentuated in settings with informed trading, where clients with reputations for being uninformed would have an incentive to reveal their identities (e.g., Holden, Lu, Lugovskyy and Puzzello, 2021).

sided rather than one-sided quotes). These behaviors mitigate information leakage. Our analysis rationalizes them by demonstrating that in settings with price impact, information leakage may lead to endogenous front-running.

In our model, a client, who is either a buyer or a seller of a security, contacts either one or two dealers for a quote, potentially providing information about her trading direction. Dealers are initially either long or short. The winning dealer observes the client's trading direction and may then trade on the market in two periods. The other dealer does not directly observe the client's trading direction, yet he may also trade on the market, either to provide liquidity or to front-run.

In determining how many dealers to contact, the client faces a tradeoff. On the one hand, a wider search both intensifies competition for the client's order and improves the odds of finding a good match, i.e., a dealer who can internalize parts of the trade. But as mentioned above, information leakage is a countervailing force: a wider search alerts more dealers to the existence of her trade, which can lead to costly front-running. In certain cases, the costs of information leakage are large enough that, with positive probability, the client contacts only a single dealer. In that sense, front-running constitutes an endogenous search friction, which may limit the scope of counterparty search.

We obtain stark implications for information design. Secrecy is best for the client in the following sense: she optimally discloses no additional information about whether she intends to buy or sell—in other words, she requests quotes for a two-sided market.

Beyond contributing a model of this search problem, our analysis also has important implications for the regulation of over-the-counter securities trading and for the design of request-for-quote platforms. Rules in these settings should be evaluated in light of endogenous front-running.

References

Allen, Jason and Milena Wittwer, "Centralizing Over-the-Counter Markets?," *Journal of Political Economy*, forthcoming. https://doi.org/10.1086/725361.

Baldauf, Markus, Christoph Frei, and Joshua Mollner, "Principal Trading Arrangements: When Are Common Contracts Optimal?," *Management Science*, 2022, *68* (4), 3112–3128.

_ , _ , and _ , "Principal Trading Arrangements: Optimality under Temporary and Permanent Price Impact," *Working Paper*, 2023. https://ssrn.com/abstract=3778956.

Bank for International Settlements, "Electronic Trading in Fixed Income Markets," *Markets Committee Report*, January 2016. https://www.bis.org/publ/mktc07.pdf.

Bergemann, Dirk and Martin Pesendorfer, "Information Structures in Optimal Auctions," *Journal of Economic Theory*, 2007, *137* (1), 580–609.

_ , Tibor Heumann, Stephen Morris, Constantine Sorokin, and Eyal Winter, "Optimal Information Disclosure in Classic Auctions," *American Economic Review: Insights*, 2022, *4* (3), 371–388.

Bernhardt, Dan, Vladimir Dvoracek, Eric Hughson, and Ingrid M. Werner, "Why Do Larger Orders Receive Discounts on the London Stock Exchange?," *The Review of Financial Studies*, 2004, *18* (4), 1343–1368.

Bertsimas, Dimitris and Andrew W. Lo, "Optimal Control of Execution Costs," *Journal of Financial Markets*, 1998, *1* (1), 1–50.

Bikhchandani, Sushil and Chi-fu Huang, "Auctions with Resale Markets: An Exploratory Model of Treasury Bill Markets," *The Review of Financial Studies*, 1989, *2* (3), 311–339.

Board, Simon, "Relational Contracts and the Value of Loyalty," *American Economic Review*, 2011, *101* (7), 3349–3367.

Brunnermeier, Markus K., "Information Leakage and Market Efficiency," *The Review of Financial Studies*, 2005, *18* (2), 417–457.

_ and Lasse Heje Pedersen, "Predatory Trading," *The Journal of Finance*, 2005, *60* (4), 1825–1863.

Budish, Eric, Peter Cramton, and John Shim, "Implementation Details for Frequent Batch Auctions: Slowing Down Markets to the Blink of an Eye," *American Economic Review: Papers & Proceedings*, 2014, *104* (5), 418–424.

_ , _ , and _ , "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response," *The Quarterly Journal of Economics*, 2015, *130* (4), 1547–1621.

_ , Robin S. Lee, and John J. Shim, "A Theory of Stock Exchange Competition and Innovation: Will the Market Fix the Market?," *Journal of Political Economy*, forthcoming. https://www.nber.org/papers/w25855.

Bulow, Jeremy and Paul Klemperer, "Auctions versus Negotiations," *American Economic*

Review, 1996, *86* (1), 180–194.

— **and** — , "Prices and the Winner's Curse," *The RAND Journal of Economics*, 2002, *33* (1), 1–21.

Burdett, Kenneth and Maureen O'Hara, "Building Blocks: An Introduction to Block Trading," *Journal of Banking & Finance*, 1987, *11* (2), 193–212.

Calzolari, Giacomo and Alessandro Pavan, "Monopoly with Resale," *The RAND Journal of Economics*, 2006, *37* (2), 362–375.

CanDeal, "CanDeal's 2019 Year in Review," February 2020. https://www.candeal.com/en/news/candeals-2019-year-review.

Carlin, Bruce Ian, Miguel Sousa Lobo, and S. Viswanathan, "Episodic Liquidity Crises: Cooperative and Predatory Trading," *The Journal of Finance*, 2007, *62* (5), 2235–2274.

Carroll, Gabriel and Ilya Segal, "Robustly Optimal Auctions with Unknown Resale Opportunities," *The Review of Economic Studies*, 2019, *86* (4), 1527–1555.

Che, Yeon-Koo and Ian Gale, "Optimal Design of Research Contests," *American Economic Review*, 2003, *93* (3), 646–671.

Cho, In-Koo and David M. Kreps, "Signaling Games and Stable Equilibria," *The Quarterly Journal of Economics*, 1987, *102* (2), 179–221.

Collin-Dufresne, Pierre, Benjamin Junge, and Anders B. Trolle, "Market Structure and Transaction Costs of Index CDSs," *The Journal of Finance*, 2020, *75* (5), 2719–2763.

Commodity Futures Trading Commission, "Core Principles and Other Requirements for Swap Execution Facilities," *Federal Register*, January 2011, *76* (5), 1213–1259.

— , "Core Principles and Other Requirements for Swap Execution Facilities," *Federal Register*, June 2013, *78* (107), 33475–33604.

— , "Swaps Regulation Version 2.0: An Assessment of the Current Implementation of Reform and Proposals for Next Steps," *White Paper*, April 2018. https://bit.ly/3eebTio.

Crémer, Jacques and Richard P. McLean, "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 1988, *56* (6), 1247–1257.

Di Maggio, Marco, Amir Kermani, and Zhaogang Song, "The Value of Trading Relations in Turbulent Times," *Journal of Financial Economics*, 2017, *124* (2), 266–284.

Duffie, Darrell, Nicolae Gârleanu, and Lasse Heje Pedersen, "Over-the-Counter Markets," *Econometrica*, 2005, *73* (6), 1815–1847.

Dugast, Jérôme, Semih Üslü, and Pierre-Olivier Weill, "A Theory of Participation in OTC and Centralized Markets," *The Review of Economic Studies*, 2022, *89* (6), 3223–3266.

Electronic Debt Markets Association – Europe, "The Value of RFQ," December 2018. https://bit.ly/3z4iLbX.

Eső, Péter and Balazs Szentes, "Optimal Information Disclosure in Auctions and the Handicap Auction," *The Review of Economic Studies*, 2007, *74* (3), 705–731.

European Securities and Markets Authority, "Final Report: Call for Evidence on Periodic Auctions," June 2019. https://bit.ly/2PaECcE.

Fullerton, Richard L. and R. Preston McAfee, "Auctioning Entry into Tournaments," *Journal of Political Economy*, 1999, *107* (3), 573–605.

FX Markets, "For FX Liquidity Buyers, Fewer Providers Mean Better Execution," October 2021. https://fx-markets.com/7886951.

Ganuza, Juan-José and Jose S. Penalva, "Signal Orderings Based on Dispersion and the Supply of Private Information in Auctions," *Econometrica*, 2010, *78* (3), 1007–1030.

Garratt, Rod and Thomas Tröger, "Speculation in Standard Auctions with Resale," *Econometrica*, 2006, *74* (3), 753–769.

Garratt, Rodney J., Thomas Tröger, and Charles Z. Zheng, "Collusion via Resale," *Econometrica*, 2009, *77* (4), 1095–1136.

Glebkin, Sergei, Bart Zhou Yueshen, and Ji Shen, "Simultaneous Multilateral Search," *The Review of Financial Studies*, 2023, *36* (2), 571–614.

Grossman, Sanford J., "The Informational Role of Warranties and Private Disclosure about Product Quality," *The Journal of Law and Economics*, 1981, *24* (3), 461–483.

Gupta, Madhurima and Bernard Lebrun, "First Price Auctions with Resale," *Economics Letters*, 1999, *64* (2), 181–185.

Hafalir, Isa and Vijay Krishna, "Asymmetric Auctions with Resale," *American Economic Review*, 2008, *98* (1), 87–112.

Haile, Philip A., "Auctions with Private Uncertainty and Resale Opportunities," *Journal of Economic Theory*, 2003, *108* (1), 72–110.

Hau, Harald, Peter Hoffmann, Sam Langfield, and Yannick Timmer, "Discriminatory Pricing of Over-the-Counter Derivatives," *Management Science*, 2021, *67* (11), 6660–6677.

Hendershott, Terrence and Ananth Madhavan, "Click or Call? Auction versus Search in the Over-The-Counter Market," *The Journal of Finance*, 2015, *70* (1), 419–447.

_ , Dan Li, Dmitry Livdan, and Norman Schürhoff, "Relationship Trading in Over-the-Counter Markets," *The Journal of Finance*, 2020, *75* (2), 683–734.

_ , _ , _ , and Norman Schürhoff, "True Cost of Immediacy," *Working Paper*, 2022. https://ssrn.com/abstract=3680035.

_ , Dmitry Livdan, and Norman Schürhoff, "All-to-All Liquidity in Corporate Bonds," *Working Paper*, 2021. https://ssrn.com/abstract=3895270.

Holden, Craig W., Dong Lu, Volodymyr Lugovskyy, and Daniela Puzzello, "What is the Impact of Introducing a Parallel OTC Market? Theory and Evidence from the Chinese Interbank FX Market," *Journal of Financial Economics*, 2021, *140* (1), 270–291.

Jehiel, Philippe and Benny Moldovanu, "Strategic Nonparticipation," *The RAND Journal of*

Economics, 1996, *27* (1), 84–98.

_ **and** _ , "Auctions with Downstream Interaction among Buyers," *The RAND Journal of Economics*, 2000, *31* (4), 768–791.

_ **,** _ **, and Ennio Stacchetti**, "How (Not) to Sell Nuclear Weapons," *American Economic Review*, 1996, *86* (4), 814–829.

Kamenica, Emir and Matthew Gentzkow, "Bayesian Persuasion," *American Economic Review*, 2011, *101* (6), 2590–2615.

Kondor, Péter and Gábor Pintér, "Clients' Connections: Measuring the Role of Private Information in Decentralized Markets," *The Journal of Finance*, 2022, *77* (1), 505–544.

Lauermann, Stephan and Asher Wolinsky, "Bidder Solicitation, Adverse Selection, and the Failure of Competition," *American Economic Review*, 2017, *107* (6), 1399–1429.

Lee, Tomy and Chaojun Wang, "Why Trade Over-the-Counter?," *Working Paper*, 2022. https://finance.wharton.upenn.edu/~wangchj/papers/OTC_EX.pdf.

Levin, Dan and James L. Smith, "Equilibrium in Auctions with Entry," *American Economic Review*, 1994, *84* (3), 585–599.

Li, Hao and Xianwen Shi, "Discriminatory Information Disclosure," *American Economic Review*, 2017, *107* (11), 3363–3385.

Liu, Ying, Sebastian Vogel, and Yuan Zhang, "Electronic Trading in OTC Markets vs. Centralized Exchange," *Working Paper*, 2018. https://ssrn.com/abstract=3189834.

Menezes, Flavio M. and Paulo K. Monteiro, "Auctions with Endogenous Participation," *Review of Economic Design*, 2000, *5* (1), 71–89.

Milgrom, Paul R., "Good News and Bad News: Representation Theorems and Applications," *The Bell Journal of Economics*, 1981, *12* (2), 380–391.

_ **and Robert J. Weber**, "A Theory of Auctions and Competitive Bidding," *Econometrica*, 1982, *50* (5), 1089–1122.

Myerson, Roger B., "Optimal Auction Design," *Mathematics of Operations Research*, 1981, *6* (1), 58–73.

Nasdaq, Inc., "How Much Does Trading Cost the Buy Side?," February 2022. https://bit.ly/3U1O2Gw.

O'Hara, Maureen and Xing Alex Zhou, "The Electronic Evolution of Corporate Bond Dealers," *Journal of Financial Economics*, 2021, *140* (2), 368–390.

O'Hara, Maureen, Yihui Wang, and Xing Alex Zhou, "The Execution Quality of Corporate Bonds," *Journal of Financial Economics*, 2018, *130* (2), 308–326.

Riggs, Lynn, Esen Onur, David Reiffen, and Haoxiang Zhu, "Swap Trading after Dodd-Frank: Evidence from Index CDS," *Journal of Financial Economics*, 2020, *137* (3), 857–886.

Risk.net, "Buy Side Using Two-Way Prices in Bid to Hide Trade Intent," August 2018. https:

//risk.net/5858751.

Sannikov, Yuliy and Andrzej Skrzypacz, "Dynamic Trading: Price Inertia and Front-Running," *Working Paper*, 2016. https://ssrn.com/abstract=2882809.

Securities and Exchange Commission, "Self-Regulatory Organizations; Cboe BYX Exchange, Inc.; Notice of Filing of Amendments No. 3 and No. 4, and Order Granting Accelerated Approval of a Proposed Rule Change, as Modified by Amendments No. 3 and No. 4, To Introduce Periodic Auctions for the Trading of U.S. Equity Securities," March 2021. https://www.sec.gov/rules/sro/cboebyx/2021/34-91423.pdf.

_ , "Amendments Regarding the Definition of "Exchange" and Alternative Trading Systems (ATSs) That Trade U.S. Treasury and Agency Securities, National Market System (NMS) Stocks, and Other Securities," *Federal Register*, March 2022, *87* (53), 15496–15696.

Securities Industry and Financial Markets Association, "Global Equity Markets Primer: Market Costs," November 2021. https://bit.ly/3eGr5IL.

Seppi, Duane J., "Equilibrium Block Trading and Asymmetric Information," *The Journal of Finance*, 1990, *45* (1), 73–94.

Stigler, George J., "The Economics of Information," *Journal of Political Economy*, 1961, *69* (3), 213–225.

Taylor, Curtis R., "Digging for Golden Carrots: An Analysis of Research Tournaments," *American Economic Review*, 1995, *85* (4), 872–890.

Viswanathan, S. and James J. D. Wang, "Inter-Dealer Trading in Financial Markets," *The Journal of Business*, 2004, *77* (4), 987–1040.

Wang, Chaojun, "The Limits of Multi-Dealer Platforms," *Journal of Financial Economics*, 2023, *149* (3), 434–450.

Weitzman, Martin L., "Optimal Search for the Best Alternative," *Econometrica*, 1979, *47* (3), 641–654.

Wilson, Robert, "Auctions of Shares," *The Quarterly Journal of Economics*, 1979, *93* (4), 675–689.

Yueshen, Bart Zhou and Junyuan Zou, "Less is More," *Working Paper*, 2023. https://ssrn.com/abstract=4274063.

Zheng, Charles Zhoucheng, "Optimal Auction with Resale," *Econometrica*, 2002, *70* (6), 2197–2224.

Zhu, Haoxiang, "Finding a Good Price in Opaque Over-the-Counter Markets," *The Review of Financial Studies*, 2012, *25* (4), 1255–1285.

Appendix

What if dealers cannot front-run?

To develop a benchmark, this appendix analyzes a version of the model in which losing dealers cannot front-run. Although the risk of front-running is both real and significant, it might be mitigated by a variety of considerations—including repeated interaction (as discussed in Section 2.2), reputational considerations, or legal concerns. The extreme case is that no front-running occurs at all. To model it, we assume for purposes of this appendix that the losing dealer cannot trade in the first period.

This changes nothing for the case in which only one dealer is contacted. Letting $\tilde{c}_1$ denote the minimum procurement cost for that case, we have $\tilde{c}_1 = \hat{c}_1 = \frac{3\bar{s}^2}{4}$. Turning to the case in which two dealers are contacted, we obtain the following analogue of Lemma 2:

Lemma 2′. *In the class of RFQs that contact $M = 2$ dealers, induce beliefs ϕ, and lead to execution with probability one, the minimum expected procurement cost is a linear function $\tilde{c}_2(\phi)$ that satisfies $\tilde{c}_2(\phi) \leq \tilde{c}_1$ and $\tilde{c}_2(\phi) \leq \hat{c}_2(\phi)$. It is achieved by RFQs featuring reserve prices $\bar{b} = (\bar{s}^2, \bar{s}^2)$.*

Proof of Lemma 2′. In this scenario, equilibrium in the continuation trading game is precisely as in the case where one dealer is contacted (*cf.* Lemma A1). Hence, trading costs are the same as well. It therefore only remains to derive equilibrium bids, which can be done following the logic described in Section 3.3: A's bid should reflect the difference in his trading costs between the scenario where he wins and the scenario where B wins. Thus, A's equilibrium bids are

$$(b^A_{-\bar{s}}, b^A_{\bar{s}}) = \begin{cases} \left(\frac{3\bar{s}^2}{4}, 0\right) & \text{if } (e^A, e^B) = (1, 1), \\ \left(\frac{7\bar{s}^2}{16}, \frac{\bar{s}^2}{4}\right) & \text{if } (e^A, e^B) = (1, -1), \\ \left(\frac{\bar{s}^2}{4}, \frac{7\bar{s}^2}{16}\right) & \text{if } (e^A, e^B) = (-1, 1), \\ \left(0, \frac{3\bar{s}^2}{4}\right) & \text{if } (e^A, e^B) = (-1, -1). \end{cases}$$

From this, we can derive the client's procurement costs. With probability $\phi\psi[1 - (1 - \psi)(1 - \rho)]$, $(s, e^A, e^B) = (\bar{s}, 1, 1)$, and the client's procurement cost will be 0. With probability $(1-\phi)\psi[1-(1-\psi)(1 - \rho)]$, $(s, e^A, e^B) = (-\bar{s}, 1, 1)$, and the client's procurement cost will be $\frac{3\bar{s}^2}{4}$. With probability $\phi(1-\psi)[1-\psi(1-\rho)]$, $(s, e^A, e^B) = (\bar{s}, -1, -1)$, and the client's procurement cost will be $\frac{3\bar{s}^2}{4}$. With probability $(1-\phi)(1-\psi)[1-\psi(1-\rho)]$, $(s, e^A, e^B) = (-\bar{s}, -1, -1)$, and the client's procurement cost will be 0. With the remaining probability $2\psi(1 - \psi)(1 - \rho)$, $e^A \neq e^B$, and the client's procurement cost will be $\max\{\frac{7\bar{s}^2}{16}, \frac{\bar{s}^2}{4}\} = \frac{7\bar{s}^2}{16}$. In each case, the client's procurement cost is weakly less than $\tilde{c}_1 = \frac{3\bar{s}^2}{4}$. Hence, the client's expected procurement cost $\tilde{c}_2(\phi)$ is also weakly less than $\tilde{c}_1$. In fact, we can compute

$$\tilde{c}_2(\phi) = \frac{3\bar{s}^2}{4}\Big((1-\phi)\psi\big[1-(1-\psi)(1-\rho)\big] + \phi(1-\psi)\big[1-\psi(1-\rho)\big]\Big) + \frac{7\bar{s}^2}{8}\psi(1-\psi)(1-\rho), \quad (3)$$

which is indeed linear in ϕ. Comparing this expression against that for $\hat{c}_2(\phi)$ provided by equation (10), we also see that $\tilde{c}_2(\phi) \leq \hat{c}_2(\phi)$. □

With this result in hand, we can immediately characterize the optimal RFQ policy—and see how it differs from the optimal policy when front-running is a concern.

- *Is the client better off if front-running is not a concern?* Yes. This follows from $\tilde{c}_2(\phi)$ being everywhere weakly less than $\hat{c}_2(\phi)$, together with $\tilde{c}_1 = \hat{c}_1$.

- *If front-running is not a concern, should the client always call both dealers?* Yes. This follows from $\tilde{c}_2(\phi)$ being everywhere weakly less than $\tilde{c}_1$. In contrast, she might limit herself to contacting only one dealer when front-running is a concern (*cf.* Proposition 5).

- *If front-running is not a concern, should the client reveal more information?* Yes, weakly. This follows from $\tilde{c}_2(\phi)$ being linear in ϕ, which implies that the client is indifferent over all information structures. In contrast, she has a strict preference for revealing as little information as possible when front-running is a concern (*cf.* Proposition 3).

  The intuition is that information design does not matter for the winning dealer, since he will anyway observe s before he starts trading. Remember also that dealers submit contingent bids (one bid for the $s = \bar{s}$ case and another for $s = -\bar{s}$), so the fact that the winning dealer may not know s while he is bidding does not matter: for each case, he can simply bid what he would bid if he knew that case was the correct one. Similarly, information design also does not matter for the losing dealer here: not for informing his first-period trade, since he cannot make one in this version of the model, nor for informing his second-period trade, since by that point he will have already inferred s from the winning dealer's first-period trade.

One interpretation of this analysis is that it constitutes an argument that front-running is a real concern in practice. Indeed, the optimal RFQ policy described above is inconsistent with certain stylized facts: investors seeking to trade sometimes conduct limited searches, and these investors display a strict preference for revealing as little information as possible. These stylized facts are much better explained by our analysis of what is optimal when front-running is a concern.

Online Appendix

A Proofs

A.1 Proof of Lemma 1

To establish Lemma 1, we in fact prove the stronger result stated in Lemma A1.

Lemma A1. *In a subgame following an RFQ that contacts $M = 1$ dealer, the unique on-path equilibrium behavior is as follows. Dealer A bids*

$$(b^A_{-\bar{s}}, b^A_{\bar{s}}) = \begin{cases} \left(\frac{3\bar{s}^2}{4}, 0\right) & \textit{if } (e^A, e^B) = (1, 1), \\ \left(\frac{7\bar{s}^2}{16}, 0\right) & \textit{if } (e^A, e^B) = (1, -1), \\ \left(0, \frac{7\bar{s}^2}{16}\right) & \textit{if } (e^A, e^B) = (-1, 1), \\ \left(0, \frac{3\bar{s}^2}{4}\right) & \textit{if } (e^A, e^B) = (-1, -1). \end{cases}$$

If dealer A wins, the on-market trades are

$$(x^A_1, x^A_2, x^B_2) = \begin{cases} (0, 0, 0) & \textit{if } (s, e^A, e^B) = (\bar{s}, 1, 1), \\ (-\frac{\bar{s}}{2}, -\frac{\bar{s}}{2}, 0) & \textit{if } (s, e^A, e^B) = (-\bar{s}, 1, 1), \\ (0, 0, 0) & \textit{if } (s, e^A, e^B) = (\bar{s}, 1, -1), \\ (-\frac{\bar{s}}{4}, -\frac{3\bar{s}}{4}, \frac{\bar{s}}{2}) & \textit{if } (s, e^A, e^B) = (-\bar{s}, 1, -1), \\ (\frac{\bar{s}}{4}, \frac{3\bar{s}}{4}, -\frac{\bar{s}}{2}) & \textit{if } (s, e^A, e^B) = (\bar{s}, -1, 1), \\ (0, 0, 0) & \textit{if } (s, e^A, e^B) = (-\bar{s}, -1, 1), \\ (\frac{\bar{s}}{2}, \frac{\bar{s}}{2}, 0) & \textit{if } (s, e^A, e^B) = (\bar{s}, -1, -1), \\ (0, 0, 0) & \textit{if } (s, e^A, e^B) = (-\bar{s}, -1, -1). \end{cases}$$

Lemma 1 follows from Lemma A1 for reasons discussed in the main text. To obtain execution with probability one, the client's reserve must be at least the dealer's bid in the worst case. If the client sells, the worst case is $(e^A, e^B) = (1, 1)$, where $b^A_{-\bar{s}} = \frac{3\bar{s}^2}{4}$, according to Lemma A1. Symmetrically, if the client buys, the worst case is $(e^A, e^B) = (-1, -1)$, where $b^A_{\bar{s}} = \frac{3\bar{s}^2}{4}$. Therefore, it follows that the expected procurement cost resulting from an RFQ that contacts one dealer, induces a belief that ϕ is the probability of $s = \bar{s}$, and leads to execution with probability one is at least

$$\hat{c}_1 = \phi\frac{3\bar{s}^2}{4} + (1 - \phi)\frac{3\bar{s}^2}{4} = \frac{3\bar{s}^2}{4},$$

as claimed by Lemma 1. Moreover, the RFQ that uses reserve prices $\bar{b} = (\frac{3\bar{s}^2}{4}, \frac{3\bar{s}^2}{4})$ achieves this bound.

Proof of Lemma A1. Because both dealers observe the entire vector (e^A, e^B), the four possible realizations of that vector can be analyzed separately. Below, we analyze the cases of $(1,1)$ and $(1,-1)$; the remaining cases can be handled symmetrically.

Case 1: $(e^A, e^B) = (1,1)$. Here is a full specification of an equilibrium for this case. Dealer A bids $\left(b^A_{-\bar{s}}, b^A_{\bar{s}}\right) = \left(\frac{3\bar{s}^2}{4}, 0\right)$. If he loses, the game ends. Therefore, suppose henceforth that he wins. If $s = \bar{s}$, dealer A sets $x_1^A = 0$ and

$$x_2^A = \begin{cases} -\frac{x_1^A}{3} & \text{if } 0 \le x_1^A \le \bar{s}, \\ -\frac{x_1^A}{2} & \text{if } \max\{-\bar{s}, 2\bar{s}-4\} \le x_1^A < 0, \\ \bar{s} - 2 - x_1^A & \text{if } -\bar{s} \le x_1^A \le 2\bar{s} - 4. \end{cases} \tag{4}$$

If $s = -\bar{s}$, dealer A sets $x_1^A = -\frac{\bar{s}}{2}$ and

$$x_2^A = \begin{cases} -\bar{s} - x_1^A & \text{if } -\bar{s} \le x_1^A \le \bar{s}. \end{cases} \tag{5}$$

Dealer B sets

$$x_2^B = \begin{cases} -\frac{x_1^A}{3} & \text{if } 0 \le x_1^A \le \bar{s}, \\ 0 & \text{if } -\bar{s} \le x_1^A < 0. \end{cases} \tag{6}$$

Dealers' beliefs prior to bidding are $\mu_0^A = \mu_0^B = \phi$. Dealer B's beliefs prior to second-period trading are

$$\mu_2^B = \begin{cases} 1 & \text{if } 0 \le x_1^A \le \bar{s}, \\ 0 & \text{if } -\bar{s} \le x_1^A < 0. \end{cases}$$

We claim that the specified strategies and beliefs satisfy the solution concept described in Section 3.1 and moreover that anything else also satisfying the solution concept must feature the same on-path behavior. The argument consists of three parts.

Part (i): We check the consistency of dealer B's beliefs. Given the specified strategy for dealer A, Bayes' rule requires only that

$$\mu_2^B = \begin{cases} 1 & \text{if } x_1^A = 0, \\ 0 & \text{if } x_1^A = -\frac{\bar{s}}{2}. \end{cases}$$

This is indeed consistent with the specified beliefs.

Part (ii): Given the specified beliefs, we check that the solution concept uniquely pins down the specified strategies. We proceed by backward induction:

- *Period-2 reaction functions.* Dealer A's trading costs are $x_1^A x_1^A + (x_1^A + x_2^A + x_2^B)x_2^A$. For $s \in \{-\bar{s}, \bar{s}\}$, dealer A best responds with $x_2^A = \left[-\frac{x_1^A + x_2^B}{2}\right]_{s-2-x_1^A}^{s-x_1^A}$.[33] Dealer B's trading costs are $(x_1^A + x_2^A + x_2^B)x_2^B$. Dealer B best responds with $x_2^B = \left[-\frac{x_1^A + x_2^A}{2}\right]_{-2}^{0}$.

[33] We use the notation $[x]_a^b$ to denote truncation to the interval $[a,b]$: $[x]_a^b = \max\left(a, \min(b,x)\right)$.

- *Dealer B's period-2 action.* If $0 \leq x_1^A \leq \bar{s}$ so that $\mu_2^B = 1$, then x_2^B is pinned down by the intersection of $x_2^A = \left[-\frac{x_1^A + x_2^B}{2}\right]_{\bar{s}-2-x_1^A}^{\bar{s}-x_1^A}$ and $x_2^B = \left[-\frac{x_1^A + x_2^A}{2}\right]_{-2}^{0}$, so that we indeed have $x_2^B = -\frac{x_1^A}{3}$. If $-\bar{s} \leq x_1^A < 0$ so that $\mu_2^B = 0$, then x_2^B is pinned down by the intersection of $x_2^A = \left[-\frac{x_1^A + x_2^B}{2}\right]_{-\bar{s}-2-x_1^A}^{-\bar{s}-x_1^A}$ and $x_2^B = \left[-\frac{x_1^A + x_2^A}{2}\right]_{-2}^{0}$, so that we indeed have $x_2^B = 0$. Together, these two cases verify (6).

- *Dealer A's period-2 action.* If $s = \bar{s}$, then x_2^A is pinned down by the intersection of $x_2^A = \left[-\frac{x_1^A + x_2^B}{2}\right]_{\bar{s}-2-x_1^A}^{\bar{s}-x_1^A}$ and (6), which verifies (4).

  If $s = -\bar{s}$, then x_2^A is pinned down by the intersection of $x_2^A = \left[-\frac{x_1^A + x_2^B}{2}\right]_{-\bar{s}-2-x_1^A}^{-\bar{s}-x_1^A}$ and (6), which verifies (5).

- *Dealer A's period-1 action.* Dealer A's trading costs are $x_1^A x_1^A + (x_1^A + x_2^A + x_2^B)x_2^A$. If $s = \bar{s}$, then we can plug in (4) and (6) to express dealer A's trading costs as a function of x_1^A:

$$\begin{cases} x_1^A x_1^A - \left(\frac{x_1^A}{3}\right)^2 & \text{if } 0 \leq x_1^A \leq \bar{s}, \\ x_1^A x_1^A - \left(\frac{x_1^A}{2}\right)^2 & \text{if } \max\{s, 2\bar{s}-4\} \leq x_1^A < 0, \\ x_1^A x_1^A + (\bar{s}-2)(\bar{s}-2-x_1^A) & \text{if } -\bar{s} \leq x_1^A < 2\bar{s}-4, \end{cases}$$

  which is indeed minimized by $x_1^A = 0$.

  Alternatively, if $s = -\bar{s}$, then we can plug in (5) and (6) to express dealer A's trading costs as a function of x_1^A:

$$\begin{cases} x_1^A x_1^A + (-\bar{s} - \frac{x_1^A}{3})(-\bar{s} - x_1^A) & \text{if } 0 \leq x_1^A \leq \bar{s}, \\ x_1^A x_1^A + (-\bar{s})(-\bar{s} - x_1^A) & \text{if } -\bar{s} \leq x_1^A < 0, \end{cases}$$

  which is indeed minimized by $x_1^A = -\frac{\bar{s}}{2}$.

- *Dealer A's bid.* Plugging in the trading behavior derived above, we have the following. If $s = \bar{s}$, then dealer A's continuation utility is c if he wins and 0 if he loses. If $s = -\bar{s}$, then dealer A's continuation utility is $c - \frac{3\bar{s}^2}{4}$ if he wins and 0 if he loses. Based on the solution concept described in Section 3.1, dealer A must therefore bid $\left(b_{-\bar{s}}^A, b_{\bar{s}}^A\right) = \left(\frac{3\bar{s}^2}{4}, 0\right)$.

Part (iii): We check that the equilibrium specified above satisfies the restrictions on beliefs described in Section 3.1. And we also show that any equilibrium satisfying those conditions must feature on-path behavior coinciding with that of the equilibrium specified above. Let $\tilde{\mu}_2^B$ be candidate beliefs. One requirement is that $\tilde{\mu}_2^B$ must have the threshold structure described in the text.[34]

[34] That is, (*i*) for all $x \in [-\bar{s}, \bar{s}]$, $\tilde{\mu}_2^B(x) \in \{0, 1\}$, and (*ii*) for all $x' < x''$, $\tilde{\mu}_2^B(x') = 1 \implies \tilde{\mu}_2^B(x'') = 1$ and $\tilde{\mu}_2^B(x'') = 0 \implies \tilde{\mu}_2^B(x') = 0$.

- We start by establishing an upper bound for the location of the jump. Suppose, by way of contradiction, that we have an equilibrium with beliefs such that $\tilde{\mu}_2^B(0) = 0$. By the threshold structure of $\tilde{\mu}_2^B$ and because beliefs must be correct on path, this means that on path, dealer A sets $x_1^A > 0$ when $s = \bar{s}$. Arguments similar to those given above imply that we would subsequently have $(x_2^A, x_2^B) = (-\frac{x_1^A}{3}, -\frac{x_1^A}{3})$, leading to total trading costs for dealer A of $(x_1^A)^2 - \left(\frac{x_1^A}{3}\right)^2 = \frac{8}{9}(x_1^A)^2$. On the other hand, suppose that dealer A deviated to set $x_1^A = 0$ when $s = \bar{s}$. Arguments similar to those given above imply that we would subsequently have $(x_2^A, x_2^B) = (0, 0)$, leading to trading costs of 0. This constitutes a profitable deviation, contradicting the putative equilibrium.

- We next establish a corresponding lower bound. Suppose, by way of contradiction, that we have an equilibrium with beliefs such that $\tilde{\mu}_2^B(-\frac{\bar{s}}{2}) = 1$. By the threshold structure of $\tilde{\mu}_2^B$ and because beliefs must be correct on path, this means that on path, dealer A sets $x_1^A < -\frac{\bar{s}}{2}$ when $s = -\bar{s}$. Arguments similar to those given above imply that we would subsequently have $(x_2^A, x_2^B) = (-\bar{s} - x_1^A, 0)$, leading to trading costs for dealer A of $(x_1^A)^2 + (-\bar{s})(-\bar{s} - x_1^A) = \frac{3\bar{s}^2}{4} + (\frac{\bar{s}}{2} + x_1^A)^2$. On the other hand, suppose that dealer A deviated to set $x_1^A = -\frac{\bar{s}}{2}$ when $s = \bar{s}$. Arguments similar to those given above imply that we would subsequently have $(x_2^A, x_2^B) = (-\frac{\bar{s}}{2}, 0)$, leading to trading costs of $\frac{3\bar{s}^2}{4}$. This constitutes a profitable deviation, contradicting the putative equilibrium.

- Finally, we turn to the remaining possibilities. Suppose we have an equilibrium with beliefs such that both $\tilde{\mu}_2^B(0) = 1$ and $\tilde{\mu}_2^B(-\frac{\bar{s}}{2}) = 0$. Using arguments similar to those given above, we can show that any such equilibrium induces the same on-path behavior as the equilibrium specified above. In particular, dealer A's equilibrium trading costs are as above: $C_*^A(\bar{s}) = 0$ and $C_*^A(-\bar{s}) = \frac{3\bar{s}^2}{4}$. We can also use arguments similar to those given above to compute

$$C^A(\bar{s}, x_1^A, 1) = \begin{cases} (x_1^A)^2 - \left(\frac{x_1^A}{3}\right)^2 & \text{if } 0 \le x_1^A \le \bar{s}, \\ (x_1^A)^2 - \left(\frac{x_1^A}{2}\right)^2 & \text{if } 2\bar{s} - 4 \le x_1^A \le 0, \\ (x_1^A)^2 + (\bar{s} - 2)(\bar{s} - 2 - x_1^A) & \text{if } -\bar{s} \le x_1^A \le 2\bar{s} - 4, \end{cases}$$
$$C^A(\bar{s}, x_1^A, 0) = \begin{cases} (x_1^A)^2 + (-\bar{s})(-\bar{s} - x_1^A) & \text{if } -\bar{s} \le x_1^A \le \bar{s}. \end{cases}$$

  And so we conclude that for all $x_1^A \in [-\bar{s}, \bar{s}]$, $C^A(\bar{s}, x_1^A, 1) \ge 0 = C_*^A(\bar{s})$ and $C^A(-\bar{s}, x_1^A, 0) \ge \frac{3\bar{s}^2}{4} = C_*^A(-\bar{s})$. Hence, the intuitive criterion does not rule out any such equilibrium (e.g., the equilibrium specified above).

Case 2: $(e^A, e^B) = (1, -1)$. Here is a full specification of a WPBE for this case. Dealer A bids $\left(b_{-\bar{s}}^A, b_{\bar{s}}^A\right) = \left(\frac{7\bar{s}^2}{16}, 0\right)$. If he loses, the game ends. Therefore, suppose henceforth that he wins. If

$s = \bar{s}$, dealer A sets $x_1^A = 0$ and

$$x_2^A = \begin{cases} -\frac{x_1^A}{2} & \text{if } 0 \le x_1^A \le \bar{s}, \\ -\frac{x_1^A}{3} & \text{if } \max\{\frac{3\bar{s}}{2} - 3, -\frac{\bar{s}}{6}\} \le x_1^A < 0, \\ \bar{s} - 2 - x_1^A & \text{if } -\frac{\bar{s}}{6} \le x_1^A < \frac{3\bar{s}}{2} - 3, \\ -\frac{x_1^A}{2} - \frac{\bar{s}}{4} & \text{if } \max\{-\bar{s}, \frac{5\bar{s}}{2} - 4\} \le x_1^A < -\frac{\bar{s}}{6}, \\ \bar{s} - 2 - x_1^A & \text{if } -\bar{s} \le x_1^A < \min\{\frac{5\bar{s}}{2} - 4, -\frac{\bar{s}}{6}\}. \end{cases} \tag{7}$$

If $s = -\bar{s}$, dealer A sets $x_1^A = -\frac{\bar{s}}{4}$ and

$$x_2^A = \begin{cases} -\bar{s} - x_1^A & \text{if } -\bar{s} \le x_1^A \le \bar{s}. \end{cases} \tag{8}$$

Dealer B sets

$$x_2^B = \begin{cases} 0 & \text{if } x_1^A \ge 0, \\ -\frac{x_1^A}{3} & \text{if } \max\{\frac{3\bar{s}}{2} - 3, -\frac{\bar{s}}{6}\} \le x_1^A < 0, \\ -\frac{\bar{s}}{2} + 1 & \text{if } -\frac{\bar{s}}{6} \le x_1^A < \frac{3\bar{s}}{2} - 3, \\ \frac{\bar{s}}{2} & \text{if } -\bar{s} \le x_1^A < -\frac{\bar{s}}{6}. \end{cases} \tag{9}$$

Dealers' beliefs prior to bidding are $\mu_0^A = \mu_0^B = \phi$. Dealer B's beliefs prior to second-period trading are

$$\mu_2^B = \begin{cases} 1 & \text{if } -\frac{\bar{s}}{6} \le x_1^A \le \bar{s}, \\ 0 & \text{if } -\bar{s} \le x_1^A < -\frac{\bar{s}}{6}. \end{cases}$$

We claim that the specified strategies and beliefs satisfy the solution concept described in Section 3.1 and moreover that anything else also satisfying the solution concept must feature the same on-path behavior. The argument consists of three parts.

Part (i): We check the consistency of dealer B's beliefs. Given the specified strategy for dealer A, Bayes' rule requires only that

$$\mu_2^B = \begin{cases} 1 & \text{if } x_1^A = 0, \\ 0 & \text{if } x_1^A = -\frac{\bar{s}}{4}. \end{cases}$$

This is indeed consistent with the specified beliefs.

Part (ii): Given the specified beliefs, we check that the solution concept uniquely pins down the specified strategies. We proceed by backward induction:

- *Period-2 reaction functions.* Dealer A's trading costs are $x_1^A x_1^A + (x_1^A + x_2^A + x_2^B)x_2^A$. For $s \in \{-\bar{s}, \bar{s}\}$, dealer A best responds with $x_2^A = \left[-\frac{x_1^A + x_2^B}{2}\right]_{s-2-x_1^A}^{s-x_1^A}$. Dealer B's trading costs are $(x_1^A + x_2^A + x_2^B)x_2^B$. Dealer B best responds with $x_2^B = \left[-\frac{x_1^A + x_2^A}{2}\right]_0^2$.
- *Dealer B's period-2 action.* If $-\frac{\bar{s}}{6} \le x_1^A \le \bar{s}$ so that $\mu_2^B = 1$, then x_2^B is pinned down by the

intersection of $x_2^A = \left[-\frac{x_1^A+x_2^B}{2}\right]_{\bar{s}-2-x_1^A}^{\bar{s}-x_1^A}$ and $x_2^B = \left[-\frac{x_1^A+x_2^A}{2}\right]_0^2$, so that we indeed have

$$x_2^B = \begin{cases} 0 & \text{if } x_1^A \geq 0, \\ -\frac{x_1^A}{3} & \text{if } \max\{\frac{3\bar{s}}{2}-3, -\frac{\bar{s}}{6}\} \leq x_1^A < 0, \\ -\frac{\bar{s}}{2}+1 & \text{if } -\frac{\bar{s}}{6} \leq x_1^A < \frac{3\bar{s}}{2}-3. \end{cases}$$

If $-\bar{s} \leq x_1^A < -\frac{\bar{s}}{6}$ so that $\mu_2^B = 0$, then x_2^B is pinned down by the intersection of $x_2^A = \left[-\frac{x_1^A+x_2^B}{2}\right]_{-\bar{s}-2-x_1^A}^{-\bar{s}-x_1^A}$ and $x_2^B = \left[-\frac{x_1^A+x_2^A}{2}\right]_0^2$, so that we indeed have $x_2^B = \frac{\bar{s}}{2}$. Together, these two cases verify (9).

- *Dealer A's period-2 action.* If $s = \bar{s}$, then x_2^A is pinned down by the intersection of $x_2^A = \left[-\frac{x_1^A+x_2^B}{2}\right]_{\bar{s}-2-x_1^A}^{\bar{s}-x_1^A}$ and (9), which verifies (7).

  If $s = -\bar{s}$, then x_2^A is pinned down by the intersection of $x_2^A = \left[-\frac{x_1^A+x_2^B}{2}\right]_{-\bar{s}-2-x_1^A}^{-\bar{s}-x_1^A}$ and (9), which verifies (8).

- *Dealer A's period-1 action.* Dealer A's trading costs are $x_1^A x_1^A + (x_1^A + x_2^A + x_2^B)x_2^A$. If $s = \bar{s}$, then we can plug in (7) and (9) to express dealer A's trading costs as a function of x_1^A:

$$\begin{cases} x_1^A x_1^A - \left(\frac{x_1^A}{2}\right)^2 & \text{if } 0 \leq x_1^A \leq \bar{s}, \\ x_1^A x_1^A - \left(\frac{x_1^A}{3}\right)^2 & \text{if } \max\{\frac{3\bar{s}}{2}-3, -\frac{\bar{s}}{6}\} \leq x_1^A < 0, \\ x_1^A x_1^A + (\frac{\bar{s}}{2}-1)(\bar{s}-2-x_1^A) & \text{if } -\frac{\bar{s}}{6} \leq x_1^A < \frac{3\bar{s}}{2}-3, \\ x_1^A x_1^A - \left(\frac{x_1^A}{2}+\frac{\bar{s}}{4}\right)^2 & \text{if } \max\{-\bar{s}, \frac{5\bar{s}}{2}-4\} \leq x_1^A < -\frac{\bar{s}}{6}, \\ x_1^A x_1^A + (\frac{3\bar{s}}{2}-2)(\bar{s}-2-x_1^A) & \text{if } -\bar{s} \leq x_1^A < \min\{\frac{5\bar{s}}{2}-4, -\frac{\bar{s}}{6}\}, \end{cases}$$

  which is indeed minimized by $x_1^A = 0$.

  Alternatively, if $s = -\bar{s}$, then we can plug in (8) and (9) to express dealer A's trading costs as a function of x_1^A:

$$\begin{cases} x_1^A x_1^A + (-\bar{s})(-\bar{s}-x_1^A) & \text{if } x_1^A \geq 0, \\ x_1^A x_1^A + (-\bar{s}-\frac{x_1^A}{3})(-\bar{s}-x_1^A) & \text{if } \max\{\frac{3\bar{s}}{2}-3, -\frac{\bar{s}}{6}\} \leq x_1^A < 0, \\ x_1^A x_1^A + (-\frac{3\bar{s}}{2}+1)(-\bar{s}-x_1^A) & \text{if } -\frac{\bar{s}}{6} \leq x_1^A < \frac{3\bar{s}}{2}-3, \\ x_1^A x_1^A + (-\frac{\bar{s}}{2})(-\bar{s}-x_1^A) & \text{if } -\bar{s} \leq x_1^A < -\frac{\bar{s}}{6}, \end{cases}$$

  which is indeed minimized by $x_1^A = -\frac{\bar{s}}{4}$.

- *Dealer A's bid.* Plugging in the trading behavior derived above, we have the following. If $s = \bar{s}$, then dealer A's continuation utility is c if he wins and 0 if he loses. If $s = -\bar{s}$, then dealer A's continuation utility is $c - \frac{7\bar{s}^2}{16}$ if he wins and 0 if he loses. Based on the solution concept described in Section 3.1, dealer A must therefore bid $\left(b^A_{-\bar{s}}, b^A_{\bar{s}}\right) = \left(\frac{7\bar{s}^2}{16}, 0\right)$.

Part (iii): We check that the equilibrium specified above satisfies the restrictions on beliefs described in Section 3.1. And we also show that any equilibrium satisfying those conditions must feature on-path behavior coinciding with that of the equilibrium specified above. Let $\tilde{\mu}^B_2$ be candidate beliefs. One requirement is that $\tilde{\mu}^B_2$ must have the threshold structure described in the text.

- We start by establishing an upper bound for the location of the jump. Suppose, by way of contradiction, that we have an equilibrium with beliefs such that $\tilde{\mu}^B_2(-\frac{\bar{s}}{6} + \varepsilon) = 0$ for some $\varepsilon > 0$. By the threshold structure of $\tilde{\mu}^B_2$ and because beliefs must be correct on path, we can choose $\varepsilon > 0$ arbitrarily small (and in particular less than $\frac{2\bar{s}}{3}$). The threshold structure of $\tilde{\mu}^B_2$, together with the fact that beliefs must be correct on path, also implies that on path, dealer A sets $x^A_1 > -\frac{\bar{s}}{6} + \varepsilon$ when $s = \bar{s}$. Arguments similar to those given above imply that we would subsequently have
$$(x^A_2, x^B_2) = \begin{cases} (-\frac{x^A_1}{2}, 0) & \text{if } 0 \leq x^A_1 \leq \bar{s}, \\ (-\frac{x^A_1}{3}, -\frac{x^A_1}{3}) & \text{if } -\frac{\bar{s}}{6} + \varepsilon < x^A_1 \leq 0, \end{cases}$$
leading to total trading costs for dealer A of
$$\begin{cases} \frac{3}{4}(x^A_1)^2 & \text{if } 0 \leq x^A_1 \leq \bar{s}, \\ \frac{8}{9}(x^A_1)^2 & \text{if } -\frac{\bar{s}}{6} + \varepsilon < x^A_1 \leq 0, \end{cases}$$
which are bounded below by 0. On the other hand, suppose that dealer A deviated to set $x^A_1 = -\frac{\bar{s}}{6} + \varepsilon$ when $s = \bar{s}$. Arguments similar to those given above imply that we would subsequently have $(x^A_2, x^B_2) = (-\frac{\bar{s}}{6} - \frac{\varepsilon}{2}, \frac{\bar{s}}{2})$, leading to trading costs of $-\frac{\varepsilon}{4}(2\bar{s} - 3\varepsilon) < 0$. This constitutes a profitable deviation, contradicting the putative equilibrium.

- We next establish a corresponding lower bound. Suppose, by way of contradiction, that we have an equilibrium with beliefs such that $\tilde{\mu}^B_2(-\frac{\bar{s}}{4}) = 1$. By the threshold structure of $\tilde{\mu}^B_2$ and because beliefs must be correct on path, this means that on path, dealer A sets $x^A_1 < -\frac{\bar{s}}{4}$ when $s = -\bar{s}$. Arguments similar to those given above imply that we would subsequently have $x^A_2 = -\bar{s} - x^A_1$ and $x^B_2 = \frac{\bar{s}}{2}$, leading to trading costs for dealer A of $C^A_*(-\bar{s}) = (x^A_1)^2 + \frac{\bar{s}}{2}(\bar{s} + x^A_1) = \frac{7\bar{s}^2}{16} + (\frac{\bar{s}}{4} + x^A_1)^2$. On the other hand, suppose that dealer A deviated to set $x^A_1 = -\frac{\bar{s}}{4}$ when $s = -\bar{s}$ and that—contrary to the putative beliefs—this were to induce a belief $\mu^B_2 = 0$. Arguments similar to those given above imply that we would subsequently have $(x^A_2, x^B_2) = (-\frac{3\bar{s}}{4}, \frac{\bar{s}}{2})$, leading to trading costs of $C^A(-\bar{s}, -\frac{\bar{s}}{4}, 0) = \frac{7\bar{s}^2}{16} < C^A_*(-\bar{s})$.

  Given that $\tilde{\mu}^B_2(-\frac{\bar{s}}{4}) = 1$, we can use arguments similar to those given above to show that on-path behavior coincides with that in the equilibrium specified above when $s = \bar{s}$. In particular, dealer A's equilibrium trading costs are as above: $C^A_*(\bar{s}) = 0$. On the other hand, suppose that dealer A deviated to set $x^A_1 = -\frac{\bar{s}}{4}$ when $s = \bar{s}$ and that this were to induce a belief

$\mu_2^B \in [0,1]$. Arguments similar to those given above imply that we would subsequently have $(x_2^A, x_2^B) = \left(\frac{(3\mu_2^B-2)\bar{s}}{4(4-\mu_2^B)}, \frac{(8-7\mu_2^B)\bar{s}}{4(4-\mu_2^B)}\right)$, leading to trading costs of $C^A(\bar{s}, -\frac{\bar{s}}{4}, \mu_2^B) = \frac{(3-2\mu_2^B)(1+\mu_2^B)\bar{s}^2}{4(4-\mu_2^B)^2}$. This is minimized at $\mu_2^B = 0$, and hence $\min_{\mu_2^B \in [0,1]} C^A(\bar{s}, -\frac{\bar{s}}{4}, \mu_2^B) = \frac{3\bar{s}^2}{64} > C_*^A(\bar{s})$. The putative equilibrium therefore fails our intuitive criterion test.

- Finally, we turn to the remaining possibilities. Suppose we have an equilibrium with beliefs such that both $\tilde{\mu}_2^B(-\frac{\bar{s}}{6} + \varepsilon) = 1$ for all $\varepsilon > 0$ and $\tilde{\mu}_2^B(-\frac{\bar{s}}{4}) = 0$. Using arguments similar to those given above, we can show that any such equilibrium induces the same on-path behavior as the equilibrium specified above. In particular, dealer A's equilibrium trading costs are as above: $C_*^A(\bar{s}) = 0$ and $C_*^A(-\bar{s}) = \frac{7\bar{s}^2}{16}$. We can also use arguments similar to those given above to compute

$$C^A(\bar{s}, x_1^A, 1) = \begin{cases} (x_1^A)^2 - \left(\frac{x_1^A}{2}\right)^2 & \text{if } 0 \le x_1^A \le \bar{s}, \\ (x_1^A)^2 - \left(\frac{x_1^A}{3}\right)^2 & \text{if } \frac{3\bar{s}}{2} - 3 \le x_1^A \le 0, \\ (x_1^A)^2 + (\frac{\bar{s}}{2} - 1)(\bar{s} - 2 - x_1^A) & \text{if } -\bar{s} \le x_1^A \le \frac{3\bar{s}}{2} - 3, \end{cases}$$
$$C^A(\bar{s}, x_1^A, 0) = \begin{cases} (x_1^A)^2 + (-\frac{\bar{s}}{2})(-\bar{s} - x_1^A) & \text{if } -\bar{s} \le x_1^A \le \bar{s}. \end{cases}$$

And so we conclude that for all $x_1^A \in [-\bar{s}, \bar{s}]$, $C^A(\bar{s}, x_1^A, 1) \ge 0 = C_*^A(\bar{s})$ and $C^A(-\bar{s}, x_1^A, 0) \ge \frac{7\bar{s}^2}{16} = C_*^A(-\bar{s})$. Hence, the intuitive criterion does not rule out any such equilibrium (e.g., the equilibrium specified above).

This completes the proof. □

A.2 Proof of Lemma 2

To establish Lemma 2, we in fact prove the stronger result stated in Lemma A2.

Lemma A2. *In a subgame following an RFQ that contacts $M = 2$ dealers, induces dealer beliefs ϕ, and entails reserve prices $\bar{b}_{-\bar{s}} \ge \left[1 - \frac{23\phi(16+7\phi)}{4(24-\phi)^2}\right]\bar{s}^2$ and $\bar{b}_{\bar{s}} \ge \left[1 - \frac{23(1-\phi)(23-7\phi)}{4(23+\phi)^2}\right]\bar{s}^2$, the unique on-path equilibrium behavior is as follows. Dealer A bids*

$$(b_{-\bar{s}}^A, b_{\bar{s}}^A) = \begin{cases} \left(\left[1 - \frac{23\phi(16+7\phi)}{4(24-\phi)^2}\right]\bar{s}^2, -23\left[\frac{3(1-\phi)}{2(24-\phi)}\right]^2\bar{s}^2\right) & \text{if } (e^A, e^B) = (1,1), \\ \left(\frac{7\bar{s}^2}{16}, \frac{\bar{s}^2}{4}\right) & \text{if } (e^A, e^B) = (1,-1), \\ \left(\frac{\bar{s}^2}{4}, \frac{7\bar{s}^2}{16}\right) & \text{if } (e^A, e^B) = (-1,1), \\ \left(-23\left[\frac{3\phi}{2(23+\phi)}\right]^2\bar{s}^2, \left[1 - \frac{23(1-\phi)(23-7\phi)}{4(23+\phi)^2}\right]\bar{s}^2\right) & \text{if } (e^A, e^B) = (-1,-1). \end{cases}$$

If dealer A wins, the on-market trades are

$$(x_1^A, x_2^A, x_1^B, x_2^B) = \begin{cases} \left(\frac{7(1-\phi)\bar{s}}{2(24-\phi)}, \frac{3(1-\phi)\bar{s}}{2(24-\phi)}, -\frac{8(1-\phi)\bar{s}}{24-\phi}, \frac{3(1-\phi)\bar{s}}{2(24-\phi)}\right) & \text{if } (s, e^A, e^B) = (\bar{s}, 1, 1), \\ \left(-\frac{(16+7\phi)\bar{s}}{2(24-\phi)}, -\frac{(32-9\phi)\bar{s}}{2(24-\phi)}, -\frac{8(1-\phi)\bar{s}}{24-\phi}, \frac{8(1-\phi)\bar{s}}{24-\phi}\right) & \text{if } (s, e^A, e^B) = (-\bar{s}, 1, 1), \\ (0, 0, 0, 0) & \text{if } (s, e^A, e^B) = (\bar{s}, 1, -1), \\ \left(-\frac{\bar{s}}{4}, -\frac{3\bar{s}}{4}, 0, \frac{\bar{s}}{2}\right) & \text{if } (s, e^A, e^B) = (-\bar{s}, 1, -1), \\ \left(\frac{\bar{s}}{4}, \frac{3\bar{s}}{4}, 0, -\frac{\bar{s}}{2}\right) & \text{if } (s, e^A, e^B) = (\bar{s}, -1, 1), \\ (0, 0, 0, 0) & \text{if } (s, e^A, e^B) = (-\bar{s}, -1, 1), \\ \left(\frac{(23-7\phi)\bar{s}}{2(23+\phi)}, \frac{(23+9\phi)\bar{s}}{2(23+\phi)}, \frac{8\phi\bar{s}}{23+\phi}, -\frac{8\phi\bar{s}}{23+\phi}\right) & \text{if } (s, e^A, e^B) = (\bar{s}, -1, -1), \\ \left(-\frac{7\phi\bar{s}}{2(23+\phi)}, -\frac{3\phi\bar{s}}{2(23+\phi)}, \frac{8\phi\bar{s}}{23+\phi}, -\frac{3\phi\bar{s}}{2(23+\phi)}\right) & \text{if } (s, e^A, e^B) = (-\bar{s}, -1, -1). \end{cases}$$

Dealer B's bids and the on-market trades if dealer B wins are specified symmetrically.

Lemma 2 follows from Lemma A2 for reasons discussed in the main text. For an RFQ that contacts two dealers to obtain execution with probability one, the client's reserve must be high enough to never set the price in equilibrium. Mathematically, this means $\bar{b}_{-\bar{s}} \geq \left[1 - \frac{23\phi(16+7\phi)}{4(24-\phi)^2}\right]\bar{s}^2$ and $\bar{b}_{\bar{s}} \geq \left[1 - \frac{23(1-\phi)(23-7\phi)}{4(23+\phi)^2}\right]\bar{s}^2$, where ϕ denotes the probability of $s = \bar{s}$ induced by the RFQ. These inequalities are in particular satisfied by $\bar{b} = (\bar{s}^2, \bar{s}^2)$.

It therefore follows from Lemma A2 that the expected cost of procurement achieved by an RFQ that contacts $M = 2$ dealers, induces a belief that ϕ is the probability of $s = \bar{s}$, and leads to execution with probability one can be computed as

$$\begin{aligned} \hat{c}_2(\phi) = &\left[\frac{(1-\phi)(4-\phi)\bar{s}^2}{4} + \frac{\phi(1-\phi)^3\bar{s}^2}{4(24-\phi)^2}\right]\psi\big[1-(1-\psi)(1-\rho)\big] \\ &+ \left[\frac{\phi(3+\phi)\bar{s}^2}{4} + \frac{(1-\phi)\phi^3\bar{s}^2}{4(23+\phi)^2}\right](1-\psi)\big[1-\psi(1-\rho)\big] + \frac{7\bar{s}^2}{8}\psi(1-\psi)(1-\rho). \end{aligned} \tag{10}$$

Indeed, with probability $\phi\psi[1-(1-\psi)(1-\rho)]$, we have $(s, e^A, e^B) = (\bar{s}, 1, 1)$, and the auction's clearing price is $-23\left[\frac{3(1-\phi)}{2(24-\phi)}\right]^2\bar{s}^2$. With probability $(1-\phi)\psi[1-(1-\psi)(1-\rho)]$, we have $(s, e^A, e^B) = (-\bar{s}, 1, 1)$, and the auction's clearing price is $\left[1 - \frac{23\phi(16+7\phi)}{4(24-\phi)^2}\right]\bar{s}^2$. With probability $\phi(1-\psi)[1-\psi(1-\rho)]$, we have $(s, e^A, e^B) = (\bar{s}, -1, -1)$, and the auction's clearing price is $\left[1 - \frac{23(1-\phi)(23-7\phi)}{4(23+\phi)^2}\right]\bar{s}^2$. With probability $(1-\phi)(1-\psi)[1-\psi(1-\rho)]$, we have $(s, e^A, e^B) = (-\bar{s}, -1, -1)$, and the auction's clearing price is $-23\left[\frac{3\phi}{2(23+\phi)}\right]^2\bar{s}^2$. With the remaining probability $2\psi(1-\psi)(1-\rho)$, the auction's clearing price is $\max\{\frac{7\bar{s}^2}{16}, \frac{\bar{s}^2}{4}\} = \frac{7\bar{s}^2}{16}$. It follows that the expected procurement cost is as in (10).

The claims made by Lemma 2 about $\hat{c}_2(\phi)$ follow readily from (10). Indeed, we can see that it is differentiable and we can moreover compute

$$\hat{c}_2''(\phi) = \frac{529(624-95\phi)\bar{s}^2}{2(24-\phi)^4}\psi[1-(1-\psi)(1-\rho)] + \frac{529(529+95\phi)\bar{s}^2}{2(23+\phi)^4}(1-\psi)[1-\psi(1-\rho)],$$

which is positive on the domain $\phi \in [0,1]$. We can also compute $\hat{c}_2(\frac{1}{2}) = \frac{1933\bar{s}^2}{4418}\psi[1-(1-\psi)(1-\rho)] + \frac{1933\bar{s}^2}{4418}(1-\psi)[1-\psi(1-\rho)] + \frac{7\bar{s}^2}{8}\psi(1-\psi)(1-\rho) < \frac{3\bar{s}^2}{4} = \hat{c}_1$.

Proof of Lemma A2. Because both dealers observe the entire vector (e^A, e^B), the four possible realizations of that vector can be analyzed separately. Below, we analyze the cases of $(1,1)$ and $(1,-1)$; the remaining cases can be handled symmetrically (i.e., by flipping signs and exchanging the roles of ϕ and $1-\phi$). Within each case, we focus on the event in which dealer A wins; the events in which dealer B wins can be handled symmetrically.

Case 1: $(e^A, e^B) = (1,1)$. Here is a full specification of a WPBE for this case. Dealer A bids $\left(b^A_{-\bar{s}}, b^A_{\bar{s}}\right) = \left(\left[1 - \frac{23\phi(16+7\phi)}{4(24-\phi)^2}\right]\bar{s}^2, -23\left[\frac{3(1-\phi)}{2(24-\phi)}\right]^2\bar{s}^2\right)$. Henceforth, suppose that dealer A wins. If $s = \bar{s}$, dealer A sets $x^A_1 = \frac{7(1-\phi)\bar{s}}{2(24-\phi)}$ and

$$x^A_2 = \begin{cases} -\frac{x^A_1}{2} & \text{if } 0 \le x^A_1 \le \bar{s} \text{ and } \frac{x^A_1}{2} \le x^B_1 \le \bar{s}, \\ -\frac{x^A_1 + x^B_1}{3} & \text{if } 0 \le x^A_1 \le \bar{s} \text{ and } -\bar{s} \le x^B_1 \le \frac{x^A_1}{2}, \\ -\frac{x^A_1}{2} & \text{if } \max\{-\bar{s}, 2\bar{s}-4\} \le x^A_1 < 0 \text{ and } -\bar{s} \le x^B_1 \le \bar{s}, \\ \bar{s} - 2 - x^A_1 & \text{if } -\bar{s} \le x^A_1 < 2\bar{s} - 4 \text{ and } -\bar{s} \le x^B_1 \le \bar{s}. \end{cases} \tag{11}$$

If $s = -\bar{s}$, dealer A sets $x^A_1 = -\frac{(16+7\phi)\bar{s}}{2(24-\phi)}$ and

$$x^A_2 = \begin{cases} -\bar{s} - x^A_1 & \text{if } -\bar{s} \le x^A_1 \le \bar{s} \text{ and } -\bar{s} \le x^B_1 \le \bar{s}. \end{cases} \tag{12}$$

Dealer B sets $x^B_1 = -\frac{8(1-\phi)\bar{s}}{24-\phi}$ and

$$x^B_2 = \begin{cases} -x^B_1 & \text{if } 0 \le x^A_1 \le \bar{s} \text{ and } \frac{x^A_1}{2} \le x^B_1 \le \bar{s}, \\ -\frac{x^A_1 + x^B_1}{3} & \text{if } 0 \le x^A_1 \le \bar{s} \text{ and } -\bar{s} \le x^B_1 \le \frac{x^A_1}{2}, \\ -x^B_1 & \text{if } -\bar{s} \le x^A_1 < 0 \text{ and } -\bar{s} \le x^B_1 \le \bar{s}. \end{cases} \tag{13}$$

Dealers' beliefs prior to bidding are $\mu^A_0 = \mu^B_0 = \phi$. Dealer B's beliefs prior to first-period trading are $\mu^B_1 = \phi$. Dealer B's beliefs prior to second-period trading are

$$\mu^B_2 = \begin{cases} 1 & \text{if } 0 \le x^A_1 \le \bar{s}, \\ 0 & \text{if } -\bar{s} \le x^A_1 < 0. \end{cases}$$

We claim that the specified strategies and beliefs satisfy the solution concept described in Section 3.1 and moreover that anything else also satisfying the solution concept must feature the same on-path behavior. The argument consists of three parts.

Part (i): We check the consistency of dealer B's beliefs. First, consider dealer B's beliefs (conditional on losing) at the point just after having observed the auction's outcome. Given symmetry of the specified bidding strategies and the fact that the auction's tie-breaking rule does not depend on the realized s, the auction's outcome is uninformative so that posterior beliefs must equal the prior. Thus, we indeed have $\mu_1^B = \phi$. Second, consider dealer B's beliefs (conditional on losing) at the point just after having observed the first trading period's outcome. Given the specified strategy for dealer A, Bayes' rule requires only that

$$\mu_2^B = \begin{cases} 1 & \text{if } x_1^A = \frac{7(1-\phi)\bar{s}}{2(24-\phi)}, \\ 0 & \text{if } x_1^A = -\frac{(16+7\phi)\bar{s}}{2(24-\phi)}. \end{cases}$$

This is indeed consistent with the specified beliefs.

Part (ii): Given the specified beliefs, we check that the solution concept uniquely pins down the specified strategies. We proceed by backward induction:

- *Period-2 reaction functions.* Dealer A's trading costs are $(x_1^A + x_1^B)x_1^A + (x_1^A + x_1^B + x_2^A + x_2^B)x_2^A$. For $s \in \{-\bar{s}, \bar{s}\}$, dealer A best responds with $x_2^A = \left[-\frac{x_1^A + x_1^B + x_2^B}{2}\right]_{s-2-x_1^A}^{s-x_1^A}$. Dealer B's trading costs are $(x_1^A + x_1^B)x_1^B + (x_1^A + x_1^B + x_2^A + x_2^B)x_2^B$. Dealer B best responds with $x_2^B = \left[-\frac{x_1^A + x_1^B + x_2^A}{2}\right]_{-2-x_1^B}^{-x_1^B}$.

- *Dealer B's period-2 action.* If $0 \le x_1^A \le \bar{s}$ so that $\mu_2^B = 1$, then x_2^B is pinned down by the intersection of $x_2^A = \left[-\frac{x_1^A + x_1^B + x_2^B}{2}\right]_{\bar{s}-2-x_1^A}^{\bar{s}-x_1^A}$ and $x_2^B = \left[-\frac{x_1^A + x_1^B + x_2^A}{2}\right]_{-2-x_1^B}^{-x_1^B}$, so that we indeed have:

  $$x_2^B = \begin{cases} -x_1^B & \text{if } 0 \le x_1^A \le \bar{s} \text{ and } \frac{x_1^A}{2} \le x_1^B \le \bar{s}, \\ -\frac{x_1^A + x_1^B}{3} & \text{if } 0 \le x_1^A \le \bar{s} \text{ and } -\bar{s} \le x_1^B \le \frac{x_1^A}{2}. \end{cases}$$

  If $-\bar{s} \le x_1^A < 0$ so that $\mu_2^B = 0$, then x_2^B is pinned down by the intersection of $x_2^A = \left[-\frac{x_1^A + x_1^B + x_2^B}{2}\right]_{-\bar{s}-2-x_1^A}^{-\bar{s}-x_1^A}$ and $x_2^B = \left[-\frac{x_1^A + x_1^B + x_2^A}{2}\right]_{-2-x_1^B}^{-x_1^B}$, so that we indeed have:

  $$x_2^B = \begin{cases} -x_1^B & \text{if } -\bar{s} \le x_1^A < 0 \text{ and } -\bar{s} \le x_1^B \le \bar{s}. \end{cases}$$

  Together, these two cases verify (13).

- *Dealer A's period-2 action.* If $s = \bar{s}$, then x_2^A is pinned down by the intersection of $x_2^A = \left[-\frac{x_1^A + x_1^B + x_2^B}{2}\right]_{\bar{s}-2-x_1^A}^{\bar{s}-x_1^A}$ and (13), which verifies (11). If $s = -\bar{s}$, then x_2^A is pinned down by the intersection of $x_2^A = \left[-\frac{x_1^A + x_1^B + x_2^B}{2}\right]_{-\bar{s}-2-x_1^A}^{-\bar{s}-x_1^A}$ and (13), which verifies (12).

- *Period-1 actions.* Dealer A's trading costs are $(x_1^A + x_1^B)x_1^A + (x_1^A + x_1^B + x_2^A + x_2^B)x_2^A$. If $s = \bar{s}$, then we can plug in (11) and (13) to express dealer A's trading costs as a function of (x_1^A, x_1^B):

$$\begin{cases} (x_1^A+x_1^B)x_1^A - \left(\frac{x_1^A}{2}\right)^2 & \text{if } 0 \le x_1^A \le \bar{s} \text{ and } \frac{x_1^A}{2} \le x_1^B \le \bar{s}, \\ (x_1^A+x_1^B)x_1^A - \left(\frac{x_1^A+x_1^B}{3}\right)^2 & \text{if } 0 \le x_1^A \le \bar{s} \text{ and } -\bar{s} \le x_1^B \le \frac{x_1^A}{2}, \\ (x_1^A+x_1^B)x_1^A - \left(\frac{x_1^A}{2}\right)^2 & \text{if } \max\{-\bar{s}, 2\bar{s}-4\} \le x_1^A < 0 \text{ and } -\bar{s} \le x_1^B \le \bar{s}, \\ (x_1^A+x_1^B)x_1^A + (\bar{s}-2)(\bar{s}-2-x_1^A) & \text{if } -\bar{s} \le x_1^A < 2\bar{s}-4 \text{ and } -\bar{s} \le x_1^B \le \bar{s}. \end{cases}$$

Optimizing, we can express $x_1^A(\bar{s})$ in terms of x_1^B:

$$x_1^A(\bar{s}) = \begin{cases} -\frac{7x_1^B}{16} & \text{if } -\bar{s} \le x_1^B \le 0, \\ -\frac{2x_1^B}{3} & \text{if } 0 \le x_1^B \le \min\{-3\bar{s}+6, \bar{s}\}, \\ \frac{\bar{s}}{2} - 1 - \frac{x_1^B}{2} & \text{if } -3\bar{s}+6 \le x_1^B \le \bar{s}. \end{cases} \tag{14}$$

Alternatively, if $s = -\bar{s}$, then we can plug in (12) and (13) to express dealer A's trading costs as a function of (x_1^A, x_1^B):

$$\begin{cases} (x_1^A+x_1^B)x_1^A + (-\bar{s})(-\bar{s}-x_1^A) & \text{if } 0 \le x_1^A \le \bar{s} \text{ and } \frac{x_1^A}{2} \le x_1^B \le \bar{s}, \\ (x_1^A+x_1^B)x_1^A + (-\bar{s} - \frac{x_1^A}{3} + \frac{2x_1^B}{3})(-\bar{s}-x_1^A) & \text{if } 0 \le x_1^A \le \bar{s} \text{ and } -\bar{s} \le x_1^B \le \frac{x_1^A}{2}, \\ (x_1^A+x_1^B)x_1^A + (-\bar{s})(-\bar{s}-x_1^A) & \text{if } -\bar{s} \le x_1^A < 0 \text{ and } -\bar{s} \le x_1^B \le \bar{s}. \end{cases}$$

Optimizing, we can express $x_1^A(-\bar{s})$ in terms of x_1^B:

$$x_1^A(-\bar{s}) = \begin{cases} -\frac{\bar{s}}{2} - \frac{x_1^B}{2} & \text{if } -\bar{s} \le x_1^B \le \bar{s}. \end{cases} \tag{15}$$

The derivation of dealer B's equilibrium period-1 action consists of two parts.

– First, we show that there is a unique equilibrium in which $x_1^B \in [-\bar{s}, 0]$. Suppose $s = \bar{s}$. If x_1^A is a best response to some $\hat{x}_1^B \in [-\bar{s}, 0]$, then by (14), we have $x_1^A = -\frac{7\hat{x}_1^B}{16}$, so that $x_1^A \in [0, \frac{7\bar{s}}{16}]$. Dealer B's trading costs are $(x_1^A + x_1^B)x_1^B + (x_1^A + x_1^B + x_2^A + x_2^B)x_2^B$. We can then plug in (11) and (13) to express dealer B's trading costs, conditional on $s = \bar{s}$, as a function of $(x_1^B, \hat{x}_1^B)$:

$$C_{\bar{s}}^B(x_1^B, \hat{x}_1^B) = \begin{cases} (-\frac{7\hat{x}_1^B}{16} + x_1^B)x_1^B + (-\frac{7\hat{x}_1^B}{32})(-x_1^B) & \text{if } -\frac{7\hat{x}_1^B}{32} \le x_1^B \le \bar{s}, \\ (-\frac{7\hat{x}_1^B}{16} + x_1^B)x_1^B - \left(-\frac{7\hat{x}_1^B}{48} + \frac{x_1^B}{3}\right)^2 & \text{if } -\bar{s} \le x_1^B \le -\frac{7\hat{x}_1^B}{32}. \end{cases}$$

Suppose $s = -\bar{s}$. If x_1^A is a best response to the same $\hat{x}_1^B \in [-\bar{s}, 0]$, then by (15), we have $x_1^A = -\frac{\bar{s}}{2} - \frac{\hat{x}_1^B}{2}$, so that $x_1^A \in [-\frac{\bar{s}}{2}, 0]$. We can then plug in (12) and (13) to express dealer B's trading costs, conditional on $s = -\bar{s}$, as a function of $(x_1^B, \hat{x}_1^B)$:

$$C_{-\bar{s}}^B(x_1^B, \hat{x}_1^B) = \begin{cases} (-\frac{\bar{s}}{2} - \frac{\hat{x}_1^B}{2} + x_1^B)x_1^B + (-\bar{s})(-x_1^B) & \text{if } -\bar{s} \le x_1^B \le \bar{s}. \end{cases}$$

Because ϕ represents the probability of $\bar{s}$, dealer B's unconditional expected trading costs as a function of $(x_1^B, \hat{x}_1^B)$ are

$$\phi C_{\bar{s}}^B(x_1^B, \hat{x}_1^B) + (1-\phi) C_{-\bar{s}}^B(x_1^B, \hat{x}_1^B),$$

which is minimized by $x_1^B = -\frac{9(1-\phi)\bar{s}}{4(9-\phi)} + \frac{(72-23\phi)\hat{x}_1^B}{32(9-\phi)}$. Equilibrium occurs when this optimal value of x_1^B coincides with $\hat{x}_1^B$. This obtains at $x_1^B = -\frac{8(1-\phi)\bar{s}}{24-\phi}$, which is indeed an element of $[-\bar{s}, 0]$.

– Second, we show that there is no equilibrium in which $x_1^B > 0$. Suppose $s = \bar{s}$. Let $\hat{x}_1^B > 0$. Assume that x_1^A best responds to $\hat{x}_1^B$; it can be derived from (14). We can then plug in (11) and (13) to express dealer B's trading costs, conditional on $s = \bar{s}$, as a function of $(x_1^B, \hat{x}_1^B)$:

$$C_{\bar{s}}^B(\hat{x}_1^B, \hat{x}_1^B) = \begin{cases} \frac{2}{3}(\hat{x}_1^B)^2 & \text{if } 0 \leq \hat{x}_1^B \leq \min\{-3\bar{s}+6, \bar{s}\}, \\ \frac{(-\bar{s}+2+\hat{x}_1^B)\hat{x}_1^B}{2} & \text{if } -3\bar{s}+6 \leq \hat{x}_1^B \leq \bar{s}. \end{cases}$$

Similarly, suppose $s = -\bar{s}$. Let $\hat{x}_1^B > 0$ as above. Assume that x_1^A best responds to $\hat{x}_1^B$; it can be derived from (15). We can then plug in (12) and (13) to express dealer B's trading costs, conditional on $s = \bar{s}$, as a function of $(x_1^B, \hat{x}_1^B)$:

$$C_{-\bar{s}}^B(\hat{x}_1^B, \hat{x}_1^B) = \begin{cases} \frac{(\bar{s}+\hat{x}_1^B)\hat{x}_1^B}{2} & \text{if } 0 \leq \hat{x}_1^B \leq \bar{s}. \end{cases}$$

Because ϕ represents the probability of $\bar{s}$, dealer B's unconditional expected trading costs from setting $x_1^B = \hat{x}_1^B$ would be $\phi C_{\bar{s}}^B(\hat{x}_1^B, \hat{x}_1^B) + (1-\phi) C_{-\bar{s}}^B(\hat{x}_1^B, \hat{x}_1^B)$, which by the above expressions is strictly positive. But this cannot correspond to an equilibrium because dealer B could do strictly better by deviating: setting $x_1^B = x_2^B = 0$ guarantees trading costs of zero.

In conclusion, we have derived dealer B's period-1 action as $x_1^B = -\frac{8(1-\phi)\bar{s}}{24-\phi}$. Plugging this into (14) and (15), we indeed obtain $x_1^A(\bar{s}) = \frac{7(1-\phi)\bar{s}}{2(24-\phi)}$ and $x_1^A(-\bar{s}) = -\frac{(16+7\phi)\bar{s}}{2(24-\phi)}$, respectively. Finally, given these period-1 actions, (11)–(13) imply that the following period-2 actions will occur on the equilibrium path

$$(x_2^A, x_2^B) = \begin{cases} \left(\frac{3(1-\phi)\bar{s}}{2(24-\phi)}, \frac{3(1-\phi)\bar{s}}{2(24-\phi)}\right) & \text{if } s = \bar{s}, \\ \left(-\frac{(32-9\phi)\bar{s}}{2(24-\phi)}, \frac{8(1-\phi)\bar{s}}{24-\phi}\right) & \text{if } s = -\bar{s}. \end{cases}$$

- *Bids.* Plugging in the trading behavior derived above, we have the following. If $s = \bar{s}$ and if dealer A wins, then dealer A's continuation utility is $c + \frac{18(1-\phi)^2\bar{s}^2}{(24-\phi)^2}$, and dealer B's continuation utility is $-\frac{135(1-\phi)^2\bar{s}^2}{4(24-\phi)^2}$. By symmetry, if dealer B wins, then dealer A's continuation utility would be $-\frac{135(1-\phi)^2\bar{s}^2}{4(24-\phi)^2}$.

  If $s = -\bar{s}$ and if dealer A wins, then dealer A's continuation utility is $c - \frac{(64+5\phi)(32-9\phi)\bar{s}^2}{4(24-\phi)^2}$, and dealer B's continuation utility is $\frac{4(1-\phi)(16+7\phi)\bar{s}^2}{(24-\phi)^2}$. By symmetry, if dealer B wins, then dealer A's continuation utility would be $\frac{4(1-\phi)(16+7\phi)\bar{s}^2}{(24-\phi)^2}$.

Because the dealer sets a non-binding reserve, the event in which neither dealer wins is not relevant. Based on the solution concept described in Section 3.1, dealer A must therefore bid

$$(b^A_{-\bar{s}}, b^A_{\bar{s}}) = \left(\underbrace{\frac{(64+5\phi)(32-9\phi)\bar{s}^2}{4(24-\phi)^2} + \frac{4(1-\phi)(16+7\phi)\bar{s}^2}{(24-\phi)^2}}_{=\left[1-\frac{23\phi(16+7\phi)}{4(24-\phi)^2}\right]\bar{s}^2}, \underbrace{-\frac{18(1-\phi)^2\bar{s}^2}{(24-\phi)^2} - \frac{135(1-\phi)^2\bar{s}^2}{4(24-\phi)^2}}_{=-23\left[\frac{3(1-\phi)}{2(24-\phi)}\right]^2\bar{s}^2} \right).$$

Part (iii): We check that the equilibrium specified above satisfies the restrictions on beliefs described in Section 3.1. And we also show that any equilibrium satisfying those conditions must feature on-path behavior coinciding with that of the equilibrium specified above. Let $\tilde{\mu}^B_2$ be candidate beliefs. One requirement is that $\tilde{\mu}^B_2$ must have the threshold structure described in the text.

- We start by establishing an upper bound for the location of the jump. Suppose, by way of contradiction, that we have an equilibrium with beliefs such that $\tilde{\mu}^B_2\big(\frac{7(1-\phi)\bar{s}}{2(24-\phi)}\big) = 0$. By the threshold structure of $\tilde{\mu}^B_2$ and because beliefs must be correct on path, this means that on path, dealer A sets $x^A_1 = x^* > \frac{7(1-\phi)\bar{s}}{2(24-\phi)}$ when $s = \bar{s}$. Arguments similar to those given above imply that we have $x^A_1 = \frac{7\phi x^* - (9+7\phi)\bar{s}}{27+5\phi}$ when $s = -\bar{s}$ and $x^B_1 = -\frac{14\phi x^* + 9(1-\phi)\bar{s}}{27+5\phi}$. And when $s = \bar{s}$, we subsequently have $x^A_2 = x^B_2 = -\frac{x^* + x^B_1}{3}$, leading to total trading costs for dealer A of
$$(x^* + x^B_1)x^* - \frac{1}{9}(x^* + x^B_1)^2, \tag{16}$$
evaluated at $x^B_1 = -\frac{14\phi x^* + 9(1-\phi)\bar{s}}{27+5\phi}$. Now consider two cases:
  - First, suppose there exists a $\varepsilon > 0$ such that $\tilde{\mu}^B_2(x^* - \varepsilon) = 1$. In that case, it follows from equation (16) that x^* can be a locally optimal choice for dealer A only if $x^A_1 = -\frac{7x^B_1}{16}$. Given that $x^B_1 = -\frac{14\phi x^* + 9(1-\phi)\bar{s}}{27+5\phi}$, this requires $x^* = \frac{7(1-\phi)\bar{s}}{2(24-\phi)}$, a contradiction.
  - Second, suppose that for all $\varepsilon > 0$, $\tilde{\mu}^B_2(x^* - \varepsilon) = 0$. In that case, if dealer A deviates to $x^A_1 = x^* - \varepsilon$ when $s = \bar{s}$, then we subsequently have $x^B_2 = -x^B_1$ and $x^A_2 = -\frac{x^A_1}{2}$, leading to total trading costs for dealer A that, for small ε, are well approximated by
$$(x^* + x^B_1)x^* - \frac{1}{4}(x^*)^2, \tag{17}$$
evaluated at $x^B_1 = -\frac{14\phi x^* + 9(1-\phi)\bar{s}}{27+5\phi}$. Comparing (17) to (16) and using $x^B_1 = -\frac{14\phi x^* + 9(1-\phi)\bar{s}}{27+5\phi}$, we see that this is a profitable deviation if $x^* > \frac{6(1-\phi)\bar{s}}{45-\phi}$, which is implied by $x^* > \frac{7(1-\phi)\bar{s}}{2(24-\phi)}$.
- We then strengthen this upper bound. We know from above that $\tilde{\mu}^B_2\big(\frac{7(1-\phi)\bar{s}}{2(24-\phi)}\big) = 1$. Arguments

similar to those given above imply that our only candidate equilibrium entails $x_1^A = \frac{7(1-\phi)\bar{s}}{2(24-\phi)}$ when $s = \bar{s}$, $x_1^A = -\frac{(16+7\phi)\bar{s}}{2(24-\phi)}$ when $s = -\bar{s}$, and $x_1^B = -\frac{8(1-\phi)\bar{s}}{24-\phi}$. And when $s = \bar{s}$, we subsequently have $x_2^A = x_2^B = -\frac{x_1^A + x_1^B}{3}$, leading to total trading costs for dealer A of

$$-\frac{18(1-\phi)^2\bar{s}^2}{(24-\phi)^2}. \tag{18}$$

Now suppose, by way of contradiction, that we have an equilibrium with beliefs such that $\tilde{\mu}_2^B\big(\frac{2(1-\phi)(8-\sqrt{10})\bar{s}}{3(24-\phi)} + \varepsilon\big) = 0$ for some $\varepsilon > 0$. Suppose then that dealer A deviated to set $x_1^A = \frac{2(1-\phi)(8-\sqrt{10})\bar{s}}{3(24-\phi)} + \varepsilon$ when $s = \bar{s}$. Arguments similar to those given above imply that we would subsequently have $x_2^B = -\frac{x_1^B}{2}$ and $x_2^A = -\frac{x_1^A}{2}$, leading to total trading costs for dealer A of

$$(x_1^A + x_1^B)x_1^A - \frac{(x_1^A)^2}{4}. \tag{19}$$

Comparing (19) to (18), we see that this is a profitable deviation for sufficiently small $\varepsilon > 0$.

- We next establish a corresponding lower bound. Suppose, by way of contradiction, that we have an equilibrium with beliefs such that $\tilde{\mu}_2^B\big(-\frac{(16+7\phi)\bar{s}}{2(24-\phi)}\big) = 1$. By the threshold structure of $\tilde{\mu}_2^B$ and because beliefs must be correct on path, this means that on path, dealer A sets $x_1^A = x^* < -\frac{(16+7\phi)\bar{s}}{2(24-\phi)}$ when $s = -\bar{s}$. Arguments similar to those given above imply that we have $x_1^A = \frac{7(1-\phi)(\bar{s}+x^*)}{32-9\phi}$ when $s = \bar{s}$ and $x_1^B = -\frac{16(1-\phi)(\bar{s}+x^*)}{32-9\phi}$. When $s = -\bar{s}$, we subsequently have $x_2^A = -\bar{s} - x_1^A$ and $x_2^B = -x_1^B$, leading to equilibrium trading costs for dealer A, denoted $C_*^A(-\bar{s})$, equal to

$$\left(x_1^A - \frac{16(1-\phi)(\bar{s}+x^*)}{32-9\phi}\right)x_1^A + \bar{s}(\bar{s} + x_1^A) \tag{20}$$

evaluated at $x_1^A = x^*$. Alternatively, when $s = \bar{s}$, we subsequently have $x_2^A = x_2^B = -\frac{x_1^A + x_1^B}{3}$. Plugging in, dealer A's equilibrium trading costs are

$$C_*^A(\bar{s}) = -\frac{72(1-\phi)^2(\bar{s}+x^*)^2}{(32-9\phi)^2}. \tag{21}$$

Now define $\hat{x} := -\frac{(16+7\phi)\bar{s} - 16(1-\phi)x^*}{2(32-9\phi)}$. Note that $\hat{x} > x^*$, which follows from $x^* < -\frac{(16+7\phi)\bar{s}}{2(24-\phi)}$. Suppose that dealer A deviated to set $x_1^A = \hat{x}$ when $s = -\bar{s}$ and that—perhaps contrary to the putative beliefs—this were to induce a belief $\mu_2^B = 0$. Arguments similar to those given above imply that we would subsequently have $x_2^A = -\bar{s} - x_1^A$ and $x_2^B = -x_1^B$, leading to trading costs for dealer A, denoted $C^A(-\bar{s}, \hat{x}, 0)$, equal to (20) evaluated at $x_1^A = \hat{x}$. Note that this choice optimizes (20). On the other hand, because $x^* \neq \hat{x}$, x^* does not optimize (20). We therefore conclude $C^A(-\bar{s}, \hat{x}, 0) < C_*^A(-\bar{s})$.

On the other hand, suppose that dealer A deviated to set $x_1^A = \hat{x}$ when $s = \bar{s}$ and that this were to induce a belief $\mu_2^B \in [0, 1]$. Arguments similar to those given above imply that we would subsequently have

$$x_2^A = \frac{[16(1-\phi)x^* + (48-25\phi)\bar{s}]\mu_2^B}{2(32-9\phi)(4-\mu_2^B)} \qquad x_2^B = \frac{[(48-25\phi)\bar{s} + 16(1-\phi)x^*](4-3\mu_2^B)}{2(32-9\phi)(4-\mu_2^B)},$$

leading to trading costs of

$$C^A\left(\bar{s},\hat{x},\mu_2^B\right) = \frac{2[(48-25\phi)\bar{s}+16(1-\phi)x^*]\left([2(16+7\phi)-(16+7\phi)\mu_2^B-4(1-\phi)(\mu_2^B)^2]\bar{s}-4(1-\phi)x^*[(\mu_2^B)^2-4\mu_2^B+8]\right)}{(32-9\phi)^2(4-\mu_2^B)^2}$$

This is minimized at $\mu_2^B = 0$, and hence

$$\min_{\mu_2^B\in[0,1]} C^A\left(\bar{s},\hat{x},\mu_2^B\right) = \frac{[(48-25\phi)\bar{s}+16(1-\phi)x^*][(16+7\phi)\bar{s}-16(1-\phi)x^*]}{4(32-9\phi)^2}. \tag{22}$$

Comparing (22) to (21), it can be shown that $C_*^A(\bar{s}) < \min_{\mu_2^B\in[0,1]} C^A\left(\bar{s},\hat{x},\mu_2^B\right)$. The putative equilibrium therefore fails our intuitive criterion test.

- Finally, we turn to the remaining possibilities. Suppose we have an equilibrium with beliefs such that both $\tilde{\mu}_2^B\left(\frac{2(1-\phi)(8-\sqrt{10})\bar{s}}{3(24-\phi)}+\varepsilon\right) = 1$ for all $\varepsilon > 0$ and $\tilde{\mu}_2^B\left(-\frac{(16+7\phi)\bar{s}}{2(24-\phi)}\right) = 0$. Using arguments similar to those given above, we can show that any such equilibrium induces the same on-path behavior as the equilibrium specified above. In particular, dealer A's equilibrium trading costs are as above: $C_*^A(\bar{s}) = -\frac{18(1-\phi)^2\bar{s}^2}{24-\phi)^2}$ and $C_*^A(-\bar{s}) = \frac{(62+5\phi)(32-9\phi)\bar{s}^2}{4(24-\phi)^2}$. We can also use arguments similar to those given above to compute

$$C^A(\bar{s},x_1^A,1) = \begin{cases} \left(x_1^A - \frac{8(1-\phi)\bar{s}}{24-\phi}\right)(x_1^A) - \frac{1}{9}\left(x_1^A - \frac{8(1-\phi)\bar{s}}{24-\phi}\right)^2 & \text{if } -\frac{16(1-\phi)\bar{s}}{24-\phi} \le x_1^A \le \bar{s}, \\ \left(x_1^A - \frac{8(1-\phi)\bar{s}}{24-\phi}\right)(x_1^A) - \left(\frac{x_1^A}{2}\right)^2 & \text{if } 2\bar{s}-4 \le x_1^A \le -\frac{16(1-\phi)\bar{s}}{24-\phi}, \\ \left(x_1^A - \frac{8(1-\phi)\bar{s}}{24-\phi}\right)(x_1^A) + (\bar{s}-2)(\bar{s}-2-x_1^A) & \text{if } -\bar{s} \le x_1^A \le 2\bar{s}-4, \end{cases}$$

$$C^A(\bar{s},x_1^A,0) = \left\{\left(x_1^A - \frac{8(1-\phi)\bar{s}}{24-\phi}\right)(x_1^A) + \bar{s}(\bar{s}+x_1^A) \quad \text{if } -\bar{s} \le x_1^A \le \bar{s}.\right.$$

And so we conclude that for all $x_1^A \in [-\bar{s},\bar{s}]$, $C^A(\bar{s},x_1^A,1) \ge C_*^A(\bar{s})$ and $C^A(-\bar{s},x_1^A,0) \ge C_*^A(-\bar{s})$. Hence, the intuitive criterion does not rule out any such equilibrium (e.g., the equilibrium specified above).

Case 2: $(e^A,e^B) = (1,-1)$. Here is a full specification of a WPBE for this case. Dealer A bids

$\left(b^A_{-\bar{s}}, b^A_{\bar{s}}\right) = \left(\frac{7\bar{s}^2}{16}, \frac{\bar{s}^2}{4}\right)$. Henceforth, suppose that dealer A wins. If $s = \bar{s}$, dealer A sets $x_1^A = 0$ and

$$x_2^A = \begin{cases} -\frac{x_1^A + x_1^B}{3} & \text{if } \max\{\frac{3\bar{s}}{2} + \frac{x_1^B}{2} - 3, -\frac{\bar{s}}{6}\} \leq x_1^A \leq \bar{s} \text{ and } \frac{x_1^A}{2} \leq x_1^B \leq \bar{s}, \\ \bar{s} - 2 - x_1^A & \text{if } -\frac{\bar{s}}{6} \leq x_1^A \leq \frac{3\bar{s}}{2} + \frac{x_1^B}{2} - 3 \text{ and } \max\{-\bar{s}, \bar{s} - 2\} \leq x_1^B \leq \bar{s}, \\ -\frac{x_1^A}{2} & \max\{-\frac{\bar{s}}{6}, 2\bar{s} - 4\} \leq x_1^A \leq \bar{s} \text{ and } -\bar{s} \leq x_1^B \leq \frac{x_1^A}{2}, \\ \bar{s} - 2 - x_1^A & \text{if } -\frac{\bar{s}}{6} \leq x_1^A \leq 2\bar{s} - 4 \text{ and } -\bar{s} \leq x_1^B \leq \bar{s} - 2, \\ -\frac{\bar{s}}{4} - \frac{x_1^A}{2} - \frac{x_1^B}{4} & \text{if } \max\{-\bar{s}, \frac{5\bar{s} + x_1^B}{2} - 4\} \leq x_1^A < -\frac{\bar{s}}{6} \text{ and } -\bar{s} \leq x_1^B \leq \bar{s}, \\ \bar{s} - 2 - x_1^A & \text{if } -\bar{s} \leq x_1^A < \min\{\frac{5\bar{s} + x_1^B}{2} - 4, -\frac{\bar{s}}{6}\} \text{ and } -\bar{s} \leq x_1^B \leq \bar{s}. \end{cases} \tag{23}$$

If $s = -\bar{s}$, dealer A sets $x_1^A = -\frac{\bar{s}}{4}$ and

$$x_2^A = \begin{cases} -\bar{s} - x_1^A & \text{if } -\bar{s} \leq x_1^A \leq \bar{s} \text{ and } -\bar{s} \leq x_1^B \leq \bar{s}. \end{cases} \tag{24}$$

Dealer B sets $x_1^B = 0$ and

$$x_2^B = \begin{cases} -\frac{x_1^A + x_1^B}{3} & \text{if } \max\{\frac{3\bar{s}}{2} + \frac{x_1^B}{2} - 3, -\frac{\bar{s}}{6}\} \leq x_1^A \leq \bar{s} \text{ and } \frac{x_1^A}{2} \leq x_1^B \leq \bar{s}, \\ -\frac{\bar{s}}{2} + 1 - \frac{x_1^B}{2} & \text{if } -\frac{\bar{s}}{6} \leq x_1^A \leq \frac{3\bar{s}}{2} + \frac{x_1^B}{2} - 3 \text{ and } \max\{-\bar{s}, \bar{s} - 2\} \leq x_1^B \leq \bar{s}, \\ -x_1^B & \text{if } \max\{-\frac{\bar{s}}{6}, 2\bar{s} - 4\} \leq x_1^A \leq \bar{s} \text{ and } -\bar{s} \leq x_1^B \leq \frac{x_1^A}{2}, \\ -x_1^B & \text{if } -\frac{\bar{s}}{6} \leq x_1^A \leq 2\bar{s} - 4 \text{ and } -\bar{s} \leq x_1^B \leq \bar{s} - 2, \\ \frac{\bar{s}}{2} - \frac{x_1^B}{2} & \text{if } -\bar{s} \leq x_1^A < -\frac{\bar{s}}{6} \text{ and } -\bar{s} \leq x_1^B \leq \bar{s}. \end{cases} \tag{25}$$

Dealers' beliefs prior to bidding are $\mu_0^A = \mu_0^B = \phi$. Dealer B's beliefs prior to first-period trading are $\mu_1^B = 1$. Dealer B's beliefs prior to second-period trading are

$$\mu_2^B = \begin{cases} 1 & \text{if } -\frac{\bar{s}}{6} \leq x_1^A \leq \bar{s}, \\ 0 & \text{if } -\bar{s} \leq x_1^A < -\frac{\bar{s}}{6}. \end{cases}$$

We claim that the specified strategies and beliefs satisfy the solution concept described in Section 3.1 and moreover that anything else also satisfying the solution concept must feature the same on-path behavior. The argument consists of three parts.

Part (i): We check the consistency of dealer B's beliefs. First, consider dealer B's beliefs (conditional on losing) at the point just after having observed the auction's outcome. In this case, dealer A bids $\left(b^A_{-\bar{s}}, b^A_{\bar{s}}\right) = \left(\frac{7\bar{s}^2}{16}, \frac{\bar{s}^2}{4}\right)$, while dealer B bids $\left(b^B_{-\bar{s}}, b^B_{\bar{s}}\right) = \left(\frac{\bar{s}^2}{4}, \frac{7\bar{s}^2}{16}\right)$. It follows that dealer B loses iff $s = \bar{s}$. Thus, we indeed have $\mu_1^B = 1$. Second, consider dealer B's beliefs (conditional on losing) at the point just after having observed the first trading period's outcome. Given the specified strategy for dealer A, Bayes' rule requires only that

$$\mu_2^B = \begin{cases} 1 & \text{if } x_1^A = 0, \\ 0 & \text{if } x_1^A = -\frac{\bar{s}}{4}. \end{cases}$$

This is indeed consistent with the specified beliefs.

Part (ii): Given the specified beliefs, we check that the solution concept uniquely pins down the specified strategies. We proceed by backward induction:

- *Period-2 reaction functions.* Dealer A's trading costs are $x_1^A x_1^A + (x_1^A + x_2^A + x_2^B)x_2^A$. For $s \in \{-\bar{s}, \bar{s}\}$, dealer A best responds with $x_2^A = \left[-\frac{x_1^A + x_1^B + x_2^B}{2}\right]_{s-2-x_1^A}^{s-x_1^A}$. Dealer B's trading costs are $(x_1^A + x_2^A + x_2^B)x_2^B$. Dealer B best responds with $x_2^B = \left[-\frac{x_1^A + x_1^B + x_2^A}{2}\right]_{-x_1^B}^{2-x_1^B}$.

- *Dealer B's period-2 action.* If $-\frac{\bar{s}}{6} \le x_1^A \le \bar{s}$ so that $\mu_2^B = 1$, then x_2^B is pinned down by the intersection of $x_2^A = \left[-\frac{x_1^A + x_1^B + x_2^B}{2}\right]_{\bar{s}-2-x_1^A}^{\bar{s}-x_1^A}$ and $x_2^B = \left[-\frac{x_1^A + x_1^B + x_2^A}{2}\right]_{-x_1^B}^{2-x_1^B}$, so that we indeed have:

$$x_2^B = \begin{cases} -\frac{x_1^A + x_1^B}{3} & \text{if } \max\{\frac{3\bar{s}}{2} + \frac{x_1^B}{2} - 3, -\frac{\bar{s}}{6}\} \le x_1^A \le \bar{s} \text{ and } \frac{x_1^A}{2} \le x_1^B \le \bar{s}, \\ -\frac{\bar{s}}{2} + 1 - \frac{x_1^B}{2} & \text{if } -\frac{\bar{s}}{6} \le x_1^A \le \frac{3\bar{s}}{2} + \frac{x_1^B}{2} - 3 \text{ and } \max\{-\bar{s}, \bar{s} - 2\} \le x_1^B \le \bar{s}, \\ -x_1^B & \text{if } \max\{-\frac{\bar{s}}{6}, 2\bar{s} - 4\} \le x_1^A \le \bar{s} \text{ and } -\bar{s} \le x_1^B \le \frac{x_1^A}{2}, \\ -x_1^B & \text{if } -\frac{\bar{s}}{6} \le x_1^A \le 2\bar{s} - 4 \text{ and } -\bar{s} \le x_1^B \le \bar{s} - 2. \end{cases}$$

  If $-\bar{s} \le x_1^A < -\frac{\bar{s}}{6}$ so that $\mu_2^B = 0$, then x_2^B is pinned down by the intersection of $x_2^A = \left[-\frac{x_1^A + x_1^B + x_2^B}{2}\right]_{-\bar{s}-2-x_1^A}^{-\bar{s}-x_1^A}$ and $x_2^B = \left[-\frac{x_1^A + x_1^B + x_2^A}{2}\right]_{-x_1^B}^{2-x_1^B}$, so that we indeed have:

$$x_2^B = \begin{cases} \frac{\bar{s}}{2} - \frac{x_1^B}{2} & \text{if } -\bar{s} \le x_1^A < -\frac{\bar{s}}{6} \text{ and } -\bar{s} \le x_1^B \le \bar{s}. \end{cases}$$

  Together, these two cases verify (25).

- *Dealer A's period-2 action.* If $s = \bar{s}$, then x_2^A is pinned down by the intersection of $x_2^A = \left[-\frac{x_1^A + x_1^B + x_2^B}{2}\right]_{\bar{s}-2-x_1^A}^{\bar{s}-x_1^A}$ and (25), which verifies (23). If $s = -\bar{s}$, then x_2^A is pinned down by the intersection of $x_2^A = \left[-\frac{x_1^A + x_1^B + x_2^B}{2}\right]_{-\bar{s}-2-x_1^A}^{-\bar{s}-x_1^A}$ and (25), which verifies (24).

- *Period-1 actions.* Dealer A's trading costs are $(x_1^A + x_1^B)x_1^A + (x_1^A + x_1^B + x_2^A + x_2^B)x_2^A$. If $s = \bar{s}$, then we can plug in (23) and (25) to express dealer A's trading costs as a function of (x_1^A, x_1^B):

$$\begin{cases} (x_1^A+x_1^B)x_1^A - \left(\frac{x_1^A+x_1^B}{3}\right)^2 & \text{if } \max\{\frac{3\bar{s}}{2}+\frac{x_1^B}{2}-3, -\frac{\bar{s}}{6}\} \le x_1^A \le \bar{s} \text{ and } \frac{x_1^A}{2} \le x_1^B \le \bar{s}, \\ (x_1^A+x_1^B)x_1^A + (\frac{\bar{s}}{2}-1+\frac{x_1^B}{2})(\bar{s}-2-x_1^A) & \text{if } -\frac{\bar{s}}{6} \le x_1^A \le \frac{3\bar{s}}{2}+\frac{x_1^B}{2}-3 \text{ and } \max\{-\bar{s}, \bar{s}-2\} \le x_1^B \le \bar{s}, \\ (x_1^A+x_1^B)x_1^A - \left(\frac{x_1^A}{2}\right)^2 & \max\{-\frac{\bar{s}}{6}, 2\bar{s}-4\} \le x_1^A \le \bar{s} \text{ and } -\bar{s} \le x_1^B \le \frac{x_1^A}{2}, \\ (x_1^A+x_1^B)x_1^A + (\bar{s}-2)(\bar{s}-2-x_1^A) & \text{if } -\frac{\bar{s}}{6} \le x_1^A \le 2\bar{s}-4 \text{ and } -\bar{s} \le x_1^B \le \bar{s}-2, \\ (x_1^A+x_1^B)x_1^A - \left(\frac{\bar{s}}{4}+\frac{x_1^A}{2}+\frac{x_1^B}{4}\right)^2 & \text{if } \max\{-\bar{s}, \frac{5\bar{s}+x_1^B}{2}-4\} \le x_1^A < -\frac{\bar{s}}{6} \text{ and } -\bar{s} \le x_1^B \le \bar{s}, \\ (x_1^A+x_1^B)x_1^A + (\frac{3\bar{s}}{2}-2+\frac{x_1^B}{2})(\bar{s}-2-x_1^A) & \text{if } -\bar{s} \le x_1^A < \min\{\frac{5\bar{s}+x_1^B}{2}-4, -\frac{\bar{s}}{6}\} \text{ and } -\bar{s} \le x_1^B \le \bar{s}. \end{cases}$$

Rather than derive dealer A's function of best responses to x_1^B, we instead derive two auxiliary functions. First, we derive dealer A's best response among first-period trades satisfying $-\frac{\bar{s}}{6} \le x_1^A \le \bar{s}$:

$$x_1^A(x_1^B) = \begin{cases} -\frac{2x_1^B}{3} & \text{if } -\bar{s} \le x_1^B \le 0, \\ -\frac{7x_1^B}{16} & \text{if } 0 \le x_1^B \le \min\{\frac{8\bar{s}}{21}, \frac{8(2-\bar{s})}{5}\}, \\ -\frac{\bar{s}}{6} & \text{if } \frac{8\bar{s}}{21} \le x_1^B \le -\frac{10\bar{s}}{3}+6, \\ \frac{3\bar{s}}{2}+\frac{x_1^B}{2}-3 & \text{if } \max\{\frac{8(2-\bar{s})}{5}, -\frac{10\bar{s}}{3}+6\} \le x_1^B \le \frac{5(2-\bar{s})}{3}, \\ \frac{\bar{s}}{4}-\frac{1}{2}-\frac{x_1^B}{4} & \text{if } \frac{5(2-\bar{s})}{3} \le x_1^B \le \frac{5\bar{s}}{3}-2, \\ -\frac{\bar{s}}{6} & \text{if } \max\{-\frac{10\bar{s}}{3}+6, \frac{5\bar{s}}{3}-2\} \le x_1^B \le \bar{s}. \end{cases} \tag{26}$$

Second, we derive dealer A's best response among first-period trades satisfying $x_1^A < -\frac{\bar{s}}{6}$. Because the domain is not compact, this function is not defined everywhere:

$$x_1^A(x_1^B) = \begin{cases} \text{undefined} & \text{if } -\bar{s} \le x_1^B \le \max\{\frac{2\bar{s}}{3}, \frac{11\bar{s}}{3}-4\}, \\ \frac{\bar{s}}{6}-\frac{x_1^B}{2} & \text{if } \frac{2\bar{s}}{3} < x_1^B \le -\frac{7\bar{s}}{3}+4, \\ \frac{3\bar{s}}{2}+\frac{x_1^B}{2}-3 & \text{if } \max\{-\frac{7\bar{s}}{3}+4, \frac{11\bar{s}}{3}-4\} < x_1^B \le \bar{s}. \end{cases} \tag{27}$$

Dealer B's trading costs are $(x_1^A+x_1^B)x_1^B + (x_1^A+x_1^B+x_2^A+x_2^B)x_2^B$. If $s=\bar{s}$, then we can plug in (23) and (25) to express dealer B's trading costs as a function of (x_1^A, x_1^B):

$$\begin{cases} (x_1^A+x_1^B)x_1^A - \left(\frac{x_1^A+x_1^B}{3}\right)^2 & \text{if } \max\{\frac{3\bar{s}}{2}+\frac{x_1^B}{2}-3, -\frac{\bar{s}}{6}\} \le x_1^A \le \bar{s} \text{ and } \frac{x_1^A}{2} \le x_1^B \le \bar{s}, \\ (x_1^A+x_1^B)x_1^B - \left(\frac{\bar{s}}{2}-1+\frac{x_1^B}{2}\right)^2 & \text{if } -\frac{\bar{s}}{6} \le x_1^A \le \frac{3\bar{s}}{2}+\frac{x_1^B}{2}-3 \text{ and } \max\{-\bar{s}, \bar{s}-2\} \le x_1^B \le \bar{s}, \\ (x_1^A+x_1^B)x_1^B + (\frac{x_1^A}{2})(-x_1^B) & \max\{-\frac{\bar{s}}{6}, 2\bar{s}-4\} \le x_1^A \le \bar{s} \text{ and } -\bar{s} \le x_1^B \le \frac{x_1^A}{2}, \\ (x_1^A+x_1^B)x_1^B + (\bar{s}-2)(-x_1^B) & \text{if } -\frac{\bar{s}}{6} \le x_1^A \le 2\bar{s}-4 \text{ and } -\bar{s} \le x_1^B \le \bar{s}-2, \\ (x_1^A+x_1^B)x_1^B + (\frac{\bar{s}}{4}+\frac{x_1^A}{2}+\frac{x_1^B}{4})(\frac{\bar{s}}{2}-\frac{x_1^B}{2}) & \text{if } \max\{-\bar{s}, \frac{5\bar{s}+x_1^B}{2}-4\} \le x_1^A < -\frac{\bar{s}}{6} \text{ and } -\bar{s} \le x_1^B \le \bar{s}, \\ (x_1^A+x_1^B)x_1^B + (\frac{3\bar{s}}{2}-2+\frac{x_1^B}{2})(\frac{\bar{s}}{2}-\frac{x_1^B}{2}) & \text{if } -\bar{s} \le x_1^A < \min\{\frac{5\bar{s}+x_1^B}{2}-4, -\frac{\bar{s}}{6}\} \text{ and } -\bar{s} \le x_1^B \le \bar{s}. \end{cases}$$

Because dealer B attaches probability one to $s=\bar{s}$ (i.e., $\mu_1^B=1$), dealer B selects x_1^B to minimize this objective. Optimizing, we can express x_1^B in terms of x_1^A:

$$x_1^B(x_1^A) = \begin{cases} \frac{\bar{s}-2}{3} - \frac{2x_1^A}{3} & \text{if } -\bar{s} \le x_1^A \le \max\{-\frac{\bar{s}}{6}, \frac{16(\bar{s}-2)}{13}\}, \\ -\frac{7x_1^A}{16} & \text{if } \max\{-\frac{\bar{s}}{6}, \frac{16(\bar{s}-2)}{13}\} \le x_1^A \le 0, \\ -\frac{x_1^A}{4} & \text{if } 0 \le x_1^A \le \bar{s}. \end{cases} \tag{28}$$

To determine equilibrium period-1 actions, we argue as follows. First, observe that (27) and (28) do not intersect. This implies that there is no equilibrium in which $x_1^A \in [-\bar{s}, -\frac{\bar{s}}{6})$ when $s = \bar{s}$. Second, observe that (26) and (28) have a unique intersection at $(x_1^A, x_1^B) = (0, 0)$. This is the unique candidate for an equilibrium involving a choice of $x_1^A \in [-\frac{\bar{s}}{6}, \bar{s}]$ when $s = \bar{s}$. To show that this is in fact an equilibrium, we simply need to additionally verify that no choice of $x_1^A \in [-\bar{s}, -\frac{\bar{s}}{6})$ yields smaller trading costs for dealer A than $x_1^A = 0$ when $s = \bar{s}$ and $x_1^B = 0$, which is easily shown.[35]

Having pinned down $x_1^B = 0$ and that $x_1^A = 0$ if $s = \bar{s}$, the final step is to derive what x_1^A would be if $s = -\bar{s}$. Using $x_1^B = 0$, as derived above, we can plug in (24) and (25) to express dealer A's trading costs as a function of x_1^A:

$$x_2^B = \begin{cases} x_1^A x_1^A + (-\bar{s})(-\bar{s} - x_1^A) & \text{if } 0 \le x_1^A \le \bar{s}, \\ x_1^A x_1^A + (-\bar{s} - \frac{x_1^A}{3})(-\bar{s} - x_1^A) & \text{if } \max\{\frac{3\bar{s}}{2} - 3, -\frac{\bar{s}}{6}\} \le x_1^A \le 0, \\ x_1^A x_1^A + (-\frac{3\bar{s}}{2} + 1)(-\bar{s} - x_1^A) & \text{if } -\frac{\bar{s}}{6} \le x_1^A \le \frac{3\bar{s}}{2} - 3, \\ x_1^A x_1^A + (-\frac{\bar{s}}{2})(-\bar{s} - x_1^A) & \text{if } -\bar{s} \le x_1^A < -\frac{\bar{s}}{6}, \end{cases}$$

which is indeed minimized by $x_1^A = -\frac{\bar{s}}{4}$.

- *Bids.* Plugging in the trading behavior derived above, we have the following. If $s = \bar{s}$ and if dealer A wins, then dealer A's continuation utility is c, and dealer B's continuation utility is 0. If $s = -\bar{s}$ and if dealer A wins, then dealer A's continuation utility is $c - \frac{7\bar{s}^2}{16}$, and dealer B's continuation utility is $\frac{\bar{s}^2}{4}$.

  By symmetry, if $s = \bar{s}$ and if dealer B wins, then dealer A's continuation utility is $\frac{\bar{s}^2}{4}$. If $s = -\bar{s}$ and if dealer B wins, then dealer A's continuation utility is 0. Because the dealer sets a non-binding reserve, the event in which neither dealer wins is not relevant. Based on the solution concept described in Section 3.1, dealer A must therefore bid

$$(b_{-\bar{s}}^A, b_{\bar{s}}^A) = \left(\frac{7\bar{s}^2}{16}, \frac{\bar{s}^2}{4}\right).$$

[35] Plugging $x_1^B = 0$ into the previously-derived expression for dealer A's trading costs when $s = \bar{s}$, we obtain that a choice of $x_1^A = 0$ yields trading costs of zero and that choices of $x_1^A \in [-\bar{s}, -\frac{\bar{s}}{6})$ yield the strictly positive trading costs

$$\begin{cases} \frac{1}{16}(2x_1^A - \bar{s})(6x_1^A + \bar{s}) & \text{if } \max\{-\bar{s}, \frac{5\bar{s}}{2} - 4\} \le x_1^A < -\frac{\bar{s}}{6}, \\ \frac{1}{16}(2x_1^A - \bar{s})(6x_1^A + \bar{s}) + \frac{1}{16}(5\bar{s} - 2x_1^A - 8)^2 & \text{if } -\bar{s} \le x_1^A < \min\{\frac{5\bar{s}}{2} - 4, -\frac{\bar{s}}{6}\}. \end{cases}$$

Part (iii): We check that the equilibrium specified above satisfies the restrictions on beliefs described in Section 3.1. And we also show that any equilibrium satisfying those conditions must feature on-path behavior coinciding with that of the equilibrium specified above. Let $\tilde{\mu}_2^B$ be candidate beliefs. One requirement is that $\tilde{\mu}_2^B$ must have the threshold structure described in the text.

- We start by establishing an upper bound for the location of the jump. Suppose, by way of contradiction, that we have an equilibrium with beliefs such that $\tilde{\mu}_2^B(0) = 0$. By the threshold structure of $\tilde{\mu}_2^B$ and because beliefs must be correct on path, this means that on path, dealer A sets $x_1^A = x^* > 0$ when $s = \bar{s}$. Arguments similar to those given above imply that we have $x_1^B = -\frac{x^*}{4}$, and subsequently, $x_2^A = -\frac{x^*}{2}$ and $x_2^B = -x_1^B$, leading to total trading costs for dealer A of
$$(x^* + x_1^B)x^* - \frac{1}{4}(x^*)^2, \tag{29}$$
evaluated at $x_1^B = -\frac{x^*}{4}$. Now consider two cases:
  - First, suppose there exists a $\varepsilon > 0$ such that $\tilde{\mu}_2^B(x^* - \varepsilon) = 1$. In that case, it follows from equation (29) that x^* can be a locally optimal choice for dealer A only if $x_1^A = -\frac{2x_1^B}{3}$. Given that $x_1^B = -\frac{x^*}{4}$, this requires $x^* = 0$, a contradiction.
  - Second, suppose that for all $\varepsilon > 0$, $\tilde{\mu}_2^B(x^* - \varepsilon) = 0$. In that case, if dealer A deviates to $x_1^A = x^* - \varepsilon$ when $s = \bar{s}$, then we subsequently have $x_2^B = \frac{\bar{s}}{2} - \frac{x_1^B}{2}$ and $x_2^A = -\frac{\bar{s}+2x_1^A+x_1^B}{4}$, leading to total trading costs for dealer A that, for small ε, are well approximated by
$$(x^* + x_1^B)x^* - \left(\frac{\bar{s}}{4} + \frac{7x^*}{16}\right)^2, \tag{30}$$
evaluated at $x_1^B = -\frac{x^*}{4}$. Comparing (30) to (29) and using $x^* > 0$, we see that this is a profitable deviation.

- We then strengthen this upper bound. We know from above that $\tilde{\mu}_2^B(0) = 1$. Arguments similar to those given above imply that our only candidate equilibrium entails that when $s = \bar{s}$, $x_1^A = 0$, $x_1^B = 0$, and subsequently, $x_2^A = x_2^B = 0$, leading to total trading costs for dealer A of 0. Now suppose, by way of contradiction, that we have an equilibrium with beliefs such that $\tilde{\mu}_2^B(-\frac{\bar{s}}{6} + \varepsilon) = 0$ for some $\varepsilon > 0$. By the threshold structure of $\tilde{\mu}_2^B$ and because beliefs must be correct on path, we can choose $\varepsilon > 0$ arbitrarily small (and in particular less than $\frac{2\bar{s}}{3}$). Suppose then that dealer A deviated to set $x_1^A = -\frac{\bar{s}}{6} + \varepsilon$ when $s = \bar{s}$. Arguments similar to those given above imply that we would subsequently have $x_2^B = \frac{\bar{s}}{2}$ and $x_2^A = -\frac{\bar{s}}{4} - \frac{x_1^A}{2}$, leading to total trading costs for dealer A of $-\frac{\varepsilon}{4}(2\bar{s} - 3\varepsilon)$, which is a profitable deviation for sufficiently small $\varepsilon > 0$.

- We next establish a corresponding lower bound. Suppose, by way of contradiction, that we have an equilibrium with beliefs such that $\tilde{\mu}_2^B(-\frac{\bar{s}}{4}) = 1$. The same arguments used in the proof of Lemma 1 can be used to show that the putative equilibrium fails our intuitive criterion test.

- Finally, we turn to the remaining possibilities. Suppose we have an equilibrium with beliefs such that both $\tilde{\mu}_2^B(-\frac{\bar{s}}{6}+\varepsilon)=1$ for all $\varepsilon>0$ and $\tilde{\mu}_2^B(-\frac{\bar{s}}{4})=0$. The same arguments used in the proof of Lemma 1 can be used to show that the intuitive criterion does not rule out any such equilibrium (e.g., the equilibrium specified above).

This completes the proof. $\square$

A.3 Proof of Proposition 3

Proof. Begin with an arbitrary RFQ policy that contacts one dealer with probabilities $(q_{s'})_{s'\in\{-\bar{s},\bar{s}\}}$ and two dealers with the complementary probabilities $(1-q_{s'})_{s'\in\{-\bar{s},\bar{s}\}}$. It suffices to show that the following RFQ policy does no worse: (i) define Σ as the singleton $\{\sigma_0\}$; (ii) for $s'\in\{-\bar{s},\bar{s}\}$, define the distribution $\pi_{s'}$ to attach probability $q_{s'}$ to $\big(\sigma_0,1,(\frac{3\bar{s}^2}{4},\frac{3\bar{s}^2}{4})\big)$ and complementary probability $1-q_{s'}$ to $\big(\sigma_0,2,(\bar{s}^2,\bar{s}^2)\big)$. Indeed:

- This new policy does no worse conditional on contacting a single dealer, since, by Lemma 1, it leads to execution with probability one and achieves the cost lower bound $\hat{c}_1$. Note that the information about s communicated by this new policy is a garbling of that communicated by the original policy. However, there is no role for information design when only one dealer is contacted, because $\hat{c}_1$ is constant in ϕ for the reasons discussed at the end of Section 3.2.

- This new policy also does no worse conditional on contacting two dealers—and may in fact do strictly better. By Lemma 2, this policy leads to execution with probability one and achieves the cost $\hat{c}_2$ evaluated at the belief $\frac{\phi_0(1-q_{\bar{s}})}{(1-\phi_0)(1-q_{-\bar{s}})+\phi_0(1-q_{\bar{s}})}$. As mentioned, the information about s communicated by this new policy is a garbling of that communicated by the original policy. Thus, by standard arguments from the Bayesian persuasion literature (e.g., Kamenica and Gentzkow, 2011), the new policy is guaranteed to reduce the client's cost if $\hat{c}_2(\phi)$ is convex, which by Lemma 2 is indeed the case.

It only remains to argue that such an RFQ policy achieves optimality. To see this, note that policies of this form are described by two numbers: $q_{-\bar{s}}\in[0,1]$ and $q_{\bar{s}}\in[0,1]$. Thus, the claim follows because (i) this class of policies is compact, and (ii) within this class of policies, the client's procurement cost is a continuous function of $(q_{-\bar{s}},q_{\bar{s}})$. $\square$

A.4 Proof of Proposition 4

Lemma A3. *If $\underline{\phi}>0$ then $\hat{c}_2(\underline{\phi})<\hat{c}_1$. Symmetrically, if $\overline{\phi}<1$ then $\hat{c}_2(\overline{\phi})<\hat{c}_1$.*

Proof of Lemma A3. Assuming $\underline{\phi}>0$, we can define $L(\cdot)$ as the following linear function:

$$L(\phi)=\hat{c}_1+\frac{\hat{c}_2(\underline{\phi})-\hat{c}_1}{\underline{\phi}}\phi.$$

We begin by showing that $L(\phi)\le\hat{c}_2(\phi)$ for all $\phi\in[0,1]$. We have assumed $\underline{\phi}>0$. On the one hand, if $\underline{\phi}\in(0,1)$, then $L(\phi)$ is tangent to $\hat{c}_2(\phi)$ at $\phi=\underline{\phi}$. Indeed, we have both $L(\underline{\phi})=\hat{c}_2(\underline{\phi})$

and also $L'(\underline{\phi}) = \frac{\hat{c}_2(\underline{\phi})-\hat{c}_1}{\underline{\phi}} = \hat{c}_2'(\underline{\phi})$, where the last equality is because by Definition 1, $\underline{\phi} \in (0,1)$ implies $\hat{c}_2(\underline{\phi}) - \underline{\phi}\hat{c}_2'(\underline{\phi}) = \hat{c}_1$. On the other hand, if $\underline{\phi} = 1$, then by analogous arguments we have both $L(\underline{\phi}) = \hat{c}_2(\underline{\phi})$ and $L'(\underline{\phi}) \geq \hat{c}_2'(\underline{\phi})$. In either case, convexity of $\hat{c}_2(\phi)$ implies that $L(\phi) \leq \hat{c}_2(\phi)$ for all $\phi \in [0,1]$.

In particular, we have $L(\frac{1}{2}) \leq \hat{c}_2(\frac{1}{2})$. According to Lemma 2, $\hat{c}_2(\frac{1}{2}) < \hat{c}_1$. Hence, $L(\frac{1}{2}) < \hat{c}_1$. Because $L(\cdot)$ is a linear function with vertical intercept at $\hat{c}_1$, the fact that it is strictly less than $\hat{c}_1$ at one point on the domain $(0,1]$ means that it is strictly less than $\hat{c}_1$ on the entire domain $(0,1]$. In particular, $\hat{c}_2(\underline{\phi}) = L(\underline{\phi}) < \hat{c}_1$. □

Proof of Proposition 4. *Proof of claim (i).* As noted in footnote 28, $\hat{c}_2(\phi) - \phi\hat{c}_2'(\phi)$ is strictly decreasing on the unit interval. Hence, the claim about $\underline{\phi}$—that it is weakly increasing in ψ—would follow from Definition 1 if we can establish that $\frac{d}{d\psi}\left[\hat{c}_2(\phi) - \phi\hat{c}_2'(\phi)\right] \geq 0$. Because this derivative is linear in $\bar{s}^2$, ψ, and ρ, it suffices to check the inequality at $\bar{s} = 1$ and at the four possibilities $(\psi, \rho) \in \{(0,-1),(0,1),(1,-1),(1,1)\}$, which yield

$$\frac{d}{d\psi}\left[\hat{c}_2(\phi) - \phi\hat{c}_2'(\phi)\right]_{|\psi=0,\rho=-1} = \frac{7}{4} - \left[1 - \frac{529\phi^2(26-3\phi)}{4(24-\phi)^3}\right] + \frac{69\phi^2(71\phi+529)}{4(23+\phi)^3}$$

$$\left.\begin{aligned}&\frac{d}{d\psi}\left[\hat{c}_2(\phi) - \phi\hat{c}_2'(\phi)\right]_{|\psi=0,\rho=1}\\ &\frac{d}{d\psi}\left[\hat{c}_2(\phi) - \phi\hat{c}_2'(\phi)\right]_{|\psi=1,\rho=1}\end{aligned}\right\} = \left[1 - \frac{529\phi^2(26-3\phi)}{4(24-\phi)^3}\right] + \frac{23\phi^2(71\phi+529)}{4(23+\phi)^3}$$

$$\frac{d}{d\psi}\left[\hat{c}_2(\phi) - \phi\hat{c}_2'(\phi)\right]_{|\psi=1,\rho=-1} = \left[3 - \frac{1587\phi^2(26-3\phi)}{4(24-\phi)^3}\right] - \frac{23\phi^2(71\phi+529)}{4(23+\phi)^3} - \frac{7}{4}$$

It is easily checked that each of these are positive for all $\phi \in [0,1]$, as required. By symmetric arguments, $\overline{\phi}$ is also weakly increasing in ψ.

Proof of claim (ii). First, suppose that both $\underline{\phi} > 0$ and $\overline{\phi} < 1$. By Definition 1, $\underline{\phi} > 0$ implies $\hat{c}_2(\underline{\phi}) - \underline{\phi}\hat{c}_2'(\underline{\phi}) \geq \hat{c}_1$. By Lemma A3, $\underline{\phi} > 0$ also implies $\hat{c}_2(\underline{\phi}) < \hat{c}_1$. Together, these imply $\hat{c}_2'(\underline{\phi}) < 0$. Analogously, $\overline{\phi} < 1$ implies $\hat{c}_2'(\overline{\phi}) > 0$. By Lemma 2, $\hat{c}_2(\cdot)$ is strictly convex, which therefore delivers $\underline{\phi} < \overline{\phi}$. Second, suppose that either $\underline{\phi} = 0$ or $\overline{\phi} = 1$. In such cases, $\underline{\phi} \leq \overline{\phi}$ holds automatically. □

A.5 Proof of Proposition 5

Proof. To begin, we observe that it follows from Definition 1 that $\underline{\phi} = 0$ implies $\hat{c}_2(\underline{\phi}) - \underline{\phi}\hat{c}_2'(\underline{\phi}) \leq \hat{c}_1$; $\underline{\phi} \in (0,1)$ implies $\hat{c}_2(\underline{\phi}) - \underline{\phi}\hat{c}_2'(\underline{\phi}) = \hat{c}_1$; and $\underline{\phi} = 1$ implies $\hat{c}_2(\underline{\phi}) - \underline{\phi}\hat{c}_2'(\underline{\phi}) \geq \hat{c}_1$.

As noted in the text, it suffices to focus on RFQ policies of the form described in Proposition 3. Such RFQ policies are described by two numbers: $q_{-\bar{s}}$ and $q_{\bar{s}}$, which capture the probability with which the client contacts only one dealer when $s = -\bar{s}$ and $s = \bar{s}$, respectively. This implies that if two dealers are contacted, then the posterior probability of $s = \bar{s}$ is $\frac{\phi_0(1-q_{\bar{s}})}{(1-\phi_0)(1-q_{-\bar{s}})+\phi_0(1-q_{\bar{s}})}$. Hence, the client's expected procurement cost is

$$\left[(1-\phi_0)q_{-\bar{s}} + \phi_0 q_{\bar{s}}\right]\hat{c}_1 + \left[(1-\phi_0)(1-q_{-\bar{s}}) + \phi_0(1-q_{\bar{s}})\right]\hat{c}_2\left(\frac{\phi_0(1-q_{\bar{s}})}{(1-\phi_0)(1-q_{-\bar{s}})+\phi_0(1-q_{\bar{s}})}\right). \tag{31}$$

We establish the proposition's claim through two steps. First, we show that $C(\phi_0)$ is a lower bound on the client's expected cost of procurement under any RFQ policy that leads to execution with probability one. Second, we show (considering separately the various cases) that the RFQ policy described in the proposition achieves this lower bound. From these facts, it will then follow both (*i*) that the RFQ policy described by the proposition is optimal, and (*ii*) that $C(\phi_0)$ is the expected cost of procurement under any optimal RFQ policy.

Proof of that $C(\phi_0)$ is a lower bound. $\hat{c}_1 \geq \min\{\hat{c}_1, \hat{c}_2(\phi)\} \geq C(\phi)$ and $\hat{c}_2(\phi) \geq \min\{\hat{c}_1, \hat{c}_2(\phi)\} \geq C(\phi)$ for all $\phi \in [0,1]$ imply that (31) is bounded below by

$$\left[(1-\phi_0)q_{-\bar{s}} + \phi_0 q_{\bar{s}}\right] C\left(\tfrac{\phi_0 q_{\bar{s}}}{(1-\phi_0)q_{-\bar{s}}+\phi_0 q_{\bar{s}}}\right) + \left[(1-\phi_0)(1-q_{-\bar{s}}) + \phi_0(1-q_{\bar{s}})\right] C\left(\tfrac{\phi_0(1-q_{\bar{s}})}{(1-\phi_0)(1-q_{-\bar{s}})+\phi_0(1-q_{\bar{s}})}\right),$$

which is bounded below by $C(\phi_0)$, by the convexity of $C(\cdot)$.

The case of $\phi_0 \in [\underline{\phi}, \overline{\phi}]$ and $\underline{\phi} < \overline{\phi}$. Plugging $q_{-\bar{s}} = 0$ and $q_{\bar{s}} = 0$ into (31) yields $\hat{c}_2(\phi_0)$. Therefore, we need to show that $C(\phi_0) \geq \hat{c}_2(\phi_0)$. Define the following linear function:

$$L(\phi) = \hat{c}_2(\phi_0) - \phi_0\hat{c}_2'(\phi_0) + \hat{c}_2'(\phi_0)\phi.$$

We claim that for all $\phi \in [0,1]$, $L(\phi) \leq \min\{\hat{c}_1, \hat{c}_2(\phi)\}$:

- First, we show that $L(\phi) \leq \hat{c}_1$ for all $\phi \in [0,1]$. By linearity, it suffices to check the endpoints. By assumption, $\underline{\phi} < \overline{\phi}$, which because $\overline{\phi} \leq 1$, implies $\underline{\phi} < 1$. Hence, as observed at the beginning of the proof, we have $\hat{c}_2(\underline{\phi}) - \underline{\phi}\hat{c}_2'(\underline{\phi}) \leq \hat{c}_1$. Additionally, as established in footnote 28, $\hat{c}_2(\phi) - \phi\hat{c}_2'(\phi)$ is strictly decreasing. Thus, because $\phi_0 \geq \underline{\phi}$, it follows that

  $$L(0) = \hat{c}_2(\phi_0) - \phi_0\hat{c}_2'(\phi_0) \leq \hat{c}_2(\underline{\phi}) - \underline{\phi}\hat{c}_2'(\underline{\phi}) \leq \hat{c}_1.$$

  An analogous argument establishes that $L(1) \leq \hat{c}_1$.

- Second, note that $L(\phi)$ is tangent to $\hat{c}_2(\phi)$ at $\phi = \phi_0$. Indeed, we have both $L(\phi_0) = \hat{c}_2(\phi_0)$ and also $L'(\phi_0) = \hat{c}_2'(\phi_0)$. Thus, convexity of $\hat{c}_2(\phi)$ implies that $L(\phi) \leq \hat{c}_2(\phi)$ for all $\phi \in [0,1]$.

It follows that $C(\phi_0) \geq L(\phi_0) = \hat{c}_2(\phi_0)$, as desired.

The case of $\phi_0 = \underline{\phi} = \overline{\phi}$. Plugging $q_{-\bar{s}} = 0$ and $q_{\bar{s}} = 0$ into (31) yields $\hat{c}_2(\phi_0)$. Therefore, we need to show that $C(\phi_0) \geq \hat{c}_2(\phi_0)$. To begin, note that by Proposition 4*(ii)*, $\underline{\phi} = \overline{\phi}$ is possible only if either $\underline{\phi} = 0$ or $\overline{\phi} = 1$. Let us assume $\phi_0 = \underline{\phi} = \overline{\phi} = 0$; the case of $\phi_0 = \underline{\phi} = \overline{\phi} = 1$ is symmetric. For this part of the proof, redefine $L(\cdot)$ as a new linear function:

$$L(\phi) = \hat{c}_2(0)(1-\phi) + \hat{c}_1\phi.$$

We claim that for all $\phi \in [0,1]$, $L(\phi) \leq \min\{\hat{c}_1, \hat{c}_2(\phi)\}$:

- First, by Definition 1, $\underline{\phi} = 0$ implies that $\hat{c}_2(0) \leq \hat{c}_1$. It follows that $L(\phi) = \hat{c}_2(0)(1-\phi) + \hat{c}_1\phi \leq \hat{c}_1$ for all $\phi \in [0,1]$.

- Second, by Definition 1, $\overline{\phi} = 0$ implies that $\hat{c}_1 - \hat{c}_2(0) \leq \hat{c}_2'(0)$. It follows that $L(\phi)$ is subtangent to $\hat{c}_2(\phi)$ at $\phi = 0$. Indeed, we have both $L(0) = \hat{c}_2(0)$ and also $L'(0) = \hat{c}_1 - \hat{c}_2(0) \leq \hat{c}_2'(0)$. Thus, convexity of $\hat{c}_2(\phi)$ implies that $L(\phi) \leq \hat{c}_2(\phi)$ for all $\phi \in [0, 1]$.

It follows that $C(\phi_0) \geq L(\phi_0) = \hat{c}_2(\phi_0)$, as desired.

The case of $\phi_0 \in [0, \underline{\phi})$. Plugging $q_{-\bar{s}} = \frac{\underline{\phi} - \phi_0}{\underline{\phi}(1-\phi_0)}$ and $q_{\bar{s}} = 0$ into (31) yields $\hat{c}_1 + \frac{\phi_0}{\underline{\phi}}[\hat{c}_2(\underline{\phi}) - \hat{c}_1]$. Therefore, we need to show that $C(\phi_0) \geq \hat{c}_1 + \frac{\phi_0}{\underline{\phi}}[\hat{c}_2(\underline{\phi}) - \hat{c}_1]$. For this part of the proof, redefine $L(\cdot)$ as a new linear function:

$$L(\phi) = \hat{c}_1 + \frac{\hat{c}_2(\underline{\phi}) - \hat{c}_1}{\underline{\phi}} \phi.$$

Note that the previous quantities are all well-defined because, by assumption, $\phi_0 < \underline{\phi}$, which implies that $\underline{\phi} > 0$. Next, we claim that for all $\phi \in [0, 1]$, $L(\phi) \leq \min\{\hat{c}_1, \hat{c}_2(\phi)\}$. In fact, that claim was established in the course of the proof of Lemma A3. It follows that

$$C(\phi_0) \geq L(\phi_0) = \hat{c}_1 + \frac{\hat{c}_2(\underline{\phi}) - \hat{c}_1}{\underline{\phi}} \phi_0,$$

as desired. The case of $\phi_0 \in (\overline{\phi}, 1]$ is symmetric. $\square$

B Robustness to the commitment assumption

This appendix investigates the extent to which our results are robust to the assumption that the client can commit to an RFQ policy. We find that the result of Proposition 3 (on the optimality of minimal disclosure) does not hinge on the client's ability to commit: she would not want to deviate from a commitment to minimal disclosure *even if she could*. In contrast, the result of Proposition 5 (on the optimal policy for determining the number of dealers to contact) does hinge on the commitment assumption. We then investigate how that result would change if the client were to lack commitment power.

B.1 Does the commitment assumption matter?

To investigate these issues, we define interim analogues of $\hat{c}_1$ and $\hat{c}_2(\phi)$. To begin, suppose an RFQ contacts one dealer. Arguments analogous to those given before pin down the client's procurement cost conditional on her type: it is $\frac{3\bar{s}^2}{4}$, regardless of her realized type $s \in \{-\bar{s}, \bar{s}\}$ and regardless of the belief ϕ induced by the RFQ. Hence, we write $\hat{c}_{1,-\bar{s}} = \hat{c}_{1,\bar{s}} = \frac{3\bar{s}^2}{4}$. Note that these interim expressions are consistent with the earlier ex ante expression in that they satisfy $\hat{c}_1 = (1-\phi)\hat{c}_{1,-\bar{s}} + \phi\hat{c}_{1,\bar{s}}$.

Next, suppose an RFQ contacts two dealers and induces a belief ϕ. A derivation analogous to that in the proof of Lemma 2 pins down the client's expected procurement cost conditional on her type, $\hat{c}_{2,-\bar{s}}(\phi)$ and $\hat{c}_{2,\bar{s}}(\phi)$. These expressions are similarly consistent with the ex ante cost in that they satisfy $\hat{c}_2(\phi) = (1-\phi)\hat{c}_{2,-\bar{s}}(\phi) + \phi\hat{c}_{2,\bar{s}}(\phi)$. These expressions moreover imply the following result, which says that conditional on the client's type, her utility decreases in the probability dealers assign to that type. Proofs are in Online Appendix B.3.

Lemma B6. *$\hat{c}'_{2,-\bar{s}}(\phi) \leq 0$ and $\hat{c}'_{2,\bar{s}}(\phi) \geq 0$ on the domain $\phi \in [0,1]$.*

Optimality of minimal disclosure. The optimality of minimal disclosure does not depend on the client's ability to commit. To argue this point, we first recall that information design does not matter when only one dealer is contacted. Hence, it suffices to focus on the case of $M = 2$. Given an arbitrary RFQ policy of the form described in Proposition 3 (i.e., entailing minimal disclosure), let ϕ_2 denote the posterior beliefs about the client's type induced by a realization of $M = 2$. An implication of the monotonicity described in Lemma B6 is that there is no scope to communicate any further information via cheap talk (Crawford and Sobel, 1982). Indeed, in any equilibrium of the corresponding cheap talk game, Lemma B6 implies that a selling (buying) client must always send the message(s) inducing the highest (lowest) belief that the client is a buyer. These beliefs can then be consistent with Bayes rule only if all messages induce the same belief (necessarily ϕ_2).

Number of dealers to contact. In contrast, what may hinge on the client's ability to commit is her policy for determining the number of dealers to contact. To illustrate with an example, let us consider the parametrization corresponding to the fourth panel of Figure 3: $\psi = 0.85$, $\rho = 1$, and $\phi_0 = 0.2$. The optimal policy has the feature that under the realization $s = -\bar{s}$, the client randomizes between contacting one and two dealers—and in such a way that contacting two dealers induces a belief $\underline{\phi} = 0.632$. Yet given this belief, the client would not be indifferent between contacting one or two dealers at this interim stage. Indeed, we can compute $\hat{c}_{1,-\bar{s}} = 0.75\bar{s}^2$ and $\hat{c}_{2,-\bar{s}}(\underline{\phi}) \approx 0.7289\bar{s}^2$. Thus, the client would wish to deviate from her commitment to sometimes contact only a single dealer when her realized type is $s = -\bar{s}$.

B.2 If the client could not commit

A natural question is what would happen in our model if the client lacked the ability to commit.[36] Based on the analysis in Online Appendix B.1, the client will continue to use a minimally-disclosing RFQ policy—any difference will be only in terms of the number of dealers she contacts.

Analogous to before, we begin by defining two cutoffs: $\underset{\sim}{\phi}$ and $\tilde{\phi}$.

Definition 1′. Define $\underset{\sim}{\phi}, \tilde{\phi} \in [0,1]$ as follows. If $\hat{c}_{2,-\bar{s}}(0) \leq \hat{c}_1$, define $\underset{\sim}{\phi} = 0$; if $\hat{c}_{2,-\bar{s}}(1) \geq \hat{c}_1$, define $\underset{\sim}{\phi} = 1$; otherwise, define it implicitly as the unique $\underset{\sim}{\phi} \in (0,1)$ that solves $\hat{c}_{2,-\bar{s}}(\phi) = \hat{c}_1$. If $\hat{c}_{2,\bar{s}}(1) \leq \hat{c}_1$, define $\tilde{\phi} = 1$; if $\hat{c}_{2,\bar{s}}(0) \geq \hat{c}_1$, define $\tilde{\phi} = 0$; otherwise, define it implicitly as the unique $\tilde{\phi} \in (0,1)$ that solves $\hat{c}_{2,\bar{s}}(\phi) = \hat{c}_1$.[37]

As the notation suggests, $\underset{\sim}{\phi} \leq \tilde{\phi}$. In fact, an even stronger result holds:

Proposition 4′. *$\underset{\sim}{\phi} = 0$ or $\tilde{\phi} = 1$ (or both).*

To analyze what would happen without commitment, we characterize the client's constrained-optimal RFQ policy, where the constraint is that she would not want to deviate from the policy after learning her type. By arguments similar to those given earlier, it again suffices to focus on RFQ policies of the form described in Proposition 3. But we now select a different member of that class, one characterized by the following counterpart to Proposition 5:

Proposition 5′. *Without commitment,*

$$\begin{cases} (1-\phi_0)\hat{c}_1 + \phi_0 \hat{c}_{2,\bar{s}}(\underset{\sim}{\phi}) & \text{if } \phi_0 \in [0, \underset{\sim}{\phi}), \\ (1-\phi_0)\hat{c}_{2,-\bar{s}}(\phi_0) + \phi_0 \hat{c}_{2,\bar{s}}(\phi_0) & \text{if } \phi_0 \in [\underset{\sim}{\phi}, \tilde{\phi}], \\ (1-\phi_0)\hat{c}_{2,-\bar{s}}(\tilde{\phi}) + \phi_0 \hat{c}_1 & \text{if } \phi_0 \in (\tilde{\phi}, 1], \end{cases}$$

is the expected cost of procurement under any optimal RFQ policy. Moreover, one optimal policy is

[36] Recall that our model assumes that each dealer observes the number of other dealers that the client has contacted. Without such observability it is easy to show that the no-commitment outcome would entail unraveling: in equilibrium, the client would always contact all available dealers. With observability, the client must do no worse. This suggests that, at least in the no-commitment regime, clients benefit from RFQ platforms that reveal the number of market participants that an RFQ has been sent to, as is the case for many such platforms in practice (*cf.* footnote 16).

[37] That this provides a well-defined definition for $\underset{\sim}{\phi}$ and $\tilde{\phi}$ follows from the monotonicity of $\hat{c}_{2,-\bar{s}}(\phi)$ and $\hat{c}_{2,\bar{s}}(\phi)$, as established by Lemma B6.

in the class described by Proposition 3, where the probabilities of contacting only a single dealer are

$$(q_{-\bar{s}}, q_{\bar{s}}) = \begin{cases} \left(\frac{\phi - \phi_0}{\underset{\sim}{\phi}(1-\phi_0)}, 0\right) & \textit{if } \phi_0 \in [0, \underset{\sim}{\phi}), \\ (0,0) & \textit{if } \phi_0 \in [\underset{\sim}{\phi}, \tilde{\phi}], \\ \left(0, \frac{\phi_0 - \tilde{\phi}}{\phi_0(1-\tilde{\phi})}\right) & \textit{if } \phi_0 \in (\tilde{\phi}, 1]. \end{cases}$$

The takeaway of our main-text analysis is that information leakage can act as a search friction, leading the client to limit the number of dealers that she contacts. Proposition 5′ demonstrates that this conclusion does not depend on the commitment power we had assumed for the client. Even without commitment, the client might sometimes contact only a single dealer.

Example. We next illustrate Proposition 5′ by revisiting the example underlying Figure 3. The parametrization is as before: $\psi = 0.85$ and $\rho = 1$. Then $\hat{c}_1$, $\hat{c}_{2,-\bar{s}}(\phi)$, and $\hat{c}_{2,\bar{s}}(\phi)$ are as depicted in the first panel of Figure 3′. For comparability with the first panel of Figure 3, we also depict $\hat{c}_2(\phi) = (1-\phi)\hat{c}_{2,-\bar{s}}(\phi) + \phi\hat{c}_{2,\bar{s}}(\phi)$ but suppress it in the subsequent panels.

The second panel of Figure 3′ depicts $\underset{\sim}{\phi}$, which is defined as the intersection between $\hat{c}_{2,-\bar{s}}(\phi)$ and $\hat{c}_1$. We also have $\tilde{\phi} = 1$ in this example, but we suppress it in the figure because $\tilde{\phi}$ plays no role what follows.

The third panel depicts the case of $\phi_0 \in [\underset{\sim}{\phi}, \tilde{\phi}]$. According to Proposition 5′, the optimal RFQ policy always contacts two dealers and discloses no information about the client's order. Under this policy, dealers' beliefs therefore always coincide with the prior, so that the client's expected procurement cost is $(1-\phi_0)\hat{c}_{2,-\bar{s}}(\phi_0) + \phi_0\hat{c}_{2,\bar{s}}(\phi_0)$, or simply $\hat{c}_2(\phi_0)$.

Finally, the fourth panel depicts the case of $\phi_0 \in [0, \underset{\sim}{\phi})$. According to Proposition 5′, the optimal RFQ policy always contacts two dealers when $s = \bar{s}$; and it mixes between one and two dealers when $s = -\bar{s}$. Moreover, this mixing is designed to ensure that, conditional on two dealers being contacted, they are induced to believe that $s = \bar{s}$ with probability $\underset{\sim}{\phi}$, which renders a selling client indifferent between contacting one or two dealers, since $\hat{c}_{2,-\bar{s}}(\underset{\sim}{\phi}) = \hat{c}_1$. It is this indifference that permits the policy to entail randomization even in this no-commitment version of the model. Under this policy, the client's expected procurement cost is therefore $(1-\phi_0)\hat{c}_1 + \phi_0\hat{c}_{2,\bar{s}}(\underset{\sim}{\phi})$.

Figure 3′: Optimal RFQ policies under limited commitment power

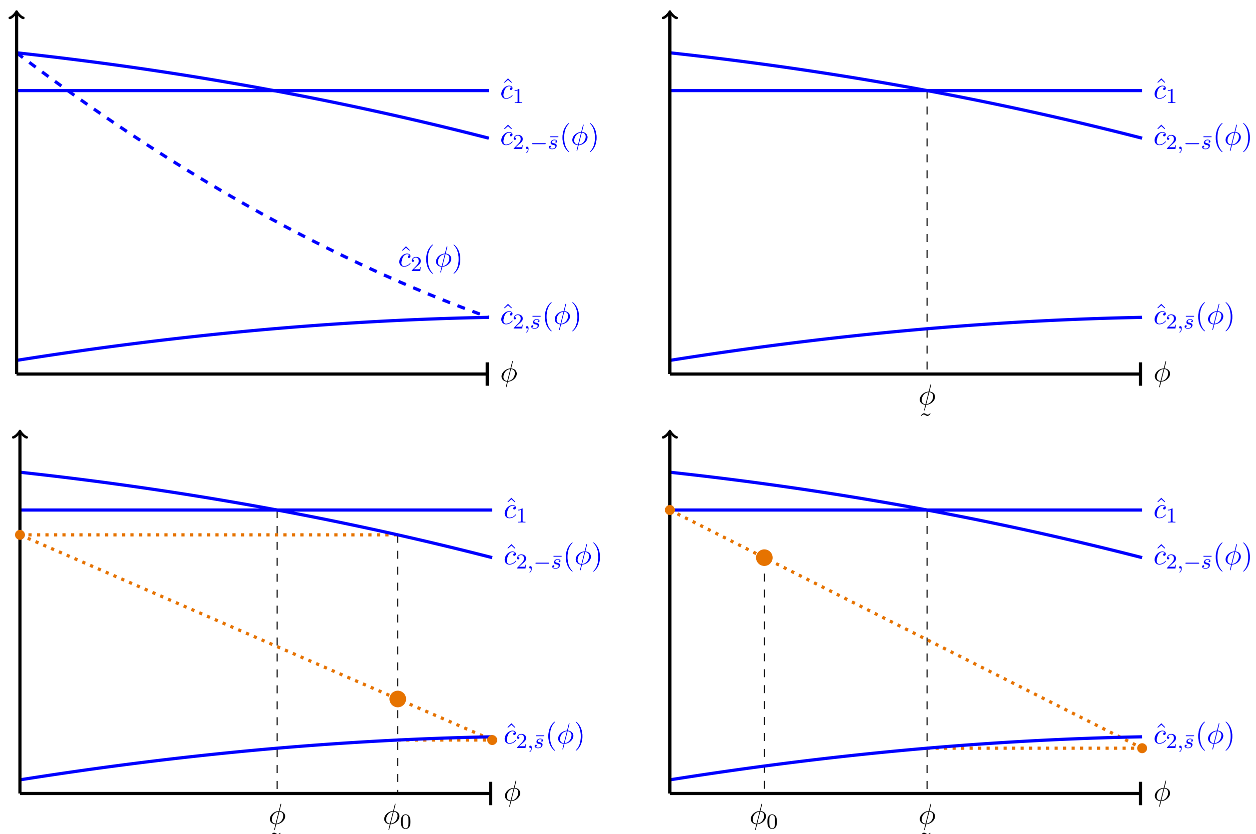

Dealer inventories are distributed according to the parameters $\psi = 0.85$ and $\rho = 1$. In all panels, $\hat{c}_1$ (in blue) depicts expected procurement cost when contacting one dealers; likewise, $\hat{c}_{2,-\bar{s}}(\phi)$ and $\hat{c}_{2,-\bar{s}}(\phi)$ (in blue) depict expected procurement cost when contacting two dealers, conditional on the client being a seller or a buyer, respectively. The first panel also depicts $\hat{c}_2(\phi)$ (in dashed blue), which is the unconditional expected procurement cost when contacting two dealers. Starting from the second panel, $\underset{\sim}{\phi}$ is as in Definition 1′. There are two cases that depend on ϕ_0, the prior probability of the client buying. If $\phi_0 \geq \underset{\sim}{\phi}$ (as in the third panel), then the commitment-free RFQ policy always contacts two dealers, inducing a posterior belief that the client is a buyer which coincides with the prior ϕ_0. If $\phi_0 < \underset{\sim}{\phi}$ (as in the fourth panel), then the commitment-free RFQ policy is such that contacting two dealers induces a posterior belief that the client is a buyer with probability $\underset{\sim}{\phi}$. In both cases, the unconditional expected procurement cost realized by the policy is represented by the large orange dot. Vertical axes do not depict units because all costs scale linearly with $\bar{s}^2$. Horizontal axes measure probability and depict the unit interval.

B.3 Proofs

Proof sketch of Lemma B6. A derivation analogous to that in the proof of Lemma 2 allows us to characterize the client's expected procurement cost conditional on her type:

$$\begin{aligned}
\hat{c}_{2,-\bar{s}}(\phi) &= \left[1 - \frac{23\phi(16+7\phi)}{4(24-\phi)^2}\right]\bar{s}^2\psi[1-(1-\psi)(1-\rho)] \\
&\quad - 23\left[\frac{3\phi}{2(23+\phi)}\right]^2\bar{s}^2(1-\psi)[1-\psi(1-\rho)] + \frac{7\bar{s}^2}{8}\psi(1-\psi)(1-\rho), \\
\hat{c}_{2,\bar{s}}(\phi) &= \left[1 - \frac{23(1-\phi)(23-7\phi)}{4(23+\phi)^2}\right]\bar{s}^2(1-\psi)[1-\psi(1-\rho)] \\
&\quad - 23\left[\frac{3(1-\phi)}{2(24-\phi)}\right]^2\bar{s}^2\psi[1-(1-\psi)(1-\rho)] + \frac{7\bar{s}^2}{8}\psi(1-\psi)(1-\rho).
\end{aligned}$$

To establish the claim, we simply compute

$$\hat{c}'_{2,-\bar{s}}(\phi) = -\frac{184(12+11\phi)\bar{s}^2}{(24-\phi)^3}\psi[1-(1-\psi)(1-\rho)] - \frac{4761\phi\bar{s}^2}{2(23+\phi)^3}(1-\psi)[1-\psi(1-\rho)],$$

which is indeed weakly negative on the domain $\phi \in [0,1]$, and

$$\hat{c}'_{2,\bar{s}}(\phi) = \frac{184(23-11\phi)\bar{s}^2}{(23+\phi)^3}(1-\psi)[1-\psi(1-\rho)] + \frac{4761(1-\phi)\bar{s}^2}{2(24-\phi)^3}\psi[1-(1-\psi)(1-\rho)],$$

which is indeed weakly positive on the domain $\phi \in [0,1]$. □

Proof of Proposition 4′. Computing, $c_{2,-\bar{s}}(0) + \hat{c}_{2,\bar{s}}(1) = \bar{s}^2\Big[\psi[1-(1-\psi)(1-\rho)] + (1-\psi)[1-\psi(1-\rho)] + \frac{7}{8}\psi(1-\psi)(1-\rho)\Big] \leq \bar{s}^2 < \frac{3}{2}\bar{s}^2 = 2\hat{c}_1$. Thus, we cannot have both $\hat{c}_{2,-\bar{s}}(0) \geq \hat{c}_1$ and $\hat{c}_{2,\bar{s}}(1) \geq \hat{c}_1$. Applying Definition 1′, we therefore cannot have both $\underaccent{\tilde}{\phi} > 0$ and $\tilde{\phi} < 1$. □

Proof of Proposition 5′. *The case of* $\phi_0 \in [\underaccent{\tilde}{\phi}, \tilde{\phi}]$. Under the stated policy $(q_{-\bar{s}}, q_{\bar{s}}) = (0,0)$, when the client contacts two dealers, she induces a belief ϕ_0. Because $\phi_0 \in [\underaccent{\tilde}{\phi}, \tilde{\phi}]$, we have both $\hat{c}_{2,-\bar{s}}(\phi_0) < \hat{c}_1$ and $\hat{c}_{2,\bar{s}}(\phi_0) < \hat{c}_1$, so regardless of the realized s, the client does not want to deviate from the stated policy.

Do there exist any other policies that she would similarly not want to deviate from? The answer is no. In the proof of Proposition 4′, we showed that $\hat{c}_{2,-\bar{s}}(0) < \hat{c}_1$ or $\hat{c}_{2,\bar{s}}(1) < \hat{c}_1$ (or both). Without loss of generality, suppose $\hat{c}_{2,\bar{s}}(1) < \hat{c}_1$. By the monotonicity of $\hat{c}_{2,\bar{s}}(\cdot)$ (i.e., Lemma B6), we have $\hat{c}_{2,\bar{s}}(\phi) < \hat{c}_1$ for all $\phi \in [0,1]$. Therefore, if the client will not deviate when $s = \bar{s}$, then the policy must entail $q_{\bar{s}} = 0$. Under any such policy, when the client contacts two dealers, she induces a belief $\phi_2 \geq \phi_0$. By monotonicity of $\hat{c}_{2,-\bar{s}}(\cdot)$ (i.e., Lemma B6), we have $\hat{c}_{2,-\bar{s}}(\phi_2) \leq \hat{c}_{2,-\bar{s}}(\phi_0) < \hat{c}_1$. Therefore, if the client will not deviate when $s = -\bar{s}$, then the policy must entail $q_{-\bar{s}} = 0$.

The case of $\phi_0 \in [0, \underaccent{\tilde}{\phi})$. Under the stated policy $(q_{-\bar{s}}, q_{\bar{s}}) = \left(\frac{\underaccent{\tilde}{\phi}-\phi_0}{\underaccent{\tilde}{\phi}(1-\phi_0)}, 0\right)$, when the client contacts two dealers, she induces a belief $\underaccent{\tilde}{\phi}$. Here, we have $\underaccent{\tilde}{\phi} > 0$, which by Proposition 4′ implies $\tilde{\phi} = 1$, and hence $\hat{c}_{2,\bar{s}}(1) \leq \hat{c}_1$. We also have $\hat{c}_{2,-\bar{s}}(\underset{\sim}{\phi}) = \hat{c}_1$, so that the client (weakly) does not want to deviate from the stated policy when $s = -\bar{s}$. And by monotonicity of $\hat{c}_{2,\bar{s}}(\cdot)$ (i.e., Lemma B6), we have $\hat{c}_{2,\bar{s}}(\underset{\sim}{\phi}) \leq \hat{c}_{2,\bar{s}}(1) \leq \hat{c}_1$, so that the client also does not want to deviate when $s = \bar{s}$.

Do there exist any other policies that she would similarly not want to deviate from? Again, the answer is no. As argued in the previous case, so that the client will not deviate when $s = \bar{s}$, the policy must entail $q_{\bar{s}} = 0$. Next, suppose by way of contradiction that $q_{-\bar{s}} < \frac{\underset{\sim}{\phi} - \phi_0}{\underset{\sim}{\phi}(1-\phi_0)}$. Then when the client contacts two dealers, she induces a belief $\phi_2 < \underset{\sim}{\phi}$. By monotonicity of $\hat{c}_{2,-\bar{s}}(\cdot)$ (i.e., Lemma B6), we have $\hat{c}_{2,-\bar{s}}(\phi_2) > \hat{c}_{2,-\bar{s}}(\underset{\sim}{\phi}) = \hat{c}_1$, which means that the client will deviate when $s = -\bar{s}$, unless $q_{-\bar{s}} = 1$. But $q_{-\bar{s}} = 1$ is incompatible with $q_{-\bar{s}} < \frac{\underset{\sim}{\phi} - \phi_0}{\underset{\sim}{\phi}(1-\phi_0)}$. An analogous argument rules out policies in which $q_{-\bar{s}} > \frac{\underset{\sim}{\phi} - \phi_0}{\underset{\sim}{\phi}(1-\phi_0)}$.

The case of $\phi_0 \in (\tilde{\phi}, 1]$ can be treated similarly. $\square$

C The competition, sampling, and front-running effects

The discussion following Lemma 2 asserted that the difference between $\hat{c}_1$ and $\hat{c}_2(\phi)$ can be decomposed into three effects: a competition effect, a sampling effect, and a front-running effect. To support that claim, this appendix provides a formalization of these three effects. To that end, define the following quantities:

- Let $C_{IV}(\phi)$ be the client's cost following an RFQ that contacts two dealers and induces belief ϕ. $C_{IV}(\phi)$ is simply $\hat{c}_2(\phi)$.

- Let $C_{III}(\phi)$ be the client's cost following an RFQ that contacts two dealers and induces belief ϕ and if ρ were to equal one. $C_{III}(\phi)$ is simply $\hat{c}_2(\phi)$, except evaluated at $\rho = 1$ rather than the true ρ.

- Let $C_{II}(\phi)$ be the client's cost following an RFQ that contacts two dealers and induces belief ϕ, if ρ were to equal one, and if the losing dealer were prohibited from trading in the first period. The appendix considers a version of the model in which the losing dealer is prohibited from trading in the first period. The equation for $\tilde{c}_2(\phi)$ derived there is the analogue of $\hat{c}_2(\phi)$. $C_{II}(\phi)$ is simply $\tilde{c}_2(\phi)$, except evaluated at $\rho = 1$ rather than the true ρ.

- Let $C_I(\phi)$ be the client's cost following an RFQ that contacts one dealer and induces belief ϕ, if ρ were to equal one, and if the losing dealer were prohibited from trading in the first period. Because only one dealer is contacted, ρ is irrelevant, as is the prohibition on first-period trading for a losing dealer. It follows that $C_I(\phi)$ is simply $\hat{c}_1$.

Given these definitions, we can interpret $C_{IV}(\phi) - C_{III}(\phi)$ as the *sampling effect*, $C_{III}(\phi) - C_{II}(\phi)$ as the *front-running effect*, and $C_{II}(\phi) - C_I(\phi)$ as the *competition effect*. Mathematically,

$$\hat{c}_2(\phi) - \hat{c}_1 = \underbrace{C_{IV}(\phi) - C_{III}(\phi)}_{\text{sampling effect}} + \underbrace{C_{III}(\phi) - C_{II}(\phi)}_{\text{front-running effect}} + \underbrace{C_{II}(\phi) - C_I(\phi)}_{\text{competition effect}}.$$

Moreover—and consistent with intuition—the competition and sampling effects benefit the client, while the front-running effect harms the client, in the sense of the following result:

Proposition C1. *For all $\phi \in [0, 1]$, $C_{IV}(\phi) \leq C_{III}(\phi)$, $C_{III}(\phi) > C_{II}(\phi)$, and $C_{II}(\phi) < C_I(\phi)$.*

The proof of this proposition uses the expression for $C_{II}(\phi)$ provided by the following lemma.

Proof of Proposition C1. To show that $C_{IV}(\phi) \leq C_{III}(\phi)$, it suffices to show that the expression for $\hat{c}_2(\phi)$ given in (10) is increasing in ρ. Taking the derivative of that expression with respect to ρ, we obtain

$$\frac{(304704 - 1216654\phi + 1210215\phi^2 + 12878\phi^3 - 6439\phi^4)\psi(1-\psi)\bar{s}^2}{8(24-\phi)^2(23+\phi)^2},$$

which is positive on the domain $\phi \in [0,1]$ and $\psi \in (0,1)$. To obtain $C_{III}(\phi)$, we evaluate the expression for $\hat{c}_2(\phi)$ given in (10) at $\rho = 1$. To obtain $C_{II}(\phi)$, we evaluate the expression for $\tilde{c}_2(\phi)$ given in (3) at $\rho = 1$. To show that $C_{III}(\phi) > C_{II}(\phi)$, we compute directly:

$$C_{III}(\phi) - C_{II}(\phi) = (1-\phi)\psi\bar{s}^2 \underbrace{\left[\frac{(4-\phi)\bar{s}^2}{4} + \frac{\phi(1-\phi)^2\bar{s}^2}{4(24-\phi)^2} - \frac{3}{4}\right]}_{>0 \text{ for } \phi\in[0,1]} + \phi(1-\psi)\bar{s}^2 \underbrace{\left[\frac{(3+\phi)\bar{s}^2}{4} + \frac{(1-\phi)\phi^2\bar{s}^2}{4(23+\phi)^2} - \frac{3}{4}\right]}_{>0 \text{ for } \phi\in[0,1]}.$$

To show that $C_{II}(\phi) < C_I(\phi)$, we also compute directly:

$$\begin{aligned} C_I(\phi) - C_{II}(\phi) &= \frac{3\bar{s}^2}{4} - \frac{3\bar{s}^2}{4}\left[(1-\phi)\psi + \phi(1-\psi)\right] \\ &= \frac{3\bar{s}^2}{4}\left[\phi\psi + (1-\phi)(1-\psi)\right], \end{aligned}$$

which is positive on the domain $\phi \in [0,1]$ and $\psi \in (0,1)$. □

D Order splitting

In this appendix, we consider what would happen if the client were to contact both dealers and—rather than auction her entire order as a single indivisible unit—instead permits each dealer to win only half of the total order. We show that a client would never optimally split an order in this way, because it leads to a problem akin to double marginalization.

For simplicity, suppose the client continues to treat the order as indivisible in the sense that she allocates the order only if the two dealers' bids both meet her reservation price. In other words, the game ends if one or both dealers fail to meet the reserve. Given the structure of the model, it is common knowledge that the client seeks to trade $\bar{s}$ shares. Thus, if a dealer is awarded an order to buy or sell $\frac{\bar{s}}{2}$ shares, then he knows that the other dealer is also being awarded an order of the same direction and size.

The following result characterizes the continuation equilibrium following such an RFQ.

Lemma D1. *There is a WPBE in which the following occurs on path. Dealer A bids*

$$(b^A_{-\bar{s}}, b^A_{\bar{s}}) = \begin{cases} \left(\frac{7\bar{s}^2}{18}, 0\right) & \textit{if } (e^A, e^B) = (1, 1), \\ \left(\frac{7\bar{s}^2}{50}, -\frac{7\bar{s}^2}{100}\right) & \textit{if } (e^A, e^B) = (1, -1), \\ \left(-\frac{7\bar{s}^2}{100}, \frac{7\bar{s}^2}{50}\right) & \textit{if } (e^A, e^B) = (-1, 1), \\ \left(0, \frac{7\bar{s}^2}{18}\right) & \textit{if } (e^A, e^B) = (-1, -1). \end{cases}$$

If both dealers win, dealer A's on-market trades are

$$(x^A_1, x^A_2) = \begin{cases} (0, 0) & \textit{if } (s, e^A, e^B) = (\bar{s}, 1, 1), \\ (-\frac{\bar{s}}{3}, -\frac{\bar{s}}{6}) & \textit{if } (s, e^A, e^B) = (-\bar{s}, 1, 1), \\ (\frac{\bar{s}}{10}, -\frac{3\bar{s}}{10}) & \textit{if } (s, e^A, e^B) = (\bar{s}, 1, -1), \\ (-\frac{\bar{s}}{10}, -\frac{2\bar{s}}{5}) & \textit{if } (s, e^A, e^B) = (-\bar{s}, 1, -1), \\ (\frac{\bar{s}}{10}, \frac{2\bar{s}}{5}) & \textit{if } (s, e^A, e^B) = (\bar{s}, -1, 1), \\ (-\frac{\bar{s}}{10}, \frac{3\bar{s}}{10}) & \textit{if } (s, e^A, e^B) = (-\bar{s}, -1, 1), \\ (\frac{\bar{s}}{3}, \frac{\bar{s}}{6}) & \textit{if } (s, e^A, e^B) = (\bar{s}, -1, -1), \\ (0, 0) & \textit{if } (s, e^A, e^B) = (-\bar{s}, -1, -1). \end{cases}$$

Dealer B's bids and on-market trades are specified symmetrically.

Proof sketch of Lemma D1. The proof sketch that we provide here is informal in that we (*i*) directly plug in the constraints that will bind on the equilibrium path, and (*ii*) ignore the constraints that do not bind on the equilibrium path. These simplifications do not affect the result. Indeed, the proof could be made more formal in the same way that the proofs of the analogous main results (Lemmas 1 and 2) are fully formal.

Because both dealers observe the entire vector (e^A, e^B), the four possible realizations of that vector can be analyzed separately. Below, we analyze the cases of $(1,1)$ and $(1,-1)$; the remaining cases can be handled symmetrically. We also note that in this case of order splitting, both dealers observe s directly before trading takes place, and so there is no need to keep track of beliefs. Thus, the cases of $s = -\bar{s}$ and $s = \bar{s}$ can also be analyzed separately.

Case 1: $(e^A, e^B) = (1,1)$ and $s = \bar{s}$. Ignoring the constraints on final inventory (which will not bind in the equilibrium), dealers A and B respectively minimize

$$(x_1^A + x_1^B)x_1^A + (x_1^A + x_1^B + x_2^A + x_2^B)x_2^A$$
$$(x_1^A + x_1^B)x_1^B + (x_1^A + x_1^B + x_2^A + x_2^B)x_2^B,$$

leading to $x_2^A = x_2^B = -\frac{x_1^A + x_1^B}{3}$. Inducting backward, we obtain $(x_1^A, x_1^B) = (0,0)$, so that $(x_2^A, x_2^B) = (0,0)$ on path. Plugging in these trades, dealer A incurs no trading costs if he wins. So the refinement described in Section 3.1 requires $b_{\bar{s}}^A = 0$ to be his bid.

Case 2: $(e^A, e^B) = (1,1)$ and $s = -\bar{s}$. Assuming that $x_2^A = -\frac{\bar{s}}{2} - x_1^A$ (which ensures that dealer A's final inventory just meets the constraint $e^A + x_1^A + x_2^A - s \leq 1$), assuming also that $x_2^B = -\frac{\bar{s}}{2} - x_1^B$ (symmetrically), and ignoring all other constraints on final inventory, dealers A and B respectively minimize

$$(x_1^A + x_1^B)x_1^A + (-\bar{s})\left(-\frac{\bar{s}}{2} - x_1^A\right)$$
$$(x_1^A + x_1^B)x_1^B + (-\bar{s})\left(-\frac{\bar{s}}{2} - x_1^B\right),$$

leading to $(x_1^A, x_1^B) = (-\frac{\bar{s}}{3}, -\frac{\bar{s}}{3})$, which implies $(x_2^A, x_2^B) = (-\frac{\bar{s}}{6}, -\frac{\bar{s}}{6})$. Plugging in these trades, dealer A incurs trading costs of $\frac{7\bar{s}^2}{18}$ if he wins. So the refinement described in Section 3.1 requires $b_{-\bar{s}}^A = \frac{7\bar{s}^2}{18}$ to be his bid.

Case 3: $(e^A, e^B) = (1,-1)$ and $s = \bar{s}$. Assuming that $x_2^B = \frac{\bar{s}}{2} - x_1^B$ (which ensures that dealer B's final inventory just meets the constraint $e^B + x_1^B + x_2^B - s \geq -1$), and ignoring all other constraints on final inventory, dealers A and B respectively minimize

$$(x_1^A + x_1^B)x_1^A + \left(x_1^A + x_2^A + \frac{\bar{s}}{2}\right)x_2^A$$
$$(x_1^A + x_1^B)x_1^B + \left(x_1^A + x_2^A + \frac{\bar{s}}{2}\right)\left(\frac{\bar{s}}{2} - x_1^B\right),$$

leading to $x_2^A = -\frac{x_1^A}{2} - \frac{\bar{s}}{4}$. Inducting backward, we obtain $(x_1^A, x_1^B) = \left(\frac{\bar{s}}{10}, \frac{\bar{s}}{10}\right)$, so that $(x_2^A, x_2^B) = \left(-\frac{3\bar{s}}{10}, \frac{2\bar{s}}{5}\right)$ on path. Plugging in these trades, dealer A incurs trading costs of $-\frac{7\bar{s}^2}{100}$ if he wins. So the refinement described in Section 3.1 requires $b_{\bar{s}}^A = \frac{7\bar{s}^2}{100}$ to be his bid.

Case 4: $(e^A, e^B) = (1,-1)$ and $s = -\bar{s}$. Assuming that $x_2^A = -\frac{\bar{s}}{2} - x_1^A$ (which ensures that dealer A's final inventory just meets the constraint $e^A + x_1^A + x_2^A - s \leq 1$), and ignoring all other constraints on final inventory, dealers A and B respectively minimize

$$(x_1^A + x_1^B)x_1^A + \left(x_1^B + x_2^B - \frac{\bar{s}}{2}\right)\left(-\frac{\bar{s}}{2} - x_1^A\right)$$
$$(x_1^A + x_1^B)x_1^B + \left(x_1^B + x_2^B - \frac{\bar{s}}{2}\right)x_2^B,$$

leading to $x_2^B = -\frac{x_1^B}{2} + \frac{\bar{s}}{4}$. Inducting backward, we obtain $(x_1^A, x_1^B) = \left(-\frac{\bar{s}}{10}, -\frac{\bar{s}}{10}\right)$, so that $(x_2^A, x_2^B) = \left(-\frac{2\bar{s}}{5}, \frac{3\bar{s}}{10}\right)$ on path. Plugging in these trades, dealer A incurs trading costs of $\frac{7\bar{s}^2}{50}$ if he wins. So the refinement described in Section 3.1 requires $b_{-\bar{s}}^A = \frac{7\bar{s}^2}{50}$ to be his bid. □

To obtain execution with probability one, Lemma D1 implies that an RFQ that splits the order in this way must entail reserve prices $\bar{b}_{-\bar{s}} \geq \frac{7\bar{s}^2}{18}$ and $\bar{b}_{\bar{s}} \geq \frac{7\bar{s}^2}{18}$. Hence, the minimum procurement cost that can be achieved by such an RFQ is twice this lower bound:

$$\hat{c}_{split} := \frac{7\bar{s}^2}{9}.$$

Finally, observe that $\hat{c}_{split} > \hat{c}_1$. The intuition is that—just as when one dealer is contacted—the client's procurement cost is driven by the worst case. When, for example, the client seeks to buy (i.e., $s = \bar{s}$) the worst case is for both dealers to be initially short (i.e., $e^A = e^B = -1$). In this worst case:

- If one dealer is contacted, then he buys a total of $\bar{s}$ on the market. In doing so, he trades at an even rate, buying $\frac{\bar{s}}{2}$ in each of the two periods, which is the cost-minimizing way to trade under permanent price impact (e.g., Bertsimas and Lo, 1998).

- On the other hand, if the order is split among the two dealers, then each buys a total of $\frac{\bar{s}}{2}$ on the market. However, each front-loads their trading, buying $\frac{\bar{s}}{3}$ in the first period and $\frac{\bar{s}}{6}$ in the second. Thus, the dealers do not collectively act to minimize their aggregate trading cost. This increase in cost is ultimately passed on to the client.

  Why do the dealers not collectively act to minimize their aggregate trading cost? If an individual dealer shifts some volume from the second period into the first period, that raises p_1 but has no effect on p_2. This affects that dealer's trading costs in two ways: (i) it reduces the trading cost for those marginal shares (because of permanent price impact, $p_2 > p_1$), but (ii) it increases the trading cost for the inframarginal shares that were already traded in the first period. In addition, there is a negative externality on the first-period trading costs of the other dealer. But because this externality is not internalized, each dealer frontloads his trading.

  By splitting the order, the client loses the coordination benefits that she would obtain by allocating the entire order to a single dealer and instead "competes against herself."

Because $\hat{c}_{split} > \hat{c}_1$, the client optimally never splits her order in this way. Any RFQ policy that does entail such order splitting could be improved by contacting only one dealer and offering the entire order to him whenever the policy would have called for order splitting. Thus, it was without loss of generality that we did not allow for such order splitting in our baseline analysis.

This result contrasts with empirical findings of Kondor and Pintér (2022) and Czech and Pintér (2022), who investigate how traders split orders (in government bonds and corporate bonds, respectively), particularly when in possession of private information about fundamentals. One difference between our model and their setting is that our model does not feature asymmetric information about fundamentals.

E Uncertainty about trade size

Our baseline analysis assumes away all uncertainty about the size of the client's desired trade (featuring uncertainty only about trade side). As mentioned, this assumption is consistent with the operation of many RFQ platforms, as well as typical practices in much off-platform trading. Nevertheless, trade size may be uncertain in other settings. The simplest model with uncertainty about trade size is one without uncertainty about trade side.[38] This appendix analyzes a version of our model along these lines. Most results are qualitatively similar to what we obtain in our baseline analysis.

Model. Suppose the client desires to sell s shares, where $s \in \{1, 2\}$ and where ϕ_0 is the prior probability that $s = 2$.[39] For simplicity, suppose dealers' initial inventories are perfectly correlated: $e^A = e^B = 0$ with probability $1 - \psi$ and $e^A = e^B = 1$ with probability ψ. And suppose the dealers' position limits require them to keep their final inventory in $[0, 1]$. In all other respects, let the model be as described in the main text.

At a high level, the analysis follows the steps used to prove our main results. We therefore only sketch the analysis below and omit supporting details.

Contacting one dealer. Without loss of generality, suppose the client contacts only dealer A. Suppose the RFQ induces a belief that $s = 2$ with probability ϕ. As in the main text, the client's procurement cost $\hat{c}_1(\phi)$ is set by the worst-case initial inventories for the dealers. This is the case of $e^A = e^B = 1$. In this case, we can derive the on-path trades as

$$(x_1^A, x_2^A, x_2^B) = \begin{cases} (-\frac{1}{2}, -\frac{1}{2}, 0) & \text{if } s = 1, \\ (-1, -1, 0) & \text{if } s = 2, \end{cases}$$

and bids as

$$(b_{s=1}^A, b_{s=2}^A) = \left(\frac{3}{4}, 3\right).$$

Hence, the client optimally sets reserves $(\bar{b}_{s=1}, \bar{b}_{s=2}) = \left(\frac{3}{4}, 3\right)$, and her procurement cost is

$$\hat{c}_1(\phi) = \frac{3}{4}(1 - \phi) + 3\phi.$$

Contacting two dealers. Without loss of generality, suppose dealer A is the winner. Suppose the RFQ induces a belief that $s = 2$ with probability ϕ. In the case where $e^A = e^B = 0$, we can derive the on-path trades as

[38] Although it might be theoretically interesting to have both trade size and side be simultaneously uncertain, doing so would require going beyond a two-point support for s and is beyond the scope of this paper.

[39] Note that this notation differs from the convention that we adopted for the baseline model, where a realization $s > 0$ means the client is a buyer. Here, the client is always a seller.

$$
(x_1^A, x_2^A, x_1^B, x_2^B) = \begin{cases} \left(\frac{\phi}{2(6-\phi)}, -\frac{\phi}{2(6-\phi)}, -\frac{\phi}{6-\phi}, \frac{\phi}{6-\phi}\right) & \text{if } s = 1, \\ \left(-\frac{3-\phi}{2(6-\phi)}, -\frac{9-\phi}{2(6-\phi)}, -\frac{\phi}{6-\phi}, \frac{3}{6-\phi}\right) & \text{if } s = 2, \end{cases}
$$

and bids as

$$
(b_{s=1}^A, b_{s=2}^A) = (b_{s=1}^B, b_{s=2}^B) = \left(-\frac{3\phi^2}{4(6-\phi)^2}, \frac{3(33-4\phi-\phi^2)}{4(6-\phi)^2}\right).
$$

Alternatively, in the case where $e^A = e^B = 1$, we can derive the on-path trades as

$$
(x_1^A, x_2^A, x_1^B, x_2^B) = \begin{cases} \left(-\frac{2-\phi}{6}, -\frac{4+\phi}{6}, -\frac{1+\phi}{3}, \frac{1+\phi}{3}\right) & \text{if } s = 1, \\ \left(-\frac{5-\phi}{6}, -\frac{7+\phi}{6}, -\frac{1+\phi}{3}, \frac{1+\phi}{3}\right) & \text{if } s = 2, \end{cases}
$$

and bids as

$$
(b_{s=1}^A, b_{s=2}^A) = (b_{s=1}^B, b_{s=2}^B) = \left(\frac{12+2\phi-\phi^2}{12}, \frac{43+6\phi-\phi^2}{12}\right).
$$

Hence, the client's procurement cost is

$$
\hat{c}_2(\phi) = \frac{3\phi(33-5\phi)}{4(6-\phi)^2}(1-\psi) + \frac{4+11\phi+\phi^2}{4}\psi.
$$

The client's optimal policy. There are both similarities and differences between the client's optimal policy here (where uncertainty is about size) and her optimal policy in the baseline model (where uncertainty is about side).

The primary difference is that the optimal policy here might take the form of always contacting only a single dealer. For example, this obtains if $\psi > \frac{3}{4}$, because $\hat{c}_2(\phi)$ is then everywhere greater than $\hat{c}_1(\phi)$. This differs from the optimal policy when uncertainty is only about the side of the client's trade, where according to Proposition 5, the client would never optimally use such a policy: there is always some realization of s under which she contacts two dealers with positive probability.

In other respects, the optimal policy is qualitatively similar to what we find in our baseline analysis:

- As in the baseline model, the optimal policy is minimally-disclosing, which follows from the convexity of $\hat{c}_2(\phi)$.
- As in the baseline model, the optimal policy might take the form of always contacting both dealers. For example, this obtains if $\psi < \frac{54}{79}$, because $\hat{c}_2(\phi)$ is then everywhere less than $\hat{c}_1(\phi)$.
- As in the baseline model, the optimal policy might involve mixing over the number of contacted dealers. For example, this may obtain if $\psi \in \left(\frac{54}{79}, \frac{3}{4}\right)$, because $\hat{c}_2(\phi)$ then crosses $\hat{c}_1(\phi)$ (once, from below).
- And if the optimal policy does involve mixing over the number of contacted dealers, then it takes the form of a policy that contacts only a single dealer only if the realized s is such that front-running appears especially likely—in this case, only if $s = 2$ is realized.

F $N > 2$ dealers

Elsewhere in the paper, we assumed that there are only two dealers. In this appendix, we investigate what happens if there are more. For this appendix only, we use N for the total number of dealers, and $N_{rel} \leq N$ for the number of dealers with whom the client has a relationship.[40] We also use superscripted numbers (rather than superscripted letters) to index dealers. As elsewhere in the paper, M is the number of dealers that the client chooses to contact, and she is constrained to choose $M \in \{1, \ldots, N_{rel}\}$. Our baseline analysis is the case of $N = N_{rel} = 2$. Here, we consider arbitrary $N \geq N_{rel} \geq 2$.

Our baseline analysis shows that when $N = N_{rel} = 2$, the optimal policy does not always entail contacting all of the available dealers. One purpose of this appendix is to demonstrate that this extends. That is, given any $N \geq N_{rel} \geq 2$, we show that there exist values for the parameters (ϕ_0, ψ, ρ) under which the optimal policy entails choosing $M < N_{rel}$ with positive probability.

F.1 The case of $\rho = 1$

We can provide a full characterization for the case when the initial dealer inventories are perfectly correlated (i.e., $\rho = 1$). At a high level, the analysis follows the steps used to prove our main results. We therefore only sketch the analysis below and omit supporting details.

Suppose the client contacts the dealers $\{1, \ldots, M\}$. Without loss of generality, suppose dealer 1 is the winner. Suppose the RFQ induces a belief that $s = \bar{s}$ with probability ϕ. There are two cases for the dealer inventories.

All dealers are initially long. In the case where all N dealers have initial inventory of 1, we can derive the on-path trades as:

$$(x_1^1, x_2^1) = \begin{cases} \left(\frac{(M^3+M^2-3M+1)(1-\phi)\bar{s}}{2[M(M+1)(M+2)-(M-1)\phi]}, \frac{(M^2-1)(1-\phi)\bar{s}}{2[M(M+1)(M+2)-(M-1)\phi]} \right) & \text{if } s = \bar{s}, \\ \left(-\frac{[2M^2+4M+(M^3+M^2-3M+1)\phi)]\bar{s}}{2[M(M+1)(M+2)-(M-1)\phi]}, -\frac{[2M^3+4M^2-(M^3+M^2-M-1)\phi)]\bar{s}}{2[M(M+1)(M+2)-(M-1)\phi]} \right) & \text{if } s = -\bar{s}, \end{cases}$$

for $i \in \{2, \ldots, M\}$:

$$(x_1^i, x_2^i) = \begin{cases} \left(-\frac{M(M+2)(1-\phi)\bar{s}}{M(M+1)(M+2)-(M-1)\phi}, \frac{(M^2-1)(1-\phi)\bar{s}}{2[M(M+1)(M+2)-(M-1)\phi]} \right) & \text{if } s = \bar{s}, \\ \left(-\frac{M(M+2)(1-\phi)\bar{s}}{M(M+1)(M+2)-(M-1)\phi}, \frac{M(M+2)(1-\phi)\bar{s}}{M(M+1)(M+2)-(M-1)\phi} \right) & \text{if } s = -\bar{s}, \end{cases}$$

for $i \in \{M+1, \ldots, N\}$:

$$(x_1^i, x_2^i) = \begin{cases} (0,0) & \text{if } s = \bar{s}, \\ (0,0) & \text{if } s = -\bar{s}. \end{cases}$$

[40] In principle, N_{rel} could be endogenized via relationship-formation costs, but for purposes of this discussion, treat N_{rel} as an exogenous parameter.

If $M = 1$, then dealer 1 bids $(b^1_{-\bar{s}}, b^1_{\bar{s}}) = \left(\frac{3\bar{s}^2}{4}, 0\right)$. If $M \geq 2$, then all contacted dealers $i \in \{1, \ldots, M\}$ bid

$$b^i_{-\bar{s}} = \left(1 - \frac{(M^3 + 3M^2 + M + 1)\phi\big[2M(M+2) + (M-1)(M^2 + 2M - 1)\phi\big]}{4\big[M(M+1)(M+2) - (M-1)\phi\big]^2}\right)\bar{s}^2,$$

$$b^i_{\bar{s}} = -(M^3 + 3M^2 + M + 1)(M-1)\left(\frac{(M+1)(1-\phi)}{2\big[M(M+1)(M+2) - (M-1)\phi\big]}\right)^2 \bar{s}^2.$$

All dealers are initially short. The continuation equilibrium is symmetric to the case where all dealers are initially long. In particular, if $M = 1$, then dealer 1 bids $(b^1_{-\bar{s}}, b^1_{\bar{s}}) = \left(0, \frac{3\bar{s}^2}{4}\right)$. If $M \geq 2$, then all contacted dealers $i \in \{1, \ldots, M\}$ bid

$$b^i_{-\bar{s}} = -(M^3 + 3M^2 + M + 1)(M-1)\left(\frac{(M+1)\phi}{2\big[M(M+1)(M+2) - (M-1)(1-\phi)\big]}\right)^2 \bar{s}^2,$$

$$b^i_{\bar{s}} = \left(1 - \frac{(M^3 + 3M^2 + M + 1)(1-\phi)\big[2M(M+2) + (M-1)(M^2 + 2M - 1)(1-\phi)\big]}{4\big[M(M+1)(M+2) - (M-1)(1-\phi)\big]^2}\right)\bar{s}^2.$$

The procurement cost $\hat{c}_M(\phi)$. From these bids, we can compute the client's procurement cost. As in our baseline analysis, $\hat{c}_1 = \frac{3\bar{s}^2}{4}$. For $M \geq 2$,

$$\hat{c}_M(\phi) = \left[\frac{(1-\phi)(4-\phi)\bar{s}^2}{4} + \frac{(M-1)^2\phi(1-\phi)^3\bar{s}^2}{4\big[M(M+1)(M+2) - (M-1)\phi\big]^2}\right]\psi + \left[\frac{\phi(3+\phi)\bar{s}^2}{4} + \frac{(M-1)^2(1-\phi)\phi^3\bar{s}^2}{4\big[M(M+1)(M+2) - (M-1)(1-\phi)\big]^2}\right](1-\psi).$$

If $M = 2$, this expression for $\hat{c}_2(\phi)$ reduces to the one obtained in our baseline analysis, which was given in equation (10).

To determine the client's optimal policy, it is important to understand the relationship among the functions $\big(\hat{c}_1, \hat{c}_2(\phi), \hat{c}_3(\phi), \ldots\big)$. It follows from the expression above that we have the following:

- The sequence of functions $\big(\hat{c}_2(\phi), \hat{c}_3(\phi), \ldots\big)$ is monotonically decreasing. That is, for all $M \geq 2$ and all $\phi \in [0, 1]$, $\hat{c}_M(\phi) > \hat{c}_{M+1}(\phi)$.
- Quantitively, the functions $\big(\hat{c}_2(\phi), \hat{c}_3(\phi), \ldots\big)$ are all extraordinarily close to each other, in this case of $\rho = 1$.

  To quantify this, we calculate $\hat{c}_\infty(\phi) := \lim_{M\to\infty} \hat{c}_M(\phi)$.[41] Comparing against $\hat{c}_2(\phi)$, as given

[41] We can compute

$$\hat{c}_\infty(\phi) = \frac{(4-\phi)(1-\phi)\bar{s}^2}{4}\psi + \frac{\phi(3+\phi)\bar{s}^2}{4}(1-\psi).$$

by equation (10), we obtain

$$\hat{c}_2(\phi) - \hat{c}_\infty(\phi) = \frac{\phi(1-\phi)^3 \bar{s}^2}{4(24-\phi)^2}\psi + \frac{(1-\phi)\phi^3 \bar{s}^2}{4(23+\phi)^2}(1-\psi),$$

which implies the bound

$$\sup_{\phi \in [0,1]} |\hat{c}_2(\phi) - \hat{c}_\infty(\phi)| \leq 0.000047\bar{s}^2,$$

which is small relative to the potential difference between $\hat{c}_1$ and $\hat{c}_2(\phi)$, as well as the stakes involved in information design.

- All elements of the sequence $\big(\hat{c}_2(\phi), \hat{c}_3(\phi), \dots\big)$ are strictly convex.

The client's optimal policy. From the facts listed above, we can deduce the structure of the client's optimal policy, simply by extending arguments made in our baseline analysis.

- First, as in the baseline analysis, convexity of the functions $\big(\hat{c}_1, \hat{c}_2(\phi), \hat{c}_3(\phi), \dots\big)$ implies that a minimally-disclosing RFQ policy is optimal.

- Second, the fact that $\big(\hat{c}_2(\phi), \hat{c}_3(\phi), \dots\big)$ is a monotonically decreasing sequence implies that the client will assign zero probability to all choices of M with the possible exceptions of $M = 1$ and $M = N_{rel}$.

  We stress, however, that this implication is quite fragile, precisely because the functions $\big(\hat{c}_2(\phi), \hat{c}_3(\phi), \dots\big)$ are all extraordinarily close. In particular, suppose we were to augment our model with an explicit marginal cost that the client must pay for each additional dealer that she contacts (in addition to our already-modeled, implicit costs of information leakage). Even if very small, such a cost could be enough to completely reverse the client's preference ranking over $(2, 3, \dots, N_{rel})$ dealers in this case of $\rho = 1$.[42] The client would then assign zero probability to all choices of $M > 2$, even when (potentially many) more dealers than two are available to her.

- Third, define $\underline{\phi}$ and $\overline{\phi}$ as in Definition 1, except using $\hat{c}_{N_{rel}}(\phi)$ in place of $\hat{c}_2(\phi)$.[43]

- Fourth, under the optimal policy, the probabilities of contacting only a single dealer are

$$(q_{-\bar{s}}, q_{\bar{s}}) = \begin{cases} \left(\frac{\underline{\phi}-\phi_0}{\underline{\phi}(1-\phi_0)}, 0\right) & \text{if } \phi_0 \in [0, \underline{\phi}), \\ (0,0) & \text{if } \phi_0 \in [\underline{\phi}, \overline{\phi}], \\ \left(0, \frac{\phi_0 - \overline{\phi}}{\phi_0(1-\overline{\phi})}\right) & \text{if } \phi_0 \in (\overline{\phi}, 1], \end{cases}$$

[42] In a similar spirit, Riggs, Onur, Reiffen and Zhu (2020) employ an explicit marginal cost to generate an interior optimum in their model.

[43] That is: If $\hat{c}_{N_{rel}}(0) \leq \hat{c}_1$, define $\underline{\phi} = 0$; if $\hat{c}_{N_{rel}}(1) - \hat{c}'_{N_{rel}}(1) \geq \hat{c}_1$, define $\underline{\phi} = 1$; otherwise, define it implicitly as the unique $\underline{\phi} \in (0,1)$ that solves $\hat{c}_{N_{rel}}(\phi) - \phi\hat{c}'_{N_{rel}}(\phi) = \hat{c}_1$. If $\hat{c}_{N_{rel}}(1) \leq \hat{c}_1$, define $\overline{\phi} = 1$; if $\hat{c}'_{N_{rel}}(0) + \hat{c}_{N_{rel}}(0) \geq \hat{c}_1$, define $\overline{\phi} = 0$; otherwise, define it implicitly as the unique $\overline{\phi} \in (0,1)$ that solves $(1-\phi)\hat{c}'_{N_{rel}}(\phi) + \hat{c}_{N_{rel}}(\phi) = \hat{c}_1$.

with all of the available N_{rel} dealers being contacted otherwise.

Summarizing. We set out to show that, given any $N \geq N_{rel} \geq 2$, there exist values for the parameters (ϕ_0, ψ, ρ) under which the optimal policy entails choosing $M < N_{rel}$ with positive probability. We have now shown this. Given any $N \geq N_{rel} \geq 2$, let, for example, $\psi = 0.85$ and $\rho = 1$ (as in Figure 3). With these values, we obtain $\underline{\phi} > 0$, hence we can choose a $\phi_0 \in [0, \underline{\phi})$. Under this parametrization, the optimal policy entails contacting only a single dealer with positive probability.

Discussion. The discussion following Lemma 2 observed that the difference between $\hat{c}_1$ and $\hat{c}_2(\phi)$ can be decomposed into three effects: a competition effect, a sampling effect, and a front-running effect. Similar thinking applies to the implications of increasing M above two. Given the Bertrand nature of competition in our setting, the competition effect has no additional bite when M is increased above two. Moreover, in this case of $\rho = 1$, the sampling effect is not at play. Nevertheless, a version of the front-running effect does remain at play. Front-running, however, creates two effects, which work in opposite directions:

1. As M increases, the aggregate amount of front-running increases, which causes the winning dealer's trading costs to increase. All else equal, this raises bids.
2. As M increases above two, competition to front-run intensifies among the losing dealers, which causes each losing dealer's trading profits to decrease. This lowers the opportunity cost of winning and, all else equal, lowers bids.

Quantitatively, these two effects come extremely close to cancelling each other out, so that an increase in M has a nearly-negligible effect on the client's transaction costs. Yet, from the fact that $\big(\hat{c}_2(\phi), \hat{c}_3(\phi), \dots\big)$ is a monotonically decreasing sequence, we conclude that the second effect slightly dominates the first.

F.2 The case of $\rho < 1$

Another question concerns what would be different if $\rho < 1$. Unfortunately, the analysis is significantly less tractable, and we have therefore not been able to provide a full characterization for this case. Nevertheless, we can provide some intuition.

If $\rho < 1$, then the sampling effect would come into play. Intuitively, this would lead $\hat{c}_M(\phi)$ to decrease more rapidly in M. Hence, we would no longer expect the functions $\big(\hat{c}_2(\phi), \hat{c}_3(\phi), \dots\big)$ to be all extraordinarily close to each other. In fact, we would expect $\lim_{M\to\infty} \hat{c}_M(\phi) = 0$. Of course, N_{rel} imposes a cap on M, so the client cannot literally send M to infinity.

G The tradeoff between execution probability and expected procurement cost

In our main analysis, we characterized the "optimal" RFQ policy, which we defined as the policy that minimizes the client's expected procurement cost subject to obtaining execution with probability one. As observed in footnote 11, this policy can be thought of as one point on the frontier of the tradeoff between execution probability and expected procurement cost. In particular, it is the point on the frontier that the client would prefer if she had lexicographic preferences (or, equivalently, infinite disutility in the event that execution fails). This appendix provides a further discussion of this frontier.

An additional assumption. For this appendix, suppose $\phi_0 = 1$. (The case of $\phi_0 = 0$ would be symmetric.) In this boundary case, the analysis simplifies in several ways. First, we need only concern ourselves with the reserve price for the scenario in which the client seeks to buy (i.e., $\bar{b}_{\bar{s}}$, henceforth simply $\bar{b}$), as well as what the dealers bid for that scenario. Second, we need not consider how the losing dealer updates his beliefs about the client's trading direction in response to, e.g., which dealer wins the RFQ (as in our main analysis), or whether the RFQ features a winner at all (which would be new relative to our main analysis).

If the client contacts $M = 1$ dealers. In this case, the procurement cost (conditional on execution) is simply $\bar{b}$. Moreover, it follows from Lemma A1 that the probability of execution as a function of $\bar{b}$ is

$$\begin{cases} 1 & \text{if } \bar{b} \geq \frac{3\bar{s}^2}{4}, \\ 1 - (1-\psi)[1-\psi(1-\rho)] & \text{if } \bar{b} \in [\frac{7\bar{s}^2}{16}, \frac{3\bar{s}^2}{4}), \\ \psi & \text{if } \bar{b} \in [0, \frac{7\bar{s}^2}{16}), \\ 0 & \text{if } \bar{b} < 0. \end{cases}$$

If the client contacts $M = 2$ dealers. We begin by separately considering the various cases for the dealers' initial inventories.

- Suppose $(e^A, e^B) = (1, 1)$. In this case, each dealer's trading cost is zero if he wins, zero if the other dealer wins, and zero if neither wins. Thus, each dealer has a dominant strategy—to bid zero—regardless of the reserve price.

- Suppose $(e^A, e^B) = (-1, -1)$. In this case, each dealer's trading cost is $\frac{3\bar{s}^2}{4}$ if he wins, $-\frac{\bar{s}^2}{4}$ if the other dealer wins, and zero if neither wins. Although neither dealer has a dominant strategy, we can still make a prediction about bidding behavior. In our main analysis, we focused on the case in which the reserve price is $\bar{b} \geq \bar{s}^2$—in this case, our refinement tells us that each bidder should bid $\bar{s}^2$. We must now consider what happens for lower reserve prices. If the reserve price is $\bar{b} < \frac{3\bar{s}^2}{4}$, then neither bidder ever meets the reserve in equilibrium (as bidding below the reserve price is a weakly dominated strategy). Somewhat more complex is the situation in which the reserve price is $\bar{b} \in [\frac{3\bar{s}^2}{4}, \bar{s}^2)$—in this case, symmetric equilibrium entails bids of

$$\begin{cases} \bar{b} & \text{w.p. } \frac{4\bar{b}-3\bar{s}}{2\bar{b}-\bar{s}}, \\ \bar{s}^2 & \text{w.p. } \frac{2\bar{s}-2\bar{b}}{2\bar{b}-\bar{s}}. \end{cases} \tag{32}$$

Note that when bidding above the reserve price, the exact bid does not matter. In equation (32) we have put $\bar{s}^2$, but any other bid above $\bar{b}$ would work equally well. Note also that asymmetric equilibria exist as well. Yet, owing to our solution concept (*cf.* Section 3.1), we focus on the symmetric equilibrium.

- Suppose $(e^A, e^B) = (1, -1)$. The case of $(e^A, e^B) = (-1, 1)$ is symmetric. In this case, dealer A's trading costs are zero if he wins, $-\frac{\bar{s}^2}{4}$ if dealer B wins, and zero if neither dealer wins. Dealer B's trading costs are $\frac{7\bar{s}^2}{16}$ if he wins, zero if dealer A wins, and zero if neither dealer wins. Thus, dealer B has a dominant strategy—to bid $\frac{7\bar{s}^2}{16}$—regardless of the reserve price. Given that, the natural outcome is that dealer A bids $\frac{\bar{s}^2}{4}$ if the reserve price is $\bar{b} \geq \frac{7\bar{s}^2}{16}$, and he bids zero if the reserve price is $\bar{b} < \frac{7\bar{s}^2}{16}$.

Assembling the cases above, the client's probability of execution when contacting $M = 2$ dealers is

$$\begin{cases} 1 & \text{if } \bar{b} \geq \bar{s}^2, \\ 1 - (1-\psi)[1-\psi(1-\rho)]\left(\frac{2\bar{s}-2\bar{b}}{2\bar{b}-\bar{s}}\right)^2 & \text{if } \bar{b} \in [\frac{3\bar{s}^2}{4}, \bar{s}^2), \\ 1 - (1-\psi)[1-\psi(1-\rho)] & \text{if } \bar{b} \in [0, \frac{3\bar{s}^2}{4}), \\ 0 & \text{if } \bar{b} < 0. \end{cases}$$

And her expected procurement cost (conditional on execution) is

$$\begin{cases} \bar{s}^2(1-\psi)[1-\psi(1-\rho)] + \frac{7\bar{s}^2}{16}[2\psi(1-\psi)(1-\rho)] & \text{if } \bar{b} \geq \bar{s}^2, \\ \frac{\bar{b}(1-\psi)[1-\psi(1-\rho)]\left[1-\left(\frac{2\bar{s}-2\bar{b}}{2\bar{b}-\bar{s}}\right)^2\right]+\frac{7\bar{s}^2}{16}[2\psi(1-\psi)(1-\rho)]}{1-(1-\psi)[1-\psi(1-\rho)]\left(\frac{2\bar{s}-2\bar{b}}{2\bar{b}-\bar{s}}\right)^2} & \text{if } \bar{b} \in [\frac{3\bar{s}^2}{4}, \bar{s}^2), \\ \frac{\frac{7\bar{s}^2}{16}[2\psi(1-\psi)(1-\rho)]}{1-(1-\psi)[1-\psi(1-\rho)]} & \text{if } \bar{b} \in [\frac{7\bar{s}^2}{16}, \frac{3\bar{s}^2}{4}), \\ \frac{\bar{b}[2\psi(1-\psi)(1-\rho)]}{1-(1-\psi)[1-\psi(1-\rho)]} & \text{if } \bar{b} \in [0, \frac{7\bar{s}^2}{16}). \end{cases}$$

The frontier of the tradeoff. In general, here is how to construct the frontier of the tradeoff between execution probability and expected procurement cost. Draw a graph with the probability of execution on the horizontal axis and the negative of expected procurement cost (conditional on execution) on the vertical axis. Plot all the pairs that can be obtained by varying $\bar{b}$, both in the $M = 1$ case and the $M = 2$ case. The upper-right envelope is the frontier of the tradeoff. As mentioned, our main analysis can be thought of as characterizing the point on the extreme right of the frontier. Note that the nature of these graphs will depend on ψ and ρ.

One special case worth considering is the one corresponding to Figure 3: $\psi = 0.85$ and $\rho = 1$. Figure 5(a) presents the aforementioned plots for that case. The left panel depicts the pairs that can be obtained by varying $\bar{b}$, both in the $M = 1$ case and the $M = 2$ case. The right panel depicts the frontier, computed as the upper-right envelope of the pairs depicted in the left panel. To provide a sense of the other possibilities, Figure 5(b) presents what would happen if we instead had $\psi = 0.15$ and $\rho = 1$.[44]

Figure 5: Execution probability versus expected procurement cost

(a) Dealers likely long, while the client is a buyer ($\psi = 0.85$, $\rho = 1$, $\phi_0 = 1$)

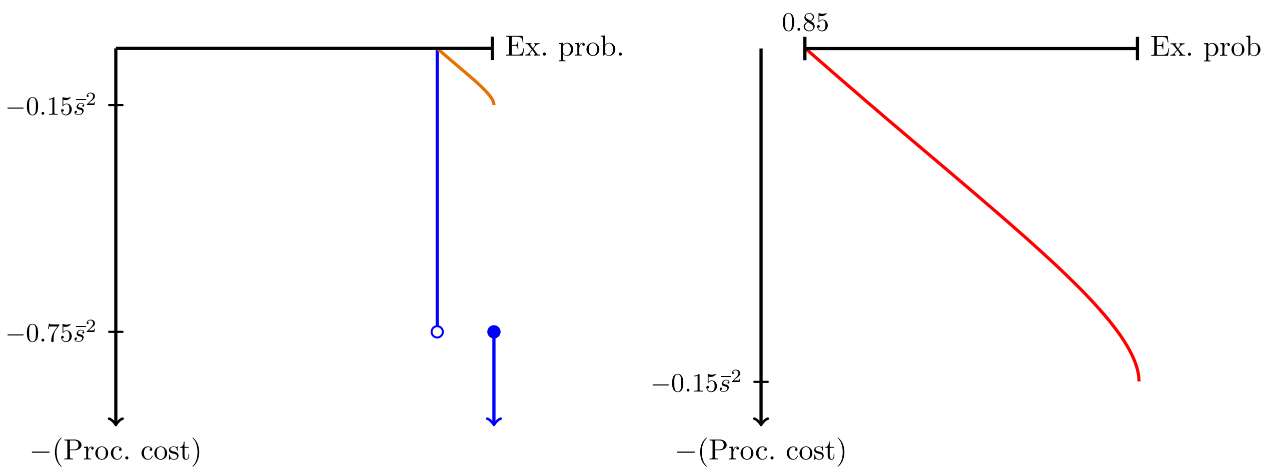

(b) Dealers likely short, while the client is a buyer ($\psi = 0.15$, $\rho = 1$, $\phi_0 = 1$)

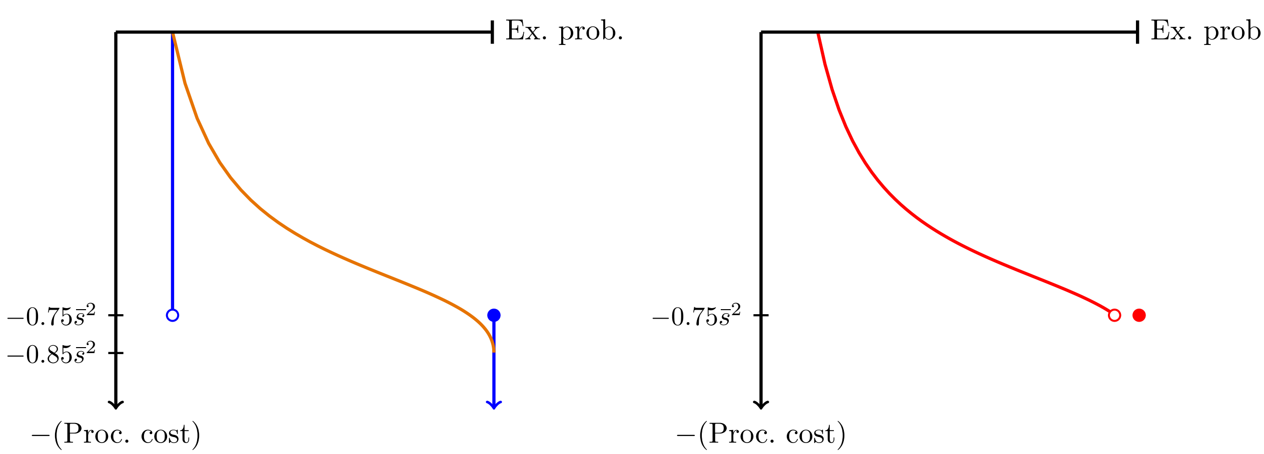

Vertical axes measure the negative of expected procurement cost (conditional on execution). Horizontal axes measure execution probability and depict the unit interval (except for the upper-right panel, in which the interval is truncated). In the left panels, the blue (orange) depicts the outcomes achievable by varying $\bar{b}$ when contacting $M = 1$ ($M = 2$) dealers. The right panels depict the frontier of the tradeoff (in red).

[44] Note that by symmetry, this is equivalent to what we would have if we had maintained the parametrization from Figures 3 and 3′ (i.e., $\psi = 0.85$ and $\rho = 1$) but had instead assumed $\phi_0 = 0$ at the outset.

The client's preferred point on the frontier. Clearly, feasible points off the frontier are dominated for the client—but what is the client's preferred point among those on the frontier? A natural way to think about this would be to suppose that the client has preferences over the two dimensions depicted: execution probability and procurement cost (conditional on execution). These preferences would generate indifference curves in this space, the slope of which would dictate the client's preferred point.

Our main analysis characterizes the cost-minimizing policy among those that lead to execution with probability one, i.e., the rightmost point on this frontier. As observed in footnote 11, this is the point on the frontier that the client would prefer if she had lexicographic preferences: first maximizing execution probability, then minimizing procurement costs. Figure 5(a) illustrates that lexicographic preferences may be necessary for this argument: with the parametrization underlying Figure 5(a), the frontier's slope diverges as we approach an execution probability of one. This is because when contacting $M = 2$ dealers, starting from $\bar{b} = 1$, a small decrease in $\bar{b}$ generates a first-order improvement in procurement cost but only a second-order loss in execution probability.[45] On the other hand, Figure 5(b) illustrates that lexicographic preferences are not always necessary to ensure that the client prefers the rightmost point on the frontier.

[45]To explain further, focus on what happens when $(e^A, e^B) = (-1, -1)$. Starting from $\bar{b} = 1$, a small decrease in $\bar{b}$ leads to a first-order change in the probability that each individual dealer bids above the reserve price (*cf.* equation (32)). Hence, we obtain a second-order change in the probability that both dealers bid above the reserve price.

H Supporting details for footnote 3

In footnote 3, we write the following:

> Based on a variety of sources, we conservatively compute the fraction of trading that takes place via RFQ platforms to be 31% in U.S. treasuries, 22.5% in U.S. corporate bonds, 14.3% in swaps, 60% in options, 21.6% in foreign exchange, and 50% in European exchange-traded products. Many of these asset classes feature separate dealer-to-client and dealer-to-dealer segments; in such cases, RFQ platforms typically account for even larger fractions of the dealer-to-client trading. See Online Appendix H for details.

This appendix details how we obtained the percentages mentioned in that footnote.

Treasuries. According to SIFMA (2023), 31% of U.S. Treasury trading is done via RFQ.

Corporate bonds. According to SIFMA (2023), average daily volume is \$26B for U.S. investment-grade corporate bonds and \$10B for U.S. high-yield corporate bonds. Moreover, the percentage traded electronically is 40% for investment-grade bonds and 31% for high-yield bonds. Taking a weighted average of investment-grade and high-yield, this implies that the percentage traded electronically is 37.5% for corporate bonds overall.

SIFMA (2023) also reports that RFQ (i.e., disclosed RFQ plus anonymous RFQ) accounts for 60% of electronic trading in corporate bonds.

It follows from these statistics that RFQ accounts for $37.5\% \times 60\% = 22.5\%$ of all U.S. corporate bond volume.

Swaps. Table 1 of Riggs, Onur, Reiffen and Zhu (2020) reports that the average daily trading volume of index credit default swaps was \$18.6B in May 2016. Multiplying by the 21 trading days of that month, we obtain a total monthly volume of \$390.1B.

Table 2 of Riggs, Onur, Reiffen and Zhu (2020) then focuses on a subset of the index CDS contracts, a subset of the trading venues (i.e., only the Bloomberg and Tradeweb SEFs), and a subset of trades. In this subsample, the total monthly volume accounted for by RFQs was \$56B.

These statistics imply a lower bound on the fraction of overall volume accounted for by RFQs of $\frac{\$56\text{B}}{\$390.1\text{B}} = 14.3\%$.

Options. Traders Magazine (2019) says, "roughly 60% of all executed options – regardless of asset class – are traded as spreads that were RFQd."

Foreign exchange. Hau, Hoffmann, Langfield and Timmer (2021) say, "we examine trades on multidealer electronic trading platforms (henceforth "platforms"), which enable clients to request quotes from multiple dealers simultaneously rather than individual dealers sequentially. . . just under 40% of all trades are executed on a platform." In their working paper version, their language is slightly less ambiguous; they write "just under 40% of all trades are executed through RFQ platforms" (Hau, Hoffmann, Langfield and Timmer, 2019). More specifically, their sample encompasses dealer-to-client trades of USD/EUR forwards between April 1, 2016 and March 31, 2017.

In addition, the dealer-to-client segment was approximately 54% of the overall foreign exchange market in 2016, according to Panel C in Graph 1 of BIS (2022).

Together, these two datapoints suggest that RFQ platforms account for $54\% \times 40\% = 21.6\%$ of overall volume. Note also that there is a sense in which this is a lower bound, because this calculation presumes that RFQ platforms do not account for any volume in the dealer-to-dealer segment.

European exchange-traded products. Bloomberg Intelligence (2023) says, "Request for quote (RFQ) platforms have grown into the dominant mechanism for ETP trading in Europe. RFQ services via multilateral trading facilities (MTFs) from Bloomberg and Tradeweb combined for about 50% of total traded value in 2022."

I Testable implications: measurement

Section 3.6 describes several empirical implications of our theory. This appendix discusses how to construct the variables that would be needed for testing those implications in the context of RFQ platforms.

The client's procurement cost. To measure procurement costs, one would compare the client's transaction price against a benchmark price. As appropriate, this benchmark could be based on the most recent trade on the opposite side (Riggs, Onur, Reiffen and Zhu, 2020), the last trade in the interdealer market (Hendershott and Madhavan, 2015; O'Hara and Zhou, 2021), the Bloomberg mid-quote (Allen and Wittwer, forthcoming), the previous day's Bank of America Merrill Lynch end-of-day quote (Hendershott, Li, Livdan, Schürhoff and Venkataraman, 2022), or the Markit mid-quote (Collin-Dufresne, Junge and Trolle, 2020).

Information disclosure. One measure of information disclosure is whether one-sided quotes (more disclosure) or two-sided quotes (less disclosure) are requested. Variation could exist within-platform (for platforms that offer a choice between one-sided or two-sided RFQs), across-platform (for platforms that do not offer such a choice), or within-platform-and-across-time (for platforms that changed their protocols).[46]

Number of dealers contacted. The number of dealers contacted is recorded in many datasets on RFQs (e.g., Hendershott and Madhavan, 2015; Riggs, Onur, Reiffen and Zhu, 2020).

The client's trading direction. The prior distribution of the client's trading direction could be proxied by her empirical frequency of buying/selling in historical data. Variation could exist either across-client or within-client-and-across-time (in which case potentially exogenous shocks could originate from mergers or changes in investment mandates). Realized trading directions are recorded in most datasets on RFQs.

Dealers' initial inventories. In some cases, dealers' positions could be observed (e.g., Lyons, 1995; Bjønnes and Rime, 2005; Di Maggio, Kermani and Song, 2017; Coen and Coen, 2022) or estimated (e.g., using the methods of Hansch, Naik and Viswanathan, 1998; Friewald and Nagler, 2016).

Dealers' on-market trading. Dealer trades are recorded in the regulatory TRACE data (used by, e.g., O'Hara and Zhou, 2021).

[46] In some contexts, "RFQ" exclusively indicates a request for one-sided quotes. In those contexts, a request for two-sided quotes is instead called a "request-for-market (RFM)."

References for the online appendix

Allen, Jason and Milena Wittwer, "Centralizing Over-the-Counter Markets?," *Journal of Political Economy*, forthcoming. https://doi.org/10.1086/725361.

Bank for International Settlements, "International Banking and Financial Market Developments," *BIS Quarterly Review*, December 2022. https://www.bis.org/publ/qtrpdf/r_qt2212.pdf.

Bertsimas, Dimitris and Andrew W. Lo, "Optimal Control of Execution Costs," *Journal of Financial Markets*, 1998, *1* (1), 1–50.

Bjønnes, Geir Høidal and Dagfinn Rime, "Dealer Behavior and Trading Systems in Foreign Exchange Markets," *Journal of Financial Economics*, 2005, *75* (3), 571–605.

Bloomberg Intelligence, "Europe's Record ETP Trading Year Leaves Retail on Sidelines," January 2023.

Coen, Jamie and Patrick Coen, "A Structural Model of Liquidity in Over-The-Counter Markets," *Working Paper*, 2022. https://ssrn.com/abstract=4130489.

Collin-Dufresne, Pierre, Benjamin Junge, and Anders B. Trolle, "Market Structure and Transaction Costs of Index CDSs," *The Journal of Finance*, 2020, *75* (5), 2719–2763.

Crawford, Vincent P. and Joel Sobel, "Strategic Information Transmission," *Econometrica*, 1982, *50* (6), 1431–1451.

Czech, Robert and Gábor Pintér, "Informed Trading and the Dynamics of Client-dealer Connections in Corporate Bond Markets," *Working Paper*, 2022. https://ssrn.com/abstract=3744325.

Di Maggio, Marco, Amir Kermani, and Zhaogang Song, "The Value of Trading Relations in Turbulent Times," *Journal of Financial Economics*, 2017, *124* (2), 266–284.

Friewald, Nils and Florian Nagler, "Dealer Inventory and the Cross-Section of Corporate Bond Returns," *Working Paper*, 2016. https://ssrn.com/abstract=2526291.

Hansch, Oliver, Narayan Y. Naik, and S. Viswanathan, "Do Inventories Matter in Dealership Markets? Evidence from the London Stock Exchange," *The Journal of Finance*, 1998, *53* (5), 1623–1656.

Hau, Harald, Peter Hoffmann, Sam Langfield, and Yannick Timmer, "Discriminatory Pricing of Over-the-Counter Derivatives," *Working Paper*, 2019. https://ssrn.com/abstract=3404056.

_ , _ , _ , and _ , "Discriminatory Pricing of Over-the-Counter Derivatives," *Management Science*, 2021, *67* (11), 6660–6677.

Hendershott, Terrence and Ananth Madhavan, "Click or Call? Auction versus Search in the Over-The-Counter Market," *The Journal of Finance*, 2015, *70* (1), 419–447.

_ , Dan Li, Dmitry Livdan, Norman Schürhoff, and Kumar Venkataraman, "Quote Competition in Corporate Bonds," *Working Paper*, 2022. https://ssrn.com/abstract=3982686.

Kamenica, Emir and Matthew Gentzkow, "Bayesian Persuasion," *American Economic Review*, 2011, *101* (6), 2590–2615.

Kondor, Péter and Gábor Pintér, "Clients' Connections: Measuring the Role of Private Information in Decentralized Markets," *The Journal of Finance*, 2022, *77* (1), 505–544.

Lyons, Richard K., "Tests of Microstructural Hypotheses in the Foreign Exchange Market," *Journal of Financial Economics*, 1995, *39* (2), 321–351.

O'Hara, Maureen and Xing Alex Zhou, "The Electronic Evolution of Corporate Bond Dealers," *Journal of Financial Economics*, 2021, *140* (2), 368–390.

Riggs, Lynn, Esen Onur, David Reiffen, and Haoxiang Zhu, "Swap Trading after Dodd-Frank: Evidence from Index CDS," *Journal of Financial Economics*, 2020, *137* (3), 857–886.

Securities Industry and Financial Markets Association, "Understanding Fixed-Income Markets in 2023," May 2023. https://bit.ly/3WrRiMJ.

Traders Magazine, "Are You Ready for RFQs in Electronic Trading?," July 2019. https://bit.ly/3Bic8p1.