The Limits of Multi-Dealer Platforms

Chaojun Wang*

The Wharton School, University of Pennsylvania

Journal of Financial Economics, Volume 149, Issue 3, September 2023

Abstract

On many important multi-dealer platforms, customers mostly request quotes from very few dealers. I build a model of multi-dealer platforms, where dealers strategically choose to respond to or ignore a request. If the customer contacts more dealers, each dealer responds with a lower probability and offers a stochastically worse price when responding. Dealers' strategic avoidance of competition overturns the customer's benefit from potentially receiving more quotes, worsening her best-overall price. In equilibrium, the customer contacts only two dealers. Multi-dealer platforms have limited ability to promote price competition: No design of information disclosure can improve the customer's payoff above this outcome.

JEL Classifications: D44, D47, D83

Keywords: Over-the-counter, multi-dealer platforms, request for quote, pre-trade transparency

*Email: wangchj@wharton.upenn.edu I am grateful for helpful suggestions from Yu An, Briana Chang, William Diamond, Itamar Drechsler, Darrell Duffie, Selman Erol, Daniel Garrett, Simon Gervais, Sergei Glebkin, Vincent Glode, Itay Goldstein, Ben Golub, Joao Gomes, Jing Huang, John Kuong, David Musto, Christian Opp, Nikolai Roussanov, Rakesh Vohra, Lucy White, Yiqing Xing, Yan Xiong, Ming Yang, and Hongda Zhong. Special thanks to Tomy Lee for extensive discussions that helped me better understand the marginal contribution of my paper. I thank Jiaqi Zou for excellent professional editing. I also thank Clarise Huang and Dylan Marchlinski for proofreading the paper.

1 Introduction

Many over-the-counter (OTC) markets now feature trading platforms that allow a customer to simultaneously request quotes from multiple dealers.[1] These multi-dealer platforms have the potential to greatly intensify competition among dealers, provided that customers request quotes from a large number of dealers. On many important multi-dealer platforms, however, customers mostly request quotes from very few dealers, for example from only three on Swap Execution Facilities (SEFs) for index credit default swaps—the minimum required by regulations.[2] Why do the customers contact so few dealers that even a lower bound of three seems to be binding? Because the benefit of contacting more dealers to incite greater competition proposed by existing theory ceases to exist, and in fact becomes a cost, when the dealers are allowed to strategically ignore a customer's request. I add exactly one feature to an otherwise standard model: Dealers can endogenously choose to respond to or ignore a request for quote (RFQ). I show that contacting more dealers, rather than spurring price competition among dealers, actually suppresses competition and leads to worse prices. In equilibrium, the customer chooses to contact only two dealers. More generally, no design of information disclosure about the number of contacted dealers can improve the customer's payoff above this outcome. In this sense, multi-dealer platforms are limited in their ability to promote price competition.

Dealers' ability to ignore a trade request is a natural yet often overlooked feature of OTC trading. Whether an OTC trade is requested on a platform or not, dealers are not forced to

[1] Examples include Bloomberg and Tradeweb for swaps, MarketAxess for bonds, and Refinitiv for currencies.

[2] At the same time, the volume traded on platforms is relatively small in many OTC markets where platform trading is not mandatory. For example, MarketAxess had 15% of the total trade volume for U.S. corporate bonds in 2016Q1, and 20% as of November 2022. Most of its market share growth was gained during the COVID pandemic in 2020.

respond, and indeed, quite often do not. Dealers' revealed preferences show that responding requires effort and cannot be completely cost-free. In my model, the cost of such effort to properly formulate a response can be arbitrarily small. Yet, that dealers endogenously decide whether to respond fundamentally transforms the intended benefit from contacting more dealers into a cost.

The underlying economics comes from the following observation: It is more cost-efficient to concentrate response probabilities among fewer dealers. I consider a simple numerical example to illustrate the cost efficiency of *response concentration*. I suppose that three dealers are contacted in an RFQ, and they each respond with a probability of 70%, 60%, and 50%, respectively. When at least one dealer responds, a trade occurs. Thus, the aggregate expected gain from trade depends on the individual dealer response probabilities through a sufficient statistic $(1-70\%)(1-60\%)(1-50\%)$, which is the probability that no dealer responds to the RFQ. On the other hand, the aggregate expected cost of responding to the RFQ depends on the sum of the response probabilities, $70\%+60\%+50\%$. Keeping the aggregate gain constant, one can reduce the aggregate cost by reducing one dealer's response probability, say, from 50% down to 0%, and raising another dealer's response probability, say, from 60% up to 80%. The adjustment does not change the aggregate gain, because the probability that no dealer responds remains the same, $(1-70\%)(1-80\%)(1-0\%) = (1-70\%)(1-60\%)(1-50\%)$. Yet the aggregate cost declines, $70\% + 80\% + 0\% < 70\% + 60\% + 50\%$. Therefore, it is more cost-efficient to shut down one dealer and concentrate the response probabilities into the remaining dealers.

My benchmark model has one of the most basic structures in economics: A customer chooses to contact a number n of dealers simultaneously in order to purchase an asset. Observing the customer's choice of n, each dealer chooses whether to respond at a cost and what price to offer.

A dealer naturally trades off the response cost and its expected trading profit: It is willing to respond only if the expected profit justifies responding. In the unique symmetric subgame perfect equilibrium, each dealer mixes between responding or not, and offers a distribution of prices conditional on responding. If the customer were to contact more dealers, every contacted dealer would strategically avoid competition by *(i)* responding with a lower probability, and *(ii)* offering a stochastically higher price when it responds. Dealers' strategic avoidance of competition more than offsets the benefit to the customer from potentially receiving more quotes, resulting in *(iii)* a *strictly less* competitive best overall price for the customer. Anticipating this dealer behavior, the customer contacts precisely two dealers in equilibrium.

Moreover, the customer contacts two dealers not only in the symmetric subgame perfect equilibrium, but also in *any* subgame perfect equilibrium. The driving force comes from the cost savings of response concentration. The cost savings, as illustrated in the above numerical example, does not require symmetry in the dealers' strategies.

An assumption of the model is that the number n of dealers contacted in a customer RFQ is disclosed to those n dealers, as is the case on SEFs. I examine alternative platform designs of information disclosure about the number n. Not disclosing any information about the number n would make the customer contact as many dealers as is feasible. However, the customer's payoff decreases relative to her status quo payoff where the number n is fully disclosed, as the customer expects a lower response probability and a worse price offer from each dealer. These predictions match the patterns for corporate bonds on MarketAxess, which does not disclose the number of contacted dealers by default. There, a customer contacts more than 25 dealers per RFQ, on average, and dealers' response rate is only around 25% (Hendershott and Madhavan, 2015). In comparison, on SEFs, dealers' response rate is almost 90% (Riggs, Onur, Reiffen, and Zhu, 2020).

Although making the number n of contacted dealers undisclosed does not lead to more competitive prices, it does cause the customer to contact more dealers in equilibrium. More generally, one may wonder whether some alternative information design that partially discloses the number n could reach the sweet spot—where the customer contacts a high enough number of dealers and the dealers engage in a sufficient level of competition—and make the dealers' best overall price more competitive than the status quo of full disclosure. Unfortunately, the answer is no. No alternative information disclosure design can improve the customer's payoff above the status quo. In this sense, multi-dealer platforms are limited in their ability to promote price competition. The driving force once again comes from the cost savings of response concentration. The cost savings, as illustrated in the above numerical example, does not depend on what information is disclosed about the number n.

The benchmark model also predicts that a dealer responds with a higher probability when facing a larger order or when perceiving a larger gain from trade. Conditional on responding, the dealer offers a stochastically lower price when facing a larger order or when perceiving a smaller gain from trade. The model's predictions are largely consistent with empirical patterns documented in the literature.

The model is also applicable to Bid Wanted in Competition (BWIC).[3] BWIC is a common method for a customer to sell collateralized loan obligations, asset-backed securities, or mortgage-backed securities through an invited auction. Invited dealers do not see how many or which other dealers are invited. Unlike RFQ on a multi-dealer platform, BWIC is not a standardized process with automated information disclosure. As a result, the customer

[3] I thank an anonymous referee for suggesting BWIC as an application.

cannot credibly commit to limit her number of invited dealers. The New York Fed conducted a series of BWIC to sell its residential mortgage-backed securities. Consistent with my model's prediction, a large number of dealers (57 on average) were invited to bid at a BWIC. Armantier and Sporn (2013) provides a detailed description of how BWIC works in practice.

1.1 Literature

I make three contributions. *(i)* My paper introduces response concentration, a new mechanism that fundamentally overturns the competitiveness of the best overall price among more potential contestants. *(ii)* My model proposes a simple explanation for why customers contact very few dealers when the number of contacted dealers is disclosed and many when the number is not disclosed. *(iii)* My results reveal the limits of multi-dealer platforms in promoting price competition: No design of information disclosure can improve a customer's payoff above the status quo of full disclosure, where the customer contacts only two dealers and prices are noncompetitive in equilibrium.

This paper belongs to the recent literature on multi-dealer platforms.[4] Existing theories cannot explain why customers contact very few dealers on some platforms like SEFs and many more on others. Most closely related to this paper are the theories of Riggs et al. (2020) and Baldauf and Mollner (2022). Both papers generate an interior solution for the optimal number of dealers to contact by trading off the benefit of dealer competition against either a direct "relationship cost" of contacting each additional dealer (Riggs et al., 2020)[5] or an indirect cost of front-running (Baldauf and Mollner, 2022). To reconcile the large

[4]Examples include Collin-Dufresne, Junge, and Trolle (2020), Hau, Hoffmann, Langfield, and Timmer (2021), O'Hara and Alex Zhou (2021), Hendershott, Livdan, and Schürhoff (2021), Liu, Vogel, and Zhang (2017), Vogel (2019), and Allen and Wittwer (2023).

[5]Riggs et al. (2020) also features the winner's curse. In their model, "[t]he relationship channel generates an interior solution for the optimal number of dealers requested, and the winner's curse channel generates the comparative statics that [they] eventually test."

disparity in customer behavior on SEFs versus MarketAxess, those papers would require the relationship or the front-running cost to be sufficiently high on one type of platform yet extremely low on the other. My paper questions whether dealers' best overall price indeed becomes more competitive when more potential contestants are present in the first place. Letting dealers strategically decide whether to respond fundamentally overturns their standard pattern of competition in the presence of more rivals. As a result, my paper does not need any added trade-offs as those in Riggs et al. (2020) and Baldauf and Mollner (2022). I provide a parsimonious model that unambiguously explains why customers contact very few dealers when the number of contacted dealers is disclosed to those dealers and many when the number is not disclosed. Moreover, unlike existing theories, my paper goes beyond a specific design of multi-dealer platforms. By exploring general platform designs of information disclosure, my paper contributes a novel result on the limits of multi-dealer platforms in promoting price competition.

Glode and Opp (2020) shows how interacting with more dealers reduces the dealers' incentives to acquire costly expertise. They then show that an OTC market, which limits competition, might socially dominate an exchange market. Dealers in their model do not offer quotes—they receive quotes. The mechanism that pins down pricing in my paper—dealers' strategic avoidance of competition—requires dealers' ability to set quotes, so that they can offer stochastically worse prices if a customer contacts more dealers.

With different focuses, Glebkin, Yueshen, and Shen (2022) and Yueshen (2017) also feature an uncertain number of dealers who respond to a trade request. My model differs from theirs in two respects: *(1)* the number of contacted dealers is endogenously chosen by the customer instead of being exogenously fixed; *(2)* each dealer endogenously mixes between responding or not, instead of having response probabilities that do not depend on any agent's endogenous strategy. These two modeling distinctions are crucial for obtaining my paper's results. In particular, if a dealer's response probability were exogenously fixed at some constant, as in Glebkin et al. (2022) and Yueshen (2017), the customer would contact as many dealers as is feasible, and her price would approach the competitive limit when the pool of potential dealers is large (Appendix A).

This paper is broadly related to the literature on search frictions,[6] market concentration,[7] and sticky relationships[8] in OTC trading. Can these general features of OTC markets explain why a customer contacts very few dealers on multi-dealer platforms? They cannot. First, the very objective of multi-dealer platforms is to reduce search frictions by making it easy for customers to reach out to many dealers at once. Second, these general arguments cannot explain why a customer contacts many more dealers on platforms that by default do not disclose the number of contacted dealers, such as MarketAxess. Therefore, these general features of OTC markets cannot be the driving force that determines the extensive margin of *how many* dealers a customer would contact on a multi-dealer platform.[9]

[6]Duffie, Gârleanu, and Pedersen (2005) pioneered the OTC search literature. Examples include Atkeson, Eisfeldt, and Weill (2015), Bethune, Sultanum, and Trachter (2021), Dugast, Üslü, and Weill (2022), Hugonnier, Lester, and Weill (2020), Li, Rocheteau, and Weill (2012), Maurin (2022), Praz (2014), Tsoy (2021), Vayanos and Weill (2008), and Wang (2022) among many others. Weill (2020) reviews the literature of search models in OTC markets.

[7]Most OTC markets exhibit a highly concentrated core-periphery trading network (Abad, Aldasoro, Aymanns, D'Errico, Fache Rousová, Hoffmann, Langfield, Neychev, and Roukny, 2016; Afonso, Kovner, and Schoar, 2014; Bech and Atalay, 2010; Craig and von Peter, 2014; Hollifield, Neklyudov, and Spatt, 2017; in't Veld and van Lelyveld, 2014; King, Osler, and Rime, 2012; Li and Schürhoff, 2019; Peltonen, Scheicher, and Vuillemey, 2014). Theoretical explanations include Chang and Zhang (2022), Farboodi, Jarosch, and Shimer (2022), Sambalaibat (2022), Üslü (2019), and Wang (2016).

[8]Examples include Di Maggio, Kermani, and Song (2017) and Hendershott, Li, Livdan, and Schürhoff (2020).

[9]Riggs et al. (2020) provides evidence that the relationship channel matters for the intensive margin of *which* dealers a customer would contact and *which* dealer would be more likely to offer her a better price, although no evidence is provided for whether relationship quality drives the extensive margin of *how many* dealers a customer would contact.

To the consumer search literature,[10] which also features mixed pricing strategies by firms, this paper contributes a novel model, where the number of responding firms is endogenously determined by the firms' own decisions of whether to respond instead of being exogenously chosen by nature. If one replaces the firms' endogenous decisions by an exogenous availability constraint, then the customer would contact as many firms as is feasible (Appendix A). Another distinction is the absence of a search cost in my paper. The consumer search literature assumes a positive search cost for at least some customers. In my model, the customer chooses to contact only two dealers for prices in the absence of any search cost.

My paper also belongs to the literature on auctions with entry.[11] My paper uncovers the more general underlying economics of response concentration, which works under different model setups. That the seller's revenue decreases with the number of potential bidders in a symmetric equilibrium is one special implication of response concentration. Thanks to its generality, response concentration also drives all my other results: that the seller invites very few potential bidders in *any* equilibrium, and that *no* design of information disclosure can improve the seller's revenue. Further, existing papers are concerned about the auction format—such as second-price auctions and the commitment to a reservation price—which are far away from practical implementation on multi-dealer platforms in real-world financial markets. My model differs in that the number of potential bidders is chosen by the seller, and may be fully disclosed, partially disclosed, or not disclosed to those potential bidders. Thereby, my paper contributes a novel result on the optimal information disclosure about the number of potential bidders, which is assumed to be exogenously fixed in this literature.

Yueshen and Zou (2022) assumes the expected response cost to be convex—as opposed to

[10] Stigler (1961) pioneered the consumer search literature. Examples include Varian (1980), Burdett and Judd (1983), Stahl (1989), and Lester (2011).

[11] Examples include McAfee and McMillan (1987), Engelbrecht-Wiggans (1987), Levin and Smith (1994), Menezes and Monteiro (2000), and Jovanovic and Menkveld (2022).

linear—in the response probability, and shows that the customer continues to contact very few dealers. They do not consider the general design of information disclosure.

1.2 Institutional Details

Multi-dealer platforms allow customers to simultaneously request quotes from multiple dealers. On SEFs, the Commodity Futures Trading Commission used to require a customer to contact at least 2 dealers in an RFQ. On average, a customer contacted 2.9 dealers per request (McPartland, 2014). The lower bound then increased by 1, to 3 dealers per request, in 2014. After the change, Riggs et al. (2020) finds that customers most frequently contact only 3 dealers, on average contact 4 (one more than before the change), and rarely contact more than 5. If weighted by notional quantity, customers contact even fewer dealers on average, because larger orders tend to be exposed to fewer dealers. These facts suggest that the lower bound of 3 is most often binding for customers. Dealers' response rate is slightly below 90% on SEFs (Riggs et al., 2020, Table 3). These are in contrast to MarketAxess' dramatically higher (25+) number of contacted dealers—which one naturally expects—and a much lower response rate (around 25%) (Hendershott and Madhavan, 2015, Table VI).

In an RFQ, SEFs disclose the number of dealers contacted to those dealers (Riggs et al., 2020). MarketAxess does not disclose that number by default. A customer can set it to disclose her number of contacted dealers in a global setting, although not on a request-by-request basis. Even when a customer chooses to do so, the disclosed number is greatly contaminated by the customer's possible selection of Open Trading (OT). Since 2012, MarketAxess' RFQ platform has an OT option that allows a customer to simultaneously send an RFQ to buy-side firms and non-permissioned dealers, in addition to the dealers that the customer chooses to contact (MarketAxess, 2022). Selecting OT would send the RFQ to hundreds if not thousands of other platform participants. However, MarketAxess increases the number of contacted participants by only one if the customer selects OT, and does not indicate whether the customer selected OT. In practice, OT is selected more than 90% of the time (Hendershott et al., 2021).

An advantage of my model is that it does not require me to take a stance on how large dealers' response cost is empirically. All my results hold, even with an infinitesimally small response cost. A dealer's cost of responding to a customer's RFQ can arise from the effort and the resources to evaluate the asset, the customer, market conditions, and the dealer's own inventory—activities that are necessary prior to forming a price offer. Each RFQ is unique, requesting to trade a specific asset at a specific time under a specific market condition. Even if the same dealer receives two RFQs for the same asset within a short period of time, the two RFQs would most likely come from different customers, request to trade different sizes, and have other distinctive characteristics. Such RFQs would still receive different prices, as price discrimination is a prominent feature of OTC trading: The same dealer typically offers different prices to a hedge fund versus to an insurance company (Ramadorai, 2008; Hau et al., 2021; Bjønnes, Kathitziotis, and Osler, 2015; Pinter, Wang, and Zou, 2021), to customers and trades of different sizes (Pinter, Wang, and Zou, 2020), to customers of different activeness (O'Hara, Wang, and Zhou, 2018), and price discrimination persists on multi-dealer platforms (Hau et al., 2021). Price discrimination can arise from differences in information asymmetry (Lee and Wang, 2018; Pinter et al., 2020, 2021), dealer-customer relationships (Di Maggio et al., 2017; Hendershott et al., 2020), customers' outside options (Hendershott et al., 2020; Pinter et al., 2020), and customers' willingness to pay, etc. As a result, a dealer needs to exert costly effort and resources to separately evaluate each RFQ from each customer, even when the RFQ arrives within a short period of time after another

RFQ for the same asset.

The remainder of the paper is organized as follows. Section 2 sets up the benchmark model. Section 3 solves for the unique symmetric equilibrium in closed form and obtains the main results in any equilibrium. Section 4 examines alternative designs of information disclosure and establishes the limits of multi-dealer platforms. Section 5 derives the models' other empirical predictions. Section 6 concludes.

2 Benchmark Model

2.1 Trading game

The trading game proceeds in three stages. In Stage 1, a customer seeking to buy[12] one unit of an asset chooses a number n of *ex-ante* identical dealers to contact in an RFQ. Observing the customer's choice n, each dealer j chooses whether to respond and what price p_j to offer in Stage 2. The asset's expected payoff is normalized to 0, and the customer has an additional private value v of owning the asset. Responding to the RFQ incurs a cost $c > 0$, which is assumed to be less than the value v ($c < v$) so that there is a positive net surplus from trade. In Stage 3, the customer chooses whether and against which dealer's price to trade. I do not impose any tie-breaking rule in the case of indifference. All agents are risk-neutral. Figure 1 summarizes the timeline of the model.

[12]The case where the customer is seeking to sell is symmetric.

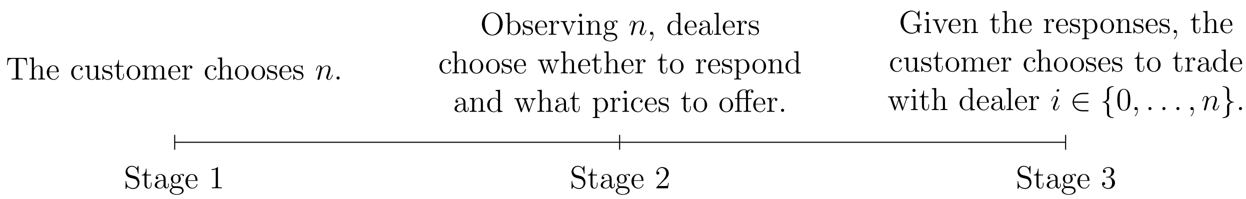

Figure 1: Timeline

2.2 Strategies and equilibrium concept

The customer's strategy consists of a couple (n, i), following which the customer contacts n dealers in Stage 1 and trades with dealer $i \in \{0, 1, \ldots, n\}$ in Stage 3 after receiving the dealers' responses $(p_1, \ldots, p_n)$ in Stage 2. If dealer j chooses to ignore the customer's RFQ in Stage 2, then $p_j = \text{NA}$ by convention. If the customer chooses not to trade in Stage 3, then $i = 0$ by convention. Mixed strategies are allowed. The strategy of dealer j consists of a couple $(a_{j,n}, F_{j,n})$ for each number of dealers n chosen by the customer, where $a_{j,n} \in [0, 1]$ is the probability with which the dealer responds to the RFQ, and $F_{j,n}$ is the CDF of the dealer's price offer if the dealer does respond.

The solution concept is subgame perfect equilibrium. I first solve for the unique *symmetric* subgame perfect equilibrium, where all dealers employ the same strategy (a^*, F^*). Then, I show that the same results hold in *all* subgame perfect equilibria (Theorem 2). In a symmetric subgame perfect equilibrium,

$$\text{(symmetry)} \quad a^*_{j,n} = a^*_n \text{ and } F^*_{j,n} = F^*_n \text{ for every } n \in \mathbb{Z}^{++} \text{ and } j = 1, \ldots, n.$$

The agents' optimality conditions are derived as follows. In Stage 3, the customer chooses to trade with the dealer who offers the lowest price if that price is less than the customer's private value v, and does not trade otherwise:

$$i^* \begin{cases} \in \underset{j=1,\dots,n}{\operatorname{argmin}}\ p_j & \text{if } \min_{j=1,\dots,n} p_j < v, \\ = 0 & \text{if } \min_{j=1,\dots,n} p_j > v, \\ \in \{\underset{j=1,\dots,n}{\operatorname{argmin}}\ p_j, 0\} & \text{if } \min_{j=1,\dots,n} p_j = v. \end{cases} \tag{1}$$

In Stage 2, every price p that belongs to the support of the equilibrium price distribution F_n^* maximizes a dealer's expected payoff given the other dealers' pricing strategy for every $n \geq 1$,

$$p \in \underset{\tilde{p} \in \mathbb{R}}{\operatorname{argmax}} \left(\tilde{p}\ \mathbb{1}_{\tilde{p} \leq v}\ [1 - a_n^* F_n^*(\tilde{p})]^{n-1} \right) \qquad \forall p \in \operatorname{supp} F_n^*. \tag{2}$$

The right-hand side of (2) is the expected trading profit that the dealer maximizes: When offering a price $\tilde{p} \leq v$, the dealer trades with the customer if and only if no other dealer offers a price lower than $\tilde{p}$, an event that occurs with probability $[1 - a_n^* F_n^*(\tilde{p})]^{n-1}$.

The dealer's individual rationality is given by

$$p\ \mathbb{1}_{p \leq v}\ [1 - a_n^* F_n^*(p)]^{n-1} \geq c \qquad \forall p \in \operatorname{supp} F_n^*. \tag{3}$$

That is, the dealer's expected trading profit must be at least as large as its cost c of responding to the RFQ. If the dealer responds with probability $a_n^* < 1$, the dealer must be indifferent between responding or not. In this case, (3) must hold as an equality,

$$\text{if } a_n^* < 1, \qquad p\ \mathbb{1}_{p \leq v}\ [1 - a_n^* F_n^*(p)]^{n-1} = c \qquad \forall p \in \operatorname{supp} F_n^*. \tag{4}$$

In Stage 1, the customer chooses the number of dealers n to maximize her expected payoff,

$$n^* \in \underset{n \in \mathbb{Z}^+}{\operatorname{argmax}} \left[v - \mathbb{E}_{G_n^*}(p \wedge v) \right], \tag{5}$$

$$\text{where } 1 - G_n^*(p) = \left[1 - a_n^* F_n^*(p)\right]^n.$$

Here, $v - \mathbb{E}_{G_n^*}(p \wedge v)$ is the customer's expected payoff upon contacting n dealers, and G_n^* is the CDF of the dealers' best price offer $p = \min_{j=1,\ldots,n} p_j$.

Proposition 0. *A symmetric subgame perfect equilibrium is a strategy profile* (n^*, i^*, a^*, F^*) *such that*

- *the customer's strategy* (n^*, i^*) *satisfies the optimality conditions* (1) *and* (5), *and*

- *all dealers employ the same strategy* (a^*, F^*) *that satisfies the optimality conditions* (2) *to* (4).

2.3 Discussion

Instead of a response cost on dealers, Riggs et al. (2020) imposes a cost on the customer to *contact* each additional dealer. The contact cost mechanically appears in the customer's utility function and directly limits her optimal number of contacted dealers to an interior solution, which depends on the magnitude of the contact cost. The contact cost has to be sufficiently large to explain why the customer contacts very few dealers. In my model, there is no additional cost to contact more dealers. The response cost c is materialized only when a dealer strategically decides to *respond* to the RFQ. Further, the response cost c can be arbitrarily close to 0 or heterogeneous across dealers (as an extension in Appendix C). In both situations, the customer always contacts $n^* = 2$ dealers in equilibrium.

How, then, does my model make the customer contact generically fewer dealers than in Riggs et al. (2020) without using costs that operate directly on the final outcome variable? Because the dealers' cost leads to their *strategic avoidance of competition*, whereas the customer's cost does not. The customer's contact cost is sunk by the time dealers receive an RFQ. As a result, the customer's cost does not appear in any dealer's utility function, and thus does not affect how dealers compete with each other. Therefore, the conventional competition argument applies: Contacting more dealers induces a more competitive best overall price. Contrarily, the dealers' response cost, together with their ability to ignore the RFQ, overturns their standard pattern of competition. The cost incentivizes and the RFQ's ignorability enables the dealers to strategically avoid competition when the customer contacts more dealers, resulting in a *strictly less* competitive best overall price for the customer.

The dealer's response cost also allows my paper to explain the large disparity in customer behavior on SEFs versus MarketAxess based on their different disclosure designs regarding the number n of dealers contacted. With the customer's contact cost alone and not the dealers' response cost, the disclosure design is less relevant.[13] The customer's contact cost serves as a natural commitment device because dealers understand that it is not incentive compatible for the customer to contact many dealers due to her contact cost. To reconcile the large gap between SEFs and MarketAxess, the contact cost would have to be sufficiently high on one type of platform yet extremely low on the other. In contrast, *any* response cost on the dealers—even an infinitesimal one—would push the customer to contact only two dealers when the number n is disclosed, and many dealers when n is not. The disclosure design determines whether the customer can commit to contact few dealers when facing dealers' strategic avoidance of competition. When the number n is disclosed, the customer

[13] I thank an anonymous referee for this sharp observation.

can and will commit to contact very few dealers. Once n becomes undisclosed, the customer loses the ability to credibly commit to do so. Thus, the customer contacts as many dealers as is feasible, and the dealers correctly conjecture so (Section 4.1).

An indispensable ingredient for modeling dealers' strategic avoidance of competition is their ability to endogenously decide whether to respond to an RFQ. Existing works[14] instead assume that every dealer is available with an exogenously fixed probability α and responds whenever available. With such an exogenous response probability, Appendix A shows that the customer would contact as many dealers as is feasible (Proposition 5). Moreover, allowing the available dealers to endogenously decide whether to respond would restore an interior solution for the number n (Proposition 6). These results illustrate the opposing effects of an exogenous versus an endogenous response probability: Letting dealers respond with an exogenous probability pushes the customer to contact more dealers, whereas endogenizing their response probability pushes the customer to contact fewer dealers. The exogenous availability constraint captures a case of heterogeneous dealer valuations, in that the dealers with a high valuation for the asset are not available to sell the asset. Equilibrium results remain qualitatively identical if some dealers, instead of being exogenously unavailable to respond, are subject to an exogenous budget constraint limiting their ability to respond with a competitive price (Internet Appendix B).

What drives my result is not the deadweight cost of responding to an RFQ *per se*, because the equilibrium is completely unaffected even if the cost becomes a mere transfer from the responding dealers to the customer (Appendix D). The key force is dealers' strategic avoidance of competition, arising from the cost and enabled by the dealers' ability to ignore the RFQ.

[14]Examples include Glebkin et al. (2022), Yueshen (2017), and papers in the consumer search literature.

3 Equilibrium

This section establishes that *(1)* the customer contacts only $n^* = 2$ dealers in the unique symmetric subgame perfect equilibrium (Theorem 1), *(2)* all subgame perfect equilibria are payoff equivalent (Theorem 2 Part *(i)*), and *(3)* with a mild tie-breaking rule, the customer always contacts $n^* = 2$ dealers in *any* subgame perfect equilibrium (Theorem 2 Part *(ii)*).

3.1 Symmetric subgame perfect equilibrium

Theorem 1. *The benchmark model has a unique symmetric subgame perfect equilibrium* (n^*, i^*, a^*, F^*)*, where*

$$n^* = 2, \qquad i^* \text{ satisfies (1)},$$

$$a_n^* = \begin{cases} 1 - \left(\dfrac{c}{v}\right)^{\frac{1}{n-1}} & \text{if } n > 1, \\ 1 & \text{if } n = 1, \end{cases} \qquad F_n^*(p) = \begin{cases} \dfrac{1 - \left(\frac{c}{p}\right)^{\frac{1}{n-1}}}{1 - \left(\frac{c}{v}\right)^{\frac{1}{n-1}}}, \text{ and supp } F_n^* = [c, v] & \text{if } n > 1, \\ \mathbb{1}_{p \geq v} & \text{if } n = 1. \end{cases}$$

Theorem 1 shows that it is strictly optimal for the customer to contact only $n^* = 2$ dealers in the unique symmetric subgame perfect equilibrium. Next, I proceed with backward induction to solve for the symmetric equilibrium.

In Stage 3, the customer's dealer choice i^* is directly given by her optimality condition (1).

In Stage 2, if the customer contacts $n = 1$ dealer, then it would be strictly optimal for the dealer to respond with probability $a_n^* = 1$ and offer the monopoly price v deterministically, $F_n^*(p) = \mathbb{1}_{p \geq v}$. When the customer contacts $n > 1$ dealers, the price distribution F_n^* cannot have any atom: Given that responding with any non-positive price is strictly dominated by not responding, then $F_n^*(0) = 0$; If F_n^* had an atom at some price $p^0 > 0$, slightly undercutting it by offering a price $p^0 - \varepsilon$ would yield a strictly higher payoff than offering p^0 for at least one dealer.[15] Given that responding with any price higher than the customer's value v is strictly worse than not responding, then the upper bound $\bar{p} \coloneqq \sup(\operatorname{supp} F_n^*)$ of a dealer's price support is at most v, $\bar{p} \leq v$. When a dealer offers a price $p \in \operatorname{supp} F_n^*$ within its support that arbitrarily approaches its upper bound $\bar{p}$, it gets to trade with the customer if and only if no other dealers respond, an event that occurs with probability $(1 - a_n^*)^{n-1}$. Thus, the dealer's expected trading profit approaches $\bar{p}(1 - a_n^*)^{n-1}$. If the dealer offers price v, its expected trading profit equals $v(1 - a_n^*)^{n-1}$. The dealer's optimality condition (2) implies that $\bar{p} = v$.

The dealer's individual rationality (3) implies that its expected trading profit upon responding is no lower than its response cost, $v(1 - a_n^*)^{n-1} \geq c$. Hence, $a_n^* < 1$. The dealer's indifference condition (4) is equivalent to

$$p\left[1 - a_n^* F_n^*(p)\right]^{n-1} = c, \quad \forall p \in \operatorname{supp} F_n^*. \tag{6}$$

Setting $p = v$ in the above equation yields

$$v(1 - a_n^*)^{n-1} = c \iff a_n^* = 1 - \left(\frac{c}{v}\right)^{\frac{1}{n-1}}. \tag{7}$$

[15] When every dealer offers the same price p^0, at least one dealer trades with the customer with a probability of less than 1 regardless of how the customer breaks her tie. Then, that dealer is strictly better off offering the slightly lower price $p^0 - \varepsilon$ to undercut the other dealers.

Then, equation (6) uniquely determines the price distribution F_n^*,

$$F_n^*(p) = \frac{1 - \left(\frac{c}{p}\right)^{\frac{1}{n-1}}}{1 - \left(\frac{c}{v}\right)^{\frac{1}{n-1}}} \quad \text{and supp } F_n^* = [c, v]\,. \tag{8}$$

Proposition 1 (Dealers' strategic avoidance of competition)**.** *(i) The probability a_n^* that each dealer responds is strictly decreasing in the number of dealers $n > 1$ contacted by the customer. (ii) Conditional on responding, each dealer's price distribution $F_{n'}^*$ first-order stochastically dominates F_n^* for $n' > n > 1$, $F_{n'}^* \succ_{(1)} F_n^*$. (iii) The distribution $G_{n'}^*$ of the dealers' best overall price offer $p = \min_{j=1,\ldots,n} p_j$ first-order stochastically dominates G_n^* for $n' > n > 1$, $G_{n'}^* \succ_{(1)} G_n^*$.*

When the customer contacts one more dealer, dealers strategically avoid competition by *(i)* responding with a lower probability, $a_{n+1}^* < a_n^*$, and *(ii)* offering *strictly less* competitive prices when responding, $F_{n+1}^* \succ_{(1)} F_n^*$. Dealers' strategic avoidance of competition more than offsets the benefit to the customer from potentially receiving one more quote, and results in *(iii)* a *strictly less* competitive best overall price for the customer, $G_{n+1}^* \succ_{(1)} G_n^*$. That the dealers can endogenously decide whether to respond is crucial for enabling their strategic avoidance of competition. If the response probability were exogenously fixed at some constant α instead of varying endogenously with the number of dealers n contacted by the customer, Appendix A shows that the dealer's price distribution F_n^α would become stochastically *smaller* (that is, *more competitive*) if the customer contacts a larger number n of dealers. So would the distribution G_n^α of the dealers' best overall price. With such an exogenous response probability, the customer would contact as many dealers as is feasible in equilibrium (Proposition 5).

Since the distribution G_n^* of the dealers' best overall price is a sufficient statistic for the customer's choice of the number n in Stage 1, Proposition 1's ultimate implication is its Part *(iii)*. Does Part *(iii)* rely on any parametric assumption? No. To show this, I derive Part *(iii)* from a dealer's individual rationality without resorting to the closed-form solutions (7) and (8) for a_n^* and F_n^*. When a dealer j offers a given price p in its price support, its expected trading profit

$$\underbrace{\gamma_n^*(p)}_{\mathbb{P}(\text{Dealer } j \text{ wins the trade when offering a given price } p)} \times \underbrace{p}_{\text{the trade's profit}}$$

equals its response cost c. Then, $\mathbb{P}\,(\text{Dealer } j \text{ wins the trade when offering a given price } p)$ should remain constant when the customer contacts one more dealer, $\gamma_{n+1}^*(p) = \gamma_n^*(p)$. Dealer j wins the trade if and only if no other dealer offers a price lower than p. Then, $\gamma_n^*(p) = [\theta_n^*(p)]^{n-1}$, where

$$\theta_n^*(p) \coloneqq \mathbb{P}\,(\text{a given dealer does not offer a price lower than } p)\,.$$

Because the probability that j wins the trade remains constant, $[\theta_{n+1}^*(p)]^n = [\theta_n^*(p)]^{n-1}$. Thus, $\mathbb{P}\,(\text{a given dealer does not offer a price lower than } p)$ should rise, $\theta_{n+1}^*(p) > \theta_n^*(p)$. Hence, $\mathbb{P}\,(\text{no dealer offers a price lower than } p)$ should rise too, $[\theta_{n+1}^*(p)]^{n+1} > [\theta_n^*(p)]^n$. Therefore, the dealers' best overall price becomes stochastically larger, $G_{n+1}^* \succ_{(1)} G_n^*$.

Therefore, Proposition 1 provides the basis for the customer to contact fewer dealers in Stage 1. I now turn to solve the customer's problem in Stage 1.

In Stage 1, the customer's payoff is 0 upon contacting $n = 1$ dealer. If the customer contacts $n > 1$ dealers, the best price offer $p = \min_{j=1,\ldots,n} p_j$ for her becomes first-order stochastically *higher* when she contacts more dealers (Proposition 1). Therefore, it is strictly optimal for the customer to contact $n^* = 2$ dealers. This establishes Theorem 1.

Theorem 1 continues to hold in the exact same form when the customer is risk-averse: In Stage 3, the customer continues to trade with the dealer who offers the lowest price if that price is less than the customer's reservation value v. Thus, in Stage 2, the dealers' subgame remains unaffected, and the dealers continue to follow the equilibrium strategy (a^*, F^*). In Stage 1, given that the stochastic dominance $G_2^* \prec_{(1)} G_3^* \prec_{(1)} \ldots$ is *first-order*, risk aversion does not affect the customer's choice of the number n. Therefore, a risk-averse customer continues to contact $n^* = 2$ dealers in the unique symmetric subgame perfect equilibrium.

Theorem 1 also holds in the exact same form when dealers are risk-averse. Appendix E shows that the above intuition for Proposition 1 (dealers' strategic avoidance of competition) remains valid when dealers are risk-averse and pushes the customer to contact $n^* = 2$ dealers. Appendix E then formally solves for the equilibrium.

Responding to multiple RFQs that occur within a short period of time may provide some economies of scale for a given dealer. Internet Appendix A solves an extension with three sequential RFQs, where each dealer incurs a cost c only once upon responding to one or more RFQs. This cost assumption captures the extreme version of economies of scale. In equilibrium, the customer in each RFQ continues to contact two dealers. Proposition A.1 establishes the unique symmetric Markov perfect equilibrium of this extension.

3.2 Any subgame perfect equilibrium

The above intuition for Proposition 1 (dealers' strategic avoidance of competition) requires symmetry in the dealers' strategies. This section uncovers the underlying economics, which does not require the symmetry and generalizes Theorem 1 to *any* subgame perfect equilibrium.

There exist subgame perfect equilibria other than the symmetric one. For example, one can modify the symmetric equilibrium as follows to obtain another subgame perfect equilibrium: When the customer contacts $n > 2$ dealers, one can let $n - 2$ of them not respond, and let the other two respond with probability a_2^* and offer the price distribution F_2^* conditional on responding. In such an equilibrium, the customer still earns the same *ex-ante* payoff of π_2^* as in the symmetric equilibrium, where

$$\pi_n^* := v - \mathbb{E}_{G_n^*}(p \wedge v). \tag{9}$$

Generalizing this example, the next result establishes that *(i)* all subgame perfect equilibria are payoff equivalent, and the customer's *ex-ante* payoff equals π_2^* across all subgame perfect equilibria, and *(ii)* subject to a mild tie-breaking rule, the customer always contacts $n^* = 2$ dealers in any subgame perfect equilibrium.

Theorem 2. *(i) All subgame perfect equilibria are payoff equivalent. In any subgame perfect equilibrium, the customer's ex-ante payoff equals π_2^*, and all dealers' ex-ante payoffs equal* 0*. (ii) If one imposes a tie-breaking rule that the customer contacts fewer dealers whenever she is indifferent, then the customer always contacts $n^* = 2$ dealers in any subgame perfect equilibrium.*

The proof, provided in Appendix B, generalizes that of Theorem 1. The underlying economics—I call it *response concentration*—comes from the observation that it is more *cost-efficient* to concentrate response probabilities among fewer dealers. To illustrate this observation, I fix an arbitrary subgame perfect equilibrium. For a given number of contacted dealers n, I let $a_{1,n}^*, \ldots, a_{n,n}^*$ be the n dealers' equilibrium response probabilities. Then the aggregate expected gain from trade is $v[1 - (1 - a_{1,n}^*) \ldots (1 - a_{n,n}^*)]$. That is, the aggregate gain depends on the individual response probabilities $a^*_{1,n}, \ldots, a^*_{n,n}$ only through the sufficient statistic $(1 - a^*_{1,n}) \ldots (1 - a^*_{n-1,n})(1 - a^*_{n,n})$, which is the probability that no dealer responds to the customer's RFQ. Keeping this aggregate gain constant, one can reduce the expected response cost $c(a^*_{1,n} + \ldots + a^*_{n-1,n} + a^*_{n,n})$ by reducing one dealer's response probability, say $a^*_{n,n}$, down to 0 and raising another dealer's response probability, say $a^*_{n-1,n}$, appropriately. Specifically, the minimization problem

$$\min \; a^*_{n-1,n} + a^*_{n,n}$$
$$\text{subject to } \; (1 - a^*_{n-1,n})(1 - a^*_{n,n}) = \text{constant}$$

is solved when either $a^*_{n-1,n} = 0$ or $a^*_{n,n} = 0$. Therefore, it is more *cost-efficient* to shut down one dealer and concentrate the response probabilities into the remaining dealers. Applying this argument inductively, one obtains that it is more cost-efficient to let at most 2 dealers respond with a positive probability.

When the customer contacts exactly 2 dealers, her expected payoff is shown to be equal to the aggregate net payoff, which is π^*_2. Thus, in any subgame perfect equilibrium, the customer's *ex-ante* payoff always equals π^*_2, which she can achieve by contacting only 2 dealers.

I further exploit the cost efficiency of response concentration to generalize Theorem 1 with alternative platform designs. An assumption of the model is that the number n of contacted dealers is disclosed to the n dealers, as is the case on SEFs. The next section shows that no alternative design of information disclosure about the number n can improve the customer's payoff.

4 Alternative Designs

In practice, many multi-dealer platforms (such as SEFs) disclose the number n of dealers contacted in a customer RFQ to those n dealers (Riggs et al., 2020), perhaps as a way to motivate the dealers to offer more competitive prices. Proposition 1 shows that upon observing a larger number n of contacted dealers, the dealers' strategic avoidance of competition more than offsets the benefit to the customer from potentially receiving more quotes, leading to a less competitive best overall price for the customer. This section examines alternative platform designs of the information disclosure about the number n. I first consider the other extreme case where the platform designer does not disclose any information about the number n. Then, I search for the optimal design of information disclosure that maximizes price competitiveness. Overall, no alternative design of information disclosure can improve the customer's payoff above that of her status quo, where the number of contacted dealers n is fully disclosed.

4.1 No disclosure

I modify the benchmark model as follows: (1) The dealers cannot observe how many other dealers are contacted by the customer. (2) The customer can contact at most $\bar{n}$ dealers ($\bar{n} > 2$), because it turns out that the customer will contact as many dealers as is feasible in equilibrium. (3) I assume the tie-breaking rule that the customer contacts fewer dealers whenever she is indifferent. (4) To account for imperfect information as the number n becomes unobservable, I use the solution concept of symmetric perfect Bayesian equilibrium (PBE), where all dealers employ the same strategy $(a^{\text{unobs}}, F^{\text{unobs}})$. I do not impose any restriction on off-path beliefs. The remaining setup is identical to the benchmark model in

Section 2.

With the number of contacted dealers n being unobservable, a dealer's response strategy $(a^{\text{unobs}}, F^{\text{unobs}})$ can no longer depend on the number n. There are two symmetric PBE, one of which is degenerate in that the customer submits no RFQ at all. I first solve for the unique non-degenerate equilibrium, then spell out the degenerate one.

Formally, a symmetric PBE is *non-degenerate* if the customer submits an RFQ with a positive probability. In a non-degenerate equilibrium, the customer's *ex-ante* payoff must be strictly positive. Thus, the dealers must respond and offer prices strictly less than the monopoly price v with a positive probability, $a^{\text{unobs}} > 0$ and $F^{\text{unobs}}(v^-) > 0$. Hence, it is strictly optimal for the customer to contact as many dealers as is feasible, $n^{\text{unobs}} = \bar{n}$. In equilibrium, the dealers have the correct conjecture about the customer's equilibrium choice, $n^{\text{unobs}} = \bar{n}$, leading to the following equilibrium result.

Proposition 2. *I consider the modified model where the number of contacted dealers is not disclosed. (i) There exists a unique non-degenerate symmetric PBE,* $(n^{\text{unobs}}, i^*, a^{\text{unobs}}, F^{\text{unobs}})$*, where* i^* *is the same as in Theorem 1, and*

$$n^{\text{unobs}} = \bar{n}, \quad a^{\text{unobs}} = a^*_{\bar{n}}, \quad F^{\text{unobs}} = F^*_{\bar{n}}.$$

A dealer believes that the customer contacted $\bar{n}$ *dealers whenever it receives an RFQ. (ii) Under* $(n^{\text{unobs}}, i^*, a^{\text{unobs}}, F^{\text{unobs}})$*, the customer's ex-ante payoff is strictly lower than her status quo payoff* π^*_2 *given by (9) wherein her number of contacted dealers,* n*, is fully disclosed.*

Once her number n becomes undisclosed, the customer can no longer commit to contact fewer than $\bar{n}$ dealers, and thereby receives a lower equilibrium payoff. Specifically, the customer's equilibrium payoff becomes $\pi^*_{\bar{n}}$, which is what she would have earned if she had contacted $n = \bar{n}$ dealers in the benchmark model. This payoff is strictly less than the customer's equilibrium payoff π_2^* in the benchmark model where she contacts 2 dealers in equilibrium, $\pi_{\bar{n}}^* < \pi_2^*$. That is, the customer receives a worse overall price and a lower equilibrium payoff despite her contacting more dealers.

In the U.S. corporate bond market, MarketAxess does not disclose customers' number of contacted dealers by default; voluntary disclosure is ineffectual due to the platform's lack of detailed transparency (i.e. dealers cannot determine whether a customer has selected Open Trading, as it would merely count as one additional contacted participant). Consistent with Proposition 2, a customer contacts more than 25 dealers in an RFQ, on average, and dealers' response rate is only around 25% (Hendershott and Madhavan, 2015, Table VI). In comparison, the response rate on SEFs is almost 90% (Riggs et al., 2020, Table 3). To jointly explain these numbers, a rough computation based the relationship $1 - (c/v)^{1/(n-1)} = a_n^*$ estimates that the cost-to-value ratios c/v that I would need for SEFs and MarketAxess are nearly identical, $(1 - 90\%)^{4-1} \approx (1 - 25\%)^{25-1} \approx 10^{-3}$. The trading cost in basis points is also much higher on MarketAxess than on SEFs, although such a comparison may be confounded by other factors, such as the lower liquidity of corporate bonds relative to index credit default swaps.

Next, I spell out the degenerate equilibrium:

- In Stage 1, the customer submits no RFQ at all;
- In Stage 2, a dealer believes it is the only dealer contacted by the customer whenever it receives an RFQ. The dealer responds with probability 1 and offers the monopoly price v deterministically;
- In Stage 3, the customer's dealer choice remains to be i^*, as given by her optimality

condition (1), if the customer had received any quote(s).

It is easy to verify that the above constitutes a symmetric PBE. Proposition 7 (Appendix B) shows that there exists no other symmetric PBE.

4.2 Optimal information disclosure

Although making the number n of contacted dealers undisclosed does not lead to more competitive prices, it does cause the customer to contact more dealers in equilibrium. One may wonder whether there exists an alternative information design with partial disclosure that could reach the sweet spot—where the customer contacts a high enough number of dealers and the dealers engage in a sufficient level of competition—and make the dealers' best overall price more competitive than the status quo of full disclosure. Unfortunately, the answer is no.

Formally, a *design of information disclosure* (S, μ) consists of a countable realization space S and a family of distributions $\{\mu(\cdot|n)\}_{n\in\mathbb{Z}^{++}}$ over the space S. The design of information disclosure is common knowledge among all market participants. Here are three examples of information design.

- Example 1 (full disclosure): When $S = \mathbb{Z}^{++}$ and $\mu(s|n) = \mathbb{1}_{\{s=n\}}$, $\forall s \in S$, the information design fully discloses the number n.

- Example 2 (no disclosure): When S is a singleton, the information design discloses nothing about the number n.

- Example 3 (partial disclosure): When $S = \{\text{odd}, \text{even}\}$ and $\mu(\text{odd}|n) = \mathbb{1}_{\{n \text{ is odd}\}}$, the information design discloses only whether n is odd or even.

I generalize the benchmark model to allow for any arbitrary design of information disclosure (S, μ). In Stage 1, the customer chooses the number n of dealers to contact. A signal s is drawn from the distribution $\mu(\cdot|n)$ and is observed by the contacted dealers, who then choose whether to respond and which price to offer in Stage 2. In Stage 3, the customer chooses whether and against which dealer's price to trade. As in the benchmark model, I do not impose any tie-breaking rule in the case of indifference. All agents are risk-neutral. The solution concept is symmetric PBE, where all dealers employ the same strategy (a^{μ}, F^{μ}). I do not impose any restriction on off-path beliefs. Figure 2 summarizes the timeline of the generalized model.

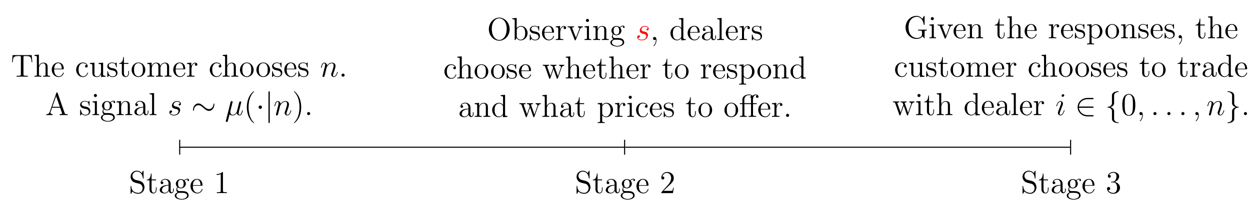

Figure 2: Timeline: Dealers now observe a signal s instead of the number n directly.

The way that information is optimally disclosed is identical to that of Bayesian Persuation (Kamenica and Gentzkow, 2011) in that the platform designer sends a signal about the state of the world n, and the contacted dealers receive the signal. There is one fundamental distinction: Here, the receivers' prior belief about the state of the world n is endogenously determined by the customer's equilibrium choice. In Bayesian Persuasion, the receiver's prior belief about the state of the world is exogenously given.

In my setup, fully disclosing the state of the world n is optimal.

Theorem 3. *Given any design of information disclosure (S, μ), the customer's ex-ante payoff in any symmetric PBE is less than or equal to her status quo payoff π_2^* given by (9).*

Although the customer contacts only two dealers when the number of her contacted dealers is fully disclosed, no alternative design of information disclosure can improve her payoff above this outcome. Theorem 3 establishes the limits of multi-dealer platforms in promoting price competition.

The proof of Theorem 3, provided in Appendix B, generalizes that of Theorem 1. The underlying economics again comes from the cost savings of response concentration. Specifically, conditional on any signal realization $s \in S$ that is drawn with a positive probability under a given equilibrium, I let $\chi_s(n)$ be the posterior probability that the customer contacted n dealers. The conditional expectations of the aggregate gain from trade and the aggregate response cost simply average the gain and the cost across all on-path choices n with their respective posterior probabilities $\chi_s(n)$. Given each on-path choice $n \geq 2$, it is more cost-efficient to concentrate response probabilities among fewer dealers. Taking the expectations across all on-path choices n naturally preserves this property of cost savings.

5 Empirical Predictions

The benchmark model in Section 2 can be easily extended to include an order size q and to yield more testable predictions. These predictions are largely consistent with facts documented in the literature.

I extend the benchmark model as follows: In Stage 1, a customer seeking to buy q units of the asset chooses a number n of *ex-ante* identical dealers to contact in an RFQ. The order size q is an exogenous parameter and thus is common knowledge among market participants. Stages 2 and 3 remain identical to those in the benchmark model. To ensure a positive net surplus from trade, I assume that $c < vq$.

Proposition 3. *The extended model with an order size q has a unique symmetric subgame perfect equilibrium (n^*, i^*, a^q, F^q), where n^* and i^* are the same as in Theorem 1, and*

$$a_n^q = \begin{cases} 1 - \left(\dfrac{c}{vq}\right)^{\frac{1}{n-1}} & \text{if } n > 1, \\ 1 & \text{if } n = 1, \end{cases} \quad F_n^q(p) = \begin{cases} \dfrac{1 - \left(\frac{c}{pq}\right)^{\frac{1}{n-1}}}{1 - \left(\frac{c}{vq}\right)^{\frac{1}{n-1}}}, \text{ and supp } F_n^q = \left[\dfrac{c}{q}, v\right] & \text{if } n > 1, \\ \mathbb{1}_{p \geq v} & \text{if } n = 1. \end{cases}$$

Compared to Theorem 1, the customer's equilibrium strategy (n^*, i^*) remains unchanged. The only difference lies in the expressions of the dealer's equilibrium strategy (a_n^q, F_n^q) in that the cost c is normalized by the quantity q and becomes the per-unit cost c/q. This difference arises from the fact that the dealer's indifference condition (6) becomes

$$pq\left[1 - a_n^q F_n^q(p)\right]^{n-1} = c, \quad \forall p \in \text{supp } F_n^q.$$

The remaining proof is otherwise identical to that of Theorem 1.

Based on the above equilibrium result, the next proposition provides testable predictions on a dealer's response probability and price distribution. I write $a_2^{q,v}$ for the response probability a_2^q, and $F_2^{q,v}$ for the price distribution F_2^q, to state the effects of the value v.

Proposition 4. *(i) The equilibrium response probability $a_2^{q,v}$ is strictly increasing in order size q and value v. (ii) Conditional on responding, each dealer's equilibrium price distribution $F_2^{q,v}$ becomes first-order stochastically smaller with a larger size q or a smaller value v,*

$$F_2^{q,v} \succ_{(1)} F_2^{q',v} \textit{for } q < q' \textit{ and } F_2^{q,v'} \succ_{(1)} F_2^{q,v} \textit{ for } v < v'.$$

On SEFs for index credit default swaps, Riggs et al. (2020) documents patterns that are largely consistent with the predictions of Propositions 1 and 4. Specifically, they find that a dealer's likelihood of responding to an RFQ decreases in the number of contacted dealers (Proposition 1 Part *(i)*) and increases in notional quantity (Proposition 4 Part *(i)*). A customer's RFQ is more likely to result in an actual trade if the order size is larger or nonstandard, which is consistent with the interpretation that those orders imply larger gains from trade between customers and dealers (Proposition 4 Part *(i)*). Conditional on dealer(s) responding to an RFQ, they find that dealers' quoted spreads and customers' transaction costs become larger if more dealers are contacted in the RFQ (Proposition 1 Part *(ii)*) or if order sizes are nonstandard (Proposition 4 Part *(ii)*), although the effects are mild. For larger order sizes, however, they find that dealers' quoted spreads are slightly larger by an economically and statistically insignificant magnitude.

6 Conclusion

On many important multi-dealer platforms such as SEFs, customers mostly request quotes from very few dealers. I build a model of multi-dealer platforms, where a customer can simultaneously request quotes from any number of dealers, and each dealer strategically chooses to respond to or ignore the request. In this otherwise standard model of price competition, letting dealers endogenously decide whether to respond overturns the competitiveness of the dealers' best overall price. If the customer contacts more dealers, every contacted dealer strategically avoids competition by *(i)* responding with a lower probability, and *(ii)* offering a stochastically worse price when it responds. Dealers' strategic avoidance of competition more than offsets the benefit to the customer from potentially receiving more quotes, and results in *(iii)* a strictly less competitive best overall price for the customer. The more general underlying economics is “response concentration”: It is more cost-efficient to concentrate response probabilities among fewer dealers. In equilibrium, it is strictly optimal for the customer to contact only two dealers. More generally, no design of information disclosure about the number of contacted dealers can improve the customer’s payoff above this outcome. In this sense, multi-dealer platforms are limited in their ability to promote price competition.

Appendices

A Exogenous Availability Constraint

The dealers' ability to endogenously decide whether to respond is a driving feature of the model. To illustrate its role, this appendix considers two variants of the model that include an exogenous availability constraint.

Variant A1

Variant A1 differs from the benchmark model in two respects: (1) Instead of deciding whether to respond at a cost, each contacted dealer is available with an exogenously fixed probability $\alpha < 1$ and responds whenever available, and (2) the customer can contact at most $\bar{n}$ dealers ($\bar{n} > 1$), because it turns out that the customer will contact as many dealers as is feasible in equilibrium. Variant A1 is otherwise identical to the benchmark model. That is, this variant differs from the benchmark model in that the dealers' exogenous availability constraint *replaces* their ability to endogenously decide whether to respond.

With such an exogenous response probability α, the next proposition establishes that each dealer's pricing becomes *more* competitive when the customer contacts more dealers (as in a standard model of price competition), and the customer contacts $\bar{n}$ dealers in equilibrium.

Proposition 5. *Variant A1 has a unique symmetric subgame perfect equilibrium* $(n^\alpha, i^*, F^\alpha)$*, where* i^* *is the same as in Theorem 1, and*

$$n^\alpha = \bar{n}, \qquad F_n^\alpha(p) = \begin{cases} \dfrac{1-(1-\alpha)\left(\frac{v}{p}\right)^{\frac{1}{n-1}}}{\alpha}, \text{ and supp } F_n^\alpha = \left[v(1-\alpha)^{n-1}, v\right] & \text{if } n > 1, \\ \mathbb{1}_{p \geq v} & \text{if } n = 1. \end{cases}$$

In particular, each dealer's price distribution F_n^α becomes first-order stochastically smaller as n increases. When $\bar{n} \to \infty$, the equilibrium price distribution $F_{\bar{n}}^\alpha$ converges in distribution to the competitive limit 0.

Proof. The backward induction for Variant A1 is similar to that for the benchmark model. The only change is that a contacted dealer no longer needs to be indifferent between responding or not, as it now responds with an exogenous probability instead of endogenously deciding whether to respond. That is, the indifference condition (4) need not hold.

In Stage 3, the customer's dealer choice remains to be i^*, as given by her optimality condition (1).

In Stage 2, if the customer contacts $n = 1$ dealer, then it would be strictly optimal for the dealer to offer the monopoly price v deterministically whenever the dealer is available, $F_n^\alpha(p) = \mathbb{1}_{p \geq v}$. When the customer contacts $n > 1$ dealers, the price distribution F_n^α cannot have an atom at any positive price: If F_n^α had an atom at some price $p^0 > 0$, slightly undercutting it by offering a price $p^0 - \varepsilon$ would yield a strictly higher payoff than offering p^0 for at least one dealer. When a dealer offers a price $p \in \operatorname{supp} F_n^\alpha$ within its support that arbitrarily approaches its support's upper bound $\tilde{p} \coloneqq \sup(\operatorname{supp} F_n^\alpha)$, the dealer's expected trading profit is at most $\tilde{p}\mathbb{1}_{\tilde{p} \leq v}(1-\alpha)^{n-1}$. If the dealer offers price v, its expected trading profit is at least $v(1-\alpha)^{n-1}$. The dealer's optimality condition (2) implies that $\tilde{p} = v$.

Then, (2) is equivalent to

$$p\left[1 - \alpha F_n^\alpha(p)\right]^{n-1} = v\,(1-\alpha)^{n-1}, \quad \forall p \in \operatorname{supp} F_n^\alpha$$
$$\iff F_n^\alpha(p) = \frac{1 - (1-\alpha)\left(\frac{v}{p}\right)^{\frac{1}{n-1}}}{\alpha}, \quad \text{and } \operatorname{supp} F_n^\alpha = \left[v(1-\alpha)^{n-1}, v\right].$$

In Stage 1, the customer's payoff is 0 upon contacting $n = 1$ dealer. If $n > 1$, the distribution G_n^α of the dealers' best overall price $p = \min_{j=1,\ldots,n} p_j$ is given by

$$1 - G_n^\alpha(p) = [1 - \alpha F_n^\alpha(p)]^n, \quad \forall p \in \left[v(1-\alpha)^{n-1}, v\right].$$

If n increases, $F_n^\alpha(p)$ strictly increases, and thus $1 - G_n^\alpha(p)$ strictly decreases. That is, the dealers' best overall price $p = \min_{j=1,\ldots,n} p_j$ becomes first-order stochastically *smaller* when the customer contacts more dealers. Therefore, the customer's unique optimal choice is $n^\alpha = \bar{n}$. $\square$

Variant A2

Variant A2 differs from Variant A1 in two respects: (1) Each available dealer can endogenously decide whether to respond at a cost c, which is assumed to be less than v ($c < v$) so that there is a positive net surplus from trade, while a non-available dealer simply does not respond. (2) The customer could contact any arbitrary number of dealers. Variant A2 is otherwise identical to Variant A1. That is, Variant A2 *reintroduces* dealers' ability to ignore the RFQ to Variant A1.

This feature restores an interior solution for the equilibrium number of contacted dealers.

Proposition 6. *Variant A2 has a unique symmetric subgame perfect equilibrium* $(n^{\alpha,c}, i^*, a^{\alpha,c}, F^{\alpha,c})$*, where* i^* *is the same as in Theorem 1, and*

$$n^{\alpha,c} = m \text{ or } m-1, \qquad \text{where } m \text{ is uniquely determined by } a^*_{m-1} > \alpha \geq a^*_m,$$

$$a_n^{\alpha,c} = \begin{cases} 1 & \text{if } n < m, \\ \dfrac{a_n^*}{\alpha} & \text{if } n \geq m, \end{cases} \qquad F_n^{\alpha,c} = \begin{cases} F_n^\alpha & \text{if } n < m, \\ F_n^* & \text{if } n \geq m. \end{cases}$$

When the contacted dealers are able to endogenously decide whether to respond, the exogenous availability constraint becomes non-binding when $a_n^* \leq \alpha$, because each dealer would respond with a probability a_n^* that is lower than α anyway. Since the endogenous response probability a_n^* decreases to 0 as n increases, it declines below the exogenous probability α when n is above a certain threshold m. When $n \geq m$, the exogenous availability constraint becomes irrelevant. Thus, the dealers behave as in the benchmark model of Section 2: Each dealer's effective response probability remains at $\alpha a_n^{\alpha,c} = a_n^*$, and its price distribution remains at $F_n^{\alpha,c} = F_n^*$. Hence, the customer strictly prefers to contact fewer dealers in this range. When $n < m$, the exogenous availability constraint is binding while the dealers' ability to strategically ignore the RFQ becomes irrelevant. Thus, the dealers behave as in Variant A1. Hence, the customer strictly prefers to contact more dealers in this range. Overall, the customer's optimal choice is $n^{\alpha,c} = m$ or $m - 1$, depending on how close the two probabilities a_m^* and α are when a_m^* declines below α.

The backward induction for Variant A2 is similar to that for the benchmark model. The only change is that the dealers need not be indifferent between responding or not when the exogenous availability constraint is strictly binding. That is, the indifference condition (4) need not hold when $n < m$. I do not repeat the formal proof.

B Proofs

Proof of Theorem 1. The proof is given immediately after Theorem 1. $\square$

Proof of Proposition 1. *Part (i)*: Since $c < v$, then $(c/v)^{1/(n-1)}$ is strictly increasing in n. Thus, the probability a_n^* is strictly decreasing in n.

Part (ii): Fixing any $p \in [c, v]$, I let $c/p = \eta$ and $c/v = \delta$. Then, $\delta < \eta < 1$ and

$$\ln F_n^*(p) = \ln\left(1 - \eta^{\frac{1}{n-1}}\right) - \ln\left(1 - \delta^{\frac{1}{n-1}}\right).$$

It suffices to show that $\ln F_n^*(p)$ is strictly decreasing in $n > 1$. I view n as a continuous variable and take the partial derivative of $\ln F_n^*(p)$ with respect to n to obtain

$$\frac{\partial}{\partial n} \ln F_n^*(p) = \frac{1}{n-1}\left[\frac{\tilde{\eta}\ln(\tilde{\eta})}{1-\tilde{\eta}} - \frac{\tilde{\delta}\ln(\tilde{\delta})}{1-\tilde{\delta}}\right],$$

where $\tilde{\delta} := \delta^{1/(n-1)} < \eta^{1/(n-1)} =: \tilde{\eta}$. Since the function $x \mapsto (x\ln x)/(1-x)$ is strictly decreasing in $x \in (0,1)$, then $\frac{\partial}{\partial n}\ln F_n^*(p) < 0$. Hence for any $p \in [c, v]$, $F_n^*(p)$ is strictly decreasing in $n > 1$. Therefore, $F_{n'}^* \succ_{(1)} F_n^*$ for $n' > n > 1$.

Part (iii): If the customer contacts $n > 1$ dealers, the distribution G_n^* of the dealers' best price offer $p = \min_{j=1,\ldots,n} p_j$ is given by

$$1 - G_n^*(p) = [1 - a_n^* F_n^*(p)]^n = \left(\frac{c}{p}\right)^{\frac{n}{n-1}}, \quad \forall p \in [c, v].$$

If n increases, $1 - G_n^*(p)$ strictly increases. Therefore, $G_{n'}^* \succ_{(1)} G_n^*$ for $n' > n > 1$. $\square$

The next two lemmas are useful to prove Theorem 2.

Lemma 1. *Given a price p^0, after the customer contacts any number n of dealers, at most one contacted dealer's price distribution can have an atom at price p^0 in any subgame perfect equilibrium of the benchmark model.*

Proof. I suppose that two price distributions $F_{j,n}^*$ and $F_{j',n}^*$ have an atom at p^0. Since dealer j and j' must earn strictly positive expected trading profits when offering price p^0 in compensation for their response cost, then either j or j' is strictly better off undercutting by offering some price $p^0 - \varepsilon$. This contradicts the optimality of the price p^0 for dealer j and j'. Lemma 1 follows. □

Lemma 2. *Given a contacted dealer j, if there exits some price $p^0 < v$ and $\varepsilon > 0$ such that $(p^0, p^0 + \varepsilon) \cap \text{supp } F^*_{j',n} = \emptyset$ for any other contacted dealer $j' \neq j$, then the price p^0 cannot be in the price support of dealer j, $p^0 \notin \text{supp } F^*_{j,n}$.*

PROOF. I suppose that the condition of Lemma 2 holds.

Step 1: By offering any price $p \in (p^0, p^0 + \varepsilon)$, dealer j gets to trade with the customer with a constant probability. Conditional on such a trade, the trading profit earned by j equals the price p, which is strictly increasing in p. Thus, the expected trading profit earned by j is either 0, which is insufficient to cover the response cost c, or strictly increasing in $p \in (p^0, p^0 + \varepsilon)$. Hence, no price $p \in (p^0, p^0 + \varepsilon)$ can be in the price support of dealer j,

$$p \notin \text{supp } F^*_{j,n}, \qquad \forall p \in (p^0, p^0 + \varepsilon).$$

Step 2: If the distribution $F^*_{j,n}$ has an atom at p^0, no other price distribution $F^*_{j',n}$ can have an atom at p^0 (Lemma 1). Then, dealer j is strictly worse off offering price p^0 than some price $p \in (p^0, p^0 + \varepsilon)$. Hence, the distribution $F^*_{j,n}$ cannot have an atom at $p = p^0$. By the same argument, any other price distribution $F^*_{j',n}$ cannot have an atom at $p = p^0$ either.

Step 3: When dealer j offers a price $p^0 - \varepsilon'$ that arbitrarily approaches the price p^0 from below, j either trades with the customer with probability 0 or is strictly better off offering some price $p \in (p^0, p^0 + \varepsilon)$, as no other price distributions have an atom at p^0. Hence, for some $\varepsilon' > 0$, no price $p \in (p^0, p^0 - \varepsilon')$ can be in the price support of dealer j,

$$p \notin \operatorname{supp} F^*_{j,n}, \qquad \forall p \in (p^0, p^0 - \varepsilon').$$

The conclusions of Steps *1-3* together imply that p^0 cannot be in the price support of dealer j, $p^0 \notin \operatorname{supp} F^*_{j,n}$. □

Proof of Theorem 2. I fix an arbitrary subgame perfect equilibrium. If the customer contacts $n = 1$ dealer, the dealer would respond with probability 1 and offer the monopoly price v deterministically. Thus, the customer's payoff is 0. If the customer contacts some given number $n \geq 2$ of dealers, I show that the customer's expected payoff does not exceed π^*_2.

Given a contacted dealer j, I let $\bar{p}_j$ be the upper bound of the dealer's price support, and let $\bar{p}_{-j}$ be the highest upper bound of the other contacted dealers' price supports:

$$\bar{p}_j := \operatorname{supp} F^*_{j,n}, \qquad \bar{p}_{-j} := \max_{j' \neq j} \bar{p}_{j'}.$$

If a dealer j' responds with probability 0, then $\bar{p}_{j'} := -\infty$ by convention. If $\bar{p}_{-j} < v$, then Lemma 2 implies that $\bar{p}_{-j} \notin \operatorname{supp} F^*_{j,n}$. Hence, Lemma 2 further implies that $\bar{p}_{-j} \notin \operatorname{supp} F^*_{j',n}$ for any other contacted dealer $j' \neq j$. Then, $\bar{p}_{-j}$ cannot be the upper bound of any dealer's price support, which contradicts the definition of $\bar{p}_{-j}$. Thus, $\bar{p}_{-j} = v$ for any contacted dealer j. Hence, at least two contacted dealers' price supports have an upper bound of v.

At most one contacted dealer can have an atom at v in its price support (Lemma 1). Without loss of generality, I let dealer 1 be such that $\bar{p}_1 = v$ and all other dealers do not have an atom at v. When dealer 1 offers a price $p \in \operatorname{supp} F^*_{1,n}$ within its support that arbitrarily approaches v, the dealer's expected trading profit approaches $v\left(1-a^*_{2,n}\right)\ldots\left(1-a^*_{n,n}\right)$. This limiting profit must be at least c,

$$v\left(1-a^*_{2,n}\right)\ldots\left(1-a^*_{n,n}\right)\geq c \iff \left(1-a^*_{2,n}\right)\ldots\left(1-a^*_{n,n}\right)\geq \frac{c}{v}. \tag{10}$$

Thus, $a^*_{j,n}<1$ for dealer $j=2,\ldots,n$. Hence, all those dealers' expected payoffs must equal 0.

It then follows that the customer's expected payoff does not exceed π^*_2,

$$\begin{aligned}
&\underbrace{v\left[1-\prod_{j=1}^{n}\left(1-a^*_{j,n}\right)\right]-c\sum_{j=1}^{n}a^*_{j,n}}_{\text{aggregate expected payoff}}-\underbrace{\left[v\prod_{j'=2}^{n}\left(1-a^*_{j',n}\right)-c\right]a^*_{1,n}}_{\text{expected payoff of dealer 1}}\\
&=v\left[1-\left(1-a^*_{2,n}\right)\ldots\left(1-a^*_{n,n}\right)\right]-c\left[a^*_{2,n}+\ldots+a^*_{n,n}\right]\\
&\leq v\left[1-\left(1-a^*_{2,n}\right)\ldots\left(1-a^*_{n,n}\right)\right]-c\left[1-\left(1-a^*_{2,n}\right)\ldots\left(1-a^*_{n,n}\right)\right] &(11)\\
&\leq v\left[1-\frac{c}{v}\right]-c\left[1-\frac{c}{v}\right] \qquad \text{(following from (10))} &(12)\\
&=v\left[1-(1-a^*_2)^2\right]-c\left[a^*_2+a^*_2\right]=\pi^*_2.
\end{aligned}$$

Inequality (11) above follows from $x+y\leq xy+1$ for $0\leq x,y\leq 1$ and an induction over n:

$$\begin{aligned}
&a^*_{2,n}+\ldots+a^*_{n,n}\\
=\;&n-1-\left[\left(1-a^*_{2,n}\right)+\ldots+\left(1-a^*_{n-2,n}\right)+\left(1-a^*_{n-1,n}\right)+\left(1-a^*_{n,n}\right)\right]\\
\geq\;&n-1-\left[\left(1-a^*_{2,n}\right)+\ldots+\left(1-a^*_{n-2,n}\right)+\left(1-a^*_{n-1,n}\right)\left(1-a^*_{n,n}\right)+1\right]\\
\geq\;&\ldots\\
\geq\;&n-1-\left[\left(1-a^*_{2,n}\right)\ldots\left(1-a^*_{n,n}\right)+n-2\right]\\
=\;&1-\left(1-a^*_{2,n}\right)\ldots\left(1-a^*_{n,n}\right).
\end{aligned}$$

In the special case where the customer contacts 2 dealers, I show that the customer's expected payoff equals π_2^*. First, inequality (11) becomes an equality. Further, Lemma 2 implies that both contacted dealers' price supports must share the same lower bound $\underline{p} \geq c$. If $\underline{p} > c$, then undercutting by offering some price $\underline{p} - \varepsilon$ would yield a strictly positive payoff to dealer 2. Thus, $\underline{p} = c$. When dealer 1 offers a price $p \in \text{supp } F_{1,2}^*$ within its support that arbitrarily approaches c, the dealer's limiting payoff is non-positive. Thus, the expected payoff of dealer 1 must also be 0. Thus, (10) must hold as an equality, $1 - a_{2,2}^* = c/v$. Hence, (12) becomes an equality too.

Then, the customer's *ex-ante* payoff equals π_2^*, which she can achieve by contacting only 2 dealers. Thus, for any n that is chosen by the customer with a strictly positive probability under the equilibrium, (12) must hold as an equality. So must (10), which implies that the *ex-ante* payoff of dealer 1 must be 0. Hence, all dealers' *ex-ante* payoffs equal 0. Theorem 2 follows. □

PROOF OF PROPOSITION 2. The proof is given immediately before Proposition 2. □

Proposition 7. *I consider the modified model where the number of contacted dealers is undisclosed. All symmetric PBE are provided in Section 4.1, and there exists no other symmetric PBE.*

PROOF. Proposition 2 establishes the unique non-degenerate symmetric PBE. It is easy to verify that the other candidate degenerate equilibrium provided in Section 4.1 indeed constitutes a PBE. It suffices to show that there exists no other degenerate symmetric PBE.

In a given degenerate symmetric PBE, the customer's *ex-ante* payoff is 0. Then, every dealer must offer the monopoly price v deterministically whenever it receives an RFQ and decides to respond. Thus, upon receiving an RFQ and under any belief about how many other dealers are contacted by the customer, a given dealer j can secure a strictly positive payoff by offering some price $p \in (c, v)$. Hence, dealer j optimally responds with probability 1 and offers the monopoly price v deterministically. For such a response strategy to be optimal for dealer j against other dealers' response strategy, dealer j has to believe that the customer contacted no other dealers. Proposition 7 follows. $\square$

Proof of Theorem 3. Given a design of information disclosure (S, μ) and a symmetric PBE, I fix any signal realization $s \in S$ that is drawn with a positive probability under the PBE. That is, $\sum_{n=1}^{\infty} \mu(s|n)\xi(n) > 0$, where $\xi(n)$ is the prior probability that the customer contacts n dealers under the PBE. It suffices to show that the customer's expected payoff conditional on the signal realization s does not exceed π_2^*.

I let χ_s denote a dealer's posterior belief about the number n of contacted dealers, and let (a_s^μ, F_s^μ) denote the dealer's equilibrium strategy upon observing the signal s. If $\chi_s(n) = 1$ for $n = 1$, then it would be strictly optimal for the dealer to respond with probability 1 and offer the monopoly price v deterministically. Thus, the customer's expected payoff conditional on the signal realization s equals 0. If $\chi_s(n) < 1$ for $n = 1$, the price distribution F_s^μ cannot have any atom: Given that responding with any non-positive price is strictly dominated by not responding, then $F_s^\mu(0) = 0$; If F_s^μ had an atom at some price $p^0 > 0$, slightly undercutting it by offering a price $p^0 - \varepsilon$ would yield a strictly higher payoff than offering p^0 for at least one contacted dealer. Given that responding with any price higher than the customer's value v is strictly worse than not responding, then the upper bound $\hat{p} \coloneqq \sup(\operatorname{supp} F_s^\mu)$ of a dealer's price support is at most v, $\hat{p} \leq v$. When a dealer offers a price $p \in \operatorname{supp} F_s^\mu$ within its support that arbitrarily approaches its upper bound $\hat{p}$, the dealer's expected trading profit approaches $\hat{p} \sum_{n \geq 1} \chi_s(n) \left(1 - a_s^\mu\right)^{n-1}$. If the dealer offers price v, its expected trading profit equals $v \sum_{n \geq 1} \chi_s(n) \left(1 - a_s^\mu\right)^{n-1}$. The dealer's optimality condition implies that $\hat{p} = v$. Then, the dealer's individual rationality is given by

$$v \sum_{n \geq 1} \chi_s(n) \left(1 - a_s^{\mu}\right)^{n-1} \geq c \iff \sum_{n \geq 1} \chi_s(n) \left(1 - a_s^{\mu}\right)^{n-1} \geq \frac{c}{v}. \tag{13}$$

It then follows that the customer's expected payoff conditional on the signal realization s does not exceed

$$\begin{aligned}
&\underbrace{v \left[1 - \sum_{n \geq 1} \chi_s(n) \left(1 - a_s^{\mu}\right)^{n}\right] - c \sum_{n \geq 1} \chi_s(n) n a_s^{\mu}}_{\text{aggregate conditional expected payoff}} - \underbrace{\left[v \sum_{n \geq 1} \chi_s(n) \left(1 - a_s^{\mu}\right)^{n-1} - c\right] a_s^{\mu}}_{\text{one dealer's conditional expected payoff}} \\
&= v \left[1 - \sum_{n \geq 1} \chi_s(n) \left(1 - a_s^{\mu}\right)^{n-1}\right] - c \sum_{n \geq 1} \chi_s(n)(n-1) a_s^{\mu} \\
&\leq v \left[1 - \sum_{n \geq 1} \chi_s(n) \left(1 - a_s^{\mu}\right)^{n-1}\right] - c \sum_{n \geq 1} \chi_s(n) \left[1 - \left(1 - a_s^{\mu}\right)^{n-1}\right] \\
&\leq v \left[1 - \frac{c}{v}\right] - c \left[1 - \frac{c}{v}\right] \qquad \text{(following from (13))} \\
&= \pi_2^*.
\end{aligned}$$

□

Proof of Proposition 3. The proof is given immediately after Proposition 3. □

Proof of Proposition 4. *Part (i)*: Given that $c/(vq)$ is strictly decreasing in order size q and value v, probability a_2^q is thus strictly increasing in size q and value v.

Part (ii): Fixing any $p \in [c/q, v]$,

$$1 - F_2^q(p) = \frac{\frac{c}{p} - \frac{c}{v}}{q - \frac{c}{v}}$$

is strictly decreasing in size q and strictly increasing in value v. Therefore, $F_2^q \succ_{(1)} F_2^{q'}$ for $q < q'$, and $F_2^{v'} \succ_{(1)} F_2^{v}$ for $v < v'$. □

C Heterogeneous Response Costs

The benchmark model in Section 2 can be easily extended to allow for heterogeneous response costs across dealers. This appendix shows that the customer continues to contact $n^* = 2$ dealers in the extended model.

To the benchmark model in Section 2, I add a Stage 0 in which each dealer j privately observes its response cost c_j. The response costs are independently and identically distributed with a CDF H whose support is within $[0, v]$, $c_j \overset{\text{iid}}{\sim} H$, $H(0) = 0$ and $H(v^-) = 1$. The remaining setup is identical to the benchmark model in Section 2.

Proposition 8. *In any symmetric subgame perfect equilibrium of the extended model with heterogeneous response costs, the customer contacts $n^* = 2$ dealers. Upon observing the customer's choice of n, each dealer j responds with probability $a_n^H(c_j)$, which is a function of its response cost c_j,*

$$a_n^H(c_j) = \begin{cases} 1 & \text{if } n > 1 \text{ and } c_j < c^*, \\ 0 & \text{if } n > 1 \text{ and } c_j > c^*, \\ \dfrac{a_n^*(c^*) - H(c_j^-)}{H(c_j) - H(c_j^-)} & \text{if } n > 1, c_j = c^*, \text{ and } H(c_j) > H(c_j^-), \\ 1 & \text{if } n = 1, \end{cases}$$

where $c^* \in (0, v)$ is uniquely determined by

$$v[1 - H(c^*)]^{n-1} \leq c^* \leq v[1 - H(c^{*-})]^{n-1},$$

and $a_n^*(c^*)$ has the same expression as a_n^* given by (7) with cost c being replaced by c^*,

$$a_n^*(c^*) = 1 - \left(\frac{c^*}{v}\right)^{\frac{1}{n-1}}.$$

When the dealer responds, its price distribution F^H, unconditional on the realization of its response cost, has the same expression as F^* given by (8) with cost c being replaced by c^*,

$$F_n^H(p) = \frac{1 - \left(\frac{c^*}{p}\right)^{\frac{1}{n-1}}}{1 - \left(\frac{c^*}{v}\right)^{\frac{1}{n-1}}}, \quad \textit{and} \ \operatorname{supp} F_n^H = [c^*, v].$$

Compared to Theorem 1, the only difference is how a dealer decides whether to respond in equilibrium. Instead of mixing between responding or not, dealer j decides to respond based on its response cost c_j: Dealer j responds with probability 1 if its response cost c_j is below threshold c^*, and does not respond if its cost c_j is above the threshold; if its cost c_j equals the threshold, j is indifferent and may mix between responding or not. When the cost distribution H has no atom, allowing heterogeneous costs purifies the dealer's decision to respond. Unconditionally, a dealer's response probability is $a_n^*(c^*)$ as if all dealers' response costs were homogeneous and equal to c^*. Conditional on responding, the response cost c_j is sunk. Hence, the dealer's problem of what price to offer remains the same as in the case with homogeneous cost c^*. Therefore, the customer continues to contact $n^* = 2$ dealers as in the homogeneous cost case. The proof is otherwise identical to that for Theorem 1.

D The Deadweight Nature of the Response Cost

What happens if the dealers' response cost c, rather than being a deadweight loss, is a fee levied by the platform and redistributed to the customer? I show that (1) the customer continues to contact two dealers, and (2) her *ex-ante* payoff is $\pi_2^* + 2a_2^* c$, which is her equilibrium payoff π_2^* in the benchmark model of Section 2 plus the expected fee $2a_2^* c$ paid by the two dealers.

The benchmark model is modified in one way: Each dealer pays a fee $c > 0$ to respond to the RFQ in Stage 2. The fee paid by all responding dealers is transferred to the customer upon a successful trade and refunded back to the responding dealers if the RFQ fails. The refund prevents the customer from sending "fake" RFQs only to collect the fee.

Proposition 9. *I consider the modified model with a response fee replacing a deadweight response cost. Its unique symmetric subgame perfect equilibrium is identical to that of the benchmark model,* (n^*, i^*, a^*, F^*)*, given in Theorem 1.*

Proof. In Stage 3, the customer continues to trade with the dealer offering the lowest price if that price does not exceed the customer's reservation price v. Thus, in Stage 2, the dealers' subgame remains unaffected, and the dealers continue to follow the equilibrium strategy (a^*, F^*). In Stage 1, the customer's payoff is 0 upon contacting $n = 1$ dealer. If the customer contacts $n > 1$ dealers, her expected payoff is given by

$$
\begin{aligned}
\pi_n^* + na_n^* c &= v - \underbrace{\mathbb{E}_{G_n^*}(p \wedge v)}_{\text{suppression of competition}} + \underbrace{na_n^* c}_{\text{fee revenue}} \qquad \text{(following from (9))} \\
&= v - \left(c + \int_c^v [1 - G_n^*(p)] \mathrm{d}p \right) + na_n^* c \\
&= v - \left(c + \int_c^v [1 - a_n^* F_n^*(p)]^n \,\mathrm{d}p \right) + na_n^* c \\
&= v - c - \int_c^v \left(\frac{c}{p} \right)^{\frac{n}{n-1}} \mathrm{d}p + n \left[1 - \left(\frac{c}{v} \right)^{\frac{1}{n-1}} \right] c \qquad \text{(following from (7) and (8))} \\
&= v - c - (n-1)c \left[1 - \left(\frac{c}{v} \right)^{\frac{1}{n-1}} \right] + nc \left[1 - \left(\frac{c}{v} \right)^{\frac{1}{n-1}} \right] \\
&= v - c \left(\frac{c}{v} \right)^{\frac{1}{n-1}} .
\end{aligned}
$$

The customer's expected payoff is thus strictly decreasing in $n \geq 2$. Therefore, it is strictly optimal for the customer to contact $n^* = 2$ dealers. □

For some fee-to-value ratio c/v, the aggregate fee revenue $na_n^* c$ expected by the customer could be strictly increasing in the number n of contacted dealers (Figure 3, when $c/v = 0.1$). However, dealers' strategic avoidance of competition causes the customer's trading profit to decline more steeply with n, dominating any fee incentive.

This variant shuts down the deadweight loss aspect of the response cost while preserving the suppression of dealer competition caused by the cost. What drives my result is not the deadweight cost of responding to an RFQ *per se*, because the equilibrium is completely unaffected even if the cost becomes a mere transfer from the dealers to the customer. The key force is dealers' strategic avoidance of competition, arising from the cost and enabled by the dealers' ability to ignore the RFQ.

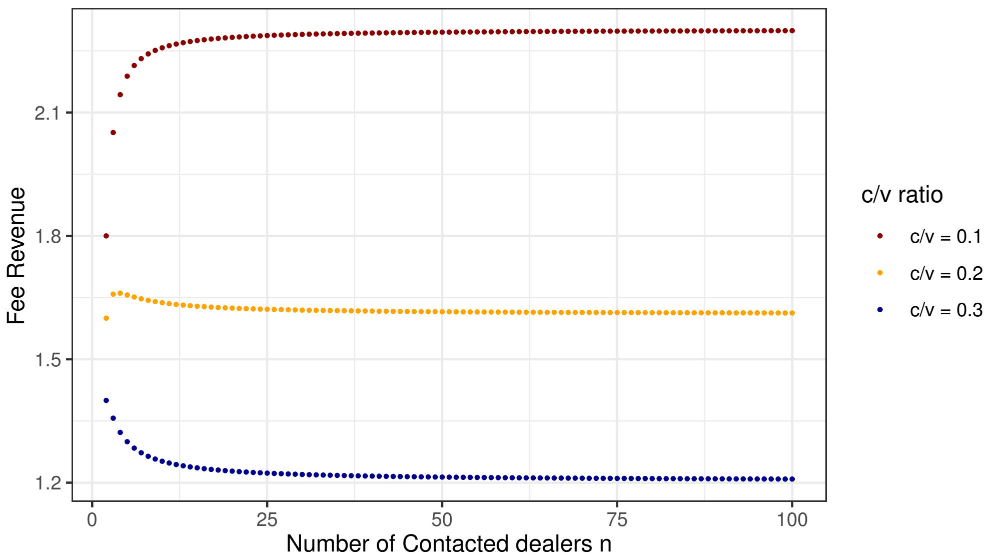

Figure 3: Aggregated Expected Fee Revenue

E Risk-Averse Dealers

This appendix solves the case where dealers are risk-averse. I modify the benchmark model by letting the dealer's utility function u be concave. The remaining setup remains identical to the benchmark model. It remains strictly optimal for the customer to contact $n^* = 2$ dealers in the unique symmetric subgame perfect equilibrium of this modified model. I first obtain this result by applying the intuition for Proposition 1 (dealers' strategic avoidance of competition). Then, I formally solve for the equilibrium.

The intuition for Proposition 1 (dealers' strategic avoidance of competition) remains valid.

When a dealer j offers a given price p in its price support, its expected utility

$$\underbrace{\gamma_n^u(p)}_{\mathbb{P}(\text{Dealer } j \text{ wins the trade when offering a given price } p)} \times u(p-c) + [1-\gamma_n^u(p)] \times u(-c)$$

equals that of its outside option, $u(0)$. Then, the remaining steps follow in the exact same manner. When the customer contacts one more dealer, $\gamma_n^u(p)$ should remain constant. Given that $\gamma_n^u(p) \coloneqq \mathbb{P}\left(\text{Dealer } j \text{ wins the trade when offering a given price } p\right) = [\theta_n^u(p)]^{n-1}$, where

$$\theta_n^u(p) \coloneqq \mathbb{P}\left(\text{a given dealer does not offer a price lower than } p\right),$$

then $[\theta_{n+1}^u(p)]^n = [\theta_n^u(p)]^{n-1}$. Thus, $\mathbb{P}\left(\text{a given dealer does not offer a price lower than } p\right)$ should rise, $\theta_{n+1}^*(p) > \theta_n^*(p)$. Hence, $\mathbb{P}\left(\text{no dealer offers a price lower than } p\right)$ should rise too, $[\theta_{n+1}^*(p)]^{n+1} > [\theta_n^*(p)]^n$. Therefore, dealers' best overall price offer becomes stochastically larger, $G_{n+1}^* \succ_{(1)} G_n^*$. The intuition for Proposition 1 (dealers' strategic avoidance of competition) remains valid when dealers are risk-averse, and results in a strictly less competitive best overall price $p = \min_{j=1,\dots,n} p_j$ for the customer if the customer contacts one more dealer. Consequently, it remains strictly optimal for the customer to contact two dealers.

Formal solution for the equilibrium

Proposition 10. *The modified model where dealers are risk-averse has a unique symmetric subgame perfect equilibrium (n^*, i^*, a^u, F^u), where n^* and i^* are the same as in Theorem 1, and*

$$a_n^u = \begin{cases} 1 - \left(\dfrac{u(0) - u(-c)}{u(v-c) - u(-c)}\right)^{\frac{1}{n-1}} & \text{if } n > 1, \\ 1 & \text{if } n = 1, \end{cases}$$

$$F_n^u(p) = \begin{cases} \dfrac{1 - \left(\frac{u(0)-u(-c)}{u(p-c)-u(-c)}\right)^{\frac{1}{n-1}}}{1 - \left(\frac{u(0)-u(-c)}{u(v-c)-u(-c)}\right)^{\frac{1}{n-1}}}, \text{ and } \operatorname{supp} F_n^u = [c, v] & \text{if } n > 1, \\ \mathbb{1}_{p \geq v} & \text{if } n = 1. \end{cases}$$

Compared to Theorem 1, the customer's equilibrium strategy (n^*, i^*) remains unchanged.

The only difference lies in the expressions of the dealer's equilibrium strategy (a_n^u, F_n^u), which arises from the fact that the dealer's indifference condition (6) becomes

$$\underbrace{[1 - a_n^u F_n^u(p)]^{n-1}}_{\mathbb{P}(\text{Dealer } j \text{ wins the trade})} \; u(p - c) + \underbrace{\left(1 - [1 - a_n^u F_n^u(p)]^{n-1}\right)}_{\mathbb{P}(\text{Dealer } j \text{ loses the trade})} u(-c) = u(0), \quad \forall p \in \operatorname{supp} F_n^u.$$

The remaining proof is otherwise identical to that of Theorem 1. Some algebra shows that $a^u = a^*$ and $F^u = F^*$ if u is linear, and $a_n^u < a_n^*$ and $F_n^u \succ_{(1)} F_n^*$ for all $n > 1$ if u is strictly concave: Risk aversion causes every dealer to respond with a lower probability and offer a stochastically worse price when responding.

Internet Appendices

A Sequential Requests for Quote

This section solves an extension with three sequential RFQs, where each dealer incurs a cost c only once upon responding to one or more RFQs. Through a backward induction, I show that the customer in each RFQ continues to contact two dealers, in the unique symmetric Markov perfect equilibrium of this extension.

Trading Game. Formally, the three-period trading game works as follows. A dealer incurs no cost upon responding to an RFQ and is thus called a "free dealer" if it had already responded to a prior RFQ. Otherwise, the dealer is called a "non-free dealer" and incurs a response cost $c > 0$. In each period k ($k = 1, 2, 3$), an RFQ proceeds in three stages. In Stage 1, customer k chooses a number n_1^k of free dealers and a number n_0^k of non-free dealers to contact. Observing the total number f^k of free dealers at the beginning of period k and the customer's choice (n_0^k, n_1^k), each contacted dealer j chooses whether to respond and what price p_j^k to offer in Stage 2. In Stage 3, the customer chooses whether and against which dealer's price to trade. I assume two tie-breaking rules: Whenever the customer is indifferent among several options, she first minimizes the number of dealers she contacts, and then maximizes the probability of a trade.[16] The asset's expected payoff is normalized to 0, and the customer has an additional value v of owning the asset. The response cost c is assumed to be less than the value v ($c < v$) so that the RFQ has a positive net surplus from trade at all possible contingencies. As in the benchmark model of Section 2, all agents are risk-neutral with no time discounting.

[16]The priority of these two tie-breaking rules is irrelevant. All equilibrium results continue to hold if a customer first maximizes the probability of a trade, and then minimizes the number of dealers she contacts.

Strategies and equilibrium concept. The strategy of customer k consists of a triple $\left(n_0^k, n_1^k, i^k\right)$, following which the customer contacts n_1^k free dealers and n_0^k non-free dealers in Stage 1, and trades with dealer $i^k \in \{0, 1, \ldots, n^k\}$ in Stage 3, where $n^k := n_0^k + n_1^k$. The strategy of dealer j consists of a couple $\left(a_{j,h_j^k}^k, F_{j,h_j^k}^k\right)$ for each period $k \in \{1, 2, 3\}$, which can depend on all information h_j^k available to the dealer right before it responds to the k^{th} RFQ.

A symmetric Markov perfect equilibrium[17] is a symmetric subgame perfect equilibrium such that every dealer's equilibrium strategy depends on its information set only through the payoff-relevant state $(f^k, f_j^k, n_0^k, n_1^k)$, where $f_j^k := \mathbb{1}\{j \text{ is free at the beginning of period } k\}$:

$$\text{(Markov symmetry)}\ a_{j,h_j^k}^{*k} = a_{f_j^k, f^k, n_0^k, n_1^k}^{*k} \text{ and } F_{j,h_j^k}^{*k} = F_{f_j^k, f^k, n_0^k, n_1^k}^{*k}$$
$$\text{for } k \in \{1, 2, 3\},\ f^k, n_0^k \in \mathbb{Z}^+,\ n_1^k \le f^k,\ j \le n^k, \text{ and } f_j^k \le \max\{f^k, 1\}.$$

Equilibrium

Proposition A.1. *The three-period model has a unique symmetric Markov perfect equilibrium consisting of the following strategies:*

$$n_1^{*k} = \min\left\{f^k, 2\right\},\ n_0^{*k} = 2 - n_1^{*k},\ i^{*k} \begin{cases} \in \operatorname*{argmin}_{j=1,\ldots,n} p_j & \textit{if } \min_{j=1,\ldots,n} p_j \le v, \\ = 0 & \textit{if } \min_{j=1,\ldots,n} p_j > v, \end{cases} \quad \textit{for } k = 1, 2, 3,$$

[17] The solution concept of Markov perfect equilibrium has been widely applied in publications starting about 1988 in the work of Jean Tirole and Eric Maskin (Maskin and Tirole, 1988b,a; Fudenberg and Tirole, 1991). It is formally defined in Maskin and Tirole (2001).

	$n_1^k > 1$	$k = 3$ *or* $f^k > 0,\ n_1^k \leq 1$	$k = 2,\ f^k = 0$	$k = 1$
$a^{*k}_{1,f^k,n_0^k,n_1^k}$	*1*	*1*	NA	NA
$F^{*k}_{1,f^k,n_0^k,n_1^k}$	$\mathbb{1}_{p\geq 0}$	$a^c_{n^k} F^c_{n^k} + \left(1 - a^c_{n^k}\right) \mathbb{1}_{p\geq v}$	NA	NA
$a^{*k}_{0,f^k,n_0^k,n_1^k}$	*0*	$a^c_{n^k} := a^*_{n^k}$	$a^{c''}_{n^k}$	$a^{c'}_{n^k}$
$F^{*k}_{0,f^k,n_0^k,n_1^k}$	NA	$F^c_{n^k} := F^*_{n^k}$	$F^{c''}_{n^k}$	$F^{c'}_{n^k}$

Here, I write a^c for a^* and F^c for F^* (both are given in Theorem 1) to underline their dependence on the cost c, and

$$c'' := \frac{cv}{c+v}, \qquad c' := \frac{cv^2}{v^2 + cv + c^2}.$$

If all dealers are non-free at the beginning of an RFQ ($f^k = 0$), responding to the RFQ spares a dealer the cost of responding to subsequent RFQs, giving that dealer a potential advantage over other dealers and thus a higher continuation payoff. The anticipation of this gain in continuation payoff reduces every dealer's effective cost of responding to the current RFQ down to some c'' in period 2 and c' in period 1, respectively. I compute the effective response costs c'' and c' in the formal proof below. Hence, the dealers respond with probability $a^{c''}_{n^k}$ and offer the price distribution $F^{c''}_{n^k}$ in period $k = 2$, and follow $(a^{c'}_{n^k}, F^{c'}_{n^k})$ in period $k = 1$. Therefore, it remains strictly optimal for the customer in each of periods 2 and 1 to contact two dealers.

If some dealers are free at the beginning of an RFQ ($f^k > 0$), responding to the RFQ does not change the continuation payoff of any dealer j, regardless of whether j was free before responding. Hence, one can solve for the current RFQ's equilibrium strategies as if it were in the last period and thus insulated from other RFQs. Following the steps above, I show that every customer continues to contact two dealers in all circumstances.

Now, I proceed with a backward induction to formally solve for the symmetric Markov perfect equilibrium.

Proof of Proposition A.1. In Stage 3 of any customer RFQ, it is always optimal for the customer to trade with the dealer who offers the lowest price if that price does not exceed the customer's private value v, and does not trade otherwise.

In Stage 2, if the customer contacts more than one free dealer ($n_1^k > 1$), those free dealers respond with probability $a^{*k}_{1,f^k,n_0^k,n_1^k} = 1$ and offer the competitive price 0 deterministically, while any non-free dealer responds with probability $a^{*k}_{0,f^k,n_0^k,n_1^k} = 0$. If the customer contacts a total of $n^k = 1$ dealer, then that dealer responds with probability 1 and offers the monopoly price v deterministically. It remains to consider cases where $n_1^k \leq 1, n^k > 1$.

- Period 3 ($k = 3$):

When the customer contacts only non-free dealers ($n_1^k = 0$), those non-free dealers' equilibrium response strategy consists and only consists of a symmetric subgame perfect equilibrium being played in the benchmark model's second stage. Thus, those non-free dealers respond with probability $a^c_{n^k}$ and offer the price distribution $F^c_{n^k}$ conditional on responding (Theorem 1).

If the customer contacts $n_1^k = 1$ free dealer, that free dealer responds with probability $a^{*k}_{1,f^k,n_0^k,n_1^k} = 1$, because it can undercut all non-free dealers and thus secure a strictly positive payoff by offering a price slightly below c. The price distributions $F^{*k}_{1,f^k,n_0^k,n_1^k}$ and $F^{*k}_{0,f^k,n_0^k,n_1^k}$ cannot both have an atom at the same price: Given that responding with any non-positive price is strictly dominated by not responding for a non-free dealer, then $F^{*k}_{0,f^k,n_0^k,n_1^k}(0) = 0$. If $F^{*k}_{1,f^k,n_0^k,n_1^k}$ and $F^{*k}_{0,f^k,n_0^k,n_1^k}$ both had an atom at some price $p^0 > 0$, undercutting it by offering a price $p^0 - \varepsilon$ would yield a strictly higher payoff than offering p^0 for at least one dealer.[18] Letting the upper bounds of the free and a non-free dealers' price supports be $\hat{p}_1 := \sup(\text{supp } F^{*k}_{1,f^k,n_0^k,n_1^k})$ and $\hat{p}_0 := \sup(\text{supp } F^{*k}_{0,f^k,n_0^k,n_1^k})$, respectively, I show that $\hat{p}_1 = v$. When a non-free dealer offers a price $p \in \text{supp } F^{*k}_{0,f^k,n_0^k,n_1^k}$ within its support that arbitrarily approaches its upper bound $\hat{p}_0$, it gets to trade with the customer with a probability of at most $1 - F^{*k}_{1,f^k,n_0^k,n_1^k}(\hat{p}_0^-)$. Then, $F^{*k}_{1,f^k,n_0^k,n_1^k}(\hat{p}_0^-) < 1$. Thus, $\hat{p}_0 = \hat{p}_1$ and $F^{*k}_{1,f^k,n_0^k,n_1^k}$ has an atom at $\hat{p}_1$, or $\hat{p}_0 < \hat{p}_1$. In either case,

$$F^{*k}_{0,f^k,n_0^k,n_1^k}(\hat{p}_1^-) = 1, \tag{A.1}$$

as $F^{*k}_{0,f^k,n_0^k,n_1^k}$ and $F^{*k}_{1,f^k,n_0^k,n_1^k}$ cannot both have an atom at $\hat{p}_1$. Given that responding with any price higher than the customer's value v would earn the free dealer a payoff of zero, then $\hat{p}_1 \leq v$. When the free dealer offers a price $p \in \text{supp } F^{*k}_{1,f^k,n_0^k,n_1^k}$ within its support that arbitrarily approaches its upper bound $\hat{p}_1$, it gets to trade with the customer if and only if all non-free dealers do not respond, an event that occurs with probability $(1 - a^{*k}_{0,f^k,n_0^k,n_1^k})^{n_0^k}$. Thus, the free dealer's expected trading profit in period 3 approaches $\hat{p}_1(1 - a^{*k}_{0,f^k,n_0^k,n_1^k})^{n_0^k}$. If the free dealer offers price v, its expected trading profit is at least $v(1 - a^{*k}_{0,f^k,n_0^k,n_1^k})^{n_0^k}$. The free dealer's optimality condition implies that $a^{*k}_{0,f^k,n_0^k,n_1^k} < 1$ and $\hat{p}_1 = v$.

When a non-free dealer offers a price $\underline{p}_1 - \varepsilon$ that arbitrarily approaches the free dealer's lower bound $\underline{p}_1 := \inf(\text{supp } F^{*k}_{1,f^k,n_0^k,n_1^k})$ from below, the non-free dealer's expected trading profit approaches $\underline{p}_1[1 - a^{*k}_{0,f^k,n_0^k,n_1^k} F^{*k}_{0,f^k,n_0^k,n_1^k}(\underline{p}_1^-)]^{n_0^k - 1}$. This limiting gain is at most c, since the non-free dealer's response probability is less than 1, $a^{*k}_{0,f^k,n_0^k,n_1^k} < 1$. When the free dealer offers a price $p \in F^{*k}_{1,f^k,n_0^k,n_1^k}$ within its support that arbitrarily approaches its lower bound $\underline{p}_1$, its expected trading profit is at most $\underline{p}_1[1 - a^{*k}_{0,f^k,n_0^k,n_1^k} F^{*k}_{0,f^k,n_0^k,n_1^k}(\underline{p}_1^-)]^{n_0^k}$. Then, the free

[18] When every dealer offers the same price p^0, at least one dealer trades with the customer with a probability of less than 1 regardless of how the customer breaks her tie. Then, that dealer is strictly better off offering the slightly lower price $p^0 - \varepsilon$ to undercut the other dealers.

dealer's expected trading profit is at most c and thus must equal c,

$$p\left[1-a^{*k}_{0,f^k,n_0^k,n_1^k}F^{*k}_{0,f^k,n_0^k,n_1^k}(p^-)\right]^{n_0^k}=c,\quad \forall p\in \operatorname{supp} F^{*k}_{1,f^k,n_0^k,n_1^k}. \tag{A.2}$$

Setting $p=\hat{p}_1=v$ in the equation above together with equation (A.1) yields

$$v\left[1-a^{*k}_{0,f^k,n_0^k,n_1^k}\right]^{n_0^k}=c\iff a^{*k}_{0,f^k,n_0^k,n_1^k}=1-\left(\frac{c}{v}\right)^{\frac{1}{n^k-1}}.$$

A non-free dealer's expected trading profit also equals c,

$$p\left[1-a^{*k}_{0,f^k,n_0^k,n_1^k}F^{*k}_{0,f^k,n_0^k,n_1^k}(p^-)\right]^{n_0^k-1}\left[1-F^{*k}_{1,f^k,n_0^k,n_1^k}(p^-)\right]=c,\quad \forall p\in \operatorname{supp} F^{*k}_{0,f^k,n_0^k,n_1^k}. \tag{A.3}$$

Then, equations (A.2) and (A.3) uniquely determine the price distributions $F^{*k}_{0,f^k,n_0^k,n_1^k}$ and $F^{*k}_{1,f^k,n_0^k,n_1^k}$,

$$F^{*k}_{0,f^k,n_0^k,n_1^k}=F^c_{n^k},\qquad F^{*k}_{1,f^k,n_0^k,n_1^k}=a^c_{n^k}F^c_{n^k}+(1-a^c_{n^k})\,\mathbb{1}_{p\geq v}.$$

Instead of ignoring the RFQ with probability $1-a^c_{n^k}$, the free dealer offers the monopoly price v with the same probability $1-a^c_{n^k}$. Otherwise, the n dealers' response strategies are identical to those in the benchmark model of Section 2.

In Stage 1, when there exist at least two free dealers ($f^k\geq 2$), the customer's payoff is v, which she achieves by contacting exactly $n_1^{*k}=2$ free dealers; when there exists at most one free dealer ($f^k\leq 1$), the customer's expected payoff is π_2^*, which she achieves by contacting $n_1^{*k}=f^k$ free dealer and $n_0^{*k}=2-n_1^{*k}$ non-free dealer(s) due to the two tie-breaking rules. As a result, a dealer's expected payoff in period 3 is summarized as follows.

Lemma A.1. *In all cases, a dealer expects a payoff of c in period 3 if it is the only free* *dealer at the beginning of period 3, and a payoff of* 0 *otherwise.*

- Period 2 ($k = 2$):

If some dealers are free at the beginning of period 2 ($f^k > 0$), then given any contacted dealer j, responding in period 2 does not change its period-3 continuation payoff, regardless of whether j was free before responding: If j was free, j remains free and thus expects the same continuation payoff in period 3, regardless of whether j responds in period 2. If j was non-free, j cannot be the only dealer who is free at the beginning of period 3 as some other dealers were already free. Thus, j expects a continuation payoff of 0 in period 3, regardless of whether j responds in period 2 (Lemma A.1). Hence, one can solve for the current RFQ's equilibrium strategies as if it were in the last period and thus insulated from other RFQs. Therefore,

$$a^{*2}_{f^2_j, f^2, n^2_0, n^2_1} = a^{*3}_{f^2_j, f^2, n^2_0, n^2_1}, \quad F^{*2}_{f^2_j, f^2, n^2_0, n^2_1} = F^{*3}_{f^2_j, f^2, n^2_0, n^2_1}, \quad n^{*2}_0 = n^{*3}_0, \quad n^{*2}_1 = n^{*3}_1, \quad \text{if } f^2 > 0.$$

In particular, the customer contacts a total number of $n^{*2} = n^{*3} = 2$ dealers.

If all dealers are non-free at the beginning of period 2 ($f^k = 0$), then the price distribution $F^{*k}_{0, f^k, n^k_0, n^k_1}$ cannot have an atom at any positive price: If $F^{*k}_{0, f^k, n^k_0, n^k_1}$ had an atom at some price $p^0 > 0$, undercutting it by offering a price $p^0 - \varepsilon$ would yield a strictly higher payoff than offering p^0 for at least one dealer.

Responding in period 2 earns a dealer an additional expected payoff of c in period 3 if no other dealer responds in period 2. Thus, upon offering a price $p \in \operatorname{supp} F^{*k}_{0, f^k, n^k_0, n^k_1}$ within its support that arbitrarily approaches its support's upper bound $\bar{p} := \sup \left(\operatorname{supp} F^{*k}_{0, f^k, n^k_0, n^k_1} \right)$, a dealer's continuation payoff is at most $(\bar{p} \mathbb{1}_{\bar{p} \leq v} + c)(1 - a^{*k}_{0, f^k, n^k_0, n^k_1})^{n^k - 1}$. If the dealer offers price v, its continuation payoff is at least $(v + c)(1 - a^{*k}_{0, f^k, n^k_0, n^k_1})^{n^k - 1}$. The dealer's optimality condition implies that $\bar{p} = v$. The dealer's individual rationality implies that its continuation payoff upon responding is no lower than its response cost, $(v+c)(1 - a^{*k}_{0,f^k,n^k_0,n^k_1})^{n^k-1} \geq c$. Hence, $a^{*k}_{0,f^k,n^k_0,n^k_1} < 1$. Then the dealer's indifference condition implies

$$(c+v)\left[1 - a^{*k}_{0,f^k,n^k_0,n^k_1}\right]^{n^k-1} = c \iff a^{*k}_{0,f^k,n^k_0,n^k_1} = 1 - \left(\frac{c}{c+v}\right)^{\frac{1}{n^k-1}}.$$

Thus, responding in period 2 earns a dealer an additional expected payoff of $c(1-a^{*k}_{0,f^k,n^k_0,n^k_1})^{n^k-1}$ $= c^2/(c+v)$ in period 3, which reduces every dealer's "effective response cost" in period 2 to

$$c'' = c - \frac{c^2}{c+v} = \frac{cv}{c+v}.$$

Therefore, every dealer responds with probability $a^{*k}_{0,f^k,n^k_0,n^k_1} = a^{c''}_{n^k}$ and offers the price distribution $F^{*k}_{0,f^k,n^k_0,n^k_1} = F^{c''}_{n^k}$ conditional on responding (Theorem 1), and the customer contacts a total number of $n^{*k} = 2$ dealers.

- Period 1 ($k = 1$):

As in period 2, the price distribution $F^{*k}_{0,f^k,n^k_0,n^k_1}$ cannot have an atom at any positive price. Responding in period 1 earns a dealer an additional expected payoff of c in period 2 if no other dealer responds in period 1, and another payoff of c in period 3 if no other dealer responds in either period 1 or period 2. Thus, upon offering a price $p \in \operatorname{supp} F^{*k}_{0,f^k,n^k_0,n^k_1}$ within its support that arbitrarily approaches its support's upper bound $\tilde{p} := \sup\left(\operatorname{supp} F^{*k}_{0,f^k,n^k_0,n^k_1}\right)$, a dealer's continuation payoff is at most $[\tilde{p}\mathbb{1}_{\tilde{p}\leq v} + c + c(1-a^{*2}_{0,1,n^{*2}_0,n^{*2}_1})^{n^{*2}-1}](1-a^{*k}_{0,f^k,n^k_0,n^k_1})^{n^k-1} =$ $(\tilde{p}\mathbb{1}_{\tilde{p}\leq v} + c + c^2/v)(1-a^{*k}_{0,f^k,n^k_0,n^k_1})^{n^k-1}$. If the dealer offers price v, its continuation payoff is at least $(v + c + c^2/v)(1 - a^{*k}_{0,f^k,n^k_0,n^k_1})^{n^k-1}$. The dealer's optimality condition implies that $\tilde{p} = v$. The dealer's individual rationality implies that its continuation payoff upon responding is no lower than its response cost, $(v + c + c^2/v)(1 - a^{*k}_{0,f^k,n^k_0,n^k_1})^{n^k-1} \geq c$. Hence, $a^{*k}_{0,f^k,n^k_0,n^k_1} < 1$.

Then, the dealer's indifference condition implies

$$\left(v+c+\frac{c^2}{v}\right)\left(1-a^{*k}_{0,f^k,n_0^k,n_1^k}\right)^{n^k-1}=c \iff \left(1-a^{*k}_{0,f^k,n_0^k,n_1^k}\right)^{n^k-1}=\frac{c}{v+c+\frac{c^2}{v}}.$$

Thus, responding in period 2 earns a dealer an additional combined payoff of $(c+c^2/v)(1-a^{*k}_{0,f^k,n_0^k,n_1^k})^{n^k-1}=(c^2v+c^3)/(v^2+vc+c^2)$ in periods 2 and 3, which reduces every dealer's effective response cost in period 1 to

$$c'=c-\frac{c^2v+c^3}{v^2+cv+c^2}=\frac{cv^2}{v^2+cv+c^2}.$$

Therefore, every dealer responds with probability $a^{*k}_{0,f^k,n_0^k,n_1^k}=a^{c'}_{n^k}$ and offers the price distribution $F^{*k}_{0,f^k,n_0^k,n_1^k}=F^{c'}_{n^k}$ conditional on responding (Theorem 1), and the customer contacts a total number of $n^{*k}=2$ dealers. □

B Exogenous Budget Constraint

If some dealers are subject to an exogenous budget constraint, the equilibrium results remain qualitatively identical to those in the case where some dealers are subject to an exogenous availability constraint (Appendix A). This section solves two variants of the model that include an exogenous budget constraint.

Variant B1

Variant B1 alters Variant A1 of Appendix A as follows. In Stage 0, every dealer j privately observes whether it is subject to a budget constraint $p_j \leq z$, which occurs with probability $1-\beta$ independently across dealers ($\beta < 1$). In Stage 1, a customer seeking to *sell* one unit of an asset chooses a number n of dealers to contact in an RFQ. The customer can contact at most $\bar{n}$ dealers ($\bar{n} > 1$), because it turns out that the customer will contact as many dealers as is feasible in equilibrium. Observing the customer's choice n, each dealer j chooses what price p_j to *bid* subject to its potential budget constraint $p_j \leq z$ in Stage 2. The asset's expected payoff is normalized to 0, and the customer has a private *disutility* $-v$ of owning the asset due to a liquidity shock. The budget constraint z is assumed to be $-v \leq z < 0$. In Stage 3, the customer chooses whether and against which dealer's price to trade. I do not impose any tie-breaking rule in the case of indifference. All agents are risk-neutral.

Variant B1 differs from Variant A1 in Appendix A in two respects: (1) Variant B1 flips the direction of trade by letting the customer be a seller and the dealers be bidders, in order to introduce the budget constraint for the dealers, and (2) the budget constraint *replaces* the exogenous availability constraint, which would be equivalent to a budget constraint with $z < -v$ such that a constrained dealer would be unable to bid any price acceptable to the customer. (On the other hand, the benchmark model of Section 2 with no budget constraint would be equivalent to a budget constraint $z \geq 0$ that is non-binding as a dealer would not bid a price exceeding the asset's expected payoff 0 anyway.)

As in Variant A1 of Appendix A, the customer would contact as many dealers as is feasible. The next proposition establishes the counterpart of Proposition 5 (Appendix A) for the budget constraint.

Proposition B.1. *Variant B1 has a unique symmetric subgame perfect equilibrium* (n^z, i^{bid},

$F^{\text{cons}}, F^{\text{uncons}}$), *where the customer's strategy is*

$$n^z = \bar{n}, \qquad i^{\text{bid}} \begin{cases} \in \underset{j=1,\dots,n}{\operatorname{argmax}}\; p_j & \text{if } \underset{j=1,\dots,n}{\max}\; p_j > -v, \\ = 0 & \text{if } \underset{j=1,\dots,n}{\max}\; p_j < -v, \\ \in \{\underset{j=1,\dots,n}{\operatorname{argmax}}\; p_j, 0\} & \text{if } \underset{j=1,\dots,n}{\max}\; p_j = -v, \end{cases}$$

and the price distributions F^{cons} *and* F^{uncons} *of a constrained and an unconstrained dealer are*

	$n > 1$	$n = 1$
$F_n^{\text{cons}}(p)$	$\mathbb{1}_{p \geq z}$	$\mathbb{1}_{p \geq -v}$
$F_n^{\text{uncons}}(p)$	$\frac{1-\beta}{\beta}\left[\left(\frac{z}{p}\right)^{\frac{1}{n-1}} - 1\right]$, and supp $F_n^{\text{uncons}} = [z, z(1-\beta)^{n-1}]$	$\mathbb{1}_{p \geq -v}$

In particular, the best overall bid $p = \max_{j=1,\dots,n} p_j$ *becomes first-order stochastically larger as* n *increases. When* $\bar{n} \to \infty$, *the equilibrium best overall bid* $p = \max_{j=1,\dots,\bar{n}} p_j$ *converges in distribution to the competitive limit 0.*

Proof. The backward induction for Variant B1 closely parallels that for Variant A1 (Proposition 5, Appendix A).

In Stage 3, the customer optimally chooses to sell to the dealer with the *highest* bid if that bid *exceeds* the customer's disutility $-v$, and does not sell otherwise.

In Stage 2, if the customer contacts $n = 1$ dealer, then it would be strictly optimal for the dealer to bid the monopoly price $-v$ deterministically, $F_n^{\text{cons}}(p) = F_n^{\text{uncons}}(p) = \mathbb{1}_{p \geq -v}$. When the customer contacts $n > 1$ dealers, the price distribution F_n^{cons} cannot have an atom at any price $p^0 < z$: If F_n^{cons} had an atom at some price $p^0 < z$, slightly overbidding it with a price $p^0 + \varepsilon$ would yield a strictly higher payoff than bidding p^0 for at least one constrained dealer. At least one constrained dealer can secure an expected trading profit of $-z(1-\beta)^{n-1}/n > 0$

or higher by bidding price z. When that dealer bids a price $p \in \text{supp } F_n^{\text{cons}}$ within its support that arbitrarily approaches its support's lower bound $\underline{p}^{\text{cons}} := \inf(\text{supp } F_n^{\text{cons}})$, the dealer's expected trading profit approaches 0 if $\underline{p}^{\text{cons}} < z$. Hence, $\underline{p}^{\text{cons}} = z$. Therefore, a constrained dealer bids price z deterministically, $F_n^{\text{cons}}(p) = \mathbb{1}_{p \geq z}$.

Similar to F_n^{cons}, the price distribution F_n^{uncons} cannot have an atom at any negative price. When an unconstrained dealer bids a price $p \in \text{supp } F_n^{\text{uncons}}$ within its support that arbitrarily approaches its support's lower bound $\underline{p}^{\text{uncons}} := \inf(\text{supp } F_n^{\text{uncons}})$, the dealer's expected trading profit is at most $-\underline{p}^{\text{uncons}} \mathbb{1}_{\{\underline{p}^{\text{uncons}} \geq z\}}(1-\beta)^{n-1}$. If the dealer bids a price $z+\varepsilon$ that arbitrarily approaches z from above, its expected trading profit is at least $-z(1-\beta)^{n-1}$. The dealer's optimality condition implies that $\underline{p}^{\text{uncons}} = z$. Then, the dealer's optimality condition is equivalent to

$$
\begin{aligned}
& - p\left[\beta F_n^{\text{uncons}}(p) + 1 - \beta\right]^{n-1} = -z(1-\beta)^{n-1}, \quad \forall p \in \text{supp } F_n^{\text{uncons}}, \\
\iff & F_n^{\text{uncons}}(p) = \frac{1-\beta}{\beta}\left[\left(\frac{z}{p}\right)^{\frac{1}{n-1}} - 1\right], \quad \text{and supp } F_n^{\text{uncons}} = \left[z, z(1-\beta)^{n-1}\right].
\end{aligned}
$$

In Stage 1, the customer's payoff is 0 upon contacting $n = 1$ dealer. If $n > 1$, the distribution G_n^z of the dealers' best overall bid $p = \max_{j=1,\ldots,n} p_j$ is given by

$$
G_n^z(p) = \left[\beta F_n^{\text{uncons}}(p) + 1 - \beta\right]^n, \quad \forall p \in \left[z, z(1-\beta)^{n-1}\right].
$$

If n increases, $F_n^{\text{uncons}}(p)$ strictly decreases, and thus $G_n^z(p)$ strictly decreases. That is, the dealers' best overall bid $p = \max_{j=1,\ldots,n} p_j$ becomes first-order stochastically *larger* when the customer contacts more dealers. Therefore, the customer's unique optimal choice is $n^z = \bar{n}$. $\square$

Variant B2

Variant B2 differs from Variant B1 in two respects: (1) Each dealer can endogenously decide whether to respond (subject to its potential budget constraint) at a cost c, which is assumed to be less than v ($c < v$) so that there is a positive net surplus from trade. (2) The customer could contact any arbitrary number of dealers. Variant B2 is otherwise identical to Variant B1. That is, Variant B2 *reintroduces* dealers' ability to ignore the RFQ to Variant B1.

Variant B2 is the counterpart of Variant A2 in Appendix A. As in Variant A2, an interior solution for the equilibrium number of contacted dealers resurfaces. The next proposition establishes the counterpart of Proposition 6 (Appendix A) for the budget constraint.

Proposition B.2. *All symmetric subgame perfect equilibria* $(n^{z,c}, i^{\text{bid}}, a^{\text{cons},c}, F^{\text{cons},c}, a^{\text{uncons},c}, F^{\text{uncons},c})$ *of Variant B2 are such that* i^{bid} *is the same as in Proposition B.1,*

$$n^{z,c} = \underline{m}, \ldots, \text{ or } \bar{m},$$

and the strategies $(a^{\text{cons},c}, F^{\text{cons},c})$ *and* $(a^{\text{uncons},c}, F^{\text{uncons},c})$ *of a constrained and an unconstrained dealer satisfy*

$$(1-\beta)\, a_n^{\text{cons},c} \left[1 - F_n^{\text{cons},c}(p)\right] + \beta\, a_n^{\text{uncons},c} \left[1 - F_n^{\text{uncons},c}(p)\right] = a_n^* F_n^*(-p) \text{ if } n \geq \bar{m}, \tag{B.1}$$

	$n \leq \underline{m}$	$\underline{m} < n < \bar{m}$
$a_n^{\text{cons},c}(p)$	*1*	$\max\left\{\tilde{a},\ 1 - \frac{\left(\frac{c}{v}\right)^{\frac{1}{n-1}}}{1-\beta}\right\}$
$F_n^{\text{cons},c}(p)$	$F_n^{\text{cons}}(p)$	$\min\left\{1 - \frac{1-\frac{1}{1-\beta}\left(\frac{c}{-p}\right)^{\frac{1}{n-1}}}{1-\frac{1}{1-\beta}\left(\frac{c}{v}\right)^{\frac{1}{n-1}}},\ 1-\delta\right\} \mathbb{1}_{p \geq -v} + \delta \mathbb{1}_{p \geq z}$
$a_n^{\text{uncons},c}(p)$	*1*	*1*
$F_n^{\text{uncons},c}(p)$	$F_n^{\text{uncons}}(p)$	$F_n^{\text{uncons}}(p)$.

Here, $(F^{\text{cons}}, F^{\text{uncons}})$ are given in Proposition B.1; (a^*, F^*) are given in Theorem 1; $\underline{m}$ and $\bar{m}$ are uniquely determined by

$$a^*_{\underline{m}} F^*_{\underline{m}} \left(\frac{-z}{\underline{m}}\right) \geq \beta > a^*_{\underline{m}+1} F^*_{\underline{m}+1} \left(\frac{-z}{\underline{m}+1}\right) \text{ and}$$
$$a^*_{\bar{m}-1} F^*_{\bar{m}-1}(-z) > \beta \geq a^*_{\bar{m}} F^*_{\bar{m}}(-z), \text{ respectively;}$$

and $\tilde{a} \in (0,1)$ is uniquely determined by

$$\frac{1 - [1-\tilde{a}]^n}{\tilde{a}} = \frac{nc}{-z(1-\beta)^{n-1}}, \tag{B.2}$$

and

$$\delta = \frac{\tilde{a}}{a_n^{\text{cons},c}}.$$

An interior solution on the number n resurfaces in Variant B2 following an argument similar to that for Variant A2 in Appendix A. When the n contacted dealers are able to strategically ignore the RFQ, the budget constraint becomes non-binding when $m \geq \bar{m}$, because each dealer would bid above the constraint z with a probability $a_n^* F_n^*(-z)$ that is lower than β anyway. Thus, the dealers behave as if the budget constraint were absent: Each dealer's unconditional bid distribution, given by the left-hand side of equation (B.1), remains to be $a_n^* F_n^*(-p)$. Hence, the customer strictly prefers to contact fewer dealers in this range. Therefore, an interior solution on the equilibrium number of contacted dealers resurfaces, as in Proposition 6 (Appendix A).

The equilibrium is formally solved as follows.

Proof of Proposition B.2. In Stage 3, the customer's optimal dealer choice remains to be i^{bid}, as given in Proposition B.1.

In Stage 2, when the customer contacts $n > 1$ dealers, the price distribution $F_n^{\text{uncons},c}$ cannot have any atom. Otherwise, bidding slightly above the atom would yield a strictly higher payoff for at least one unconstrained dealer. Likewise, the price distribution $F_n^{\text{cons},c}$ cannot have an atom at any price $p < z$.

- If the customer contacts $n \geq \bar{m}$ dealers: If an unconstrained dealer j had a strictly positive expected payoff, j would respond with probability 1. When j bids a price $p \in \operatorname{supp} F_n^{\text{uncons},c}$ within its support that arbitrarily approaches its support's lower bound $\underline{p}^{\text{uncons},c} := \inf\left(\operatorname{supp} F_n^{\text{uncons},c}\right)$, its expected trading profit is at most $-\underline{p}^{\text{uncons},c}(1-\beta)^{n-1}$. Given that $n \geq \bar{m}$ is equivalent to

$$-z(1-\beta)^{n-1} \leq c,$$

it must be that $\underline{p}^{\text{uncons},c} < z$. Then a constrained j must have the same expected payoff as an unconstrained j, which was assumed to be strictly positive. These would imply $a_n^{\text{cons},c} = 1$ and $a_n^{\text{uncons},c} = 1$. Thus, sparing the n dealers the response cost c would not make any dealer's equilibrium response strategy suboptimal. Hence, the n dealers' response strategy would constitute a subgame perfect equilibrium being played in the second stage of Variant B1, where every dealer responds (subject to its potential budget constraint). However, in the unique subgame perfect equilibrium of Variant B1 (given in Proposition B.1), the lower bound $\inf(\operatorname{supp} F_n^{\text{uncons}})$ of an unconstrained dealer's support equals z, contracting $\underline{p}^{\text{uncons},c} < z$. Therefore, an unconstrained dealer's expected payoff must equal 0. The optimality of an unconstrained dealer implies that the n dealers' unconditional response strategy constitutes a subgame perfect equilibrium being played in the second stage of an RFQ with no budget constraint. Equation (B.1) thus follows.

- If the customer contacts $n \le \underline{m}$ dealers: Equivalently,

$$\frac{-z(1-\beta)^{n-1}}{n} \ge c.$$

If a constrained dealer responds with a probability of less than 1 ($a_n^{\text{cons},c} < 1$), then at least one constrained dealer can secure an expected payoff that strictly exceeds $-z(1-\beta)^{n-1}/n - c \ge 0$ by bidding price z, contradicting the assumption that the dealer responds with a probability of less than 1. Thus, $a_n^{\text{cons},c} = 1$. An unconstrained dealer can secure an expected payoff of $-z(1-\beta)^{n-1} - c > 0$ or higher by bidding a price $z + \varepsilon$ that arbitrarily approaches z from above. Thus, an unconstrained dealer also responds with probability 1, $a_n^{\text{uncons},c} = 1$. Then, sparing the n dealers the response cost c would not make any dealer's equilibrium response strategy suboptimal. Hence, the n dealers' response strategy constitutes a subgame perfect equilibrium being played in the second stage of Variant B1, where every dealer responds (subject to its potential budget constraint). Therefore,

$$F^{\text{cons},c} = F^{\text{cons}}, \qquad F^{\text{uncons},c} = F^{\text{uncons}}, \qquad a_n^{\text{cons},c} = 1, \qquad a_n^{\text{uncons},c} = 1.$$

- When the customer contacts $\underline{m} < n < \bar{m}$ dealers: If a constrained dealer j had a strictly positive expected payoff, so must an unconstrained j. These would imply that $a_n^{\text{cons},c} = 1$ and $a_n^{\text{uncons},c} = 1$. Then, sparing the n dealers the response cost c would not make any dealer's equilibrium response strategy suboptimal. Hence, the n dealers' response strategy would constitute a subgame perfect equilibrium being played in the second stage of Variant B1, where every dealer responds (subject to its potential budget constraint). However, under the equilibrium given in Proposition B.1, at least one constrained dealer's expected trading profit is bounded from above by $-z(1-\beta)^{n-1}/n$, which is strictly less than the response cost c for $n > \underline{m}$. This contradicts the constrained dealer's individual rationality. Therefore, a constrained dealer's expected payoff must equal 0. Given that an unconstrained dealer can secure an expected payoff of at least $-z(1-\beta)^{n-1} - c > 0$ by bidding a price $z + \varepsilon$ that arbitrarily approaches z from above, then the dealer never bids below z, $\underline{p}^{\text{uncons},c} \geq z$. That the supports of constrained and unconstrained dealers do not overlap allows me to solve for their response strategies $(a^{\text{cons},c}, F^{\text{cons},c})$ and $(a^{\text{uncons},c}, F^{\text{uncons},c})$ separately.

An unconstrained dealer responds with probability 1, $a^{\text{uncons},c} = 1$. Following identical steps that solve F^{uncons} in the proof of Proposition B.1, one obtains the same solution for $F^{\text{uncons},c}$, $F^{\text{uncons},c} = F^{\text{uncons}}$.

A constrained dealer cannot have an atom at any price $p < z$, and could have an atom at price z. I let $\delta := 1 - F_n^{\text{cons},c}(z^-)$ be the dealer's mass at price z. Then $\delta > 0$, otherwise bidding z would yield an expected payoff of $-z(1-\beta)^{n-1} - c > 0$ for the constrained dealer. Thus,

$$-z \sum_{k=1}^{n} \binom{n}{k} [(1-\beta)a_n^{\text{cons},c}\delta]^{k-1}[1-\beta-(1-\beta)a_n^{\text{cons},c}\delta]^{n-k} = nc.$$

The left-hand side above is a given dealer's expected trading profit upon bidding z, summed across the n dealers. This sum equals nc, the response cost summed across the n dealers. Letting $\tilde{a} \coloneqq a_n^{\text{cons},c}\delta$, the above equation is equivalent to equation (B.2), which is a polynomial that admits a unique solution $\tilde{a} \in (0, 1)$.

Conditional on responding, if a constrained dealer does not bid z deterministically ($\delta < 1$), then $-v \in \text{supp } F_n^{\text{cons},c}$: Given that responding with any bid lower than the customer's value $-v$ is strictly worse than not responding, then the lower bound $\underline{p}^{\text{cons},c} \coloneqq \inf\left(\text{supp } F_n^{\text{cons},c}\right)$ of the dealer's support is at least $-v$, $\underline{p}^{\text{cons},c} \geq -v$. When the dealer bids a price $p \in \text{supp } F_n^{\text{cons},c}$ within its support that arbitrarily approaches its lower bound $\underline{p}^{\text{cons},c}$, the dealer's expected trading profit approaches $-\underline{p}^{\text{cons},c}[(1-\beta)(1-a_n^{\text{cons},c})]^{n-1}$. If the dealer offers price $-v$, its expected trading profit equals $v\left[(1-\beta)(1-a_n^{\text{cons},c})\right]^{n-1}$. The dealer's optimality condition implies that $\underline{p}^{\text{cons},c} = -v$.

A constrained dealer's expected trading profit upon bidding $-v$ does not exceed its response cost c,

$$v\left[(1-\beta)(1-a_n^{\text{cons},c})\right]^{n-1} \leq c.$$

If $v\left[(1-\beta)(1-\tilde{a})\right]^{n-1} > c$, then $\tilde{a} < a_n^{\text{cons},c}$ and thus $\delta < 1$. Hence, $-v \in \text{supp } F_n^{\text{cons},c}$, and

$$v\left[(1-\beta)(1-a_n^{\text{cons},c})\right]^{n-1} = c \iff a_n^{\text{cons},c} = 1 - \frac{\left(\frac{c}{v}\right)^{\frac{1}{n-1}}}{1-\beta}.$$

If $v\left[(1-\beta)(1-\tilde{a})\right]^{n-1} \leq c$, then either $\delta = 1$ or $v\left[(1-\beta)(1-a_n^{\text{cons},c})\right]^{n-1} < c$. In either case, $-v \notin \text{supp } F_n^{\text{cons},c}$. Thus,

$$\delta = 1, \text{ and } a_n^{\text{cons},c} = \tilde{a}.$$

Overall,

$$a_n^{\text{cons},c} = \max\left\{\tilde{a},\ 1 - \frac{\left(\frac{c}{v}\right)^{\frac{1}{n-1}}}{1-\beta}\right\}, \qquad \delta = \frac{\tilde{a}}{a_n^{\text{cons},c}}.$$

Finally, I solve for $F_n^{\text{cons},c}$ by adding a constrained dealer's indifference condition upon bidding any price $p < z$ within its support:

$$\begin{cases} -p\left[(1-\beta)(1-a_n^{\text{cons},c}(1-F_n^{\text{cons},c}(p)))\right]^{n-1} = c, \quad \forall p < z \text{ and } p \in \operatorname{supp} F_n^{\text{cons},c}, \\ F_n^{\text{cons},c}(z^-) = 1-\delta, \qquad F_n^{\text{cons},c}(z) = 1 \end{cases}$$

$$\iff F_n^{\text{cons},c}(p) = \min\left\{1 - \frac{1 - \frac{1}{1-\beta}\left(\frac{c}{-p}\right)^{\frac{1}{n-1}}}{1 - \frac{1}{1-\beta}\left(\frac{c}{v}\right)^{\frac{1}{n-1}}},\ 1-\delta\right\}\mathbb{1}_{p\geq -v} + \delta\mathbb{1}_{p\geq z}. \qquad \square$$

References

Abad, J., I. Aldasoro, C. Aymanns, M. D'Errico, L. Fache Rousová, P. Hoffmann, S. Langfield, M. Neychev, and T. Roukny (2016): "Shedding Light on Dark Markets: First Insights from the New EU-wide OTC Derivatives Dataset," ESRB Occasional Paper Series 11, European Systemic Risk Board.

Afonso, G., A. Kovner, and A. Schoar (2014): "Trading Partners in the Interbank Lending Market," Working Paper, Federal Reserve Bank of New York and MIT Sloan.

Allen, J. and M. Wittwer (2023): "Centralizing Over-The-Counter Markets?" *Journal of Political Economy*.

Armantier, O. and J. Sporn (2013): "Auctions Implemented by the Federal Reserve Bank of New York during the Great Recession," Staff Reports.

Atkeson, A. G., A. L. Eisfeldt, and P.-O. Weill (2015): "ENTRY AND EXIT IN OTC DERIVATIVES MARKETS," *Econometrica*, 83, 2231–2292.

Baldauf, M. and J. Mollner (2022): "Competition and Information Leakage," Working Paper.

Bech, M. L. and E. Atalay (2010): "The Topology of the Federal Funds Market," *Physica A: Statistical Mechanics and its Applications*, 389, 5223–5246.

Bethune, Z., B. Sultanum, and N. Trachter (2021): "An Information-based Theory of Financial Intermediation," *The Review of Economic Studies*, rdab092.

Bjønnes, G., N. Kathitziotis, and C. Osler (2015): "Price Discrimination and the Cost of Liquidity in OTC Markets," Working Paper.

Burdett, K. and K. L. Judd (1983): "Equilibrium Price Dispersion," *Econometrica*, 51, 955–969.

Chang, B. and S. Zhang (2022): "Financial Market Structure and Risk Concentration," Working Paper.

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

Craig, B. and G. von Peter (2014): "Interbank Tiering and Money Center Banks," *Journal of Financial Intermediation*, 23, 322–347.

Di Maggio, M., A. Kermani, and Z. Song (2017): "The Value of Trading Relations in Turbulent Times," *Journal of Financial Economics*, 124, 266–284.

Duffie, D., N. Gârleanu, and L. H. Pedersen (2005): "Over-the-Counter Markets," *Econometrica*, 73, 1815–1847.

Dugast, J., S. Üslü, and P.-O. Weill (2022): "A Theory of Participation in OTC and Centralized Markets," *The Review of Economic Studies*, rdac010.

Engelbrecht-Wiggans, R. (1987): "On Optimal Reservation Prices in Auctions," *Management Science*, 33, 763–770.

Farboodi, M., G. Jarosch, and R. Shimer (2022): "The Emergence of Market Structure," *The Review of Economic Studies*, rdac014.

Fudenberg, D. and J. Tirole (1991): *Game Theory*, vol. 1, The MIT Press.

Glebkin, S., B. Z. Yueshen, and J. Shen (2022): "Simultaneous Multilateral Search," *The Review of Financial Studies*.

Glode, V. and C. C. Opp (2020): "Over-the-Counter versus Limit-Order Markets: The Role of Traders' Expertise," *The Review of Financial Studies*, 33, 866–915.

Hau, H., P. Hoffmann, S. Langfield, and Y. Timmer (2021): "Discriminatory Pricing of Over-the-Counter Derivatives," *Management Science*, 67, 6660–6677.

Hendershott, T., D. Li, D. Livdan, and N. Schürhoff (2020): "Relationship Trading in Over-the-Counter Markets," *The Journal of Finance*, 75, 683–734.

Hendershott, T., D. Livdan, and N. Schürhoff (2021): "All-to-All Liquidity in Corporate Bonds," Working Paper.

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

Hollifield, B., A. Neklyudov, and C. Spatt (2017): "Bid-Ask Spreads, Trading Networks, and the Pricing of Securitizations," *The Review of Financial Studies*, 30, 3048–3085.

Hugonnier, J., B. R. Lester, and P.-O. Weill (2020): "Heterogeneity in Decentralized Asset Markets," Working Paper.

in't Veld, D. and I. van Lelyveld (2014): "Finding the Core: Network Structure in Interbank Markets," *Journal of Banking & Finance*, 49, 27–40.

Jovanovic, B. and A. J. Menkveld (2022): "Equilibrium Bid-Price Dispersion," *Journal of Political Economy*, 130, 426–461.

Kamenica, E. and M. Gentzkow (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.

King, M. R., C. Osler, and D. Rime (2012): "Foreign Exchange Market Structure, Players, and Evolution," in *Handbook of Exchange Rates*, John Wiley & Sons, Ltd, chap. 1, 1–44.

Lee, T. and C. Wang (2018): "Why Trade Over-the-Counter?" Working Paper.

Lester, B. (2011): "Information and Prices with Capacity Constraints," *American Economic Review*, 101, 1591–1600.

Levin, D. and J. L. Smith (1994): "Equilibrium in Auctions with Entry," *The American Economic Review*, 84, 585–599.

Li, D. and N. Schürhoff (2019): "Dealer Networks," *The Journal of Finance*, 74, 91–144.

Li, Y., G. Rocheteau, and P.-O. Weill (2012): "Liquidity and the Threat of Fraudulent Assets," *Journal of Political Economy*, 120, 815–846.

Liu, Y., S. Vogel, and Y. Zhang (2017): "Electronic Trading in OTC Markets vs. Centralized Exchange," Working Paper.

MarketAxess, C. (2022): "MarketAxess Corporation Participant System Guidelines," https://www.marketaxess.com/pdf/Participant-System-Guidelines.pdf.

Maskin, E. and J. Tirole (1988a): "A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs," *Econometrica*, 56, 549–569.

——— (1988b): "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles," *Econometrica*, 56, 571–599.

——— (2001): "Markov Perfect Equilibrium: I. Observable Actions," *Journal of Economic Theory*, 100, 191–219.

Maurin, V. (2022): "Liquidity Fluctuations in Over-the-Counter Markets," *The Journal of Finance*, 77, 1325–1369.

McAfee, R. P. and J. McMillan (1987): "Auctions with Entry," *Economics Letters*, 23, 343–347.

McPartland, K. (2014): "The SEF RFQ Minimum Is Moving to 3. Does It Matter? Nope." https://www.greenwich.com/blog/sef-rfq-minimum-moving-3-does-it-matter-nope.

Menezes, F. M. and P. K. Monteiro (2000): “Auctions with Endogenous Participation,” *Review of Economic Design*, 5, 71–89.

O’Hara, M. and X. Alex Zhou (2021): “The Electronic Evolution of Corporate Bond Dealers,” *Journal of Financial Economics*, 140, 368–390.

O’Hara, M., Y. Wang, and X. Zhou (2018): “The Execution Quality of Corporate Bonds,” *Journal of Financial Economics*, 130, 308–326.

Peltonen, T. A., M. Scheicher, and G. Vuillemey (2014): “The Network Structure of the CDS Market and Its Determinants,” *Journal of Financial Stability*, 13, 118–133.

Pinter, G., C. Wang, and J. Zou (2020): “Information Chasing versus Adverse Selection,” Working Paper.

——— (2021): “Size Discount and Size Penalty: Trading Costs in Bond Markets,” Working Paper.

Praz, R. (2014): “Equilibrium Asset Pricing with Both Liquid and Illiquid Markets,” Working Paper.

Ramadorai, T. (2008): “What Determines Transaction Costs in Foreign Exchange Markets?” *International Journal of Finance and Economics*, 13, 14–25.

Riggs, L., E. Onur, D. Reiffen, and H. Zhu (2020): “Swap Trading after Dodd-Frank: Evidence from Index CDS,” *Journal of Financial Economics*, 137, 857–886.

Sambalaibat, B. (2022): “Endogenous Specialization and Dealer Networks,” Working Paper.

Stahl, II, D. O. (1989): “Oligopolistic Pricing with Sequential Consumer Search,” *American Economic Review*, 79, 700–712.

Stigler, G. J. (1961): “The Economics of Information,” *Journal of Political Economy*, 69, 213–225.

Tsoy, A. (2021): “Over-the-Counter Markets for Non-Standardized Assets,” Working Paper.

Üslü, S. (2019): “Pricing and Liquidity in Decentralized Asset Markets,” *Econometrica*, 87, 2079–2140.

Varian, H. R. (1980): “A Model of Sales,” *The American Economic Review*, 70, 651–659.

VAYANOS, D. AND P.-O. WEILL (2008): "A Search-Based Theory of the On-the-Run Phenomenon," *The Journal of Finance*, 63, 1361–1398.

VOGEL, S. (2019): "When to Introduce Electronic Trading Platforms in Over-the-Counter Markets?" Working Paper.

WANG, C. (2016): "Core-Periphery Trading Networks," Working Paper.

——— (2022): "Search in Over-the-Counter Markets," Working Paper.

WEILL, P.-O. (2020): "The Search Theory of Over-the-Counter Markets," *Annual Review of Economics*, 12, 747–773.

YUESHEN, B. Z. (2017): "Uncertain Market Making," Working Paper.

YUESHEN, B. Z. AND J. ZOU (2022): "Less Is More," Working Paper.