Trading Against Expert Dealers Under Limited Information Spillovers *

Maria Chaderina Vincent Glode

February 2, 2026

Abstract

We model investors' allocation of order flow across over-the-counter dealers jointly with dealers' acquisition of expertise, used to take advantage of investors across transactions. Whereas investors have incentives to flock toward the dealer expected to have the lowest level of expertise, a dealer that expects to attract many investors has incentives to acquire additional expertise. In contrast with standard models, we allow dealers' expertise to exhibit limited spillovers across transactions. As a result, investors prefer dealers that intermediate large volumes of transactions and, in equilibrium, order flow may concentrate around the dealer making the largest investments in expertise.

*We thank Huiyao Chen, Yunzhi Hu, Wei Li, Martin Oehmke, Cecilia Parlatore, Thomas Philippon, Thomas Rivera, Ruslan Sverchkov, and seminar participants at Carnegie Mellon University, Indiana University, University of Washington, the European Winter Finance Conference, the Financial Intermediation Research Society Meeting, the Midwest Finance Association Meeting, the Northern Finance Association Meeting, WAPFIN, and the Pacific Northwest Finance Conference for their helpful comments. Chaderina is at the Lundquist College of Business, University of Oregon, 1208 University Street, Eugene, OR 97403 and can be reached by e-mail at mchaderi@uoregon.edu. Glode is at the Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia, PA 19104 and can be reached by e-mail at vglode@wharton.upenn.edu.

1 Introduction

When a financial institution acts as an over-the-counter (OTC) dealer for investors, it buys and sells securities on its own behalf, using its expertise (e.g., data, technology, and skills) to make decisions that boost the institution's profits (see, e.g., Manaster and Mann 1996, Chae and Wang 2003, van der Wel, Menkveld, and Sarkar 2009).[1] Investors should thus be wary of trading against any dealer possessing superior expertise that could be used to take advantage of them (see, e.g., Akerlof 1970). Nevertheless, the dealers handling the most order flow in OTC markets tend to be highly sophisticated institutions consistently making large investments in expertise. For example, Goldman Sachs spends $\$400M$ per year to acquire financial data from third-party sources and $\$400K$ on average per employee to attract, compensate, and retain the best and brightest, yet it remains the second most active OTC dealer of derivatives among U.S. banks.[2] Unsurprisingly, a large fraction of OTC dealers' profits comes from trading against less informed investors (see Green, Hollifield, and Schürhoff 2007a, 2007b).

In this paper, we jointly model dealers' expertise acquisition and investors' order-flow allocation in OTC markets. Dealers acquire costly expertise to improve their ability to value the assets they trade with investors, and thereby appropriate a larger fraction of the surplus from trade. Investors, however, allocate their order flow across dealers, taking into account each dealer's informational advantage. Our model features a paradox regarding dealers' expertise levels: whereas investors have incentives to flock toward a dealer expected to have a low level of expertise, a dealer that expects to attract many investors has incentives to acquire additional expertise. As a

[1]This role contrasts with that of a "broker," which executes orders on behalf of its clients and earns a flat commission for the service.

[2]See OCC (2021), Campbell (2018), and Goldman Sachs' 2021-Q4 earnings report for the specific numbers in the sentence. For additional evidence of concentrated OTC intermediation for different types of assets, see, e.g., Di Maggio, Kermani, and Song (2017), Munyan and Watugala (2018), Goldstein and Hotchkiss (2020), and Hendershott et al. (2020) for corporate bonds, Li and Schürhoff (2019) for municipal bonds, Atkeson, Eisfeldt, and Weill (2014) and Siriwardane (2019) for credit default swaps, Cetorelli et al. (2007) for equity derivatives, Begenau, Piazzesi, and Schneider (2015) for interest-rate derivatives, Joseph and Vasios (2022) for currency derivatives, Hollifield, Neklyudov, and Spatt (2017) for securitized products, and Hagströmer and Menkveld (2019) and Hasbrouck and Levich (2021) for currencies.

result, how dealers' expertise acquisition and investors' order-flow allocation interact to reach an equilibrium depends on market conditions.

Indeed, an important and novel feature of our analysis consists of how we model dealers' expertise to flexibly capture the degree of information spillovers across all the transactions each dealer intermediates. In most models of information acquisition, a trader's expertise level dictates the informational *quality* of the signals this trader observes for all transactions under consideration. While this assumption might be realistic for highly specialized dealers or commodity markets, in many other settings a dealer's acquisition of superior information about a specific security, deal, or counterparty is likely to have limited pricing implications for other securities, deals, and counterparties. For example, Li and Schürhoff (2019) identify 1.5 million separate municipal bonds issued by 50,000 different U.S. entities over the 1998-2012 period and after issuance, the overwhelming majority of municipal bonds are traded less than once a month, according to Green, Hollifield, and Schürhoff (2007b). Similarly, the majority of single-entity credit default swaps and corporate bonds are rarely traded (see, e.g., Chen et al. 2011, Chaderina, Muermann, and Scheuch 2025), with differences in contract terms rendering even securities linked to the same entity imperfect substitutes. Thus, if a dealer were to gain an informational advantage for a specific transaction, it does not mean that this dealer would also benefit from the same advantages for all its other transactions.

In our model, we parameterize how a dealer's expertise level impacts the *quantity* of transactions for which this dealer will have the attention and resources required to gain an informational advantage over its counterparties. By hiring smart traders, purchasing fast computers, and acquiring proprietary databases, a dealer ends up increasing its capacity or bandwidth to take advantage of counterparties across multiple transactions involving different assets within a short time period. When investors are interested in trading various assets (e.g., corporate bonds, municipal bonds, or derivative products linked to different entities), a dealer's resources must be split to assess the terms of trade of all proposed transactions. Once the dealer allocates some scarce resources to value an interest-rate futures contract for example, those same resources may no longer be available to value a currency swap as part of a different transaction. As more order flow gets directed toward this dealer, its resources are spread out more thinly across all proposed transactions, thereby weakening this dealer's ability to gain an informational advantage and assess the terms of trade associated with each transaction. These "liquidity externalities" are reminiscent of models of rational inattention in which an agent's attention budget must be spread across a set of information and decisions (see, e.g., Sims 2003, Kacperczyk, Van Nieuwerburgh, and Veldkamp 2016, Maćkowiak, Matějka, and Wiederholt 2023). They are also consistent with empirical findings by Corwin and Coughenour (2008) and Chakrabarty and Moulton (2012) suggesting that dealers have limited attention and resources they can dedicate to assessing their various trades. These liquidity externalities are the first of two key forces in our analysis.

As alluded to above, our analysis shows that market conditions significantly impact the equilibrium trading network and the associated levels of dealer expertise, order-flow concentration, and trading liquidity. When the cost of expertise is low, the degree of information spillovers is low, and the volume of order flow is large, investors concentrate their trading around the one dealer that acquires the highest expertise capacity in equilibrium. This "central" dealer invests in expertise until investors are indifferent between trading with this dealer whose expertise capacity is thinly spread among a large number of transactions (resulting in high liquidity) and trading with less sophisticated dealers whose scarcer resources can be used to gain informational advantages in a small number of transactions (resulting in low liquidity). Thus, paradoxically, in equilibrium order flow is concentrated toward a dealer that employs smarter traders, owns faster computers, and has better data. This outcome arises because dealers' expertise investments are endogenous responses to investors' expected allocations of order flow. A dealer expecting to intermediate a high volume of transactions cannot credibly commit to not acquiring the expertise needed to take advantage of its unsophisticated counterparties. An outcome where most order flow is concentrated toward one unsophisticated intermediary cannot be sustained in equilibrium as the large number of investors expected to do business with this central intermediary would render expertise acquisition too profitable for this intermediary. This endogeneity of dealers' expertise investments is the second of two key forces in our analysis.

While order-flow concentration toward one dealer providing liquid trading is a salient prediction in the parametric region just described and is consistent with empirical findings from Hollifield, Neklyudov, and Spatt (2017), Li and Schürhoff (2019), and Joseph and Vasios (2022), alternative trading networks can arise in the other parametric regions of our model. In parametric regions where the liquidity externalities are weak, multiple equilibria can arise, including ones for which order flow is dispersed across several dealers providing illiquid trading. Our model's comparative statics can thus shed light on the evolution of OTC trading networks. Specifically, if over time expertise were to get cheaper, its spillovers across more diversified asset bases were to weaken, and/or the number of "independent" investors (who optimally pick their dealers) were to increase, we would be more likely to observe in equilibrium that independent investors concentrate their order flow toward one dealer making large investments in expertise while also providing liquid trading. Additionally, if the growth in independent investors were to come from investors who were previously loyal to their local dealers, it would imply that these dealers' remaining loyal investors would obtain lower liquidity as a result of these changes.

Our paper contributes to the literature that studies the allocation of order flow in OTC markets. In contrast with theoretical predictions from the large literature that assumes random matching in OTC markets, Hagströmer and Menkveld (2019) and Hendershott et al. (2020) empirically show that order-flow matching patterns are highly persistent, thereby highlighting the need to understand how traders select their counterparties. Green (2007) theoretically analyzes dealers' incentives to take advantage of investors who cannot compare terms of trade across dealers without incurring search costs. Unlike us, Green (2007) does not study the endogenous acquisition of expertise by dealers in response to the expected volume of transactions or the liquidity externalities of order-flow concentration in light of adverse selection concerns. Chacko, Jurek, and Stafford (2008), Lester, Rocheteau, and Weill (2015), Neklyudov (2019), and Sambalaibat (2025) highlight the execution-speed implications of order-flow concentration without modeling dealers' optimal response in terms of expertise acquisition. Pagano (1989) and Chaderina and Green (2014) endogenize market participation in light of liquidity externalities, but do not consider the role played by dealers' endogenous expertise.

Our paper shows how order-flow concentration can endogenously reduce trading inefficiencies due to dealers' superior information, thereby rationalizing why Hollifield, Neklyudov, and Spatt (2017), and Joseph and Vasios (2022) find that, for various OTC markets, a core of dealers intermediates most transactions and the transactions intermediated by these central dealers commands better terms of trade than the transactions intermediated by more "peripheral" dealers. While Li and Schürhoff (2019) show that, for the average municipal bond transaction, peripheral dealers offer better terms of trade, this advantage disappears once they isolate the component of pricing attributable to asymmetric information by focusing on marginal transactions. In that setting, central dealers offer more attractive terms of trade than peripheral dealers, consistent with the predictions of our model. Baumann et al. (2023) find that following corporate default, bond investors concentrate their order flow toward a few select dealers with superior expertise, and receive better terms of trade than what is offered by other dealers. These findings are consistent with our model's insight that, despite higher expertise investments, a central dealer can offer better terms of trade than competing peripheral dealers.

A few recent papers have rationalized order-flow concentration toward intermediaries that possess socially valuable information. Babus and Hu (2017) argue that trading through a central dealer provides this dealer with information that is useful in policing counterparties and disciplining them in case of misbehavior. Li and Song (2024) jointly model dealers' expertise acquisition and investors' order-flow allocation in OTC markets when a dealer's superior information is directly shared with its customers, who can then make better portfolio decisions thanks to the expertise acquired by their chosen dealers. They show that concentration toward one dealer is optimal when dealers effectively act as brokers or advisors. Bethune, Sultanum, and Trachter (2022) show that if central dealers know more about the private valuations of their clients (e.g., their liquidity needs or inventory levels), concentrating order flow toward them improves trade efficiency. Relatedly, Chang and Zhang (2021) show how trading concentration arises in equilibrium, thereby enhancing trade efficiency, if traders are uncertain about their counterparties' trading needs when they must form their trading network. In light of these insights, one might presume that, if endowing a dealer with surplus-creating information increases the order flow allocated to this dealer as shown in these papers, endowing a dealer with surplus-appropriating information might repel traders from sending their order flow to such dealer. In many financial markets, superior information about the common value of traded securities, such as their future cash flows or default probabilities, is likely to be what most dealers aim to acquire. The resulting asymmetric information gives rise to adverse selection and is therefore welfare destroying. In this paper, we focus on how the acquisition of superior information about the common value of financial securities and the adverse selection it creates affect order-flow allocations. We show that when dealers' informational advantages are used against their investors but have limited spillovers across transactions/securities, pooling order flow around a central dealer is optimal for investors, as it protects them from adverse selection. At a broader level, our paper contributes to the large literature shedding light on recent trends in U.S. industrial concentration (see, e.g., Covarrubias, Gutiérrez, and Philippon 2019, and the references therein) by highlighting a novel channel specific to the financial sector.

Finally, our paper contributes to the literature that studies information acquisition in OTC markets. Glode, Green, and Lowery (2012) highlight OTC traders' incentives to overinvest in expertise prior to their trading interactions in order to take advantage of their counterparties, but assume an exogenous order-flow allocation. Glode and Opp (2020) show how predictable trading interactions can incentivize costly information acquisition by OTC traders. Galeotti and Goyal (2010) and Herskovic and Ramos (2020) study network formation games where information is shared within connections, leading to complementarities and concentration. Boyarchenko, Lucca, and Veldkamp

(2021) study information sharing among OTC dealers who learn from taking their customers' orders. All these papers do not model limited information spillovers and are therefore silent about the associated liquidity externalities of order-flow concentration in light of adverse selection concerns, which are the focus of our analysis.

2 Model

Our model has two stages. In the first stage, L dealers acquire costly expertise while N investors choose the dealers with whom they will trade. In the second stage, trade takes place bilaterally among each investor and its chosen dealer. Below, we will first introduce the trading game and derive agents' optimal trading behavior in the second stage. Using these results, we will then solve for the optimal amounts of expertise that dealers acquire and the optimal allocations of order flow that investors choose in the first stage.

2.1 Trading Stage

Since the focus of our paper is on how dealers and investors behave in the first stage, we keep our model of the second stage purposefully simple. While the model makes stylized assumptions about how trading occurs among agents in the second stage, it will become clear when analyzing the first stage that our paper's main insights only rely on the intuitive and robust property that a trader's informational advantage increases its trading profit at the expense of its counterparty's trading profit (subject to adverse selection limitations).

Consider a trading game between the current owner of a financial asset and a prospective buyer. The current owner values the asset for its common value v, whose realization can either be v_l or v_h ($> v_l$) with equal probabilities based on public information. The expected value of the asset is then $E(v) \equiv \frac{v_h+v_l}{2}$. The prospective buyer reaches out to the current owner of the asset (a.k.a., the seller) because, in addition to its common value v, the buyer would collect a private benefit

$b > 0$ from acquiring and holding the asset (e.g., caused by unmodeled diversification, investment horizon, or liquidity benefits). The existence of gains to trade b is public knowledge and implies that trading the asset would improve welfare.

Before a transaction occurs, the seller receives a private signal $s \in \{v_l, v_h\}$ about the value of the asset and this signal is accurate with probability $\alpha = \frac{1}{2} + a$. We use $a \in \left[0, \frac{1}{2}\right]$ to denote the seller's informational advantage over the buyer about the value of this specific asset. If $a = 0$, the seller's private signal is uninformative about the value of the asset, whereas if $a = \frac{1}{2}$, the seller learns perfectly the value of the asset. We assume that, during the trading stage, both the seller and the buyer know the level of a. While the seller's informational advantage a is taken as given in this stage, we will later formalize dealers' incentives to invest in expertise (e.g., technology, human capital, and data) and boost their a, based on how they expect investors' order flow to be allocated.

To avoid signaling and equilibrium multiplicity concerns, we assume that the prospective buyer makes a take-it-or-leave-it offer to purchase the asset from the (privately informed) seller at a price P. When deciding which price P to offer, the buyer faces an intuitive tradeoff. Offering a higher price means that the buyer is more likely to get the asset and realize the gains to trade b. However, offering a higher price also means that, conditional on getting the asset, the buyer shares more of its surplus with the seller.

Specifically, the buyer considers offering a price:

$$P_h = \alpha v_h + (1 - \alpha) v_l = E(v) + a(v_h - v_l), \tag{1}$$

which is equal to how much the seller would value the asset after observing a signal $s = v_h$. If this price is offered, the seller accepts to trade with the buyer regardless of whether the signal is $s = v_h$ or $s = v_h$, which is the socially optimal trading outcome as the gains to trade b are realized with probability one.

The buyer alternatively considers offering a price:

$$P_l = \alpha v_l + (1-\alpha)v_h = E(v) - a(v_h - v_l), \tag{2}$$

which is equal to how much the seller would value the asset after observing a signal $s = v_l$. The seller accepts this price after observing a signal $s = v_l$ but not after observing a signal $s = v_h$, which means that the gains to trade b are destroyed half of the time.

Overall, the buyer's expected surplus from offering a price $P \in \{P_l, P_h\}$ is:

$$\begin{aligned}\Pi(\alpha, b, P) &\equiv \begin{cases} \left[\frac{1}{2}\alpha + \frac{1}{2}(1-\alpha)\right](v_h + b - P_h) + \left[\frac{1}{2}(1-\alpha) + \frac{1}{2}\alpha\right](v_l + b - P_h) & \text{if } P = P_h, \\ \frac{1}{2}(1-\alpha)(v_h + b - P_l) + \frac{1}{2}\alpha(v_l + b - P_l) & \text{if } P = P_l, \end{cases} \\ &= \begin{cases} b - a(v_h - v_l) & \text{if } P = P_h, \\ \frac{b}{2} & \text{if } P = P_l. \end{cases}\end{aligned} \tag{3}$$

A buyer never finds it optimal to offer either $P > P_h$ (which is dominated by offering $P = P_h$) or $P < P_l$ (which is dominated by offering $P = P_l$), nor does the buyer ever offer $P \in (P_l, P_h)$ (which is dominated by offering P_l). We now characterize the buyer's optimal bidding strategy when facing a seller with an informational advantage a.

Lemma 1. *In equilibrium, the buyer offers the price:*

$$P^* = \begin{cases} P_h = E(v) + a(v_h - v_l) & \textit{if } a \leq \bar{a} \equiv \frac{b}{2(v_h - v_l)}, \\ P_l = E(v) - a(v_h - v_l) & \textit{if } a > \bar{a} \equiv \frac{b}{2(v_h - v_l)}. \end{cases} \tag{4}$$

Throughout the paper, proofs of our formal results are relegated to the Appendix. As Lemma 1 shows, the solution for P^* allows for two cases that differ in how adverse selection impacts the liquidity of trade. For low levels of $a \leq \bar{a}$, adverse selection concerns are mild and, as a result, the buyer finds it optimal to make a high offer P_h that the seller accepts with probability one. In this case, liquidity is high and the full surplus from trade is split between the buyer and the seller through the optimal bid P^*, which is increasing in the seller's informational advantage a. However, for high levels of $a > \bar{a}$, the seller's informational advantage becomes too concerning for the buyer who then prefers to offer the low price P_l, making the optimal bid P^* decreasing in the seller's informational advantage a. In this case, liquidity is low as half of the surplus from trade is destroyed due to adverse selection concerns. Altogether, the seller benefits from a higher informational advantage a as long as the buyer is willing to offer P_h, but having too big of an informational advantage can impede trade and result in the seller being worse off. In addition, the buyer's incentives to offer a high price that sustains liquid trading are increasing in the private benefit b of holding the asset. These comparative statics are illustrated in Figure 1, which plots the buyer's optimal bid as a function of the seller's informational advantage a for two levels of private benefit b.

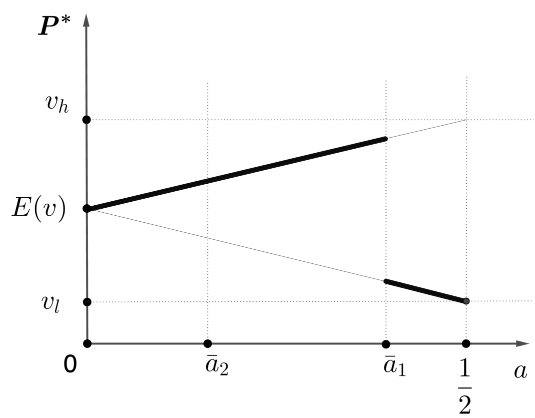

(a) Gains to trade b_1

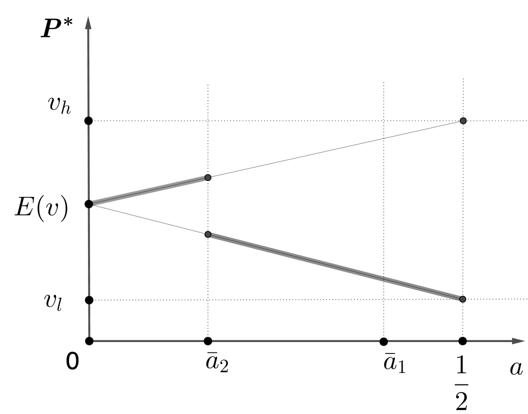

(b) Gains to trade b_2

Figure 1

Buyer's Optimal Bidding Function. This graph plots the buyer's optimal price offer P^* as a function of the seller's informational advantage a, as formalized in Proposition 1, for two levels of private benefit from trade b_1 and b_2, where $b_1 > b_2$. Each panel highlights how the buyer's optimal price offer P^* drops when a crosses either $\bar{a}_1 \equiv \frac{b_1}{2(v_h - v_l)}$ in panel (a) or $\bar{a}_2 \equiv \frac{b_2}{2(v_h - v_l)}$ in panel (b).

The next corollary establishes that the buyer is made better off by being paired with a seller with a smaller informational advantage a.

Corollary 1. *If* $a \leq \bar{a} \equiv \frac{b}{2(v_h - v_l)}$, *the buyer's expected surplus* Π *is decreasing in the seller's informational advantage* a. *If* $a > \bar{a} \equiv \frac{b}{2(v_h - v_l)}$, *the buyer's expected surplus* Π *is unaffected by the seller's informational advantage* a.

For the remainder of the baseline analysis, we impose a parametric restriction on b that results in adverse selection impeding trade for high enough levels of a:

Assumption 1. $0 \leq b < v_h - v_l$.

Imposing $b < v_h - v_l$ guarantees that the liquid-trading threshold $\bar{a}$ satisfies $\bar{a} < \frac{1}{2}$ and that there exists a region $\bar{a} < a \leq \frac{1}{2}$ where $P^* = P_l$ is the optimal bidding strategy and trade is therefore illiquid. This restriction will help capture, in the analysis of the first stage, the notion that dealers must limit how much expertise they acquire in anticipation of investors' adverse selection concerns. We will relax this assumption in Section 4.

Discussion: Robustness of Trading Game. When endogenizing dealers' expertise acquisition and investors' order-flow allocation in the next subsection, we will use the trading profit functions derived above. While the analysis of the trading stage above focused on the seller being privately informed, the payoff functions we derived would not change if we instead assumed that an uninformed seller was quoting an ultimatum price to a buyer with an informational advantage a.[3] Thus, we can interpret the buyer's expected profit Π derived above more broadly as the payoff an unsophisticated investor collects from trading against a better informed dealer, regardless of who is long and who is short the asset. Moreover, using Grossman and Perry's (1986) equilibrium refinement, Glode, Green, and Lowery (2012) derive qualitatively consistent trading outcomes in an analogous trading game where it is the privately informed trader that makes the take-it-or-leave-it offer (thereby triggering a signaling game). As will be clear in the next subsection, our paper's

[3]Using mechanism design, Samuelson (1984) shows that an ultimatum (i.e., take-it-or-leave-it) offer can be the optimal incentive-compatible mechanism for an uninformed trader facing a privately informed counterparty.

main insights would survive in alternative trading games as long as in equilibrium a trader's informational advantage increases its trading profit at the expense of its counterparty's trading profit (with adverse selection potentially impeding trade).

2.2 Expertise Acquisition and Order-Flow Allocation Stage

Using the trading payoff functions derived above, we can now study how dealers choose the level of expertise to acquire and how investors choose the dealers with whom they trade. To capture the idea that dealers' investments in expertise are made in anticipation of investors' order-flow allocation and investors' order flow is allocated in anticipation of dealers' expertise levels, we solve for a Nash equilibrium when investors and dealers move simultaneously.

Each (ex-ante homogenous) dealer is assumed to start with a loyal investor base of measure $\underline{n} > 0$. These are unsophisticated clients who always trade with the same (e.g., local) dealer, in the spirit of Green (2007). The economy is also populated by N investors who are not loyal to a specific dealer and who optimize based on the expected terms of trade, captured by a_j for a given dealer j. We label these investors as independent. A dealer j's total order flow is normalized to be equal to the number of investors it transacts with: $n_j \equiv \underline{n} + n_j^*$, combining the transaction counts from its loyal investors and from the independent investors who choose this particular dealer. Since the derived trading profit functions are independent of the trade direction in our stylized model without inventory concerns, we do not need to keep track of whether investors want to buy or sell their asset. When choosing a dealer in the first stage, each independent investor must solely conjecture the informational advantage a_j that will dictate the terms of trade offered by each dealer j in the second stage.

As briefly explained in the introduction, our model is designed to capture the notion that a dealer's investments in expertise improve its capacity to take advantage, within a short period of time, of various counterparties across multiple transactions involving different securities. As highlighted by the empirical findings in Corwin and Coughenour (2008) and Chakrabarty and Moulton

(2012), a dealer cannot value all existing securities and assess the terms of trade of all proposed deals simply by hiring one trader, purchasing one computer, or acquiring one proprietary database. Instead, the dealer needs to spread its resources across the various transactions it performs, in a way reminiscent of Sims' (2003) attention allocation constraint. The more resources the dealer invests to boost its expertise, the higher is that dealer's bandwidth for assessing the profitability of a variety of potential transactions. To depict in a tractable manner this idea that more scalability requires additional investments, we assume that the market is fragmented in the sense that it is populated by multiple investors each interested in trading a slightly different asset (e.g., a specific corporate bond, municipal bond, or derivative product).[4] Each dealer starts with $\underline{k}$ units of expertise capacity/bandwidth available to value assets and gain an informational advantage over investors. This initial level of expertise capacity can be thought of as originating from the various financial activities that dealers participate in, such as investment banking and sell-side analysis, that may boost dealers' informational advantages when trading but are not motivated by dealers' trading profits per se.[5] Each dealer j can then expand its expertise capacity to $k_j > \underline{k}$ at a cost $c(k_j - \underline{k})$, where $c > 0$, if such investment (e.g., in data, human capital, or computing power) is deemed to increase trading profits. A dealer's optimal investment in expertise will depend both on how order flow is expected to be allocated in the first stage and on how trading is expected to take place in the second stage.

To analyze how expertise capacity spreads across transactions, we assume that dealer j's informational advantage in each transaction is given by:

$$a_j = \frac{1}{2} \min\left(\frac{k_j}{n_j^{1-\gamma}}, 1\right), \tag{5}$$

[4]See Chen et al. (2011), Li and Schürhoff (2019), and Chaderina, Muermann, and Scheuch (2025) who document the high levels of asset heterogeneity within the credit default swap market, the municipal bond market, and the corporate bond market, respectively.

[5]See, for example, Chung and Cho (2005) for empirical evidence of expertise-related interactions between sell-side analysis and market making.

where k_j denotes the dealer's total expertise capacity and n_j denotes its total order flow. The parameter $\gamma \in [0, 1]$ allows us to vary the level of information spillovers across all transactions dealers intermediate (e.g., based on whether dealers in a market are specialized or trade highly heterogenous assets). When $\gamma = 1$, there are full spillovers (i.e., the benchmark often assumed in models such as Kyle 1985), implying that a single piece of information or analytical capability applies equally to all transactions, making dealer j's informational advantage in each transaction independent of the number of transactions n_j. When $\gamma = 0$, there are no spillovers, implying that each transaction requires its own separate investment in expertise, making the dealer's informational advantage in each transaction proportional to $\frac{1}{n_j}$. Intermediate values $0 < \gamma < 1$ capture partial spillovers, implying that some elements of expertise (e.g., shared research, inventory analysis) are reusable across transactions, while others are transaction-specific. In all cases, the factor $\frac{1}{2}$ acts solely as a normalization to reflects that the signal's probability of being correct is $\frac{1}{2} + a_j$ and cannot exceed one. When $n_j^{1-\gamma} \leq k_j$, the dealer achieves the upper bound $a_j = \frac{1}{2}$, corresponding to a perfect signal for each transaction. If $n_j^{1-\gamma} > k_j$, the dealer's expertise capacity is spread over more transactions, reducing the accuracy of each signal. Moreover, the functional form for a_j assumes that dealer j ends up spreading its expertise equally across transactions, which would be optimal in a relaxed environment since the dealer would want to avoid crossing the adverse selection threshold $a_j \leq \bar{a}$ in any transaction (as further discussed in Section 4).

For the remainder of the analysis, we impose the following assumption that information spillovers are at least partially limited, thereby ensuring that increasing order flow weakens a dealer's informational advantage:

Assumption 2. $\gamma \in [0, 1)$.

To shed light on a dealer's incentives to invest in expertise, we substitute the optimal bid P^* derived in Proposition 1 into the dealer's profit function. We can then write the dealer's total expected profit as:

$$\Delta\left(k_j, n_j\right) = \begin{cases} n_j a_j (v_h - v_l) - c(k_j - \underline{k}) & \text{if } a_j \leq \bar{a}, \\ -c(k_j - \underline{k}) & \text{if } a_j > \bar{a}. \end{cases} \tag{6}$$

The next lemma shows that when expertise is sufficiently cheap to acquire dealer j is better off increasing k_j until $a_j = \bar{a}$.

Lemma 2. *As long as* $a_j < \bar{a} \equiv \frac{b}{2(v_h - v_l)}$*, we can write* $\frac{\partial \Delta(k_j, n_j)}{\partial k_j} = \frac{1}{2} n_j^{\gamma} (v_h - v_l) - c$*, which is strictly positive if and only if* $c < \frac{1}{2} n_j^{\gamma} (v_h - v_l)$*.*

As documented by Green, Hollifield, and Schürhoff (2007a, 2007b), trading against less informed investors is one of the main sources of OTC dealers' profits. Manaster and Mann (1996), Chae and Wang (2003), and van der Wel, Menkveld, and Sarkar (2009) empirically show that dealers and market makers in various markets use superior information to boost their profits. In our model, a dealer benefits from its expertise in two ways. A first benefit comes from making better decisions when responding to investors' offers. This benefit is most evident when $a_j > \bar{a}$ as investors are then so concerned about adverse selection that they make offers that result in illiquid trading. For a given price offer, the more accurate the signal is, the more profitable is the dealer's decision whether to refuse or execute the transaction. However, the improved-decision benefit is offset by worse prices being offered by investors who grow more concerned about adverse selection. This situation helps capture how trade failures negatively impact liquidity and immediacy in OTC markets, as documented by Hendershott et al. (2024). A second benefit of expertise exists when $a_j \leq \bar{a}$. As the dealer's signal precision increases within that region, investors respond by making more generous offers to ensure that the dealer agrees to trade with probability one.

Overall, there is nothing unethical or illegal in how dealers interact with investors in our model. Dealers are simply using their superior information when bargaining with investors, occasionally rejecting offers expected to yield negative profits (see Hendershott et al. 2024). In equilibrium, investors are not misled or exploited as they accurately anticipate dealers' expertise levels when trading with them. In reality, investors' knowledge of dealers' expertise could originate from a more complicated (dynamic) game than what we currently model. Yet, it is important to emphasize that the central insights we will derive below solely rely on the generic property from our trading game that, ceteris paribus, a dealer's expertise increases the surplus from trade it can appropriate and decreases the surplus from trade its counterparties can retain (subject to adverse selection limitations).

3 Equilibrium Analysis

An equilibrium of the expertise acquisition and order-flow allocation stage is defined by dealers' expertise choices $k_j^* \geq \underline{k}$ and investors' dealer choices $Dealer_i^*$ such that:

- for any investor i that chooses to rout its trade to dealer j, i.e., $Dealer^(*)_i = j$, we have:

$$\Pi\left(\frac{1}{2}+\frac{1}{2}\min\left(\frac{k_j^*}{n_j^{1-\gamma}},1\right),b,P_j^*\right) \geq \Pi\left(\frac{1}{2}+\frac{1}{2}\min\left(\frac{k_{j'}^*}{(n_{j'}+1)^{1-\gamma}},1\right),b,P_{j'}^*\right) \qquad \forall j' \neq j, \tag{7}$$

  where P_j^* and $P_{j'}^*$ denote the optimal price offers from investors when trading with dealer j and dealer $j^{'}$ respectively (derived in Proposition 1);

- for any dealer j that chooses to acquire an expertise capacity k_j^* and expects to receive order flow n_j as a result of investors' dealer choices $Dealer_i^*$, we have:

$$\Delta\left(k_j^*,n_j\right) \geq \Delta\left(k_j',n_j\right) \qquad \forall k_j' \neq k_j^*. \tag{8}$$

As already stated, we model dealers' and investors' decisions in the first stage as a simultaneous-move game. This timeline captures the notion that expertise acquisition and order-flow allocation are endogenous to each other and none of these decisions unilaterally drives the other (as would be the case if these decisions were sequential). While the initial level of expertise as well as the loyal client base of each dealer are well-known to all market participants, simultaneity implies that at this stage investors do not know for sure all dealers' future levels of expertise, whereas dealers do not know for sure all investors' future allocations of order flow. Our equilibrium definition, however, imposes that there is no systematic deviation between conjectured and realized outcomes when agents make their decisions in the first stage. Thereafter, all agents observe the allocation of investors' order flow and dealers' acquisition of expertise, allowing them to enter the trading stage knowing their dealer's transaction-specific informational advantage a_j. (We will discuss the robustness of our main insights to alternative timelines in Section 4.)

To help characterize possible equilibria across parametric regions, we define a threshold η with the following condition:

$$\frac{1}{2}\min\left(\frac{k}{(\underline{n}+\eta+1)^{1-\gamma}},\,1\right)=\bar{a}. \tag{9}$$

When η exists, it represents the volume of independent order flow that makes an additional/deviating transaction sufficient to drive a dealer's information advantage (without expertise investments) to cross the adverse selection cutoff $\bar{a}\equiv\frac{b}{2(v_h-v_l)}$. Under Assumptions 1 and 2, such η always exists and simplifies to:

$$\eta\equiv\left(\frac{k}{2\bar{a}}\right)^{\frac{1}{1-\gamma}}-\underline{n}-1. \tag{10}$$

How the number N of independent investors in the market compares to this threshold η will dictate which possible deviation can rule out various types of conjectured equilibria in our model. But before we can fully characterize the possible equilibria in all parametric regions, we need to establish a result that highlights how concentrating investors' order flow toward one dealer can reduce the harm that dealers' expertise does to investors' trading profits.

Lemma 3. *If at least one dealer has an informational advantage that is weak enough to permit* *liquid trading* $\left(i.e., a_j \leq \bar{a} \equiv \frac{b}{2(v_h - v_l)}\right)$*, then an allocation of independent investors' order flow to more than one dealer cannot sustain an equilibrium.*

Recall from Corollary 1 that an investor's trading profit is decreasing in its dealer's informational advantage a_j as long as $a_j \leq \bar{a}$ and trading thereby remains liquid. Intuitively, Lemma 3 shows that if independent investors were expected to split their transaction volume across more than one dealer, then any independent investor would benefit from a deviation that consists of re-routing order flow to a different dealer. By doing so, a deviating investor would decrease how much expertise the targeted dealer can use for each transaction (i.e., this dealer's a_j). Therefore, it would be in the best interest of every independent investor to allocate its order flow to a single dealer in order to stretch this dealer's expertise capacity as much as possible and minimize the dealer's informational advantage in each transaction.

Lemma 3 emphasizes how "liquidity externalities" can incentivize independent investors to concentrate their order flow toward one dealer. Each investor prefers to trade with a dealer that is attracting a lot of independent investors, because this dealer is left with less expertise capacity to take advantage of this particular investor. In response to the dealer's weaker informational advantage, investors are willing to make more generous offers that result in trading with this dealer being more liquid (i.e., more likely to happen). However, as we show when we formalize the full equilibrium in the six propositions below, dealers take order-flow allocations into account when choosing how much expertise to acquire. Different parametric regions will give rise to different types of deviations to rule out, thereby affecting the possible equilibria in those regions. Once our formal equilibrium characterizations are established, we will combine the various propositions and numerically analyze their comparative statics.

3.1 High Volume of Independent Order Flow

To facilitate the intuition behind our equilibrium characterizations, we first analyze parametric regions where $0 \leq \eta < \frac{N}{L}$. When the volume of independent order flow is sufficiently large (i.e., $N > L \cdot \eta$), order-flow concentration triggers the liquidity externalities highlighted in Lemma 3. If we also assume (for now) that expertise is prohibitively costly, we can shut down dealers' endogenous response to order-flow concentration and establish the following result.

Proposition 1. *When $0 \leq \eta < \frac{N}{L}$ and $c > \frac{1}{2}(\underline{n} + N)^{\gamma}(v_h - v_l)$, in equilibrium one dealer, say j^*, receives order flow $\underline{n} + N$ whereas all other dealers only receive $\underline{n}$. Moreover, no dealer acquires expertise capacity above the initial level $\underline{k}$.*

The condition $c > \frac{1}{2}(\underline{n} + N)^{\gamma}(v_h - v_l)$ guarantees that no dealer, even if it were to receive all independent order flow N, would find it optimal to acquire additional expertise capacity. Thus, all dealers have $k_j = \underline{k}$ in equilibrium. Since $\frac{N}{L} > \eta$, we know that at least one dealer would provide liquid trading if an investor were to unilaterally deviate from the order-flow allocation most likely to sustain illiquid trading among all dealers (i.e., equal allocation across dealers). Due to the liquidity externalities highlighted through Lemma 3, each independent investor prefers to trade with the most popular dealer as this dealer ends up with the weakest informational advantage (i.e., a_j is decreasing in n_j). In equilibrium, we thus have all independent investors concentrating their order flow toward the same (arbitrarily chosen) dealer, say j^*, who offers strictly better terms of trade than all other dealers. Indeed, the condition that $\eta > 0$ implies that the "peripheral" dealers who only receive their loyal investors' order flow have plenty of expertise capacity they could use against any deviating independent investor — a deviation that reallocates an investor's order flow toward a peripheral dealer would therefore result in illiquid trading and make the deviating investor strictly worse off. All investors who trade with the "central" dealer are thus strictly better off than the loyal investors of peripheral dealers. Consistent with this prediction, Hollifield, Neklyudov, and

Spatt (2017), Li and Schürhoff (2019), and Joseph and Vasios (2022) find that, in the securitized product market, the municipal bond market, and the currency derivative market, marginal investors receive better terms of trade from central dealers than from peripheral dealers.

By first assuming that $c > \frac{1}{2}(\underline{n} + N)^{\gamma}(v_h - v_l)$, we effectively neutralized dealers' optimal response to order-flow concentration. From Lemma 2, we know however that assuming $c < \frac{1}{2}(\underline{n} + N)^{\gamma}(v_h - v_l)$ would make it profitable for a central dealer to acquire additional expertise as long as $a_j^* \leq \bar{a}$ in equilibrium. We can then establish the following result.

Proposition 2. *When $0 \leq \eta < \frac{N}{L}$ and $c \leq \frac{1}{2}(\underline{n} + N)^{\gamma}(v_h - v_l)$, in equilibrium one dealer, say j^*, receives order flow $\underline{n} + N$ whereas all other dealers only receive $\underline{n}$. Moreover, dealer j^* acquires additional expertise capacity $(k_{j^*} - \underline{k}) > 0$ such that in equilibrium $a_{j^*} = \bar{a}$, whereas all other dealers do not acquire expertise capacity above the initial level $\underline{k}$.*

Proposition 2 establishes that, when the cost of expertise is below the previously-stated threshold, in equilibrium all independent investors concentrate their order flow toward one central dealer acquiring the most expertise among all dealers. While its investors would prefer this dealer to have weaker informational advantages when trading, alternative dealers who have less expertise capacity than the central dealer also happen to participate in fewer transactions. As was the case in Proposition 1, the condition that $\eta \geq 0$ implies that a deviating investor who reallocates its order flow away from the central dealer would end up participating in illiquid trading. Thus, the central dealer is able to acquire expertise capacity that increases its trading profit up without repelling independent order flow. The liquidity externalities that exist when $\gamma \in [0, 1)$ allow the most popular dealer j^* to acquire additional expertise to be profitably used against its investors and, since this dealer provides liquid trading (by targeting $a_{j^*} = \bar{a}$), its expertise acquisition does not hinder its popularity. Investors recognize that trade is more liquid with the central dealer than with peripheral dealers who possess too much expertise for the limited order flow they receive. In equilibrium, investors trading with the central dealer have their offers accepted with probability one, implying maximal levels of liquidity and surplus from trade, whereas peripheral dealers accept to trade with their investors only half of the time, implying inefficiently low levels of liquidity and surplus from trade. Independent investors are thus indifferent between doing business with the expert dealer j^* that attracts all independent investors' order flow and doing business with any of the other dealers whose levels of expertise capacity and order flow are limited — in both cases, investors collect an expected profit of $\frac{b}{2}$, as shown in subsection 2.1.

This indifference contrasts with the equilibrium outcome formalized in Proposition 1, where the $(\underline{n} + N)$ investors who trade with the central dealer were strictly better off than the $(L-1)\underline{n}$ investors who are loyal to their (peripheral) dealers. Yet, order-flow concentration still allows for more liquid trading than what peripheral dealers offer, consistent with findings by Joseph and Vasios (2022) for the currency derivative market. The central dealer captures a large fraction of investors' surplus from trade, consistent with empirical findings by Munyan and Watugala (2018) and Hendershott et al. (2020) for the corporate and municipal bond markets.[6] In fact, Munyan and Watugala (2018) find that more central dealers earn higher, more volatile profits across their transactions than peripheral dealers, which is consistent with our prediction that the central dealer is the only dealer facing uncertain per-transaction profit as all peripheral dealers make zero profit for each transaction.

Overall, we can interpret the equilibrium outcomes from Propositions 1 and 2 through the lenses of Green (2007). In equilibrium, a sophisticated dealer (i.e., with weakly higher expertise capacity) attracts the transaction volume of attentive (i.e., independent) investors while multiple less sophisticated dealers benefit from the inattentiveness of their loyal investor clienteles. The central dealer provides liquid trading and earns a positive expected profit, whereas peripheral dealers provide illiquid trading and earn no profit. Yet, for sufficiently low c, investors are indifferent

[6]Our model cannot generate heterogeneous terms of trade across investors trading with the same dealer, as documented by Hendershott et al. (2020). This limitation is due to the fact that, once order flow is allocated across dealers, all investors choose what they offer to a given dealer j solely as a function of this dealer's equilibrium level a_j. To resolve this issue, we could complicate the model by assuming that investors are heterogeneous not only in their degree of loyalty to a specific dealer, but also in their bargaining power, their information, or their liquidity needs.

about which dealer to trade with in equilibrium because all dealers end up with $a_j \geq \bar{a}$.

3.2 Low Volume of Independent Order Flow

We now analyze parametric regions where $\frac{N}{L} \leq \eta$. When the volume of independent order flow is smaller than previously assumed (i.e., $N \leq L \cdot \eta$), liquidity externalities do not guarantee that at least one dealer provides liquid trading. As a result, the sufficient condition for Lemma 3 to hold might be violated and equilibria with dispersed order-flow allocation can arise. We start by considering a case where the volume of independent order flow is moderate, such that the condition $\frac{N}{L} \leq \eta \leq N - 1$ holds.

Proposition 3. *When $\frac{N}{L} \leq \eta \leq N - 1$, multiple types of equilibria exist:*

1. *concentrated order-flow allocation consistent with either Proposition 1 or Proposition 2 (depending on c);*
2. *any dispersed allocation is an equilibrium as long as $n_j < \eta + \underline{n}$ for all dealers (i.e., $a_j > \bar{a}$).*

The condition that $\eta \leq N - 1$ implies that a concentrated allocation of independent order flow would satisfy the requirement for Lemma 3 to hold. In this case, the liquidity externalities when a dealer j^* receives all independent order flow N would be strong enough to sustain an equilibrium. The condition that $\frac{N}{L} \leq \eta$, however, also admits the possibility that a dispersed allocation of independent order flow where each dealer receives between $\frac{N}{L}$ and η of independent order flow would result in all dealers providing illiquid trading in equilibrium. In that case, liquidity externalities do not bite locally as no unilateral deviation by an independent investor is sufficient to drive any dealer's informational advantage to $a_j < \bar{a}$.

These two types of equilibria both exist when $\frac{N}{L} \leq \eta \leq N - 1$, yet they greatly differ in terms of their welfare implications. The concentrated equilibrium allows all independent investors to obtain liquid trading with the central dealer. The full trade surplus is thus preserved and distributed among agents. The dispersed equilibrium prevents all investors from obtaining liquid trading. The surplus from trade is therefore destroyed whenever trade breaks down, which happens with probability $1/2$. The remaining surplus is all captured by investors.

The alternative case where $N - 1 < \eta$ rules out the possibility that concentrated order flow would produce sufficient liquidity externalities that sustain an equilibrium with liquid trading. Thus, we have the following result.

Proposition 4. *When $\frac{N}{L} \leq \eta$ and $N - 1 < \eta$, any allocation of independent investors can be sustained as an equilibrium as long as $k_j = \underline{k}$ and $a_j > \bar{a}$ for all dealers.*

In this case, all dealers provide illiquid trading to their investors despite not acquiring additional expertise capacity, irrespective of how order flow is allocated among dealers.

3.3 High Volume of Loyal Order Flow

The only parametric regions left to analyze have $\eta < 0$. In all of our previous propositions, the condition $\eta \geq 0$ ensured that peripheral dealers' informational advantage was too high (i.e., $a_{-j^*} > \bar{a}$) to convince an independent investor to unilaterally reallocate its order flow away from the central dealer, even if this central dealer was an expert whose $a_{j^*} = \bar{a}$. The condition $\eta < 0$ instead implies that peripherals dealers have enough loyal investors to be able to promise liquid trading to any independent investor, as long as these dealers do not acquire too much additional expertise capacity. We can thus establish the following result.

Proposition 5. *When $\eta < 0$ and $c > \frac{1}{2}(\underline{n}+N)^{\gamma}(v_h - v_l)$, in equilibrium one dealer, say j^*, receives order flow $\underline{n} + N$ whereas all other dealers only receive $\underline{n}$. Moreover, no dealer acquires expertise capacity above the initial level $\underline{k}$.*

As earlier, by setting $c > \frac{1}{2}(\underline{n} + N)^{\gamma}(v_h - v_l)$, we have shut down the channel of endogenous expertise acquisition, thereby ensuring that the liquidity externality channel drives the equilibrium construction. As before, by concentrating their order flow toward one dealer, independent investors are reducing this dealer's informational advantage a_{j^*} and improving their terms of trade, relative to what they would get with peripheral dealers. This equilibrium outcome thus resembles that formalized in Proposition 1, except that all dealers are now providing liquid trading to their investors.

When $c \leq \frac{1}{2}(\underline{n} + N)^{\gamma}(v_h - v_l)$, equilibrium analysis is complicated by the contrasting preferences of investors and dealers for dealers' levels of expertise.

Proposition 6. *When $\eta < 0$ and $c \leq \frac{1}{2}(\underline{n} + N)^{\gamma}(v_h - v_l)$, no pure-strategy equilibrium exists.*

In this last parametric region, there does not exist an equilibrium where dealers' expertise acquisition and investors' order-flow allocation are pure strategies. The reason behind this non-existence is that agents' best responses are cycling: high expertise acquisition by a specific dealer promotes low order-flow allocation to this dealer, which promotes low expertise acquisition, which promotes high order-flow allocation, which promotes high expertise acquisition and so on and so forth. While the analytical tractability that allowed us to fully characterize the possible equilibrium outcomes in Propositions 1-5 is now missing, the proof of Proposition 6 highlights that the economic forces at play in this region are the same ones we analyzed previously.

3.4 Numerical Analysis of Various Parametric Regions

Now that we have formalized all the possible equilibrium cases across parametric regions, we numerically illustrate how the predicted equilibrium outcomes are impacted by a few important parameters. Specifically, we investigate the consequences of varying the number of independent investors N, the cost of expertise c, and the degree of information spillovers γ. We keep other parameters fixed, making sure that the conditions for no pure-strategy equilibrium to exist (i.e., Proposition 6) are not satisfied by any of our parameterizations.

Table 1
Fixed Model Parameters

Parameter	Symbol	Value
Valuation spread	$v_h - v_l$	1
Surplus from trade	b	0.5
Number of dealers	L	5
Dealers' loyal investors	$\underline{n}$	15
Dealers' initial expertise	$\underline{k}$	10

In Figure 2, we set $\gamma = 0.4$ and split the equilibrium regions based on the number of independent investors N and the cost of expertise c. The figure shows that, as expertise gets cheaper, order-flow concentration by investors is more likely to happen around a dealer that makes investments boosting its expertise capacity. The figure also shows that, as more independent investors participate in financial markets, investors' order flow is more likely to be concentrated toward a central dealer and this central dealer is more likely to make investments boosting its expertise capacity.

In Figure 3, we set $c = 5$ and split the equilibrium regions based on the number of independent investors N and the degree of information spillovers γ. Recall from Propositions 1 and 2 that the central dealer's optimal expertise acquisition strategy depends on how the cost of expertise compares to a measure that is increasing in the total volume of order flow available to a central dealer. More precisely, if $c > \frac{1}{2}(v_h - v_l)\,(\underline{n} + N)^\gamma$, the central dealer does not invest beyond its initial level of expertise. But as γ rises, the term $(\underline{n} + N)^\gamma$ increases, which raises the threshold $\frac{1}{2}(v_h - v_l)(\underline{n} + N)^\gamma$. Consequently, it becomes more likely that $c \leq \frac{1}{2}(v_h - v_l)(\underline{n} + N)^\gamma$. In that case, order-flow concentration is more likely to feature a central dealer that acquires additional expertise up to the highest level that does not induce a breakdown of trade in equilibrium, that is, $a^* = \bar{a}$. In other words, if information spillovers weaken, a central dealer benefits less, across all transactions, from acquiring an additional unit of expertise. The benefit of acquiring additional expertise capacity is thus reduced for this central dealer.

Overall, these comparative statics shed light about the different levels of concentration in OTC

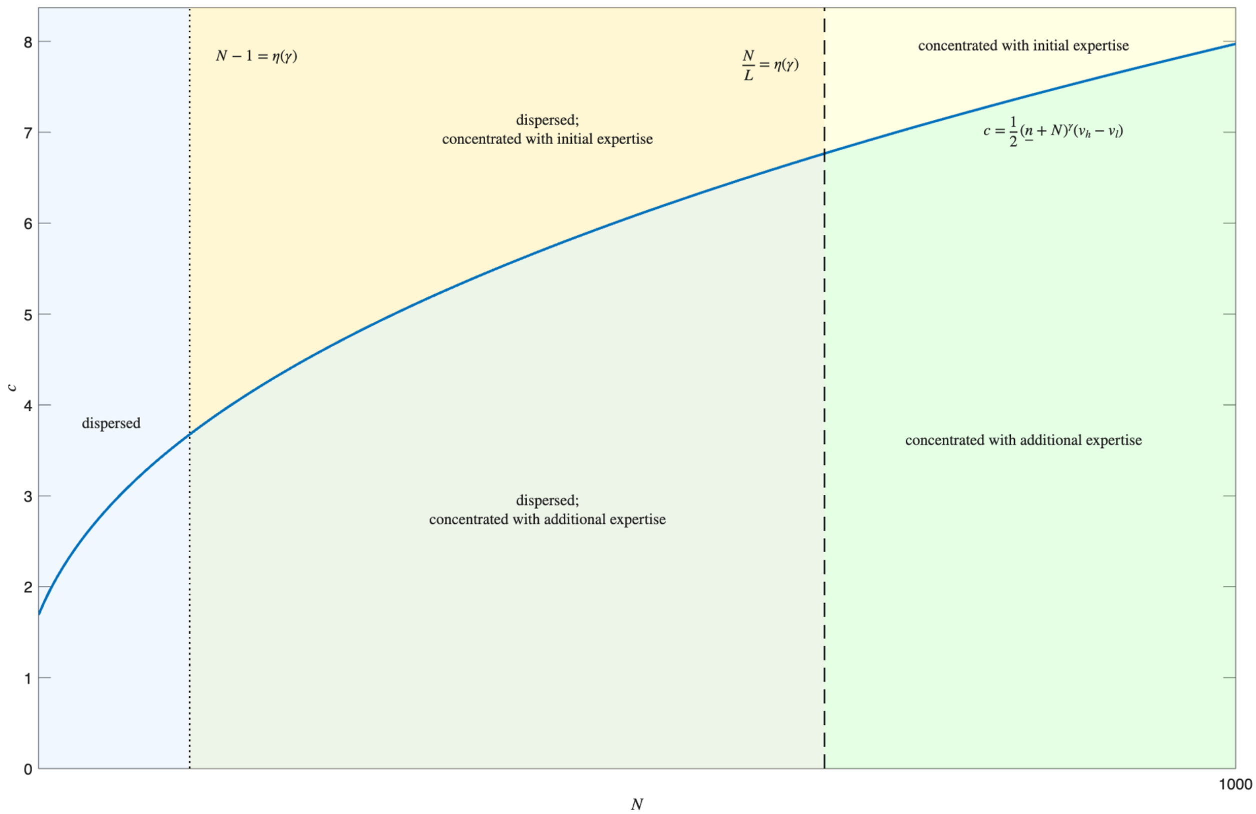

Figure 2

Equilibrium Regions For Various N and c. This graph plots the types of equilibria (characterized by Propositions 1-6) that arise when we allow the number of independent investors N and the cost of expertise c to vary, assuming that $\gamma = 0.4$ and that other parameters are held fixed according to Table 1.

markets. For example, equity-option trading concentration across dealers is significantly higher for Latin American firms than for American, Asian, and European firms, according to a 2025 BIS report. Our model suggests that this higher concentration could be due to Latin American firms being associated with weaker information spillovers (i.e., stronger liquidity externalities), higher levels of asymmetric information, or higher fractions of independent investors. Since our model takes the assets traded by investors and dealers as given, it cannot speak to the geographical specialization of dealers in the municipal bond market, as documented by Jotikasthira, Lundblad, and Xue (2025). However, the liquidity externalities generated by limited information spillovers in our model can explain why Jotikasthira, Lundblad, and Xue (2025) find that more specialized dealers provide their investors with worse terms of trade: dealers who trade more similar assets

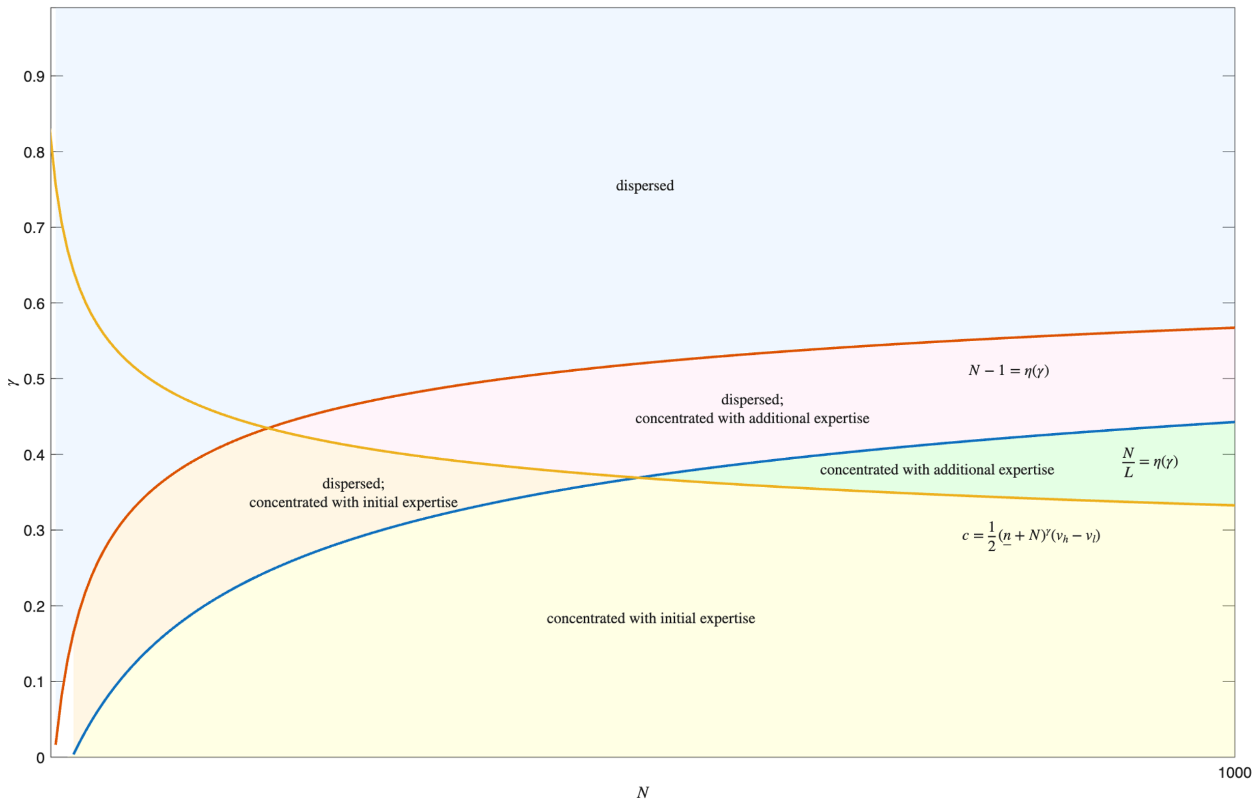

Figure 3

Equilibrium Regions For Various N and γ. This graph plots the types of equilibria (characterized by Propositions 1-6) that arise when we allow the number of independent investors N and the degree of information spillovers γ to vary, assuming that $c = 5$ and that other parameters are held fixed according to Table 1.

should have a higher γ, making it more likely that they acquire additional expertise and capture a larger share of the gains to trade away from their investors. Dealers finding it optimal to specialize in different segments (e.g., issuing entities, investor clienteles) can also help explain why, in reality, order-flow concentration often revolves around more than one dealer as, unlike in our environment, competing dealers are not perfect substitutes ex ante.

Our comparative statics also have implications for the evolution of OTC trading networks. If one were to predict how parameters are likely to evolve over time, it would be sensible to conjecture that expertise will get cheaper, its spillovers across more diversified asset bases will weaken, and/or the number of independent investors who get to choose their dealers will increase. Our model then predicts that these changes would result in a unique equilibrium outcome that is more likely to feature independent investors concentrating their order flow toward one dealer making large investments in expertise that can be used against investors. Additionally, if the growth in independent investors were to come from previously loyal investors, it would result in peripheral dealers' remaining loyal investors obtaining lower liquidity as a result of these changes.

4 Robustness and Extensions

We now discuss how our paper's main insights extend when we relax a few simplifying assumptions imposed in the baseline model.

4.1 Milder Adverse Selection

Recall that at the end of subsection 2.1 we imposed Assumption 1: $0 \leq b < v_h - v_l$. We now briefly discuss how relaxing these boundaries on the surplus from trade b would affect our predictions. First, if $b \leq 0$, there would be no gains from trade and standard no-trade theorem logic would apply: no trade happens in equilibrium and no dealer acquires additional expertise. Second, if instead $b \geq v_h - v_l$ (> 0), investors' benefit from trading would always dominate their adverse selection concerns, guaranteeing that any uninformed investor would follow a pricing strategy resulting in liquid trading (regardless of the dealers' expertise acquisition choice). All dealers would then provide liquid trading to their investors. This equilibrium characterization is reminiscent of what happened in our baseline analysis when the volume of loyal order flow was high (i.e., $\eta < 0$). In fact, depending on the cost of expertise, we would end up with equilibrium outcomes consistent with either Proposition 5 or Proposition 6 (with the minor adjustment that $\bar{a}$ is now defined as $\frac{1}{2}$).

4.2 Uneven Allocation of Expertise or Order-Flow Refusal

In our baseline model, dealers chose their level of expertise and investors chose the dealers they would trade with. Once the aggregate amount of order flow that investors sent to dealer j was known, this dealer's expertise capacity was "mechanically" spread out evenly among all transactions and trade occurred following investors' optimal bidding behavior. But what would happen if the central dealer was able to refuse order flow in order to keep its informational advantage high? Or what if the central dealer could unevenly spread its expertise across transactions?

Equation (6) from the baseline analysis showed that a dealer j's total expected profit could be written as $\Delta(k_j, n_j) = n_j a_j (v_h - v_l) - c(k_j - \underline{k})$ as long as $a_j \leq \bar{a}$. Using $a_j = \frac{1}{2}\min\left(\frac{k_j}{n_j^{1-\gamma}}, 1\right)$ and the fact that a central dealer would never acquire costly expertise capacity that pushes $\frac{k_j}{n_j^{1-\gamma}} > 1$, we can also write $\Delta(k_j, n_j) = \frac{k_j}{2} n_j^{\gamma}(v_h - v_l) - c(k_j - \underline{k})$, which means that once expertise acquisition has taken place, the dealer does not benefit from refusing order flow. For a given level of expertise capacity k_j, a central dealer that refuses to service the order flow from some investors would gain from having stronger informational advantages in the remaining (accepted) transactions, but this per-transaction gain would be more than offset by the reduction in the number of transactions the dealer extracts surplus from. Moreover, by refusing order flow, this dealer would risk violating the condition $a_j \leq \bar{a}$, which makes its expected profit negative.

Similarly, the linearity of dealer j's expected profit to k_j when $a_j \leq \bar{a}$ implies that promising to use less expertise against investor i, resulting in more expertise being used against investor i', would be weakly suboptimal for dealer j. Indeed, an uneven allocation of expertise would result either in the same dealer profit as in our baseline analysis if trading with both investors remained liquid or in a strictly lower dealer profit if the higher expertise used against investor i' were to cross the adverse selection threshold $\bar{a}$.[7]

Overall, in our setting a central dealer would collect a weakly lower profit than in our baseline analysis if it were to refuse some investors' order flow or unevenly allocate its expertise across transactions.[8]

[7] On their end, peripheral dealers would like to promise uneven allocations of expertise across transactions such that $a_j \leq \bar{a}$ for as many investors as possible, since dealers make zero profit from any transaction where their $a_j > \bar{a}$. This promise to investors would, however, not be credible as dealers would have an incentive to secretly use all their expertise on the transactions that investors expect to be liquid.

[8] For order-flow refusal or uneven allocation of expertise to become optimal, we would need the liquidity externalities to exhibit sufficient concavity. For example, if trading an additional asset with a new investor became increasingly

4.3 Alternative Timeline

Our baseline model split agents' decisions into two stages. In the first stage, investors' order-flow allocation and dealers' expertise acquisition were set simultaneously, as Nash best responses. This simultaneity allowed our model to cleanly highlight the feedback loop between investors' choice of dealers and dealers' choice of expertise investments. In the second stage, investors offered prices to their chosen dealers, using the order-flow allocations and expertise levels determined in the first stage as inputs for the optimal bargaining strategies. Given that a key contribution of our paper was to model each dealer's informational advantage a_j as a function of both its expertise capacity k_j and its transaction volume n_j, relegating the pricing and trading of assets to the second stage allowed us to micro-found payoff functions that could be used to solve for a Nash equilibrium of the first stage. Although the terms of trade originated from a bilateral trading game between one dealer and one investor in the second stage, competition among dealers was still taking place, but it was within our first stage as investors' order flow was rationally allocated toward the dealer expected to provide the best terms of trade (an equilibrium object by itself).[9]

While merging the pricing/trading stage with the simultaneous decisions that occur in the first stage would make our equilibrium analysis untractable, we can provide intuition about how our feedback loop would be affected if order-flow allocation and expertise acquisition became sequential best responses rather than simultaneous best responses, within the first stage. In particular, we now discuss how these alternative timelines would impact the salient equilibrium prediction established in Proposition 2, which best illustrates the feedback loop highlighted throughout our paper.

As in Proposition 2, consider the case where N is relatively high and c is relatively low. If independent investors were to allocate their order flow first (in stage 1a), any dealer receiving a

cumbersome in terms of the expertise capacity a dealer needs to gain an informational advantage, then the dealer's expected profit could be decreasing in its total order flow n_j.

[9] See Hendershott et al. (2020) who document the high persistence of dealer-investor relationships, suggesting that investors choose to allocate their order flow to specific dealers for the long run, without continuously observing and comparing transaction-by-transaction terms of trade.

sufficient level of order flow would then respond (in stage 1b) by acquiring additional expertise capacity k_j that sets its informational advantage equal to the adverse selection threshold: $a_j = \bar{a}$. As a result, all investors would collect a trading profit $\frac{b}{2}$, regardless of which dealer they chose. By allowing dealers to adjust their expertise acquisition based on investors' order flow, but not the converse, we would shut down the part of the feedback loop that loads on liquidity externalities. Indeed, the benefits investors get from concentrating their order flow would be neutralized by dealers' subsequent responses and investors' resulting indifference would trigger a multiplicity of possible equilibrium types, unlike in Proposition 2.

Instead, if dealers were to acquire expertise first (in stage 1a), each independent investor would then respond (in stage 1b) by allocating its order flow to the dealer expected to offer the best terms of trade (i.e., the lowest a_j, which depends on n_j and k_j). Due to liquidity externalities, order-flow concentration would still arise in equilibrium, but it could be toward any dealer with a sufficiently low k_j (not necessarily the lowest). Consistent with our baseline analysis, any dealer who anticipates to be chosen as the central dealer by investors (in stage 1b) would have an incentive (in stage 1a) to acquire expertise capacity k_j that sets its informational advantage equal to the adverse selection threshold: $a_j = \bar{a}$. Thus, despite the equilibrium multiplicity in the order-flow subgame (i.e., stage 1b), the equilibrium prediction established in Proposition 2 would still be supported as a subgame-perfect Nash equilibrium of this alternative timeline. However, if investors were able to select which dealer they coordinate on in stage 1b (remember: they are indifferent about the identity of the central dealer in the baseline analysis), they would favor the dealer that chose the lowest k_j in stage 1a. In equilibrium, we would then have no dealer acquiring additional expertise capacity above the initial level $\underline{k}$, thereby resulting in dealers appropriating the lowest possible amount of surplus, in the spirit of Bertrand competition.

Overall, it should be clear from this discussion that the economic forces at play in all these alternative timelines are the same as in our baseline analysis: popular dealers have incentives to boost their informational advantages over investors and investors have incentives to allocate their order flow toward dealers whose informational advantages are weakened by their popularity. By making expertise acquisition and order-flow allocation sequential, however, we are effectively toning down the bi-directional feedback loop at the heart of our equilibrium analysis, thereby triggering a multiplicity of potential outcomes. Our baseline timeline was chosen to illustrate our paper's main economic insights in the sharpest manner.

4.4 Ex Ante Dealer Heterogeneity and Coordination

In our baseline model, the equilibrium allocation of order flow disproportionately favored one dealer even though all dealers were assumed to be homogeneous ex ante. When order-flow concentration was the equilibrium outcome and c was small enough, a central dealer found it optimal to acquire more expertise capacity than peripheral dealers. This prediction rationalized why investors in some markets (paradoxically) allocate their order flow to dealers making large expertise investments, in light of standard adverse selection concerns.

A limitation of our equilibrium analysis was, however, that any of the ex-ante identical dealers could be selected by investors to become the central dealer — indeed, the baseline analysis was silent about how independent investors coordinated on picking one of the L dealers and how this chosen dealer knew to expect more order flow when deciding how much expertise to acquire. To shed light on how the ex-ante characteristics of the dealer that turns out ex post to be the central one impact investor welfare, we can allow for ex-ante heterogeneity in c, in $\underline{k}$, or in $\underline{n}$, and rank possible equilibria in terms of investors' welfare. Corollary 1 already showed that investors weakly prefer to trade with dealers with a lower a_j. It is therefore intuitively clear that, depending on the assumed source of heterogeneity, independent investors would weakly prefer to coordinate their order-flow allocation toward the dealer starting with the highest cost of expertise c, the lowest initial expertise bandwidth $\underline{k}$, or the highest measure of loyal order flow $\underline{n}$.

4.5 Payment for Order Flow

In contrast with models like Green (2007) in which dealers take advantage of their loyal/naive clienteles, the peripheral dealers in our baseline model collected no profit from trading with loyal investors when $\eta \geq 0$. This difference was due to the fact that, while by construction loyal investors cannot switch dealers in both settings, in our setting loyal investors could still worsen the terms of trade they offered if their local dealer violated the adverse selection threshold $\bar{a}$. As an implication of this result, if we allowed dealers to pay for order flow, peripheral dealers would be unwilling to offer any compensation aimed at attracting independent investors' order flow, unless the additional order flow was so large that it would reduce the dealer's information advantage enough to allow for liquid trading. Thus, subject to an intuitive parametric restriction, the prediction of order-flow concentration would survive in an alternative environment in which payment for order flow is allowed.

5 Conclusion

We jointly model dealers' acquisition of expertise and investors' allocation of order flow in OTC markets. An important and novel feature of our analysis consists of how we model dealers' expertise, in a way reminiscent of models of rational inattention (see, e.g., Sims 2003, Kacperczyk, Van Nieuwerburgh, and Veldkamp 2016, Maćkowiak, Matĕjka, and Wiederholt 2023) and consistent with financial intermediaries' attention constraints empirically documented by Corwin and Coughenour (2008) and Chakrabarty and Moulton (2012). In our model, a dealer's investment in expertise determines the resources that this dealer can allocate toward gaining informational advantages over its counterparties across multiple transactions involving different assets. Liquidity externalities arise in our model since order-flow concentration spreads out the resources that the central dealer can use to assess each proposed transaction, thereby weakening each investor's adverse selection concerns and increasing the liquidity of trade. Yet, expecting this high concentration of order flow, the central dealer finds it profitable to acquire the highest level of expertise that does not push investors away.

Under intuitive conditions related to the degree of information spillovers, the number of independent investors, and the cost of expertise, we show that the equilibrium allocation of order flow can be concentrated toward one dealer that invests significant resources to gain informational advantages against its investors. Despite the adverse selection concerns that could be associated with the dealer's large investments in expertise, investors prefer to funnel their transactions to this expert dealer rather than trading with less popular dealers who provide less liquidity. Our analysis sheds light on the drivers behind the concentration of OTC intermediation as well as on the welfare implications of easier access to OTC markets.

Appendix

A Proofs

Proof of Lemma 1: From (3), it is optimal to bid P_l if and only if $\Pi(\alpha, b, P_l) > \Pi(\alpha, b, P_h)$. Hence, the buyer's optimal bidding strategy can be written as:

$$P^* = \begin{cases} P_h & \text{if } a \leq \frac{b}{2(v_h - v_l)}, \\ P_l & \text{if } a > \frac{b}{2(v_h - v_l)}. \end{cases} \tag{A1}$$

□

Proof of Corollary 1: Substitute the optimal bidding policy of the buyer into the buyer's profit:

$$\Pi(\alpha, b, P) = \begin{cases} b - a(v_h - v_l) & \text{if } a \leq \frac{b}{2(v_h - v_l)}, \\ \frac{b}{2} & \text{if } a > \frac{b}{2(v_h - v_l)}. \end{cases} \tag{A2}$$

This is a weakly-decreasing function of a without jumps:

$$\frac{\partial \Pi(\alpha, b, P)}{\partial a} = \begin{cases} -(v_h - v_l) & \text{if } a \leq \frac{b}{2(v_h - v_l)}, \\ 0 & \text{if } a > \frac{b}{2(v_h - v_l)}. \end{cases} \tag{A3}$$

□

Proof of Lemma 2: First of all, for a given level of order flow n_j and expertise bandwidth k_j, a dealer's expected profit, when the investor is the buyer and the dealer is the seller, is given by:

$$\Delta(k_j, n_j) \equiv \begin{cases} \frac{n_j}{2}\left[\alpha_j(P_h - v_h) + (1-\alpha_j)(P_h - v_l) + \alpha_j(P_h - v_l) + (1-\alpha_j)(P_h - v_h)\right] - c \cdot (k_j - \underline{k}) & \text{if } P = P_h, \\ \frac{n_j}{2}\left[\alpha_j(P_l - v_l) + (1-\alpha_j)(P_l - v_h)\right] - c \cdot (k_j - \underline{k}) & \text{if } P = P_l, \end{cases} \tag{A4}$$

where $\alpha_j = \frac{1}{2} + a_j = \frac{1}{2} + \frac{1}{2}\min\left(\frac{k_j}{n_j^{1-\gamma}}, 1\right)$ and b is from the range of parameters consistent with Assumption 1. We can substitute the optimal bidding function from Proposition 1 to arrive at equation (6). Then we can show that when $a_j < \frac{b}{2(v_h - v_l)}$:

$$\frac{\partial \Delta\left(k_j, n_j\right)}{\partial k_j} = \frac{1}{2} n_j^{\gamma}(v_h - v_l) - c. \tag{A5}$$

Thus, as long as $c < \frac{1}{2} n_j^{\gamma}(v_h - v_l)$, $\Delta\left(k_j, n_j\right)$ is strictly increasing in k_j. □

Proof of Lemma 3: Suppose, by contradiction, that in a conjectured equilibrium given all dealers' choices of expertise, independent investors allocate their order flow among at least two dealers, that is, there exist dealers j' and j'' such that $n^*_{j'} > 0$ and $n^*_{j''} > 0$. Then, if $a_{j'} < a_{j''} < \frac{b}{2(v_h - v_l)}$ or $a_{j'} < \frac{b}{2(v_h - v_l)} < a_{j''}$, it is optimal for an independent investor who was planning to do business with dealer j'' to switch to dealer j' because of the weaker informational advantage that dealer j' would have when trading against this investor, following Corollary 1. Thus, this allocation of order flow cannot be part of an equilibrium. If instead $\frac{b}{2(v_h - v_l)} > a_{j'} > a_{j''}$ or $a_{j'} > \frac{b}{2(v_h - v_l)} > a_{j''}$, the reverse is true, and similarly this order-flow allocation cannot be part of an equilibrium.

Now if $a_{j'} = a_{j''} < \frac{b}{2(v_h - v_l)}$, an independent investor who was planning to do business with either dealer is strictly better off switching because $\frac{k_{j'}}{(n_{j'}+1)^{1-\gamma}} < \frac{k_{j''}}{(n_{j''}-1)^{1-\gamma}}$ and $\frac{k_{j'}}{(n_{j'}-1)^{1-\gamma}} > \frac{k_{j''}}{(n_{j''}+1)^{1-\gamma}}$. Finally, if $a_{j'} \geq \frac{b}{2(v_h - v_l)}$ and $a_{j''} \geq \frac{b}{2(v_h - v_l)}$, then an independent investor is better off switching to the dealer(s) whose $a_j < \frac{b}{2(v_h - v_l)}$ (which exist(s) given the statement of the lemma). Hence, no matter the dealers' expertise levels, it is always optimal for a subset of independent investors to reallocate their order flow whenever the order flow of independent investors is allocated among two or more dealers and at least one dealer provides liquid trading (i.e., $a_j < \frac{b}{2(v_h - v_l)}$). □

Proof of Proposition 1: To show that the stated equilibrium is indeed an equilibrium, we first rule out that no dealer would acquire additional expertise $k_j > \underline{k}$ in any circumstances, including those from the stated equilibrium. From equation (6) and Lemma 2, we know that the benefit of acquiring additional expertise is at most $\frac{1}{2}(\underline{n} + N)^{\gamma}(v_h - v_l) - c < 0$, since no dealer can receive more than

$(\underline{n}+N)$ in order flow and the proposition states that $c > \frac{1}{2}(\underline{n}+N)^{1-\gamma}(v_h - v_l)$. Thus, regardless of the equilibrium conjectured, the only set of strategies that can constitute an equilibrium are the ones where no dealer acquires additional expertise.

We next rule out that an independent investor would want to deviate and allocate its order flow to a different dealer than the central dealer j^*, within the stated equilibrium. The central dealer's informational advantage is $a_{j^*} = \frac{1}{2}\frac{k}{(\underline{n}+N)^{1-\gamma}}$. Because $N \geq 2$ and $L \geq 2$, we know that $N/L + 1 \leq N$, which then implies:

$$\frac{1}{2}\frac{k}{(\underline{n}+N)^{1-\gamma}} \leq \frac{1}{2}\frac{k}{(\underline{n}+\frac{N}{L}+1)^{1-\gamma}} < \frac{1}{2}\frac{k}{(\underline{n}+\eta+1)^{1-\gamma}} = \bar{a}, \tag{A6}$$

where the strict inequality follows from the stated condition $\eta < \frac{N}{L}$. By reallocating order flow from the central dealer whose $a_{j^*} < \bar{a}$ to a peripheral dealer whose informational advantage would become $a_j = \frac{1}{2}\min\left(\frac{k}{(\underline{n}+1)^{1-\gamma}}, 1\right) \geq \frac{b}{2(v_h - v_l)} = \bar{a}$ (since $\eta \geq 0$) after the deviation, an independent investor's profit would strictly decrease, thus it is not a profitable deviation.

Next, we show that alternative order-flow allocations cannot be sustained in equilibrium within the stated parameter range. First, since $c > \frac{1}{2}(\underline{n}+N)^{1-\gamma}(v_h - v_l)$, there cannot be an equilibrium in which any dealer acquires additional expertise, regardless of how order flow is allocated. Second, the condition $\frac{N}{L} > \eta$ rules out the possibility of independent investors splitting their order flow equally across dealers whose resulting $a_j > \bar{a}$. In particular, the condition $\frac{N}{L} > \eta$ stated in the proposition implies:

$$\frac{1}{2}\frac{k}{\left(\underline{n}+\frac{N}{L}+1\right)^{1-\gamma}} \; < \; \frac{b}{2\,(v_h - v_l)} \; = \; \bar{a}. \tag{A7}$$

Under equal allocation of independent investor flow, each dealer's informational advantage is $a_j = \frac{1}{2}\min\left(\frac{k}{\left(\underline{n}+\frac{N}{L}\right)^{1-\gamma}}\right)$. a_j can be less than $\bar{a}$ or just above it, such that if an independent investor was to deviate from another dealer and re-allocate the order flow to a dealer j', its informational advantage would drop below $\bar{a}$. Hence, such a deviation would be strictly profitable — either the investor is going from $a_j > \bar{a}$ to $a_{j'} < \bar{a} < a_j$ or from $a_j = \frac{1}{2}\frac{k}{\left(\underline{n}+\frac{N}{L}\right)^{1-\gamma}}$ to a dealer with

$a_{j'} < a_j$. Thus, the equal allocation of independent order flow across dealers when $\frac{N}{L} > \eta$ cannot be sustained in equilibrium.

Third, we rule any equilibrium in which independent order flow is allocated unevenly to at least two dealers. In particular, among the dealers receiving independent order flow, some of them receive $n_j > \underline{n} + \frac{N}{L}$ in total order flow. Because its per-transaction information advantage $a_j = \frac{1}{2}\frac{k}{n_j^{1-\gamma}} < \bar{a}$ (see condition $\frac{N}{L} > \eta$ above), the conditions of Lemma 3 are satisfied. Hence, no asymmetric allocation of independent order flow with at least two dealers receiving some independent investor flow can be part of an equilibrium. Indeed, only the concentrated allocation of order flow stated in the proposition can be sustained in equilibrium. □

Proof of Proposition 2: Some steps in this proof resembles those from the proof of Proposition 1, except this time we need to consider $c \leq \frac{1}{2}(\underline{n} + N)^{1-\gamma}(v_h - v_l)$. To show that the stated equilibrium is indeed an equilibrium, we first analyze how much expertise dealers would find optimal to acquire, under the stated allocation of order flow. From Lemma 2, we know that the lower c implies that a central dealer j^*, who receive $(\underline{n} + N)$ in order flow, would profit from increasing k_{j^*} as long as it keeps $a_{j^*} \leq \bar{a}$. Thus, it is optimal for the central dealer to invest in expertise until $k_{j^*} = n_{j^*}^{1-\gamma}\left(\frac{b}{v_h - v_l}\right)$ and $a_{j^*} = \bar{a}$.

In contrast, other dealers already have an informational advantage $a_j > \bar{a}$ before acquiring any additional expertise. Indeed, the condition $\eta \geq 0$ implies that:

$$a_j = \frac{1}{2}\frac{\underline{k}}{\underline{n}^{1-\gamma}} > \frac{1}{2}\frac{\underline{k}}{(\underline{n}+1)^{1-\gamma}} \geq \frac{1}{2}\frac{b}{v_h - v_l} \equiv \bar{a}. \tag{A8}$$

Thus, no peripheral dealer can provide liquid trade when only receiving order flow from their loyal investors. These dealers thus collect 0 in profit and have no incentive to spend c to boost their expertise capacity.

Within the stated equilibrium, independent investors do not have incentive to deviate by reallocating their order flow away from the central dealer. The central dealer's $a_{j^*} = \bar{a}$ provides investors with a surplus $\frac{b}{2}$, which is the same a deviating independent investor would make by trading with a peripheral dealer whose informational advantage became $a'_j = \frac{1}{2}\frac{k}{(\underline{n}+1)^{1-\gamma}} \geq \frac{1}{2}\frac{b}{v_h - v_l} \equiv \bar{a}$ (since $\eta \geq 0$) as a result of the deviation.

Next, we show that alternative order-flow allocations cannot be sustained in equilibrium within the stated parameter range. First, as in Proposition 1, the condition $\frac{N}{L} > \eta$ implies that, even if independent order flow was split evenly across dealers, we could not construct an equilibrium in which all dealers provide illiquid trading, that is, $a_j > \bar{a}$ for all dealers. (We do not repeat the arguments here, as they are identical.) Second, Lemma 3 allows to rule out any equilibrium where more than one dealer receives enough independent order flow such that their $a_j \leq \bar{a}$. In such case, independent investors would have an incentive to reallocate their order flow, consistent with the arguments developed in Lemma 3. Altogether, these results imply that we can only have one dealer whose $a_j \leq \bar{a}$ in equilibrium, and this central dealer receives all independent order flow. □

Proof of Proposition 3: First, consider an outcome where all independent order flow is allocated to a single dealer j^*. Condition $N - 1 > \eta$ implies that, without additional expertise, $\frac{1}{2}\frac{k}{(\underline{n}+N)^{1-\gamma}} < \bar{a}$. Since $a_{j^*} < \bar{a}$, Lemma 3 applies and ensures that order-flow concentration is sustainable in equilibrium. Depending on the cost of expertise c, the equilibrium either features no additional expertise (if c is high) or the central dealer acquires expertise (if c is low). Please note that in the latter case, the dealer acquires expertise so that $a_{j^*} = \bar{a}$, so the conditions of the Lemma 3 are still satisfied.

Next, consider an outcome where independent order flow is allocated equally across dealers, meaning that each dealer receives $n_j = \underline{n} + \frac{N}{L}$. Condition $\frac{N}{L} \leq \eta$ implies that $\frac{1}{2}\frac{k}{\left(\underline{n}+\frac{N}{L}+1\right)^{1-\gamma}} \geq \bar{a}$, meaning that no dealer would provide liquid trading even if an independent investor were to deviate and reallocate its order flow. This thus means that, without a deviation, each dealer's informational advantage satisfies:

$$a_j = \frac{1}{2}\frac{k}{\left(\underline{n} + \frac{N}{L}\right)^{1-\gamma}} > \bar{a}, \tag{A9}$$

and as a result, none of the dealers finds it optimal to invest in additional expertise capacity. No independent investor benefits from unilaterally reallocating its order flow away from the conjectured equilibrium allocation. Therefore, equal allocation of independent order flow and no additional expertise acquired by any dealers can be sustained as an equilibrium.

More generally, any allocation of independent investor order flow can be supported as an equilibrium provided that $n_j < \underline{n} + \eta$ for all dealers. Indeed, if $n_j < \underline{n} + \eta$, then $\frac{1}{2} \frac{\underline{k}}{(n_j+1)^{1-\gamma}} > \bar{a}$. Thus, if an independent investor unilaterally switches its order flow to another dealer, the latter's per-transaction informational advantage remains strictly above $\bar{a}$, leaving the investor's expected profit unchanged. Such deviation is therefore unprofitable and as a result the stated allocation can be sustained in equilibrium. In all of these allocations, all dealers have $a_j > \bar{a}$ at their initial level of expertise capacity $\underline{k}$, hence, they do not have incentives to acquire additional expertise. Altogether, any allocation of independent order flow with no additional expertise acquisition by dealers is an equilibrium as long as $n_j < \underline{n} + \eta$ for all dealers. □

Proof of Proposition 4: First, consider an outcome where all independent investors direct their order flow to a single dealer j^*, who does not acquire additional expertise. Condition $N - 1 < \eta$ implies that this dealer's informational advantage satisfies:

$$a_{j^*} = \frac{1}{2} \frac{\underline{k}}{(\underline{n} + N)^{1-\gamma}} > \bar{a}. \tag{A10}$$

Thus, even if all investors concentrate their order flow toward one dealer, the liquidity externalities are not strong enough to allow for liquid trading. If an independent investor were to deviate by reallocating its order flow to a peripheral dealer, this investor's expected payoff would remain at $\frac{b}{2}$. By our equilibrium definition, such order-flow allocation can therefore be sustained in equilibrium. The central dealer gets a trading profit of 0, regardless of the allocation and regardless of its level of expertise, and thus has no incentives to acquire additional expertise that would boost its a_j^* further above $\bar{a}$, at a cost c.

In fact, since trading remains illiquid for all dealers regardless of the allocation of independent order flow, any allocation that keeps dealers' expertise at $k_j = \underline{k}$ is an equilibrium. $\square$

Proof of Proposition 5: When $c > \frac{1}{2}(v_h - v_l)(\underline{n} + N)^{\gamma}$, we know that, regardless of the order-flow allocation, no dealer would find it optimal to acquire additional expertise bandwidth. Condition $\eta < 0$ also implies that each dealer has enough loyal investors that its informational advantage without attracting independent order flow would satisfy $a_j = \frac{1}{2}\frac{\underline{k}}{\underline{n}^{1-\gamma}} < \bar{a}$ and therefore allow for liquid trading. Because of the liquidity externalities, the central dealer's informational advantage $a_{j^*} = \frac{1}{2}\frac{\underline{k}}{(\underline{n}+N)^{1-\gamma}}$ is lower than peripheral dealers' $a_j = \frac{1}{2}\frac{\underline{k}}{\underline{n}^{1-\gamma}}$. Thus, no independent investors would find it optimal to unilaterally move its order flow to a peripheral dealer, even though this peripheral dealer now allows for liquid trading.

Although $\eta < 0$, when $c > \frac{1}{2}(v_h - v_l)(\underline{n} + N)^{\gamma}$, we can employ the same arguments used to prove Lemma 3 and show that only order-flow concentration can be supported in equilibrium. $\square$

Proof of Proposition 6: In this proof, we rule out all possible allocations of independent order flow and corresponding acquisition strategies by dealers as potential pure-strategy equilibria.

First, we note that given any allocation of independent order flow, the L dealers would provide liquid trading if they were to not acquire additional expertise capacity. Thus, in anticipation of any order-flow allocation, no dealer has incentives to acquire expertise capacity that pushes its $a_j > \bar{a}$ as it would reduce its trading profit to 0. Second, due to the liquidity externalities, independent investors would benefit from deviating and moving their order flow, unless all order flow is concentrated toward one dealer with the lowest a_j. However, if all independent investors were to concentrate their order flow toward one dealer, then this dealer would find it optimal to acquire additional expertise since $c \leq \frac{1}{2}(v_h - v_l)(\underline{n} + N)^{\gamma}$. Indeed, this dealer's optimal response to an allocation of $\underline{n} + N$ is to set $a_{j^*} = \bar{a}$. Yet, this strategy cannot support an equilibrium since a deviating investor who reallocates its order flow to a peripheral dealer would make this dealer's informational advantage $a_j < \bar{a}$ (by $\eta < 0$), thereby benefitting the deviating investor. While it would have been in the central dealer's best interest to limit its investment in expertise to a lower level that sustained the full concentration of order flow, it is not the central dealer's best response to an anticipated order $(\underline{n} + N)$ given the simultaneous-move game we study. Thus, no pure-strategy equilibrium exists in this region. $\square$

References

Akerlof, George A., 1970, “The Market for “Lemons”: Quality Uncertainty and the Market Mechanism,” *Quarterly Journal of Economics* 84, 488-500.

Atkeson, Andrew, Andrea L. Eisfeldt, and Pierre-Olivier Weill, 2014, “The Market for OTC Credit Derivatives,” Unpublished Working Paper, UCLA.

Babus, Ana, and Tai-Wei Hu, 2017, “Endogenous Intermediation in Over-the-Counter Markets,” *Journal of Financial Economics* 125, 200-215.

Baumann, Friedrich, Ali Kakhbod, Dmitry Livdan, Abdolreza Nazemi, and Norman Schürhoff, 2023, “Life after Default: Dealer Intermediation and Recovery in Defaulted Corporate Bonds,” Working Paper, Swiss Finance Institute.

Begenau, Julianne, Monika Piazzesi, and Martin Schneider, 2015, “Banks’ Risk Exposures,” Unpublished Working Paper, Stanford University.

Bethune, Zachary, Bruno Sultanum, and Nicholas Trachter, 2022, “An Information-based Theory of Financial Intermediation,” *The Review of Economic Studies* 89, 2381-2444.

Bank of International Settlements (BIS), 2025, “OTC Derivatives Statistics at End-June 2025.”

Boyarchenko, Nina, David O. Lucca, and Laura Veldkamp, 2021, “Taking Orders and Taking Notes: Dealer Information Sharing in Treasury Auctions,” *Journal of Political Economy* 129, 607-645.

Campbell, Dakin, 2018, “Financial Data is a Goldmine and this Number from Goldman Sachs Shows Just How Much,” *Business Insider*, May 9. 2018.

Cetorelli, Nicola, Beverly Hirtle, Donald P. Morgan, Stavros Peristiani, and Joao A.C. Santos, 2007, "Trends in Financial Market Concentration and their Implications for Market Stability," *Economic Policy Review* 13, 33-51.

Chacko, George C., Jakub W. Jurek, and Erik Stafford, 2008, "The Price of Immediacy," *Journal of Finance* 36, 1253-1290.

Chaderina, Maria, and Richard C. Green, 2014, "Predators and Prey on Wall Street," *Review of Asset Pricing Studies* 4, 1-38.

Chaderina, Maria, Alexander Muermann, and Christoph Scheuch, 2025, "The Dark Side of Liquid Bonds in Fire Sales," forthcoming, *Management Science*.

Chae, Joon, and Albert Wang, 2003, "Who Makes Markets? Do Dealers Provide or Take Liquidity?," Unpublished Working Paper.

Chakrabarty, Bidisha, and Pamela C. Moulton, 2012, "Earnings Announcements and Attention Constraints: The Role of Market Design," *Journal of Accounting and Economics* 53, 612-634.

Chang, Briana, and Zhengxing Zhang, 2021, "Endogenous Market Making and Network Formation," Unpublished Working Paper.

Chen, Kathryn, Michael Fleming, John Jackson, Ada Li, and Asani Sarkar, 2011, "An Analysis of CDS Transactions: Implications for Public Reporting," Unpublished Working Paper.

Chung, Kee H., and Seong-Yeon Cho, 2005, "Security Analysis and Market Making," *Journal of Financial Intermediation* 14, 114-141.

Corwin, Shane A., and Jay F. Coughenour, 2008, "Limited Attention and the Allocation of Effort in Securities Trading," *Journal of Finance* 63, 3031-3067.

Covarrubias, Matias, Germán Gutiérrez, and Thomas Philippon, 2019, "Explaining the Rising Concentration of U.S. Industries: Superstars, Intangibles, Globalization or Market Power?," *NBER Macroeconomics Annual* 34, 1-46.

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

Galeotti, Andrea, and Sanjeev Goyal, 2010, "The Law of the Few," *American Economic Review* 100, 1468-1492.

Glode, Vincent, Richard C. Green, and Richard Lowery, 2012, "Financial Expertise as an Arms Race," *Journal of Finance* 67, 1723-1759.

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

Goldman Sachs, 2021, "Full Year and Fourth Quarter 2021 Earnings Results".

Goldstein, Michael A., and Edith S. Hotchkiss, 2020, "Providing Liquidity in an Illiquid Market: Dealer Behavior in US Corporate Bonds," *Journal of Financial Economics* 135, 16-40.

Green, Richard C., 2007, "Presidential Address: Issuers, Underwriter Syndicates, and Aftermarket Transparency," *Journal of Finance* 62, 1529-1550.

Green, Richard C., Burton Hollifield, and Norman Schürhoff, 2007a, "Dealer Intermediation and Price Behavior in the Aftermarket for New Bond Issues," *Journal of Financial Economics* 86, 643-682.

Green, Richard C., Burton Hollifield, and Norman Schürhoff, 2007b, "Financial Intermediation and the Costs of Trading in an Opaque Market," *Review of Financial Studies* 20, 275-314.

Grossman, Sanford, and Motty Perry, 1986, "Perfect Sequential Equilibrium," *Journal of Economic Theory* 39, 97-119.

Hagströmer, Björn, and Albert J. Menkveld, 2019, "Information Revelation in Decentralized Markets," *Journal of Finance* 74, 2751-2787.

Hasbrouck, Joel, and Richard M. Levich, 2021, "Network Structure and Pricing in the FX Market," *Journal of Financial Economics* 141, 705-729.

Hendershott, Terry, Dan Li, Dmitry Livdan, and Norman Schürhoff, 2020, "Relationship Trading in Over-the-Counter Markets," *Journal of Finance* 75, 683-734.

Hendershott, Terry, Dan Li, Dmitry Livdan, and Norman Schürhoff, 2024, "When Failure is an Option: Fragile Liquidity in Over-the-Counter Markets," *Journal of Financial Economics* 157: 103859.

Herskovic, Bernard, and João Ramos, 2020, "Acquiring Information through Peers," *American Economic Review* 110, 2128-2152.

Hollifield, Burton, Artem Neklyudov, and Chester Spatt, 2017, "Bid-Ask Spreads, Trading Networks, and the Pricing of Securitizations," *Review of Financial Studies* 30, 3048-3085.

Jotikasthira, Chotibhak, Christian T. Lundblad, and Jinming Xue, 2025, "Dealer Specialization and Market Segmentation," Unpublished Working Paper.

Joseph, Andreas, and Michalis Vasios, 2022, "OTC Microstructure in a Period of Stress: A Multi-Layered Network Approach," *Journal of Banking & Finance* 138, 106400.

Kacperczyk, Marcin, Stijn Van Nieuwerburgh, and Laura Veldkamp, 2016, "A Rational Theory of Mutual Funds' Attention Allocation," *Econometrica* 84, 571-626.

Kyle, Albert S., 1985, "Continuous Auctions and Insider Trading," *Econometrica* 53, 1315–1335.

Lester, Benjamin, Guillaume Rocheteau, and Pierre-Olivier Weill, 2015, "Competing for Order Flow in OTC Markets," *Journal of Money, Credit and Banking* 47, 77-126.

Li, Dan, and Norman Schürhoff, 2019, "Dealer Networks," *Journal of Finance* 74, 91-144.

Li, Wei, and Zhaogang Song, 2024, "Dealer Expertise and Market Concentration in Over-the-Counter Trading," Unpublished Working Paper.

Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt, 2023, "Rational Inattention: A Review," *Journal of Economic Literature* 61, 226-273.

Manaster, Steven, and Steven C. Mann, 1996, "Life in the Pits: Competitive Market Making and Inventory Control," *Review of Financial Studies* 9, 953-975.

Munyan, Benjamin, and Sumudu W. Watugala, 2018, "What Makes Dealers Central? Evidence from Credit Interdealer Networks," Unpublished Working Paper.

Neklyudov, Artem, 2019, "Bid-Ask Spreads and the Over-the-Counter Interdealer Markets: Core and Peripheral Dealers," *Review of Economic Dynamics* 33: 57-84.

Office of the Comptroller of the Currency (OCC), 2021, "Quarterly Report on Bank Trading and Derivatives Activities – Third Quarter 2021".

Pagano, Marco, 1989, "Trading Volume and Asset Liquidity," *Quarterly Journal of Economics* 104, 255-274.

Sambalaibat, Batchimeg, 2025, "Heterogeneous Clienteles and Dealer Networks," *Journal of Financial Economics* 174:1 104185.

Samuelson, William, 1984, "Bargaining Under Asymmetric Information," *Econometrica* 52, 995-1005.

Sims, Christopher, 2003, "Implications of Rational Inattention," *Journal of Monetary Economics* 50, 665-690.

Siriwardane, Emil, 2019, "Limited Investment Capital and Credit Spreads," *Journal of Finance* 74, 2302-2347.

van der Wel, Michel, Albert J. Menkveld, and Asani Sarkar, 2009, "Are Market Makers Uninformed and Passive? Signing Trades in the Absence of Quotes?," Unpublished Working Paper.