Trading Against Expert Dealers *

Maria Chaderina Vincent Glode

November 21, 2024

Abstract

We model investors' allocation of order flow across over-the-counter dealers jointly with dealers' acquisition of expertise that can be used to take advantage of investors across transactions. Ceteris paribus, investors benefit from allocating their order flow to dealers expected to intermediate large volumes of transactions and to acquire low levels of expertise, whereas dealers profit more from acquiring expertise when intermediating large volumes of transactions. Our model's equilibrium rationalizes why the most sought-after dealers are often those with the best data, technology, and skills, despite the significant adverse selection concerns triggered by their superior 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, 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

A seminal insight from the literature on asymmetric information, dating back to at least Akerlof (1970), is that one should be wary of trading against a counterparty that possesses superior expertise (e.g., data, technology, and skills). Such counterparty can condition its trading decisions on superior information and take advantage of less informed traders. Yet, in real financial markets, the dealers that intermediate most trades happen to be extremely sophisticated institutions (see, e.g., Hagströmer and Menkveld 2019). For example, Goldman Sachs remains the second most active over-the-counter (OTC) dealer of derivatives among U.S. banks while spending $\$400M$ 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.[1] When a financial institution acts as an OTC "dealer," it buys and sells securities on its own behalf, using its expertise 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).[2] It is thus puzzling that most investors and traders opt to trade against expert dealers, despite the obvious adverse selection concerns associated with their superior expertise. Additionally, why don't we see dealers trying to build a reputation for lacking the private information and skills used to take advantage of their counterparties, thereby minimizing adverse selection concerns?

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 — expertise allows dealers to take advantage of the less informed investors that send them order flow. Investors, however, take into account the expertise each dealer is expected to use against them when they choose how to split order flow across dealers.

[1] See OCC (2021), Campbell (2018), and Goldman Sachs' 2021-Q4 earnings report for the specific numbers in the sentence. For evidence of concentrated OTC intermediation for various types of assets, see Cetorelli et al. (2007), Atkeson, Eisfeldt, and Weill (2014), Begenau, Piazzesi, and Schneider (2015), Di Maggio, Kermani, and Song (2017), Hagströmer and Menkveld (2019), Li and Schürhoff (2019), Siriwardane (2019), and Hendershott et al. (2020), among many others.

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

An important and novel feature of our analysis consists of how we model dealers' expertise to capture the limited information spillovers across assets and transactions. In most models of information acquisition, a trader's expertise level dictates the informational *quality* of the signals this trader can observe for all transactions being considered. 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 or transaction is likely to have limited pricing implications for other securities and transactions. For example, the majority of corporate bonds trade less than twice a year (see Chaderina, Muermann, and Scheuch 2022) while differences in maturity, duration, covenants, and collateral pledged render even bonds from the same issuer imperfect substitutes. Thus, if a dealer gains an informational advantage for a specific transaction, it does not mean that this dealer will also benefit from a similar advantage in all future transactions.

In our model, we assume that a dealer's expertise level dictates 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 various assets within a short time period. When investors are interested in trading different assets (e.g., various corporate bonds, municipal bonds, and 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 are no longer available to value a currency swap that is part of a different transaction. As more order flow gets directed toward this dealer, its resources are spread out more thinly across the 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 essential forces in our model.

As hinted earlier, our model features a tension between investors' and dealers' preferences for dealers' expertise capacity. Ceteris paribus, investors prefer to trade against dealers possessing little expertise capacity, as it reduces adverse selection concerns. Dealers, on the other hand, prefer to have access to high expertise capacity, as it allows them to extract a larger share of the surplus when trading with investors. We show that, when the cost of expertise is sufficiently low, trading is concentrated around one dealer that acquires a high level of expertise 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 around 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 around 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. The endogeneity of dealers' expertise levels is the second essential force in our model.

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) and Sambalaibat (2018) 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 Li and Schürhoff (2019) find that a few dealers tend to attract most OTC trade orders and that the marginal transaction intermediated by these central dealers commands a smaller markup than the marginal transaction intermediated by more peripheral dealers. Baumann et al. (2023) find that following corporate default, bond investors concentrate their order flow around 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 equilibrium prediction that, despite its high expertise investments, a central dealer offers better terms of trade than competing peripheral dealers.

A few recent papers have rationalized order-flow concentration around 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 around 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 around 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 are, however, silent about the 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, dealers acquire costly expertise while investors choose the dealers with whom they will trade. In the second stage, trade takes place between investors and dealers. We will first introduce the trading game and analyze 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 as simple as possible. While the model makes formal assumptions about how trading occurs among agents, we will later emphasize the generic properties of trading games that our central insights rely on.

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 $\alpha = \frac{1}{2} + e$. We use $e_i \in \left[0, \frac{1}{2}\right]$ to denote the seller's level of expertise for this specific transaction (relative to what the buyer knows). If $e = 0$, the seller's private signal is uninformative about the value of the asset, whereas if $e = \frac{1}{2}$, 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 e. While the seller's transaction-specific expertise e 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 e as a function of 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:

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

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:

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

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:

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

$$= \begin{cases} \frac{b}{2} & \text{if } P = P_l, \\ b - e(v_h - v_l) & \text{if } P = P_h. \end{cases} \tag{3}$$

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 transaction-specific expertise e.

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

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

Throughout the paper, proofs of our formal results are relegated to the Appendix. As Proposition 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 $e \leq \bar{e}$, 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 expertise level e. However, for high levels of expertise $e > \bar{e}$, the seller's informational advantage becomes too concerning for the buyer who prefers to offer the low price P_l, making the optimal bid P^* decreasing in the seller's expertise level e. 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 transaction-specific expertise as long as the buyer is willing to offer P_h, but having too much expertise 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 transaction-specific expertise level e for two levels of private benefit b.

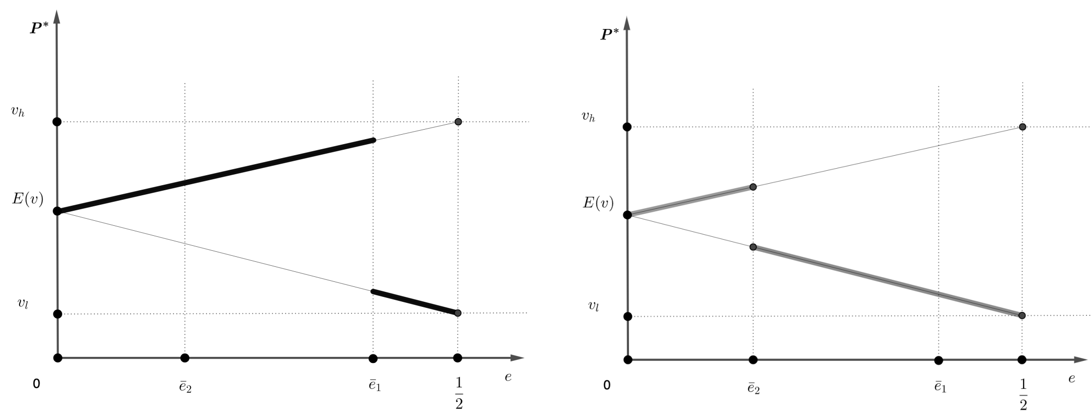

(a) Gains to trade b_1

(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 expertise level e, 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 the seller's expertise e crosses either $\bar{e}_1 \equiv \frac{b_1}{2(v_h - v_l)}$ in panel (a) or $\bar{e}_2 \equiv \frac{b_2}{2(v_h - v_l)}$ in panel (b).

The next corollary establishes that the buyer is made better off by being paired with a seller with lower transaction-specific expertise.

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

For the remainder of the paper, we impose a parametric restriction on b that results in adverse selection impeding trade for high enough levels of transaction-specific expertise:

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

Imposing $b < v_h - v_l$ guarantees that the liquid-trading threshold satisfies $\bar{e} < \frac{1}{2}$ and that there exists a region $\bar{e} < e \leq \frac{1}{2}$ where $P^* = P_l$ is the optimal bidding strategy. This restriction will help capture, in the analysis below, the notion that dealers must limit how much expertise they acquire given investors' adverse selection concerns.

2.2 Expertise Acquisition and Order-Flow Allocation Stage

Using the trading game analyzed above, we now study how dealers choose how much expertise to acquire and how investors choose the dealers with whom they trade. While the analysis of the trading stage in subsection 2.1 focused on the seller being privately informed, the payoff functions we derived would not change if we instead assumed that an uninformed seller quoted a price to a buyer with expertise level e. Thus, we can interpret the buyer's expected profit Π derived in subsection 2.1 more broadly as the payoff of an unsophisticated investor trading against an expert dealer, regardless of who is long and who is short the asset. 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' decisions are made simultaneously.

As briefly explained in the introduction, our model is designed to capture the notion that a dealer's investment in expertise improves 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 different asset (e.g., a specific corporate bond, municipal bond, or derivative product). Each dealer starts with $\underline{k}$ units of total expertise (e.g., data, human capital, computing power) available to value assets and gain an informational advantage over investors. This initial level of expertise 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 without being motivated by dealers' trading profits per se.[3] Each dealer j can then expand its total expertise to $k_j > \underline{k}$ at a cost $c(k_j - \underline{k})$ if such investment is deemed to increase the dealer's trading profits. As we show below, each dealer's optimal investment in expertise depends both on how order flow is expected to be allocated in the first stage and on how trading is expected to occur in the second stage.

Each dealer is assumed to start with a loyal client base of measure $\underline{n}$. These are unsophisticated investors who always trade with the same (e.g., local) dealer, in the spirit of Green (2007). We refer to the N investors who are not loyal to any specific dealer and who optimize based on the expected terms of trade as independent investors. A dealer j's total order flow $n_j \equiv \underline{n} + n_j^*$ combines the transaction volume from its loyal investors and from the independent investors who chose this particular dealer. In the simultaneous-move game of this stage, each independent investor chooses which dealer to do business with and each dealer chooses how much additional expertise capacity to acquire. When considering doing business with dealer j, investor i conjectures the transaction-specific expertise e_j that will allow this dealer to gain an informational advantage. Then trade

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

occurs as described in subsection 2.1, taking these decisions as given.

Dealer j's transaction-specific expertise is assumed to be $e_j = \frac{1}{2}\min\left(\frac{k_j}{n_j}, 1\right)$, reflecting the dealer j's' expertise capacity k_j normalized by its total order flow n_j while ensuring that the signal's probability of being correct, $\frac{1}{2} + e_j$, cannot be greater than one. When dealer j receives weakly less order flow n_j than its expertise capacity k_j, dealer j is able to produce a perfect signal about the value of each asset traded with investors. However, if the order flow n_j is larger than the expertise capacity k_j, dealer j's resources are spread out across all proposed transactions and the resulting signal associated with each transaction is imperfect. Dealers are allowed to be short or long any asset in their inventory (recall: the model behaves symmetrically whether the dealer is a buyer or a seller).

Substituting the optimal bid P^* derived in Proposition 1 into the dealer's profit function, we can write the dealer's total expected profit as:

$$\Delta\left(k_j, n_j\right) = \begin{cases} -c(k_j - \underline{k}) & \text{if } e_j > \bar{e}, \\ n_j e_j (v_h - v_l) - c(k_j - \underline{k}) & \text{if } e_j \leq \bar{e}. \end{cases} \tag{5}$$

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

Lemma 1. *For a given n_j, $\frac{\partial \Delta(k_j, n_j)}{\partial k_j} > 0$ as long as $e_j < \bar{e} \equiv \frac{b}{2(v_h - v_l)}$ and $c < \frac{1}{2}(v_h - v_l)$.*

Manaster and Mann (1996), Chae and Wang (2003), and van der Wel, Menkveld, and Sarkar (2009), among others, document how dealers and market makers use their superior information to boost their own profits. In our setting, a dealer benefits from increasing its expertise in two ways. A first benefit comes from making better decisions when responding to offers. This benefit is most evident when $e_j > \bar{e}$ as investors are 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 is triggered when $e_j \leq \bar{e}$. In that region, if the dealer's signal becomes marginally more precise, investors respond by making more generous offers to ensure that the dealer agrees to trade with probability one. While our model makes formal assumptions about how trading occurs among agents, the central insights we highlight throughout the paper rely on the generic properties 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.

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:

$$\Pi\left(\frac{1}{2}+\frac{1}{2}\min\left(\frac{k_j^*}{n_j},1\right),b,P_j^*\right) \geq \Pi\left(\frac{1}{2}+\frac{1}{2}\min\left(\frac{k_{j'}^*}{n_{j'}+1},1\right),b,P_{j'}^*\right) \qquad \forall j' \neq j, \tag{6}$$

  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:

$$\Delta\left(k_j^*, n_j\right) \geq \Delta\left(k_j', n_j\right) \qquad \forall k_j' \neq k_j^*. \tag{7}$$

As already stated, we model dealers' and investors' decisions as a simultaneous-move game. This timeline captures the idea 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 yet know for sure dealers' future levels of expertise, whereas dealers are not yet certain of investors' future allocation 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 and enter the trading stage knowing what each dealer's transaction-specific expertise e_j is.

We should also emphasize that 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 forecast dealers' expertise level. In reality, investors' knowledge about dealers' expertise might be acquired through a more complicated (dynamic) game than what we currently model, yet the two generic properties of the trading game that we highlighted above can survive in many alternative economic environments.

For the remainder of this section, we make the following parametric assumption about dealers and investors:

Assumption 2. *All dealers are ex-ante identical with respect to their initial levels of expertise capacity* $\underline{k}$ *and order flow from loyal investors* $\underline{n}$. *Moreover,* $\underline{k} \geq \underline{n} + 1$.

Assumption 2 ensures that if any dealer is only intermediating its loyal investors' flow, the resulting transaction-specific expertise is above the liquid trading threshold $\bar{e}$, and trade happens only half of the time. If adverse selection concerns were milder and these dealers' transaction-specific expertise was below $\bar{e}$, then adverse selection would not prevent trade to remain liquid and would solely transfer surplus. Since the focus of our paper is on the effects of adverse selection on order-flow allocation, we impose this parametric restriction that will ensure that adverse selection is a serious concern for investors when they decide how to allocate their order flow.

In what follows, we show how concentrating trade around one dealer can reduce the negative effects of dealers' endogenous adverse selection, thereby allowing independent investors to trade efficiently with their chosen dealer(s). Before fully characterizing the equilibrium, we develop the basic intuition for why order-flow concentration benefits independent investors.

Lemma 2. *If in equilibrium at least one dealer has transaction-specific expertise that does not prevent liquid trading (i.e.,* $e_j \leq \bar{e} \equiv \frac{b}{2(v_h - v_l)}$*), then all independent investors' order flow is allocated to the same dealer and, as a result, other dealers only receive order flow from their loyal investors.*

Intuitively, this lemma shows that if independent investors were expected to split their transaction volume among two or more dealers, then each independent investor would benefit from re-routing its 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 e_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 and minimize its informational advantage in each transaction. The existence of one dealer whose $e_j \leq \bar{e}$ ensures that while an additional independent investor switching to dealer j would further decrease e_j, it would also benefit this investor as well as all other investors that trade with this dealer. Recall from Corollary 1 that an investor's trading profit is decreasing in its dealer's expertise e_j as long as $e_j \leq \bar{e}$ and trading thereby remains liquid.

Lemma 2 thus highlights how "liquidity externalities" can incentivize independent investors to concentrate their order flow around one dealer. Each investor prefers to trade with a dealer that is attracting a lot of other independent investors, because this dealer uses less of its expertise capacity to take advantage of this specific investor. In response to the lower per-transaction expertise of this dealer, 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 in the next proposition, dealers take the allocation of order flow into account when choosing how much expertise to acquire. But before deriving this result, we need to rule out extreme outcomes by imposing a parametric restriction on the populations of dealers and investors.

Assumption 3. *The market is populated with L dealers and N independent investors, such that* $\frac{\underline{k}}{\underline{n}+\frac{N}{L}} \leq \frac{b}{v_h - v_l}$.

This restriction limits how much expertise dealers start with for the case where aggregate order flow is split evenly across all dealers. Dealers' initial level of expertise $\underline{k}$ is then low enough to allow for liquid trading as investors still find it optimal to make generous offers (see Proposition 1), despite their adverse selection concerns. This restriction also implies that, if all independent investors were to send their order flow to the same dealer, spreading the expertise capacity $\underline{k}$ across this larger transaction volume would result in transaction-specific expertise e_j that is still low enough to allow for liquid trading. We will study the case in which this assumption is violated in Section 3.

In the next proposition we combine the insights of Lemmas 1 and 2 with Assumption 3 and pin down the equilibrium allocations of investors' order flow and the equilibrium levels of dealers' expertise.

Proposition 2. *When $c < \frac{1}{2}(v_h - v_l)$, in equilibrium one dealer, say j^*, receives all the order flow from independent investors and acquires the highest level of expertise that allows trading to remain liquid, whereas all other dealers trade only with their loyal investors and do not acquire additional expertise above the initial level $\underline{k}$.*

Proposition 2 establishes that, when the cost of expertise is low enough, in equilibrium all independent investors' order flow is concentrated around one central dealer that makes the largest investment in expertise among all dealers. While its investors would prefer this dealer to have weaker informational advantages when trading, alternative dealers who acquire less expertise than the central dealer also happen to participate in fewer transactions. In fact, Assumption 2 implies that these "peripheral" dealers 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 result in illiquid trading. Consistent with this prediction, Li and Schürhoff (2019) find that the marginal order intermediated by peripheral dealers commands a larger markup than the marginal order intermediated by central dealers. From Lemma 1, we know that for a small enough cost parameter c a dealer benefits from increasing its expertise as long as it does not prevent trading to remain liquid. Thus, the central dealer optimally reaches its expertise cutoff given its equilibrium level of order flow, whereas peripheral dealers are in position to use their limited resources to take advantage of any investor that trades with them. As a result, independent investors are 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 and order flow are limited — in both cases, investors collect an expected profit of $\frac{b}{2}$ as shown in subsection 2.1.

Proposition 2 thus highlights that in equilibrium the large investments in expertise the central dealer makes do not deter independent investors from sending their order flow its way. 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.

In the spirit of Green (2007), we can interpret this equilibrium as featuring a sophisticated (i.e., expert) dealer attracting the transaction volume of attentive (i.e., independent) investors and multiple less sophisticated dealers benefiting from the inattentiveness of their loyal investor clienteles. Yet, for sufficiently low c, investors are indifferent about which dealer to trade with in equilibrium because all dealers have $e_j \geq \bar{e}$. The central dealer provides liquid trading and earns a positive expected profit, whereas peripheral dealers provide illiquid trading and earn no profit.

We now analyze how the cost of expertise more generally affects the central dealer's investments in expertise, while taking into account the endogeneity of order flow. As discussed above, a dealer benefits from acquiring expertise through the better terms of trade investors offer and through the ability to make better decisions in response to these offers. In order to decide how much expertise to acquire, a dealer must then compare these benefits with the cost of expertise. The proposition that follows generalizes Proposition 2 to allow for $c \geq \frac{1}{2}(v_h - v_l)$.

Proposition 3. *In equilibrium, the central dealer's level of transaction-specific expertise is:*

$$e^* = \begin{cases} \bar{e} \equiv \frac{b}{2(v_h - v_l)} & \text{if} \quad c \leq \frac{1}{2}(v_h - v_l), \\ \underline{e} \equiv \frac{1}{2}\left(\frac{k}{\underline{n} + N}\right) & \text{if} \quad c > \frac{1}{2}(v_h - v_l). \end{cases} \tag{8}$$

This proposition shows that the central dealer's equilibrium level of expertise is weakly decreasing in the cost of expertise. This prediction may seem natural at first, but the analysis must account for how the central dealer's choice of expertise is complicated by investors' optimal bidding response. If investors expect their dealer to have a low level of expertise, they offer good terms of trade to ensure that the dealer accepts them with probability one, thereby maximizing the liquidity of trade. As a result, acquiring a marginally higher level of expertise solely benefits the dealer by improving the terms of trade and thereby increasing its share of the surplus. Having said that, acquiring substantially more expertise would, however, trigger the opposite response. When the dealer's transaction-specific expertise is sufficiently high, adverse selection concerns are too severe and investors offer worse terms of trade, resulting in illiquid trading. At that point, the benefits of higher expertise are guaranteed to be lower than the costs.

Figure 2 illustrates the benefits and costs of expertise acquisition for high and low levels of c. Panel (a) shows that when acquiring expertise is cheap, the central dealer maximizes its profit by

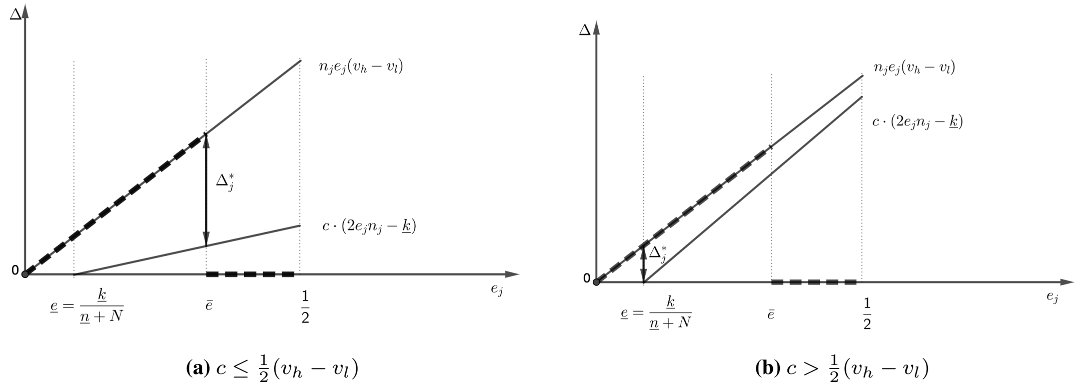

Figure 2

Benefits and Costs of Acquiring Expertise. Panel (a) compares the central dealer's gross profit from trade to the cost of expertise when the cost of expertise is low. Panel (b) compares the central dealer's gross profit from trade to the cost of expertise when the cost of expertise is high. In both cases, the central dealer's equilibrium net profit Δ_j^* as a function of its expertise e_j is the distance between the two plotted lines.

targeting the maximum level of expertise that still allows for liquid trading. The graph illustrates that in this case the dealer earns a sizable profit. When the expertise costs are high, however, the central dealer maximizes its profit by not acquiring any extra expertise. Panel (b) illustrates that in this case the central dealer still earns a positive profit, albeit smaller than in panel (a).

Overall, a central dealer may benefit from boosting its expertise capacity because it allows to receive better terms of trade from investors and to make better decisions. The dealer compares the resulting higher trade payoff with the cost of acquiring expertise. In equilibrium, the dealer acquires a level of expertise that does not repel the order flow from independent investors, who recognize that other (less popular) dealers would also take advantage of them since their limited expertise capacity is not spread as thinly as the central dealer. In all cases, the equilibrium can still be thought of featuring a central dealer providing high liquidity to a large pool of attentive investors and multiple peripheral dealers only attracting their loyal clienteles and trading less frequently with each of them, thereby providing less liquidity. Moreover, when the cost of expertise is high, independent investors are collecting strictly higher trading profits from their central dealer, who provide high liquidity, than loyal investors are collecting from their peripheral dealers, who provide low liquidity.

3 Extensions and Model Implications

We now extend the baseline model to discuss the robustness of our insights and to enrich their cross-sectional and time-series implications.

3.1 Refusing Order Flow

In our baseline model, dealers choose their level of expertise and investors choose the dealers that will receive their order flow. Once the total amount of order flow that investors send to dealer j is known, the dealer's expertise capacity is mechanically spread out among all transactions and trade occurs following investors' optimal bidding behavior. But what would happen if the central dealer was able to refuse order flow in order to keep its per-transaction expertise high?

From the baseline analysis, equation (5) shows that a dealer's total expected profit can be written as $\Delta(k_j, n_j) = n_j e_j(v_h - v_l) - c(k_j - \underline{k})$ as long as $e_j \leq \bar{e}$. Using $e_j = \frac{1}{2}\min\left(\frac{k_j}{n_j}, 1\right)$ and the fact that a central dealer will never find it optimal to acquire additional excess capacity that pushes $\frac{k_j}{n_j} > 1$, we can write $\Delta(k_j, n_j) = \frac{k_j}{2}(v_h - v_l) - c(k_j - \underline{k})$, which means that once the 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 bring more expertise to each accepted transaction, thereby earning higher profit per transaction. Yet, this effect would be neutralized by the reduction in the number of transactions from which the dealer extracts surplus. Moreover, by refusing order flow, this dealer would also risk violating the condition $e_j \leq \bar{e}$ that renders its expected profit negative (see equation (5)). 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.[4]

[4] For order-flow refusal 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 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.

3.2 Payment for Order Flow

In the equilibrium of our baseline model, peripheral dealers collect no profit from trading with their loyal investors, in contrast with models like Green (2007) where dealers take advantage of their loyal/naive clienteles. This difference is due to the fact that while by construction loyal investors cannot switch dealers in either setting, in our setting loyal investors can still worsen the terms of trade they offer if their local dealer happens to have concerning levels of expertise. As an implication of this result, if we were to allow 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 dealer expertise below the point where $e_j = \bar{e}$ and trading is liquid. Thus, subject to an intuitive parametric restriction, the concentration of order flow around a central dealer would survive in an environment where payment for order flow is allowed.

3.3 Ex Ante Dealer Heterogeneity

In our baseline model, the equilibrium allocation of order flow disproportionately favors one dealer even though all dealers are assumed to be homogeneous ex ante. When the cost of expertise is small enough, the central dealer finds it optimal to acquire more expertise than its peers. This prediction rationalizes why investors tend to (paradoxically) allocate their order flow to the most sophisticated dealers, despite standard adverse selection concerns. A natural limitation of our equilibrium predictions so far is that any of the ex-ante identical dealers can be selected by investors to become the market leader — our analysis has been silent about how independent investors may coordinate on picking a single dealer and how this chosen dealer knows 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 now allow for ex-ante heterogeneity in the per-unit cost of expertise acquisition and rank possible equilibria in terms of investors' welfare. We denote by c_j the cost of acquiring expertise for dealer j.

Proposition 4. *Independent investors weakly prefer an equilibrium in which the dealer with the highest cost of expertise c_j becomes the central one.*

Proposition 3 shows that the central dealer's level of expertise is weakly decreasing in c, while Corollary 1 shows that investors weakly prefer to trade with dealers with a lower e_j. It is therefore intuitively clear that independent investors are weakly better off when the central dealer is the one facing the highest cost of expertise.

3.4 Increasing Accessibility of OTC Markets

In our baseline model, we assume a market populated by many independent investors. In particular, we focus on the homogenous case where $\underline{k} < \frac{b}{v_h - v_l}\left(\underline{n} + \frac{N}{L}\right)$, implying that investors' total order flow (i.e., $L\underline{n} + N$) is large relative to dealers' total initial expertise capacity (i.e., $L\underline{k}$).

We now relax this restriction to shed light on recent trends in investors' accessibility to OTC markets. To do so, we analyze an equilibrium for the case with ex-ante homogenous dealers and $\underline{k} > \left(\underline{n} + \frac{N}{L} + 1\right)\left(\frac{b}{v_h - v_l}\right)$. When this condition holds, all dealers start with a level of expertise that would result in illiquid trading if independent investors were to allocate their transactions evenly among dealers. The following proposition demonstrates that an even allocation of order flow can now be observed in equilibrium.

Proposition 5. *If $\underline{k} > \frac{b}{v_h - v_l}\left(\underline{n} + \frac{N}{L} + 1\right)$, there exists an equilibrium in which all independent investors allocate their order flow evenly among dealers and no dealer acquires any additional expertise.*

Proposition 5 describes a market outcome in sharp contrast to the one discussed in the baseline analysis: we now have an equal allocation of order flow across all dealers. Independent investors know that all dealers will have informational advantages that result in illiquid trading, yet since such business model is prevalent, investors have no incentives to switch dealers.

Unlike when Assumption 3 holds, here none of the dealers expects to receive enough order flow to warrant acquiring additional expertise. The dispersed allocation of order flow in this equilibrium does not, however, benefit independent investors. That is, they are all trading with dealers who are so informed that their investors only collect $\frac{b}{2}$ of surplus in equilibrium, due to the illiquidity of trade.

Figure 3 illustrates how the informational advantage of the central dealer changes when the number of independent investors crosses the threshold:

$$\left[\underline{k}\left(\frac{v_h - v_l}{b}\right) - \underline{n} - 1\right] L. \tag{9}$$

Figure 4 highlights the dynamics of the central dealer's order-flow allocation. When the number of independent investors is small, an equal allocation of order flow across all dealers is sustainable in equilibrium. Panel (a) and (b) in Figure 3 show that $e^* > \bar{e}$ for small N, while Figure 4 shows that each dealer's order flow is $\underline{n} + \frac{N}{L}$. As the number of independent investors increases beyond the threshold above, this equilibrium is no longer sustainable. The order flow concentrates around a central dealer, which receives a total order flow $\underline{n} + N$ and either acquires enough expertise to have the most precise signal without impeding liquid trading (i.e., $e^* = \bar{e}$) or does not acquire any extra expertise (i.e., $e^* = \underline{e}$), depending on the cost of expertise.

This extension can thus shed light on the consequences of the increasing liberalization of financial markets. Our results suggest that a disproportionally faster growth in the number of investors than in the initial level of expertise for dealers (e.g., N growing faster than $\underline{k}$) could have led to more concentration of order flow. Moreover, this inflow of participants may have increased the central dealer's incentives to acquire expertise above and beyond what is or was standard in the industry. Yet, since the dispersed equilibrium described in Proposition 5 features $e^* > \bar{e}$ for all

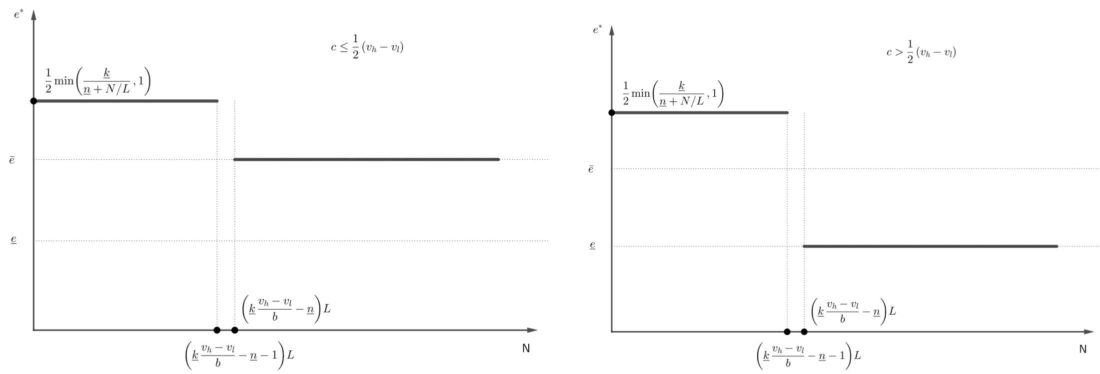

(a) $c \leq \frac{1}{2}(v_h - v_l)$ **(b)** $c > \frac{1}{2}(v_h - v_l)$

Figure 3

Impact of Number of Independent Investors on Central Dealer's Equilibrium Expertise. This figure plots the quality of the central dealer's information signal as a function of the number of independent investors, for two levels of expertise cost. Panel (a) shows the case with a low cost of expertise, while panel (b) shows the case with a high cost of expertise.

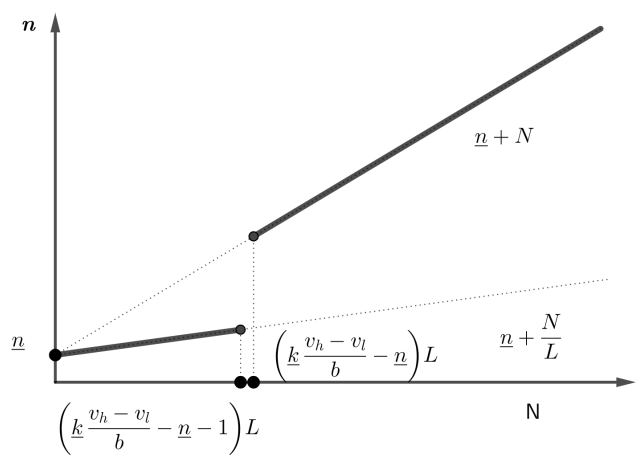

Figure 4

Impact of Number of Independent Investors on Central Dealer's Equilibrium Order Flow. This figure plots the transaction volume allocated to a central dealer as a function of the number of independent investors.

dealers whereas the concentrated equilibrium from Proposition 2 may feature $e_{j^*} \leq \bar{e}$ for some levels of expertise acquisition costs, independent investors may or may not have benefited from this liberalization. We formalize this result in the corollary that follows.

Corollary 2. *An increase in the number of independent investors from* $N < \left[\underline{k}\left(\frac{v_h - v_l}{b}\right) - \underline{n} - 1\right]L$ *to* $N > \left[\underline{k}\left(\frac{v_h - v_l}{b}\right) - \underline{n}\right]L$ *benefits independent investors if and only if* $c > \frac{1}{2}(v_h - v_l)$*.*

While new investors may greatly benefit from gaining access to OTC markets, our analysis in Corollary 2 shows that incumbent investors are made better off only if the costs of expertise are high enough. In particular, with a high enough c, the central dealer is unwilling to invest in expertise and reach $e = \bar{e}$ following an increase in OTC market accessibility. Yet, our analysis shows that potential advancements in technology, and the associated decreases in the costs of expertise dealers face, may have allowed a central dealer to extract a large fraction of the social surplus created by providing investors with easier access to OTC markets.

4 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, Matějka, 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. 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.

We show that, under intuitive conditions, the equilibrium allocation of order flow is concentrated toward one dealer that invests significant resources to gain informational advantages against its investors. Despite their adverse selection concerns 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 would 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 Omitted from the Text

Proof of Proposition 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:

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

□

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

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

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

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

□

Proof of Lemma 1: 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:

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

where $\alpha_j = \frac{1}{2} + e_j = \frac{1}{2} + \frac{1}{2}\min\left(\frac{k_j}{n_j}, 1\right)$ and b is from the range of parameters consistent with Assumption 1. We can substitute optimal bidding function from Proposition 1 to arrive at Equation (5). Then we can show that when $e_j < \frac{b}{2(v_h - v_l)}$:

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

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

Proof of Lemma 2: Suppose, by contradiction, that given all dealers' choices of expertise, independent investors allocate their order flow among at least two dealers, i.e., there exist dealers j' and j" such that $n_j^* > 0$ and $n_{j'}^* > 0$. Then, if $e_{j'} < e_{j''} < \frac{b}{2(v_h - v_l)}$ or $e_{j'} < \frac{b}{2(v_h - v_l)} < e_{j''}$, 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' has in this specific transaction with this investor, following Corollary 1. Thus, this allocation of order flow cannot be part of an equilibrium. If instead $\frac{b}{2(v_h - v_l)} > e_{j'} > e_{j''}$ or $e_{j'} > \frac{b}{2(v_h - v_l)} > e_{j''}$, the reverse is true, and similarly this order-flow allocation cannot be part of an equilibrium.

Now if $e_{j'} = e_{j''} < \frac{b}{2(v_h - v_l)}$, an independent investor who was planning to do business with either dealer is strictly better off switching because $\frac{k_{j'}}{n_{j'}+1} < \frac{k_{j''}}{n_{j''}-1}$ and $\frac{k_{j'}}{n_{j'}-1} > \frac{k_{j''}}{n_{j''}+1}$. Finally, if $e_{j'} \geq \frac{b}{2(v_h - v_l)}$ and $e_{j''} \geq \frac{b}{2(v_h - v_l)}$, then an independent investor is better off switching to the dealer(s) whose $e_j < \frac{b}{2(v_h - v_l)}$ (which exist(s) given the statement of the lemma). Hence, no matter what the dealers' expertise levels are, it is always optimal for a subset of independent investors to reallocate their order flow as long as the order flow of independent investors is allocated among two or more dealers and at least one dealer in the economy has $e_j < \frac{b}{2(v_h - v_l)}$. $\square$

Proof of Proposition 2: To show that the stated equilibrium is indeed an equilibrium, we first rule out that a dealer $j \neq j^*$ would acquire $k_j > \underline{k}$ of expertise. When $\underline{k} \geq \underline{n} + 1$, this dealer's e_j is already greater than $\frac{b}{2(v_h - v_l)}$, which means that this dealer's informational advantage creates so much adverse selection that investors find it optimal to offer prices that do not yield liquid trading (remember: dealers $j \neq j^*$ do not attract any order flow from independent investors). Thus, dealer

j's profit is zero and there is no benefit to acquiring more expertise.

We next rule out that an independent investor would leave dealer j^* and allocate order flow to a different dealer. If that was the case, this investor would go from trading with a dealer with $e_{j^*} = \frac{b}{2(v_h - v_l)}$ to trading with a dealer with $e_j = \frac{1}{2} \min\left(\frac{\underline{k}}{\underline{n}+1}, 1\right) > \frac{b}{2(v_h - v_l)}$. Since an investor's profit is equal to $\frac{b}{2}$ whenever its dealer's expertise satisfies $e > \frac{b}{2(v_h - v_l)}$, thus, there would be no benefit to switching dealers.

Moreover, when $c < \frac{1}{2}(v_h - v_l)$, dealer j^* finds it optimal to increase its expertise level until $e_{j^*} = \frac{b}{2(v_h - v_l)}$ (see Lemma 1). Once dealer j^* reaches an expertise bandwidth of $k_{j^*} = (\underline{n} + N)\,\frac{v_h - v_l}{b}$, the transaction-specific expertise reaches the maximum level that does not impede full trade. At that point, there would be no benefit to acquiring more expertise. Thus, the outcome described in the proposition is indeed an equilibrium.

Next, we show that an outcome where two dealers acquire expertise above $\underline{k}$ such that $\frac{k_j}{n_j} = \frac{k_j^*}{n_j^*}$ cannot be an equilibrium. That is, we cannot have multiple dealers receiving order flow from independent investors and acquiring enough expertise to take full advantage of their order flow. In such a situation, similar to the logic in Lemma 2, independent investors would have an incentive to go from dealer j to j^* or the other way around. By switching from dealer j to j^*, investors would increase the order flow going to dealer j^* above this dealer's expertise capacity, thereby spreading its resources further and decreasing the average quality of its signals, i.e., $\frac{k_j}{n_j - 1} > \frac{k_j^*}{n_j^* + 1}$.

Finally, Assumption 3 rules out as an equilibrium a situation where independent investors allocate their order flow equally among dealers that all have excess expertise capacity despite not making any investment, i.e., $\frac{k_j}{n_j} > \frac{b}{v_h - v_l}$ for all j. □

Proof of Proposition 3: Denote by $n_{j^*} = \underline{n} + N$ the total order flow that the dealer selected by independent investors, say j^*, expects to receive. Recall from Equation (5) that the dealer's profit is:

$$\Delta\left(k_{j^*}, n_{j^*}\right) = \begin{cases} -c \cdot (k_{j^*} - \underline{k}) & \text{if } e_{j^*} > \frac{b}{2(v_h - v_l)}, \\ n_{j^*} k_{j^*}(v_h - v_l) - c \cdot (k_{j^*} - \underline{k}) & \text{if } e_{j^*} \leq \frac{b}{2(v_h - v_l)}; \end{cases} \tag{A6}$$

or alternatively,

$$\Delta\left(k_{j^*}, n_{j^*}\right) = \begin{cases} -c \cdot (k_{j^*} - \underline{k}) & \text{if } k_{j^*} > n_{j^*} \frac{b}{v_h - v_l}, \\ n_{j^*} k_{j^*}(v_h - v_l) - c \cdot (k_{j^*} - \underline{k}) & \text{if } k_{j^*} < n_{j^*} \frac{b}{v_h - v_l}. \end{cases} \tag{A7}$$

Hence, $\Delta\left(k_{j^*}, n_{j^*}\right)$ is a piece-wise function of k_{j^*} for a given n_j^* with a jump at $k_{j^*} = n_{j^*} \frac{b}{v_h - v_l}$. The cost parameter c determines the e_{j^*} that the central dealer aims to achieve in equilibrium, for a given level of order flow n_{j^*}. In particular, this dealer's marginal benefit from acquiring more expertise is given by $\frac{\partial \Delta(k_{j^*}, n_{j^*})}{\partial k_{j^*}}$. If $k_{j^*} \leq n_{j^*}$ (thus: $e_{j^*} \leq \frac{1}{2}$), this marginal benefit can be written as:

$$\Delta\left(k_{j^*}, n_{j^*}\right) = \begin{cases} -c & \text{if } k_{j^*} > n_{j^*} \cdot \frac{b}{v_h - v_l}, \\ \frac{1}{2}(v_h - v_l) - c & \text{if } k_{j^*} < n_{j^*} \cdot \frac{b}{v_h - v_l}. \end{cases} \tag{A8}$$

If, however, $k_{j^*} > n_{j^*}$, the marginal benefit is:

$$\frac{\partial \Delta\left(k_{j^*}, n_{j^*}\right)}{\partial k_{j^*}} = -c. \tag{A9}$$

Hence, given the anticipated order flow of n_{j^*}, the optimal choice of $k_{j^*} \in \left[\underline{k}, n_{j^*} \cdot \frac{b}{v_h - v_l}\right]$. In order to find the optimal k_{j^*} we need to see if the c is above or below $\frac{1}{2}(v_h - v_l)$. If $c > \frac{1}{2}(v_h - v_l)$, then $\Delta(k_{j^*}, n_{j^*})$ as a function of k_{j^*} is globally decreasing, meaning that it is optimal for the dealer not to acquire any extra expertise. Then $e^* = \underline{e} \equiv \frac{1}{2}\left(\frac{\underline{k}}{\underline{n}+N}\right)$.

If $c \leq \frac{1}{2}(v_h - v_l)$, then the first segment of the function is increasing, meaning that it is optimal for the dealer to invest in expertise until $k_{j^*} = n_{j^*} \frac{b}{v_h - v_l}$. Then $e_j^* = \frac{b}{2(v_h - v_l)}$.

Altogether, the equilibrium features the following level of transaction-specific expertise for the central dealer:

$$e^* = \begin{cases} \underline{e} & \text{if} \quad c > \frac{1}{2}(v_h - v_l), \\ \frac{b}{2(v_h - v_l)} & \text{if} \ \ c \leq \frac{1}{2}(v_h - v_l). \end{cases} \tag{A10}$$

Since $\underline{e} < \frac{b}{2(v_h - v_l)}$ as implied by Assumption 3, e^* is weakly decreasing in c. □

Proof of Proposition 4: From Corollary 1 we know that independent investors weakly prefer to trade with a dealer with the smallest e_{j^*}, thereby minimizing exposure to the information advantage of a dealer. Since e_{j^*} is weakly decreasing in c_{j^*} according to Proposition 3, the higher the cost of acquiring expertise is for the dealer, the lower is the equilibrium level of informativeness of the central dealer. Hence, independent investors are better off if the central dealer is the dealer j with the highest c_j. □

Proof of Proposition 5: Consider a dealer j that receives $\frac{N}{L}$ trades from independent investors and does not acquire additional expertise, i.e., $k_j = \underline{k}$. The proposition statement implies that:

$$e_j = \frac{1}{2} \min \left(\frac{\underline{k}}{\underline{n} + \frac{N}{L}}, 1 \right) > \frac{b}{2(v_h - v_l)}, \tag{A11}$$

for all dealers. Hence, if an independent investor were to switch from dealer j to dealer j', dealer j''s transaction-specific expertise would become:

$$e_{j'} = \frac{1}{2} \min \left(\frac{\underline{k}}{\underline{n} + \frac{N}{L} + 1}, 1 \right) > \frac{b}{2(v_h - v_l)}. \tag{A12}$$

Hence, there is no incentive for any of the independent investors to deviate away from the current choice of dealers. Therefore, $k_j = \underline{k}$ for all j and independent investors allocating evenly their order flow among dealers is an equilibrium. □

Proof of Corollary 2: Before the increase in N, the equilibrium was consistent with Proposition 5, which means that all investors were trading with a highly informed dealer, i.e., $e_j > \frac{b}{2(v_h - v_l)}$ for all j. After the increase in N, Proposition 3 states that for a central dealer j^* that receives all independent investors' order flow, we have $e_{j^*} < \frac{b}{2(v_h - v_l)}$ as long as $c > \frac{1}{2}(v_h - v_l)$. Hence, independent investors benefit from a liberalization of OTC that reduces their dealer's transaction-specific expertise. Otherwise, if $c \leq \frac{1}{2}(v_h - v_l)$ we still have that $e_{j^*} = \frac{b}{2(v_h - v_l)}$ for the central dealer and $e_j = \frac{1}{2}\min\left(\frac{\underline{k}}{\underline{n}}, 1\right) > \frac{b}{2(v_h - v_l)}$ for all other dealers, thus none of the incumbent investors is made better off. □

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(5), 2381-2444.

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, 2022, "The Dark Side of Liquid Bonds in Fire Sales," Unpublished Working Paper, University of Oregon.

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," *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".

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

Hagströmer, Björn, and Albert J. Menkveld, 2019, "Information revelation in decentralized markets," *Journal of Finance* 74, 2751-2787.

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.

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

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.

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, 2018, "Endogenous Specialization and Dealer Networks," Unpublished Working Paper.

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.