Information Chasing versus Adverse Selection

Gabor Pinter* Chaojun Wang† Junyuan Zou‡

Bank of England The Wharton School INSEAD

March 1, 2022

Latest Version

Abstract

Contrary to the prediction of the classic adverse selection theory, a more informed trader could receive better pricing relative to a less informed trader in over-the-counter financial markets. Dealers chase informed orders to better position their future quotes and avoid winner's curse in subsequent trades. When dealers are perfectly competitive and risk averse, their incentive of information chasing dominates their fear of adverse selection. In a more general setting, information chasing can dominate adverse selection when dealers face differentially informed speculators, while adverse selection dominates when dealers face differentially informed trades from a given speculator. These two seemingly contrasting predictions are supported by empirical evidence from the UK government bond market.

JEL Classifications: G14, G18, D82

Keywords: Information chasing, adverse selection, over-the-counter, price efficiency

*Bank of England. Email: gabor.pinter@bankofengland.co.uk

†The Wharton School, University of Pennsylvania. Email: wangchj@wharton.upenn.edu

‡Finance Area, INSEAD. Email: junyuan.zou@insead.edu.

1 Introduction

The classic adverse selection theory predicts that more informed trades should receive worse pricing. However, this pattern could reverse in over-the-counter (OTC) financial markets—instead of being deterred by adverse selection risk, dealers actively chase informed orders by offering tighter bid-ask spreads to more informed traders.

We show that dealers chase informed orders to better position their future price quotes and avoid winner's curse in subsequent trades. On a multi-dealer trading platform, dealers' incentive to chase informed orders exactly offsets their fear of adverse selection. Through information chasing, dealers transform adverse selection by the informed into winner's curse when bidding for the uninformed. As a result, the adverse selection cost is entirely passed on to liquidity traders. More generally, without assuming any specific trading platform, we show that across differentially informed speculators, information chasing as a component of the bid-ask spread dominates the adverse selection component if and only if a more informed speculator receives a tighter bid-ask spread; Within a given speculator, however, adverse selection always dominates information chasing, so that a more informed trade always receives worse pricing than a less informed trade from the same speculator. These two predictions—which contrast sharply with each other—both find strong empirical support in the UK government bond market. Post-trade transparency reduces information chasing incentive and thus price efficiency.

The benchmark model works as follows. An asset with uncertain payoff is traded over-the-counter on a multi-dealer platform. In Stage 0, a speculator exerts costly effort to acquire a private signal about the asset payoff. In Stage 1, the speculator requests two-sided quotes for a selected quantity of the asset, without revealing her desired trade direction, from a number of dealers simultaneously on the multi-dealer platform. Every dealer quotes a bid and an ask to the speculator, who can then trade (buy or sell) with one dealer at that dealer's respectively quoted price. The trade is not publicly disclosed. In Stage 2, a mass of liquidity traders send quote requests to the dealers simultaneously on the multi-dealer platform to trade one unit of the asset each.

Dealers are incentivized to chase an informed order because executing such a trade allows a dealer to extract information about the asset payoff, then use this information to more accurately position its subsequent quotes to liquidity traders. If, say, the informed speculator chooses to sell to a given dealer j in Stage 1, then the asset payoff is likely to be low and Dealer j would lower its quotes to liquidity traders in Stage 2 to attract more buy orders, leaving undesired sell orders to the other dealers. This subjects the other dealers to winner's curse when competing for liquidity traders, forcing them to quote less aggressively thus allowing Dealer j to profit. While setting quotes to the informed speculator, dealers compete to narrow their bid-ask spread as long as the cost of being adversely selected does not exceed the expected gain from subjecting other dealers to winner's curse when bidding for liquidity traders. Therefore, through information chasing, dealers transform adverse selection by the informed into winner's curse when bidding for the uninformed. In equilibrium, the dealers all quote a zero bid-ask spread to the speculator in Stage 1, meaning that their incentive to chase the informed order exactly offsets their fear of adverse selection. When setting quotes to liquidity traders in Stage 2, dealers employ mixed strategies to mitigate winner's curse, giving rise to a new form of price dispersion. This type of price dispersion, induced by winner's curse, persists on a multi-dealer platform with simultaneous price competition and does not vanish even when the number of competing dealers goes to infinity or the signal about the asset payoff becomes perfectly accurate.

Direct price competition on a multi-dealer platform is not a prerequisite for information chasing. More generally, without assuming any specific trading platform, we show that across differentially informed speculators, information chasing dominates adverse selection if and only if a more informed speculator receives a tighter bid-ask spread. Within a given speculator, however, adverse selection always dominates information chasing, so that a more informed order always receives a wider bid-ask spread than a less informed order from the same speculator. The sharp contrast between these two predictions is due to an additional incentive compatibility condition required for the within-speculator comparison: a more informed speculator cannot pretend to be a less informed one and vice versa, while a given speculator can pretend to be less informed when she is actually more informed. For the within-speculator comparison, the resulting incentive compatibility condition is precisely sufficient and necessary for adverse selection to dominate information chasing.

These two predictions find strong support in the UK government bond market. To show this, we use a non-anonymous trade-level dataset which allows use to exploit both the across-client and within-client variations to test our predictions. First, we exploit the across-client variation in our data to show that, ceteris paribus, a more informed client receives on average 0.5 bps lower trading cost than a less informed client. Second, we exploit the within-client variation to show that informed clients face less favorable trading costs then their trades better predict future price movements. Importantly, we do not find evidence for such a positive relationship among less informed clients. In addition, we provide empirical evidence of dealers trading more profitably against uninformed clients as well as the inter-dealer-broker sector when giving more favorable execution costs to their informed clients. This result is suggestive of dealers acquiring valuable information from increased trading activity with informed clients and using this information to trade more profitably against the rest of its client base.

Regulators have been promoting post-trade transparency in the traditionally opaque OTC markets. FINRA and MSRB implemented real-time reporting and public dissemination of trades in corporate and municipal bonds via TRACE and RTRS since 2002 and 2005 respectively. After the 2008 financial crisis, the Dodd-Frank Act in the US expanded mandatory trade disclosures to swaps, while the more aggressive MiFID II Transparency Rules in EU cover a much wider range of fixed-income assets. In our model, trade disclosure after Stage 1 reduces information-chasing incentives thus harms information production and price efficiency. This prediction is shared with Banerjee, Davis and Gondhi (2018) and empirically supported by evidence in Lewis and Schwert (2018).

There are few theory papers demonstrating that a more informed trade may be receive better pricing. Naik, Neuberger and Viswanathan (1999), perhaps the closest paper to ours, shows that if a dealer is able to effectively "observe" the informativeness of a trade after executing the trade, then a more informed trade may receive better pricing. Our theory differs by explicitly modeling a dealer's inference of a trade's informativeness through the trader's identity and trade size. This approach is crucial in delivering the opposite predictions for our within- versus across-trader comparisons, thus providing guidance on where to locate empirical evidence of information chasing. Two empirical papers, Ramadorai (2008), Bjønnes, Kathitziotis and Osler (2015), document trading patterns that are consistent with information chasing in the foreign exchange market using independent data sources. However, those empirical patterns may also be consistent with non-informational mechanisms. Also based on data in the foreign exchange market, Hagströmer and Menkveld (2019) document that the most central dealer quotes a low bid-ask spread and learns quickest. This empirical finding is consistent with the information chasing mechanism. However, it is also consistent with the classic adverse selection theory, as pointed out by Hagströmer and Menkveld (2019). We follow our own empirical guidance and simultaneously find evidence of information chasing in the cross-trader comparison, and evidence of adverse selection in the within-trader comparison in the UK government bond market. These two opposite trading patterns provide a natural yet strong identification of the information-chasing mechanism.

There is a large literature on adverse selection in financial markets.[1] We contribute to this literature by showing that in OTC markets, dealers have an additional incentive to chase informed orders, which may very well dominate their fear of adverse selection. Consequently, we derive many pricing implications of information asymmetry that are opposite to the predictions of the classic adverse selection theory.

[1] Grossman and Stiglitz (1980), Glosten and Milgrom (1985), Kyle (1985, 1989) and Vives (2011) provide theoretical benchmarks; More recently, Lester, Shourideh, Venkateswaran and Zetlin-Jones (2018) and Chen and Wang (2020) develop dynamic models of market making under adverse selection risk. Two recent empirical papers, Collin-Dufresne, Junge and Trolle (2020a) and Collin-Dufresne, Hoffmann and Vogel (2020b), document empirical trading patterns that are consistent with the adverse selection theory in the index-CDS market and the FX Forward market respectively.

Our paper is related to the literature on information transmission in OTC markets.[2] We explicitly model a dealer's incentive to chase informed orders through aggressive pricing, which is the mechanism through which a dealer can learn and subsequently transmit information. The pricing implication of information-chasing incentive is the focus, while learning and transmission of information are merely natural consequences of information chasing. Although with a different focus, Golosov, Lorenzoni and Tsyvinski (2014) also endogenizes how information is transmitted through bilateral trading. Compared to their setting, our speculator's information precision is his private knowledge, which is signaled through the quantity he requests to trade: the better he is informed, the more he requests to trade. This rules out the possibility of entertaining an arbitrarily small trade, which is how uninformed traders elicit information from their trading partners in Golosov, Lorenzoni and Tsyvinski (2014). In our setting, dealers can only elicit information from the speculator at the expense of a non-trivial adverse selection cost. The trade off between the dealers' incentive to learn against their fear of adverse selection is the key determinant of our pricing implications.

Our model exploits two distinguishing features of an OTC market: *(i)* end-investors trade through dealers instead of directly with each other, and *(ii)* traders' identity are observed to each other before a price is formed, allowing dealers to price-discriminate based on the identity of the end-investor.[3] Abad, Aldasoro, Aymanns, D'Errico, Fache Rousová, Hoffmann, Langfield, Neychev and Roukny (2016) documents how a small number of dealers

[2] Duffie and Manso (2007), Duffie, Malamud and Manso (2009), Duffie, Giroux and Manso (2010), Duffie, Malamud and Manso (2014) and Golosov, Lorenzoni and Tsyvinski (2014) show how information percolates in OTC markets under different settings. Li and Song (2019) shows how a dealer can act as an information intermediary to channel information from informed to uninformed. With the exception of Golosov, Lorenzoni and Tsyvinski (2014), these papers separate information transmission and price formation so that adverse selection is assumed away. Empirically, Hagströmer and Menkveld (2019) showed that dealers' quotes are differentially informed and exhibits a stable correlation map, suggesting information transmission among dealers.

[3] Hau, Hoffmann, Langfield and Timmer (2021) provide evidence of discriminatory pricing against nonfinancial clients in the foreign exchange derivatives market. Lee and Wang (2018) shows that when a centralized exchange and an OTC market co-exist, the OTC dealers cream-skim liquidity traders from the exchange by offering them better pricing. Our paper considers OTC trading without an exchange running in parallel, which is the case for currencies and Treasury bonds. However, our theory would make the same prediction if both markets co-exist: When an exchange is available where trading prices are common knowledge, dealers no longer have incentive to chase informed orders in the OTC market. Therefore, adverse selection induces worse pricing for more informed speculators.

intermediate trades among a large number of buyside firms across various OTC markets.[4] More recently, a number of theoretical papers endogenize such a trading pattern.[5] We assume dealer intermediation as an exogenous feature of our OTC market, and focus on the pricing implication of asymmetric information in such a market structure. Glode and Opp (2016) provide an information-based mechanisms that endogenize intermediation chains. Our model may seem similar to Glode and Opp (2016) in its setup. The difference is, however, fundamental: in Glode and Opp (2016), an intermediary is exogenously informed, intermediating between a more informed trader and a less informed one. In our model, all dealers are uninformed *ex-ante*, and chase to become informed while providing intermediation. In this sense, Glode and Opp (2016) explains intermediation chains by assuming exogenously and heterogeneously informed agents, whereas we assume that dealers intermediate trades exogenously, and show how dealers take advantage of their role as intermediaries to become informed endogenously.

We offer a different view of how dealers become informed in OTC markets than Rüdiger and Vigier (2020), where dealers decide whether to acquire exogenous and verifiable information at a fixed cost. This approach can be interpreted as hiring research team in the same way that a speculator acquires information. Our setting departs from this three ways: a dealer learns indirectly through trading, the information is endogenously acquired by the speculator, and the precision of the information is private to the speculator. The first distinction—that dealers learn indirectly through trading—is more realistic in an environment where dealers are not allowed to directly generate its own signal. In practice, a bank holding company usually imposes a "China Wall" between its market making arm and its asset management arm, even before the Volcker Rule (which bans dealers from proprietary trading) was implemented. In addition to realism, our setting of learning endogenous information with privately chosen precision generates unique asset pricing implications.

[4]Markets for interest rate swaps, credit default swaps and foreign exchange forwards.

[5]Examples include, but are not limited to Chang and Zhang (2019, 2021), Hugonnier, Lester and Weill (2020), Neklyudov (2019), Sambalaibat (2018), Üslü (2019), Wang (2016), etc.

This paper is also related to the literature on information spillovers in trading. For example, Camargo, Kim and Lester (2016) and Asriyan, Fuchs and Green (2017) show how the transaction of one asset can reveal information about another asset with correlated value and affect the trading strategies of other market participants. In their papers, the value of information to other market participants does not enter the gains of the first trade. We focus on markets with little to none post-trade transparency in which an uninformed party privately learns from a trade with an informed party and subsequently makes profit from the private information by trading with other traders. In addition to being applicable to different market settings, this different trading environment leads to a further theoretical distinction: in our setup, the value of information to be materialized in subsequent trades enters the gains of the first trade, thus affects pricing of the first trade in a way that is not present in the literature on information spillovers.

The remaining of the paper is organized as follow: Section 2 sets up and solves the benchmark model with a multidealer platform. Section 3 derives conditions for information chasing to dominate adverse selection in a more general setting, without assuming any specific trading protocol. Section 4 provides empirical support for the testable predictions in the UK government bond market. Section 5 concludes.

2 The Benchmark Model

2.1 Setup

There are three types of risk-neutral agents—one speculator, n dealers ($n \geq 2$), and a mass m of liquidity traders—trading one common asset in the market. The asset payoff is v, which is either 1 or -1 with equal probability. Each liquidity trader needs to buy or sell, independently and with equal probability, one unit of the asset regardless of the price.[6]

[6]Since the asset price will be bounded between -1 and 1 in equilibrium, a sufficient condition for a liquidity trader to be willing to trade at any price is that she values the asset at $v + \delta$, where her liquidity benefit δ satisfies $|\delta| > 2$.

The trading game has three stages. In Stage 0, the speculator exerts costly effort to acquire information about the asset value v. Specifically, the speculator pays a cost $c(\eta)$ to acquire a binary signal s with a selected precision $\eta \in [0, 1]$. The binary signal s takes the value of 1 or -1 with equal probability, and returns the true asset value v with probability $(1 + \eta)/2$ ($\mathbb{P}(s = v) = (1 + \eta)/2$). We assume that the information acquisition cost function c satisfies

$$c(0) = 0, \quad c(1) = +\infty, \quad \lim_{\eta \to 0} c'(\eta) = 0, \quad \lim_{\eta \to 1} c'(\eta) = +\infty, \quad \text{and } c''(\eta) > 0$$

to insure a unique interior precision choice. The chosen precision and the signal realization are both private information of the speculator.[7] The dealers have no additional information about the asset value v, assigning equal probability to its potential values 1 and -1.

In this benchmark model, we assume that traders trade with the dealers on a multi-dealer platform using Request-for-Market (RFM) as the trading protocol, as follows. In Stage 1, the speculator requests two-sided quotes from the dealers simultaneously to trade a selected size $q \geq 0$ of the asset, without revealing her desired trade direction.[8] Since purchase and sale are symmetric in Stage 1, we consider the case where each given dealer j offers an ask $a_{1,j}(q)$ and a bid $-a_{1,j}(q)$ with a mid price equal to the unconditional mean of the asset, which is 0. Therefore, the dealer's pricing strategy in Stage 1 can be represented by its mid-to-bid spread $a_{1,j}(q)$ as a function of the order size q. The bid-ask quotes constitute a binding take-it-or-leave-it offer to buy or sell q units of the asset at the respective prices. The speculator can select one dealer to buy or sell at that dealer's respectively quoted price.

[7] Most existing work assume observable information precisions. In many settings, however, it is more realistic that the information acquisition cost function $c(\eta)$ is common knowledge, while the actual effort η that an investor exerts is not directly observable. Xiong and Yang (2020) reviews existing work in this modeling choice and further develops a model to study its effect.

[8] Such a request is called a "request-for-market" (RFM). In contrast, a "request-for-quote" (RFQ) indicates a desired trade direction upfront. It is common for traders to use RFM over the phone especially for larger trades. The trading protocol is also growing very quickly on electronic trading platforms. It is reported in Becker (2018) that the number of RFM-traded tickets on Tradeweb increases 510% in 2017 across interest rate swaps as traders try to hide their trading intentions. In the model, the speculator would choose to submit an RFM instead of an RFQ if she had a choice precisely to conceal her signal and thus incentivize dealers to chase her order.

There is no post-trade transparency, which means that the liquidity traders and the other dealers do not observe the direction or the execution price of the trade. In Stage 2, each liquidity trader requests a bid-ask quote $(a_{2,j}, b_{2,j})$ simultaneously from all dealers to trade 1 unit of asset. Since the liquidity trader's order is uninformed, it is irrelevant whether she indicates her desired trade direction at the time of her request. The liquidity trader then trades with the dealer who offers the best quote. At the end of Stage 2, the asset payoff is realized. The speculator and the dealers receive the realized payoff of their asset position plus the net payments they received from trading. The timeline and the market structure is summarized in Figure 1.

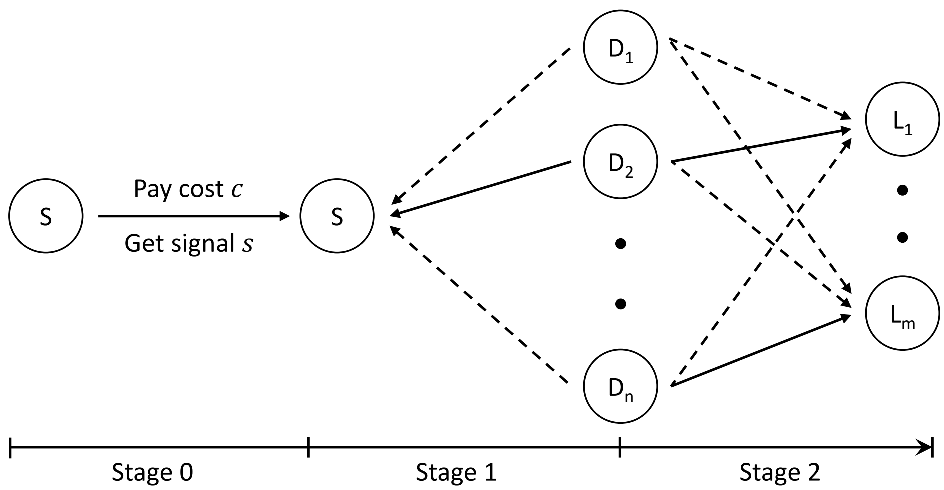

Figure 1: The Timeline.

In Stage 1, the speculator naturally chooses to trade with the dealer who offers the lowest spread. If the lowest spread is offered by multiple dealers, the speculator is indifferent toward trading with any of these dealers. We focus on equilibria in which the speculator randomly selects one of those dealers offering the lowest spread, independently from the realization of her private signal. This is not a restriction on the speculator's strategy, but rather a property that we require an equilibrium to satisfy. This property can be guaranteed by any tie-breaking rule that ranks the dealers in certain preference order in the event of a tie.

2.2 Equilibrium

Using backward induction, we show that the benchmark model has a unique perfect Bayesian equilibrium that satisfies forward induction.

Stage 2: Competing for Liquidity Trades

In equilibrium, it will turn out to be the case that at the beginning of Stage 2, one dealer is able to perfectly infer the signal realization by executing the speculator's trade in Stage 1 while other dealers remain uninformed. The Stage-2 game is equivalent to a first-price-sealed-bid auction with asymmetric information and discrete signals. In the appendix, we show that the corresponding continuation game in Stage 2 has no pure-strategy equilibrium. Intuitively, the uninformed dealers use a mixed pricing strategy to avoid being completely outbid by the informed dealer precisely when the asset is good, mitigating winner's curse. The informed dealer also mixes to avoid being completely outbid by the uninformed dealers.

Lemma 1 summarizes dealers' unique bidding strategies in Stage 2 constructed from results established in Syrgkanis, Kempe and Tardos (2019). We denote a dealer's belief regarding the signal precision by $\hat{\eta}$, which will be uniquely pined down by forward induction as a function of the order size q. Since all dealers observe q, they hold the same belief $\hat{\eta}$. We denote the bid-ask quotes of the informed dealer by (b_2^+, a_2^+) when $s = 1$, and (b_2^-, a_2^-) when $s = -1$. The bid-ask quotes of an uninformed dealer are denoted by (b_2^0, a_2^0).

Lemma 1 *In Stage 2, with one informed dealer, there is an equilibrium where*

(i) the informed dealer quotes $b_2^- = -\hat{\eta}$, $a_2^+ = \hat{\eta}$, and draw b_2^+ and $-a_2^-$ from a continuous distribution with CDF

$$G^+(b) = \frac{2}{1 - b/\hat{\eta}} - 1, \quad b \in [-\hat{\eta}, 0].$$

(ii) the uninformed dealers draw b_2^0 and $-a_2^0$ from a hybrid distribution with CDF

$$G_n(b) = \sqrt[n-1]{\frac{1}{1-b/\hat{\eta}}}, \quad b \in [-\hat{\eta}, 0].$$

The payoffs and the distribution of the winning bid in any other stage-2 equilibrium are the same as in the stated equilibrium above.

The distribution G_n describes a hybrid bidding strategy, in that an uninformed dealer bids $-\hat{\eta}$ with probability $\sqrt[n-1]{1/2}$, and otherwise draws its bid from the distribution with CDF

$$\frac{\sqrt[n-1]{\frac{1}{1-b/\hat{\eta}}} - \sqrt[n-1]{\frac{1}{2}}}{1 - \sqrt[n-1]{\frac{1}{2}}}, \quad b \in [-\hat{\eta}, 0].$$

When $n = 2$, the distribution G_2 of the uninformed dealer's bid is the same as the unconditional distribution of the informed dealer's bid. That is, the uninformed dealer "fakes" a signal by randomly flipping a coin, and bids according to the fake signal as if it was informed. When $n > 2$, the maximum bid of all the $n-1$ uninformed dealers is distributed following the CDF $G_n^{n-1} = G_2$, which is not affected by n. Consistently, the informed dealer's bidding strategy is also not affected by the number $n-1$ of competing uninformed dealers.

From dealers' bidding strategies, we can compute their Stage-2 payoffs.

Lemma 2 *In Stage 2, with one informed dealer, the expected payoff of an uninformed dealer is* 0*, and the expected payoff of the informed dealer is $m\hat{\eta}/2$.*

The uninformed dealers shade their bid-ask offers due to their fear of winner's curse, allowing the informed dealer to earn a positive profit. The value of being the only informed dealer is increasing in the mass of liquidity traders and the precision of the signal.

Given the prospect of earning a positive payoff in Stage 2 if informed, dealers are incentivized to chase the speculator's order in Stage 1.

Stage 1: Chasing the Informed Order

While setting quotes to the informed speculator in Stage 1, dealers compete to narrow their bid-ask spread until the cost of being adversely selected is about to exceed the expected gain from being able to more accurately position their quotes to liquidity traders in Stage 2. Bertrand competition implies that at least 2 dealers offer the competitive spread $a_1(q)$ satisfying the following zero profit condition:

$$q\left[a_1(q) - \hat{\eta}(q)\right] + \frac{1}{2}m\hat{\eta}(q) = 0, \tag{1}$$

where $\hat{\eta}(q)$ is dealers' belief about η given an order size q. We will pin down $\hat{\eta}(q)$ by forward induction when we solve for the speculator's choice of information acquisition in Stage 0. Therefore, given any order size q from the speculator, two or more dealers quote the same competitive mid-to-bid spread in Stage 1:

$$a_1^*(q) = \hat{\eta}(q) - \frac{m\hat{\eta}(q)}{2q}. \tag{2}$$

Upon selecting a dealer offering the best price, the speculator buys if she receives a positive signal, and sells otherwise.

The pricing function in (2) reflects the combined effect of two countervailing incentives—the fear of adverse selection and the urge of information-chasing. The first term of the spread $a_1^*(q)$ is a dealer's expected per-unit value of the asset when she receives an order of size q. Since the informed speculator always trades in the direction that is adverse to the dealer, the dealer charges the speculator this expected asset value through the spread to compensate for its expected loss from the trade. This is the classic adverse selection component of a bid-ask spread. When the speculator's information becomes more precise, dealers widen their bid-ask spread to protect themselves from the increasing adverse selection cost. The second term in $a_1^*(q)$ reflects dealers' incentive to chase informed orders. A dealer can profit from its information advantage over other dealers when competing for liquidity traders' orders in Stage 2. Anticipating this benefit, all dealers narrow their bid-ask spread to compete for the informed order in Stage 1. Given an order size q, the existence of more liquidity traders in Stage 2 gives dealers stronger incentive to chase the informed order in Stage 1 and narrows dealers' bid-ask spread.

The sign of the spread in (2) depends the relative strength of these two countervailing incentives. If $q > m/2$, the per unit value of information is smaller the associated adverse selection cost. Thus, The mid-to-bid spread $a_1^*(q)$ is positive. The reverse is true when $q < m/2$.[9] When $q = m/2$, the two incentives exactly offset each other, which will turn out to be the case in equilibrium. The order size q reveals the information acquisition effort of the speculator, and will be pinned down endogenously when we examine Stage 0.

Stage 0: Information Acquisition

Given any order size q and pricing strategies $(a_{1,j})_{j=1,\ldots,n}$ of the dealers, the speculator's optimal choice of information precision maximizes her expected payoff

$$\hat{\eta}(q) \in \underset{\eta}{\operatorname{argmax}} \; q \cdot \left[\eta - \min_j a_{1,j}(q)\right] - c(\eta). \tag{3}$$

In the speculator's payoff, the first term represents the expected profit of trading q units of the asset in the direction indicated by the signal at the best executable quote. The second term represents the cost of information acquisition.

The spreads $(a_{1,j}(q))_{j=1,\ldots,n}$ cannot depend on η since dealers cannot observe the speculator's actual choice η of information precision. Thus, when choosing η, the speculator need not consider the dealers' pricing functions. For a given trade size q, it is thus a dominant

[9] In this paper, we abstract away from other market-making costs such as inventory and operational costs. Thus, we only capture the informational component of a bid-ask spread. A negative spread should be interpreted as a negative informational component in a positive spread.

strategy for the speculator to choose the precision

$$\hat{\eta}(q) = c'^{-1}(q) \tag{4}$$

that equates the marginal benefit from trading with a more precise signal and the marginal cost of acquiring information. Forward induction thus implies that the dealers hold the same belief $\hat{\eta}(q)$ regarding the precision when receiving an order of size q. This belief $\hat{\eta}(q)$ is increasing in q because dealers understand that the speculator must have acquired a more precise signal if she requests to trade a larger size. From the speculator's perspective, she can always credibly signal her information acquisition effort through her order size.

In Stage 0, the speculator, anticipating some equilibrium pricing function a_1^* by the dealers, chooses η and q jointly to maximize its expected payoff:

$$(\eta^*, q^*) = \underset{\eta, q}{\operatorname{argmax}} \; q[\eta - a_1^*(q)] - c(\eta). \tag{5}$$

Now we can solve for the equilibrium of the game in Stage 0 and Stage 1, taking the payoff in the Stage 2 bidding game as given.

Definition 1 (Equilibrium) *A PBE of the 3-stage trading game consists of (i) the speculator's strategy (η^*, q^*), (ii) the dealers' pricing strategy a_1^* in Stage 1, and (iii) dealer's belief $\hat{\eta}(q)$ that satisfy*

(i) speculator's optimality condition (5),

(ii) dealers' zero profit condition (2), *and*

(iii) the forward induction condition (4).

Substituting dealers' equilibrium belief (4) into their zero profit condition (2), we obtain dealers' equilibrium spread offered to the speculator:

$$a_1^*(q) = c'^{-1}(q)\left(1 - \frac{m}{2q}\right). \tag{6}$$

Using the one-to-one relationship $\eta = \hat{\eta}(q) = c'^{-1}(q)$ between the optimal choices of η and q, and plugging in the expression (6) of the equilibrium spread $a_1^*(q)$, we can simplify the speculator's problem (5) into a one dimensional optimization problem over η,

$$\max_{\eta} \frac{m\eta}{2} - c(\eta). \tag{7}$$

Solving the optimization problem yields

$$\eta^* = c'^{-1}\left(\frac{m}{2}\right), \quad q^* = \frac{m}{2}. \tag{8}$$

Proposition 1 *The PBE of the 3-stage trading game can be described as follows:*

1. *In Stage 0, the speculator acquires information with precision $\eta^* = c'^{-1}(m/2)$.*

2. *In Stage 1, the speculator sends dealers a trade request of size $q^* = m/2$. At least two dealers quote a mid-to-bid spread according to (6), while other dealers quote equal or larger spreads. The speculator randomly chooses to trade with one dealer with the lowest spread, and she buys (sells) if her private signal in Stage 0 is positive (negative).*

3. *In Stage 2, the trading strategy of the dealers are described by Lemma 1.*

In equilibrium, the size of the informed order and the signal precision both increase in the mass m of liquidity traders. Intuitively, a larger amount of liquidity trades raises the profit of offering informed quotes, thus intensifies dealers' incentive to chase the informed order in Stage 1. Therefore, dealers shrink their bid-ask spread to the informed order, which in turn encourages the speculator to acquire more precision information and trade more.

Plugging the equilibrium size $q^* = m/2$ of the informed order into the dealer's Stage-1 pricing strategy (2), the dealers' equilibrium spread quoted to the speculator thus becomes

$$a_1^*(q^*) = \eta^* \left(1 - \frac{m}{2q^*}\right) = 0. \tag{9}$$

This zero spread result holds for any parametric assumption in the benchmark model. Without other trading frictions such as search frictions, inventory costs or transaction costs, Proposition 1 should be interpreted as a zero informational component in the bid-ask spread. Dealers trade off two opposing incentives when setting quotes to the speculator: their fear of adverse selection drives up their bid-ask spread, while their urge of information chasing pushes down the spread. On a multi-dealer trade platform, these two countervailing forces precisely offset each other, rendering a zero net effect of information on the spread. We will show, in a generalized model in Section 3, that this is a consequence of the market structure allowing dealers to compete directly in their pricing for liquidity trades.

2.3 Pricing Implications

The model has several testable implications on trading prices.

Bid-Ask Spreads Since liquidity traders as a whole place the same amount of sell orders and buy orders, There is no net asset transfer between dealers and liquidity traders. Thus, the average mid-to-bid spread in Stage 2 is equal to the dealers' trading profit per unit of liquidity orders. Liquidity traders face a positive expected mid-to-bid spread given by

$$\frac{1}{2}\,\eta^* = \frac{1}{2}\,c'^{-1}\left(\frac{m}{2}\right). \tag{10}$$

Comparing the expected bid-ask spread received by the informed trader versus the liquidity traders, we have the following testable implication.

Claim 1 *In OTC markets with non-anonymous trading, informed trades receive lower bid-*

ask spreads.

We can also calculate the expected bid-ask spread of all trades in both stages, weighted by their trade sizes:

$$\bar{\Delta} = \frac{q^* \Delta_1 + m\Delta_2}{q^* + m} = \frac{2}{3} c'^{-1} \left(\frac{m}{2} \right).$$

The average bid-ask spread is increasing in the amount of liquidity traders.[10]

Claim 2 *In OTC markets with non-anonymous trading, other things equal, the average bid-ask spread is larger when there are more liquidity traders.*

Claim 3 *In OTC markets with non-anonymous trading, other things equal, bid-ask spread is smaller when the cost of information acquisition is uniformly higher.*

Price Dispersions An important feature of the equilibrium is that price dispersion arises endogenously as a result of winner's curse. Without any search frictions, both the informed dealer and the uninformed dealers use mixed pricing strategies when competing for liquidity trades in Stage 2.

The price dispersion arising from winner's curse persists even when the liquidity trader has access to a large number of dealers ($n \to \infty$), and even increases when the signal of the asset payoff becomes more accurate ($\eta \to 1$). This is because dealers' pricing functions are linearly scalable in the signal precision η (Lemma 1). Letting $\sigma(\eta)$ denote some homothetic measure of price dispersion for trades in Stage 2 with a given signal precision η, then

$$\sigma(\eta) = \eta\sigma(1). \tag{11}$$

Since the equilibrium signal precision η^* increases with greater mass m of liquidity traders and lower margin cost of information acquisition, we obtain the following prediction.

[10]The comparative static doesn't change if we use different weights.

Claim 4 *In OTC market with non-anonymous trading, other things equal, price dispersion is higher when there are more liquidity traders, and when the marginal cost of information acquisition is lower.*

Price Informativeness Since the transaction price in Stage 1 is always 0, it carries no information about the asset's common value v. We define price informativeness as the proportion of variance in the asset value v explained by the observed trading prices in Stage 2. Depending on the realization of the speculator's signal s, the best bid and ask in Stage 2 follow different distributions (Lemma 1). An econ metrician can therefore precisely estimate the signal s from a large sample of transaction prices in Stage 2. The price informativeness thus equals the fraction of variance in v explained by the speculator's signal s.

$$\tau(\eta) = 1 - \frac{\mathrm{Var}[v|s]}{\mathrm{Var}[v]} = \eta^2. \tag{12}$$

Claim 5 *In OTC markets with non-anonymous trading, other things equal, price informativeness is higher when there are more liquidity traders, and when the marginal cost of information acquisition is lower.*

3 General Trading Protocols

In the benchmark model with direct price competition among dealers, information chasing exactly offsets adverse selection, resulting in the zero bid-ask spread received by the speculator. Further, the speculator's trade size doesn't depend on her information acquisition technology. In this section, we study, in a general setting without assuming any specific trading platform, how the speculator's trade size and bid-ask spread vary with the informativeness of her trade. We give a sufficient and necessary condition under which information chasing dominates adverse selection and vice versa. When information chasing dominates, a more informed trader receives a tighter bid-ask spread and trades a smaller size.

In the previous section, dealers compete on the multi-dealer platform for liquidity trades in Stage 2. Direct price competition determines the informed dealer's profit from its information. Now, we generalize the Stage 2 game by assuming that the informed dealer receives some reduced-form continuation payoff $V_I(\eta)$ from exploiting the information content of the speculator's trade.[11] Then, $V_I(\eta)$ is also the total surplus of the trade between the speculator and the dealer. Also, to generalize the assumption that dealers compete *à la Bertrand* in Stage 1, we now assume that a fraction $\varphi \in [0,1]$ of the trading surplus $V_I(\eta)$ goes to the speculator, while the dealer executing the trade gets the remaining $1-\varphi$ fraction. The split of the trading surplus can be viewed as the outcome of bilateral bargaining in a trade between the dealer and the speculator, with the case $\varphi = 1$ corresponding to Bertrand competition by dealers. Since the total trade surplus and the split of the surplus are both given in reduced form, the number of dealers becomes irrelevant. For simplicity, we will view the generalized model as one dealer trading with one speculator. In terms of the signal distribution, we assume that the expected unit value of the asset is $v(\eta)$ or $-v(\eta)$ conditional on a positive or negative signal respectively.

Therefore, the benchmark model is a special case of the generalized model with

$$v(\eta) = \eta, \quad V_I(\eta) = \frac{1}{2}m\eta, \quad \varphi = 1. \tag{13}$$

We impose the following regularity conditions on $V_I(\eta)$ and $v(\eta)$.

Assumption 1 *The functions $v(\cdot)$ and $V_I(\cdot)$ are both twice differentiable, and*

1. $v(0) = 0$, $V_I(0) = 0$, $v'(\eta) > 0$, $V_I'(\eta) > 0$, $c'(\eta) > 0$;
2. $v''(\eta) \leq 0$, $c''(\eta) > 0$, $\phi V_I''(\eta) - c''(\eta) < 0$.

It follows from (13) that Assumption 1 is a generalized version of the previously assumed

[11] Brancaccio, Li and Schürhoff (2020) estimates that in the US municipal bond market, the average value of information is worth 7 bps, providing empirical support for the premises of our mechanism.

differentiability and convexity of $c(\cdot)$. It guarantees that the model has a unique interior equilibrium.

The generalized model can be solved in the same way as the benchmark model. Receiving a trade order of size q, the dealer expects the speculator to have chosen the dominant information precision $\hat{\eta}(q)$ that equalizes the marginal change in the total value of the order and the marginal cost of information acquisition.

$$qv'(\hat{\eta}) = c'(\hat{\eta}). \tag{14}$$

The speculator receives a total payoff of $\varphi V_I(\hat{\eta})$ in the form of price discount. This replaces the dealer's zero-profit condition (2) in the benchmark model. The dealer's pricing function then becomes

$$a_1^*(q) = v(\hat{\eta}(q)) - \frac{\varphi V_I(\hat{\eta}(q))}{q}. \tag{15}$$

Taking the dealer's belief into consideration, the speculator optimally chooses the information precision η^* to equalize her share of the marginal value of information and the marginal cost of information acquisition.

$$\varphi V_I'(\eta^*) = c'(\eta^*). \tag{16}$$

This determines the equilibrium level of information precision η^*. Combining (14), (15) and (16), we establish the equilibrium relationship among information precision, order size and the bid-ask spread.

Proposition 2 *In the generalized model, there exists a unique equilibrium in which*

$$q^* = \varphi \left. \frac{dV_I}{dv} \right|_{\eta=\eta^*}, \tag{17}$$

$$a_1^*(q^*) = v(\eta^*) \left[1 - \frac{1}{\varepsilon(\eta^*)}\right], \quad \textit{where } \varepsilon(\eta) = \frac{d \ln V_I}{d \ln v}. \tag{18}$$

Here, $\varepsilon(\eta)$ measures how the percentage change in the value v of the asset affects the value of information V_I in percentage. Thus, it is the *elasticity* of V_I with respect to v. The key to understand the intuition of Proposition 2 is the trade size q. From (15), we know that the spread is the difference between the value of one unit of asset and the value of information *distributed to each unit of asset*. Given the same asset value v, if the speculator trades a larger quantity q, the value of information per unit of asset will be diluted more, and the spread will be larger. The tipping point is $q = \varphi V_I(v)/v$, when the adverse selection component and the information chasing component exactly offset each other. How is q determined in equilibrium? Equation (17) shows that in equilibrium q always equals the marginal value of information captured by the speculator $\varphi V_I'(v)$. Suppose the speculator trades $q > \varphi V_I'(v)$. To make sure the trade is placed in the right direction, the speculator has to acquire information to the point that the marginal cost of information acquisition equals q, which exceeds the marginal value of information. This means that the speculator is acquiring too much information, and at the same time, trading too much. The same reasoning can be used to show the sub-optimality of $q < \varphi V_I'(v)$. Now we only need to compare the equilibrium trade size $\varphi V_I'(v)$ with the tipping point size $\varphi V_I(v)/v$. It turns out that this comparison is equivalent to comparing ε evaluated at the equilibrium information precision to 1. The equilibrium bid-ask spread is positive if $\varepsilon(\eta^*) > 1$, and negative if $\varepsilon(\eta^*) < 1$. In the benchmark model, both $V_I(\eta)$ and $v(\eta)$ are linear function of η, so the elasticity is exactly equal to 1. Therefore, the bid-ask spread for the speculator always equals 0 in the benchmark model.

We also gives a graphical illustration of Proposition 2 in panel (a) and (b) of Figure 2.

In each panel, we plot the speculator's share of the value of information φV_I and the cost of information acquisition c as a function of v. The two dashed tangent lines marks the value of v such that $\varphi V_I(v)$ and $c(v)$ have the same slope. This is the equilibrium unit cost $v(\eta^*)$ of adverse selection, following the speculator's optimal information acquisition condition (16). The forward induction condition (14) implies that the common slope of $c(v)$ and $\varphi V_I(v)$ at $v(\eta^*)$ is the equilibrium order size q^*. The equilibrium mid-to-bid spread $a_1^*(q^*)$ can be decomposed into two parts—an adverse selection component and an information chasing component, as shown by the expression (15) of $a_1^*(q^*)$. The adverse selection component equals $v(\eta^*)$, the absolute deviation between of the asset value's ex-post mean and its ex-ante mean. The information chasing component measures the value of information *per unit* captured by the speculator. In panel (a), the elasticity ε of V_I with respect to v is greater than 1 at $v(\eta^*)$. This is equivalent to say that the instantaneous rate of change of $V_I(v)$ at $v = v(\eta^*)$ is greater than the average rate of change of $V_I(v)$ between $v = 0$ and $v = v(\eta^*)$. Therefore, the speculator trades a large quantity q^* such that the value of information per unit is smaller than the asset's value $v(\eta^*)$. The adverse selection component dominates the information chasing component, resulting in a positive bid-ask spread. In panel (b), the elasticity ε of V_I with respect to v is smaller than 1 at $v(\eta^*)$. In contrast to the first case, the speculator trades a small quantity q^* such that the value of information per unit is greater than the asset's value $v(\eta^*)$. As a result, the information chasing component dominates the adverse selection component, resulting in a negative bid-ask spread.

Across-speculator heterogeneity of trade size and bid-ask spread One sufficient condition for ε to be always greater than 1 is that V_I is a convex function of v. In fact, the convexity of V_I in v has important implications for understanding the cross-sectional pattern of order sizes, bid-ask spreads and information content of trading. Consider two cost functions of information acquisition $c_1(\cdot)$ and $c_2(\cdot)$ which satisfy the regularity conditions in Assumption 1. We say that it is more costly for a speculator to acquire information under

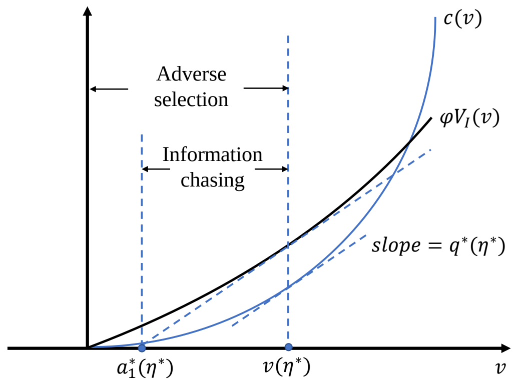

(a) Convex $V_I(v)$, high $c(v)$.

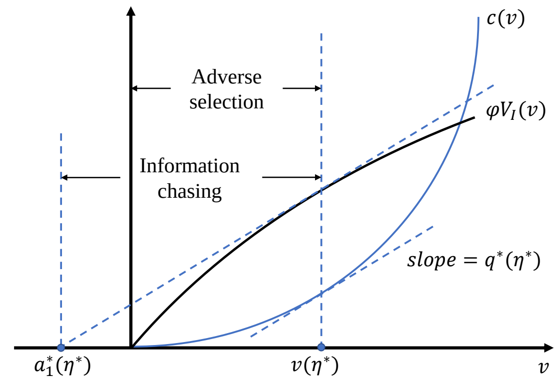

(b) Concave $V_I(v)$, high $c(v)$.

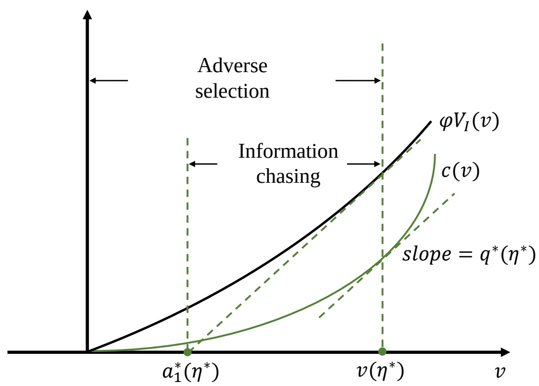

(c) Convex $V_I(v)$, low $c(v)$.

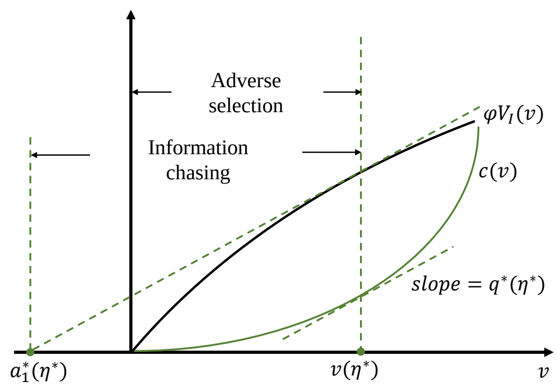

(d) Concave $V_I(v)$, low $c(v)$.

Figure 2: Decomposition of the bid-ask spread: information chasing vs adverse selection.

$c_1(\cdot)$ if $c_1'(\eta) > c_2'(\eta)$ for any $\eta \in [0, 1]$.

Proposition 3 (Across-speculator trade heterogeneity) *A speculator with lower cost of information acquisition always chooses a higher precision η^*. Moreover,*

- *If V_I is convex in v, a speculator with a lower cost of information acquisition trades a higher quantity q^* and receives a higher half spread $a_1^*(q^*)$.*
- *If V_I is concave in v, a speculator with a lower cost of information acquisition trade a lower quantity q^* and receives a lower half spread $a_1^*(q^*)$.*
- *If V_I is linear in v, all speculators trade the same quantity q^* and receives a zero bid-ask spread regardless of their cost of information acquisition.*

The proof of Proposition 3 can be found in the appendix. It is intuitive that a speculator with lower marginal cost of information acquisition acquires more information. Why does the ranking of trade size and bid-ask spread depend on the convexity of $V_I(v)$? Again, we need to start from understanding the ranking of trade size q. As we have shown previously, it is optimal for a speculator to trade q which equals the marginal value of information. If $V_I(v)$ is convex, the speculator who acquires more information has higher marginal value of information, therefore trades higher q. In fact, the trade size is so large such that the increased value of information, after being diluted by q, is smaller than the increased value of asset. Thus, the information chasing component becomes relatively weaker compared to the adverse selection component, and the bid-ask spread increases. The opposite holds when $V_I(v)$ is concave in v.

Here we give a graphic illustration in Figure 2. Panel (c) and (d) depict the equilibrium when the cost of information acquisition is uniformly lower than that in panel (a) and (b). No matter whether V_I is convex or concave, the speculator increases the level of information acquisition in response to the decline in the cost of information acquisition to capture more value of information. However, the change in q^* and $a_1^*(q^*)$ are different in the two cases. When V_I is a convex function of v, lower information acquisition induces the speculator to aggressively improve the information precision and signal this information precision with a higher q^*. The increase in q^* further spreads out the value of information and enlarges the difference between the dominating adverse selection component and the information chasing component, resulting in a larger bid-ask spread. On the contrary, when V_I is a concave function of v, the same level of decline in the cost of information acquisition only leads to a mild increase in the optimal information precision η^*. Because the information acquisition cost is lower, the speculator can signal this higher information precision with a lower q^*. Therefore, value of information becomes more concentrated in a smaller amount of traded asset and further dominates the adverse selection components. In equilibrium, the speculator receives a more negative bid-ask spread compared with panel (b).

Proposition 3 has important implications for identifying informed orders in OTC markets. Conventional wisdom generally believes that better informed traders trade larger quantity in financial markets. We show that this conventional wisdom can be reversed when information chasing effect dominates the classical adverse selection effect. This is indeed relevant in opaque markets without post-trade transparency where dealers can profit from their private information gathered from informed orders.

Within-speculator heterogeneity of order size and bid-ask spread Up till now we have focused on the heterogeneity of order size and bid-ask spread originated from speculator's heterogeneity. In fact, in the data it is quite usual to observe that the same speculator trades different quantities at different spreads. Now we relax the binary signal structure to account for this within-client trade heterogeneity.

Suppose the speculator observes a private signal x with a symmetric c.d.f $F(x)$ after incurring the information acquisition cost $c(\eta)$. The ex-post value of the asset not only depends on the precision η, but also depends on the signal x. Without loss of generality, we assume the value of asset $v(\eta, x)$ is an increasing function of both η and x. To understand this assumption intuitively, we can think of η as the preditive power of the speculator's quantitative model, and x as the predictied value based on the model. The deviation of the speculator's value from the market expectation increases in the quality of the speculator's model and the innovation predicted by the model. For the value of information, we maintain the same assumption that V_I is an increasing function of $v(\eta, x)$.

Here the speculator has two dimensions of private information, the precision η and the signal x. However, after the speculator has choosen η, the only unobservable variable that matters for the trading profit for both the speculator and the dealers is the value of the asset $v(\eta, x)$. The trading game is essentially a Bayesian game with one-dimentional private information on v. By the revelation principle, the Bayes-Nash equilbrium can be described by a quantity fundtion $q(v)$ and half-spread function $a(v)$. The trading payoff of the speculator, gross of the cost of information acquisition, is given by

$$V_S(v) = [v - a(v)]q(v). \tag{19}$$

Lemma 3 *In a symmetric Bayes-Nash equilibrium, $q(v)$ must be non-negative and weakly increasing in v for any $v > 0$, and*

$$a(v) = v - \frac{V_S(0) + \int_0^v q(z)\mathrm{d}z}{q(v)}. \tag{20}$$

We omit the proof of Lemma 3 since it is a property of bilateral trading with private information in Myerson and Satterthwaite (1983). Lemma 3 immmediately implies that the speculator trades weakly more when having more private information represented by a higher absolute value of v. Also, from (20) we know that the trading profit of the speculator $V_S(v)$, which equals $V_S(0) + \int_0^v q(z)\mathrm{d}z$, must be a convex and increasing function of v for $v > 0$. $V_S(0)$ must also be non-negative since the individual rationality constraint must hold when the speculator has $v = 0$. Given that $V_S(v)$ is a non-negative, increasing convex funciton of v, we can show that $a(v)$ is an increasing function of v using the same steps as in the proof of Proposition 3. We formally state this relationship in the proposition below.

Proposition 4 (Within-speculator trade heteogeneity) *A speculator trades more at a higher spread when receiving higher private signal x.*

Recall that the across-client relationship among trade sizes, bid-ask spreads and trade informativeness depend on the convexity of $V_I(v)$, the value of information function. In contrast, Proposition 4 shows that the within-client variation is independent of the shape of $V_I(v)$. Why is there a disconnection between within-client variation and across-client variation? This is because with non-binary signals the speculator does not always capture a fixed fraction of the realized value of information. The surplus captured by the speculator with private information v, which is now represented by $V_S(v)$, can deviate from $V_I(v)$, the value of information v to the dealer. In fact, IC constraint forces $V_S(v)$ to be weakly convex, independent of the convexity of $V_I(v)$. In the appendix, we solve a joint model that features both within-client and across-client variation in trade size, bid-ask spread and trade informativeness. We show that $V_S(v)$ and $V_I(v)$ only equal to each other in expectation with respect to the private signal x. In the joint model, all the previous results in Proposition 3 and 4 hold with slight modifications.

We adopted endogenous information acquisition, whereby the speculator chooses a signal precision η which is not observed by other traders. In equilibrium, this choice is signaled through the size q of the speculator's request to trade via forward induction $\hat{\eta} = c'^{-1}(q)$. This modeling approach is necessary to allow for a separating equilibrium when V_S is concave, in which case information chasing dominates adverse selection. Here, we discuss how some alternative modeling approaches fail to deliver the same result.

Exogenous Information Acquisition: If a speculator's information precision η was exogenously drawn by the nature while remaining private knowledge to the speculator, the speculator's distinguishing feature, its information acquisition cost function c, becomes entirely irrelevant. Therefore, the comparison between a more informed speculator (higher η) with a less informed speculator (lower η) effectively reduces to our within-trader comparison: In a separating equilibrium, the speculator signals his information precision η through the size of his trade request q. For such a separating equilibrium to exist, the information rebate function V_S has to be convex, which implies that adverse selection always dominates information chasing.

Observable Information Precision: If the speculator's information precision η is observable to the dealer, then model would lose the discipline on the trade size q, regardless of whether the information precision η is endogenously chosen by the speculator or exogenously drawn by the nature. Hence, the bid-ask spread is undetermined.

Xiong and Yang (2020) show that the observabilility of a speculator's information acquisition level lowers the speculator's equilibrium choice of information precision in the setting of Kyle (1985). In our setting, the equilibrium information precision η remains the same whether or not the choice is observable to the dealer. The difference is with the presence of noise traders whose order pool with the speculator's order. Hiding behind such noise traders, the speculator has a incentive to secretly acquire a greater information precision which is only partially inferred by the market maker. Without such noise traders, however, the speculator has no such incentive because his information precision is perfectly inferred in a separating equilibrium.

4 Empirical Evidence

4.1 Data and Client Classification

To test the predictions of the theoretical model of order chasing, one needs a detailed transaction-level dataset together with a classification scheme whereby informed and uninformed traders could be identified. To that end, we use the proprietary ZEN database maintained by the UK Financial Conduct Authority (FCA), which covers virtually the universe of secondary-market transactions in the UK government bond market. Importantly, the dataset contains information on the identity of both sides of a trade (unlike other datasets on OTC markets, such as the TRACE database). This allows us to identify informed and uninformed clients, and to keep track of the time-variation in the fraction of trading volume initiated by informed and uninformed clients at each individual dealer.

To test the predictions of our theory, we are therefore able to exploit (i) the cross-sectional variation in client types, and (ii) the time-variation in the client composition at the dealer-level. Our sample covers the period between August 2011 and December 2017. During this period, there are 21 primary dealers and 576 clients that we have identified.[12] In our baseline classification, sophisticated clients include hedge funds and asset managers;

[12]The identities of the currently active primary dealers can be found on the website of the Debt Management Office: https://www.dmo.gov.uk/responsibilities/gilt-market/market-participants/.

whereas uninformed clients include insurance companies, pension funds, government entities (e.g. central banks) and non-financial corporations. We end up with 292 sophisticated clients and 284 unsophisticated clients, accounting for approximately two thirds and one third of the total trading volume, respectively.

In addition, we apply a classification scheme that is based on clients' realised trading performance in our sample. The idea is to estimate the profit and loss (P&L) account for each client i using the realised transactions as well as evaluating any inventory outstanding at an appropriate market price. Specifically, we compute for each client i the following measure:

$$P\&L_i = \sum_{a}^{A_i} \left\{ \underbrace{\sum_{j_{i,a}^S=1}^{J_{i,a}^S} Q_{j_{i,a}^S}^S P_{j_{i,a}^S}^S - \sum_{j_{i,a}^B=1}^{J_{i,a}^B} Q_{j_{i,a}^B}^B P_{j_{i,a}^B}^B}_{Realised\ Cash-flow} + \underbrace{\left(\sum_{j_{i,a}^B=1}^{J_{i,a}^B} Q_{j_{i,a}^B}^B - \sum_{j_{i,a}^S=1}^{J_{i,a}^S} Q_{j_{i,a}^S}^S \right) \frac{1}{N_{i,a}} \sum_{m_{i,a}=1}^{M_{i,a}} P_{m_{i,a}}}_{Valuation\ of\ Inventories} \right\} \tag{21}$$

where $J_{i,a}^S$ and $J_{i,a}^B$ denote the total number of sell and buy transactions of client i in bond a, and Q and P denote the quantity and price of a given transaction of client i. The first term in 21 denotes the realised cash-flows from buying and selling bond a, and the remaining term captures the valuation effect corresponding to any negative or positive inventory the client may accumulate during the sample period. To valuate inventories, we take a conservative approach and use the average transaction price faced by given client i. If the client buys the same quantity in bond a as she sells, then this inventory term would be zero. We then sum across all the bonds that client i has traded, to arrive at the client-specific performance measure $P\&L_i$. We scale this performance measure by the total trading volume of the given client, in order to mitigate the mechanical effect of client size on performance measurement.

We use our scaled $P\&L_i$ measure to split our sophisticated clients into two groups: one with $P\&L_i$ values in the top tertile and the remaining group with $P\&L_i$ values in the bottom two tertiles. We will refer to this latter group, consisting of 97 traders, as *informed clients* in the remainder of the analysis.

4.2 Testing the Theory

4.2.1 Claim 1: Informed clients face lower average trading costs than less informed clients

To test the first claim predicted by the theory, we first construct a measure of trading costs for each trade. While trade-specific bid-ask quotes are not observed, our approximation is based on the realised price deviation of the given trade from a benchmark price in the corresponding bond (in the spirit of O'Hara, Wang and Zhou (2018), O'Hara and Zhou (2021) and Pinter, Wang and Zou (2021)). Formally, for each trade j, on day t and bond k, we construct the measure $Cost_{j,k,t}$ as follows:

$$Cost_{j,k,t} = \left[\ln\left(P^{\star}_{j,k,t}\right) - \ln\left(P_{k,t}\right)\right] \times \mathbf{1}^{B,S}_{j}, \tag{22}$$

where $P^{\star}_{j,k,t}$ is the transaction price, $P_{k,t}$ is benchmark price the daily closing quoted mid-price of the corresponding bond, and $\mathbf{1}^{B,S}_{j}$ is an indicator function equal to 1 when transaction j is a buy trade, and equal to -1 when it is a sell trade. As benchmark price, we use the average transaction price at the bond-day-dealer level. As a robustness check, we also use the daily closing quoted mid-price of the corresponding bond, obtained from Datastream. The higher the measure $Cost_{j,k,t}$ in 22, the less favourable the given client's trading costs are.

Given our measure of trading costs 22 (which proxies the bid-ask spread), we estimate the following transaction-level regression for client i, asset k, dealer m and day t:

$$Cost_{i,k,m,t} = \beta \times D^{Inf}_{i} + \gamma \times TradeSize_{i,k,m,t} + \mu_{k,t} + \delta_{m,t} + \varepsilon_{i,k,m,t}, \tag{23}$$

where D^{Inf}_{i} is a dummy taking value 1 if the client i is sophisticated and informed and 0 if the client is unsophisticated. The terms $\mu_{k,t}$ and $\delta_{m,t}$ are bond-day and dealer-day fixed effects, respectively. The object of interest in D.1 is β which captures how much more favourable the trading cost is on the trade of an informed client compared to the trade of an unsophisticated client who is trading at the same dealer on the same day (this interpretation is possible because of the inclusion of the fixed effect $\delta_{m,t}$).

Table 1: Relative Trading Costs of Informed Clients

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Informed Clients	-0.529***	-0.497***	-0.455***	-0.417***	-0.472***	-0.464***	-0.521***
	(-3.22)	(-3.54)	(-3.27)	(-2.95)	(-3.33)	(-3.17)	(-3.32)
Client Size			-0.028	-0.018	-0.094***	-0.114***	-0.111***
			(-1.06)	(-0.67)	(-2.77)	(-2.71)	(-2.73)
Dealer-Connections				-0.009	-0.036*	-0.035*	-0.028
				(-0.54)	(-1.91)	(-1.88)	(-1.46)
Client Intensity					0.228***	0.250***	0.252***
					(3.67)	(3.33)	(3.24)
Trade Size						0.019	0.018
						(0.62)	(0.57)
N	542837	538426	538426	538426	538426	538426	455403
R^2	0.010	0.158	0.158	0.158	0.158	0.158	0.340
Day FE	Yes	No	No	No	No	No	No
Bond FE	Yes	No	No	No	No	No	No
Dealer FE	Yes	No	No	No	No	No	No
Day*Dealer FE	No	Yes	Yes	Yes	Yes	Yes	No
Day*Bond FE	No	Yes	Yes	Yes	Yes	Yes	No
Day*Bond*Dealer FE	No	No	No	No	No	No	Yes

Notes: This table regresses trading costs (computed by 22 using the average transaction price at the bond-day-dealer level as the benchmark price) on an informed sophisticated client dummy, various controls and various fixed effects. "Client Size" is the log of the average monthly trading volume of a given client. "Dealer-Connections" is the total number of unique dealers (averaged across months) that a given client trades with in a given month. "Client Intensity" is the log of the average monthly number of transactions of a given client. "Trade Size" is the log of the trade size in £s. Informed clients include those asset managers and hedge funds whose average scaled P&L measure 21 is in the top tertile. To reduce noise, we winsorise the sample at the 1-99%-level. T-statistics in parentheses are based on robust standard errors, using two-way clustering at the day and client level. Asterisks denote significance levels (* p<0.1, ** p<0.05, *** p<0.01).

Table 1 shows the results from regression D.1 using various specifications. In column (1) we include the day-, bond- and dealer-level fixed effects separately (without interactions), whereas columns (2)-(6) include day-dealer and day-bond fixed effects that aim to control for the linear effect of any dealer- and bond-level shocks that might hit on a given trading day. Column (7) corresponds to the tightest specification with day-dealer-bond fixed effects, which allows for the comparison of the trading costs of different types of clients who trade the same bond, at the same dealer on the same day.

Overall, the results imply that the trading costs faced by informed clients at a dealer are about 0.5 bps lower compared to trades initiated by unsophisticated clients at the same dealer.[13] This result is robust to the inclusion of a number of additional regressors that aim to control for mechanisms related to heterogeneity in clients' bargaining power or in their exposure to search frictions.

To show that our baseline is not simply picking up these mechanisms we first include "client size" which is the average monthly trading volume of clients. Consistent with the evidence of size discount (Green, Hollifield and Schurhoff, 2007, Pinter, Wang and Zou, 2021), we find that larger traders typically get more favourable transaction prices than smaller traders. Next, we compute the total number of dealers that a client trades with in a given month. This measure of clients' "dealer-connections" aims to control for the client's position in the trading network which could affect her trading costs.[14] While this suggests that informed clients tend to have more dealer-connections, and clients with more connections tend to face more favourable trading costs[15],the informedness of clients seems to matter over and above what is captured by her trading network. Third, we add as control "client intensity" which is the total number of transactions a client carries out in a given

[13]It is important to note that, by using the average transaction price at the bond-day-dealer level as the benchmark price in 22, we implicitly control for the fact that informed and unsophisticated clients may trade with different dealers.

[14]This is motivated by recent papers that explored the cross-sectional variation of dealers' network centrality as well as dealers' relationships with clients in driving trading costs in the US corporate bond market (Maggio, Kermani and Song, 2017, Hendershott, Li, Livdan and Schurhoff, 2020).

[15]These results are consistent with Kondor and Pinter (2019).

month averaged over the sample.[16] Next, we include "client size" which is the average monthly trading volume of clients. To show that our results are not simply picking up the effect of trade size (Edwards, Harris and Piwowar, 2007, Bernhardt, Dvoracek, Hughson and Werner, 2005), column (6) includes trade size as an additional control, with little effect on the baseline results.[17]

Robustness Checks In our baseline, we used the benchmark price for our measure of trading costs that is based on realised transaction prices. Table 2 in the Appendix shows the results where we use the end-of-day mid-quote (from Datastream) as the benchmark price in 22. We find that the results become stronger with informed clients facing more favourable trading costs by about 0.6-0.8 bps compared to unsophisticated clients. A likely explanation for the stronger result is that, by using the end-of-day quote as the benchmark price, our estimation is likely capturing some of the price impact of informed trading activity. Therefore, by measuring trading costs with effective spreads (Bessembinder and Venkataraman, 2010), we get stronger results compared to our baseline in Table 1.

Moreover, we illustrate how much our baseline result is driven by our choice of control and treatment groups. First, we broaden the set of the treated group by including all sophisticated clients (i.e. all asset managers and hedge funds irrespective of their $P\&L_i$ measure) in it, and compare their trading costs to those of unsophisticated clients. Table 3 shows the results, indicating that while the difference in the trading costs of sophisticated and unsophisticated clients is still statistically significant, the economic effect drops to 0.2 bps in the most conservative specification (compared to 0.52 bps in our baseline Table 1). Next, we focus only on sophisticated clients and compare informed to uninformed clients. Table 4 shows that the trading costs of informed clients are about 0.35-0.47 bps lower than those of uninformed clients. These results show that heterogeneity both in client types and

[16] This control is motivated by O'Hara, Wang and Zhou (2018) who used a subset of the US corporate bond market to analyse the trading costs of insurance companies.

[17] While Table 1 focuses on the across-client variation, in a related paper (Pinter, Wang and Zou, 2021), we show evidence on a positive relation between trade size and trading costs once we exploit the within-client variation of the data, i.e. once we include a client fixed effect.

informedness contributes to our baseline results.

Moreover, we check how sensitive our baseline results are to changing the definition of performance that we use to place clients in the informed category. First, we use the average 20-day percentage return (instead of the $P\&L_i$ measure (21)) to identify informed clients.[18] Second, we re-estimate our baseline by using the unscaled $P\&L_i$ measure (21) to identify informed clients.[19] Tables 5–6 show that the results continue to be similar to our baseline, with informed clients facing lower average trading costs by about 0.3-0.45 bps.

4.2.2 Claim 2: A dealer anticipates larger gains against less informed client when the dealer gives better prices to informed clients

The heart of the mechanism in our theoretical model is the idea that dealers actively shape trading costs to attract trades with informed clients. Trading with informed clients allows dealers to learn from them, which can be used to make profits when the given dealer trades with less informed traders. To test this mechanism, we first construct a measure of profitability at the trade-level, in the spirit of Di Maggio, Franzoni, Kermani and Sommavilla (2019), based on the given trade's ability to predict future prices over a given horizon. Formally, for each trade j, on day t, bond k and horizon T, we construct the measure $Perf^T_{j,k,t}$ as follows:

$$Perf^T_{j,k,t} = [\ln(P_{k,t+T}) - \ln(P_{k,t})] \times \mathbf{1}^{B,S}_j, \tag{24}$$

where $P_{k,t}$ is the benchmark price of bond k on day t, $P_{k,t+T}$ is the benchmark price T days later, and $\mathbf{1}^{B,S}_j$ is an indicator function equal to 1 when transaction j is a buy trade, and equal to -1 when it is a sell trade. We then aggregate the performance measure 24 for each dealer i, month t, and horizon T, based on all the transactions against unsophisticated

[18] Specifically, we compute the 20-day return on each trade of a client (similar to formula 24 below) and compute a size-weighted average of these returns for the given client. We then rank clients based on these average returns. The choice of the 20-day horizon is motivated by recent results (Czech, Huang, Lou and Wang 2021) showing that the trades of hedge funds and asset managers have predictive power of future returns around this horizon.

[19] This is motivated by the mutual fund literature (Berk and van Binsbergen, 2015) that highlighted the important role of gross performance measures.

clients as well as inter-dealer-brokers.[20] Similarly, we aggregate the transaction performance 22 for each dealer i and month t against informed sophisticated clients. We then estimate the following panel regression at the dealer-month level:

$$UninfPerf_{i,t}^{T} = \beta \times InfCost_{i,t} + \gamma \times \log\left(Vol_{j,t}\right) + \alpha_i + \mu_t + \varepsilon_{i,t}, \tag{25}$$

where $UninfPerf_{i,t}^{T}$ is the trading performance of the group of unsophisticated clients as well as inter-dealer brokers against dealer i on day t at horizon T. The term $InfCost_{i,t}$ is average trading cost of informed sophisticated clients, against dealer i on day t. The term $Vol_{i,t}$ denotes dealer's trading volume and α_i and μ_t are dealer and month fixed effects.

Figure 3: The relationship between trading costs of informed clients at a dealer and the trading performance of unsophisticated clients against the given dealer

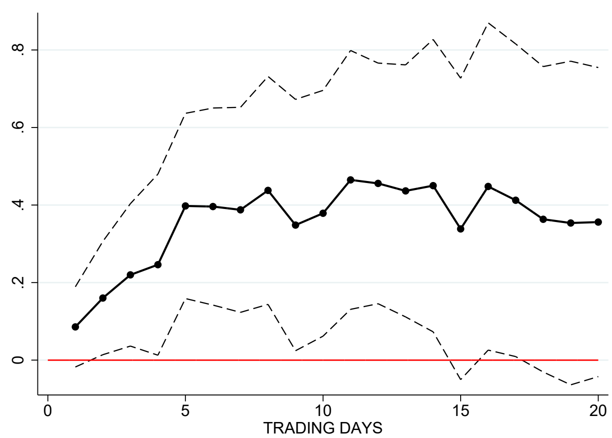

Notes: this figure plots the estimated β coefficients from our baseline monthly regression 25 up to 20-day horizon ($T = 20$), using as regressand the average value weighted performance of uninformed clients trading with dealer i in month t. We include as a control the natural logarithm of the pound trade volume of dealers. To reduce noise, we winsorise the sample at the 1%-level. The shaded area denotes the 90% confidence band, based on robust standard errors, using clustering at the month level.

The coefficient of interest in 25 is β which we expect to be positive if the suggested learning mechanism is at play: when informed sophisticated clients receive more favourable transaction prices (i.e. the dealer charging lower trading costs on these trades), then the dealer is entering into trades against unsophisticated clients that turn out to generate capital

[20]Inter-dealer-brokers (IDB) provide an important platform for inter-dealer-trades in the UK government bond market. Over 90% of inter-dealer trading volume is done through IDBs and only a small fraction directly between dealers.

gains over future horizon $T = [1, \ldots, 20]$ for the dealer (and capital losses for those clients). Figure 3 plots the estimated β coefficients in regression 25 up to 20 horizons, showing evidence that a dealer's trading performance against unsophisticated clients is higher when the given dealer offer more favourable execution prices to informed clients. We find that the effect is persistent, though it is estimated with larger uncertainty at longer horizons. The fact that the effect initially increases with the horizon suggests that the information dealers may acquire from informed clients pertains to future price movements that are less immediate.[21]

It is important to note that the inclusion of dealer fixed effects α_i means that we primarily identify the effect from the time-series, i.e. we compare months when a dealer gives more favourable execution prices to informed clients to other months when the *same dealer* gives less favourable prices. The regressions, designed to test claim 2, so far exploited monthly variation in average costs and performance. As a robustness check, we also experiment using daily, instead of monthly, variation, and find this adds additional measurement noise to our empirical tests but does not qualitatively change the result. Figure 5 in Appendix plots the estimated β coefficients in regression 25 up to 20-day horizon, using data at daily frequency, with results similar to that shown in Figure 3.

We also check whether more favourable trading costs faced by informed clients at a dealer are associated with increased trading activity of informed clients at the given dealer. To test for that, we replace the left-hand size variable in 25 with two possible proxies for informed clients' trading activity at the dealer level: the ratio of informed trading volume to total volume, and the natural logarithm of informed trading volume. Table 7 in Appendix shows a statistically significant negative relationship, confirming that dealers' trading activity with informed clients increases when those clients are offered lower trading costs by the given dealer.

[21] This is consistent with the horizon effects of informed clients documented by Kondor and Pinter (2019) and Czech, Huang, Lou and Wang (2021).

4.2.3 Claim 3: An informed client faces larger trading costs when being more informed, compared to when being less informed

To test this claim, we exploit the within-client variation in average trading costs and the level of informedness of the given client. If adverse selection dominates, we would expect trading costs of a client to be higher when the given client is more informed, i.e. when she is better at predicting future price movements. To test this prediction, we estimate the following panel regression model for each client i and month t:

$$Perf_{i,t}^{T} = \beta \times Cost_{i,t} + \mu_i + \delta_t + \varepsilon_i, \tag{26}$$

where $Perf_{i,t}^{T}$ is the average anticipation component (24) of client i over horizon T in month t. The main coefficient of interest is β which measures how much the anticipation component of a client (who is on average informed) changes when the given client's average trading costs increase by 1bp from one month to the next. Note that this time-series interpretation of the effect is possible because of the inclusion of the fixed effect μ_i which controls for the linear effects of any time-invariant cross-sectional heterogeneity in trading performance and trading costs across clients.

Figure 4 plots the estimated β coefficients in regression 26 up to a 20-day horizon, confirming that clients face higher trading costs when their trades better predict future price movements.[22] We find that a 1 bp increase in trading costs (compared to the trading costs of other clients trading at the same dealer) is associated with up to 0.5 bps increase in the anticipation component of the given client. The effect is persistent, with no sign of reversal.

For robustness checks, we re-estimate regression 26 for the group of less sophisticated clients: Figure 6 shows the results for this group (right panel) along our baseline (left panel), confirming that the positive relation between average trading costs and capital gains is only

[22] These results are consistent with informed clients trading larger amounts when being more informed (Kyle, 1985). See our companion paper (Pinter, Wang and Zou, 2021) for further empirical evidence on trade size, trading costs and informedness of clients.

Figure 4: The relationship between trading costs and future capital gains amongst informed clients: 1-20 day horizon

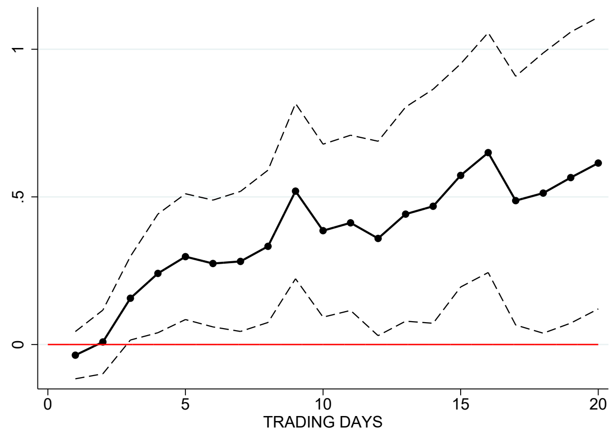

Notes: this figure plots the estimated β coefficients from our baseline monthly regression 26 up to 20-day horizon ($T = 20$), using as regressor (regressand) the unweighted average trading cost (anticipation component) of client i in month t. We include as a control the natural logarithm of the pound trade volume of clients. The sample only includes informed sophisticated clients. To reduce noise, we winsorise the sample at the 1%-level. The shaded area denotes the 90% confidence band, based on robust standard errors, using clustering at the month level.

present in our sample of informed clients. Moreover, Figure 7 shows the baseline results when we use trade size - weighted averages when constructing both the dependent and independent variables in regression 26.

5 Conclusion

Contrary to the prediction of the classic adverse selection theory, a more informed trader receives better pricing relative to a less informed trader in some over-the-counter financial markets. We show that dealers compete for information by chasing informed orders so as to better position their future price quotes. On a multi-dealer platform, dealers' incentive of information chasing exactly offsets their fear of adverse selection. As a result, the adverse selection cost is passed on to uninformed traders. Information chasing induces winner's curse among dealers, which in turn results in price dispersion and bid-ask spread for uninformed hedgers. Both price dispersion and price efficiency increase with hedging demand.

Information chasing is possible only without pre-trade anonymity. Hence, it is absent on centralized exchanges. It is also absent on a non-anonymous centralized exchange, because trades are disclosed in real time. Consistently, Theissen (2003) shows that in the non-anonymous Frankfurt Stock Exchange, trades that are more likely to be motivated by proprietary information about asset payoff tend to receive wider bid-ask spreads.

References

Abad, Jorge, Iñaki Aldasoro, Christoph Aymanns, Marco D'Errico, Linda Fache Rousová, Peter Hoffmann, Sam Langfield, Martin Neychev, and Tarik Roukny, "Shedding Light on Dark Markets: First Insights from the New EU-Wide OTC Derivatives Dataset," ESRB Occasional Paper Series 11, European Systemic Risk Board September 2016.

Asriyan, Vladimir, William Fuchs, and Brett Green, "Information Spillovers in Asset Markets with Correlated Values," *American Economic Review*, July 2017, *107* (7), 2007–2040.

Banerjee, Snehal, Jesse Davis, and Naveen Gondhi, "When Transparency Improves, Must Prices Reflect Fundamentals Better?," *The Review of Financial Studies*, June 2018, *31* (6), 2377–2414.

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

Berk, Jonathan B. and Jules van Binsbergen, "Measuring skill in the mutual fund industry," *Journal of Financial Economics*, 2015, *118* (1), 1–20.

Bernhardt, Dan, Vladimir Dvoracek, Eric Hughson, and Ingrid M. Werner, "Why Do Larger Orders Receive Discounts on the London Stock Exchange?," *Review of Financial Studies*, 2005, *18* (4), 1343–1368.

Bessembinder, Hendrik and Kumar Venkataraman, *Bid Ask Spreads*, Wiley, 2010.

Bjønnes, Geir, Neophytos Kathitziotis, and Carol Osler, "Price Discrimination and the Cost of Liquidity in OTC Markets," Technical Report 2015.

Brancaccio, Giulia, Dan Li, and Norman Schürhoff, "Learning by Trading: The Case of the US Market for Municipal Bonds," Technical Report 2020.

Brandt, Michael W. and Kenneth A. Kavajecz, "Price Discovery in the U.S. Treasury Market: The Impact of Orderflow and Liquidity on the Yield Curve," *Journal of Finance*, 2004, *59* (6), 2623–2654.

Camargo, Braz, Kyungmin Kim, and Benjamin Lester, "Information Spillovers, Gains from Trade, and Interventions in Frozen Markets," *The Review of Financial Studies*, 2016, *29* (5), 1291–1329.

Chang, Briana and Shengxing Zhang, "Endogenous Market Making and Network Formation," SSRN Scholarly Paper ID 2600242, Social Science Research Network, Rochester, NY February 2019.

_ and _ , "Risk Concentration and Interconnectedness in OTC Markets," SSRN Scholarly Paper ID 3808072, Social Science Research Network, Rochester, NY March 2021.

Chen, Wen and Yajun Wang, "Dynamic Market Making with Asymmetric Information and Market Power," SSRN Scholarly Paper ID 3224495, Social Science Research Network, Rochester, NY August 2020.

Chordia, Tarun, Richard Roll, and Avanidhar Subrahmanyam, "Order imbalance, liquidity, and market returns," *Journal of Financial Economics*, July 2002, *65* (1), 111–130.

_ , _ , and _ , "Evidence on the speed of convergence to market efficiency," *Journal of Financial Economics*, May 2005, *76* (2), 271–292.

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

Collin-Dufresne, Pierre, Peter Hoffmann, and Sebastian Vogel, "Informed Traders and Dealers in the FX Forward Market," Technical Report 2020.

Czech, Robert, Shiyang Huang, Dong Lou, and Tianyu Wang, "Informed trading in government bond markets," *Journal of Financial Economics*, 2021, *142* (3), 1253–1274.

Duffie, Darrell and Gustavo Manso, "Information Percolation in Large Markets," *American Economic Review*, April 2007, *97* (2), 203–209.

_ , Gaston Giroux, and Gustavo Manso, "Information Percolation," *American Economic Journal: Microeconomics*, February 2010, *2* (1), 100–111.

_ , Semyon Malamud, and Gustavo Manso, "Information Percolation With Equilibrium Search Dynamics," *Econometrica*, 2009, *77* (5), 1513–1574.

_ , _ , and _ , "Information Percolation in Segmented Markets," *Journal of Economic Theory*, September 2014, *153*, 1–32.

Edwards, Amy, Lawrence E. Harris, and Michael S. Piwowar, "Corporate Bond Market Transaction Costs and Transparency," *Journal of Finance*, 2007, *62* (3), 1421–1451.

Glode, Vincent and Christian Opp, "Asymmetric Information and Intermediation Chains," *American Economic Review*, September 2016, *106* (9), 2699–2721.

Glosten, Lawrence R. and Paul R. Milgrom, "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 1985, *14* (1), 71–100.

Golosov, Mikhail, Guido Lorenzoni, and Aleh Tsyvinski, "DECENTRALIZED TRADING WITH PRIVATE INFORMATION," *Econometrica*, 2014, *82* (3), 1055–1091.

Green, Richard, Burton Hollifield, and Norman Schurhoff, "Financial Intermediation and the Costs of Trading in an Opaque Market," *Review of Financial Studies*, 2007, *20* (2), 275–314.

Green, T. Clifton, "Economic News and the Impact of Trading on Bond Prices," *Journal of Finance*, 06 2004, *59* (3), 1201–1234.

Grossman, Sanford J and Joseph E Stiglitz, "On the Impossibility of Informationally Efficient Markets," *The American Economic Review*, 1980, *70* (3), 393–408.

Hagströmer, Björn and Albert J. Menkveld, "Information Revelation in Decentralized Markets," *The Journal of Finance*, 2019, *74* (6), 2751–2787.

Hau, Harald, Peter Hoffmann, Sam Langfield, and Yannick Timmer, "Discriminatory Pricing of Over-the-Counter Derivatives," *Management Science*, March 2021.

Hendershott, Terrence, Dan Li, Dmitry Livdan, and Norman Schurhoff, "Relationship Trading in Over the Counter Markets," *Journal of Finance*, April 2020, *75* (2), 683–734.

Hugonnier, Julien, Benjamin Lester, and Pierre-Olivier Weill, "Heterogeneity in Decentralized Asset Markets," Technical Report 14274, C.E.P.R. Discussion Papers January 2020.

Kondor, Peter and Gabor Pinter, "Clients' Connections: Measuring the Role of Private Information in Decentralised Markets," CEPR Discussion Paper July 2019.

Kyle, Albert S., "Continuous Auctions and Insider Trading," *Econometrica*, 1985, *53* (6), 1315.

_ , "Informed Speculation with Imperfect Competition," *The Review of Economic Studies*, 1989, *56* (3), 317.

Lee, Tomy and Chaojun Wang, "Why Trade Over-the-Counter? When Investors Want Price Discrimination," SSRN Scholarly Paper ID 3087647, Social Science Research Network, Rochester, NY February 2018.

Lester, Benjamin R., Ali Shourideh, Venky Venkateswaran, and Ariel Zetlin-Jones, "Market-Making with Search and Information Frictions," SSRN Scholarly Paper ID 3228477, Social Science Research Network, Rochester, NY July 2018.

Lewis, Ryan and Michael Schwert, "The Effects of Transparency on Trading Profits and Price Informativeness: Evidence From Corporate Bonds," *SSRN Electronic Journal*, December 2018.

Li, Wei and Zhaogang Song, "Dealers as Information Intermediaries in Over-the-Counter Market," *SSRN Electronic Journal*, April 2019.

Maggio, Marco Di, Amir Kermani, and Zhaogang Song, "The value of trading relations in turbulent times," *Journal of Financial Economics*, 2017, *124* (2), 266 – 284.

Maggio, Marco Di, Francesco Franzoni, Amir Kermani, and Carlo Sommavilla, "The relevance of broker networks for information diffusion in the stock market," *Journal of Financial Economics*, 2019, *134* (2), 419–446.

Myerson, Roger B and Mark A Satterthwaite, "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, April 1983, *29* (2), 265–281.

Naik, Narayan Y., Anthony Neuberger, and S. Viswanathan, "Trade Disclosure Regulation in Markets with Negotiated Trades," *Review of Financial Studies*, July 1999, *12* (4), 873–900.

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

O'Hara, Maureen and Xing Zhou, "Anatomy of a Liquidity Crisis: Corporate Bonds in the Covid-19 Crisis," *Journal of Financial Economics*, 2021, (xxx), xxx.

_ , Yihui Wang, and Xing Zhou, "The execution quality of corporate bonds," *Journal of Financial Economics*, 2018, *130* (2), 308–326.

Pinter, Gabor, Chaojun Wang, and Junyuan Zou, "Size Discount and Size Penalty: Trading Costs in Bond Markets," mimeo 2021.

Ramadorai, Tarun, "What Determines Transaction Costs in Foreign Exchange Markets?," *International Journal of Finance and Economics*, January 2008, *13* (1), 14–25.

Rüdiger, Jesper and Adrien Vigier, "Who Acquires Information in Dealer Markets?," *American Economic Review*, April 2020, *110* (4), 1145–1176.

Sambalaibat, Batchimeg, "Endogenous Specialization and Dealer Networks," SSRN Scholarly Paper ID 2676116, Social Science Research Network, Rochester, NY July 2018.

Syrgkanis, Vasilis, David Kempe, and Eva Tardos, "Information Asymmetries in Common-Value Auctions with Discrete Signals," *Mathematics of Operations Research*, 2019, *44* (4), 1450–1476.

Theissen, Erik, "Trader Anonymity, Price Formation and Liquidity," *European Finance Review*, 2003, *7* (1), 1–26.

Üslü, Semih, "Pricing and Liquidity in Decentralized Asset Markets," *Econometrica*, 2019, *87* (6), 2079–2140.

Vives, Xavier, "Strategic Supply Function Competition With Private Information," *Econometrica*, 2011, *79* (6), 1919–1966.

Wang, Chaojun, "Core-Periphery Trading Networks," SSRN Scholarly Paper ID 2747117, Social Science Research Network, Rochester, NY December 2016.

Wang, Junbo, Chunchi Wu, and Eden S. H. Yu, "Order Imbalance, Liquidity, and Returns of the U.S. Treasury Market," *Review of Pacific Basin Financial Markets and Policies (RPBFMP)*, 2012, *15* (02), 1–39.

Xiong, Yan and Liyan Yang, "Secret and Overt Information Acquisition in Financial Markets," SSRN Scholarly Paper ID 3579214, Social Science Research Network, Rochester, NY April 2020.

Appendices

A Proofs

The quoting game in stage 2 is equivalent to a common-value first-price sealed bid auction with discrete signals. First we state a lemma that helps us construct the unique equilibrium of this auction games with 2 players. The proof of this lemma can be found in Syrgkanis, Kempe and Tardos (2019).

Lemma A.1 *If $n = 2$, in a common-value first-price sealed bid auction with discrete signals, there exists a unique Nash equilibrium, in which*

1. *Each dealer's mixed strategy has a common support $[\underline{x}, \bar{x}]$,*
2. *For each dealer j, there exists a partition $\underline{x} = x_0^j \leq x_1^j < x_2^j < \cdots < x_{S_j}^j = \bar{x}$, where S_j is the number of all of possible realization of dealer j's signal combination. Each interval (x_k^j, x_{k+1}^j) corresponds to dealer j's signal realization ω_k^j.*
3. *There is no gap or atom in $(\underline{x}, \bar{x}]$.*
4. *At least one dealer bid $\underline{x}$ with probability 1 when receiving the worse signal realization.*
5. *Both dealers get expected payoff 0 when receiving their worst signal.*

This lemma has a direct implication.

Corollary A.1 *The expected payoff of an uninformed dealer is 0.*

Proof of Lemma 1 & 2. Since buying and selling are symmetric, $-a_2^+$, $-a_2^-$, $-a_2^0$ must follow the same distribution as b_2^+, $b_2^-$and b_2^0. In the following proof, we only focus on the bidding prices. Let G^+, G^- and G^0 be the c.d.f. of b_2^+, $b_2^-$and b_2^0.

We start with the case of $n = 2$. Lemma A.1 implies that there exists a unique Nash equilibrium in mixed strategy, where

1. The informed dealer bid $b_2^- = -\hat{\eta}$ with probability 1 if she observes a negative signal.

2. The mixed strategies of the uninformed dealer and the informed dealer who observes a positive signal have the same lower bound $-\hat{\eta}$ and the same upper bound.

3. Both *supp* G^0 and *supp* G^+ are connected sets.

4. The distribution of b_2^+ has no mass point. The distribution of b_2^0 has no mass point other than at $-\hat{\eta}$.

An uninformed dealer must be indifferent of bidding any $b \in$ *supp* G^+, therefore

$$\frac{1}{2}G^+(b)(\hat{\eta} - b) + \frac{1}{2}(-\hat{\eta} - b) = C^0. \tag{A.1}$$

Notice C^0, the expected value of being an uninformed dealer in the third stage, must equal to 0. We can solve for G^+:

$$G^+ = \frac{2}{1 - b/\hat{\eta}} - 1. \tag{A.2}$$

The upper bound of *supp* G^+ is 0. An informed dealer must be indifferent of bidding any $b \in$ *supp* G^0, therefore for any $b \in [-\hat{\eta}, 0]$

$$(\hat{\eta} - b)G^0(b) = C^+. \tag{A.3}$$

Let $b = 0$ we have $C^+ = \hat{\eta}$. Plugging into the previous equation, we have

$$G^0(b) = \frac{\hat{\eta}}{\hat{\eta} - b}, \quad b \in [-\hat{\eta}, 0]. \tag{A.4}$$

Now we turn to the case of $n > 2$. It is easy to show that in any equilibrium the informed dealer must bid $b_2^- = -\hat{\eta}$ with probability 1 when receiving a bad signal. Denote the CDF of the informed dealer' bidding strategy when receiving a good signal by $G_{n>2}^+$ and the CDF of the maximum bid among the uninformed dealers by $G^0_{n>2}$. We want to show that in equilibrium the informed dealer is guaranteed a positive profit while any uninformed dealer gets zero profit. First notice that uninformed dealer will never bid above 0, the ex-ante expected value of the asset, if there's any positive probability of winning. Therefore, the informed dealer can always bid slightly above 0 to get a positive profit. Second, if an uninformed dealer gets positive profit in equilibrium, all uninformed dealers must have the same positive profit in equilibrium. Let $\underline{b}$ be the lower bound of the support of $G^0_{n>2}$. There must be a positive probability that $\underline{b}$ is higher than the informed dealer's bid conditional on a good signal. $G^+_{n>2}(\underline{b})$ must be positive. Since the informed dealer must have positive payoff when receiving a good signal, he must win the bidding game with positive probability. Thus, $G^0_{n>2}(\underline{b}) > 0$. This means $\underline{b}$ is not only the lower bound of the support but also a mass point of $G^0_{n>2}$. However, if this is a case, $G^+_{n>2}(\underline{b})$ is not a best response to $G^0_{n>2}$, since the informed dealer will never bid less or equal to $\underline{b}$. Contradiction. Now we have shown that the uninformed dealers must have zero payoff in equilibrium.

With the above results we can proceed to show that $G^+_{n>2}$ and $G^0_{n>2}$ must be an equilibrium of the game with $n = 2$, which means, they must be the same as in (A.2) and (A.4). To prove this, first notice that $G^+_{n>2}$ must be the best response to $G^0_{n>2}$ in the game with $n = 2$, as the expected payoff of the informed dealer only depends on the distribution of the maximum bid among the uninformed dealers. Therefore, we only need to show that $G^0_{n>2}$ is the best response to $G^+_{n>2}$ when $n = 2$. If this is not the case, there exists $b' > -\hat{\eta}$ such that by bidding b', the uninformed dealer has strictly positive profit,

$$\frac{1}{2}G^+_{n>2}(b')(\hat{\eta} - b') + \frac{1}{2}(-\hat{\eta} - b') > 0. \tag{A.5}$$

This implies that $G^+_{n>2}(b') > 0$, i.e., the informed dealer in the $n > 2$ game bids $b \leq b'$ with positive probability. Since the informed dealer has a positive payoff when receiving a positive signal, he must have positive probability of winning by bidding b'. Therefore, $G^0_{n>2}(b') > 0$.

However, if this is true any uninformed dealer can bid b' in the game with $n > 2$ and gets positive profit. This contradicts the previous results that all uninformed dealers must have zero payoff in equilibrium. Thus, $G^0_{n>2}$ is the best response to $G^+_{n>2}$ when $n = 2$.

Just to give an example of an equilibrium with $n > 2$, notice that the previous G^+ in (A.2) and

$$G^0(b) = \sqrt[n-1]{\frac{\hat{\eta}}{\hat{\eta} - b}}, \quad b \in [-\hat{\eta}, 0], \tag{A.6}$$

is a mixed strategy equilibrium of the stage-2 game with $n > 2$.

From Lemma 1, we know that if the signal is positive an informed dealer can only profit from buying from but not selling to the liquidity traders. For each liquidity seller, the expected profit is $C^+ = \eta$. On the other hand, if the signal is negative, the informed dealer can only profit from selling to the liquidity traders. Therefore, the ex-ante expected profit of the informed dealer is $\eta \cdot \frac{1}{2}m$. ■

Proof of Proposition 1. The proof has been given in the main text. ■

Proof of Proposition 2. By forward induction, the belief of the dealer $\hat{\eta}(q)$ is the solution to the speculator's optimization problem:

$$\hat{\eta}(q) = \underset{\eta}{\operatorname{argmax}} \; q\left[v(\eta) - a_1(q)\right] - c(\eta). \tag{A.7}$$

By Assumption 1, the objective function is concave in η, therefore, $\hat{\eta}(q)$ is the unique solution to the first order condition

$$qv'(\eta) = c'(\eta). \tag{A.8}$$

Plugging in the mid-to-bid price (15), the speculator's problem becomes

$$\max_q \; \varphi V_I(\hat{\eta}(q)) - c(\hat{\eta}(q)). \tag{A.9}$$

Assumption 1 implies that there is a unique solution to (A.9)

$$q^* = \frac{c'(\eta^*)}{v'(\eta^*)}, \; \eta^* \text{ solves } \varphi V_I'(\eta^*) = c'(\eta^*). \tag{A.10}$$

Note that the optimal order size can be equivalently written as

$$q^* = \frac{\varphi V_I'(\eta^*)}{v'(\eta^*)} = \varphi \left.\frac{\mathrm{d}V_I}{\mathrm{d}v}\right|_{v=v(\eta^*)}. \tag{A.11}$$

The equilibrium mid-to-bid spread is

$$a_1^*(q^*) = v(\eta^*) - \frac{\varphi V_I(\eta^*)}{q^*} = v(\eta^*) - \frac{\varphi V_I(\eta^*)}{\frac{\varphi V_I'(\eta^*)}{v'(\eta^*)}}, \tag{A.12}$$

$$= v(\eta^*)\left[1 - \frac{1}{\varepsilon(\eta^*)}\right], \quad \text{where } \varepsilon(\eta) = \frac{d\ln V_I}{d\ln v}. \tag{A.13}$$

■

Proof of Proposition 3. By Assumption 1, $\varphi V_I(\eta) - c'(\eta)$ is a decreasing function in η. When $c'(v)$ becomes uniformly lower for all v, η^* increases. Thus, a speculator acquires more information when the cost of information acquisition is lower. Next, we characterize how q^* and $a_1^*(q^*)$ depend on η^* by taking derivatives.

$$\frac{\mathrm{d}q^*}{\mathrm{d}\eta^*} = \phi v'(\eta^*) \left.\frac{\mathrm{d}^2 V_I}{\mathrm{d}v^2}\right|_{\eta=\eta^*}. \tag{A.14}$$

Since $v'(\eta) > 0$, if V_I is convex (concave) in v, q^* increases (decreases) in η^* so it decreases (increases) in the cost of information acquisition. For a_1^*, with slight abuse of notations we view V_I as a funciton of v and take derivative of a_1^* with respect to $v(\eta^*)$

$$\frac{\mathrm{d}a_1^*}{\mathrm{d}v^*} = 1 - \frac{V_I'(v^*)^2 - V_I(v^*)V_I''(v^*)}{V_I'(v^*)^2} = \frac{V_I(v^*)V_I''(v^*)}{V'(v^*)^2} \tag{A.15}$$

Since $V_I(v^*)$ and $V_I'(v^*)^2$ are both positive, the sign of $\frac{\mathrm{d}a_1^*}{\mathrm{d}v^*}$ is the same as that of $V_I''(v^*)$. If V_I is convex (concave) in v, a_1^* increases (decreases) in η^* so it decreases (increases) in the cost of information acquisition. In particular, if V_I is linear in v, a_1^* and q^* are independent of the cost of information acquisition. ■

B A Joint Model

Section 3 proposed two models with contrasting predictions: One predicts that information chasing can dominate adverse selection for between-client comparisons, and the other predicts that adverse selection always dominates information chasing for within-client comparisons. To encompass both predictions in one model, we further nest the two models in Section 3 into a more general model by allowing the speculator to obtain a continuous signal $x \in \mathbb{R}$ with a selected precision $\eta \in \mathbb{R}$. Then each of the two models in Section 3 can be viewed as a section of this more general model, by fixing either the signal strength $|x|$ or the signal precision η. We next lay out the new ingredients necessary to combine the features of both models.

Signal The marginal distribution of the private signal x is F.[23] While the joint distribution of (v, x) depends on the signal precision η (which can be defined, for example, as the negative conditional entropy $-H(v \mid x)$).

[23] Assuming that the distribution F does not depend on η is without loss of generality, because it is always possible to 1-1 map a random variable to another random variable whose distribution is given by F.

Pricing Upon receiving an trade request of size q from the speculator in a particular instant, the dealer offers the speculator a mid-to-offer spread

$$a(\hat{x}, \hat{\eta}) = v(\hat{x}, \hat{\eta}) - \frac{V_S(\hat{x}, \hat{\eta})}{q}. \tag{B.1}$$

where $\hat{x}$ and $\hat{\eta}$ are the dealer's beliefs about the speculator's signal realization x and precision η respectively, $v(x, \eta)$ is the expected value of the asset conditional on the speculator's signal realization x and precision η and, and $V_S(x, \eta)$ is the speculator's reduced-form information rebate from the dealer. To infer the signal's realization x and precision η, the dealer observes not only the size q of the instant trade request, but also the underlying distribution $\mathbb{P}$ of the size $\tilde{q}$ as a random variable through the speculator's other independent instances of trade requests.[24] Therefore, an information set of the dealer consists of $(q, \mathbb{P})$ such that q is in the support of $\mathbb{P}$. It will be the case that the dealer infers the signal realization x from the realized size q, and the precision η from the underlying distribution of size $\tilde{q}$. Neither sectional model in Section 3 requires the underlying distribution of the trade size $\tilde{q}$. This is because when the signal precision η is fixed, the dealer only needs the realized size q to infer the signal x, and when the signal x is binary, the random variable $\tilde{q}$ is also binary whose distribution is degenerate.

Utility function The speculator chooses precision η and size $q(x)$ to maximize her expected utility

$$\max_{q(\cdot), \eta} \mathbb{E}_x \left(q(x)[v(x, \eta) - a(\hat{x}, \hat{\eta})] \right), \tag{B.2}$$

anticipating the dealer's belief $(\hat{x}, \hat{\eta})$ determined by forward induction, as follows.

[24]The assumption that the dealer can observe the underlying joint distribution of trade size $\tilde{q}$ can be more formally micro-founded in a repeated trading game. At the beginning, the speculator hires researchers and builds infrastructure that determine its information precision η which cannot be changed. Then, the speculator trades with the dealer repeatedly and each period is an independent copy of the one-period trading game. The dealer can use past history to construct an empirical distribution of $\tilde{q}$. In the limit as the discount rate goes to 0, any payoff received during a transition period in which the dealer is still learning about the distribution of $\tilde{q}$ can be disregarded.

Dealer's belief Given any observed distribution $\mathbb{P}$ of $\tilde{q}$ (which can be different from the equilibrium distribution of $\tilde{q}$, constituting an off-the-equilibrium-path information set for the dealer) and a realized size q of an instant trade request, the dealer will apply forward induction to infer the signal realization x and the precision η as follows: the speculator maximizes

$$\begin{aligned} \max_{x(\cdot),\eta} \quad & \mathbb{E}_{\tilde{q}}\left(\tilde{q}[v(x(\tilde{q}),\eta) - a(\tilde{q},\mathbb{P})]\right) - c(\eta), \\ \text{subject to} \quad & x(\cdot) \text{ is injective and } x(\tilde{q}) \sim F \end{aligned}$$

The speculator solves the maximization problem above because choosing size $q(x)$ as a function of its signal realization x in its original maximization problem (B.2) is equivalent to choosing its inverse function $x : \mathbb{R} \mapsto \mathbb{R}$ subject to the constraint $x(\cdot)$ is injective and $x(\tilde{q}) \sim F$. Hence, the speculator has a dominant choice $(\hat{x}(\cdot,\mathbb{P}),\hat{\eta}(\mathbb{P}))$ of signal precision:

$$\begin{aligned} (\hat{x}(\cdot,\mathbb{P}),\hat{\eta}(\mathbb{P})) = \operatorname*{argmax}_{x(\cdot),\eta} \quad & \mathbb{E}_{\tilde{q}}\left(\tilde{q}[v(x(\tilde{q}),\eta) - a(\tilde{q},\mathbb{P})]\right) - c(\eta), \\ \text{subject to} \quad & x(\cdot) \text{ is injective and } x(\tilde{q}) \sim F \end{aligned} \tag{B.3}$$

Applying forward induction, the dealer's belief is $(\hat{x}(q,\mathbb{P}),\hat{\eta}(\mathbb{P}))$ at any given information set $(q,\mathbb{P})$, on or off the equilibrium path.

Now, we first solve the speculator's constrained maximization problem (B.3) given some distribution $\mathbb{P}$ of $\tilde{q}$, then the unconstrained maximization problem (B.2). The first order condition with respect to η is

$$\mathbb{E}_{\tilde{q}}\left(\tilde{q}\,\frac{\partial v}{\partial \eta}(\hat{x}(\tilde{q},\mathbb{P}),\hat{\eta}(\mathbb{P}))\right) = c'(\hat{\eta}(\mathbb{P})). \tag{FOC($\eta,\mathbb{P}$)}$$

Given some distribution $\mathbb{P}$ of size $\tilde{q}$, since the dealer applies forward induction to form its belief $(\hat{x}(\cdot,\mathbb{P}),\hat{\eta}(\mathbb{P}))$, the speculator's expected payoff is then

$$\mathbb{E}_{\tilde{q}}\, V_S(\hat{x}(\tilde{q},\mathbb{P}),\hat{\eta}(\mathbb{P})) - c(\hat{\eta}(\mathbb{P})) = \mathbb{E}_x\, V_S(x,\hat{\eta}(\mathbb{P})) - c(\hat{\eta}(\mathbb{P})).$$

The speculator chooses distribution $\mathbb{P}^*$ and precision $\hat{\eta}(\mathbb{P}^*) = \eta^*$ that maximizes her payoff

$$\hat{\eta}(\mathbb{P}^*) = \eta^* = \underset{\eta}{\operatorname{argmax}}\, \mathbb{E}_x\, V_S(x,\eta) - c(\eta). \tag{B.4}$$

The first order condition with respect to η is

$$\mathbb{E}_x\, \frac{\partial V_S}{\partial \eta}(x,\eta^*) = c'(\eta^*) \tag{FOC(η)}$$

The speculator chooses size $q^*(\cdot)$ that maximizes

$$q^*(\cdot) = \underset{q(\cdot)}{\operatorname{argmax}}\ \mathbb{E}_x\left(q(x)[v(x,\eta^*) - v(\hat{x}(q(x),\mathbb{P}_q),\hat{\eta}(\mathbb{P}_q))] + V_S(\hat{x}(q(x),\mathbb{P}_q),\hat{\eta}(\mathbb{P}_q))\right) - c(\eta^*),$$

where $\mathbb{P}_q = F \circ q^{-1}$ is the probability distribution on $\mathbb{R}$ induced by the measurable function $q : (\mathbb{R}, \mathcal{B}, F) \mapsto (\mathbb{R}, \mathcal{B})$. We show that

$$q^*(x) \text{ is strictly increasing in } x \text{ and } \forall x, q^*(x) = \frac{\frac{\partial V_S}{\partial x}(x,\eta^*)}{\frac{\partial v}{\partial x}(x,\eta^*)}. \tag{B.5}$$

We know from (B.4) that q^* induces probability distribution $\mathbb{P}^*$ such that $\hat{\eta}(\mathbb{P}^*) = \eta^*$ and that $\hat{x}(q^*(x),\eta^*) = x$ for every x. With $\mathbb{P} = \mathbb{P}^*$, (FOC($\eta, \mathbb{P}$)) becomes

$$\mathbb{E}_x\left(q^*(x)\,\frac{\partial v}{\partial \eta}(x,\eta^*)\right) = c'(\eta^*). \tag{B.6}$$

If for some $x < x'$,

$$q^*(x')[v(x,\eta^*) - v(x',\eta^*)] + V_S(x',\eta^*) > V_S(x,\eta^*), \tag{B.7}$$

then we can modify the value of q^* at x from $q^*(x)$ to $q^*(x')$. Then the speculator's resulting expected payoff is strictly higher since (1) her payoff upon receiving receiving signal x is strictly higher because of (B.7), and (2) her expected payoff is not affected by the change in the distribution of her requested size because of (FOC(η)) and (B.6). This contradicts the optimality of q^*. Therefore, for every $x < x'$,

$$q^*(x')[v(x,\eta^*) - v(x',\eta^*)] + V_S(x',\eta^*) \leq V_S(x,\eta^*).$$

Similarly,

$$q^*(x)[v(x',\eta^*) - v(x,\eta^*)] + V_S(x,\eta^*) \leq V_S(x',\eta^*).$$

This shows that $[q(\cdot), a(\cdot,\eta^*)]$ constitutes an incentive compatible direct mechanism. The proof of Lemma 3 establishes (B.5). This completes the solution to the speculator's problem.

The within-speculator comparison immediately follows:

Proposition B.1 (Within-speculator comparison) *A speculator receives a larger spread $a(x,\eta^*)$ when receiving a stronger signal $|x|$.*

Finally, we establish a sufficient condition for information chasing to dominate adverse selection in the across-speculator comparison:

Proposition B.2 (Across-speculator comparison) *If trade size $|q(x,\eta^*)|$ is decreasing in η^*, then a more informed trader receives a lower expected spread $\mathbb{E}\, a(x,\eta^*)$.*

Proof.

$$\begin{aligned}
\frac{\partial}{\partial \eta}\mathbb{E}_x\, a(x,\eta) &= \frac{\partial}{\partial \eta}\mathbb{E}_x\left(v(x,\eta) - \frac{V_S(x,\eta)}{q(x,\eta)}\right)\\
&= \mathbb{E}_x\left(\frac{\partial v}{\partial \eta}(x,\eta) - \frac{\frac{\partial V_S}{\partial \eta}(x,\eta)}{q(x,\eta)} + \frac{V_S(x,\eta)\frac{\partial q}{\partial \eta}(x,\eta)}{q^2(x,\eta)}\right)\\
&< \mathbb{E}_x\left(\frac{q(x,\eta)\frac{\partial v}{\partial \eta}(x,\eta) - \frac{\partial V_S}{\partial \eta}(x,\eta)}{q(x,\eta)}\right).
\end{aligned}$$

It suffices to show that $q(x,\eta)\frac{\partial v}{\partial \eta}(x,\eta) - \frac{\partial V_S}{\partial \eta}(x,\eta)$ and $q(x,\eta)$ are positively correlated, since then

$$\frac{\partial}{\partial \eta}\mathbb{E}_x\, a(x,\eta) < \mathbb{E}_x\left(q(x,\eta)\frac{\partial v}{\partial \eta}(x,\eta) - \frac{\partial V_S}{\partial \eta}(x,\eta)\right)\mathbb{E}_x\left(\frac{1}{q(x,\eta)}\right) = 0 \quad \text{(FOC}(\eta)\text{) and (B.6).}$$

Since $q(x,\eta)$ is increasing in x, it suffices to show that $q(x,\eta)\frac{\partial v}{\partial \eta}(x,\eta) - \frac{\partial V_S}{\partial \eta}(x,\eta)$ is increasing in x, or equivalently,

$$q\,\frac{\partial^2 v}{\partial x\,\partial \eta} + \frac{\partial q}{\partial x}\frac{\partial v}{\partial \eta} > \frac{\partial^2 V_S}{\partial x\,\partial \eta}. \tag{B.8}$$

Since $q(x,\eta) = \frac{\frac{\partial V_S}{\partial x}(x,\eta)}{\frac{\partial v}{\partial x}(x,\eta)}$ is decreasing in η, then

$$\begin{aligned} &\frac{\frac{\partial^2 V_S}{\partial x\,\partial \eta}\frac{\partial v}{\partial x} - \frac{\partial V_S}{\partial x}\frac{\partial^2 v}{\partial x\,\partial \eta}}{\left(\frac{\partial v}{\partial x}\right)^2} < 0, \\ \Longleftrightarrow &\frac{\partial^2 V_S}{\partial s\,\partial \eta} < q\,\frac{\partial^2 v}{\partial x\,\partial \eta}. \end{aligned}$$

The desired equality (B.8) thus follows since $\frac{\partial q}{\partial x} > 0$. ■

C Additional Tables and Figures

Table 2: Relative Trading Costs of Informed Clients: Using End-of-Day Quoted Prices as Benchmark Price

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Informed Clients	-0.828***	-0.806***	-0.702***	-0.596***	-0.651***	-0.633***	-0.647***
	(-4.53)	(-5.17)	(-4.03)	(-3.30)	(-3.46)	(-3.34)	(-3.04)
Client Size			-0.066*	-0.035	-0.117**	-0.167**	-0.073
			(-1.88)	(-0.90)	(-2.11)	(-2.57)	(-1.18)
Dealer-Connections				-0.025	-0.058**	-0.057**	-0.051**
				(-1.22)	(-2.45)	(-2.38)	(-2.18)
Client Intensity					0.258***	0.311***	0.261***
					(2.88)	(3.16)	(2.62)
Trade Size						0.048	-0.003
						(1.32)	(-0.08)
N	718418	716254	716254	716254	716254	716254	455403
R^2	0.011	0.168	0.168	0.168	0.168	0.168	0.426
Day FE	Yes	No	No	No	No	No	No
Bond FE	Yes	No	No	No	No	No	No
Dealer FE	Yes	No	No	No	No	No	No
Day*Dealer FE	No	Yes	Yes	Yes	Yes	Yes	No
Day*Bond FE	No	Yes	Yes	Yes	Yes	Yes	No
Day*Bond*Dealer FE	No	No	No	No	No	No	Yes

Notes: This table regresses trading costs (computed by 22 using the end-of-day quotes from Datastream as the benchmark price) on an informed sophisticated client dummy, various controls and various fixed effects. "Client Size" is the log of the average monthly trading volume of a given client. "Dealer-Connections" is the total number of unique dealers (averaged across months) that a given client trades with in a given month. "Client Intensity" is the log of the average monthly number of transactions of a given client. "Trade Size" is the log of the trade size in £s. Informed clients include those asset managers and hedge funds whose average scaled P&L measure 21 is in the top tertile. To reduce noise, we winsorise the sample at the 1-99%-level. T-statistics in parentheses are based on robust standard errors, using two-way clustering at the day and client level. Asterisks denote significance levels (* p<0.1, ** p<0.05, *** p<0.01).

Table 3: Relative Trading Costs of Sophisticated vs Unsophisticated Clients

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Sophisticated Clients	-0.316***	-0.323***	-0.255**	-0.236**	-0.234**	-0.237**	-0.208**
	(-2.61)	(-3.11)	(-2.42)	(-2.33)	(-2.39)	(-2.37)	(-2.11)
Client Size			-0.059***	-0.051**	-0.096***	-0.153***	-0.178***
			(-2.94)	(-2.29)	(-3.42)	(-4.36)	(-5.14)
Dealer-Connections				-0.006	-0.025*	-0.020	-0.021
				(-0.53)	(-1.69)	(-1.37)	(-1.48)
Client Intensity					0.141***	0.198***	0.232***
					(2.85)	(3.56)	(4.04)
Trade Size						0.054**	0.062***
						(2.57)	(2.82)
N	901054	901054	901054	901054	901054	901054	901054
R^2	0.007	0.100	0.100	0.100	0.100	0.100	0.290
Day FE	Yes	No	No	No	No	No	No
Bond FE	Yes	No	No	No	No	No	No
Dealer FE	Yes	No	No	No	No	No	No
Day*Dealer FE	No	Yes	Yes	Yes	Yes	Yes	No
Day*Bond FE	No	Yes	Yes	Yes	Yes	Yes	No
Day*Bond*Dealer FE	No	No	No	No	No	No	Yes

Notes: This table regresses trading costs (computed by 22 using the average transaction price at the bond-day-dealer level as the benchmark price) on dummy (taking value 1 for sophisticated clients and 0 for unsophisticated clients), various controls and various fixed effects. “Client Size” is the log of the average monthly trading volume of a given client. “Dealer-Connections” is the total number of unique dealers (averaged across months) that a given client trades with in a given month. “Client Intensity” is the log of the average monthly number of transactions of a given client. “Trade Size” is the log of the trade size in £s. Informed clients include those asset managers and hedge funds whose average scaled P&L measure 21 is in the top tertile. To reduce noise, we winsorise the sample at the 1-99%-level. T-statistics in parentheses are based on robust standard errors, using two-way clustering at the day and client level. Asterisks denote significance levels (* p<0.1, ** p<0.05, *** p<0.01).

Table 4: Relative Trading Costs of Informed Sophisticated vs Uninformed Sophisticated Clients

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Informed Clients	-0.361**	-0.377***	-0.350**	-0.380***	-0.409***	-0.372***	-0.474***
	(-2.52)	(-3.05)	(-2.59)	(-2.78)	(-3.18)	(-2.74)	(-3.64)
Client Size			-0.052*	-0.081***	-0.115***	-0.225***	-0.265***
			(-1.88)	(-2.76)	(-3.64)	(-5.78)	(-5.66)
Dealer-Connections				0.017	-0.011	0.006	0.007
				(1.22)	(-0.61)	(0.31)	(0.36)
Client Intensity					0.155**	0.235***	0.298***
					(2.51)	(3.48)	(3.97)
Trade Size						0.108***	0.118***
						(4.35)	(4.26)
N	612502	609573	609573	609573	609573	609573	545524
R^2	0.008	0.132	0.132	0.132	0.132	0.132	0.320
Day FE	Yes	No	No	No	No	No	No
Bond FE	Yes	No	No	No	No	No	No
Dealer FE	Yes	No	No	No	No	No	No
Day*Dealer FE	No	Yes	Yes	Yes	Yes	Yes	No
Day*Bond FE	No	Yes	Yes	Yes	Yes	Yes	No
Day*Bond*Dealer FE	No	No	No	No	No	No	Yes

Notes: This table regresses trading costs (computed by 22 using the average transaction price at the bond-day-dealer level as the benchmark price) on dummy (taking value 1 for informed sophisticated clients and 0 for uninformed sophisticated clients), various controls and various fixed effects. "Client Size" is the log of the average monthly trading volume of a given client. "Dealer-Connections" is the total number of unique dealers (averaged across months) that a given client trades with in a given month. "Client Intensity" is the log of the average monthly number of transactions of a given client. "Trade Size" is the log of the trade size in £s. Informed clients include those asset managers and hedge funds whose average scaled P&L measure 21 is in the top tertile. To reduce noise, we winsorise the sample at the 1-99%-level. T-statistics in parentheses are based on robust standard errors, using two-way clustering at the day and client level. Asterisks denote significance levels (* p<0.1, ** p<0.05, *** p<0.01).

Table 5: Relative Trading Costs of Informed Clients: Using Average 20-day Performance Instead of P&L

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Informed Clients	-0.443**	-0.332**	-0.312**	-0.337**	-0.360***	-0.364***	-0.324**
	(-2.18)	(-2.58)	(-2.37)	(-2.55)	(-2.72)	(-2.74)	(-2.33)
Client Size			-0.022	-0.030	-0.089***	-0.067*	-0.085**
			(-0.89)	(-1.14)	(-2.63)	(-1.67)	(-2.05)
Dealer-Connections				0.007	-0.012	-0.013	-0.007
				(0.43)	(-0.69)	(-0.75)	(-0.37)
Client Intensity					0.189***	0.166**	0.192***
					(3.06)	(2.49)	(2.69)
Trade Size						-0.021	-0.029
						(-0.93)	(-1.07)
N	437794	430536	430536	430536	430536	430536	339130
R^2	0.011	0.191	0.191	0.191	0.191	0.191	0.355
Day FE	Yes	No	No	No	No	No	No
Bond FE	Yes	No	No	No	No	No	No
Dealer FE	Yes	No	No	No	No	No	No
Day*Dealer FE	No	Yes	Yes	Yes	Yes	Yes	No
Day*Bond FE	No	Yes	Yes	Yes	Yes	Yes	No
Day*Bond*Dealer FE	No	No	No	No	No	No	Yes

Notes: This table regresses trading costs (computed by 22 using the average transaction price at the bond-day-dealer level as the benchmark price) on an informed sophisticated client dummy, various controls and various fixed effects. "Client Size" is the log of the average monthly trading volume of a given client. "Dealer-Connections" is the total number of unique dealers (averaged across months) that a given client trades with in a given month. "Client Intensity" is the log of the average monthly number of transactions of a given client. "Trade Size" is the log of the trade size in £s. Informed clients include those asset managers and hedge funds whose average 20-day ahead performance measure measure is in the top tertile. To reduce noise, we winsorise the sample at the 1-99%-level. T-statistics in parentheses are based on robust standard errors, using two-way clustering at the day and client level. Asterisks denote significance levels (* p<0.1, ** p<0.05, *** p<0.01).

Table 6: Relative Trading Costs of Informed Clients: Using Unscaled P&L to Measure Performance

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Informed Clients	-0.383***	-0.383***	-0.332***	-0.292**	-0.295**	-0.290**	-0.297**
	(-3.00)	(-3.42)	(-2.77)	(-2.46)	(-2.53)	(-2.43)	(-2.36)
Client Size			-0.027	-0.013	-0.073**	-0.133***	-0.143***
			(-1.10)	(-0.49)	(-2.18)	(-3.17)	(-3.54)
Dealer-Connections				-0.012	-0.034**	-0.030*	-0.026*
				(-0.86)	(-2.08)	(-1.86)	(-1.70)
Client Intensity					0.181***	0.242***	0.243***
					(3.25)	(3.83)	(3.64)
Trade Size						0.054**	0.060**
						(2.24)	(2.40)
N	732205	730076	730076	730076	730076	730076	685705
R^2	0.007	0.118	0.118	0.118	0.118	0.118	0.310
Day FE	Yes	No	No	No	No	No	No
Bond FE	Yes	No	No	No	No	No	No
Dealer FE	Yes	No	No	No	No	No	No
Day*Dealer FE	No	Yes	Yes	Yes	Yes	Yes	No
Day*Bond FE	No	Yes	Yes	Yes	Yes	Yes	No
Day*Bond*Dealer FE	No	No	No	No	No	No	Yes

Notes: This table regresses trading costs (computed by 22 using the average transaction price at the bond-day-dealer level as the benchmark price) on an informed sophisticated client dummy, various controls and various fixed effects. “Client Size” is the log of the average monthly trading volume of a given client. “Dealer-Connections” is the total number of unique dealers (averaged across months) that a given client trades with in a given month. “Client Intensity” is the log of the average monthly number of transactions of a given client. “Trade Size” is the log of the trade size in £s. Informed clients include those asset managers and hedge funds whose average unscaled/gross P&L measure 21 is in the top tertile. To reduce noise, we winsorise the sample at the 1-99%-level. T-statistics in parentheses are based on robust standard errors, using two-way clustering at the day and client level. Asterisks denote significance levels (* p<0.1, ** p<0.05, *** p<0.01).

Table 7: Informed Trading Volume and Informed Trading Costs

	(1)	(2)	(3)	(4)
	InfVol / Vol		log(InfVol)	
$InfCost^s_{i,t}$ [Trade Prices]	-0.002**		-0.021***	
	(-2.31)		(-3.63)	
$InfCost^s_{i,t}$ [Quoted Prices]		-0.002***		-0.037**
		(-3.18)		(-2.41)
Dealer's Trading Volume	-0.013**	-0.020***	1.004***	1.006***
	(-2.45)	(-2.73)	(21.18)	(19.18)
N	1505	1551	1505	1551
R^2	0.305	0.278	0.906	0.895

Notes: This table regresses the ratio of informed trading volume to total trading volume, "InfVol / Vol", (columns 1-2) and the natural logarithm informed trading volume, "log(InfVol)", (columns 3-4) on the average trading costs faced by informed clients at the given dealer. The term "$InfCost^s_{i,t}$ [Trade Prices]" refers to using the average transaction price (at the bond-day-dealer level) as the benchmark price in 22 to compute average trading cost for informed client i in month t. The term "$InfCost^s_{i,t}$ [Quoted Prices]" refers to using the end-of-day quote from Datastream as the benchmark price. To reduce noise, we winsorise the sample at the 1%-level. T-statistics in parentheses are based on robust standard errors, using clustering at the month level. Asterisks denote significance levels (* p<0.1, ** p<0.05, *** p<0.01).

Figure 5: The relationship between trading costs of informed clients at a dealer and the trading performance of unsophisticated clients against the given dealer: Day-Dealer level

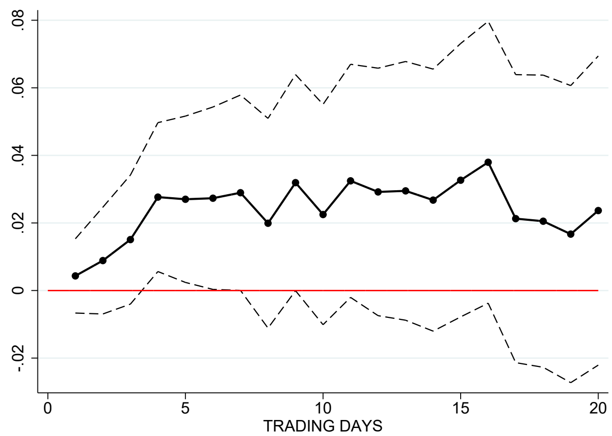

Notes: this figure plots the estimated β coefficients from our baseline daily regression 25 up to 20-day horizon ($T = 20$), using as regressand the average value weighted performance of uninformed clients trading with dealer i in month t. We include as a control the natural logarithm of the pound trade volume of dealers. To reduce noise, we winsorise the sample at the 1%-level. The shaded area denotes the 90% confidence band, based on robust standard errors, using clustering at the month level.

Figure 6: The relationship between trading costs and future capital gains amongst informed clients: Informed vs Unsophisticated Clients

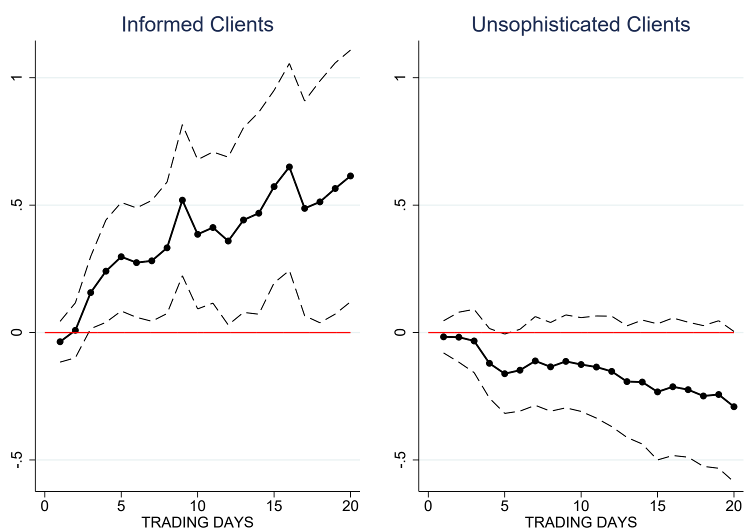

Notes: this figure plots the estimated β coefficients from our baseline monthly regression 26 up to 20-day horizon ($T = 20$), using as regressor (regressand) the unweighted average trading cost (anticipation component) of client i in month t. We include as a control the natural logarithm of the pound trade volume of clients. The left (right) panel shows the results for the sample which only includes informed sophisticated (unsophisticated) clients. To reduce noise, we winsorise the sample at the 1%-level. The shaded area denotes the 90% confidence band, based on robust standard errors, using clustering at the month level.

Figure 7: The relationship between trading costs and future capital gains amongst informed clients: trade size -weighted performance and cost measures

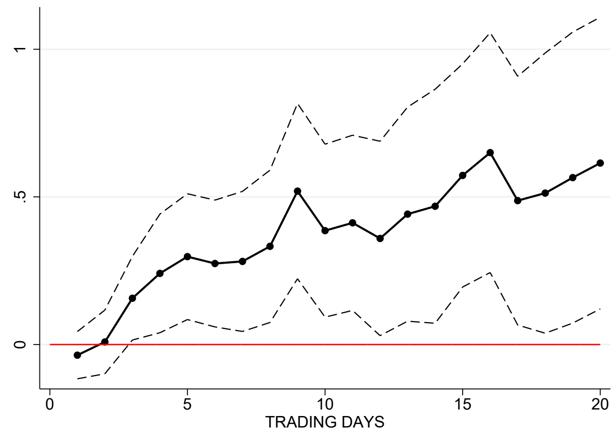

Notes: this figure plots the estimated β coefficients from our baseline monthly regression 26 up to 20-day horizon ($T = 20$), using as regressor (regressand) the size-weighted average trading cost (anticipation component) of client i in month t. We include as a control the natural logarithm of the pound trade volume of clients. The sample only includes informed sophisticated clients. To reduce noise, we winsorise the sample at the 1%-level. The shaded area denotes the 90% confidence band, based on robust standard errors, using clustering at the month level.

D Informed Trading and Trade Size

Given our baseline classification of clients, we explore whether informed clients trade in smaller quantities on average. To check this, we estimate the following transaction-level regression for client i, asset k, dealer m and day t:

$$TradeSize_{i,k,m,t} = \beta \times D_i^{Inf} + \mu_{k,t} + \delta_{m,t} + \varepsilon_{i,k,m,t}, \tag{D.1}$$

Tables 8–9 show that the average trade size of informed clients is about 0.5-1 log points lower than that of uninformed clients.

D.1 Controlling for Client Size

Table 8: Relative Trade Size of Informed vs Unsophisticated Clients

	(1)	(2)	(3)	(4)	(5)
Informed Clients	-0.588*	-0.498*	-1.267***	-0.562***	-0.615***
	(-1.81)	(-1.86)	(-5.98)	(-2.87)	(-3.54)
Client Size			0.503***	0.702***	0.604***
			(9.39)	(14.00)	(10.64)
Dealer-Connections				-0.168***	-0.129***
				(-6.05)	(-3.98)
N	730291	728093	728093	728093	469599
R^2	0.207	0.398	0.518	0.544	0.690
Day FE	Yes	No	No	No	No
Bond FE	Yes	No	No	No	No
Dealer FE	Yes	No	No	No	No
Day*Dealer FE	No	Yes	Yes	Yes	No
Day*Bond FE	No	Yes	Yes	Yes	No
Day*Bond*Dealer FE	No	No	No	No	Yes

Notes: This table regresses trade size (computed as the log of trade notional) on dummy (taking value 1 for informed sophisticated clients and 0 for unsophisticated clients), various controls and various fixed effects. "Client Size" is the log of the average monthly trading volume of a given client. "Dealer-Connections" is the total number of unique dealers (averaged across months) that a given client trades with in a given month. Informed clients include those asset managers and hedge funds whose average scaled P&L measure is in the top tertile. To reduce noise, we winsorise the sample at the 1-99%-level. T-statistics in parentheses are based on robust standard errors, using two-way clustering at the day and client level. Asterisks denote significance levels (* p<0.1, ** p<0.05, *** p<0.01).

Table 9: Relative Trade Size of More vs Less Informed Sophisticated Clients

	(1)	(2)	(3)	(4)	(5)
Informed Clients	-0.797***	-0.730***	-0.891***	-0.395**	-0.418***
	(-2.85)	(-3.06)	(-4.02)	(-2.56)	(-3.42)
Client Size			0.382***	0.844***	0.818***
			(7.58)	(17.55)	(20.86)
Dealer-Connections				-0.286***	-0.280***
				(-11.64)	(-14.34)
N	824757	823018	823018	823018	546716
R^2	0.252	0.408	0.484	0.545	0.689
Day FE	Yes	No	No	No	No
Bond FE	Yes	No	No	No	No
Dealer FE	Yes	No	No	No	No
Day*Dealer FE	No	Yes	Yes	Yes	No
Day*Bond FE	No	Yes	Yes	Yes	No
Day*Bond*Dealer FE	No	No	No	No	Yes

Notes: This table regresses trade size (computed as the log of trade notional) on dummy (taking value 1 for informed sophisticated clients and 0 for uninformed sophisticated clients), various controls and various fixed effects. "Client Size" is the log of the average monthly trading volume of a given client. "Dealer-Connections" is the total number of unique dealers (averaged across months) that a given client trades with in a given month. Informed clients include those asset managers and hedge funds whose average scaled P&L measure is in the top tertile. To reduce noise, we winsorise the sample at the 1-99%-level. T-statistics in parentheses are based on robust standard errors, using two-way clustering at the day and client level. Asterisks denote significance levels (* p<0.1, ** p<0.05, *** p<0.01).

D.2 Additional Results: Controlling for Intensity

Table 10: Relative Trade Size of Informed vs Unsophisticated Clients

	(1)	(2)	(3)	(4)	(5)
Informed Clients	-0.588*	-0.498*	-0.639**	-1.004***	-0.944***
	(-1.81)	(-1.86)	(-2.28)	(-3.29)	(-3.95)
Intensity			0.094*	-0.157	-0.128**
			(1.71)	(-1.62)	(-2.14)
Dealer-Connections				0.123***	0.104***
				(3.00)	(3.92)
N	730291	728093	728093	728093	469599
R^2	0.207	0.398	0.400	0.412	0.618
Day FE	Yes	No	No	No	No
Bond FE	Yes	No	No	No	No
Dealer FE	Yes	No	No	No	No
Day*Dealer FE	No	Yes	Yes	Yes	No
Day*Bond FE	No	Yes	Yes	Yes	No
Day*Bond*Dealer FE	No	No	No	No	Yes

Notes: This table regresses trade size (computed as the log of trade notional) on dummy (taking value 1 for informed sophisticated clients and 0 for unsophisticated clients), various controls and various fixed effects. "Intensity" is the log of the average monthly transaction number of a given client. "Dealer-Connections" is the total number of unique dealers (averaged across months) that a given client trades with in a given month. Informed clients include those asset managers and hedge funds whose average scaled P&L measure is in the top tertile. To reduce noise, we winsorise the sample at the 1-99%-level. T-statistics in parentheses are based on robust standard errors, using two-way clustering at the day and client level. Asterisks denote significance levels (* $p<0.1$, ** $p<0.05$, *** $p<0.01$).

Table 11: Relative Trade Size of More vs Less Informed Sophisticated Clients

	(1)	(2)	(3)	(4)	(5)
Informed Clients	-0.797***	-0.730***	-0.821***	-0.851***	-0.842***
	(-2.85)	(-3.06)	(-3.60)	(-3.70)	(-4.69)
Intensity			0.123***	-0.047	-0.070
			(2.74)	(-0.34)	(-0.71)
Dealer-Connections				0.065	0.065*
				(1.32)	(1.89)
N	824757	823018	823018	823018	546716
R^2	0.252	0.408	0.411	0.413	0.602
Day FE	Yes	No	No	No	No
Bond FE	Yes	No	No	No	No
Dealer FE	Yes	No	No	No	No
Day*Dealer FE	No	Yes	Yes	Yes	No
Day*Bond FE	No	Yes	Yes	Yes	No
Day*Bond*Dealer FE	No	No	No	No	Yes

Notes: This table regresses trade size (computed as the log of trade notional) on dummy (taking value 1 for informed sophisticated clients and 0 for uninformed sophisticated clients), various controls and various fixed effects. "Intensity" is the log of the average monthly transaction number of a given client. "Dealer-Connections" is the total number of unique dealers (averaged across months) that a given client trades with in a given month. Informed clients include those asset managers and hedge funds whose average scaled P&L measure is in the top tertile. To reduce noise, we winsorise the sample at the 1-99%-level. T-statistics in parentheses are based on robust standard errors, using two-way clustering at the day and client level. Asterisks denote significance levels (* p<0.1, ** p<0.05, *** p<0.01).

E Dynamics of Order Imbalance

In our theoretical model, the informed dealer is likely to mimic the trading direction of its informed clients during the dealer's subsequent trades against uniformed clients. In this section, we provide additional (albeit suggestive) evidence in support of this prediction, by studying the dynamic correlation structure of clients' order imbalance at the dealer-day level, distinguishing between the type of order imbalance initiated by informed and by uninformed clients. Specifically, we explore the idea that buy (sell) trades, initiated by informed clients are more likely to be followed by the dealer selling to (buying from) uninformed clients. To the extent that this mimicking behaviour by dealers drives the data, we expect the cross-correlation between the present informed order imbalance and future uninformed order imbalance to be negative compared to the autocorrelation in the order imbalance.

The previous literature used aggregate data to document the positive autocorrelation in aggregate order imbalance in stock markets (Chordia, Roll and Subrahmanyam 2002, 2005) and in government bond markets (Green 2004, Wang, Wu and Yu 2012).[25] Our contribution to this literature is twofold. First, we are able to analyse the persistence profile of order imbalance at a more disaggregated level, i.e. at the dealer-level. Second, we are able to decompose clients' order imbalance at the dealer-level into the part initiated by informed clients and the remaining that is related to uninformed clients.

Formally, we estimate the following dynamic regression model for dealer i, bond j, on day t:

$$Imb^k_{i,j,t} = \beta \times Imb^k_{i,j,t-1} + \varepsilon_{i,j,t}, \tag{E.1}$$

where $Imb^k_{i,j,t}$ is the order imbalance related to client type k which may include all clients, informed clients and uninformed clients, $k = \{Total, Inf, Un\}$. Informed clients are defined as the informed sophisticated clients, as defined above; uninformed clients include all other agents that the dealer trades with, including other clients as well as the inter-dealer sector. The term $Imb^k_{i,j,t}$ is measured as the net trading volume in pounds, scaled by the total trading volume (at the dealer-bond-day level).[26] The object of interest is the estimated value of β.

The upper panel of Table 12 reports the results at the bond-day-dealer level. Columns 1-3 show that there is a positive autocorrelation in total, informed and uninformed order

[25] The literature offered various expiations behind the positive autocorrelation in order imbalance, including information-induced order-splitting (Kyle 1985) and herding behaviour amongst others.

[26] Scaling the net trading volume by total trading volume serves as a standardisation device, which is motivated by the sizeable cross-sectional heterogeneity in the size of order imbalance across dealers (e.g. small vs. large dealers) and across bonds (e.g. on-the-run vs off-the-run bonds).

Table 12: The Dynamics of Order Imbalance

	$(Total_t, Total_{t-1})$	(Inf_t, Inf_{t-1})	(Un_t, Un_{t-1})	(Un_t, Inf_{t-1})	(Un_{t-1}, Inf_t)
	(1)	(2)	(3)	(4)	(5)
			Day-Dealer-Bond Level		
β	0.072***	0.119***	0.076***	-0.050***	0.013***
	(18.10)	(12.07)	(19.79)	(-7.13)	(6.46)
N	426392	90330	328060	138197	137325
R^2	0.005	0.014	0.006	0.001	0.000
			Day-Dealer-Maturity Bucket Level		
β	0.043***	0.079***	0.038***	-0.042***	0.003
	(11.73)	(11.55)	(9.73)	(-4.61)	(1.21)
N	119328	68263	115332	83350	83416
R^2	0.002	0.006	0.001	0.000	0.000

Notes: This table shows the estimation results for regression E.1. T-statistics in parentheses are based on robust standard errors, using clustering at the day-bond level (upper panel) and at the day-level (lower panel). Asterisks denote significance levels (* p<0.1, ** p<0.05, *** p<0.01).

imbalance. These results are consistent with the aggregate evidence of the existing literature (Chordia, Roll and Subrahmanyam 2002, 2005, Green 2004). The novel contribution of this empirical exercise is Column 4 of Table 12, which shows that there is a significantly negative relationship between lagged order imbalance of informed clients at a dealer and the contemporaneous order imbalance of uninformed clients at the given dealer. This is consistent with the learning mechanism in our theoretical model: a dealer is more likely to buy from (sell to) uninformed clients after it sold to (bought from) informed clients, because the transaction with informed clients revealed to the dealer that the asset value is more likely to be high (low).

Note that a negative dynamic correlation between order imbalances of different client groups at a dealer could also be driven by the dealer's need to dynamically rebalance its inventory. However, this explanation could equally apply to the case when we inspect the relationship between contemporaneous order imbalance of informed clients and the lagged order imbalance of uninformed clients. Column 5 of Table 12, however, shows that this relationship continues to be positive. This is an important cross-check, as it provides evidence (albeit suggestive) that our findings are more likely to be explained by our theoretical mechanism than simple inventory based explanations.

In the lower panel Table 12, we report the results after grouping together bonds in four maturity buckets.[27] The aggregation of bonds at different maturity segments builds on the previous literature (Brandt and Kavajecz 2004) and is motivated by the strong factor structure of government bonds. The findings continue to suggest that positive informed order imbalance against a dealer tends to be followed a negative uninformed order imbalance against the same dealer in the following trading day. Table 13 below shows that the results are similar when we include second lags as regressors in the regression.

[27]The four maturity buckets, based on years to maturity (YTM), include short-term (0-4 YTM), medium-term (4-9 YTM), long-term (9-21 YTM) and very long-term (¿21 YTM) bonds. These cut-offs are chosen to have an approximately even number of transactions in each bucket in our sample.

Table 13: The Dynamics of Order Imbalance: Including a Second Lag

Lag (days)	$(Total_t, Total_{t-1})$	(Inf_t, Inf_{t-1})	(Un_t, Un_{t-1})	(Un_t, Inf_{t-1})	(Un_{t-1}, Inf_t)
	(1)	(2)	(3)	(4)	(5)
			Day-Dealer-Bond Level		
$l = 1$	0.066***	0.117***	0.067***	-0.041***	0.011***
	(18.36)	(9.64)	(18.54)	(-5.00)	(4.76)
$l = 2$	0.037***	0.074***	0.030***	-0.022***	0.007***
	(9.64)	(6.81)	(8.12)	(-3.12)	(3.18)
N	317168	42915	234045	55686	94606
R^2	0.006	0.021	0.006	0.001	0.000
			Day-Dealer-Maturity Bucket Level		
$l = 1$	0.040***	0.067***	0.034***	-0.040***	0.003
	(10.66)	(9.66)	(8.65)	(-4.13)	(1.41)
$l = 2$	0.028***	0.050***	0.017***	-0.008	-0.000
	(7.96)	(7.27)	(4.64)	(-0.83)	(-0.12)
N	116234	57538	111100	66637	81278
R^2	0.002	0.007	0.001	0.000	0.000

Notes: This table shows the estimation results for a variant of regression E.1, where second lags too are included as regressors. T-statistics in parentheses are based on robust standard errors, using clustering at the day-bond level (upper panel) and at the day-level (lower panel). Asterisks denote significance levels (* p<0.1, ** p<0.05, *** p<0.01).