Information Chasing versus Adverse Selection in Over-the-Counter Markets

Gabor Pinter* Chaojun Wang† Junyuan Zou‡

Bank of England The Wharton School INSEAD

October 13, 2020

Latest Version

Abstract

Contrary to the prediction of the classic adverse selection theory, a more informed trader receives better pricing relative to a less informed trader in over-the-counter financial markets. Dealers aggressively chase informed orders to better position their future quotes and avoid winner's curse in subsequent trades. On a multi-dealer platform, dealers' incentive of information chasing exactly offsets their fear of adverse selection. In a more general setting of OTC markets, information chasing can dominate adverse selection when dealers face differentially informed speculators, while adverse selection always dominates when dealers face differentially informed trades from a given speculator. These two predictions—which contrast sharply with each other—both find strong empirical support in 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 reverses in over-the-counter (OTC) financial markets—instead of being deterred by adverse selection risk, dealers aggressively 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 submits a request-for-market (RFM) for a selected quantity of the asset, without revealing her desired trade direction, to a number of dealers simultaneously on the multi-dealer platform. Every dealer quotes a bid and an ask to the speculator, who then can choose one dealer to buy or sell at that dealer's respectively quoted price. The trade is not publicly disclosed. In Stage 2, a mass of liquidity traders send RFM 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 set quotes to liquidity traders. If, say, the informed speculator chooses to sell to a given dealer, then the asset payoff is likely to be low and the dealer would lower its quotes to liquidity traders 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. 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 being able to more accurately position their subsequent quotes to 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. Naturally, the dealer who wins the informed orders in Stage 1 provides the most informed quotes and earns the highest profit in Stage 2.

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 orders from a given speculator: 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. Within a given speculator, the resulting incentive compatibility condition is precisely sufficient and necessary for adverse selection to dominate information chasing. These two predictions simultaneously find strong support in the UK government bond market: *ceteris paribus*, a more informed trader receives on average 0.3 bps lower execution cost than a less informed trader, while a 1 bps increase in the execution cost of a trade of a given trader predicts that the price of the traded bond will move by 0.3 bps more in the opposite direction in 6 days.

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 and ultimately harms information production and price efficiency. This prediction is supported by empirical evidence in Lewis and Schwert (2018).

The most relevant paper is Naik, Neuberger and Viswanathan (1999), which 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 the trade's informativeness through the trader's identity and trade size. This approach yields distinctive predictions for within-trader versus across-trader comparisons, thus providing empirical guidance on where to locate evidence of information chasing. Two empirical papers, Ramadorai (2008); Bjønnes, Kathitziotis and Osler (2015), document a trading pattern that is consistent with information chasing in the foreign exchange market using independent data sources. However, the empirical pattern may also be consistent with non-informational mechanisms. We follow our own empirical guidance and simultaneously find evidence of information chasing in the cross-trader comparison, and evidence of adverse selection within trades originated from a given informed trader in the same UK government bond market. These two opposing 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 show that in OTC markets, dealers have an additional incentive to chase informed orders, which may very well dominate their fear of adverse selection. As a further distinction, Lee and Wang (2018), which 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 in parallel, which is the case for currency 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 speculators.

Our paper is related to the literature on information transmission in OTC Markets.[2] We explicitly model a dealer's incentive to chase informative 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.

The remaining of the paper is organized as follow: Section 2 sets up the benchmark model, and examines its equilibrium implications on pricing. Section 3 derives conditions for information chasing to dominate adverse selection in a more general setting, without

[1] Glosten and Milgrom (1985), Kyle (1985, 1989) and Vives (2011) provide theoretical benchmarks; A recent empirical paper, Collin-Dufresne, Junge and Trolle (2020), documents trading patterns that are consistent with adverse selection in the index-CDS market.

[2] Duffie and Manso (2007), Duffie, Malamud and Manso (2009), Duffie, Giroux and Manso (2010) and Duffie, Malamud and Manso (2014) show how information percolates in OTC Markets under different settings. Li and Song (2019) shows how a dealer can act as information intermediaries to channel information from informed to uninformed. These papers assume that information transmits in a reduced form manner.

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

This section sets up a benchmark model, and examines the pricing implications of information chasing.

2.1 Setup

There are three types of risk-neutral agents—one speculator, n dealers, 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.[3]

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 with a selected precision $\eta \in [0, 1]$. The binary signal s takes the value of 1 or -1 with equal probability, and the correlation between s and v is η. We assume that the information acquisition cost function c satisfies $c(0) = 0$, $c(1) = +\infty$, $\lim_{\eta \to 0} c'(\eta) = 0$, $\lim_{\eta \to 1} c'(\eta) = +\infty$, and $c''(\eta) > 0$ to insure a unique interior precision choice. The chosen precision and the realization of the signal are both private information of the speculator. The dealers have no additional information about the asset value v, assigning equal probability to the 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 as the trading protocol, as follows. In Stage 1, the speculator simultaneously requests a two-sided quote from the dealers to trade a selected

[3] Since the asset price will be bounded between -1 and 1, 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 $|\delta| > 2$.

size $q \geq 0$ of the asset, without indicating her desired trade direction.[4] Since purchase and sale are symmetric in the Stage 1 trading game, we consider the case where each given dealer j responds with 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. There is no post-trade transparency, which means that the liquidity traders and the other dealers do not observe the price and size of the trade. In Stage 2, each liquidity trader requests a bid-ask quote $(a_{2,j}, b_{2,j})$ simultaneously

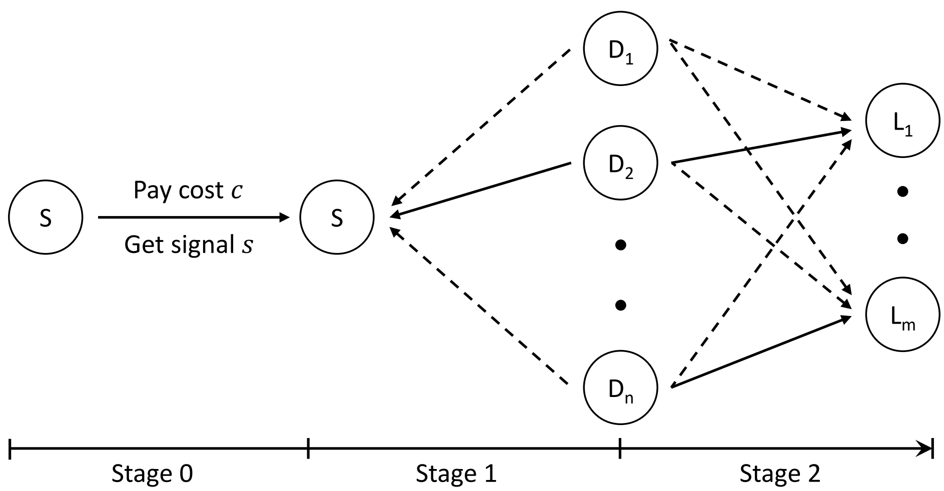

Figure 1: The Timeline.

from all dealers to trade 1 unit of asset. Since the liquidity trader's order is not informed, 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 common value of the asset is realized and the speculator and the dealers receive

[4]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.

the realized payoff of their asset position plus net payments they received from trades. The timeline and the market structure is summarized in Figure 1.

2.2 Equilibrium

Using backward induction, we show that there exists a unique perfect Bayesian equilibrium that satisfies forward induction.

Stage 2: Competing for Liquidity Trades

In Stage 1, the speculator buys upon receiving a positive signal and sells otherwise. Then at the beginning of Stage 2, one dealer is able to infer the signal realization by trading with the speculator in Stage 1 while other dealers remain uninformed about the signal realization. We let a dealer's belief regarding the signal precision be denoted by $\hat{\eta}$, which will be uniquely pined down by forward induction as a function of the order size q commonly observed by all the dealers. Therefore, all dealers hold the same belief $\hat{\eta}$.

Dealers' bidding strategies in Stage 2 depend on the number n_I of informed dealers. There are two cases: The speculator may reject all dealers' quotes in Stage 1, then $n_I = 0$; otherwise, $n_I = 1$. In equilibrium, it will turn out to be the case that there is $n_I = 1$ informed dealer in Stage 2, so that 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.

The next proposition summarizes dealers' unique mixed bidding strategies in Stage 2 constructed from results established in Syrgkanis et al. (2019). We denote the bid-ask quotes of the informed dealer by (b_2^+, a_2^+) when the signal is good, and (b_2^-, a_2^-) when the signal is bad. The bid-ask quotes of an uninformed dealer are denoted by (b_2^0, a_2^0).

Proposition 1 *If $n \geq 2$ and $n_I = 1$, the Stage-2 game has a unique equilibrium in which*

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

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

(ii) for the uninformed dealers, b_2^0 and $-a_2^0$ are drawn 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 distribution G_n describes a hybrid bidding strategy, in that an uninformed dealer bids $-\hat{\eta}$ with probability $\sqrt[n-1]{1/2}$, and with probability $1 - \sqrt[n-1]{1/2}$, it 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.

Proposition 2 *When there is one informed dealer in Stage 2, 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 becoming the only informed dealer is increasing in the mass of liquidity traders and the precision of the signal.

If $n_I = 0$, no dealer has information. Then all dealers quote the unconditional mean of the asset, setting $b_2^0 = a_2^0 = 0$. Therefore, the Stage-2 bid-ask spread is 0 and all dealers receive an expected payoff of 0.

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 $a_1(q)$ such that the following zero profit condition holds xfor dealers:

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

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.

Proposition 3 *In stage 1, given any order size q from the speculator, two or more dealers offer the same competitive mid-to-bid spread*

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

The speculator randomly selects one of the dealers who offer the smallest spread, independently from the realization of its signal. Upon selecting the dealer, the speculator buys if she receives a positive signal, and sells otherwise.

In equilibrium, the speculator's selection of the dealers cannot be correlated with its signal realization. Otherwise, those dealers who offer the smallest spread but fail to trade with the speculator can partially infer the signal realization, reducing information rent in stage 2. Thus, a dealer can profitably deviate by narrowing its spread by ϵ in order to prevent this information leakage. A formal proof of (2) is provided in the Appendix.

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 or 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 increase the 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 the orders from liquidity traders 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 their bid-ask spread.

From (2), we know the sign of the spread depends the relative strength of the two effects. If $q > m/2$, the per unit value of information of the informed order is small relatively to the adverse selection cost. Thus, The mid-to-bid spread $a_1^*(q)$ is positive. The reverse is true when $q < m/2$.[5] When $q = m/2$, the two effects exactly offset each other, which will turn out to be the case in equilibrium. Anticipating the equilibrium, the order size q will reveal the information acquisition effort of the speculator, and will be pinned down endogenously when we examine stage 0.

[5]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.

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 its expected payoff

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

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

The spreads $(a_{1,j}(q))_{j=1,\ldots,n}$ are independent of η since dealers do not 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 strategy for the speculator to choose the precision

$$\hat{\eta}(q) = c'^{-1}(q) \tag{3}$$

that equates the marginal benefit from trading with more precise information 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 communicate her signal precision through her order size.

Although the speculator sends the order size q to the dealers in Stage 1, she effectively makes a joint decision on η and q together. In Stage 0, the speculator anticipates some equilibrium pricing function a_1^* by the dealers and solves

$$(\eta^*, q^*) = \underset{\eta, q}{\operatorname{argmax}} \; q[\eta - a_1^*(q)] - c(\eta). \tag{4}$$

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.

Proposition 4 (Equilibrium) *In Statge 0 and Stage 1, a PBE satisfying forward induction consists of (1) the speculator's strategy* (η^*, q^*)*, (2) the dealers' pricing strategy* a_1^* *in Stage 1, and (3) dealer's belief* $\hat{\eta}(q)$ *that satisfy*

1. *Dealers' zero profit condition* (2)*,*
2. *Speculator's optimality condition* (4)*,*
3. *The forward induction condition* (3)*.*

Substituting dealers' equilibrium belief (3) into dealers' zero profit condition (2), we obtain dealers' equilibrium spread quoted to the speculator:

$$a_1^*(q) = c'^{-1}(q)\left(1 - \frac{m}{2q}\right). \tag{5}$$

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 (5) of the equilibrium spread $a_1^*(q)$, we can simplify the speculator's problem (4) into a one dimensional optimization problem over η,

$$\max_{\eta} \frac{m\eta}{2} - c(\eta). \tag{6}$$

Solving the optimization problem yield

$$\eta^* = c'^{-1}\left(\frac{m}{2}\right), \quad q^* = \frac{m}{2}. \tag{7}$$

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 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

$$a_1^*(q^*) = \eta^* \left(1 - \frac{m}{2q^*}\right) = 0. \tag{8}$$

Lemma 1 *The speculator receives a zero bid-ask spread in equilibrium.*

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, Lemma 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

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

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.[6]

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 (Proposition 1). Letting $\sigma(\eta)$ denote some homothetic measure of price dispersion for trades in Stage 2 with a given signal precision η,

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

then

$$\sigma(\eta) = \eta\sigma(1). \tag{10}$$

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.

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 (Proposition 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.

$$\tau(\eta) = 1 - \frac{\text{Var}[v|s]}{\text{Var}[v]} = \eta^2. \tag{11}$$

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. 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

$$v(\eta) = \eta, \quad V_I(\eta) = \frac{1}{2}m\eta, \quad \varphi = 1. \tag{12}$$

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. $v(0) = 0$, $V_I(0) = 0$, $v'(\eta) > 0$, $V_I'(\eta) > 0$,

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

It follows from (12) that Assumption 1 is a generalized version of the previously assumed 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.

$$qv'(\hat{\eta}) = c'(\hat{\eta}). \tag{13}$$

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

$$a_1^*(q) = v(\hat{\eta}(q)) - \frac{\varphi V_I(\hat{\eta}(q))}{q}. \tag{14}$$

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.

$$\varphi V_I'(\eta^*) = c'(\eta^*). \tag{15}$$

This determines the equilibrium level of information precision η^*. Combining (13), (14) and (15), we establish the equilibrium relationship among information precision, order size and the bid-ask spread.

Proposition 5 *In the generalized model, there exists a unique equilibrium in which*

$$q^* = \varphi \left. \frac{dV_I}{dv} \right|_{\eta=\eta^*}, \tag{16}$$

$$a_1^*(q^*) = v(\eta^*) \left[1 - \frac{1}{\varepsilon(\eta^*)} \right], \quad \textit{where } \varepsilon(\eta) = \frac{d \ln V_I}{d \ln v}. \tag{17}$$

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 5 is the trade size q. From (14), 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? (16) 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 5 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 (15). The forward induction condition (13) 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 (14) 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.

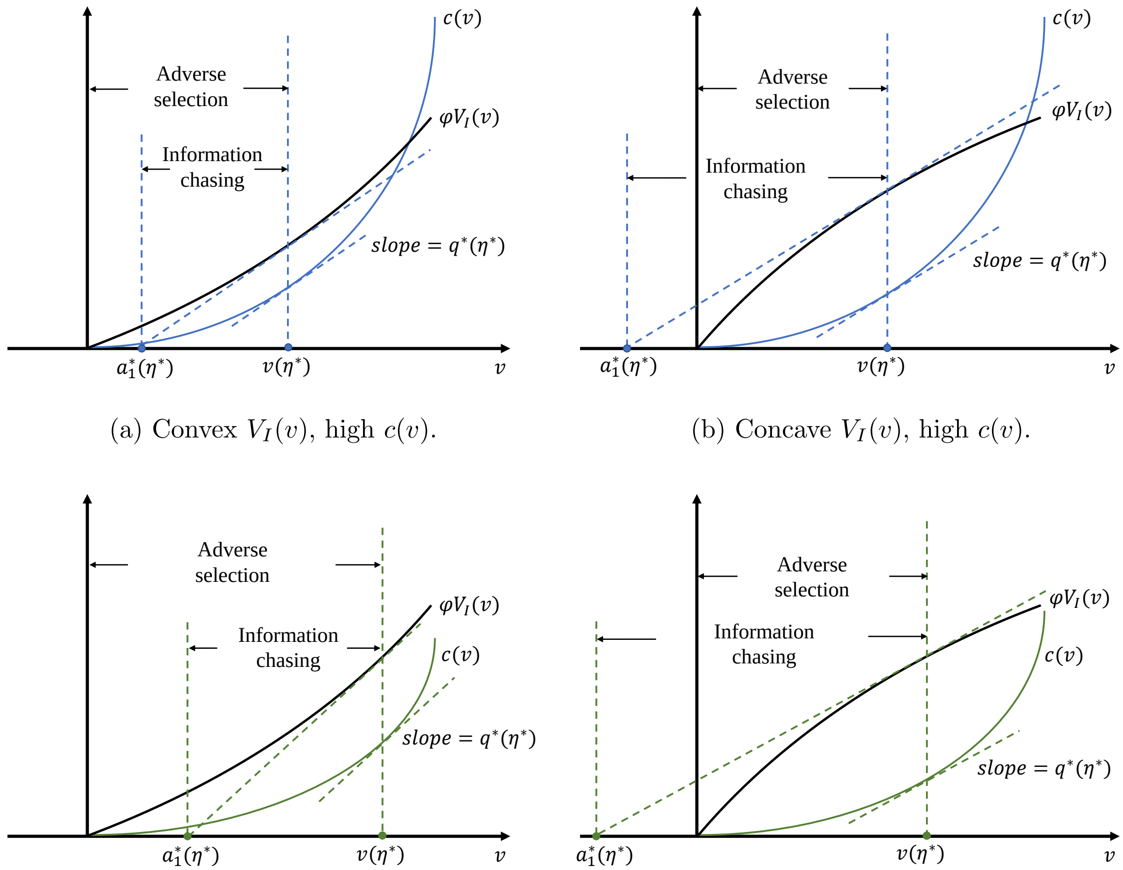

(a) Convex $V_I(v)$, high $c(v)$.

(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.

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

Proposition 6 (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 6 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).

Lemma 6 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

$$V_S(v) = [v - a(v)]q(v). \tag{18}$$

Lemma 2 (Myerson and Satterthwaite 1983) *In a symmetric Bayes-Nash equilibrium, $q(v)$ must be non-negative and weakly increasing in v for any $v > 0$, and*

$$a(v) = v - \frac{V_S(0) + \int_0^v q(z)\mathrm{d}z}{q(v)}. \tag{19}$$

Lemma 2 immmediately implies that the speculator trades weakly more when having more private information represented by a higher absolute value of v. Also, from (19) 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 6. We formally state this relationship in the proposition below.

Proposition 7 (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 7 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 informaiton v, which is now represented by $V_S(v)$, can deviate from $V_I(v)$, the value of inforamtion 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 synthetic 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 synthetic model, all the previous results in Proposition 6 and 7 hold with slight modifications.

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 October 2011 and June 2017. During this period, there are 21 primary dealers and 486 clients that we have identified. In our baseline classification, sophisticated clients include hedge funds and asset managers; whereas uninformed clients include insurance companies, pension funds, government entities (e.g. central banks) and non-financial corporations. We end up with 252 sophisticated clients and 234 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:

$$P\&L_i = \sum_{a}^{A_i} \left\{ \underbrace{\sum_{j^S_{i,a}=1}^{J^S_{i,a}} Q^S_{j^S_{i,a}} P^S_{j^S_{i,a}} - \sum_{j^B_{i,a}=1}^{J^B_{i,a}} Q^B_{j^B_{i,a}} P^B_{j^B_{i,a}}}_{Realised\ Cash-flow} + \underbrace{\left(\sum_{j^B_{i,a}=1}^{J^B_{i,a}} Q^B_{j^B_{i,a}} - \sum_{j^S_{i,a}=1}^{J^S_{i,a}} Q^S_{j^S_{i,a}} \right) \frac{1}{N_{i,a}} \sum_{m_{i,a}=1}^{M_{i,a}} P_{m_{i,a}}}_{Valuation\ of\ Inventories} \right\} \tag{20}$$

where $J^S_{i,a}$ and $J^B_{i,a}$ 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 (20) denote 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 use our measure $P\&L_i$ to split our sophisticated clients into two groups: one with $P\&L_i$ values below median and the remaining group with $P\&L_i$ values above the median.

We will refer to this latter group, consisting of 126 traders, as *informed sophisticated clients* in the remainder of the analysis.

4.2 Testing the Theory

4.2.1 Claim 1: Informed traders receive lower bid-ask spreads

To test the first claim predicted by the theory, we first construct a measure of execution 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 et al. (2018) and O'Hara and Zhou (2020)). Formally, for each trade j, on day t and bond k, we construct the measure $ExCost_{j,k,t}$ as follows:

$$ExCost_{j,k,t} = \left[\ln\left(P^{\star}_{j,k,t}\right) - \ln\left(P_{k,t}\right)\right] \times \mathbf{1}_j^{B,S}, \tag{21}$$

where $P^{\star}_{j,k,t}$ is the transaction price, $P_{k,t}$ is the benchmark price of the corresponding bond, and $\mathbf{1}_j^{B,S}$ 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 follow O'Hara and Zhou (2020) and use the average transaction price in the inter-dealer market around the time of the trade. As a robustness check, we also use the daily closing quoted mid-price of the corresponding bond, obtained from Datastream. The higher the measure $ExCost_{j,k,t}$ in (21), the less favourable the given client's execution costs are.

Given our measure of execution costs (21), which proxies the bid-ask spread, we estimate the following transaction-level regression for client i, asset k, dealer m and day t:

$$ExCost_{i,k,m,t} = \beta \times D_i^{Inf} + \gamma \times TradeSize_{i,k,m,t} + \mu_{k,t} + \delta_{m,t} + \varepsilon_{i,k,m,t}, \tag{22}$$

where D_i^{Inf} 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 key object of interest in (22) is β which captures how much more favourable the execution 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 Execution Costs of Informed Clients

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Informed Sophisticated	-0.388**	-0.367***	-0.327**	-0.309**	-0.226*	-0.252**	-0.275**
	(-2.45)	(-2.80)	(-2.21)	(-2.16)	(-1.84)	(-1.99)	(-2.22)
Client Intensity			-0.037	0.061	0.173**	0.347***	0.350***
			(-1.12)	(1.12)	(2.50)	(4.63)	(4.83)
Client Size				-0.092**	-0.053	-0.226***	-0.261***
				(-2.22)	(-1.44)	(-5.07)	(-5.59)
Dealer-Connections					-0.064***	-0.055***	-0.028*
					(-3.21)	(-2.70)	(-1.72)
Trade Size						0.147***	0.142***
						(5.52)	(5.40)
N	872763	870892	870892	870892	870892	870892	623850
r2	0.012	0.121	0.121	0.121	0.122	0.122	0.346
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 execution costs (computed by (21) using average inter-dealer transaction price as the benchmark price) on an informed sophisticated client dummy, various controls and various fixed effects. "Client Intensity" is the log of the average monthly number of transactions of a given client. "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. "Trade Size" is the log of the trade size in pounds. Informed sophisticated clients include those asset managers and hedge funds whose average P&L measure (20) is above the median. To reduce noise, we exclude client-day observations where client trades only once; and we winsorise the sample at the 2-98%-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 (22) 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 enables the comparison of the execution costs of different types of clients who trade the same bond, at the same dealer on the same day.

Overall, the results suggest that the execution costs on trades initiated by informed clients are about 0.3bp lower compared to trades initiated by unsophisticated clients.This is robust to the inclusion of a number of additional regressors that aim to control for mechanisms related to search and bargaining power. To show that our baseline is not simply picking up these mechanisms we first include "client intensity" which is the total number of transactions a client carries out in a given month averaged over the sample. This control is motivated by O'Hara et al. (2018) who used a subset of the US corporate bond market to analyse the execution costs of insurance companies.[7] Next, we include "client size" which is the average monthly trading volume of clients. Third, 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 be correlated with her execution costs as well as her type.[8] The inclusion of "dealer-connections" (column (5)) makes the largest difference to the estimated β coefficient. While this suggests that informed clients tend to have more dealer-connections, and clients with more connections tend to face more favourable execution costs[9], the informedness of clients seems to matter over and above what is captured by their trading network.

To show that our results are not simply picking up the effect of trade size (Edwards et al., 2007; Bernhardt et al., 2005), column (6) includes trade size as an additional control, with little effect on the baseline results.[10]

In our baseline, we used a benchmark price for our measure of execution costs that is

[7]They build on the search-theoretic literature (Duffie et al., 2005), and show that more active traders receive more favourable transaction prices than less active traders.

[8]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 execution costs in the US corporate bond market (Maggio et al., 2017; Hendershott et al., 2020).

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

[10]Note that the results on trade size are suggestive of evidence against the size discount, with larger trades receiving less favourable execution costs. In a related paper (Pinter et al., 2020), we show that this evidence for the size penalty is driven by the fact that client characteristics are controlled for, i.e. in a pooled regression without client-specific controls we find evidence for the size discount, consistent with the previous literature (Edwards et al., 2007; Bernhardt et al., 2005).

based on realised inter-dealer transaction prices. Table 6 in the Appendix shows the results for the case where we use the end-of-day mid-quote (from Datastream) as the benchmark price in (21). The results continue to hold up in this alternative specification.

4.2.2 Claim 2: the dealer anticipates larger gains against uninformed clients when the dealer gives tighter bid-ask spreads to informed clients

The heart of the mechanism in our theoretical model is the idea that dealers actively shape execution 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 et al. (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:

$$Perf^{T}_{j,k,t} = [\ln(P_{k,t+T}) - \ln(P_{k,t})] \times \mathbf{1}^{B,S}_{j}, \tag{23}$$

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 (23) for each dealer i, month t, and horizon T, based on all the transactions against unsophisticated clients as well as inter-dealer-brokers.[11] For aggregation, we take size-weighted averages using transaction size as weights. Similarly, we aggregate the transaction performance (21) for each dealer i and month t against informed sophisticated clients, by taking size-weighted averages. We then estimate the following panel regression at the dealer-month level:

$$Perf^{T,u}_{i,t} = \beta \times InfExCost^{s}_{i,t} + Vol_{i,t} + \alpha_i + \mu_t + \varepsilon_{i,t}, \tag{24}$$

[11] 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.

where $Perf_{i,t}^{T,u}$ is the trading performance of a group of clients u against dealer i on day t at horizon T. Group u includes unsophisticated clients as well as IDBs. The term $InfExCost_{i,t}^{s}$ is the average execution cost of group s, that includes informed sophisticated clients, against dealer i on day t. The term $Vol_{i,t}$ denotes trading volume and α_i and μ_t are dealer and month fixed effects.

Table 2: The relationship between execution costs of informed clients and trading performance of unsophisticated clients

	(1)	(2)	(3)	(4)	(5)	(6)
	1-day	2-day	3-day	4-day	5-day	6-day
Informed Execution Cost	0.079	0.134*	0.234**	0.234**	0.242***	0.262***
	(0.99)	(2.01)	(2.38)	(2.75)	(3.32)	(6.27)
Volume	0.606*	0.631	0.501	-0.346	-0.297	-0.537
	(1.84)	(1.41)	(1.14)	(-0.52)	(-0.30)	(-0.46)
N	1439	1438	1438	1438	1438	1438
R^2	0.084	0.098	0.080	0.071	0.066	0.071
Dealer FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows the estimation results for regression (24). To reduce noise, we winsorise the sample at the 1%-level. T-statistics in parentheses are based on robust standard errors, using two-way clustering at the month and dealer level. Asterisks denote significance levels (* p<0.1, ** p<0.05, *** p<0.01).

The coefficient of interest in (24) is β which we expect to be positive if our theoretical mechanism is at play: when informed sophisticated clients receive more favourable execution prices (i.e. the dealer charging lower execution costs on these trades), then the dealer is entering into trades against unsophisticated clients that turn out to generate capital gains over future horizon T for the dealer (and capital losses for those clients).

As shown in Table 2, a dealer's trading performance against unsophisticated clients over the 2-6 day horizon is higher when the given dealer offers more favourable execution prices to informed clients. 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.

4.2.3 Claim 3: an informed client has larger execution costs when she is more informed compared to when she is less informed

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

$$Perf_{i,t}^{T} = \beta \times ExCost_{i,t} + \mu_i + \delta_t + \varepsilon_i, \tag{25}$$

where $Perf_{i,t}^{T}$ is the average anticipation component (23) 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 execution cost increases 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 any time-invariant cross-sectional heterogeneity in trading performance and execution costs across clients.

Table 3: The relationship between execution costs and future capital gains of informed clients

	(1) 1-day	(2) 2-day	(3) 3-day	(4) 4-day	(5) 5-day	(6) 6-day
Informed Execution Cost	0.194***	0.244***	0.268***	0.209***	0.218***	0.262***
	(4.57)	(4.17)	(4.28)	(2.89)	(2.98)	(6.27)
N	11571	11564	11568	11569	11569	1438
R^2	0.036	0.036	0.037	0.035	0.037	0.071
Client FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows the estimation results for regression (25). To reduce noise, we winsorise the sample at the 1%-level. T-statistics in parentheses are based on robust standard errors, using two-way clustering at the month and client level. Asterisks denote significance levels (* p<0.1, ** p<0.05, *** p<0.01).

Table 3 presents the results, confirming that clients pay higher execution costs when their trades better predict future price movements over the 1-6 day horizon. One explanation for this is that clients trade larger amounts when they become more informed (Kyle, 1985).

Our companion paper (Pinter et al., 2020) provides further empirical evidence on trade size, execution costs and informativeness of clients.

4.2.4 Claim 4: average bid-ask spread is larger/smaller when there are more liquidity/informed traders

To test whether bid-ask spreads are smaller when there are more informed traders in the market, we proceed as follows. We first construct a measure of information ratio for each bond k on day t as follows:

$$InfVol_{k,t} = \frac{TradingVolume_{k,t}^{informed}}{TradingVolume_{k,t}^{total}}, \tag{26}$$

where $TradingVolume_{k,t}^{informed}$ is the trading volume generated by informed traders in bond k on day t and $TradingVolume_{k,t}^{total}$ is the total trading daily trading volume in the same bond. Given our measure, we estimate the following panel regression for bond k and trading day t:

$$BidAsk_{k,t} = \beta \times InfVol_{k,t} + controls_{k,t} + \mu_k + \delta_t + \varepsilon_{k,t}, \tag{27}$$

where $BidAsk_{k,t}$ denotes bid-ask spread quotes (obtained from Thomson Reuters Eikon), the term $controls_{k,t}$ includes the natural logarithm of daily trading volume and number of transactions in the given bond, and the terms μ_k and δ_t are bond and day fixed effects.

The results are reported in Table 4 for various specifications of regression (27). We find that on a trading day when trading volume in a bond consists entirely of transactions initiated by informed traders, the quoted bid-ask spread in the given bond is around 0.3bp smaller, compared to trading days when trading volume is generated by uninformed traders. These results are robust to the inclusion of volume and transaction number controls, i.e. the regression is not simply picking up informed traders trading different quantities and with different intensity, which could drive bid-ask spreads.

Table 4: Bid-Ask Spreads in the Presence of Informed Trading

	(1)	(2)	(3)	(4)	(5)	(6)
$InfVol$	-0.292***	-0.316***	-0.348***	-0.284***	-0.292***	-0.326***
	(-3.05)	(-3.21)	(-3.51)	(-4.17)	(-4.29)	(-4.80)
Trading Volume		-0.083**	-0.186***		-0.042***	-0.130***
		(-2.26)	(-4.47)		(-2.78)	(-6.80)
Number of Tran.			0.343***			0.328***
			(4.25)			(8.09)
N	55605	55605	55605	55602	55602	55602
R^2	0.434	0.434	0.435	0.609	0.609	0.610
Bond FE	Yes	Yes	Yes	Yes	Yes	Yes
Day FE	No	No	No	Yes	Yes	Yes

Notes: This table regresses bid-ask spreads on the measure $InfVol$, the natural logarithm of trading volume and of number of transactions and various fixed effects. Informed clients include asset managers and hedge funds. To reduce noise, we exclude client-day observations where client trades only once; and we winsorise the sample at the 1%-level. T-statistics in parentheses are based on robust standard errors, using clustering at the day level. Asterisks denote significance levels (* p<0.1, ** p<0.05, *** p<0.01).

4.2.5 Claim 5: price dispersion is higher when there are more liquidity traders in the market

To test this prediction, we first build on Jankowitsch et al. (2011) and compute the following price dispersion measure:

$$d_{k,t} = \sqrt{\frac{1}{K_{k,t}} \times \sum_{j=1}^{K_{k,t}} \left[\ln\left(P_{k,t}\right) - \ln\left(P^{\star}_{j,k,t}\right)\right]^2}, \tag{28}$$

where $d_{k,t}$ is the price dispersion measure on day t for bond k, and $P^{\star}_{j,k,t}$ is the transaction price and $P_{k,t}$ is the daily closing mid-rice of the corresponding bond, as used in definition (21) above. Given our measure, we estimate the following panel regression for bond k and trading day t:

$$d_{k,t} = \beta \times InfVol_{k,t} + controls_{k,t} + \mu_k + \delta_t + \varepsilon_{k,t}, \tag{29}$$

where the terms on the right-hand-side are the same as in specification (27).

The results are summarised in Table 5 for various specifications. Columns (1)-(3) show the case when bond fixed effects are included but time fixed effects are excluded in the regression. In this case, price dispersion is lower by about 2bp when the trading volume consists entirely of informed order flow, and the effect is little changed by including total trading volume and number of transactions as controls in the regression. Columns (4)-(6) show that the absolute value on the estimated β coefficients halves when we include time fixed effects. This is suggestive that aggregate shocks, such as the announcement of macroeconomic news, may affect both price-dispersion and the composition of trading volume.

Table 5: Price Dispersion in the Presence of Informed Trading

	(1)	(2)	(3)	(4)	(5)	(6)
$InfVol$	-2.278***	-1.885***	-1.996***	-1.134***	-1.012***	-0.987***
	(-7.63)	(-6.56)	(-6.95)	(-5.66)	(-5.13)	(-5.07)
Trading Volume		1.352***	0.992***		0.626***	0.693***
		(16.44)	(12.43)		(11.99)	(11.09)
Number of Tran.			1.197***			-0.249*
			(5.02)			(-1.67)
N	55605	55605	55605	55602	55602	55602
R^2	0.319	0.331	0.332	0.572	0.574	0.574
Bond FE	Yes	Yes	Yes	Yes	Yes	Yes
Day FE	No	No	No	Yes	Yes	Yes

Notes: This table regresses our measure of price dispersion on $InfVol$, the natural logarithm of trading volume and of number of transactions and various fixed effects. Informed clients include asset managers and hedge funds. To reduce noise, we exclude client-day observations where client trades only once; and we winsorise the sample at the 1%-level. T-statistics in parentheses are based on robust standard errors, using clustering at the day level. Asterisks denote significance levels (* p<0.1, ** p<0.05, *** p<0.01).

5 Conclusion

Contrary to the prediction of the classic adverse selection theory, informed trades receive better pricing relative to uninformed trades 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

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

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.

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

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.

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.

_ , Nicolae Garleanu, and Lasse Heje Pedersen, "Over-the-Counter Markets," *Econometrica*, November 2005, *73* (6), 1815–1847.

_ , 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.

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.

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.

Jankowitsch, Rainer, Amrut Nashikkar, and Marti G. Subrahmanyam, "Price dispersion in OTC markets: A new measure of liquidity," *Journal of Banking and Finance*, 2011, *35* (2), 343–357.

Kondor, Péter and Gabor Pinter, "Clients' Connections: Measuring the Role of Private Information in Decentralised Markets," CEPR Discussion Papers 13880, C.E.P.R. Discussion Papers 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.

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.

O'Hara, Maureen and Xing Zhou, "Anatomy of a Liquidity Crisis: Corporate Bonds in the Covid-19 Crisis," *Journal of Financial Economics*, 2020, (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, "Execution Costs in Government Bond Markets," mimeo October 2020.

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

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.

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

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 construct an equilibrium of such auction games without a proof. This lemma is established in Syrgkanis et al. (2019).

Lemma A.1 *If $n = 2$, in a common-value first-price sealed bid auction with discrete signals, there exists a unique mixed 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) correspond 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 Proposition 1. 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. Let G^+, G^- and G^0 be the c.d.f. of b_2^+, b_2^- and b_2^0. Lemma A.1 implies that

1. The informed dealer bid $b_2^- = -\eta$ 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 $-\eta$ 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 $-\eta$.

Let g^0 be the probability that $b_2^0 = -\eta$. The c.d.f. of b_2^+ is denoted by G^+, while the c.d.f. of b_2^0 conditional on $b_2^0 > -\eta$ is denoted by G^0. An uninformed dealer must be indifferent of bidding any $b \in supp\ G^+$, therefore

$$\frac{1}{2}G^+(b)[g^0 + (1-g^0)G^0(b)]^{n-2}(\eta - b) + \frac{1}{2}[g^0 + (1-g^0)G^0(b)]^{n-2}(-\eta - b) = C^0.$$

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

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

The upper bound of *supp* G^+ is 0. An informed dealer must be indifferent of bidding any $b \in supp\ G^0$, therefore

$$(\eta - b)[g^0 + (1-g^0)G^0(b)]^{n-1} = C^+.$$

Let $b = 0$ we have $C^+ = \eta$. Let $b \to -\eta$, we have $g^0 = \sqrt[n-1]{\frac{1}{2}}$. Plugging into the previous equation, we have

$$G^0(b) = \frac{\sqrt[n-1]{\frac{1}{1-b/\eta}} - \sqrt[n-1]{\frac{1}{2}}}{1 - \sqrt[n-1]{\frac{1}{2}}}.$$

■

Proof of Proposition 6. [To be completed. . .] ■

Proof of Proposition 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 $\eta \cdot \frac{1}{2}m$. It is obvious that uniformed dealers has zero expected profit from trading. ■

B A Synthetic Model

Suppose the expected value of asset is $v(\eta, x)$, a function of the speculator's effort of information acquisition η and her private signal x. The private signal is draw from a continuous distribution $F(x; \lambda)$.[12] The value of information to the winning dealer is $V_I(v(\eta, x))$, i.e., v is a sufficient statistic of η and x in determining the value of information. Each speculator has a convex cost function of information acquisition $c(\eta; \lambda)$, parameterized by $\lambda \in \Lambda$. Assume for any η, $c'(\eta; \lambda)$ is strictly decreasing in λ.

Now consider a trading game with ex-ante competition—each dealer offers a menu to the speculator before the speculator does anything, and the speculator chooses the best menu and commit to trading with that dealer.[13] Notice a speculator with the same v but different combination of (η, x) combination will trade exactly the same way. Therefore, a direct mechanism specifies the trading quantity $q(v; \lambda)$ and the speculator's total payment to the dealer $P(v; \lambda)$. We include λ as a parameter of the contract since λ is observable to all dealers. Given the menu, the speculator's payoff is the trading profit $V_S(v; \lambda) = vq(v; \lambda) - V_S(v; \lambda)$.

[12] Assumption: x and η are separable, or we can map the signal to a random variable x the distribution of which does not depend on η.

[13] The equilibrium is equivalent under the timing assumption that the speculator decides η first

The half spread is

$$a(v;\lambda) = v - \frac{V_S(v;\lambda)}{q(v;\lambda)}. \tag{B.1}$$

Lemma B.1 (IC constraint) *A (direct) menu must have $q(v;\lambda)$ weakly increasing in v, and*

$$V_S(v;\lambda) = V_S(0;\lambda) + \int_0^v q(z;\lambda)\mathrm{d}z. \tag{IC for x}$$

$q(v;\lambda)$ is an increasing function of v implies that $V_S(v;\lambda)$ is a convex function of v. Therefore, the half spread

$$a(v;\lambda) = v - \frac{V_S(v;\lambda)}{V_S'(v;\lambda)} \tag{B.2}$$

in increasing in v.

Proposition B.1 (Within-speculator trade heterogeneity) *A speculator receives a larger bid-ask spread when having larger $|v|$.*

Now let's consider compare speculators with different λ. A speculator solves the following problem

$$\max_{\eta}\ \mathbb{E}_x V_S(v(\eta,x);\lambda) - c(\eta;\lambda). \tag{B.3}$$

The first order condition is

$$\frac{\mathrm{d}}{\mathrm{d}\eta}\mathbb{E}_x V_S(v(\eta,x);\lambda) = c'(\eta;\lambda). \tag{B.4}$$

The IC constraint implies that

$$V_S'(v(\eta, x); \lambda) = q(v(\eta, x); \lambda). \tag{B.5}$$

Therefore

$$\mathbb{E}_x q(v(\eta, x); \lambda) \frac{\partial}{\partial \eta} v(\eta, x) = c'(\eta; \lambda). \tag{B.6}$$

Assumption B.1 $v(\eta, x) = \eta \tilde{v}(x)$.

Given Assumption B.1, the LHS of (B.6) can be viewed as a weighted average of speculator λ's trade size. It equals the speculator's marginal cost of information acquisition in equilibrium.

Competitive dealers solve the following problem

$$\max_{V_s(0;\lambda), q(\cdot;\lambda), \eta(\lambda)} \quad \mathbb{E}_x V_s(v(\eta(\lambda), x); \lambda) - c(\eta(\lambda); \lambda), \tag{B.7}$$

$$\text{s.t.} \quad \mathbb{E}_x V_S(v(\eta(\lambda), x); \lambda) = \mathbb{E}_x V_I(v(\eta(\lambda), x)), \tag{B.8}$$

$$\eta(\lambda) = \underset{\tilde{\eta}}{\operatorname{argmax}} \; \mathbb{E}_x V_S(v(\tilde{\eta}, x); \lambda) - c(\tilde{\eta}; \lambda), \tag{B.9}$$

$$V_S(v; \lambda) = V_S(0; \lambda) + \int_0^v q(z; \lambda) \mathrm{d}z, \tag{B.10}$$

$$q(\cdot; \lambda) \text{ is non-decreasing.} \tag{B.11}$$

Notice that dealers break even in expectation, which implies that the winning dealer might loss money given certain realizations of x.

We conjecture that constraints (B.9), (B.10) and (B.11) are not binding in the above problem. To show this, consider the $\eta^*(\lambda)$ which satisfies

$$\frac{d}{d\eta}\mathbb{E}_x V_I(v(\eta, x)) = c'(\eta; \lambda). \tag{B.12}$$

We want to show that there must exist a $V_s(0, \lambda)$ and $q(\cdot, \lambda)$ which satisfies (B.8), (B.9), (B.10) and (B.11). (B.9) can be replaced by

$$\mathbb{E}_x q(v(\eta, x); \lambda)\tilde{v}(x) = c'(\eta; \lambda). \tag{B.13}$$

$$\mathbb{E}_x \frac{\partial}{\partial v} q(v(\eta, x); \lambda)\tilde{v}(x)^2 \leq c''(\eta; \lambda) \tag{B.14}$$

We use guess and verify. Let

$$q(v, \lambda) = \frac{\tilde{q}(\lambda)}{\eta^*(\lambda)} v. \tag{B.15}$$

$\tilde{q}(\lambda)$ must satisfy

$$\tilde{q}(\lambda) = \frac{c'(\eta(\lambda), \lambda)}{\mathbb{E}_x \tilde{v}(x)^2}, \tag{B.16}$$

$$\tilde{q}(\lambda) \leq \frac{c''(\eta(\lambda), \lambda)\eta(\lambda)}{\mathbb{E}_x \tilde{v}(x)^2}. \tag{B.17}$$

The above two equations hold for all λ if and only if $c''(\eta, \lambda) \geq c'(\eta(\lambda), \lambda)\eta(\lambda)$ for any λ. A sufficient condition for this is that $c''' \geq 0$.

Notice that (B.8) can be written as

$$V_s(0, \lambda) = \mathbb{E}_x V_I(v(\eta^*(\lambda), x)) - \mathbb{E}_x \int_0^v q(z; \lambda)\mathrm{d}z \tag{B.18}$$

$$= \mathbb{E}_x V_I(v(\eta^*(\lambda), x)) - \frac{1}{2}\tilde{q}(\lambda)\eta^*(\lambda)\mathbb{E}_x \tilde{v}(x)^2, \tag{B.19}$$

$$= \mathbb{E}_x V_I(v(\eta^*(\lambda), x)) - \frac{1}{2}\eta^*(\lambda)c'(\eta^*(\lambda), \lambda). \tag{B.20}$$

The RHS is an increasing function in η. When λ increase, $\eta(\lambda)$ always increases. If V_I is convex in v, LHS is a increasing function. When λ increase, $c'(\eta(\lambda);\lambda)$ increases.[14] If V_I is concave in v, LHS is a decreasing function. When λ increase, $c'(\eta(\lambda);\lambda)$ decreases.

Proposition B.2 (Across-speculator trade heterogeneity) *A speculator with higher λ (lower cost of information acquisition) chooses higher η*

1. *If V_I is convex in v, on average a speculator with higher λ trades larger quantity q.*
2. *If V_I is concave in v, on average a speculator with higher λ trades smaller quantity q.*

[To be completed...]

C Additional Tables and Figures

[14] c is more convex in v.

Table 6: Relative Execution Costs of Informed Clients - Datastream Prices

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Informed Sophisticated	-0.583***	-0.573***	-0.559***	-0.532***	-0.415***	-0.432***	-0.465***
	(-3.69)	(-4.18)	(-3.52)	(-3.41)	(-2.78)	(-2.88)	(-2.92)
Client Intensity			-0.013	0.137**	0.291***	0.395***	0.361***
			(-0.33)	(1.99)	(3.80)	(4.70)	(4.10)
Client Size				-0.142**	-0.086*	-0.189***	-0.184***
				(-2.58)	(-1.71)	(-3.13)	(-2.95)
Dealer-Connections					-0.088***	-0.083***	-0.062***
					(-4.54)	(-4.23)	(-3.26)
Trade Size						0.087***	0.073**
						(2.70)	(2.10)
N	939437	938259	938259	938259	938259	938259	673745
r2	0.009	0.121	0.121	0.121	0.121	0.121	0.365
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 execution costs (computed by 21 using end-of-day Datastream mid-quote as the benchmark price) on an informed sophisticated client dummy, various controls and various fixed effects. "Client Intensity" is the log of the average monthly number of transactions of a given client. "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. "Trade Size" is the log of the trade size in pounds. Informed sophisticated clients include those asset managers and hedge funds whose average P&L measure 20 is above the median. To reduce noise, we exclude client-day observations where client trades only once; and we winsorise the sample at the 2-98%-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 (* $p<0.1$, ** $p<0.05$, *** $p<0.01$).