

Asymmetric Information and OTC Bid-ask Spreads

Geir Bjønnes^a

Neophytos Kathitziotis^b

Carol Osler^c

Abstract

This study identifies striking differences between OTC and order-driven markets in the influence of asymmetric information. Using highly detailed data from the world's largest OTC market, the paper first provides evidence that OTC bid-ask spreads are driven by two information asymmetries, rather than one. The paper also provides evidence that OTC dealers adopt three strategic responses to client information rather than one: widen spreads, consistent with adverse selection; narrow spreads, consistent with information chasing; and no response, consistent with uninformed clients. Adverse selection is more likely for clients engaged in HFT.

3 January 2026

^a geir.bjønnes@bi.no. Department of Finance, BI Norwegian Business School, Nydalsveien 37, 0484 Oslo, Norway. +47-46-410-502.

^b neophytos.kathitziotis@gmail.com. P.O. Box 33772, 5317 Paralimni, Cyprus. Tel.: +49 160 456 21 81

^c cosler@brandeis.edu. Corresponding author. School of Business and Economics, Mailstop 32, Brandeis University, 415 South Street, Waltham, MA 02454. +1 781-736-4826.

This paper examines the influence of asymmetric information on OTC liquidity prices, building on existing evidence of major differences between OTC and order-driven markets. Abundant research establishes that asymmetric information generates adverse selection risk in order-driven markets and widens the bid-ask spread by an amount that is common to all traders (e.g., Glosten and Milgrom, 1985; Glosten and Harris, 1988; Huang and Stoll, 1997). The growing literature on liquidity pricing in OTC markets¹ as yet includes few studies of asymmetric information, but these papers highlight major differences between OTC and order-driven markets. First, the influence of information on OTC spreads should vary across clients (Naik et al., 1999). Second, OTC dealers quote narrower spreads to better-informed clients, a strategy known as “information chasing,” rather than quoting wider spreads to these clients as predicted by adverse selection (Ramadorai, 2008; Pintér et al., 2021). Indeed, researchers have been unable to find evidence for adverse selection in OTC markets (Bernhardt et al., 2005; Osler et al., 2011).

We provide evidence for two new contrasts between OTC and order-driven markets. First, OTC client execution costs are determined by two information asymmetries rather than one. The client’s markup is influenced by asymmetric information between dealer and client while the rest of the client’s half-spread is influenced by asymmetric information among dealers and other professional traders. Second, OTC dealers employ not one but three strategies in response to client information. For some client types they narrow spreads in response to client information, consistent with information chasing; for other client types they widen spreads with client information, consistent with adverse selection; and for uninformed clients dealers make no adjustment for asymmetric information. Taken as a whole, the evidence now indicates that the influence of asymmetric information on OTC spreads has but little in common with its influence on order-driven spreads.

We exploit the complete trading record of a major dealing institution in one of the world’s most active OTC contracts, EUR-USD. Daily trading in EUR-USD was worth \$578 billion in April 2025 (B.I.S., 2025), which exceeds daily trading in all US and European equities plus all US corporate and municipal bonds.^{2,3} Our analysis thus sheds light on a contract of global importance as well as OTC markets as a class. The data, which include roughly 500K trades worth close to €400 billion, enable us to distinguish clients from professional traders because they include counterparty identifiers and ten

¹E.g., Edwards et al. (2007); Goldstein et al. (2007); Green et al. (2007); Harris and Piwowar (2007); Hotchkiss et al. (2008); O’Hara et al. (2018); Hendershott et al. (2020); O’Hara and Zhou (2021); Hau et al. (2021); Feldhütter and Poulsen (2018); Friewald and Nagler (2024); Collin-Dufresne and Fos (2015); Collin-Dufresne, et al. (2020).

² Bank for International Settlements (2022). [Triennial Survey of Foreign Exchange Turnover in April 2025](#).

³ US equities, \$486 billion, [CBOE](#). European equities: €59 billion, CBOE [Global Markets](#). US corporate bonds, \$57 billion, [MarketAxess](#). US municipal bonds, roughly \$15 billion, [MarketAxess](#).

counterparty types. The data also enable us to distinguish the influence of client information from the influence of other important spreads determinants because they include additional unusual details such as trading platform, precise markup, and the involvement of a prime broker.

OTC client execution costs are influenced by two different information asymmetries due to the two-tier structure of OTC markets (Naik et al., 1999). In the primary tier dealers trade with each other and, in FX, with other professional traders under prime brokerage contracts – an expansion that motivates the broader label, “core market.” In the second tier, dealers provide liquidity to clients, setting each client quote as a markup relative to the same-side core-market price. Each client half-spread is thus the sum of the core-market half-spread and a client-specific markup. Asymmetric information among professional traders influences the core-market half-spread. Asymmetric information between client and dealer influences the markup.

We estimate the adverse-selection component of core-market spreads using prices from the bank’s liquidity-providing trades on core-market order-driven platforms. The Huang and Stoll (1997) spread decomposition model indicates that adverse selection in the core market contributes 25% of the mean client half-spread. For the clients most sensitive to execution costs – specifically hedge funds, commodity trading advisors, and brokers – this share exceeds 70%.

OTC dealers can adopt multiple strategies in response to client information due to three distinguishing features of the OTC trading game. First, liquidity demanders move first, requesting quotes from a specific dealer and thereby revealing their identity; in order-driven markets this sequence is reversed and counterparties are anonymous. Second, dealer-client relationships are typically long-lasting (Hendershott et al., 2020; Hau et al., 2021), which enables dealers to become familiar with each client’s tendency to be informed. Third, clients face non-trivial search costs (Duffie et al, 2005), so dealers have instantaneous market power. Together, these imply that OTC dealers can condition client spreads on client characteristics including information (Naik et al., 1999; Ramadorai, 2008).

Client information inevitably creates adverse-selection risk for dealers, which could motivate wider markups for better-informed clients (Glosten and Milgrom, 1985; Easley and O’Hara, 1987). However, dealers can sometimes profit from client information, which could motivate them to quote narrower markups for better-informed clients. Naik et al. (1999) show that dealers can profit by taking parallel proprietary positions in the core market; Pintér et al. (2022) show that dealers can profit by astutely revising their quotes to other agents. If a client is uninformed or if the benefits and costs of a client’s information are roughly equal, rational dealers could make no information-based adjustment to the client markup (Naik et al., 1999).

We develop and estimate a comprehensive markup decomposition model tailored to the OTC context. The model includes client information, additional relevant client characteristics, order-handling and inventory costs (Glosten and Harris, 1988; Huang and Stoll, 1997), plus intertemporal smoothing, a strategy for maintaining client loyalty. Dealers can vary their strategic response to client information across client types. Client information is measured as mean signed post-trade returns (Anand and Subrahmanyam, 2008); the model is estimated using GMM with robust standard errors clustered by date.

Results indicate that dealers adopt all three of the potentially rational responses to client information. Consistent with information chasing they narrow the spread for better-informed brokers. Consistent with adverse selection they widen the spread for better-informed hedge funds, commodity trading firms, and client banks. For clients they consider uninformed – specifically, non-financial clients and low-leverage asset managers such as pension funds or mutual funds – dealers make no information adjustments. Results are consistent across multiple robustness tests.

The paper finishes by investigating the dealers' choice between adverse selection and information chasing. We hypothesize that it depends in part on client trades latency. FX dealers consistently sought to attract better-informed clients until the early 2000s (Clyde, 1996; Ramadorai, 2008). Soon after HFT reached FX in the early 2000s, however, they began expressing concern (Markets Committee, 2011). HFT strategies such as latency arbitrage added new dimensions of adverse selection for dealers, especially those with slow quote-updating technologies (Menkveld and Zoican, 2017). We hypothesize that trading with informed low-latency clients brings relatively high costs and low benefits, leading dealers to quote wider spreads. We test this hypothesis by re-estimating the markup decomposition model with hedge funds partitioned into those that do and do not engage in HFT. Consistent with our hypothesis, dealers widen spreads with information far more aggressively for hedge funds that engage in HFT than for other hedge funds.

Indirect evidence has long existed suggesting that adverse selection could be irrelevant to OTC client spreads (e.g., Reiss and Werner, 1996; Bernhardt et al., 2005; Edwards et al., 2007; Osler et al., 2011; Bjonnes et al., 2021). This finding was initially puzzling given the first-order importance of adverse selection in order-driven markets (e.g., Huang and Stoll, 1997). The theoretical possibility that rational dealers engage in information chasing instead of adverse selection was introduced by Naik et al. (1997) and first confirmed by Ramadorai (2008) for asset managers in spot FX. Pintér et al. (2021) have recently confirmed the relevance of information chasing with respect to asset managers in the UK government bond market.

We test the generality of information chasing in OTC markets by examining how dealers respond to the full spectrum of OTC clients: individuals, SMEs, MNCs, real-money funds, retail forex brokers, institutional brokers, CTFs, client banks, and hedge funds. Results reveal that dealers only chase information for brokers; for other informed clients they widen the spread with information consistent with adverse selection. To our knowledge this is the first rigorous evidence for the relevance of adverse selection in OTC markets. Our analysis further indicates that the contribution of adverse selection to client spreads is economically small and readily overshadowed by other sources of markup variation. This may explain its relevance was not previously detected.

The rest of this paper has five sections and a conclusion. Section I describes our data. Section II analyzes how client execution costs are influenced by information asymmetries in the FX core market. Sections III and IV provide evidence that dealers adopt three responses to client information. Section V examines the dealers' choice between adverse selection and information chasing. Section VI concludes.

1. Data

Our primary dataset comprises the complete front-office record of spot EUR-USD trades at a top-10 dealing bank over the 68 trading days from 2 January to 20 April 2012. A secondary dataset comprises second-by-second bid and ask quotes from the Reuters Matching platform.

Trading in FX, arguably the world's largest OTC market, takes place non-stop from around midnight GMT on Monday mornings to around 21 GMT on Fridays and at lower volumes on weekends. Our data include the standard fields – transaction price, quantity, date, and time (to the second) – plus counterparty ID, many client types assigned by the bank, trading platform, exact markup for client trades, trade initiator, trade direction, a flag for trades priced by human interbank dealers, and full details on prime-broker trades. All trades have complete information.

Our sample period includes a fairly important event, the resolution of private-sector contributions to a major Greek debt bailout. Following lengthy negotiations, Greek terms were announced on 21 February and endorsed by the European Commission on 1 March; the bond exchange took place in April (Bank for International Settlements, 2020). Though dramatic, these events did not generate unusual exchange-rate dynamics. As shown in Figure 1, exchange rates and the realized volatility of daily returns (20-day, centered) did not differ noticeably between our sample period and the preceding or subsequent 12 months.

1.1 Core-market Trades

We identify core-market participants as the top 50 *Euromoney* dealing banks plus counterparties who trade via prime broker agreements. Our bank acts as a liquidity provider for the vast majority (89%) of its core-market trades.

The FX core market was originally a quote-driven interdealer market similar to those today in less liquid assets such as US municipal bonds (Harris and Piwowar, 2006), precious metals, CDSs (Collin-Dufresne et al., 2020), or other OTC derivatives (Hau et al., 2021).⁴ Order-driven trading quickly became dominant in the early 1990s, after EBS and Reuters introduced matching engines. In the late 1990s a few large and sophisticated hedge funds were allowed to participate under prime brokerage contracts, a privilege subsequently extended to proprietary trading firms. During 2012 the traditional matching platforms enforced strict price-time priority, prohibited hidden orders, and set minimum tick sizes at 0.5 pip (EBS) and 1 pip (Reuters). (A pip or “price improvement point” is \$0.0001 in EUR-USD.) Newer matching engines, by contrast, generally offer last-look protection, hidden orders, and decimal-pip pricing.

As articulated by the Municipal Securities Rulemaking Board (MSRB, 2018), OTC client prices are set relative to the “prevailing market price,” which in FX necessarily comes from the primary market. Logically this will be taken from the order-driven platforms whose quotes are continuously displayed on dealer screens. Other prices, such as those for core-market trades executed in dark pools or after requests for quotes are opaque to all but the immediate counterparties and cannot be fed into pricing engines. We therefore focus on the 131,269 liquidity-providing order-driven trades in our core-market sample, worth in total €117 billion (Table 1)⁵. It is standard practice for FX dealers to split large trades into smaller ones to minimize execution costs (Bertsimas and Lo, 1999) or to mask their information content (Barclay and Warner, 1993) so we identify trade clusters having the same date and time (Huang and Stoll, 1997). This reduces the number of core-market trades to 88,376; the traditional order-driven platforms, EBS and Reuters Matching, jointly account for roughly half by trade number and three quarters by value.

1.2 Client Trades

The 2,388 clients in our sample made 257,241 trades worth €126.0 billion in aggregate.⁶ Our data include seven types of financial clients and three types of non-financial clients, a level of disaggregation well beyond that of previous studies (e.g., Osler et al., 2011; Menkhoff et al., 2016; Ranaldo and

⁴ Unlike OTC bond markets, however, dealers in liquid currency pairs never act as brokers.

⁵ Remaining core-market trades comprise 30,446 aggressive trades on order-driven platforms and 59,813 trades in dark pools.

⁶ We exclude the handful trades in which a client provided liquidity to our bank over an order-driven platform.

Somogyi, 2021). The 1,141 financial clients include 43 hedge funds, 19 commodity trading firms, 731 client banks, 56 retail brokers, 9 institutional brokers, 25 additional firms the bank originally categorized as “brokers,” and 257 real-money funds. Commodity trading firms (CTFs) engage in proprietary trading, like hedge funds, but focus on commodity futures contracts (Pirrong, 2015). Client banks are small or regional and generally focus on providing liquidity to their own clients; a few niche banks use HFT. Retail brokers offer narrow spreads over online platforms and trade with our bank to manage inventory. Institutional brokers provide advice and execution services for institutional clients. The “other brokers” identified by the bank include otherwise-uncategorized financial agents such as voice brokers, agency trading firms, factoring firms, a port operator, and a real-estate firm. Real-money funds are low-leverage asset managers such as mutual funds, pension funds, insurance firms, and endowments. The sample’s non-financial clients include are 169 MNCs, defined by the bank as having a specialized treasury unit or at least €100 million in annual sales; 990 SMEs, defined as firms that are not MNCs; and 88 individuals or “private clients.”

The mean markup across all client trades is 0.44 pips, which equates to 0.34 bps at the sample-mean exchange rate of \$1.31/€. As shown in Table 2, mean markups vary by four orders of magnitude across client types, from 0.001 pip for other brokers to 38.4 pips for private clients. Trade frequency also varies widely; the average real-money fund or non-financial client trades just once or twice per month; the average hedge fund or client bank trades once or twice daily; the average CTF or retail broker trades over 15 times daily.

Large FX dealing banks like the one we study do not eliminate most inventory positions immediately in the core market, as was the norm until the early 1990s (Bjønnes and Rime, 2005). Instead they allow client trades to offset each other and turn to the core market only for substantial inventory imbalances. These banks now price most trades electronically, using human dealers to choose quotes only for potentially toxic trades and other special circumstances. Every OTC quote is set as a markup relative to a composite same-side core-market price constructed by the bank’s pricing engine, which uses input from several order-driven platforms. The markup is tailored to the individual client, the behavior and preferences of which are familiar to their human salesperson and digitally scrutinized by a team of e-traders. The norm in FX is for clients to request two-way quotes.

Clients trade on more than 20 platforms that we partition into four groups: direct trading, single-bank platforms (SBPs), multi-bank request-for-quote platforms (MRFQs), and application programming interfaces (APIs). For a direct trade the client reaches out to a human salesperson by telephone, email, txt, etc., to request quotes. An SBP is a desktop interface over which clients can request quotes from, and

trade with, a specific bank. The simplest SBPs allow the client to place orders in advance for trades at a daily fix; the most complex provide click-and-trade functionality with streaming actionable quotes. On MRFQ platforms clients solicit simultaneous quotes from multiple dealing institutions (Hendershott and Madhavan, 2015). MRFQs enhance competition among dealers, which can bring narrower spreads, but they can also reveal information about a client’s intended trades. The installation and maintenance of MRFQ platforms requires considerable technical sophistication and usually takes at least two months. An API allows client algorithms to trade electronically at prices streamed by the dealer and requires by far the most technological sophistication from clients.

2. Client Spreads and Core-Market Information Asymmetries

This paper’s first contribution is to estimate the influence of core-market information asymmetries on OTC client spreads. We rely on the familiar Huang and Stoll (1997) spread decomposition model in which the price for trade t , P_t , equals the asset’s perceived true value, V_t , adjusted for an effective half-spread, $S/2$; for the gap between current and desired inventory, $I_t - I^*$; and for an i.i.d. mean-zero noise term, η_t : $P_t = V_t + \frac{S}{2}D_t - \beta \frac{S}{2}(I_t - I^*) + \eta_t$, $\beta > 0$. Trade direction, D_t , is +1 (-1) when the bank’s counterparty buys (sells) the base currency, EUR. V_t moves each period in response to the information conveyed by the most recent trade, $\alpha \frac{S}{2}D_{t-1}$, and to an i.i.d mean-zero public news shock, ε_t : $V_t = V_{t-1} + \alpha \frac{S}{2}D_{t-1} + \varepsilon_t$. Together, these assumptions imply the following expression for the current traded price:

$$P_t = V_{t-1} + \alpha \frac{S}{2}D_{t-1} - \beta \frac{S}{2}(I_t - I^*) + \frac{S}{2}D_t + \varepsilon_t + \eta_t. \quad (1)$$

The unknown V_{t-1} is eliminated by using the return, $\Delta P_t = P_t - P_{t-1}$, as dependent variable:

$$\Delta P_t = \frac{S}{2}\Delta D_t + \alpha \frac{S}{2}D_{t-1} - \beta \frac{S}{2}\Delta I_t + e_t, e_t \equiv \eta_t + \varepsilon_t - \varepsilon_{t-1}. \quad (2)$$

We allow S and α to vary between traditional interdealer matching engines, which require firm quotes (FQ), and the newer platforms, which should have lower adverse-selection risk because they offer “last-look” optionality (LL). With last look, liquidity providers can decline to trade during a fraction of a second after a counterparty accepts a quote, typically 80 –100 milliseconds (Phillips et al., 2020; Chaboud et al., 2023). Following Huang and Stoll (1997) We limit the baseline sample to trades of at least €250,000 given the model’s implicit assumption that trades have a consistent size. The full sample is considered in a robustness test.

As shown in Table 3, column 1, our baseline estimate of the effective core half-spread, $S/2$, is 0.139 pip for last-look platforms and 0.611 pip for firm-quote platforms; the difference is highly significant.⁷ The adverse-selection premium is smaller for last-look than firm-quote platforms, at 0.13 pip and 0.37 pip, respectively, as one would expect. The proportionate contributions of adverse selection to half-spreads – 90% for last-look platforms and 60% for traditional platforms – generally exceed the proportionate contributions of adverse-selection to US equity spreads, which average 46% and range across stocks from 21% to 64% (Huang and Stoll, 1997). The importance of information asymmetries for core FX spreads is logical given evidence in Bjønnes et al. (2021) that FX dealers that trade the most with clients have a highly-profitable information advantage relative to dealers with fewer clients. This can be explained by Evans and Lyons’ hypothesis that dealing banks infer currency fundamentals from broad patterns in client trades (2012) and by the complementary hypothesis that certain FX clients are relatively well-informed. Bjønnes et al. (2021) find that a dealer’s information rises with its hedge-fund trading and is unaffected by its trading with non-financial clients and real-money funds, findings that gain further empirical support in Section 4.

The adverse-selection premium from last-look platforms represents 22% of the mean client half-spread, 0.58 pip. This contribution is 5% or less for non-financial clients and real-money funds but reaches 74% for hedge funds, 84% for CTFs, and 90% brokers.⁸ The adverse-selection risk premium is larger on last-look platforms than firm-quote platforms, at 0.13 pip and 0.37 pip, respectively. This provides reassurance that the model is indeed capturing adverse selection in this setting (Van Ness et al., 2001; Collin-Dufresne and Fos, 2015). The difference between adverse-selection components on last-look and traditional platforms, 0.25 pip, represents the adverse selection costs avoided by the last-look option. If that option were unavailable core-market adverse selection would represent 64% of mean half-spreads across all clients.

Our finding that core-market adverse selection makes a substantial contribution to OTC client spreads is robust. The estimates of S and α change little when the sample is limited to trades during the relatively active European hours of 8 to 22 GMT (Table 3, column 2) or if price shading is ruled out (Table 3, column 3). When the sample is expanded to include trades below €250,000, the effective half-spread changes little for firm-quote platforms, because tiny trades are concentrated on last-look platforms

⁷ The interpretation of coefficient magnitudes as pips is appropriate given the median core-market trade size of €1 million on both platform types.

⁸ The mean client half-spread is the mean client markup, 0.44 pip, plus the 0.14-pip effective last-look core half-spread.

but roughly doubles for last-look platforms (Table 3, column 4). The share of adverse selection on last-look platforms in the mean client half-spread rises to 29%.

The bank generates a bespoke core-market price for each client trade, in addition to the bespoke markups we analyze in the next section. We finish this analysis by considering the extent to which these client-specific core-market spreads reflect those in our core-market sample that are actually paid by professional traders. Client-specific core half-spreads are approximately the same size as those on the last-look platforms. As shown in the last line of Table 3, the average core half-spread for client trades of at least €250,000 is 0.135 pip (calculated as the mean gap between the implied core-market price and the contemporaneous Reuters mid-quote). There is no significant difference between this and the corresponding last-look value for $S/2$, 0.139 pip. The bank does not appear to adjust these spreads for client information. We consider three groups of clients identified in Section IV: (1) hedge funds, CTAs, client banks; (2) retail, institutional, and other brokers; (3) real-money funds, non-financial firms, and private clients. The mean half-spreads for these groups are 0.17 pip, 0.16 pip, and 0.15 pip, respectively; none are statistically different from the others or from the last-look $S/2$, 0.139 pip.

Before moving on to analyze client markups we note an apparent shift over time in FX dealer behavior indicated by the statistical significance of the price shading coefficients statistically significant. Previous research has found little evidence for price shading in FX, though it is known to be important in equity markets (e.g., Huang and Stoll, 1997) and bond markets (e.g., Friewald and Nagler, 2024). Historically, FX dealers offloaded inventory positions quickly via outgoing interdealer trades (Bjønnes and Rime, 2005), and explained that choice in terms of the potential cost of conveying information about their position to the market. A shift among FX dealers towards price shading, if it occurred, could reflect the newer platforms' acceptance of hidden orders.

3. Asymmetric Information Between Dealers and Clients: An Empirical Model

This paper's second contribution is to provide evidence that OTC dealers adopt three different strategies for responding to client information. This section explains our empirical model:⁹

$$\begin{aligned} Markup = & \sum_j InfoCT'_j \pi_j + \beta TV + \delta_1 DTShr + \delta_2 DTShr^2 + \theta Skew + PCP \psi' \\ & \varphi_0 + (\kappa_1 + \lambda_1) HL + Time_i \kappa'_2 + (\kappa_3 + \lambda_2) CSprd + (\lambda_3 + \varphi_1) Sz + DT(\varphi_2 + \varphi_3 Sz) + \zeta, \end{aligned} \quad (3)$$

$$j = \{HF, CTF, ClBk, RetBrk, InstBrk, OthBrk, RM, MNC, SME, PC\}.$$

Each observation is a client trade and the dependent variable is the trade markup. The first line of Equation (3) captures client information, *Info*, and other client characteristics described below. The

⁹ A version of this model is also used in Bjønnes et al. (2026)

second line captures order-handling and inventory costs plus a new OTC-specific strategy, intertemporal smoothing.

3.1 Client Information

We use mean signed post-trade returns to capture client's c 's information following Anand and Subrahmanyam (2008): $Info_c \equiv \text{Mean}[D_{t,c}(m_{t+k} - m_t)]$, where $D_{t,c}$ is the direction of trade t , c is the client, m_t is the contemporaneous mid-quote, and k is either 1 minute or 30 minutes. Mid-quotes are the average of contemporaneous best bid and ask quotes from Reuters, which are better suited to calculating post-trade returns than the irregularly-spaced transaction prices in our main dataset. The impact of client c 's information on the markup, $Info_c \pi_{j(c)}$, can vary by client type, $j(c)$; $\pi_{j(c)}$ will be positive under adverse selection, negative under information chasing, and insignificant otherwise.

Table 4 reports descriptive statistics for $Info_c$. Mean $Info_c$ with 1-minute returns is positive and significant for hedge funds, retail brokers, and other brokers at 0.27 bp, 0.10 bp, and 0.34 bp, respectively. For 30-minute post-trade returns, standard errors are far larger and mean $Info_c$ is only significant for hedge funds.

The information in client trades could vary with trade size: Easley and O'Hara (1987) predict that informed agents will make large trades while Barclay and Warner (1993) predict that informed agents hide information by splitting large trades into smaller ones. We recalculate $Info_c$ after partitioning trades into four size categories (with amounts in millions): tiny $\in [\text{€}0, \text{€}0.5)$; small $\in [\text{€}0.5, \text{€}1.0)$; medium $\in [\text{€}1.0, \text{€}5.0)$; and large $\in [\text{€}5.0, \text{€}\infty]$. No consistent relation emerges between trade size and information. Note, however, that trades by CTFs and client banks in the largest three size categories carry information despite the overall insignificance of their post-trade returns.

3.2 Other Client Characteristics

The model comprehensively includes the other relevant client characteristics highlighted in the literature: trading activity, sophistication, and access to additional counterparties.

Client trading volume, TV , is measured as the log of the client's daily average trading in EUR, $TV_c \equiv \text{Ln}(\text{TrdVol}_c)$. We expect $\beta < 0$ given Bernhardt et al.'s (2005) hypothesis that rational dealers will quote narrower spreads to relatively active clients, seeking to attract their future business. Evidence for such activity discounts is provided by Bernhardt et al. (2005) for the LSE, Hau et al. (2021) for FX derivatives, and O'Hara et al. (2018) for U.S. corporate bonds.

The variables $DTShr$ and $Skew$ capture client "sophistication," a summary term for negotiating skills, market knowledge, and familiarity with different trading platforms (Green et al., 2007). The most

sophisticated clients hire professional FX traders; the least sophisticated traders may not even know that FX dealers trade as principals rather than as agents (Galanek, 2010). Sophistication increases a client's sensitivity to execution costs, limits the dealer's market power, and brings narrower equilibrium spreads (Ho and Stoll, 1981; Duffie et al., 2005; Green et al., 2007).

$DTShr_c$, the share of client c 's trades handled directly, serves as an inverse measure of sophistication following Liu et al. (2023), given the way clients learn about OTC markets over time. A small client needing to exchange currencies for the first time will seek guidance from its relationship bank, which will explain how to make a direct trade. Direct trades have relatively high costs because they require the time of a human dealer, unlike the other platforms. A client that trades infrequently may never learn about other platforms, however, because it has little incentive to invest the time and effort required to minimize execution costs. All of the private clients in our sample trade infrequently (at most once per week) and rely exclusively on direct trading. Similarly, 96% of the SMEs that trade infrequently rely exclusively on direct trading.

Larger clients are more likely to have in-house expertise, particularly clients with a treasury function, and are thus more likely to know about less costly platforms. Among MNCs – which are defined to be larger than SMEs or to have a treasury function – only 62% of the infrequent traders rely on direct trading. Among financial clients that trade infrequently an even smaller share rely on direct trading, 22%. The relatively heavy reliance on direct trading of small clients that trade infrequently is apparent in Figure 2A, which shows the share of different client types that trade directly for five different trade frequencies.

As clients trade more frequently, it becomes rational to invest in sophistication and they progressively shift from direct trading to SBPs and MRFQs. The shift out of direct trading is apparent in Figure 2A. The shifts into SBPs or MRFQs are apparent in 2B and 2C, respectively.

We allow the relationship between markups and $DTShr_c$ to be non-linear by including $DTShr_c^2$. We hypothesize that heavy reliance on direct trading signals a lack of sophistication, leading dealers to quote wider spreads on all of a client's trades. By contrast, modest reliance on direct trading could signal sophistication because the client takes advantage of the human dealers' skill at executing large trades.

Skew captures a second strategy through which dealers earn higher markups from unsophisticated clients: quoting a normal spread but raising both quotes relative to the mid-quote for a client likely to buy (and vice versa). An unsophisticated client accepts the trade given the normal spread, without recognizing that a trade in their normal direction will cost more than normal. We measure the predictability of a client's trade direction with $BuySell_c \equiv |#Buys_c - Sells_c| / (#Buys_c + #Sales_c) \in [0, 1]$; 0

indicates a clients that is equally likely to buy or sell; 1 indicates a client that always buys or always sells. *Skew* interacts *DTShr* and *BuySell*: $Skew \equiv DTShr \times BuySell$. We expect $\theta > 0$.

Access to Potential Counterparties. The variable *PCP* captures cross-client variation in access to potential counterparties. A client with multiple dealing relationships simply declines to trade when quoted a wide spread and seeks a better spread with another dealer. Access to multiple dealing institutions thus brings a high price elasticity of demand (Ho and Stoll, 1981), low search costs (Duffie et al., 2005), and strong negotiating leverage (Green et al., 2007) and thus lower spreads (Hendershott et al., 2020; Hau et al., 2021). Nonetheless, it is common for OTC clients to have few dealing relationships. Over half of clients trading FX derivatives (Hau et al., 2021) and roughly 30% of insurance firms trading US corporate bonds (Hendershott et al., 2020) trade with just one dealer in a given year. Establishing multiple relationships may be discouraged by its high cost. A client must qualify for a line of credit at a new bank, which requires posting a large, permanent deposit or demonstrating credit-worthiness. The latter requires gathering and presenting a mass of detailed information, which can be challenging for individuals or small firms where employees wear multiple hats and procedures are not well-established.

PCP comprises indicator variables for clients that are likely to have either a multitude of dealing relationships or very few. The first corresponds to clients that trade via prime broker, *PB*, which have access to the entire professional trading community as alternative counterparties. We expect $\psi_{PB} < 0$.

Two *PCP* indicators correspond to clients that are likely to have few dealing relationships by choice: private clients and SMEs that do not trade over MRFQs; we expect $\psi_{Priv}, \psi_{SME-NM} > 0$. MRFQs, which are intended for clients with multiple dealing relationships, are used by none of the private clients in our sample and just 1% of the SMEs. By contrast, MRFQs are used by 27% of MNCs and 50% of financial clients. Consistent with the view that costs discourage smaller firms from establishing many dealing relationships, SMEs switch to SBPs but not MRFQs as trading frequency rises; MNCs, by contrast, switch directly to MRFQs (Figures 2B and 2C).¹⁰

PCP also includes an indicator for a set of clients that are likely constrained from having multiple dealing relationships: banks from countries with limited financial development (LFD). Major dealing banks are generally reluctant to trade with such banks, based on concerns about settlement risks and financial instability in their base countries. These concerns can be overcome, however, when the dealing

¹⁰ According to market participants, the MNCs' preference for MRFQs often reflects internal control mandates that require bids from multiple service providers. MRFQs can generate narrow spreads by intensifying competition among liquidity providers, as noted in Section I. However, a representative of a major MRFQ reports that these platforms are generally not cost-effective for clients that trade less than €100 million annually due to their high set-up costs. The process generally requires at least two months and specialized expertise.

bank serves as correspondent and therefore holds a large deposit from the LFD bank; the deposit largely eliminates FX settlement risk and generates profits for the dealing bank. We define LFD banks as those based in countries with 2012 IMF Financial Development Index in the lowest tercile. Just 14% of these banks rely on MRFQs, well below the 50% share for other client banks. Indeed, MRFQ reliance is lower for LFD banks than other banks at every trade frequency, as shown in Figure 2D. (Results are similar if LFD is defined as the lowest quartile.) We expect $\psi_{LFD} > 0$.

3.3 Order-handling Costs, Inventory Risk, and Smoothing

The variables in the second line of Equation (3) capture order-handling costs, inventory costs, and intertemporal smoothing. The order-handling cost of a fully-electronic trade is captured by $\varphi_0 + \varphi_1 Sz$, where $Sz = \ln(TradeSz)$ for $TradeSz$ in EUR. Because FX trading has fixed costs we expect $\varphi_0 > 0$ and $\varphi_1 < 0$. The additional fixed cost of a direct trade is captured by $(\varphi_2 + \varphi_3 Sz)DT$, where DT indicates a direct trade; we expect $\varphi_2 > 0$ and $\varphi_3 < 0$.

We model inventory risk as $\lambda_1 Sz + \lambda_2 HL + \lambda_3 CSprd$. A larger trade and higher volatility, HL , bring greater prospective inventory risk so we expect $\lambda_1, \lambda_2 > 0$ (Ho and Stoll, 1981). HL is the high-low range of the previous trading hour. $CSprd$, the contemporaneous core spread from Reuters, captures market-wide illiquidity (Fleming, 2003) and we expect $\lambda_3 > 0$. Inventory costs can theoretically include carry as well as risk, but carry is negligible in spot FX because dealers infrequently hold positions overnight (Bjønnes and Rime, 2005).

We hypothesize that dealers try to maintain a steady spread for each client. They recognize that a client's search for best execution is not random, undirected, or nonstrategic. Instead, clients choose dealers in part for “consistency of spreads across ... market conditions.” and “resilience during volatile periods or stress scenarios,” two of *Euromoney*'s annual dealer award categories.¹¹ According to the dealer survey of Cheung and Chinn (1998), “the ability to consistently offer quotes with ... conventional spreads in a hectic market is regarded as an essential characteristic of a market leader” (p. 444). Smoothing also spares clients the confusion and annoyance of frequent price changes (Dholakia, 2016, p. 1) while reducing price noise that can mask a dealer's competitiveness.

We model smoothing as follows: $\kappa_1 HL + Time \kappa_2' + \kappa_3 CSprd$. Smoothing implies an inverse relation between markups and the core spread, so we expect $\kappa_3 < 0$. $CSprd$ is unlikely to respond to subtle variations in market conditions, however, given Reuters' 1-pip minimum tick. The model therefore

¹¹ *Euromoney*, [Foreign Exchange Awards, 2025](#), p. 23.

includes two additional influences on core spreads: volatility and strong intraday seasonals in trading volume. *Time* comprises three time-of-day indicators: European hours, 8–16 GMT, when volume is highest; overnight hours, 22–02 GMT, when volume is lowest; and New York trading, 17–21 GMT (Asian hours are excluded). We expect $\kappa_1 < 0$, $\kappa_{2-Europe} > 0$, $\kappa_{2-Overnight} < 0$.

4. Asymmetric Information Between Dealers and Clients: Results

We estimate the model using GMM with Newey-West standard errors clustered by date.¹² The results, presented in this section, provide extensive support. All coefficients are consistent with their motivating theories and the adjusted R^2 is encouragingly high at 0.58. Though one cannot directly compare explanatory power across datasets and models, it is nonetheless notable that explanatory power in previous studies of OTC client execution costs has been far lower, with a maximum of 0.33 (Hau et al., 2021).

4.1 Dealer Responses to Client Information

The estimated coefficients indicate that dealers adopt all three of the potentially rational responses to client information. The *Info* coefficient is negative and significant for all three types of brokers, which jointly account for 58% of trades and 30% of trade value. This confirms the importance of information chasing documented by Ramadorai (2008) and Pintér et al. (2022). A one-standard-deviation decrease in *Info* raises the markup for retail, institutional, and other brokers by 0.88 pip, 0.08 pip, and 0.52 pip, respectively.

The *Info* coefficient is positive and significant for hedge funds, CTFs, and client banks, which jointly account for 33% of clients and 67% of client trade value. This finding provides the first rigorous evidence for adverse selection between OTC dealers and their clients. These three client types. A one-standard-deviation increase in *Info* is estimated to raise markups for hedge funds, CTFs, and client banks by 0.59 pip, 0.87 pip, and 0.37 pip, respectively. The impact on client-bank markups may be relatively small because many are likely *uninformed*, specifically those focused on client service rather than proprietary trading, which trade with our bank almost exclusively for inventory management. Those banks' own clients tend to be uninformed because the best-informed clients prefer to trade with the dominant banks (Bjønnes et al., 2021).

Our analysis is able to identify adverse selection in part because the data include a comprehensive set of clients and the model includes a comprehensive set of FX markup determinants. As shown in Bjønnes

¹² An autocovariance lag length of 23 trades, which corresponds roughly to one hour, is chosen by the Newey and West (1994) automatic process. Results are qualitatively unchanged if lag length is chosen using other standard criteria including AIC.

et al. (2026), the originally puzzling inverse relation between OTC spreads and trade size (e.g., Bernhardt et al., 2005) can be traced to differences in client sophistication and access to potential counterparties, characteristics that cannot be identified without client characteristics. Clients such as hedge funds are better informed than average but pay narrower spreads because they are relatively skilled negotiators, they trade over low-cost platforms, and they maintain dealing relationships with multiple banks. Our new evidence for adverse selection may also reflect the emergence of HFT in FX, as discussed in Section V.

The *Info* coefficients for real-money funds, MNCs, SMEs, and private clients are statistically insignificant, from which we infer that their spreads are not influenced by client information. This group includes 1,504 clients, 62% of the total, though they account for just 5% of trade value.

The option to ignore client information has not previously been documented empirically, to our knowledge. Based on post-trade returns reported in Table 4 we infer that non-financial clients and real-money funds are simply uninformed, consistent with evidence that FX dealer information does not rise with their trading with these client types (Bjønnes et al., 2021). A lack of information could be logical among non-financial firms because they primarily use FX as a medium of exchange and internal governance mandates treat it as a cost within the finance function rather than a source of profits. Indeed, the bylaws of many non-financial firms discourage speculative trading because it brings rogue-trader risk, which is costly to control (Osler, 2006).

The lack of information among real-money funds is not consistent with standard assumptions in economic theory though it is consistent, however, with the view of industry experts. They report that real-money funds tend to treat currency trading as an unavoidable administrative cost rather than a source of returns (Taylor and Farstrup, 2006; Galanek, 2010). The low frequency with which these clients trade (Table 2) and the high cost of currency analysts may limit their incentive to invest in information.

Theory does not dictate a specific time horizon for calculating the post-trade returns of *Info_c*, so we experiment with an alternative, 30 minutes, a period that is quite long given the speed and market penetration of HFT during our sample period (Chaboud et al., 2014). The information coefficients for hedge funds, CTFs, retail brokers and institutional brokers retain their original signs but are substantially smaller, which is logical because the *Info* standard deviations are many times higher (Table 4). The *Info* coefficients for CTFs and client banks become insignificant, suggesting that their trades primarily carry information about high-frequency returns. The estimated impact on markups of a one-standard-deviation decline in retail and institutional brokers' *Info* increases to 2.5 pips and 2.0 pips, respectively. The corresponding figure for hedge-fund *Info* increases to 4.2 pips. With the 30-minute time horizon *Info* remains insignificant for real-money funds, private clients, MNCs, and SMEs.

In summary, the evidence presented in this section confirms and expands Ramadorai's (2008) evidence for vast differences between order-driven and quote-driven markets with respect to asymmetric information across trade counterparties. In order-driven markets, information asymmetries generate adverse selection that inevitably raises the spread and does so by an amount that is common across traders (Glosten and Milgrom, 1985). In OTC markets, the sign of a dealer's rational response to client information depends on the client's type and the magnitude varies with the client's information.

4.2 Other Model Implications

The remaining coefficients confirm that relatively active and sophisticated clients benefit from narrower markups. The response of markups to $DTShr$ is nonlinear, as expected: modest reliance on direct trading brings slightly narrower markups and heavy reliance brings sharply higher markups. Consistent with quote skewing, markups are positively related to the predictability of a client's trade direction and the effect diminishes with sophistication. Markups are larger (smaller) for clients with the fewest (most) alternative trading relationships.¹³

The results also confirm that dealers pass on the extra order-handling costs of direct trades: the markup rises by 4.1 pips if an average-sized trade is made directly rather than electronically. Markups on all trades are inversely related to trade size, consistent with fixed costs. Intertemporal smoothing is evident from the inverse relation between markups and volatility – which does not conform to standard inventory theory (Ho and Stoll, 1981) – and from the positive relation between markups and trading volume.

4.3 Extensions and Robustness Tests

These results prove robust to six modifications of the baseline model, as shown in Table 6.

Price Shading. We investigate the possibility that dealers manage inventory by shading prices to clients. Inventory is approximated by the dealers' net EUR accumulation over the most recent 100 EUR-USD trades, which corresponds to about 20 minutes. This is consistent with prior research showing that active FX dealers return inventories to the desired level within five minutes of a shock (Bjønnes and Rime, 2005). This measure is not ideal because the dealing floor's inventory of EUR will also be influenced by trades against other currencies; however, those trades are not included in our data. Consistent with earlier research discussed in Section II, the inventory measure proves insignificant and the model's qualitative conclusions are unchanged.¹⁴

¹³ We review non-*Info* coefficient estimates briefly because they are not central to our analysis; a thorough discussion with economic impacts is available in Bjønnes et al. (2026).

¹⁴ We also examined the robustness of our findings for CTFs by excluding their five largest trades, which had an extraordinarily large mean 1-minute return, 4.0 pips. This brought no meaningful changes in the results.

Trading hours. We limit the sample to European and New York business hours, 8 – 21 GMT, when EUR-USD trading is most active. This brings no qualitative change in dealer responses to client information though the *Info* responses strengthen for hedge funds and brokers.

Measuring sophistication. We add an indicator variable for hedge funds and non-LFD client banks, which are known to be expert FX traders. The model's qualitative conclusions remain unchanged though the previously-significant dealer responses to information increase in magnitude. Most of these changes are modest (less than 50%), though the response to institutional-broker information increases by a factor of five. A one-standard-deviation decline in an institutional broker's information widens the markup by 4.7 pips.

Other Platform Shares. To investigate the possibility that dealers learn about client sophistication from their tendency to rely on SBPs or MRFQs we add the associated client platform shares, *SBPShr* and *MRFQShr*. Coefficients on both platform shares are negative and significant, which suggests that their use communicates sophistication, but their economic effects are tiny. A rise *SBPShr* or *MRFQShr* by 50 percentage points reduces the markup on all trades by less than 0.01 pip.

Trade Direction. We add an indicator for client sales of EUR. The indicator is negative and statistically significant but economically trivial; spreads are estimated to be 0.07 pip for client EUR sales than client EUR purchases. This is unsurprising because the FX market is generally symmetrical and short sales are not defined because the sale of one currency is the purchase of another.

Market volatility. The model's qualitative conclusions about dealer responses to client information are unchanged if the realized volatility of one-minute returns over the previous hour replaces the high-low range.

5. Adverse Selection versus Information Chasing

This section takes a closer look at the dealers' choice between adverse selection and information chasing. The choice between narrowing and widening the markup for better-informed clients is determined by the balance between adverse-selection costs and information-chasing benefits. Trades with informed clients incur adverse selection losses, on average (Glosten and Milgrom, 1985), but dealers might offset those through quick adjustments in their quotes to others (Pintér et al., 2022) or parallel proprietary trades in the core market (Naik et al., 1999). Information chasing was standard operating procedure in FX during the 1980s and 1990s. Indeed, the physical layout of an FX trading floor was designed to support this strategy according to William Clyde, an academic researcher who managed an FX trading desk early in his career (Clyde, 1996).

Our dataset's 253 trades with negative markups appear to capture information chasing in action, as shown in Table 7. These trades tend to be larger than other trades, with mean size €2.4, and better informed. One-minute post-trade returns tend to be positive but their mean is insignificant at 0.22 pip. The mean two-minute post-trade returns remains positive and is significant at 0.47 pip (0.36 bp). Two minutes is a long time in liquid 21st-century markets, long enough for dealers to adjust quotes to other clients or to profit from a parallel core-market position. The remaining trades in our sample provide a striking contrast, with post-trade returns that are two orders of magnitude smaller and statistically insignificant.

We hypothesize that trade latency is a critical determinant of the dealers' choice between adverse selection and information chasing. The attitude of FX dealers towards informed hedge funds shifted markedly when high-frequency trading arrived in the 2000s (Markets Committee, 2011). This is unsurprising because dealers were vulnerable to at least three standard HFT strategies. The first, latency arbitrage, intensifies the risk that a dealer's stale quotes will be picked off (Menkveld and Zoican, 2017). The second, anticipating large trades as they move through the market (Hirshy, 2021), compromises dealer profitability by accelerating price discovery and shortening the already-brief period during which dealers can exploit client information. Finally, HFT profits from riskless arbitrage opportunities, which were previously available only to dealers (Budish et al., 2015; Chaboud et al., 2014).

We suggest that HFT shifted the dealers' cost-benefit balance on trades with hedge funds, leading them to shift from information chasing to adverse selection. We test this hypothesis by re-estimating the model with hedge funds partitioned into those that do and do not engage in HFT according to information in their websites.¹⁵ The latency hypothesis predicts a larger *Info* coefficient for HFT hedge funds than other hedge funds.

The results strongly support the relevance of latency to the dealers' choice between adverse selection and information chasing (Table 8). The *Info* coefficient for HFT funds is large, at 1.35 pip, and over three times the *Info* coefficient for non-HFT hedge funds, 0.38. For a one-standard-deviation rise in hedge-fund information, the markup widens by 1.3 pip for an HFT hedge fund and just 0.4 pip for a non-HFT hedge fund.

This finding may clarify why dealers widen the spread for better-informed hedge funds but narrow the spread for better-informed institutional brokers, which often trade for hedge funds. The information in the trades of institutional brokers comes from their clients with an inescapable lag due to client-broker

¹⁵ It proved infeasible to identify HFT trading among CTFs or client banks.

communication. By implication such trades involve substantial latency, which reduces those traders' adverse-selection costs.

6. Concluding Remarks

This paper provides evidence for additional striking differences between OTC and order-driven markets in the influence of asymmetric information using highly detailed data from the world's largest OTC market. Previous research shows that the response of OTC spreads to client information will vary by client (Naik et al., 1999; Ramadorai, 2008); in order-driven markets, by contrast, it is common to all agents. Previous research also finds that OTC spreads are narrower for better-informed clients,

We first propose and documents that OTC spreads are influenced by not one but two distinct information asymmetries: (i) the asymmetry between dealer and client and (ii) the asymmetry among professional traders in the core (or interdealer) market. This conclusion reflects the two-tier structure of OTC markets and the standard practice of setting OTC client prices as a markup over the "prevailing market price" (MSRB, 2018). Using the Huang and Stoll (1997) spread decomposition model we estimate that adverse selection in the FX core market contributes 25% of the average client spread, a share that can exceed 74% for the clients most sensitive to execution costs.

The paper also provides evidence that OTC dealers adopt not one but three different strategies for responding to client information. This provides a further striking contrast with order-driven markets, where adverse selection dominates (Glosten and Milgrom, 1985; Bernhardt et al., 2005). We estimate an OTC-tailored model of the dealer's choice variable, the markup. The model comprehensively captures the spread determinants highlighted in the literature that are relevant in FX, including four client characteristics, order-handling costs, inventory costs, and intertemporal smoothing.

The model indicates that OTC dealers quote wider spreads to better-informed hedge funds, commodity trading firms, and client banks, consistent with adverse selection (Glosten and Milgrom, 1985). To our knowledge this is the first clear evidence for adverse selection with respect to OTC clients. The model further indicates that dealers quote narrower spreads for better-informed brokers, consistent with information chasing (Naik et al., 1999; Pintér et al., 2022). Finally, the model indicates that dealers make no information-related adjustment to spreads for clients they consider uninformed, specifically low-leverage asset managers ("real-money funds"), private clients, and non-financial firms.

We hypothesize that trade latency is a critical determinant of the dealers' choice between information chasing and adverse selection. We partition hedge funds into those that do and do not engage in HFT and

re-estimate the model. Consistent with the latency hypothesis, the model indicates that dealers widen spreads much more aggressively for informed HFT hedge funds than for other informed hedge funds.

Future research could usefully analyze the extent to which interdealer information asymmetries influence client spreads in other OTC markets. It would also be valuable to examine the relevance of adverse selection and information chasing in less liquid OTC markets, such as those for US corporate or municipal bonds.

REFERENCES

- Anand, Amber, and Avandhar Subrahmanyam, 2008. Information and the intermediary: Are market intermediaries informed traders in electronic markets? *Journal of Financial and Quantitative Analysis* 43(1):1-28.
- Bank for International Settlements, 2020. The 2012 private sector involvement (PSI) in Greece. *Discussion Paper Series* No. 11.
- Bank for International Settlements, 2025. *Triennial Central Bank Survey, Foreign Exchange and Derivatives Market Activity in 2025*. B.I.S., Basel.
- Barclay, Michael J., and Jerold B. Warner, 1993. Stealth trading and volatility: Which trades move prices? *Journal of Financial Economics* 34(3): 281-305.
- Bernhardt, D., Vladimir Dvoracek, Eric Hughson, and Ingrid Werner, 2005. Why do larger orders receive discounts on the London Stock Exchange? *The Review of Financial Studies* 18: 1343-1368.
- Bjønnes, Geir Høidal, and Dagfinn Rime, 2005. Dealer behavior and trading systems in foreign exchange markets. *Journal of Financial Economics* 75: 571–605.
- Bjønnes, Geir Høidal, Carol Osler, and Dagfinn Rime, 2021. Price discovery in two-tier markets. *International Journal of Finance and Economics* 26(2): 3109-3133.
- Bjønnes, Geir Høidal, Neophytos Kathiziotis, and Carol Osler, 2026. Who pays the most to trade? Cross-client variation in OTC execution costs. Working paper.
- Budish, Eric, Peter Cramton, and John Shim, 2015. The high-frequency trading arms race: Frequent batch auctions as a market design response. *Quarterly Journal of Economics* 130(4): 1547-1621.
- Chaboud, Alain P., Benjamin Chiquoine, Erik Hjalmarsson, and Clara Vega, 2014. Rise of the machines: Algorithmic trading in the foreign exchange market. *The Journal of Finance* 69(5): 2045-2084.
- Cheung, Y.-W., Chinn, M.D., 2001. Currency traders and exchange rate dynamics: A survey of the U.S. market. *Journal of International Money and Finance* 20: 439-471.
- Clyde, William, 1996. Intraday foreign exchange market anomalies and microstructure. [Working Paper](#), Quinnipiac University.
- Collin-Dufresne, Pierre, and Vyacheslav Fos, 2015. Do prices reveal the presence of informed trading? *Journal of Finance* 70(4): 1555-1582.
- Collin-Dufresne, Pierre, Benjamin Junge, and Anders B. Trolle, 2020. Market structure and transaction costs of Index CDSs. *Journal of Finance* 75(5): 2719-2763.
- Dhokakia, Utpal M., 2015. The risks of changing your prices too often. *Harvard Business Review* (6 July). <https://hbr.org/2015/07/the-risks-of-changing-your-prices-too-often>.
- Duffie, Darrell, Nicolae Gârleanu, and Lasse Heje Pedersen, 2005. Over-the-counter markets. *Econometrica* 73: 1815-1847.
- Easley, David, and Maureen O'Hara, 1987. Price, trade size, and information in securities markets. *Journal of Financial Economics* 19: 69-90.
- Edwards, Amy, Lawrence Harris, and Michael Piwowar, 2007. Corporate bond market transaction costs - and transparency. *The Journal of Finance* 62(3): 1421-1451.
- Evans, Martin, and Richard K. Lyons, 2012. Exchange rate fundamentals and order flow. *Quarterly Journal of Finance* 2(4): 1-63.
- Feldhütter, Peter, and Thomas Kjaer Poulsen, 2018. What determines bid-ask spreads in over-the-counter markets? <https://ssrn.com/abstract=3286557>.
- Fleming, Michael, 2003. Measuring Treasury Market Liquidity. Federal Reserve Bank of New York *Economic Policy Review* 9(3): 83-108
- Friewald, Nils, and Florian Nagler 2024. Dealer inventory and the cross-section of corporate bond returns. *Economics Letters* 239. <https://doi.org/10.1016/j.econlet.2024.111710>.

- Galanek, John 2010. What Can Usain Bolt Teach Us About FX Transaction Cost Analysis? FX Transparency, Inc. <https://fxtransparency.com/2010/04/14/what-can-usain-bolt-teach-us-about-fx-transaction-cost-analysis/>
- Glosten, Lawrence, and Paul Milgrom, 1985. Bid, ask, and transaction prices in a specialist market with heterogeneously informed agents. *Journal of Financial Economics* 14: 71–100.
- Glosten, Lawrence, and Lawrence Harris, 1988. Estimating the components of the bid/ask spread. *Journal of Financial Economics* 21: 123–42.
- Goldstein, Michael, Edith S. Hotchkiss, and Erik Sirri, 2007. Transparency and liquidity: A controlled experiment on corporate bonds. *Review of Financial Studies* 20(2): 235–273.
- Green, Richard C., Burton Hollifield, and Norman Schurhoff, 2007. Financial intermediation and the costs of trading in an opaque market. *Review of Financial Studies* 20: 275–314.
- Harris, Lawrence, and Michael Piwowar, 2006. Secondary trading costs in the municipal bond market. *The Journal of Finance* 61(3): 1361–1397.
- Hau, Harald, Peter Hoffmann, Sam Langfield, and Yannick Timmer, 2021. Discriminatory pricing of over-the-counter derivatives. *Management Science* 67(11): 6660–6677.
- Heimer, Rawley, and David Simon, 2014. Facebook finance: How social interaction propagates active investing. Presented at American Finance Association Annual Meetings, June.
- Hendershott, Terrance, and Ananth Madhavan, 2015. Click or Call? Auction versus Search in the Over-the-Counter Market. *Journal of Finance* 70(1): 419–447.
- Hendershott, Terrance, Dan Li, Dmitry Livdan, and Norman Schurhoff, 2020. Relationship trading in OTC markets. *The Journal of Finance* 75(2): 683–734.
- Hirshy, Nicholas, 2021. Do high-frequency traders anticipate buying and selling pressure? *Management Science* 67(6): 3321–3345.
- Ho, Thomas S.Y. and Hans R. Stoll, 1981. Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics* 9: 47–73.
- Huang, Roger D. and Hans R. Stoll, 1997. The components of the bid-ask spread: A general approach. *Review of Financial Studies* 10: 995–1034.
- Ito, Takatoshi, and Yuko Hashimoto, 2006. Intraday seasonality in activities of the foreign exchange markets: Evidence from the electronic broking system. *Journal of Japanese International Economics* 20: 637–664.
- Liu, Jiageng, Igor Makarov, Antoinette Schoar, 2023. Anatomy of a run: The Terra Luna crash. *NBER Working Paper* 31,160 (April 2023).
- Markets Committee, 2011. *High-frequency trading in the foreign exchange market*. Basel: Bank for International Settlements.
- Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimp, 2016. Information flows in foreign exchange markets: Dissecting customer currency trades. *The Journal of Finance* 71(2): 601–634.
- MSRB (Municipal Securities Rulemaking Board), 2018. *How are municipal bonds quoted and priced?* <https://www.msrb.org/sites/default/files/how-are-municipal-bonds-quoted-and-priced.pdf>
- Naik, Narayan Y., Anthony Neuberger, and S. Viswanathan, 1999. Trade disclosure regulation in markets with negotiated trades. *Review of Financial Studies* 12: 873–900.
- O’Hara, Maureen, Yihui Wang, and Xing Alex Zhou, 2018. The execution quality of corporate bonds. *Journal of Financial Economics* 130: 308–326.
- O’Hara, Maureen, and Xing Alex Zhou, 2021. The electronic evolution of corporate bond dealers. *Journal of Financial Economics* 140: 368–390.
- Osler, Carol, 2006. Macro lessons from microstructure. *International Journal of Finance and Economics* 11: 55–80.

- Osler, Carol, Alexander Mende, and Lukas Menkhoff, 2011. Price discovery in currency markets. *Journal of International Money and Finance* 30: 1696–1718.
- Phillips, Andrew, Andrew Stewart, and Sam Adams, 2020. TCA and fair execution: The metrics that the FX industry must use. LMAX Exchange *White Paper*. <https://www.lmax.com/blog/business-and-technology/2017/05/22/tca-fair-execution-metrics-fx-industry-must-use/>
- Pintér, Gábor, Chaojun Wang, and Junyuan Zou, 2022. Information chasing versus adverse selection. Bank of England *Staff Working Paper* 971.
- Pirrong, Craig, 2015. The economics of commodity trading firms. *Trafigura White Paper*. <https://www.bauer.uh.edu/centers/uhgemi/casedocs/The-Economics-of-Commodity-Trading-Firms-2.pdf>.
- Ramadorai, Tarun, 2008. What determines transactions costs in foreign exchange markets? *International Journal of Finance and Economics* 13: 14-25.
- Ranaldo, Angelo, and Fabricius Somogyi, 2021. Asymmetric information risk in FX markets. *Journal of Financial Economics* 140: 391-311.
- Reiss, Peter C., and Ingrid M. Werner, 1996. Transaction costs in dealer markets: Evidence from the London Stock Exchange. In Andrew W. Lo, Ed., *The Industrial Organization and Regulation of the Securities Industry*. University of Chicago Press.
- Taylor, Andrew, and Adam Farstrup, 2006. *Active Currency Management: Arguments, Considerations, and Performance for Institutional Investors* (CRA RogersCasey International Equity Research, Darien Connecticut).
- Van Ness, Bonnie, Robert Van Ness, and Richard Warr, 2001. How well do adverse selection components measure adverse selection? *Financial Management* 30(3): 77-98.

Table 1: Descriptive Statistics for Core-market Trades

Table shows summary statistics for liquidity-providing order-driven core-market trades. A cluster of trades have the same date and time. Data include all EUR-USD trades by a top-10 EUR-USD dealing bank during the 68 trading days from 2 January through 20 April 2012.

	All	All	Clustered Last look	Firm- quote
All trades				
Number trades	131,269	117,805	56,925	60,899
Trade size (€ mn)				
Mean	0.891	0.993	0.465	1.485
Median	1.000	1.000	0.200	1.000
St. dev.	0.854	1.127	0.618	1.263
All trades $\geq$ €250K				
Number trades	101,820	88,376	27,936	60,440
Trade size (€ mn)				
Mean	1.134	1.307	0.898	1.496
Median	1.000	1.000	1.000	1.000
St. dev.	0.823	1.138	0.637	1.263

Table 2: Client Trades, Descriptive Statistics

Table shows summary statistics for all market-making EUR-USD trades by a top-10 EUR-USD dealing bank with non-dealing-bank counterparties during the 68 trading days from 2 January through 20 April 2012. Markups measured in pips, where one pip is \$0.0001/€. For client c , $BuySell_c \equiv |\#Buys_c - Sells_c| / (\#Buys_c + \#Sales_c)$, $BuySell_c \in [0, 1]$.

	All	Hedge Fund	CTF	Client bank	Retail Broker	Institution Broker	Other Broker	Real-money	MNC	SME	Private Clients
A. Basic properties											
Number clients	2,388	43	19	731	56	9	25	257	169	990	88
Number trades	257,241	6,624	19,792	73,730	79,397	12,703	17,660	1,185	1,261	3,660	113
Mn. trades/client	108	154	1,042	101	1418	1411	706	5	7	4	1
Trade size (€ mn)											
Mean	0.49	1.20	0.58	0.88	0.23	0.81	0.29	2.08	2.53	0.23	0.17
St. dev.	1.91	1.81	0.49	3.15	0.62	0.85	0.70	6.60	5.48	0.84	0.30
B. Markups (pips)											
Mean	0.44	0.03	0.01	0.26	0.15 (e-1)	0.02 (e-1)	0.01 (e-1)	2.59	3.09	22.26	38.35
St. Dev.	3.96	0.48	0.37	2.51	0.69	0.04	0.04	11.59	10.80	18.39	17.19
C. Daily Trading (€ mn)											
Mean	1647.46	102.46	148.01	832.16	236.16	132.45	66.41	31.69	40.93	10.61	0.24
St. Dev.	400.39	60.06	77.40	258.46	85.99	64.53	24.54	29.66	23.93	6.79	0.72
D. Platform shares (% of trades)											
Direct	1.87	1.97	0.03	2.04	0.02	0.02	0.46	13.80	24.57	67.59	100.00
SBP	29.61	1.17	0.03	68.57	17.37	77.99	0.61	32.83	17.68	30.52	0.00
MRFQ	27.15	47.00	57.03	26.11	27.79	3.61	69.09	53.18	57.75	0.94	0.00
API	41.37	49.86	42.92	3.29	54.81	18.37	29.84	0.19	0.00	0.96	0.00
E. BuySell											
Mean	0.11	0.09	0.03	0.25	0.04	0.08	0.05	0.29	0.29	0.70	0.23
St. Dev.	0.20	0.09	0.05	0.27	0.04	0.06	0.05	0.34	0.33	0.41	0.39

Table 3: Effective Core Spreads

Table reports coefficients from estimating Equation (2), the decomposition model for core spreads:

$$\Delta P_t = \frac{S}{2}(D_t - D_{t-1}) + \alpha \frac{S}{2}D_{t-1} + \beta \frac{S}{2}\Delta I_t + e_t, \quad e_t \equiv \eta_t + \varepsilon_t - \varepsilon_{t-1}.$$

P_t : traded price in pips from an order-driven core-market platform; D_t : signed trade direction, +1 (-1) when the counterparty buys (sells); I_t : dealer inventory; S : effective spread in pips; α : adverse selection contribution as share of effective spread; β : price shading as share of effective spread. Data comprise complete spot EUR-USD trading record by a top-10 dealing bank during the first 68 trading days of 2012. Trades at the same price, date, and time are clustered as one transaction. Estimation is by GMM, robust standard errors clustered by time and trading platform. Significance at 10%, 5%, 1% and 0.1% levels indicated by *, **, *** and †, respectively.

Table also reports mean core half-spreads on client-facing trades, calculated as $10000D_t(\text{client core price} - \text{midquote})$, using Reuters midquotes. Client core price = Client trade price - $D_t \text{Markup}$.

<i>Platform Type</i>	All trades ≥ €250K	Active hours only	No price shading	All sizes
Firm-quote order-driven				
Effective half-spread (e-4), $S/2$	0.611†	0.553†	0.575†	0.580†
Adverse selection, α (share of $S/2$)	0.608†	0.616†	0.576†	0.636†
Price shading, β (share of $S/2$)	0.040†	0.035†		0.039†
Last-look order-driven				
Effective half-spread (e-4), $S/2$	0.139†	0.137†	0.133†	0.229†
Adverse selection, α (share of $S/2$)	0.900†	0.874†	0.892†	0.738†
Price shading, β (share of $S/2$)	0.054*	0.050		0.039***
Adj. R^2	0.036	0.033	0.035	0.011
<i>N. Obs.</i>	88,375	72,071	88,375	117,805
Mean core half-spread, client trades	0.135	0.147	0.135	0.161

Table 4: Client Information

Table reports information measures by client type. Data include all market-making EUR-USD trades by a top-10 dealing bank with non-dealing-bank counterparties during the 68 trading days from 2 January through 20 April 2012.

	Hedge Fund	CTF	Client Bank	Retail Broker	Institution Broker	Other Broker	Real Money	MNC	SME	Private Clients
A. Post-trade Returns, 1 min. (pips)										
All Trades, Mean	0.349†	-0.017	-0.013	0.136†	0.003	0.110†	0.066	0.046	0.054	-0.381
Standard Deviation by Trade	3.069	3.277	3.044	3.441	2.940	3.521	2.783	2.957	3.077	2.720
Standard Deviation by Client	0.95	0.80	1.01	1.05	0.324	0.962	1.33	1.54	2.505	2.564
Tiny: [€0,€0.5) mn. Mean	0.357†	-0.100***	-0.056†	0.125†	0.126**	0.051	0.065	-0.240	0.406	-0.375
Standard Deviation	2.848	3.170	3.017	3.492	2.607	3.647	2.514	2.593	3.071	2.775
Small: [€0.5,€1.0) mn. Mean	0.430†	0.046*	0.068**	0.141†	-0.033	0.149**	-0.039	-0.256	0.552	0.500
Standard Deviation	3.178	3.380	3.173	3.167	2.692	2.965	2.961	2.855	3.197	1.924
Medium: [€1.0,€5.0) mn. Mean	0.271†	0.042*	0.079**	0.274†	-0.008	0.436†	0.100	0.069	-0.152	-2.333
Standard Deviation	3.324	3.346	3.024	3.260	3.359	3.197	3.000	3.111	3.078	0.577
Large: [€5.0,€∞] mn. Mean	0.284*	4.000**	0.178**	0.038	0.269	0.087	0.233	-0.207	-0.367	—
Standard Deviation	2.896	4.528	0.813	2.3068	4.023	2.727	2.655	2.940	2.781	—
B. Post-trade Returns, 30 min. (pips)										
All Trades, Mean	0.577†	-0.005	-0.011	0.042	-0.062	0.035	-0.234	-0.304	0.054	0.602
Standard Deviation	13.852	15.175	14.266	14.750	13.640	14.688	12.908	14.320	15.073	16.17
C. Num. Obs.										
All trades	6,624	19,792	73,730	79,397	12,703	17,660	1,185	1,261	3,660	113
[€0,€0.5) mn	1,392	8,409	41,429	65,520	1,793	13,392	416	248	3,335	104
[€0.5,€1.0) mn	1,838	1,526	11,419	8,754	6,613	2,035	296	133	165	6
[€1.0,€5.0) mn	3,090	9,852	17,874	4,979	4,046	2,181	341	763	145	3
[€5.0,€∞] mn	304	5	1,038	144	251	52	132	117	15	—

Table 5: Client Information and Client Markups

Table reports coefficients from estimating the markup decomposition model, Equation (3). Each observation is a client trade. Dependent variable is the absolute difference between the client's price and the prevailing core price or "markup." *Client Information*: client's average 1-minute post-trade return. *Trading Volume*: client's log trading volume over the sample period. *DTShr*: clients' share of trades handled directly. *Skew*: the predictability of client c 's trade direction interacted with *DTShr*. *Prime Broker*: indicates clients that trade under a prime brokerage agreement; *LFD_{NM}*: indicates client banks from countries with low financial development that do not trade over MRFQs. *SME_{NM}*: SMEs that do not trade over MRFQs. *Time*: indicators for European hours, New York hours, or overnight trading (Asian hours excluded). *Volatility*: (log) high-low range over the previous hour. *Sz*: trade t 's (log) absolute amount in EUR. *DT*: indicator for direct trades. Data include all market-making EUR-USD client trades by a top-10 dealing bank during the first 68 trading days of 2012. Robust standard errors clustered by date. Significance at the 5%, 1%, and 0.1% levels indicated by **, ***, and †, respectively.

	1-min Info	30-min Info
Client Information		
<i>x Hedge Fund</i>	0.622†	0.303†
<i>x CTF</i>	1.084†	0.301
<i>x Client Bank</i>	0.371†	-0.004
<i>x Retail Broker</i>	-0.837†	-0.167†
<i>x Institutional Broker</i>	-0.233†	-0.149†
<i>x Other Broker</i>	-0.539†	0.045†
<i>x Real-money</i>	0.242	-0.049
<i>x MNC</i>	-0.726	-0.135
<i>x SME</i>	0.299	-0.045
<i>x Private Client</i>	-1.193	-0.139
Controls		
Client Trading Volume	-0.130†	-0.122†
Client Sophistication		
<i>Direct</i> (e-2)	-6.577†	-6.371†
<i>DTShr</i> ² (e-2)	0.117†	0.115†
<i>Skew</i> (e-2)	4.447†	4.435†
Client Access to Counterparties		
<i>Prime Broker</i>	-0.395†	-0.410†
<i>LFD_{NM}</i>	0.585†	0.583†
<i>SME_{NM}</i>	10.263†	10.303†
<i>Priv. Client</i>	22.714†	23.129†
Order-handling, Inventory, Smoothing		
<i>Constant</i>	2.167†	1.972†
<i>Sz</i> (e-2)	-0.699†	-0.098
<i>DT</i>	29.513†	29.635†
<i>DT x Sz</i>	-1.938†	-1.946†
<i>Volatility</i> (e-2)	-0.073**	-0.086†
<i>Time: Europe</i>	0.101†	0.101†
<i>NYC</i>	0.080†	0.065†
<i>Overnight</i>	-0.065**	-0.071†
<i>Core Spread</i>	0.007	0.007
Adj. R^2	0.583	0.582
N. Observations	257,241	257,241

Table 6. Robustness Tests

Table reports coefficients from estimating variants of the markup decomposition model, Equation (3). Each observation is a client trade. Dependent variable is the absolute difference between the client's price and the prevailing core price or "markup." Variables in the baseline regression defined in Table 5. Price shading: adds a measure of current inventory. European Hours: sample restricted to trades during the most active trading hours. FX Experts: adds an indicator for hedge funds and non-LFD client banks. *Client sale*: adds indicator for client EUR sales. *Alternative Volatility*: Realized volatility of 1-minute returns over the previous hour replaces the high-low range. Data include all market-making EUR-USD client trades by a top-10 dealing bank during the first 68 trading days of 2012. Robust standard errors clustered by date. Significance at the 10%, 5%, 1%, and 0.1% levels indicated by *, **, ***, and †, respectively.

	Price Shading	European Hours	FX Experts	SMEShr & MRFQShr	Client Sale Indicator	Alternative Volatility
	1	2	3	4		5
Client Information						
<i>x Hedge Fund</i>	0.644†	1.055†	0.874†	0.635†	0.630†	0.616†
<i>x CTF</i>	1.092†	1.380†	1.272†	1.225†	1.082†	1.086†
<i>x Client Bank</i>	0.371***	0.418***	0.419†	0.378***	0.369***	0.373***
<i>x Retail Broker</i>	-0.527†	-0.630†	-0.865†	-0.583†	-0.537†	-0.543†
<i>x Inst. Broker</i>	-0.822†	-1.160†	-1.100†	-1.123 †	-0.841†	-0.832†
<i>x Other Broker</i>	-0.222†	-0.267†	-0.772†	-0.332†	-0.223†	-0.226†
<i>x Real-money</i>	0.241	0.299	0.241	0.246	0.242	0.241
<i>x MNC</i>	-0.728	-0.692	-0.713	-0.728	-0.728	-0.734
<i>x SME</i>	-0.300	0.226	-0.300	-0.299	0.300	0.300
<i>x Private Client</i>	-1.193	-0.788	1.193	-1.193	-1.192	-1.193
Controls						
Client Trading Vol.	-0.130†	-0.163†	-0.180†	-0.144†	-0.130†	-0.130†
Client Sophistication						
<i>DTShr</i> (e-2)	-6.575†	-0.079†	-6.239†	-6.642†	-6.566†	-6.577†
<i>DTShr</i> ² (e-2)	0.117	0.142†	0.113†	0.116†	0.116†	0.117†
<i>Skew</i> (e-2)	4.448†	3.900†	4.433†	4.449†	4.455†	4.448†
<i>SBPShr</i> (e-2)				-0.126†		
<i>MRFQShr</i> (e-2)				-0.147†		
<i>FX Experts</i>			-0.371 †			
Client Access to Counterparties						
<i>Prime Broker</i>	-0.389†	-0.610†	-0.276†	-0.453†	-0.395†	-0.392†
<i>LFD_{NM}</i>	0.584†	0.439†	0.717†	0.594†	0.581†	0.585†
<i>SME_{NM}</i>	10.259†	9.989†	9.927†	10.241†	10.284†	10.257†
<i>Private Client</i>	22.709†	21.761†	22.301†	22.668†	22.724†	22.706†
Order-handling, Inventory risk, Smoothing						
<i>Constant</i>	2.207†	2.817†	2.784†	2.338†	2.212†	2.154†
<i>Sz</i> (e-2)	-0.684†	-0.830†	2.031†	0.514**	-0.731†	-0.671†
<i>DT</i>	29.509†	26.854†	29.359†	29.518†	29.483†	29.516†
<i>DT x Sz</i>	-1.937†	-1.693†	-1.932†	-1.939†	-1.936†	-1.938†
<i>Inventory</i> (e-2)	-0.021*					
<i>Volatility</i> (e-2)	-0.081**	-0.060	-0.078**	-0.073**	0.071**	0.093
<i>Client Sale Indicator</i>					-0.074†	
<i>Time: Europe</i>	0.102†	—	0.104†	0.106†	0.100†	0.092†
<i>NYC</i>	0.078†	—	0.037†	0.084†	0.080†	0.079†
<i>Overnight</i>	-0.065**	—	-0.111†	-0.062†	-0.694***	-0.056**
<i>CSprd</i>	0.006	0.008	0.007	0.00	0.007	0.007
Adj. <i>R</i> ²	0.583	0.5894	0.584	0.591	0.583	0.583
N. Observations	257,241	257,141	257,241	181,098	257,141	257,141

Table 7. Trades with Negative Markups

Data include all market-making client trades in EUR-USD by a top-10 dealing bank during the first 68 trading days of 2012. Significance at the 5% level indicated by **.

	Mean Size	Client's Post-trade Return (pips)		
		1 min	2 min	5 min
Markup < 0	2.40	0.217	0.472**	0.522
Other trades by clients with negative markups	2.08	0.117	0.096	0.201
Markup = 0	0.47	0.004	0.005	0.007
Markup > 0	0.67	-0.003	-0.002	0.000

Table 8: Adverse Selection vs. Information Chasing

Table reports coefficients from estimating the markup decomposition model, Equation (3). Each observation is a client trade. Dependent variable is the absolute difference between the client's price and the prevailing core price or "markup." Variables in the baseline regression defined in Table 5. Data include all market-making client trades by a top-10 EUR-USD dealing bank during the first 68 trading days of 2012. Robust standard errors clustered by date. Significance at the 5%, 1%, and 0.1% levels indicated by **, ***, and †, respectively.

	Info: 1-min.	Info: 30 min.
Client Information		
<i>x Hedge Fund</i>		
<i>Low Latency</i>	1.352†	0.476†
<i>Other</i>	0.383†	0.235†
<i>x CTF</i>	1.083†	0.300
<i>x Client Bank</i>	0.371***	-0.004
<i>x Retail Broker</i>	-0.540†	-0.168†
<i>x Institutional Broker</i>	-0.838†	-0.149†
<i>x Other Broker</i>	-0.237†	-0.045†
<i>x Real-money</i>	-0.241	-0.048
<i>x MNC</i>	-0.726	-0.135
<i>x SME</i>	0.299	-0.045
<i>x Private Client</i>	-1.193	-0.139
Controls		
Client Trading Volume	-0.130†	-0.122†
Client Sophistication		
<i>DTShr</i> (e-2)	-6.575†	-6.367†
<i>DTShr</i> ² (e-2)	0.117†	0.115†
<i>Skew</i> (e-2)	4.437†	4.423†
Client Access to Counterparties		
<i>Prime Broker</i>	-0.515†	-0.501†
<i>LFD_{NM}</i>	0.583†	0.583†
<i>SME_{NM}</i>	10.258†	10.299†
<i>Private Client</i>	22.701†	23.117†
Order-handling, Inventory risk, Smoothing		
<i>Constant</i>	2.164†	1.968†
<i>Sz</i> (e-2)	-0.634†	-0.075
<i>DT</i>	29.495†	29.614†
<i>DTxSz</i>	-1.937†	-1.944†
<i>Volatility</i> (e-2)	-0.072**	-0.086***
<i>Time: Europe</i>	0.101†	0.101†
<i>NYC</i>	0.079†	0.064†
<i>Overnight</i>	-0.065**	-0.071***
<i>Core Spread</i> (e-2)	0.007	0.007
Adj. <i>R</i> ²	0.583	0.582
N. Obs.	257,241	257,241

Figure 1. EUR-USD, 2 January 2011 through 4 May, 2013

EUR-USD level and realized volatility. Daily data. Realized volatility calculated from daily returns over centered 20-day periods.




Figure 2. Client Sophistication and Institutional Constraints

Share of clients that rely on a given trading platform, by client trading frequency. *Rare*, *Infrequent*, *Moderate*, *Active*, and *Most Active* are buckets of increasing trade frequencies partitioned at once monthly, once weekly, once daily, and ten times daily. *LFD client banks* come from countries in the lowest third of the IMF's Financial Development Index for 2012. Data include all liquidity-providing EUR-USD client trades by a top-10 dealing bank during the first 68 trading days of 2012.

