Price Discrimination in OTC Markets

Abstract

This study develops and tests a spread decomposition model tailored to dealer-client trades in OTC markets. The model incorporates two strategic options available primarily to OTC dealers: first-degree price discrimination and intertemporal optimization within client relationships. Our data comprise the complete record of client trades at a top-10 dealing bank in a major OTC market. Dealers price discriminate in response to client trading activity, market sophistication, and private information. Dealers optimize intertemporally by providing volume discounts, dealing strategically, and smoothing client spreads. Average spreads vary substantially across clients and price discrimination accounts for over 80% of that variation.

This paper develops and tests a spread decomposition model tailored to dealer-client trades in two-tier OTC markets. The analysis focuses on three major hypotheses. First: OTC dealers engage in first-degree price discrimination, meaning they discriminate based on client characteristics (Waldfogel, 2005). We hypothesize, more specifically, that dealers discriminate based on three client characteristics: trading activity, market sophistication, and private information. Second: OTC dealers optimize intertemporally within client relationships (Bernhardt et al., 2005), in part by stabilizing D2C spreads over time. Third: price discrimination generates substantial differences across clients in average OTC execution costs.

The strategic options of first-degree price discrimination and intertemporal optimization for each client are scarce or impossible in order-driven markets or call markets. We trace this difference to differences in trading-game structures. In OTC markets liquidity demanders move first by requesting quotes. This removes client anonymity; the cost of search gives dealers non-zero market power during the client interaction, which enables the dealer to condition quotes on counterparty characteristics. In order-driven markets, liquidity suppliers move first by posting quotes; in call markets, all agents post quotes simultaneously. Trading in both cases is anonymous and liquidity suppliers generally cannot tailor prices to individual counterparties.

The literature highlights three client characteristics that could motivate rational price discrimination. Bernhardt et al. (2005) show that rational dealers should quote narrower spreads to clients that trade more actively, to encourage those clients' repeat business. Strong evidence consistent with this strategy, which is akin to a volume discount, exists for the London Stock Exchange (Reiss and Werner, 1996; Hansch and Neuberger, 1996; Bernhardt et al., 2005), the US corporate bond market (Goldstein et al., 2007), and FX dealer-client trades (Osler et al. 2011).

Market sophistication refers to a suite of client features such as familiarity with different trading platforms, number of dealing relationships, and negotiating skills. Some features reduce search costs and should, in theory, induce dealers to set narrower spreads (Duffie et al., 2005). Green et al. (2007) present a model in which negotiating skill brings narrower D2C spreads, and provide evidence for the relevance of negotiating skill in the U.S. municipal bond market.

A client's tendency to be privately informed could elicit one of three rational dealer responses. A dealer could simply ignore information for clients who do not seek to be informed, such as SMEs. A dealer could rationally quote narrower spreads to clients whose information can be profitably exploited via parallel trades in the core market (Naik et al., 1999). The goal of this practice, known as strategic dealing, is to encourage the repeat business of informed clients. Evidence consistent with strategic dealing in FX is presented by Osler et al. (2011). Finally, a dealer could quote wider spreads to clients whose information he cannot profitably exploit because that information poses adverse-selection risk (Glosten and Milgrom, 1985; Easley and O'Hara, 1987). Evidence for adverse-selection risk in D2C spreads has been notable for its absence (e.g., Bernhardt et al., 2005; Goldstein et al., 2007; Osler et al., 2011).

The hypothesis that OTC dealers optimize intertemporally within client relationships emerges logically from the fact that OTC clients maintain a limited number of long-term dealing relationships. Hendershott et al. (2020) show that most U.S. insurance firms maintain between five and 20 municipal-bond dealing relationships. Volume discounts (Bernhardt et al., 2005) and strategic dealing (Naik et al., 1999), discussed above, are both forms of intertemporal optimization: setting narrower spreads now to encourage attractive clients to bring more of their future trades. Another form would be smoothing spreads for a given client over time, which could be rational because customers in any market tend to find frequent price changes "confusing, frustrating, and annoying" (Dholakia, 2016, p. 1). Consistent with smoothing, FX dealers surveyed by Cheung and Chinn (2001) report that their quoted spreads primarily depend on the spread's conventional level; maximizing anticipated profits on a given trade is of secondary importance.

Our third main hypothesis is that price discrimination adds substantial cross-client variation to execution costs. In order-driven or call markets, variation across clients arises solely from observable features of the trade or the market environment: trade size, trade venues, time of day, volatility, etc. In OTC markets dealers can observe client identities as well as these trade features, so they can vary spreads along additional dimensions.

We develop a spread decomposition model tailored to OTC markets that enables us to test these three hypotheses. The model disaggregates the D2C half-spread, *ClientHS,* into two equilibrium prices. The markup, $Markup_t$, meaning the gap between the client price and the same-side core-market price; and the contemporaneous core-market half-spread, *CoreHS*:

$$ClientHS = Markup + CoreHS. \tag{1}$$

The dealer determines the markup after viewing the core FX spread, which is determined by the professional trading community in order-driven trading.[1]

Our data comprise the complete record of dealer-client (D2C) transactions in the world's most active OTC market, FX. We consider at a top-10 dealing bank in one of the most actively traded contracts in that market, spot EUR-USD. For each trade our data include not just the usual basics – price, quantity, direction, and time – but also the client's identity, precise markup, trading platform, and one of six client types: Hedge Fund, Client Bank, Broker, Real-money Fund (low-leverage institution such as a pension or mutual fund), Multinational Corporation (MNC), or Small and Medium Enterprise (SME).

The model's dependent variable, the trade's markup, helps us understand dealer pricing strategies because the markup is entirely under the dealer's control. Client identifiers enable us to capture the separate influences of each client characteristic. Client trading activity is measured directly. Market sophistication is captured by the frequency with which a client chooses specific trading venues. Traditional direct trades are the least sophisticated, single-bank platforms are intermediate, and APIs and request-for-quote systems are the most sophisticated. Private information is captured by a client's average post-trade returns (Anand and Subrahmanyam, 2008).

The absence of client identifiers has constrained previous studies of OTC spreads to focus on one client characteristic at a time and to use just one variable, trade size, to proxy for all three characteristics. Trade size proxies for client trading activity in Reiss and Werner (1996), Hansch and Neuberger (1996), and Bernhardt et al. (2005); trade size proxies for market sophistication in Green

[1] The "core" market in FX was originally the interdealer market. Today non-dealers participate via prime brokerage contracts, and the relevant platforms include some in which non-dealers and dealers trade all-to-all. The professional trading community comprises active and sophisticated traders under prime brokerage contracts.

et al. (2007); and trade size proxies for client information in Osler et al. (2011). To our knowledge, the absence of client identifiers has also impeded previous research from documenting the magnitude of cross-client variation in spreads or investigating the sources of that variation.

To capture intertemporal smoothing the model includes market volatility, the interdealer half-spread, and dummies for periods in the 24-hour trading day. The model also includes controls for operating and inventory costs. Estimation relies on OLS with robust standard errors.

The model achieves an R^2 s between 0.5 and 0.6 when estimated on the full sample. This far exceeds R^2s from related studies, which range from 0.10 to 0.33 (e.g., Hau et al., 2021; Goldstein et al., 2007; Bernhardt et al., 2005). The results falsify none of our hypotheses and provide strong support for the model, support that is sustained over numerous robustness tests.

As hypothesized, OTC dealers price discriminate in favor of active clients and sophisticated clients (Bernhardt et al., 2005; Duffie et al., 2005; Green et al., 2007). Dealers respond to client information in multiple ways. The information variable is insignificant for SMEs, MNCs, and Real-money Funds. Information has a significantly negative coefficient for Brokers, consistent with strategic dealing (Naik et al., 1999), and a significantly positive coefficient for Hedge Funds and Client Banks, consistent with adverse selection (Glosten and Milgrom, 1985).

The insignificance of information for Utilitarian clients is consistent with findings in Bjønnes et al. (2021) that such clients' trades tend to be uninformed. To explain the choice between strategic dealing and adverse selection we consider whether a client's information is incorporated into prices swiftly or slowly. We hypothesize that strategic dealing is optimal if price discovery is slow enough for the dealer to exploit the client's information via parallel trades in the core market (Naik et al., 1999). With other informed clients dealers rationally protect themselves from adverse selection. To test this hypothesis we partition Hedge Funds into those that do and do not engage in high-frequency trading (HFT). HFT in FX is of captures violations of parity conditions that disappear rapidly, possibly too rapidly for profitable exploitation (Budish, Cramton, and Shim, 2015; Farooq et al., 2008). Consistent with our hypothesis, the dealers' adverse-selection response is 50% stronger for HFT Funds than other Hedge Funds.

Taken together, our findings on price discrimination indicate a striking difference between order-driven and OTC markets in the role of adverse selection. In order-driven markets adverse selection between counterparties widens every agent's spread and does so by the same amount. (Glosten and Milgrom, 1985; Easley and O'Hara, 1987). In dealer-client trades, by contrast, adverse selection between those counterparties is irrelevant for some clients, specifically those that are uninformed or whose trades the dealer selects to attract. In our sample adverse selection between dealer and client is irrelevant for Brokers, Real-money funds, MNCs, and SMEs. When adverse selection between dealer and client does influence markups its influence is not fixed but varies according to the client's information. Nonetheless, adverse selection for a different set of counterparties is relevant for dealer-client trades: adverse selection among dealers influences the core half-spread (McGroarty et al., 2007), which is included in every client half-spread.

The evidence we provide for volume discounts and strategic dealing supports our hypothesis that dealers optimize intertemporally. The coefficient on volatility is negative, which could seem surprising given the consistently positive influence of volatility on spreads in order-driven markets. In dealer-client trades, however, the negative coefficient is predicted by smoothing given the positive influence of volatility on core spreads. Similarly, the coefficients on intraday time dummies imply that dealers quote larger markups during London and NY trading hours, when core spreads are narrow, and larger markups during the brief "overnight" period when core spreads are wide.

The hypothesis that first-degree price discrimination generates substantial cross-client variation in dealer-client spreads gains strong support from the empirical estimates of our model. The variance directly attributable to price discrimination is 59.4 pips^2, roughly twice the variance directly attributable to all other factors, 31.6 pips^2. The influence of price discrimination on cross-client variation in dealer-client spreads also encompasses the strong positive correlation between price discrimination and the other factors, which contributes an additional 72.1 pips^2 to variance. Overall, price discrimination contributes 80.1% of the cross-client variance in markups. Within the markup's price discrimination component, sophistication dominates the other client characteristics, contributing 90.1% of cross-client variation.

Our flexible spread decomposition model could readily incorporate additional client characteristics. Nonetheless, the three characteristics already included may appropriately span the other characteristics currently suggested by the literature. The number of a client's dealing relationships (Hau et al., 2021; Hendershott et al., 2020), should be encompassed by client sophistication. The strength of the dealer-client trading relationship (Chernenko and Sunderam, 2014; Cocco et al., 2009; Han and Nikolau, 2016) should be encompassed by client trading volume with a given bank. Client credit risk could be relevant for other OTC contracts, but spot FX dealers report that credit risk is controlled with credit limits, prime brokerage contracts, and margins. A final variable of potential interest, country risk, could influence client markups because it reflects in part the risk of financial instability or the imposition of capital controls. However, this variable is not client-specific and it would contribute to inventory risk rather than price discrimination. We investigate the relevance of country risk in a robustness test.

Our conceptual framework can be interpreted as an extension of a key finding from Duffie et al. (2005): in equilibrium, a given client will be quoted the same spread by every identical dealer, and those client-specific spreads vary according to the client's search costs. We generalize search costs to market sophistication and allow dealers to vary spreads in response to additional client characteristics. We then label the aggregate of dealer responses to client characteristics as the price-discrimination component of the markup.

Our analysis brings into focus four fundamental features of dealer-client trades that are not shared by order-driven trades. First, dealer-client execution costs include two separate equilibrium prices, the markup and the core half-spread, that are determined by distinct sets of agents. Second, liquidity demanders in dealer-client trades are not anonymous and dealers have instantaneous market power. Client characteristics influence dealer-client spreads through first-degree price discrimination. Third, dealers intertemporally optimize within client relationships, in part by smoothing spreads over time. Fourth, adverse selection between trade counterparties can be irrelevant for some clients and its influence for other clients varies with their private information; further, private information can motivate dealers to narrow spreads rather than widen them. Because of these features, the most familiar spread decomposition models (e,g,, Glosten and Harris, 1988; Huang and Stoll, 1997; Madhavan et al., 1997) may be inappropriate for dealer-client trades.

Variation in the spreads set by one dealing institution, the focus of this study, is distinct from "dispersion" in spreads, a label that applies to variation in prices quoted to one client by multiple dealing institutions (Burdett and Judd, 1983). Dispersion in OTC spreads is well-documented (Garbade and Silber, 1976; Reinganum, 1979; Jankowitsch et al., 2011). Garbade and Silber attribute such dispersion to differences across market makers in inventory levels, expectations of future prices, trading costs, familiarity with other dealers' current quotes, and information.

Our study extends a long tradition of research on comparative financial market structures (e.g., Garbade and Silber, 1979; Madhavan, 1992; Glosten, 1994; Biais et al., 1998; Viswanathan, 2002; Collin-Dufresne et al., 2020). It is also relevant to the stability of market structures. In FX, the rise of retail trading platforms was largely a response to the cavernous spreads dealers set for small traders.[2] In the corporate bond market of the early 1900s, trading shifted from order-driven to OTC (Biais and Green, 2007). This shift could have been initiated by the most active and sophisticated clients, who could benefit from discounts unavailable in order-driven trading. The departure would tend to reduce exchange trading volume and raise average exchange spreads, thereby encouraging more market participants to seek narrower spreads in OTC trading. Under certain conditions this self-reinforcing process could have proceeded until exchange-based trading was no longer economically viable.

The rest of this paper has five sections and a conclusion. Section I describes our data. Section II develops our spread decomposition model for OTC markets. Section III examines the estimates of implications for first-degree price discrimination and smoothing. Section IV examines cross-client variation in markups. Section V presents robustness tests. Section VI summarizes our findings.

I. DATA

Our data come from the world's largest OTC market, spot FX, where trading averages \$2.0 trillion per day (Bank for International Settlements, 2019). They comprise the complete front-office

[2] In the early 1990s the standard OTC spread for tiny trades was 3%, according to the traders interviewed by Osler during that period.

trade record of a top-10 dealing bank over 68 trading days from 2 January to 20 April 2012. We extract from this all spot trades in EUR-USD, the most liquid currency pair and perhaps the world's most liquid single financial contract. With daily turnover on the order of $500 billion (Bank for International Settlements, 2019), EUR/USD trading is similar in magnitude to daily trading in all US Treasury securities or all US stocks.[3]

The data are extraordinarily detailed. They include standard information about each trade – date, time to the second, transaction price as $/€, and quantity traded in euros – plus client ID, client type, client country of origin, exact markup, trade initiator, trade direction (bank buys or sells), trading platform, a flag for trades priced by humans, and full details for trades on which the bank serves as prime broker. To focus on dealer-client trades, we exclude the bank's internal trades and trades that are generally carried out on matching engines, with the latter specified as prime broker trades and trades with other top 50 *Euromoney* dealing banks. We also exclude the few dealer-client trades on matching engines or for which our bank is a price taker. After exclusions the sample includes 257,421 deals worth roughly €126.0 billion in total (Table 1, Panel A).

The data have many advantages for examining OTC liquidity prices, foremost among them the inclusion of client identifiers and many client types. OTC datasets with dealer-client trades often provide no information about clients (e.g., Green et al., 2007) or group clients into a few broad types, such as financial vs. commercial (e.g. Evans and Lyons, 2005; Osler et al., 2011) or retail vs. institutional (e.g., Anand and Subrahmanyam, 2008). Also advantageous is the identification of over 20 trading platforms, which help us identify the influence of sophistication. Finally, our data provide a precise measure of the dealer's strategic variable, the client's markup over the same-side core-market price. Most spread decomposition models approximate execution costs as price changes across trades (e,g,, Glosten and Harris, 1988; Huang and Stoll, 1997; Madhavan et al., 1997).

Until the early 2000s most FX markups were tailored to the client by a salesperson. With electronic trading most markups are now set by an algorithm designed and parameterized by the

[3]US Treasuries: https://www.statista.com/statistics/189302/trading-volume-of-us-treasury-securities-since-1990/ (accessed 1/26/2022.). US equities: https://www.cboe.com/us/equities/market_share/ (accessed 2/23/2022). https://www.nyse.com/markets/us-equity-volumes.

salespeople together with a dedicated team of financial engineers or "e-traders." The bank we study, like most large dealing banks today, internalizes most of its dealer-client trades.

A. FX Spot Clients

As shown in Table 1, our sample includes 2,286 active clients of which 43 are Hedge Funds (*HF*).[4] There are 731 Client Banks (*CBk*) in two groups: small and/or regional banks that focus on client service and sophisticated niche banks that develop high-frequency trading algorithms. There are 110 Brokers, which also come in two groups: Broker-dealers acting as agents for certain hedge funds and retail FX Brokers, which are online platforms for trading small amounts that lay off their own customer trades with the dealing bank. The sample includes 257 Real-money Funds (*RM*), meaning low-leverage asset managers such as mutual funds, pension funds, insurance firms, or endowments. The sample also includes 169 MNCs and 1,078 SMEs. For expositional convenience we follow Harris (2003) and partition clients into Profit-seeking traders, meaning Hedge Funds, Brokers, and Client Banks, and Utilitarian Traders, meaning Real-money Funds, MNCs, and SMEs.

The average client makes 108 deals during our three-month sample, a figure that ranges across client types from 4 for SMEs to 1,552 for Brokers. Because SMEs trade infrequently they represent over half of the clients but account for just 1.5% of trades and 0.7% of traded value. Trading volume is high for Brokers because FX retail platforms trade small amounts at great frequency. Brokers represent just 5% of clients but 66% of trades and 38% of traded value.

To describe markups and their variation we calculate each client's average markup across deals (client-average markup), and take the mean across clients within client types. The mean client-average markup across the entire sample is 1.8 pips (1.4 bps at the average exchange rate of $1.3/€). Given the mean core half-spread of 0.2 pips, the mean client-average half-spread is 2.0 pips.[5] Client markups range widely, with coefficient of variation 11.8, and tend to be lower for Profit-seeking than Utilitarian clients. Among Profit-seeking clients the mean client-average markup ranges from 0.4 for Hedge Funds to 1.3 pips for Brokers. By contrast, means for Utilitarian clients range from 3.6 pips

[4] The SME label is assigned to private clients and firms with annual sales below €100 mn. An MNC will have a specialized treasury unit or at least €100 mn in annual sales.

[5] Core half-spreads are calculated as the distance from core price to the midpoint of the relevant trading platform,

for Real-money Funds to 32.9 pips for SMEs. The tiny fraction of trades by SMEs represent 30.6% of the bank's markup revenue.

B. Trading Platforms

The clients of our bank trade on more than 20 platforms that we partition into four groups. Table 2 presents each client type's mean client trading shares for each venue group.

Direct trading involves contacting a human dealer directly for prices and execution. The telephone was once the only direct connection; today these connections also include email, fax, and chat rooms. Direct trading remains relatively easy to understand and involves the fewest technological challenges, so it is preferred by the least sophisticated FX clients. However, it is expensive because it involves high operating costs, as shown later. Further, it commits the client to more time interacting with dealers, which increases the dealer's instantaneous market power. The mean client-average share of direct trading ranges from 4% for Brokers to 95% for SMEs.

Single-bank platforms (SBPs) were developed in the late 1990s and each major FX dealing bank now offers a menu of them. On the simplest SBPs a client can instruct the bank to trade a certain quantity at a daily fix. On the most sophisticated SBPs a client can click-and-deal on live-streaming quotes or trade algorithmically over an application programming interface. SBPs require only modest client sophistication because they are installed and maintained by the dealing institution. They are in most frequent use among Client Banks, for which they account for 55% of trades, as well as Real-money Funds and MNCs.

Application Programming Interface (API) connections embedded in some SBPs enable clients to execute algorithmic and low-latency trades. Substantial technical sophistication is required: clients must establish and maintain their own connections to the API; many also create their own trading algorithms. These connections are used relatively frequently by Hedge Funds (25% of trades, on average) and Brokers (23% of trades, on average), though for different reasons. Hedge Funds use APIs for HFT and for the algorithmic execution of larger trades. Retail FX brokers rely on APIs for the efficient execution of myriad tiny trades.

Multibank request-for-quote (MRFQ) platforms allow clients to request and receive quotes from multiple dealing banks simultaneously. This minimizes client search costs and maximizes the competitive pressures on quoting banks. This venue group also includes a few trades over voice brokers, which also have low search costs and high competitive pressure. MRFQ platforms require substantial sophistication to install and maintain. They are used most frequently by Hedge Funds (50% of trades), Brokers (56%), and Real-money Funds (52%).

II. MARKUP DECOMPOSITION MODEL

This section presents our empirical decomposition model for dealer-client execution costs. The model relies on trade-by-trade observations of execution cost with the markup as dependent variable. The markup is ideal because it is the dealer's choice variable. Further, each markup is reported precisely by the bank. An alternative measure of execution cost, the client half-spread, can only be measured imprecisely. To measure it precisely would require the contemporaneous same-side core price, but these are derived from several platforms. Our closest approximation is the difference between the traded price and the nearest same-side price on just one platform, Reuters Dealing. The most common measure for spread decomposition models, trade-by-trade price changes, would be less precise still.

For trade t by client c we disaggregate the markup into four components:

$$Markup_t = PriceDiscimination_{c(t)} + Smootht,_{c(t)} + OpCost_t + Inventory_t + \varepsilon_t \ . \tag{2}$$

The rest of this section introduces sub-models for each markup component.

A. Price Discrimination

The sub-model for $PriceDiscimination_c$ includes each client characteristic separately:

$$PriceDiscimination_c = VolumeDiscount_c + Sophistication_c + Information_c \ . \tag{3}$$

Volume discounts: Client c's volume discount is estimated as a function of its log trading volume, $CTV_c \equiv Ln(TrdVol_c)$, with $TrdVol_c$ measured directly from the data in euros:

$$VolumeDiscount_c = \beta CTV_{c(t)} \ . \tag{4}$$

Bernhardt et al. (2005) predict $\beta < 0$ because dealers rationally encourage the repeat business of active clients by quoting them narrower spreads, a form of intertemporal optimization. Empirical evidence consistent with this hypothesis is presented in three studies of the London Stock Exchange: Reiss and Werner (1996), Hansch and Neuberger (1996), and Bernhardt et al. (2005). Reiss and Werner consider these findings to be notable because they "appear[] at odds with asymmetric information models that predict small uninformed trades will receive more favorable execution" (p. 144). However, these studies rely on a potentially unreliable proxy for client trading volume, and trade size. In our data, for example, average trading volume is almost 600 times higher for Brokers than SMEs but average trade size is small for both, at €0.29 million for Brokers and €0.22 million for SMEs. Concerns about reliability also arise because trade size is influenced by the other client characteristics and because trade size influences inventory risk.

Market sophistication. Existing theory suggests that OTC dealers set narrower spreads for relatively sophisticated clients (Duffie et al., 2005; Green et al., 2007). Highly sophisticated clients are astute negotiators who can bargain hard and they are comfortable with all trading technologies. Evidence for the importance of sophistication for OTC client spreads is presented in Green et al. (2007), Hendershott et al. (2020), and Hau et al. (2021), among others.

The model includes two variables to capture the influence of sophistication:

$$Sophistication_c = \delta_{1,j(c)} + Soph_c\delta_2' \quad . \tag{5}$$

The first variable, $\delta_{1,j(c)}$, $j = \{HF, Broker, CBk, RM, MNC, SME\}$, is a client-type constant intended to capture systematic differences in the incentive to invest in sophistication. We conjecture that Utilitarian clients, which trade rarely, have few incentives to invest and the highest values of $\delta_{1,j(c)}$.[6] $Soph_c$ is a vector of the share of client c's deals conducted on each of four platforms, p = {*Direct, SBP, API-Brokers, API-Others*}; MRFQ is omitted. The Brokers' reliance on APIs is distinguished from that of other clients because APIs provide unique benefits for retail brokers. Based on our earlier observations about cross-client differences we conjecture that $\delta_{2,Direct} > \delta_{2,SBP} > \delta_{2,API\text{-}Others}$.

[6] A handful of MNCs worldwide take currency trading seriously but far more discourage it to avoid important agency problems that are expensive to control (Osler, 2006). Further, dealers are unlikely to condition behaviour on MNC information, which likely applies to a time horizon far longer than the time horizon relevant to OTC dealers.

Client information: The information subcomponent of client c's price discrimination component is estimated as follows:

$$Information_c = Info_c \pi_{j(c)}\ , \qquad j = \{HF, Broker, CBk, RM, MNC, SME\}. \tag{6}$$

We measure $Info_c$, the client's private information, as the client's mean signed one-minute post-trade return, following Anand and Subrahmanyam (2008). Reuters mid-quotes are used to provide exact one-minute horizons, which cannot be achieved with the bank's trade-by-trade data. Results are qualitatively unchanged with a thirty-minute time horizon. Using a yet-longer time horizon, such as a day, might capture private information held by clients such as Real-money Funds. However, such information is unlikely to motivate dealers to price discriminate because they generally close positions within the day.

The mean of $Info_c$ is statistically zero for all client types (Table 2, Panel C) except Hedge Funds, for which it is positive and economically meaningful. For Hedge Funds with $Info_c$ one standard deviation above or below the group mean the difference in returns would be 3.7% if they traded once per day for a full year. If they traded every minute that difference would exceed 12,000%.

We estimate separate responses to $Info_c$ by client type, in part because $Info_c$ is imprecisely measured for clients that trade infrequently. This choice also reflects systematic differences across client types in the incentive to become informed. Most Profit-seeking clients have strong incentives because they rely on foreign currency as a store of value, the exception being retail FX Brokers. SMEs and MNCs have weak incentives because they primarily rely on currencies as mediums of exchange. The coefficient on $Info_c$ will be positive under adverse selection (Glosten and Milgrom, 1985; Easley and O'Hara, 1987), negative under strategic dealing (Naik et al., 1999), and insignificant if dealers consider clients uninformed.

We conjecture that information will be insignificant for Utilitarian traders because they have few incentives to incur the costs of information gathering. As outlined in Osler (2006), SMEs and MNCs trade to pursue profits in real-side commerce and primarily use currency as a medium of exchange. Real-money Funds use currency as a store of value, which motivates the common theoretical assumption that they rationally forecast exchange rates. In However, Real-money Funds generally treat currency trading as an administrative necessity rather than a source of returns (Taylor and Farstrup, 2006; Galanek, 2010).[7]

In sum, client c's price-discrimination component is modeled as follows:

$$PriceDiscrimination_c = \beta CTV_c + \delta_{j(c)} + Soph_c\delta_2' + Info_c\pi_{j(c)} \quad . \tag{7}$$

Other client characteristics studied in the literature, specifically the number or strength of the client's trading relationships (e.g., Hendershott et al., 2020; Hau et al., 2021), should be encompassed by client sophistication and trading volume, respectively. In theory dealers could also price discriminate based on credit risk; in practice, however, FX spreads are not used to manage credit risk, according to dealers. This may partially reflect the inherently low credit risk of such trades, given T+2 settlement in (and just T+1 in North America), but credit risk is a concern nonetheless. Dealing institutions rely instead on credit limits assigned before deals can begin; margins for clients whose credit quality is difficult to ascertain, such as young Hedge Funds; and prime brokerage contracts for Hedge Funds and proprietary trading firms that trade in the core market.

B. Intertemporal Optimization

Our markup decomposition model also incorporates dealers' efforts to optimize intertemporally within client relationships. This strategy is already embedded in $PriceDiscrimination_c$ because it motivates volume discounts and strategic dealing. We hypothesize further that dealers smooth spreads over time for each client, a strategy that supports client relationships in two ways: it eliminates noise that impedes clients from recognizing that a dealer's prices are competitive and it accommodates the clients' normal preference for predictability (Dholakia, 2016). In addition, smoothing could have reputational advantages similar to those that motivate smoothing in the interbank market according to the dealer survey in Cheung and Chinn (1998). Two-thirds of responding dealers report that a spread's the conventional level is the dominant influence on their quoted spreads to each other; the profitability of a given trade is secondary. Dealers explain that "the ability to consistently offer quotes with … conventional spreads in a hectic market is regarded as an essential characteristic of a market leader" (Cheung and Chinn, p. 444).

[7] This could be rational because major exchange rates approximate a random walk and currency analysts are costly.

To smooth client spreads the dealer must move the markup inversely to the core spread, which in turn is determined by factors including market volatility and trading volume (Ho and Stoll, 1981; Chaboud et al., 2004). We estimate $Smooth_{t,c(t)}$ as follows:

$$Smooth_{t,c(t)} = \kappa_1 HL_t + Time_t \kappa_2' + \kappa_3 IBSprd_t \,. \tag{8}$$

HL_t, our main measure of volatility, is the high-low range of the previous trading hour; unlike realized volatility HL_t is always apparent to dealers (results are robust to using realized volatility). HL_t has mean and standard deviation of 31.7 pips and 17.5 pips, respectively. Core FX spreads are positively related to volatility so smoothing implies $\kappa_1 < 0$. $Time_t$ is a vector of time-of-day indicators to capture regular movements in core-market spreads over the global trading day, movements that mirror market-wide trading volume (Chaboud et al., 2004). *Time* includes three indicators corresponding to European trading hours, 8-17 GMT, when volume (the core spread) is highest (lowest); the New York afternoon, 17-22 GMT, when volume (the core spread) is falling (rising); and overnight, 22-3 GMT, when volume (the core spread) is lowest (highest); Asian hours are excluded. Under smoothing, $\kappa_{2Europe} > 0$, $\kappa_{2Overnight} < 0$, and the sign of κ_{2NY} is indeterminate.

$IBSprd_t$ is the interbank spread from Reuters Dealing. Smoothing implies $\kappa_3 < 0$.

C. Operating Costs and Inventory Costs

Market makers must cover operating costs and inventory costs in every trading structure. Sub-models for these components are based on existing theory (e.g., Ho and Stoll, 1981; Lagos and Rocheteau, 2009).

We estimate operating costs as a function of $Sz_t = ln(TrdSz_t)$ and DT_t, a direct-trade indicator:

$$Operating_t = \varphi_1 Sz_t + DT_t \;(\varphi_2 + \varphi_3 Sz_t) \;\;. \tag{9}$$

Theory suggests $\varphi_1 < 0$ and $\varphi_3 < 0$ because larger trades allow dealers to cover fixed costs with a smaller markup. We expect $\varphi_2 > 0$ because direct trades involve costly dealer time.

Inventory Costs. Inventories incur costs of carry and risk. Cost of carry is negligible in spot FX because dealers infrequently hold positions overnight (Bjønnes and Rime, 2005). To capture inventory risk we include volatility, trade size, and the interbank spread, with the latter serving as a measure of market illiquidity (Lagos and Rocheteau, 2009). We estimate the inventory component as follows:

$$Inventory_t = \lambda_1 HL_t + \lambda_2 Sz_t + \lambda_3 IBSprd_t \,. \tag{10}$$

Ho and Stoll (1981) predict λ_2 , $\lambda_2 > 0$ and we expect $\lambda_3 > 0$.

C. Summary: Markup Decomposition Model for Dealer-client Trades

Equations (7), (8), (9), and (10), when combined according to Equation (2), generate the following markup decomposition model for OTC client trades:

$$\begin{aligned} Markup &= \beta CTV + CT\delta_1' + Soph\delta_2' + Info\,\pi' \\ &+ (\kappa_1 + \lambda_1)HL + Time\,\kappa_2' + (\kappa_3 + \lambda_3)IBSprd + (\varphi_1 + \lambda_2)Sz + \varphi_2 DT + \varphi_3 DT\,Sz + \varepsilon. \end{aligned} \tag{11}$$

The first line captures *PriceDiscrimination* (see Equation (7)); *CT* is a 257,241 x 6 vector of client-type dummies. *Smooth*, *OpCost*, and *Inventory* jointly comprise the second line; these three markup sub-components cannot be identified separately given their shared determination by *HL*, *Sz*, and *IBSprd*. The model is estimated using OLS with Newey-West standard errors clustered by date.

Multicollinearity among client characteristics might be a challenge if, for example, higher trading volume motivates investments in both sophistication and private information. To examine this possibility we calculate bilateral correlations among the three client descriptors, using the client's share of direct trading as a scalar measure of sophistication. These correlations indicate that multicollinearity is unlikely to compromise the results: $\rho(Info_c, Soph_c) = 0.10$; $\rho(Info_c, TrdVol_c) = -0.00$; $\rho(Soph_c, TrdVol_c) = 0.27$.

III. PRICE DISCRIMINATION AND SMOOTHING IN OTC MARKETS

The results of estimating the model, shown in Table 3, provide strong support. The model's adjusted R^2 of 0.57 far exceeds R^2s in comparable studies of dealer-client trading, which range from nearly zero to 0.33 (Hau et al., 2021; Goldstein et al., 2007; Bernhardt et al., 2005). Harris and Piwowar (2006) find a similar lack of explanatory power for client spreads in the U.S. municipal bond market, and suggest two alternative explanations. "The [large unexplained] variation may be idiosyncratic or due to an inability of the cost function to well represent average trade costs for all trade sizes" (p. 1378). Our estimates indicate that such variation reflects neither idiosyncrasies nor missing costs and instead reflects first-degree price discrimination.

This section discusses the implications of the estimated coefficients for our hypotheses that OTC dealers engage in first-degree price discrimination and intertemporally optimize within client relationships. The following section discusses implications for our hypothesis that first-degree price discrimination magnifies cross-client variation in execution costs.

A. First-degree Price Discrimination

The results indicate that dealers price discriminate based on all three client characteristics.

Client Trading Volume: The coefficient on client trading volume is negative and significant both statistically and economically, consistent with volume discounts (Bernhardt et al., 2005). A rise in client trading volume from the SME mean to the Hedge Fund mean reduces the markup by one pip, or twice the mean for Hedge Funds. This supports both of the pricing strategies under consideration: first-order price discrimination and intertemporal optimization.

Sophistication: All but one of the ten sophistication coefficients is statistically significant and their signs support our hypothesis that dealers price discriminate against less sophisticated clients. The sophistication intercepts are consistently lower for Profit-seeking clients than Utilitarian clients, as conjectured given differences in the incentive to invest in sophistication. For Profit-seeking clients these range from 2.3 pips to 3.0 pips; for Utilitarian clients they range upward from 3.1 pips. Consistent with our conjecture that SMEs are least sophisticated, their client-type intercept, at 13.5 pips, is over three times the next-largest client-type intercept of 3.7 for Real-money Funds. The coefficients on client platform shares also indicate that dealers price discriminate against less sophisticated clients. As conjectured, $\delta_{2,}Direct > \delta_{2,}SBP > \delta_{2,}API\text{-}Others$, with each coefficient an order of magnitude larger than the next (all differences are statistically significant). Notably, $\delta_{2,}API\text{-}Brokers$ is negative, which implies that APIs are of greater sophistication than MRFQ platforms for Brokers.

Overall, sophistication – or the lack thereof – raises the average markup of SMEs, MNCs, and Real-money Funds by 18.1 pips, 5.6 pips, and 4.5 pips, respectively, relative to their level on an

MRFQ platform. The increases for Profit-seeking clients are consistently smaller, ranging from 2.9 pips for Hedge Funds to 3.2 pips for Brokers.

We evaluate our assumption that client-type dummies capture market sophistication by re-estimating Equation (11) without them (Table 3, second column). The assumption's validity gains strong support. The coefficients on platform shares all increase substantially and the regression R^2 declines only modestly, from 0.57 to 0.52.

Private information. The coefficients on client information, π_j for $j = \{HF, Broker, CBk, RM, MNC, SME\}$, are small and insignificant for Utilitarian traders, from which we infer that dealers consider most such traders to be uninformed. The $Info_c$ coefficients are significant for both Hedge Funds and Brokers, indicating that dealers consider these clients to be informed. However, π_{HF} is positive, consistent with adverse selection (Easley and O'Hara, 1987), and $\pi_{Brokers}$ is negative, consistent with strategic dealing (Naik et al., 1999). A 1-pip increase in $Info_c$ for a Hedge Fund raises its markup by 1.0 pip, roughly twice the Hedge Funds' average markup. A 1-pip increase in $Info_c$ for a Broker lowers its markup by 0.6 pips, roughly half the Brokers' average markup. $\pi_{ClientBanks}$ is positive, consistent with adverse selection, but substantially smaller than π_{HF} and statistically insignificant; in robustness tests $\pi_{ClientBanks}$ remains positive and is sometimes significant.

The dealers' varied responses to client information highlight a striking contrast between order-driven markets and OTC markets. In order-driven markets adverse selection widens the spread for every liquidity demander, informed or uninformed, and does so by an amount that does not vary by counterparty (Glosten and Milgrom, 1985). In OTC markets adverse selection does not influence spreads for uninformed clients and for some informed clients. In our data adverse selection only widens spreads for a subset of clients, specifically Hedge Funds, which account for just 2% of the bank's clients, and possibly for Client Banks, which are 31% of clients. Further, the contribution of adverse selection to markups varies with a client's information: for Hedge Funds with $Info_c$ one standard deviation above and below the Hedge-fund mean, markups will differ by 1.9 pips, four times these clients' mean markup of 0.4 pips.

B. Strategic Dealing vs. Adverse Selection

To further examine the dealers' choice between adverse selection and strategic dealing, we consider the 253 trades with negative markups. Table 4 highlights four features of these trades. First, the trades are large: their mean size of EUR 2.4 million far exceeds the mean size for all clients, EUR 0.5 million. Second, clients that benefit from negative markups tend to make relatively large trades: even their other (positive-markup) trades have mean size of EUR 2.1 million. Third, negative-markup clients appear to be informed: returns to negative-markup trades are significantly positive and annualize to 759%. Returns to these clients' positive-markup trades are also relatively high, averaging between 0.1 pip and 0.2 pips at horizons from 1 to 3 minutes. Though not statistically significant, such returns far exceed mean post-trade returns for all clients, 0.01 pip. These findings indicate that informed clients sometimes deal directly for large trades, presumably to take advantage of the human dealers' skill at minimizing execution costs. Fourth, the positive returns to negative-markup trades are realized slowly: they only become statistically significant at the two-minute horizon. Given today's trading speeds, two minutes is sufficient for dealers to make parallel trades in the core market and exploit their knowledge of an informed client's trade direction. To illustrate, after an informed client buys a dealer would immediately turn to the core market and purchase the same amount (to cover the inventory) or more (to take a proprietary position), knowing that on average the price discovery process would continue in a profitable direction thereafter.

These features of negative-markup trades suggest the following hypothesis regarding the dealers' choice between adverse selection and strategic dealing: Dealers narrow spreads to attract the repeat business of clients whose information influences price slowly enough to permit profitable exploitation via parallel trades in the core market (Naik et al., 1997). For other informed clients dealers widen spreads to protect themselves from adverse selection.

This hypothesis could explain the difference between the dealers' treatment of Hedge Funds and Brokers. Many Hedge Funds in our sample engage in HFT, which in FX commonly involves violations of triangular or covered-interest parity relations. The half-lives of such violations are measured in microseconds (Farooq et al., 2008), so dealers have no opportunity to make parallel trades. Some Brokers, by contrast, trade on an agency basis for Hedge Funds, a process that is too slow for HFT.

To test this time-to-price-discovery hypothesis we partition Hedge Funds into those that do and do not engage in HFT, using the information on their websites. If our hypothesis is correct, the coefficient on $Info_c$ should be larger for HFT Hedge Funds funds than other Hedge Funds. Table 3, final column, shows the results. Consistent with our hypothesis, the coefficient for HFT Hedge funds is roughly 50% larger than the coefficient for other Hedge Funds and the difference is significant.

C. Intertemporal Optimization

The evidence for volume discounts and strategic dealing presented above support our second main hypothesis, that dealers optimize intertemporally within client relationships. Further evidence for intertemporal optimization comes from the coefficients on volatility, the time dummies, and the interdealer spread, which support smoothing. The negative and significant coefficient on volatility supports smoothing because it is the sum of $\kappa_1 < 0$, the influence of smoothing, and $\lambda_1 > 0$, the influence of inventory risk. Time-of-day dummies provide further evidence for smoothing. Coefficients are positive for European and New York trading hours, when core spreads are low, and negative for the overnight hors, when core spreads are high. The coefficient on *IBSprd* may also provide evidence for smoothing. It represents the sum of $\kappa_3 < 0$, the influence of smoothing, and $\lambda_3 > 0$, the influence of overall market illiquidity.[8] Its lack of significance implies that these influences negate each other.

D. Operating Costs and Inventory Costs

Coefficients on the remaining markup determinants – trade size, the direct-trade dummy, and *Time* – can be sensibly interpreted in terms of operating or inventory costs. The coefficient on trade size is insignificant, which indicates that the opposing influences of operating costs, $\varphi_1 < 0$, and inventory costs, $\lambda_2 > 0$, are roughly balanced. By contrast, trade-size coefficients are strong negative when traditional spread decomposition models are applied to OTC markets (e.g., Reiss and Werner,

[8] This coefficient could also be increased by information if informed trades tend to be large (Easley and O'Hara, 1987).

1996; Hansch and Neuberger, 1996; Bernhardt et al., 2005; Green et al., 2007; Gtifa and Liouane, 2013). As conjectured, the direct-trade dummy has a positive coefficient, indicating that direct trades have higher fixed costs. The interaction between the direct-trade dummy and trade size has a negative sign, consistent with the hypothesis that the fixed costs of direct trades are covered by smaller markups on larger trades.

IV. MARKUP LEVELS AND CROSS-CLIENT VARIATION

This section uses the model's estimated coefficients to analyze markup levels and their cross-client variation. The analysis begins with the fitted markup for each trade, $\widehat{Markup_t}$ (a hat denotes fitted values). This includes a fitted price discrimination component and a fitted sum of the smoothing, operating cost, and inventory-cost components, $\widehat{OthMkp_t}$ (which cannot be disaggregated, as noted earlier). The price discrimination component is constant for client *c*, unlike the rest of the markup:

$$\widehat{PD_c} = \hat{\beta} CTV_c + [\widehat{\delta_{1,J(c)}} + Soph_c \widehat{\delta_2}] + Info_c \widehat{\pi_{J(c)}}, \quad j = \{HF, Broker, CBk, RM, MNC, SME\}. \quad (12)$$

$$\widehat{OthMkp_t} = (\widehat{\kappa_1} + \widehat{\lambda_2}) HL_t + Time_t \widehat{\kappa_2} + (\widehat{\kappa_3} + \widehat{\lambda_3}) IBSprd_t + (\widehat{\varphi_1} + \widehat{\lambda_2}) Sz_t + DT_t (\widehat{\varphi_2} + \widehat{\varphi_3} Sz_t). \quad (13)$$

We take the average of these for each client and means across clients within client types.

A. Markup Levels

Figure 1 presents client-type means for $\widehat{Markup}$, $\widehat{PD}$, and $\widehat{Other}$. It is immediately apparent that first-degree price discrimination is the main contributor to markups: it represents 64.1% of the markup, on average, with a minimum share across client types of 55% for MNCs.

Figure 2 presents the separate contributions to price discrimination of client trading volume, sophistication, and information. Volume discounts have mean of -1.4 pips across clients. This is substantial insofar as it exceeds client-average markups for all Profit-seeking client types and is roughly half the mean client-average markup for Real-money Funds. Because volume discounts are negative they reduce the apparent contribution of price discrimination to markup levels. The largest contributor to markup levels is the sophistication sub-component – or "simplicity penalty," for brevity – which averages 10.1 pips for the whole sample and ranges from 3.0 pips for Hedge Funds and Client Banks to 18.1 pips for SMEs. Information contributes almost nothing to mean client-average markups, which is unsurprising because the mean of $Info_c$ within client types is essentially zero for all but Hedge Funds. For Hedge Funds the information sub-component raises markups by 0.1 pip on average.

Figure 1: Markup Components

Figure shows the magnitude in pips of the price-discrimination component, the remaining components, and their sum. Equation (11) is estimated on all EUR-USD client trades at a top-10 dealing bank during the first 68 trading days of 2012. Fitted subcomponents are calculated for each trade, averaged for each client, and then averaged within client categories.

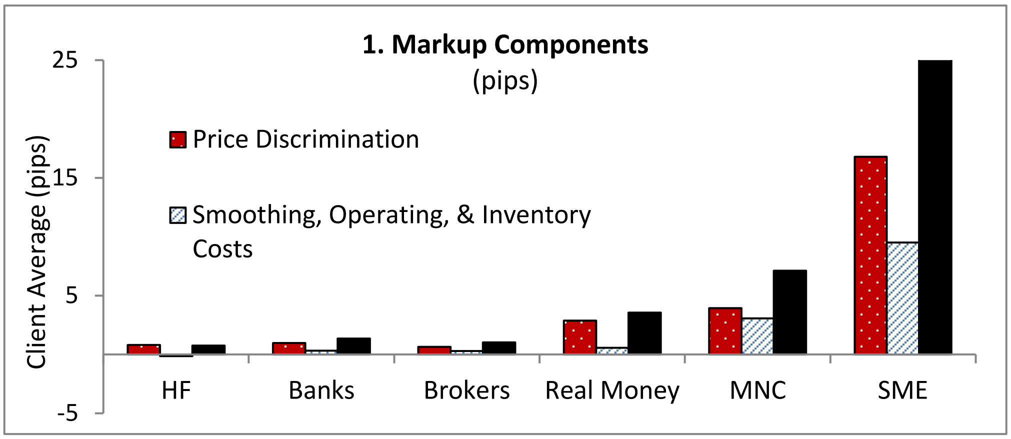

Figure 2: Magnitude of Price Discrimination Component and Sub-components

Figure shows the magnitude in pips of the price-discrimination component and its three subcomponents. Equation (11) is estimated on all EUR-USD client trades at a top-10 dealing bank during the first 68 trading days of 2012. Fitted subcomponents are calculated for each trade, averaged for each client, and then averaged within client categories.

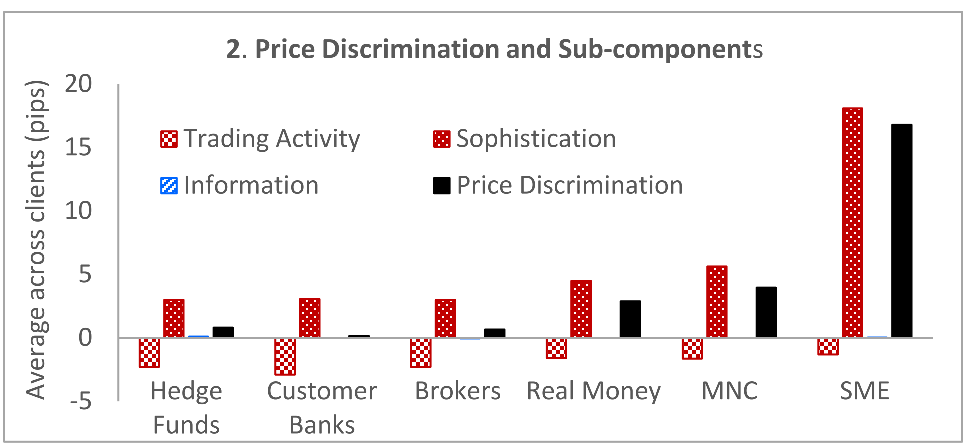

We digress briefly to discuss the relevance of these results for structural instability in financial markets. U.S. corporate bond trading shifted endogenously from the New York Stock Exchange to OTC roughly a century ago (Biais and Green, 2007), and internalization generates a shift in the same direction for equity trading. Biais and Green (2007) suggest that the shift in bonds was driven by institutional traders, for whom spreads were narrower on OTC trades than trades on the NYSE. Our analysis supports this interpretation and highlights a potentially critical role for first-degree price discrimination in the process of structural change. An example will clarify. Assume that OTC dealers price discriminate as hypothesized here and that the initial level of exchange spreads exceeds the level that the most sophisticated traders would pay to OTC dealers. When those agents shift, lower volume on the exchange will raise spreads for the remaining agents, potentially motivating the next-most sophisticated agents to shift to OTC trading. This raises exchange spreads further, etc. If a large enough share of agents shifts to OTC trading, dealers could eliminate exchange trading altogether by setting spreads for the remaining agents slightly below now-high exchange spreads. Thereafter, in the absence of exchange trading, the dealers could potentially raise spreads for the least sophisticated traders yet further.

B. Cross-client Variation in Markups and Half-spreads

We finish our core analysis by examining our third main hypothesis, that first-degree price discrimination substantially increases cross-client variation in dealer-client markups. Table 5 provides descriptive statistics for $\widehat{Mkp_c}$, $\widehat{PD_c}$, and $Oth\widehat{Mkp_c}$. As suggested by Figure 1, cross-client variation in $\widehat{Mkp_c}$ is substantial, with standard deviation 12.8 pips (σ_{Mkp}). Markups vary less among Profit-seeking clients than among Utilitarian clients, with standard deviations of 2.3 pips and 11.0 pips, respectively. Nonetheless, proportionate variation is higher for Profit-seeking clients than Utilitarian clients as revealed by their respective coefficients of variation, 2.1 and 0.5.

To examine the contribution of price discrimination to markup variance we exploit the formula for variance of a sum: $\sigma^2_{Mkp} = \sigma^2_{PD} + \sigma^2_{OthMkp} + 2Cov(PD, OthMkp)$. For the full sample, σ^2_{OthMkp} = 31.6 pips^2 and $\sigma^2_{PD} =$ 59.4 pips^2, indicating that price discrimination more than doubles cross-client variance relative to its level with just operating costs, inventory costs, and smoothing. (Note that the net contribution of smoothing for a given client should average to zero over time.) $Cov(PD, OthMkp)$ is positive, in part because a client's reliance on direct trading raises both the sophistication sub-component of $\widehat{PD}$ and the operating-cost sub-component of $\widehat{OthMkp}$. Covariance represents an indirect contribution of price discrimination because it would be zero in the absence of that strategy.

Taken together, the direct and indirect contributions of price discrimination add 131.5 pips^2 to the cross-client variance of dealer-client markups, raising it by a factor of four and accounting for 80.1% of total variance. As shown in Figure 3, first-degree price discrimination also contributes the bulk of cross-client variation for both Profit-seeking and Utilitarian clients, accounting for 71.4% and 74.4% of variance, respectively.

Figure 3: Sources of Cross-Client Variation in Markups

Figure compares the extent to which price discrimination, *PD*, and other markup determinants, *OthMkp*, contribute to cross-client variance in client-average markups. Contributions shown in percent of total markup variance to facilitate comparisons across groups. The direct contribution for component *i* is $\sigma_i^2/\sigma_{Mkp}^2$, *i* = {*PD*, *OthMkp*}. "Other" markup components are smoothing, operating costs, and inventory costs. Markups and its components are fitted values from estimating Equation (11) on all EUR-USD client trades through a top-10 dealing bank during the first 68 trading days of 2012.

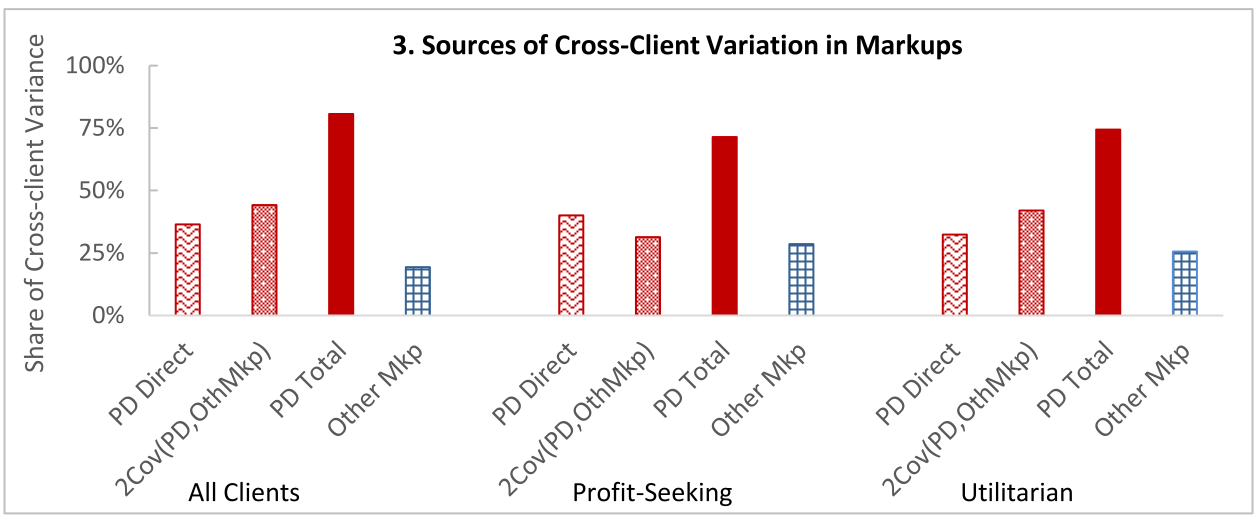

Price discrimination also dominates the variance of a more complete measure of clients' execution costs, the client half-spread. We construct fitted client half-spreads by adding estimated core half-spreads, $\widehat{CoreHS}_t$, to fitted markups:

$$\widehat{ClientHS}_t = \widehat{Mkp}_{c,t} + \widehat{CoreHS}_t = \widehat{PD}_{c(t)} + \widehat{OthMkp}_t + \widehat{CoreHS}_t. \tag{14}$$

$\widehat{CoreHS}_t$ is the near-contemporaneous interbank half-spread on the Reuters Dealing platform. This overestimates the core spread because it comes from one trading platform rather than an aggregator. In consequence it may also over-estimate the variance of the core half-spread and provide conservative estimates of the contribution of price discrimination to cross-client variance.

As shown in Figure 4, which provides the variance decomposition for $ClientHS_t$, price discrimination is the dominant source of cross-client variation for client half-spreads as it is for client markups. The core half-spread has extremely low variance and its covariances with $\widehat{PD}$ and $\widehat{OthMkp}$ are essentially zero. Though Profit-seeking and Utilitarian clients are not depicted separately in Figure 4, price discrimination dominates variation in client half-spreads for both groups, contributing 70.5% and 74.1%, respectively.

Figure 4: Cross-client Variation in Half-spreads

Figure shows the contribution of price discrimination, *PD*, and other determinants of client half-spreads to cross-client variance in client-average half-spreads. Markups, price discrimination, and other markup components ("OthMkp") are fitted values from estimating Equation (11) on all EUR-USD client trades through a top-10 dealing bank during the first 68 trading days of 2012. "Other" markup components are smoothing, operating costs, and inventory costs. Core half-spread taken from Reuters Dealing 3000.

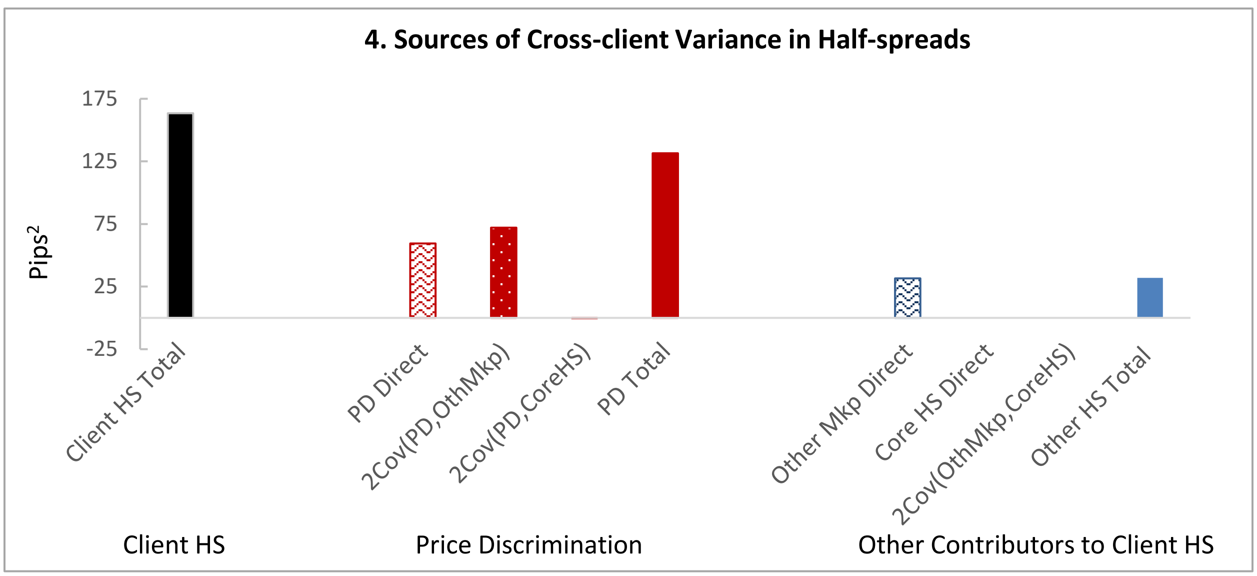

C. Cross-client Variation in Price Discrimination

We finish this analysis by comparing the contributions of individual client characteristics to the cross-client variation in price discrimination components. Theory provides no guidance on these relative contributions. Existing studies are silent as well because they examine individual characteristics in isolation. We calculate fitted values and standard deviations for the sub-components of $\widehat{PD}_c$ according to Equations (4), (5), and (6).[9] Figure 5 presents a variance decomposition.

Market sophistication clearly dominates cross-client variation in $\widehat{PD}_c$. Its direct contribution to variance is 91% across all clients, 87% for Profit-seeking clients, and 94% for Utilitarian clients. The modestly lower contribution of sophistication among Profit-seeking clients is logical given their generally high level of sophistication. The modestly higher contribution of information for these same clients is also logical, given the low incentives of retail FX Brokers to invest in information and the high incentives of other Profit-seeking clients.

Figure 5: Contributions to Cross-client Variation in Price Discrimination

Figure compares the direct contributions of client characteristics to cross-client variance of price-discrimination components, measured as σ_i^2/σ_{PD}^2 for i = {*Volume discount*, *Sophistication*, *Information*}, where σ_i^2 is the variance of the sub-component and σ_{PD}^2 is the cross-client variance of price discrimination components. Figure also shows $(\sigma_{PD}^2 - \Sigma_i \sigma_i^2)/\sigma_{PD}^2$, the total contribution of covariances among subcomponents. Markups and sub-components are fitted values from estimating Equation (11) on all EUR-USD client trades through a top-10 dealing bank during the first 68 trading days of 2012.

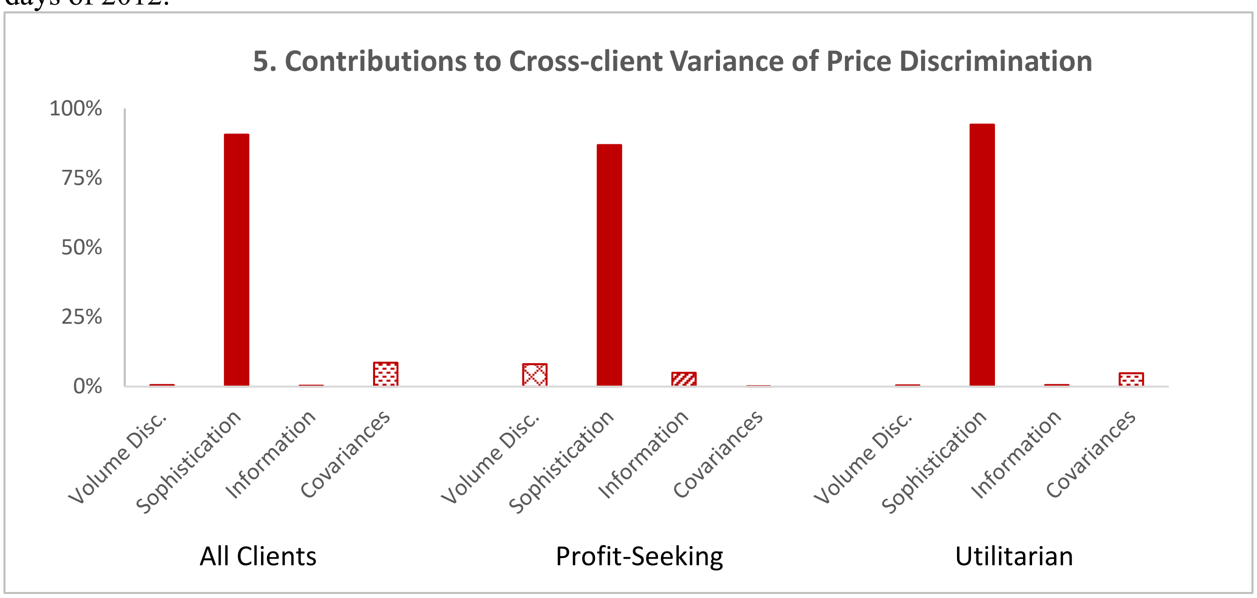

We confirm the dominance of market sophistication by calculating these same shares using the results from estimating Equation (11) without client-type intercepts (Table 3, second column). For

[9] This calculation assumes no influence of information for utilitarian clients, consistent with the econometric results.

the full sample, the direct contributions of trading volume and client information remain below 10% and the direct contribution of sophistication remains high, at 83% (results available upon request).

V. ROBUSTNESS TESTS

This section further examines the model's robustness to modifications in methodology. All tests rely on OLS estimates with Newey-West standard errors.

A. Variable Measurement

We consider alternative measures of market sophistication, private information, and volatility.

Market Sophistication: The share of trading on a given platform is a noisy measure of sophistication. For example, sophisticated clients on occasion deal directly for large trades, as noted earlier. As an alternative we identify each client with their most sophisticated trading platform, with platforms ranked according to the coefficients in our baseline regression (Table 3, first column). That is, in Equation (11) we replace *Soph* with a different 257,241 x 4 matrix, *AltSoph*, in which client *c*'s platform vector is unity for its most sophisticated platform and zero otherwise. The results from this experiment (Table 6, first column) are consistent with our main regression results. Explanatory power rises slightly but the relative magnitudes of the platform coefficients are unchanged. Coefficients on other variables are largely stable.

Private information: Theory does not dictate a specific time horizon for the post-trade returns in $Info_c$. Indeed, the relevant time horizon presumably varies across client types, with shorter horizons relevant to low-latency traders and longer horizons relevant for other clients. We check robustness by calculating post-trade returns over 30 minutes rather than one minute.

This modification leaves unchanged our qualitative findings, including the insignificance of information for Utilitarian clients (Table 6, second column). Nonetheless, it reduces the magnitude of information coefficients for all Profit-seeking client types. The smaller magnitudes could simply reflect increased noise in the information measure, but they could also be economically meaningful. The decline in the Hedge-fund coefficient from 0.98 to 0.26, for example, could reflect in part the dealers' ability to profit from client information that influences prices relatively slowly.

Market Volatility: We replace HL_t with realized volatility over the same 60-minute pre-trade interval, calculated with one-minute mid-quote returns. The new estimates of Equation (11) are almost identical to the original estimates (Table 6, third column).

B. Human Trades

Human dealers priced 3,066 trades in our sample. These trades tend to be large, with average size EUR 6.1 million, though just 1% of all trades they account for 15% of total trade value. Among these trades, 43% are direct, 22% come in through a single-bank platform, and the rest come in through MRFQ platforms. Most human-priced trades are placed by Profit-seeking clients, especially Client Banks, which seek the dealers' help for at least three reasons. First, any trade of €25 million or more is automatically flagged for human intervention, regardless of platform or client. Second, dealers are skilled at minimizing the execution costs of large trades, such as slippage. Third, flags are generated for any trade the dealers consider toxic.

Estimating Equation (11) on the subsample of human-priced trades brings substantial quantitative and qualitative changes to the coefficients (Table 6, fourth column). Most markups become far more sensitive to client characteristics. In terms of quantitative changes, the coefficient on trading volume doubles; most sophistication intercepts rise by 20, which almost triples the SME intercept and raises the others by multiples of six to eleven; most platform-share coefficients and all the information coefficients also multiply in magnitude.

In terms of qualitative changes, note first that the coefficients on SBP and API-Other shares both become negative, like the coefficient on API-Brokers. By implication, human dealers consider clients making large trades to be more sophisticated if they rely on SBPs or APIs than if they rely on an MRFQ platform. This follows logically from the fact that a client requesting quotes for a large trade on an MRFQ platform effectively invites multiple dealers to front-run.

The coefficients on client information remain insignificant for Utilitarian clients but those for Client Banks and Brokers are both positive, significant, and far larger than the original coefficient for Hedge Funds. By implication, dealers consider large trades by Profit-seeking clients to be toxic.

C. Country Risk and Client Identifier

Our fifth robustness test adds country risk to Equation (11), using contemporary risk indicators from *The Economist* that range from 1 (lowest risk) to 4. Country risk encompasses the risk of payment interruption due to financial instability, or political upheaval, or the imposition of capital controls. To illustrate, Afghanistan, Argentina, Iraq, and Venezuela are among the 12% of countries with highest risk, a group that places just 1% of trades. We hypothesize that a client's country risk contributes to the dealer's inventory risk and has a positive coefficient. We also hypothesize that country risk is primarily relevant for high-risk countries because the risk of nation-wide payment interruption is otherwise insubstantial. We add to the model a vector of country-risk indicators as independent variable, with the lowest risk level excluded.

This modification does not influence the qualitative findings discussed previously (Table 6, penultimate column). The coefficients on country risk support our hypothesis that it contributes to inventory risk and does so primarily for high-risk countries. For level-2 countries the coefficient is significant but economically trivial. For level-3 countries the coefficient is a highly significant 0.15 pip; for level-4 countries it is 0.71 pip.

Our final robustness test investigates whether dealers identify clients with their holding company rather than their specific subsidiary (Table 6, final column). This adjustment changes only one minor element of our previous results; two of the small coefficient differences associated with venue choice are no longer statistically significant.

V. CONCLUDING REMARKS

This paper shows that OTC dealers exploit two strategies vis-à-vis clients that are generally unavailable to market makers in other trading structures: first-degree price discrimination and intertemporal optimization within client relationships. It also shows that first-degree price discrimination creates substantial cross-client variation in execution costs.

First-degree price discrimination and intertemporal optimization are possible in OTC markets because liquidity demanders move first, by requesting quotes from a specific dealer, which precludes pre-trade anonymity. In order-driven and call markets, by contrast, liquidity suppliers move first, setting quotes before knowing specific potential counterparties. These trading structures also differ insofar as sequential search in OTC markets gives dealers instantaneous market power.

Our first hypothesis is that dealers price-discriminate in response to three client characteristics: trading activity (Bernhardt et al., 2005), market sophistication (Duffie et al., 2005; Green et al., 2007), and information (Easley and O'Hara, 1987; Naik et al., 1999). Our second hypothesis is that OTC dealers optimize intertemporally within long-term client relationships (Bernhardt et al., 2005), another pricing strategy that relies on the lack of anonymity in OTC trading. We hypothesize that dealers also smooth client spreads over time to accommodate the clients' natural preference for predictability in liquidity pricing (Dholakia, 2016). Our third hypothesis is that first-degree price discrimination generates substantial cross-client variation in execution costs.

To our knowledge, we are the first to examine these pricing strategies comprehensively and to identify the influence of each client characteristic separately. We develop a spread decomposition model tailored to dealer-client trading in OTC markets. In addition to operating and inventory costs, which are common to every trading structure, the model incorporates first-degree price discrimination in response to all three client characteristics and intertemporal smoothing.

We test this model on the complete trading record of a top-10 bank in one of the world's most actively traded contracts, spot EUR/USD. The analysis is feasible because the data include rarely available information such as client identity, client type, precise markup over the core-market price, trade direction, time, and trading platform. The dependent variable is the dealer's strategy variable, the markup. The model's explanatory power substantially exceeds that found in previous analyses of OTC execution costs (e.g., Hau et al., 2021; Goldstein et al., 2007; Bernhardt et al., 2005).

Estimated coefficients for all three client characteristics have signs, magnitudes, and significance consistent with theory. Dealers set narrower markups for their more-active clients, presumably to attract their repeat business as predicted by Bernhardt et al. (2005). They set wider markups for less sophisticated clients, presumably because low sophistication gives dealers greater market power as predicted by Green et al. (2007). Dealers respond in multiple ways to private information. They appear to consider Real-money Funds, MNCs, and SMEs to be uninformed. They set wider spreads for better informed Hedge Funds and Client Banks, consistent with adverse selection (Easley and O'Hara, 1987). They set narrower spreads for better-informed Brokers, presumably to encourage these clients' repeat business as predicted by the strategic dealing hypothesis (Naik et al, 1999).

We posit that dealer's choice between adverse selection and strategic dealing depends on the speed with which a client's information becomes embedded in price. When price discovery is relatively slow, dealers can exploit their knowledge of an informed client's trade direction by making parallel trades in the core market. If the core-market price will not yet embed that information so the parallel trades will be profitable and dealers can rationally set narrower spreads to attract those clients' repeat business. When price discovery is extremely fast, as it is for HFT, parallel trading cannot be profitable so dealers will rationally widen spreads to protect themselves from adverse selection. We test this hypothesis by partitioning Hedge funds into those that do and do not engage in HFT. Consistent with our hypothesis, the dealers' adverse-selection response is especially aggressive for Hedge Funds engaged in HFT.

The model's evidence for volume discounts and strategic dealing support our hypothesis that dealers optimize intertemporally. Further support for this hypothesis comes from evidence that dealers smooth spreads over time within client relationships. The empirical analysis also provides support for our hypothesis that first-degree price discrimination generates substantial cross-client variation in execution costs. Indeed, it indicates that price discrimination accounts for 80% of such variation in FX dealer-client spot trades. The analysis also shows that market sophistication is the dominant source of cross-client variation in the price discrimination component of the markup.

Our qualitative conclusions are robust to using alternative measures of market sophistication, private information, and volatility; to restricting the sample to trades priced by a human; and to identifying clients by holding company rather than local subsidiary.

Our analysis highlights at least four features of dealer-client trades that cannot be studied with the most familiar spread decomposition models (e,g,, Glosten and Harris, 1988; Huang and Stoll, 1997; Madhavan et al., 1997). First, the execution costs has two parts that are determined separately, not one. Second, dealers engage in first-order price discrimination so client characteristics are important determinants of execution costs. Third, dealers smooth spreads over time within client relationships. Finally, information asymmetries influence execution costs in different directions and different amounts for different clients, rather than one direction and amount for all clients.

There will inevitably be variation across OTC markets in the magnitude of client price discrimination components and the relative importance of each client characteristic. We conjecture that adverse selection is more important in markets with limited interdealer trading, like the municipal bond market, because parallel trading among dealers is less likely to be profitable. We also conjecture that strategic dealing is more common in markets with limited low-latency trading. Future research could fruitfully test these conjectures.

REFERENCES

Abudy, Menachem, and Avi Wohl, 2018. Corporate bond trading on a limit order Book exchange. *Review of Finance* 22(4): 1413-1440.

Akram, Farooq, Dagfinn Rime, and Lucio Sarno, 2008. Arbitrage in the Foreign Exchange Market: Turning on the Microscope. *Journal of International Economics* 76: 237-253.

Anand, Amber, and Avanidhar 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, 2019. *Triennial Central Bank Survey, Foreign Exchange and Derivatives Market Activity in 2019*. B.I.S., Basel. https://www.bis.org/statistics/rpfx19_fx_annex.pdf

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.

Biais, Bruno, Thierry Foucault, François Salinié, 1998. Floors, dealer markets and limit order markets. *Journal of Financial Markets* 1: 253-284.

Biais, Bruno, and Richard Green, 2007. The microstructure of the bond market in the 20th century. Unpublished monograph, Tepper School of Business, and Toulouse University.

Bjønnes, Geir Høidal, and Dagfinn Rime, 2005. Dealer behavior and trading systems in foreign exchange markets. *Journal of Financial Economics* 75: 571–605.

Bollerslev, Tim, and Michael Melvin, 1994. Bid-ask spreads and volatility in the foreign exchange market. J*ournal of International Economics* 36: 355-372.

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.

Burdett, K., and K. L. Judd (1983). Equilibrium price dispersion. *Econometrica* 51(4): 955-969.

Chaboud, Alain P., Sergey V. Chernenko, Edward Howorka, Raj. S. Krishnasami Iyer, David Liu, and Jonathan H. Wright, 2008. The high-frequency effects of U.S. macroeconomic data releases on prices and trading activity in the global interdealer foreign exchange market. *Journal of the European Economic Association* 6(2-3): 589-596.

Chernenko, Sergey, and Adi Sunderam, 2014. Frictions in shadow banking: Evidence from the lending behavior of money market mutual funds. *Review of Financial Studies* 27(6): 1717-1750.

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.

Cocco, João. F., Francisco J. Gomes, and Nuno C. Martins (2009). Lending relationships in the interbank market. *Journal of Financial Intermediation* 18(1): 24-48.

Dhokakia, Utpal M. (6 July 2015). The risks of changing your prices too often. *Harvard Business Review*. 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, 2002. Order flow and exchange rate dynamics. *Journal of Political Economy* 110: 170–180.

Evans, Martin, and Richard K. Lyons, 2005. Do currency markets absorb news quickly? *Journal of International Money and Finance*.24(2): 197-217

Froot, Kenneth A., Tarun Ramadorai, 2005. Currency returns, intrinsic value, and institutional-investor flows. *Journal of Finance* 60(3): 1535-66.

Galanek, John 2010. The search for clarity: Exploring new frontiers in FX TCA. *Currency Investor Magazine*. http://e-FX.net/articles/jul-2010-the-search-for-clarity-exploring-new-frontiers-in-fx-tca.html. (Accessed 4/17/2020.)

Garbade, Kenneth, and William Silber, 1976. Price dispersion in the government securities market. *Journal of Political Economy* 84(4, Part 1): 721-740.

Garbade, Kenneth, and William Silber, 1979. Structural organization of secondary markets: Clearing frequency, dealer activity and liquidity risk. *Journal of Finance* 34(3): 577-593.

Glosten, Lawrence, 1994. Is the electronic open limit order market inevitable? *Journal of Finance* 49(4): 1127-1161.

Glosten, Lawrence, and Lawrence Harris, 1988. Estimating the components of the bid-ask spread. *Journal of Financial Economics* 21: 123-142.

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.

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.

Gtifa, Saida, and Naoufel Liouane, 2013. Bid-ask spread, order size, and volatility in the foreign exchange market: An empirical investigation. *Journal of Business Management and Economics* 4(12): 267-275.

Han, Song, and Kleopatra Nikolaou, 2016. Trading relationships in the OTC market for secured claims: Evidence from triparty repos. *Finance and Economics Discussion Series* 2016-064. Washington: Board of Governors of the Federal Reserve System.

Hansch, Oliver, Neuberger, Anthony, 1996. Strategic trading by market makers on the London Stock Exchange. The Institute for Financial Research, Birkbeck College. Working Papers.

Harris, Lawrence, 2003. *Trading and Exchanges* (Oxford University Press: New York, New York).

Harris, Lawrence, and Michael Piwowar, 2006. Secondary trading costs in the municipal bond market. *The Journal of Finance* 61(3): 1361-1397.

Hartmann, Philipp, 1999. Trading volumes and transaction costs in the foreign exchange market: Evidence from daily dollar-yen spot data. *Journal of Banking and Finance* 23: 801-824.

Hau, Harald, Peter Hoffmann, Sam Langfield, and Yannick Timmer (2021). Discriminatory pricing of over-the-counter derivatives. *Management Science*. https://doi.org/10.1287/mnsc.2020.3787.

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 Albert Menkveld, 2014. Price pressures. *Journal of Financial Economics* 114(3): 405-423.

Hendershott, Terrance, Dan Li, Dmitry Livdan, and Norman Schurhoff, 2020. Relationship trading in OTC markets. *The Journal of Finance* 75(2): 683-734.

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.

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

Lagos, Ricardo, and Guillame Rocheteau, 2009. Liquidity in asset markets with search frictions. *Econometrica* 77: 403-426.

Leach, J. Chris, and Ananth N. Madhavan, 1993, Price experimentation and security market structure, *Review of Financial Studies* 6: 375-404.

Lyons, Richard K. (1997). A simultaneous trade model of the foreign exchange hot potato. *Journal of International Economics* 42: 275-98.

Madhavan, Ananth N., 1992. Trading mechanisms in securities markets. *Journal of Finance* 47(2): 607-641.

Madhavan, Ananth N., and S. Smidt. 1993. An analysis of changes in specialist inventories and quotations. *Journal of Finance* 48: 1595–1628.

Madhavan, Ananth N., Matthew Richardson, and Mark Roomans, 1997. Why do security prices change? A transaction-level analysis of NYSE stocks. *Review of Financial Studies* 10(4): 1035 00 1064.

Matejka, Filip (2015). Rigid Pricing and Rationally Inattentive Consumer. *Journal of Economic Theory* (2015), 158(B): 656–678.

McGroarty, Frank, Owain ap Gwilym, and Stephen Thomas, 2007. The components of electronic inter-dealer spot FX bid-ask spreads. *Journal of Business Finance & Accounting* 34(9 & 10): 1635 - 1650.

McInish, Thomas, and R. Wood: 1992. An analysis of intraday patterns in bid/ask spreads for NYSE stocks, *The Journal of Finance* 47(2): 753-764.

Naik, Narayan Y., Anthony Neuberger, and S. Viswanathan, 1999. Trade disclosure regulation in markets with negotiated trades. *Review of Financial Studies* 12: 873–900.

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 Inte.rnational Money and Finance* 30: 1696–1718.

Reiss, Peter C., and Ingrid Werner, 1996. Transaction costs in dealer markets: Evidence from the London Stock Exchange. In *The Industrial Organization and Regulation of the Securities Industry*, Chicago: The University of Chicago Press, 125-175.

Reinganum, Jennifer, 1979. A simple model of price dispersion. *Journal of Political Economy* 87(4): 851-858.

Rime, Dagfinn, Andreas Schrimpf, and Olav Syrstad, 2022. Covered Interest Parity Arbitrage. Forthcoming, *Review of Financial Studies.*

Shiller, Benjamin R., 2020. Approximating purchase propensities and reservation prices from broad consumer tracking. *International Economic Review,* forthcoming.

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.

Viswanathan, S., and James J.D. Wang, 2002. Market architecture: limit-order books versus dealership markets. *Journal of Financial Markets* 5: 127-167.

Waldfogel, Joel, 2015. First degree price discrimination goes to school. *The Journal of Industrial Economics* 63(4): 569-597

Table 1: Descriptive statistics

Table shows summary statistics for markups measured in pips, where one pip is \$0.0001/€. Data include 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.

	All	Hedge Funds	Client Banks	Brokers	Real-money	MNC	SME
A. Basic descriptors							
Number clients	2,388	43	731	110	257	169	1,078
Number trades	257,241	6,624	73,730	170,668	1,185	1,261	3,773
Mean trades/client	108	154	101	1,552	5	7	4
Trade size (€ mn)							
Mean	0.49	1.21	0.88	0.29	2.09	2.53	0.22
St. dev	1.91	1.81	3.15	0.63	6.61	5.48	0.82
Max	184.0	56.0	184.0	44.2	100.0	50.0	25.0
B. Markups by trade (pips)							
Mean	0.44	0.03	0.25	0.01	2.53	3.04	22.74
St. Dev	3.96	0.48	2.52	0.49	11.60	10.81	18.57
Max (pips)	325.00	15.70	325.00	53.00	96.00	130.00	104.00
C. Client-specific average markups (pips)							
Mean (pips)	1.79	0.41	1.05	1.31	3.57	9.61	32.87
St. dev	21.21	1.69	9.62	6.54	21.24	17.37	18.59
Max	239.29	8.78	239.29	46.00	40.00	120.00	104.00

`

Table 2: Client Characteristics Relevant to Price Discrimination

Table shows summary statistics for regression variables. $TrVol_c$ is client-average daily trading volume. $CTV_c = Ln(TrVol_c) - Ln(TrVol_{SME})$. Sophistication is a vector comprising the share of a client's trading over four platform types: direct, single-bank platform, API-brokers, API-other clients. Single-bank platform is a price source on the client's desktop that connects directly to the dealing bank that provided these data. API is automated programming interface which allows compute-generated trade requests to go directly to the dealing bank's e-dealing desk. MRFQ is request for quote system, which allows one client to request near-simultaneous quotes from multiple dealing banks. Information is a client's average one-minute post-trade return in pips, where one pip is \$0.0001/€. *Ln*(*TrdSize*) is the log of the client's average trade size in € mns.* Data comprise all market-making dealer-client EUR-USD spot trades by a top-10 EUR-USD dealing bank from 2 January through 20 April, 2012.

	Hedge Funds	Client Banks	Brokers	Real Money	MNC	SME
A. *Trading volume*						
$TrVol_c$ in € mns						
Mean by client	2.73	1.31	6.56	0.14	0.28	0.01
Median	0.40	0.21	0.95	0.01	0.02	0.00
St. dev.	4.63	3.89	16.61	0.41	1.09	0.04
Max	17.37	52.40	135.00	4.12	12.24	0.98
CTV_c						
Mean by client	13.28	12.1	13.25	9.17	9.48	7.57
St. dev.	2.04	2.33	2.88	2.76	3.05	2.02
B. Sophistication: Shares of trading across platforms						
Direct trading						
Mean by trades	1.97	2.04	0.07	13.80	24.57	68.56
Mean by client	14.07	10.79	4.26	16.81	52.73	95.24
St. dev "	32.61	28.62	18.05	35.74	47.42	20.78
Single-bank platform						
Mean by trades	1.17	68.57	13.95	32.83	17.68	29.61
Mean by client	9.20	54.59	16.00	30.37	22.27	3.47
St. dev "	28.55	46.97	34.88	44.88	40.48	18.03
API						
Mean by trades	49.86	3.28	58.98	0.19	0.00	0.93
Mean by client	25.24	1.28	23.19	0.02	0.00	4.31
St. dev "	43.42	11.05	40.96	0.30	0.00	0.19
MRFQ						
Mean by trades	47.00	26.11	27.00	53.18	57.75	0.91
Mean by client	49.55	32.20	55.80	52.22	24.68	2.23
St. dev "	48.64	43.79	47.63	49.53	41.79	10.29
C. Information						
Mean by client	0.10	-0.03	0.07	-0.14	0.03	0.06
Median "	0.16	0.00	0.09	0.00	0.00	0.01
St. dev "	0.95	1.01	0.76	1.33	1.54	1.92
Mean by Trade						
1-Min	0.27	-0.01	0.04	0.05	-0.03	0.03
30-Min	0.44	-0.01	-0.02	-0.18	-0.23	0.05
D. Ln(TrdSize)[§]						
Mean by client	14.37	13.08	12.45	12.68	12.73	11.21
St. dev	1.44	1.78	1.54	2.47	2.56	1.82

* *TrdSize* is measured in EUR for the regressions.

Table 3: Determinants of FX Client Markups

Table reports coefficients from estimating Equation ((11), repeated below:

$$Markup = \beta CTV + \Sigma_j \delta_{jc} + Soph\,\delta + \Sigma_j \pi_j Info$$
$$+ (\kappa_1 + \lambda_1)HL + Time\,\kappa_2 + (\kappa_3 + \lambda_3)IBSprd + (\varphi_1 + \lambda_2)Sz + \varphi_2 DT + \varphi_3 Sz\,DT + \varepsilon.$$

Markup is the absolute difference between the dealing bank's price on trade *t* and the prevailing core price. *CTV*: client's log average daily trading volume. *Soph*: vector comprising the client's share of trades handled directly, over SBP, or over APIs (MRFQs are the excluded platform). *Info* is the average 1-minute post-trade return for client *c* in transaction *t*. *HL*: (log) high-low range over the previous hour. *Time*: indicators for trades during European trading, NYC trading, or overnight trading (Asian trading hours excluded). *IBSpr*: contemporaneous core spread. *Sz*: log of trade *t*'s absolute amount in euros. *DT*: indicator for direct trades. Column 2 excludes client-type indicators to clarify their relevance for client sophistication. Column 3 disaggregates Hedge Funds into those that engage in low-latency trading and others to clarify the dealers' choice between adverse selection and strategic dealing. Data include all client trades through a top-10 EUR-USD dealing bank during the first 68 trading days of 2012. Robust standard errors clustered by date. *, ** and *** indicate 5%, 1%, and 0.01% significance, respectively. All regressions have 257,241 observations.

	Baseline	**Exclude Client-type Indicators**	**Partitioned Hedge Funds**
Volume Discounts	-0.173***	-0.255***	-0.173**
Sophistication: Intercepts			
x Hedge Fund	2.281***		2.252***
x Cust. Bk	2.470***		2.440***
x Broker	2.972***		2.760***
x Real-money	3.687***		3.656***
x MNC	3.082***		3.050***
x SME	13.538***		13.508***
Constant		4.199***	
Sophistication: Platform shares			
Direct	4.771e-2***	7.358e-2***	4.799e-2***
SBP	0.123e-2***	0.433e-2***	0.150e-2***
API-Other	0.024e-2	0.300e-2***	0.053e-2*
API-Broker	-0.1214e-2***	-0.011e-2	-0.093e-2***
Information			
3 x Hedge Fund	0.980***	0.325***	
Low Latency	–	–	1.230***
Other	–	–	0.831***
5 x Client. Bk	0.230*	0.355***	0.231*
2x Broker	-0.559***	-0.692***	-0.560***
1 x Real-money	0.202	-0.465	0.202
4 x MNC	-0.717	-0.300	-0.717
6 x SME	0.162	0.179	0.162
Spread Stabilization			
Volatility	-0.070e-2**	-0.025e-2	-0.001**
Time: Europe	0.119***	0.168***	0.119***
NYC	0.073***	0.143***	0.073***
Overnight	-0.101***	-0.110***	-0.101***
InterBankSpread	0.740e-2	0.779e-2	0.074e-2
Other Controls for Operating and Inventory Costs			
TrdSz	0.591e-2*	-2.60e-2***	0.622e-2**
Direct Trade (*DT*)	35.603***	56.848***	35.603***
DT x TrdSz	-2.300***	-3.472***	-2.300***
Adj. R^2	0.571	0.521	0.571
N. Obs.	257,241	257,241	257,241

Table 4. Negative Markups: Further Evidence for Strategic Dealing

Table examines post-trade returns for the 253 trades with negative markup and compares it to trades trades with positive markups. Data include all market-making dealer-client EUR/USD trades by a top-10 dealing bank during the 68 trading days from 2 January through 20 April, 2012. ** indicates significances at the 5% level.

	Mean Size	**Client's Post-trade return (pips)**		
		1 min	**2 min**	**5 min**
Markup < 0	2.39	0.160	0.350**	0.400
Other trades by negative-markup clients	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 5: Variation in Markups and Major Components

Table presents standard deviations and coefficients of variation for the average of fitted markups, fitted price-discrimination components, and fitted "Other" components by client. Fitted values based on estimated coefficients from Equation (10), shown in

	All Clients	Profit-seeking	Utilitarian
$\widehat{Markup}_c$			
Standard deviation	21.2	9.1	21.2
Coefficient of variation	1.3	8.6	0.8
$\widehat{PD}_c$			
Standard deviation	7.7	1.5	6.3
Coefficient of variation	0.9	1.6	0.5
$\widehat{Other}_c$			
Standard deviation	5.6	1.3	5.6
Coefficient of variation	1.2	4.0	0.8

Table 6. Robustness Tests

Table reports coefficients from estimating Equation ((11), repeated below:

$$Markup = \beta CTV + \Sigma_j \delta_{jc} + Soph\,\delta + \Sigma_j \pi_j Info$$
$$+ (\kappa_1 + \lambda_1)HL + Time\,\kappa_2 + (\kappa_3 + \lambda_3)IBSprd + (\varphi_1 + \lambda_2)Sz + \varphi_2 DT + \varphi_3 Sz\,DT + \varepsilon.$$

Alt Sophistication: Captures platform sophistication with an indicator for the most sophisticated platform used by the client. 30-min $Info_c$ replaces 1-minute with 30-minute post trade returns in $Info_c$. Alternative Volatility: uses a realized volatility to capture volatility. Human Trades: sample is limited to trades priced by a person rather than an algorithm. No Time Dummies excludes time tummies. Holding Company identifies each client with its holding company. Data include all market-making EUR-USD client trades through a top-10 dealing bank during the first 68 trading days of 2012. Robust standard errors clustered by date. *, ** and *** indicate 5%, 1%, and 0.1% significance, respectively.

	Alternative Sophistication	**30-min *Info_c***	**Alternative Volatility**	**Human Trades**	**Country Risk**	**Holding Company**
Volume Discounts	-0.165***	-0.161***	-0.173***	-0.325	-0.171***	-0.196***
Sophistication: Intercepts						
x Hedge Fund	2.295***	2.260***	2.270***	21.505***	2.081***	2.727***
x Cust. Bk	2.208***	2.313***	2.457***	26.231***	2.233***	2.805***
x Broker	2.461***	2.594***	2.778***	24.483***	2.566***	3.169***
x Real-money	3.266***	3.571***	3.674***	22.657***	3.517***	3.940***
x MNC	2.975***	2.933***	3.074***	27.442***	2.943***	3.378***
x SME	12.102***	13.411***	13.518***	37.976***	13.426***	13.793***
Sophistication: Platform Shares (e-2)						
Direct	9.959***	4.788***	4.770***	5.174**	4.753***	4.706***
SBP	0.285***	0.111**	0.124**	-1.672	0.110***	0.062
API-Other	0.169***	0.006	0.024	-2.193	0.072***	0.001
API-Broker	0.085**	-0.151***	-0.121***	-3.460**	-0.011*	-0.164***
Information						
x Hedge Fund	0.182*	0.259***	0.973***	5.606**	0.944***	0.626***
x Cust. Bk	0.341***	0.017	0.232*	2.914**	0.271**	0.241*
x Broker	-0.577***	-0.073**	-0.561***	2.601**	-0.505***	-0.737***
x Real-money	-0.026	-0.031	0.201	0.115	0.200	0.612
x MNC	-0.903	-0.147	-0.728	1.543	-0.718	-0.811
x SME	0.120	-0.053	0.162	-0.900	0.163	0.167
Spread Stabilization						
Volatility (e-2)	-0.078*	-0.078*	-0.001	-0.029	-0.072***	-0.066*
Time: Europe	0.118***	0.118***	0.110***	3.076***	0.128***	0.122***
NYC	0.063***	0.063***	0.073***	1.776	0.073***	0.080***
Overnight	-0.115***	-0.115***	-0.092***	-3.604***	-0.095***	-0.084***
IBSprd	0.007	0.007	0.008	0.106	0.007*	0.007
Other Controls for Operating and Inventory Costs						
TrdSz	0.006*	0.006*	0.006*	-1.436**	0.007***	0.004
Direct Trade	35.575**	35.575***	35.610***	–	35.622***	34.643***
DT x TrdSz	-2.297***	-2.297***	-2.299***	–	-2.302***	-2.248***
Country Risk						
Level 2	–	–	–	–	-0.036***	–
Level 3	–	–	–	–	0.160***	–
Level 4	–	–	–	–	0.710***	–
Adj. R^2	0.580	0.571	0.571	0.266	0.571	0.571
N. Observations	257,241	257,241	257,212	3,066	257,241	257,241

.