Competition and Relationship Trading in OTC Markets*

Rainer Haselmann **Ingomar Krohn** **Philippe Mueller** **Maik Schmeling**

This version: June 2026

Abstract

We study how competition and relationships shape execution costs in over-the-counter foreign exchange markets under request-for-quote (RFQ) protocols. Using proprietary data from a major multi-dealer platform, we observe dealers contacted, quotes submitted, and executed prices. Execution costs are non-monotonic in the number of dealers contacted. Bilateral requests obtain tight prices, two-dealer requests are most expensive, and costs fall only under broader competition. This reflects both client selection and within-client execution choices. High-cost clients benefit from contacting more dealers, while low-cost, high-activity clients use bilateral and competitive execution selectively. Relationships and competition act as substitutes at the margin, and bilateral trading premiums are concentrated among corporate clients with narrow dealer networks and limited outside options. These findings challenge standard search-and-bargaining models in which contacting more dealers mechanically improves execution quality.

JEL classification: F31, G14, G15, G20

Keywords: foreign exchange, request for quote, OTC markets, execution costs, financial intermediation.

*First version: 18 November 2025. We are thankful for comments and suggestions from Philippe Cote, Xiang Fang, Aditya Kaul, Dmitry Livdan, and seminar and conference participants at CICF Hong Kong, St. Gallen Financial Economics Workshop, International Finance Society Meeting Gerzensee, Frontiers in Finance Conference Banff, Vienna Graduate School of Finance, UBC Winter Finance Conference, SKEMA Business School and WBS. The views expressed in this paper are solely those of the authors and do not necessarily reflect those of the Bank for International Settlements. All errors are our own. Authors: Rainer Haselmann, Goethe University Frankfurt and CEPR, London, Email: haselmann@econ.uni-frankfurt.de; Ingomar Krohn, Bank for International Settlements, Basel, Email: Ingomar.Krohn@bis.org; Philippe Mueller, Warwick Business School, Coventry, Email: Philippe.Mueller@wbs.ac.uk; Maik Schmeling, Goethe University Frankurt and CEPR, London, Email: schmeling@finance.uni-frankfurt.de.

Modern OTC markets are built around a simple trade-off. A client can request a price from a trusted dealer, relying on an established relationship, or ask several dealers to compete for the same order. Standard search-and-bargaining logic suggests that more quotes should improve execution quality by increasing competitive pressure. Yet exposing an order to more dealers may also reveal information, alter dealers' incentives, and weaken the advantages of bilateral trading. Whether competition lowers execution costs therefore depends not only on the number of dealers contacted, but also on the structure of the client-dealer relationship and the client's outside options. This raises a central question for OTC market design: when does competition lower execution costs, and when do bilateral relationships provide better pricing? Despite the importance of FX markets for global trade, investment, funding, and risk management, we know little about how clients choose between these trading channels, how dealers respond to being placed in competition, and how these choices shape transaction costs at the moment of execution.

We study these questions using proprietary data from one of the largest global multi-dealer platforms. The data allow us to reconstruct the full life cycle of each request: the client initiating the trade, the dealers contacted, the direction and notional size of the order, the full set of submitted quotes, and the quote ultimately accepted. We match each execution to a high-frequency benchmark mid-price, which allows us to measure transaction costs at the tick level. The sample covers approximately 14 million observations and nearly 1.5 million executed trades between January 2018 and July 2022 across 28 currency pairs. To our knowledge, this is the most detailed dataset yet used to study OTC FX trading, and it allows us to observe price formation, liquidity provision, and dealer competition at the moment of execution.

We document five main results. First, execution costs are non-monotonic in the number of dealers contacted: requests sent to one dealer obtain tight prices, requests sent to two dealers are the most expensive, and costs decline only once broader competition is introduced. Second, this pattern reflects both client selection and within-client variation. High-cost clients benefit monotonically from contacting more dealers, while low-cost, high-activity clients display the strongest deviation from standard competition effects. Third, relationships and competition act as substitutes: the marginal value of an additional quote declines with the depth of the requester-provider relationship. Fourth, bilateral and competitive trades form structurally separate execution channels rather than two points on a single continuum. Fifth, the cost of bilateral execution is concentrated among high-cost corporate clients with narrow dealer networks, whereas low-cost institutional clients use both channels without paying a bilateral premium.

The first result is that execution costs are not monotonically decreasing in competition. In standard search-and-bargaining models, contacting more dealers should weakly improve execution quality because each additional quote tightens the pricing constraint faced by the winning dealer. The data depart from this prediction as requests sent to a single dealer obtain relatively tight prices, while requests sent to exactly two dealers are substantially more expensive. Only beyond this point do costs decline with the number of dealers contacted. This pattern suggests that narrow competition may weaken the advantages of bilateral trading without yet generating enough competitive pressure to discipline dealer quotes.

Second, the non-monotonic pattern cannot be explained by a single selection story. In the full sample, requester fixed effects largely remove the premium associated with asking exactly two dealers, indicating that part of the pattern reflects selection across clients. However, this conclusion changes once requesters are split by their average execution costs. High-cost clients benefit from competition in the standard way: their costs fall as more dealers are contacted. Low-cost, high-activity clients, by contrast, continue to exhibit a pronounced cost increase when moving from one to two dealers. For these clients, the choice between bilateral and competitive execution appears to reflect a more deliberate trading strategy rather than simple differences in client quality.

Third, we find that relationships and competition substitute for one another at the margin. A deeper requester-provider relationship is associated with lower execution costs, and contacting more dealers also lowers costs once the request is competitive. However, the two channels are not additive. The marginal benefit of an additional quote is smaller when the requester has a deeper trading history with the executing provider. Conversely, the value of relationship depth is strongest when competition is limited. In this sense, a deep bilateral relationship and a broad competitive auction can deliver similar execution quality through different mechanisms.

Fourth, bilateral and competitive trades are structurally separate execution channels. Within a given requester-provider relationship, trades tend to be routed predominantly through one channel or the other, rather than being smoothly distributed between bilateral and competitive execution. The choice of channel is also persistent: a requester that recently traded bilaterally is substantially more likely to do so again on the next comparable trade. This persistence suggests that clients do not simply choose the number of dealers trade by trade in response to transitory conditions. Instead, they appear to maintain distinct dealer relationships for relationship-based execution and competitive auctions.

Finally, we document that the cost of bilateral trading is highly uneven across clients. High-cost requesters pay a substantial premium for bilateral execution relative to competitive trades within the same requester, currency pair, and month. These clients are more likely to be corporate requesters with narrow provider networks and concentrated trading relationships. Low-cost requesters, by contrast, pay no comparable bilateral premium. They are more often institutional clients with broad dealer networks and appear able to use bilateral and competitive execution selectively. The bilateral premium is therefore not a uniform feature of one-sided trading, but a relationship rent borne by a narrow segment of clients with limited outside options.

Taken together, the novel facts we uncover suggest that RFQ trading in modern FX markets does not conform to simple search-and-bargaining models in which additional dealers mechanically improve execution quality. Competition matters, but its effect depends on the structure of the client–dealer match and on the client's outside option to shop more widely. A bilateral request may obtain tight pricing when it preserves inventory and information advantages that are weakened once the trade is exposed to additional dealers. By contrast, narrow competition may weaken the benefits of a relationship without yet generating enough competitive pressure to discipline quotes. Only when a sufficiently broad set of dealers is contacted do competitive mechanisms dominate and execution costs decline in line with standard intuition. Execution costs therefore reflect an interaction between relationship capital, competitive search, inventory risk, information leakage, funding constraints, and client-specific bargaining power. The electronic RFQ protocols of our platform allow these channels to be observed almost in real time, making the FX RFQ setting an ideal environment in which to study the boundary between bilateral trading and electronic competition. Overall, our evidence calls for richer models of OTC trading that combine search, inventory management, funding frictions, information design, and long-term relationships on both the dealer and client side.

Literature Review. Our paper first contributes to the emerging empirical literature on request-for-quote (RFQ) trading in over-the-counter markets. Closest to our setting, Kargar, Lester, Plante, and Weill (2025) and Hendershott, Li, Livdan, Schürhoff, and Venkataraman (2025) study search costs and quote competition in corporate bond markets, showing how dealer responses shape execution quality and transaction costs. We extend this literature to the FX spot market, where trading is more frequent, liquidity is deeper, and the trading protocol allows clients to vary not only the number of dealers contacted but also the information revealed about trade direction. To our knowledge, ours is the first paper to analyse complete quote and execution data

for FX requests, allowing us to study dealer competition at the moment prices are formed.

A related theoretical literature studies how clients and dealers behave on multi-dealer platforms. Wang (2023) models trading on platforms with RFQ features, where dealers strategically choose whether and how aggressively to respond to requests. Yueshen and Zou (2022) show that OTC customers may optimally contact only a small number of dealers and submit fewer but larger orders in order to elicit liquidity provision. Our evidence speaks directly to these mechanisms. We observe both the number of dealers contacted and the full set of quotes submitted, which allows us to test whether additional dealer competition improves execution quality and how this effect depends on relationship depth.

Second, our paper relates to the broader literature on FX market microstructure. Prior research studies price discovery, liquidity provision, and execution quality in dealer-to-dealer and dealer-to-client segments of the FX market, including Lyons (1995) and Evans and Lyons (2002). These studies typically rely on executed trades or aggregate order-flow measures. By contrast, our data contain both executed and non-executed quotes within the same request. We can therefore observe the competitive process directly: which dealers are invited, how aggressively they quote, how dispersed their quotes are, and which quote the client ultimately accepts.

Third, we build on research on price discrimination and heterogeneous pricing across clients and dealers. Hau, Hoffmann, Langfield, and Timmer (2021) show that dealers adjust OTC derivative prices depending on client sophistication and clients' connections with other dealers, while Ranaldo and Somogyi (2021) and Menkhoff, Sarno, Schmeling, and Schrimpf (2016) document systematic differences in trading behaviour and execution outcomes across customer types, including corporates, asset managers, and other financial institutions. We extend this literature by linking client heterogeneity to RFQ-level competition. This allows us to show not only that different clients face different prices, but also that they use different combinations of bilateral relationships and multi-dealer competition.

More broadly, we contribute to theoretical and empirical work on OTC market structure, search, and bargaining (e.g., Duffie, Gârleanu, and Pedersen, 2005; Duffie, Dworczak, and Zhu, 2017). A central theme in this literature is that trading costs depend not only on the number of counterparties contacted, but also on search frictions, bargaining power, and the information revealed during the trading process. Hendershott and Madhavan (2015) and Riggs, Onur, Reiffen, and Zhu (2020) provide related empirical evidence on trading-protocol choice in corporate bond and swap markets, respectively, showing that clients trade off the benefits of electronic competition

against search frictions, relationship trading, and information leakage. Baldauf and Mollner (2024) formalise this trade-off by showing that greater counterparty competition can be costly because contacting more dealers increases information leakage and exposes the trader to front-running by losing dealers. Baldauf, Mollner, and Williams (2025) further show that privacy can improve off-market execution when intermediaries compete for privately informative order flow.

Our results support the view that platform design, search costs, information leakage, and inventory management jointly shape pricing outcomes. At the same time, the evidence we present does not fit neatly into existing models. In particular, the clients obtaining the lowest transaction costs do not rely exclusively on either relationships or competition. Instead, they use both channels selectively across trades. Our evidence suggests that relationship-based and competitive trading are substitutes at the margin but not mutually exclusive in practice: the value of an established dealer relationship depends on when, and how broadly, a client exposes order flow to competition.

I. Institutional Setting and Data

To the best of our knowledge, our data comprise the most comprehensive set of granular data for FX spot quotes and transactions between dealers and end-user clients. We first briefly review the institutional setting and the mechanisms using request-for-quote terminology. Second, we provide an overview of the data available to us. Third, we highlight relevant summary statistics. Finally, we define different client groups based on how they use the request-for-quote options available to them.

I.A. Institutional Setting

Foreign exchange spot trading is segmented between an inter-dealer market, where dealers manage inventory among themselves, and a customer-dealer market, where end-user clients trade with one or more dealers. Our data come from a major electronic multi-dealer platform in the customer-dealer segment, where clients solicit executable prices from liquidity providers and decide whether and at which quote to execute.[1]

The platform supports a *request-for-stream* (RFS) protocol, whereby a client requests a continuous stream of executable prices from one or more providers, and each provider updates their stream in response to changing market conditions. RFS differs from the more widely studied

[1]This contrasts with the inter-dealer platforms typically studied in the FX microstructure literature—such as LSEG Matching and Electronic Brokerage Systems (EBS)—in which dealers trade anonymously with one another.

request-for-quote (RFQ) protocol, under which the client sends a discrete request and each provider responds with a single quote that remains executable for a short time window. From the client's perspective, the decision-relevant information at the moment of execution is the same as in an RFQ—a discrete set of competing quotes, one of which is selected—and we therefore use the more familiar RFQ terminology throughout the paper. For each client request we observe a snapshot of all quotes provided at the moment of execution. That is, we observe not only the executed price offered by a particular dealer that "wins" the trade but also the dealer-price pairs for all quotes not selected by the client at the time of execution. We refer to the client as the *requester*, to the dealers as the *providers*, and treat each executed trade together with its associated set of provider quotes as a "request."

Figure 1 illustrates a hypothetical example of a client request. The requester ("Customer X") seeks to purchase 7.5 million euro spot against the U.S. dollar, soliciting quotes from multiple providers. Each provider responds with a quote, represented by a line connecting the requester to the provider node in the figure. By construction, only one quote results in execution (thick red line), while all others are non-executed responses (gray lines).

For each request, we observe the requester's anonymized identifier and a classifier that distinguishes corporate from institutional clients. Moreover, we observe the direction of trade (buy vs. sell) as well as the size of the request. We distinguish between two request types: *one-sided* requests, in which the requester solicits a quote on one side of the market (a request to buy or to sell), and a *request for market* (RFM), whereby the requester solicits a two-way quote, thus hiding the intended trade direction from the provider. A separate flag indicates whether the request is *competitive* (COMP), with the requester soliciting quotes from one or more providers in a price-discovery process, or *non-competitive* (NONC), in which the requester directs the trade to a specific provider under a pre-agreed pricing arrangement. On the dealer side we also have anonymized identifiers, the country in which the provider is located, along with the quoted price, and the executed traded price with the exact high-frequency timestamp. This level of granularity enables us to reconstruct the complete life cycle of each RFQ, from initial request through competitive quoting to final execution. Even though provider and requester ID are re-assigned twice during our sample period, the data allow for detailed measurement of execution quality, assessment of price improvement, and analysis of liquidity provision and dealer competition at high frequency.

[INSERT FIGURE 1 HERE]

I.B. Data

RFQ Data The dataset covers the full universe of RFQs submitted to the platform, containing both quote and trade messages, along with precise time stamps for the moment of eventual execution.

Our sample spans four and a half years from January 2018 to July 2022 and covers 14 currencies traded vis-à-vis the U.S. dollar or the euro: the Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), Danish krone (DKK), British pound sterling (GBP), Hong Kong dollar (HKD), Hungarian forint (HUF), Japanese yen (JPY), Mexican peso (MXN), Norwegian krone (NOK), New Zealand dollar (NZD), Polish zloty (PLN), Swedish krona (SEK), and South African rand (ZAR). Overall, the sample comprises 28 currency pairs in total: 15 denominated in U.S. dollars (the 14 currencies above plus the euro) and 13 denominated in euros (all of the 14 above except the New Zealand dollar).

Reference Prices. We augment the RFQ dataset with high-frequency benchmark prices from LSEG Tick History, a widely-used platform that provides indicative dealer-to-client streaming quotes. For each trade i, we measure the execution cost as the signed deviation between the trade price and the nearest available mid-quote from LSEG Tick History. Specifically, we match the timestamp of each trade to the closest mid-quote available before the executed transaction, ensuring precise time alignment between trade and benchmark prices. Using the benchmark prices, we define transaction costs (TC) in basis points as

$$\text{TC}_i \equiv b_i \cdot \frac{q_{i,p_i^*} - m_i}{m_i} \times 10{,}000 \tag{1}$$

where m_i is the prevailing mid at request i, q_{i,p_i^*} is the price quoted by the executing provider p_i^*, and $b_i \in \{+1, -1\}$ equals $+1$ for buy orders and -1 for sell orders. To mitigate the influence of data errors and extreme price movements, we winsorize transaction cost observations at the 1st and 99th percentiles. Furthermore, we discard transactions if we cannot observe a benchmark quote in the 60 seconds before the time stamp of the executed transaction, thus ensuring that the benchmark reflects prevailing market conditions while remaining robust to spurious price spikes and stale quotes.[2]

In addition to the transaction cost for the executed trade, we define the *quote cost* for each

[2]In the Online Appendix we show how the sample size is affected by using a 1-minute cut-off, compared with more stringent or more relaxed matching windows.

direction-aligned non-executed quote (i, p) in request i from provider p:

$$\mathrm{QC}_{i,p} \equiv b_i \cdot \frac{q_{i,p} - m_i}{m_i} \times 10{,}000 \tag{2}$$

where $q_{i,p}$ is provider p's quoted price and m_i and b_i are as defined above. Quote costs are signed such that positive values mean the requester pays a premium relative to the mid. By construction, $TC_i = QC_{i,p_i^*}$—the transaction cost is the quote cost evaluated at the executing provider.

I.C. Summary Statistics

Table I documents the richness and granularity of our RFQ database. Across all currency pairs, the dataset comprises 14.12 million observations, of which 1.41 million correspond to executed trades that can be linked to nearly 12.70 million associated quotes. These quotes are submitted by 959 distinct requesters from 218 providers, based in 48 countries, resulting in more than 6,810 unique dealer–client relationships. The average notional value per trade is 2.46 million U.S. dollars.

[INSERT TABLE I HERE]

Across currencies, we observe the expected heterogeneity consistent with global FX market liquidity patterns. The EUR/USD pair is the most actively traded, accounting for 5.58 million observations and featuring an average trade size of 3.80 million U.S. dollar. This is followed by other major pairs such as GBP/USD, CAD/USD, CHF/USD, and AUD/USD, each representing significant shares of total activity. By contrast, emerging market (EM) currencies are traded less frequently, reflecting their lower global liquidity and more segmented dealer networks. Nonetheless, even within EM pairs, the dataset reveals rich client–dealer interaction structures.

The database also includes parallel coverage of USD- and EUR-denominated pairs, often capturing the same underlying currencies under different quotation conventions—for example, USD/CHF and EUR/CHF, or USD/SEK and EUR/SEK. Consistent with the U.S. dollar's global role as the primary vehicle currency, trading activity and notional volumes are generally higher for USD-quoted pairs. However, several EUR-based crosses—notably GBP/EUR and CHF/EUR—exhibit active two-way flows and broad dealer participation, indicating that liquidity is not confined to U.S. dollar markets alone.

One of the novelties of the database is that it provides the first evidence on the full RFQ life cycle in FX spot markets. Figure 2 illustrates these dynamics. The hashed grey bar shows that a majority of trades, i.e., 77% of requests, refer to non-competitive quotes (NONC). NONC trades

are relationship-based or execution-only requests, reflecting transactions where a client submits an RFQ to a single dealer, often under pre-agreed pricing arrangements. In such cases, the platform functions primarily as an electronic execution and settlement channel rather than a venue for price discovery or dealer competition. Distinguishing these trades from competitive RFQs is crucial, as they represent the non-price-discovery segment of the market and provide a benchmark for assessing how competition affects execution quality in the remaining sample.

[INSERT FIGURE 2 HERE]

The remaining 23% of the sample correspond to competitive trades (COMP), where clients request quotes from one or more dealers simultaneously. The light blue bars depict the distribution of these competitive requests, with percentages shown above each bar representing shares within the COMP segment.

Within competitive requests we observe substantial heterogeneity in client behaviour. Roughly a third of competitive requests are still directed to only one dealer at a time, suggesting limited dealer competition for at least a subset of the "competitive" quotes.[3] However, the remaining requests are directed to two or more dealers with the mode being observed in the 6 to 10 dealer range (this category accounts for roughly one-fifth of competitive quotes). That is, a large fraction of quote requests involve substantial multi-dealer engagement.

I.D. Client Groups

We classify each requester's trading into one of three groups based on how they use competition within a given currency pair. For each (requester, currency pair, subperiod) combination we compute n_p, the number of unique providers the requester trades with, and $\min D$ and $\max D$, the minimum and maximum number of providers they invite across their requests in that pair. A requester is classified *Single RFQ* in a given pair if they invite only one provider per request ($\max D = 1$) but use multiple providers overall ($n_p \geq 2$); *Multi RFQ* if they always invite multiple providers ($\min D \geq 2$); and *Mixed RFQ* if they switch between modes, inviting one provider on some requests and multiple on others ($\min D = 1,\ \max D \geq 2$).

[INSERT TABLE II HERE]

[3] Throughout the empirical analysis, D denotes the number of direction-aligned executable quotes observed for a request. We refer to $D = 1$ requests as bilateral executions and to $D \geq 2$ requests as competitive executions, unless otherwise stated. This distinction is separate from the platform's COMP/NONC flag, which identifies whether the request is part of a price-discovery process or a pre-arranged execution workflow.

Table II reports the size, composition, and execution-cost profile of each group. The working sample contains 14,391 distinct (requester, currency pair, subperiod) classifications: 928 in Single RFQ, 6,767 in Mixed RFQ, and 6,696 in Multi RFQ. Mixed-RFQ behavior is by far the most common: 2,044 requesters appear at least once in this group, generating 69.8% of trades and 70.7% of total notional (2.46 trillion U.S. dollars). Multi-RFQ classifications account for 17.0% of trades and 22.1% of notional, while Single-RFQ accounts for the remaining 13.2% of trades and 7.1% of notional. All three groups are active in all 28 currency pairs of the working sample.

The three groups differ markedly in client composition. Single RFQ requesters are predominantly institutional (76.3% Inst, 23.7% Corp), consistent with the use of single-dealer-directed trades as the execution channel for long-standing institutional relationships. Multi RFQ requesters are predominantly corporate (77.0% Corp, 23.0% Inst), consistent with corporate treasury desks routing flow through competitive auctions to satisfy best-execution policies. Mixed RFQ requesters sit between the two (60.8% Corp, 39.2% Inst), reflecting clients who deploy both competition modes selectively across their trading. Group membership therefore reflects both behavior and client type, with Mixed RFQ as the segment most balanced across institutional categories.

Average transaction costs are economically modest but show a clear ranking across groups: 0.69 bps for Single RFQ, 0.50 bps for Mixed RFQ, and 0.51 bps for Multi RFQ. Single-RFQ trades pay roughly 0.2 bps more than the other two groups—a gap consistent with a relationship premium paid for direct execution with a designated dealer. Median trade sizes echo this pattern: Single RFQ trades are the largest at 0.77 million U.S. dollars, followed by Mixed RFQ at 0.67 million U.S. dollars and Multi RFQ at 0.55 million U.S. dollars. The use of two-way requests, in which the requester hides the trade direction by asking for both bid and ask quotes, differs sharply across groups: Single RFQ requesters rely on RFM in only 9.4% of trades—there is little need to hide direction from a dedicated dealer—while Mixed and Multi RFQ requesters use RFM in 41.0% and 44.4% of trades, respectively.

Within each requester's trading book, the choice of competition mode tends to be applied broadly. Multi RFQ requesters are classified as Multi RFQ in 3.62 currency pairs on average per subperiod, Mixed RFQ in 3.31 pairs, and Single RFQ in 2.42 pairs—suggesting that competition behavior is largely a stable trait rather than an ad-hoc choice driven by individual currency-pair conditions. Indeed, when we re-classify requesters at the cross-currency level—aggregating n_p, $\min(Q)$, and $\max(Q)$ across all pairs the requester trades within a subperiod—we find that pure-strategy behavior is perfectly absorbing: 100% of requesters classified as Single RFQ across all

pairs are also Single RFQ in every individual pair, and 100% of cross-pair Multi RFQ requesters are also Multi RFQ in every pair. Only Mixed RFQ behavior is heterogeneous across pairs, with cross-pair Mixed requesters splitting 58% Mixed, 38% Multi, and 4% Single at the pair level.

II. Competition and Transaction Costs

In this section we document that transaction costs are not monotonically related to the number of competing quotes. While costs are generally decreasing in the degree of competition, trades where only one dealer is approached for a quote come with lower transaction costs compared to trades where the client asks two dealers for competing quotes. This pattern is robust across currency pairs, trade sizes, and market conditions, and that it is not entirely explained by client selection.

II.A. Non-Monotonic Transaction Costs Pattern

In standard models of OTC markets, transaction costs are weakly decreasing in the number of competing quotes: each additional quote tightens the winning provider's pricing constraint and the client extracts a larger share of the gains from trade. The empirical literature on dealer markets largely corroborates this prediction in corporate bonds and some other asset classes. We document a deviation from this prediction in the FX customer-dealer market at the request-for-quote level.

[INSERT FIGURE 3 HERE]

Figure 3 plots the average transaction cost by the number of dealers approached per request, D, across our full sample. Overall, the pattern is non-monotonic. Average transaction costs decline as two or more dealers are approached—from 1.32 bps at $D = 2$ to 0.62 bps, 0.37 bps and 0.29 bps for the $D = 3--4$, $D = 5--10$, and $D = 11+$ bins, respectively. At the same time, transaction costs when only one dealer is asked for a quote are a mere 0.66 bps on average, implying that costs roughly double as clients move from $D = 1$ to $D = 2$.

Note that the spike is indeed concentrated at $D = 2$. While the next category contains all transactions where three or four dealers are approached, costs at $D = 3$ alone are down to 0.74 bps, more than 40% below the costs for the $D = 2$ bin. Moreover, the pattern survives at the median, suggesting that the spike is not driven by the upper tail of the cost distribution but is a property of the central mass. Median transaction costs increase from 0.39 bps to 0.93 bps as clients move from requesting one to requesting two quotes while they are again monotonically decreasing with 0.43 bps, 0.27 bps, and 0.18 bps, respectively, for the remaining three categories.

Pairwise Welch tests reject monotonicity at very high levels of significance: $TC(D = 2) > TC(D = 1)$ at $t = 104.2$ and $TC(D = 2) > TC(D = 3\text{–}4)$ at $t = 106.9$, both with p-values indistinguishable from zero. Any monotone $TC(D)$ would require the $(D = 1, D = 2, D = 3\text{–}4)$ triple to be ordered in one of two ways; both orderings are rejected. Over the declining region $D \geq 2$, the Spearman rank correlation between D and TC is -0.19 (p indistinguishable from zero), confirming that the break in monotonicity is specifically the $D = 1 \to D = 2$ inversion and that the post-spike decline is itself monotone.

[INSERT TABLE III HERE]

Table III establishes that the pattern is not an artefact of currency-pair composition. The hump appears in every subsample we examine: EURUSD on its own ($N = 354{,}800$), the 15 U.S. dollar-base pairs, the 13 euro-base pairs, the 9 G10 pairs, and the 19 EM/other pairs. In every subsample, the mean transaction cost at $D = 2$ exceeds the means at $D = 1$ and at $D = 3\text{–}4$, and the gap between $D = 2$ and $D = 5\text{–}10$ ranges in magnitude from 0.61 bps (EURUSD) to 1.06 bps (G10), significant at the 1% level in every subsample. The pattern is also robust to splitting the sample by notional tercile—the $D = 2$ vs. $D = 5\text{–}10$ gap is 1.17 bps, 0.85 bps and 0.66 for the small, medium and large tercile—and to splitting by sub-period—the gap is 0.84 bps, 1.10 bps and 0.87 bps during the pre-COVID (2018-01 to 2020-02), COVID (2020-03 to 2021-12), and post-COVID (2022-01 to 2022-07) period, again all highly statistically significant.

[INSERT FIGURES 4 AND 5 HERE]

Before we proceed, we rule out two specific narrow explanations. First, the hump cannot be attributed to small or unsophisticated clients. Figure 4 illustrates that restricting our sample to the top decile of requesters by total volume traded over the sample period leaves the pattern intact. That is, the $D = 2$ hump persists even among the largest, most active requesters—those most likely to be sophisticated and to demand best execution. Second, the hump is not an artefact of which quote the requester chooses to execute. Within each request, we rank all direction-aligned quotes from lowest to highest implied transaction cost. The 1st-best is the cheapest offer made to the requester—by construction this is (almost always) the executed ("traded") quote—the 2nd-best is the next-cheapest offer that is made on the same request but not selected, and the 3rd-, 5th-, and 11th-best are defined analogously provided the client asked enough providers for a quote. Figure 5 plots the mean implied cost of the k-th best quote by D bin. The hump at $D = 2$ appears

at every rank where it can mechanically occur—for the executed quote, for the 2nd-best (the only other offer at $D = 2$), and for the 3rd-, 5th-, and 11th-best quotes within the larger-D bins (and where those ranks exist). Providers as a group price more conservatively at $D = 2$ and the pattern occurs over the entire quote distribution, not just in the selection by the requesters.

II.B. Client Selection

There are broadly two possible explanations for the hump-shaped pattern. First, it may reflect a selection effect at the client level whereby some clients that can obtain very low transaction costs are more likely to request a quote from a single provider only and clients that pay higher transaction costs are more inclined to ask quotes from two providers. Second, the effect may occur for trades within clients or within certain client categories and may be driven by differences across the trades themselves such as notional size, market conditions, traded currency, time of trade or other salient characteristics. In this section we show that in our full sample, both explanations are likely relevant.

[INSERT FIGURE 6 HERE]

In Figure 6 we first average the transaction costs for every trade across each requester within every D-bin. Second, we average within every bin across the different requesters that are active in the bin. The resulting average transaction costs are now monotonically decreasing as we move across the bins. Overall, the sample used is the same as in Figure 3, the only difference is whether averages are taken across all trades within a bin or across requester-averages per bin. The disappearance of the hump is therefore consistent with a selection effect.

To test this idea more formally, we run the following regression:

$$TC_i = \sum_{k \in \{2,\ 3\text{-}4,\ 5\text{-}10,\ 11+\}} \delta_k \, \mathbf{1}\{D_i = k\} + \gamma' X_i + \text{FE} + \varepsilon_i, \tag{3}$$

where i indexes a trade, $\mathbf{1}\{D_i = k\}$ is an indicator equal to one if request i falls in D-bin k (with $D = 1$ as the omitted category), X_i collects the log notional trade amount and the prevailing bid-ask spread, and FE denotes the fixed effects added progressively across columns: currency-pair (α_c), provider (α_p), requester (α_r), and requester $\times$ provider ($\alpha_{r,p}$). Standard errors are two-way clustered by requester and provider.

[INSERT TABLE IV HERE]

Table IV reports the results. Specification (1) reproduces the unconditional means by D-bin from Figure 3 in regression form and serves as the benchmark against which the fixed-effects buildup is read. The remaining columns take $D = 1$ as the omitted base, so the coefficients $\hat{\delta}_k$ measure the conditional difference between bin k and $D = 1$. Specification (2) adds currency pair fixed effects together with controls for log notional and the bid-ask spread; the $D = 2$ coefficient that measures the premium for asking exactly two providers falls only marginally, from the original difference measured as 0.66 bps to $\hat{\delta}_2 = 0.59$ (with a t-statistic of 4.9). Trade-level composition explains very little of the premium.

Specification (4) adds provider fixed effects on top of pair fixed effects and controls. The $D = 2$ coefficient falls further to 0.41 (t-statistic of 4.9), i.e., the provider identity accounts for about one-third of the unconditional premium but leaves the rest unaffected. Provider fixed effects absorb only the average level of prices charged by each individual dealer—the fact that some providers consistently quote tighter or wider than others across their entire book—but they do not absorb how a given provider's pricing varies across D-bins within their own book. The 0.41 bps residual therefore rules out a pure provider-mix explanation in which $D = 2$ trades simply flow to systematically expensive providers; the premium persists within every individual provider's customer base.

The key result appears in specification (3). Replacing provider fixed effects with requester fixed effects reduces $\hat{\delta}_2$ to 0.04 bps, economically and statistically indistinguishable from zero. The coefficients on the higher D-bins remain significantly negative for the three additional bins. In specification (5) we add provider fixed effects to the specification from specification (3) to obtain a virtually identical result. Requester fixed effects absorb the average transaction cost faced by each individual requester. Their inclusion drives the $D = 2$ premium to zero while leaving the $D \geq 3$ coefficients significantly negative.

II.C. Client-Group Selection

The results thus far suggest that the $D = 2$ premium that results in the hump-shaped pattern in transaction costs across D-bins is due to a selection effect while competition across providers is responsible for the monotonically decreasing transaction cost levels in the bins $D \geq 3$. However, average transaction costs vary tremendously across different requesters and, in fact, the cross-sectional dispersion across different client groups is much more severe than the variation across D-bins. The average transaction cost for a client in the highest cost tercile is almost 7 times more

than the average transaction cost for a client in the lowest cost tercile. Thus, while requester fixed effects absorb the average level of transaction costs, the remaining variation within a requester and across D-bins is still very high for high transaction costs clients. This may mask more subtle variation across D-bins for low transaction cost clients.

Thus, we repeat some of the regression analysis from Table IV in Table V split across client terciles based on overall average transaction costs. In Panel T1 we revisit specifications (1), (3) and (4) from Table IV for the 320 requesters in the lowest transaction cost tercile. Panels T2 and T3 repeat the same analysis for the second and third group of requesters in our sample. While all terciles have essentially the same number of requesters, they vary a lot across other dimensions. The average transaction costs per tercile are 0.29 bps, 0.62 bps and 1.95 bps, respectively. While the ordering is by construction, the spread between the high and the low transaction cost decile is very high. Moreover, there is also a clear ordering in terms of the overall number of trades with the low transaction cost tercile accounting for over 955 thousand trades compared to 320 thousand for T2 and a mere 137 thousand for T3. That is, the requesters in T1 do not only enjoy low costs but they also trade a lot more.

[INSERT TABLE V HERE]

Interestingly, the results from Table IV do not unambiguously carry over when the sample is split according to transaction costs. The pattern is only repeated for the middle tercile. Specification (1) in Panel T2 shows the hump-shaped pattern as we move from $D = 1$ to $D = 2$ and further. As in Table IV, adding requester fixed effects removes the hump in the coefficient estimates (specification (2) in Panel T2). The high transaction costs requesters on the other hand do not even exhibit the original hump-shaped pattern in the D-bins, even without any fixed effects (see specification (1) in Panel T3). When adding requester fixed effects, the decrease in costs by D-bin is more accentuated but remains monotonic.

Finally, Panel T1 delivers yet another result: specification (1) in Panel T1 displays a very pronounced hump shape with average transaction costs doubling as we move from $D = 1$ to $D = 2$ and reverting completely by the $D = 5 - -10$ bin. Adding requester fixed effects in specification (2) attenuates the spike at $D = 2$ a little bit but the increase is still statistically significant. At the same time, it also accentuates the decrease in transaction costs when moving from $D = 2$ to $D = 11+$ with the coefficient of the latter being statistically significantly negative and almost tripling compared to the same coefficient in specification (1).

Taken together, the results in Table V suggest that the increase in cost when moving from asking one dealer to asking two dealers for a quote is not due to client selection, at least not in the subset of clients that can obtain very low transaction costs and that trade disproportionately more often. While it seems that high transaction cost clients unambiguously benefit from quote competition at the provider level, a more nuanced channel must be at play for the big, active and sophisticated clients that obtain the best conditions overall in the market. These are clients that have a wide set of options when it comes to asking for quotes and they make use of all of them in an attempt to optimise their trading. This sometimes entails asking only one provider for a quote and it sometimes entails asking a large set of providers and it implies that these clients are able to treat relationships and competition as substitutes.

III. Competition and Relationships Are Substitutes

Section II establishes that competition generally lowers transaction costs within a given requester. However, the results also point to the potential importance of relationships in order to obtain good pricing, especially for clients that trade a lot and overall enjoy low transaction costs. In this section we further explore how a requester's history of prior trading with the executing provider may affect costs and how these two forces may interact. If they reinforce each other, a requester who maintains relationships with several dealers extracts more by putting them in competition. If they substitute, a deep relationship already does part of what competition would do, and adding competing by including an additional quote may add very little cost benefit. Understanding these dynamics is important because it may guide how requesters route their flow. Under substitution, they face a real trade-off between bilateral and competitive execution. The substitution hypothesis predicts that the marginal effect of competition on transaction costs weakens as the requester–provider relationship deepens, and we test it directly.

III.A. The Substitution in Levels

We begin by examining the substitution non-parametrically. Let $d_{r,p,c,t}$ denote the cumulative number of executed trades between requester r and provider p in currency pair c prior to time t, constructed without look-ahead bias and computed pair by pair, and let $D_{r,p,c,t}$ denote the number of direction-aligned quotes received for the focal request. Table VI presents mean transaction costs by relationship-depth bin (rows) and D-bin (columns).

[INSERT TABLE VI HERE]

The substitution reveals itself in two directions. Moving down within a column shows the depth gradient at each competition level. At $D = 2$, transaction costs fall from 2.64 bps in fresh relationships ($d = 0$) to 1.53 bps in deep relationships ($d \geq 100$), a relationship-driven decline of 1.11 bps. At $D = 11+$, the corresponding decline is smaller with only 0.70 bps (1.53 to 0.83). The depth gradient is roughly 1.6 times steeper when competition is scarce than when it is abundant. Reading across rows produces the mirror image. For fresh relationships ($d = 0$), the competition gradient from $D = 2$ to $D = 11+$ is 1.11 bps, but only 0.70 bps for deep relationships ($d \geq 100$).

The diagonal highlights the substitution effect most starkly. Transaction costs reach the same 1.53 bps at (deep, $D = 2$) and at (fresh, $D = 11+$). Either route delivers identical execution quality, whether through a deep relationship or through many competing quotes. Combining both routes at (deep, $D = 11+$) saves only additional 0.70 bps, far less than the 1.11 bps each route delivers individually. Each channel substitutes for the other, and their effects are not additive.

III.B. Regression Evidence

To investigate the substitution effect more formally we estimate the following regression:

$$TC_{r,p,c,t} = \beta_1 \log D_{r,p,c,t} + \beta_2 \log(1 + d_{r,p,c,t}) + \beta_3 \log D_{r,p,c,t} \times \log(1 + d_{r,p,c,t}) + \gamma' X_{r,p,c,t} + \text{FE} + \varepsilon_{r,p,c,t}, \tag{4}$$

where $X_{r,p,c,t}$ collects log notional and the prevailing bid-ask spread, and FE denotes the fixed effects added progressively across columns: currency-pair (α_c), requester (α_r), and provider (α_p). Standard errors are two-way clustered by requester and provider. The coefficient of interest is β_3. A positive β_3 indicates substitution: the marginal effect of competition shrinks in absolute value as the relationship deepens.

[INSERT TABLE VII HERE]

Table VII reports the results. Specification (1) estimates the two main effects without the interaction, including only pair fixed effects: $\hat{\beta}_1 = -0.21$ ($t = -8.2$) and $\hat{\beta}_2 = -0.06$ ($t = -8.3$). That is, each channel reduces transaction costs in isolation.

We add the interaction in specification (2). The substitution coefficient is statistically highly significant and estimated as $\hat{\beta}_3 = 0.04$, while the main effects strengthen to $\hat{\beta}_1 = -0.42$ and $\hat{\beta}_2 = -0.12$. The implied marginal effect of $\log D$ is $-0.42 + 0.04 \log(1 + d)$. That means that for

a fresh relationship ($d = 0$) the marginal value of competition is -0.42 bps per log-unit, while at a typical established relationship ($d = 100$, $\log(1 + d) \approx 4.6$) it falls in absolute term to -0.21 bps, a reduction of fifty percent!

In specifications (3) and (4) we progressively absorb time-invariant requester and provider heterogeneity. With pair, requester, and trade-level controls (specification (3)), $\hat{\beta}_3$ is 0.02 (and statistically highly significant), while adding provider fixed effects in specification (4) leaves the estimate unchanged. That is, about half of the cross-sectional substitution reflects time-invariant differences across requesters and providers, but the rest operates within these dimensions: even after fixing the identity of the requester and the provider, the marginal value of competition declines with the depth of the (requester, provider) relationship.

Specifications (5) and (6) repeat the analysis with alternative measures of relationship intensity. In specification (5) we replace $\log(1 + d)$ with $\log b_{r,t}$, the cumulative number of distinct providers requester r has traded with to date—a requester-level measure of breadth in counterparty selection. The interaction $\log D \times \log b$ is $+0.16$ (with a t-statistic of 4.1), i.e., high-breadth requesters, who routinely shop across many providers, extract less from any single competitive auction. Specification (6) replaces depth with a binary loyalty indicator equal to one if the executing provider is the requester's most-used provider on the pair. Loyal trades pay a 1.35 bps premium over non-loyal trades—the rent extracted in established bilateral relationships—and competition erodes this rent at the rate of -0.39 bps per log-unit of D. The three operationalisations of relationship intensity—cumulative depth, breadth of counterparties, and binary loyalty—all point to the same conclusion: competition and relationships substitute at the margin.

III.C. Within-Request Identification

Unlike transaction-only data typical of OTC studies, our RFQ data lets us estimate the premium by comparing providers within the same request. For example, for a given request we can compare a provider with 100 prior trades to a provider with no prior trades, both quoting at the same moment. A request fixed effect then absorbs every request-level confound, including currency pair, notional, time of day, market state, and the requester's general pricing level. This identification strategy addresses a residual concern with the cross-request regression in Table VII, where a more sophisticated requester might solicit more quotes *and* trade more often with the same dealer, thus generating spurious comovement between D and d.

For each competitive request ($D \geq 2$), we estimate

$$TC_{i,p} = \beta \log(1 + d_{i,p}) + \gamma' X_{i,p} + \alpha_i + \varepsilon_{i,p}, \tag{5}$$

where i indexes a request, p indexes the providers quoting on request i, $d_{i,p}$ is the cumulative pair-specific relationship depth between the requester on request i and provider p as of the request time, and α_i is a request fixed effect.

[INSERT TABLE VIII HERE]

Table VIII reports the depth coefficient $\hat{\beta}$ split by D-bin. In Panel A (request fixed effects only), the depth premium declines monotonically as D rises, with $\hat{\beta} = -0.38$ at $D = 2$, -0.24 at $D = 3$–4, -0.21 at $D = 5$–10, and -0.17 at $D = 11+$, all significant at the 1% level. The within-request pattern mirrors the non-parametric heatmap in Section III.A. Relationships price more aggressively when competition is scarce.

Panel B adds pair and provider fixed effects together with the trade-level controls, absorbing each provider's average pricing level on each currency pair. The coefficients fall to roughly half their Panel A magnitudes (-0.13, -0.17, -0.08, -0.08 across D-bins) but remain significant at the 1% level. The depth premium therefore survives the absorption of cross-provider heterogeneity in pricing.

Panel C adds requester–provider fixed effects, identifying the depth premium from variation *within* a stable (requester, provider) pair over time. All four coefficients become indistinguishable from zero. Panels A and B identify the depth premium across (requester, provider) matches. Providers with deeper relationships to a given requester quote more tightly than providers with shallower ones. Panel C on the other hand asks the different question of whether the same (requester, provider) pair quotes more tightly as it accumulates trades, and the answer is no. The depth premium reflects the persistent quality of which (requester, provider) matches form, not within-match learning over time. Relationship capital is a sorting outcome. Requesters with longer histories with a particular provider obtain better prices because the specific match is a higher-quality one, not because the relationship becomes more productive as it ages.

III.D. Placebo Test

A concern is that part of $\hat{\beta}_3$ could be mechanical rather than economic. Depth and $\log D$ may co-vary within currency-month cells in ways unrelated to the substitution mechanism, simply because

both vary systematically across pairs and over time. In a placebo test we isolate this mechanical component. We randomly reshuffle $\log(1 + d_{r,p,c,t})$ across trades within (pair, year-month) cells, re-estimate equation (4), and repeat the procedure 500 times. The shuffle preserves the joint variation in depth, $\log D$, and the fixed effects within each cell, so any mechanical covariation survives. At the same time, it destroys the match between depth and the actual requester-provider-pair on each trade, which is the only channel through which substitution can operate. The placebo distribution therefore captures only the mechanical component, and the gap between it and the unshuffled estimate reveals the substitution effect.

[INSERT FIGURE 7 HERE]

Figure 7 reports the placebo distribution. The 500 placebo estimates concentrate in a tight band of [0.021, 0.022]. The unshuffled point estimate is $\hat{\beta}_3 = 0.033$, above the entire placebo distribution with a p-value of 0.002. The placebo distribution is not centred at zero: it picks up a residual baseline of about 0.022 that reflects mechanical covariation between depth and $\log D$ within currency-month cells. What identifies the substitution is the incremental 0.011 that the unshuffled estimate adds on top. The substitution mechanism is not an artefact of how depth and competition vary across pairs and time but separately captures who trades with whom.

III.E. Heterogeneity

In Figure 8 we re-estimate equation (4) within four splits of the sample: Panel A G10 versus EM/other currency pairs, Panel B trade-size terciles, Panel C corporate versus institutional requesters, and Panel D bid-ask-spread terciles capturing high, medium, and low contemporaneous liquidity. Each panel displays the competition coefficient $\hat{\beta}_{\log D}$ within each subgroup, with 95% confidence-interval error bars.

[INSERT FIGURE 8 HERE]

Two patterns emerge. First, the substitution interaction $\hat{\beta}_3$ is positive and statistically significant in every one of the eight subgroup cells, in a narrow band of approximately 0.03 to 0.05. The mechanism replicates universally—across currency tier, trade size, requester type, and liquidity state—and at near-identical intensity. Second, the *level* of the competition effect $\hat{\beta}_{\log D}$ varies considerably across subgroups. Corporate requesters and small trades extract roughly two to three times more from competition than institutional requesters and large trades do, consistent with

a bargaining-power reading in which weaker requesters benefit more when forced to shop. The G10/EM and liquidity splits show no statistically distinguishable difference in $\hat{\beta}_{\log D}$.

We highlight the contrast between the universally stable interaction and the heterogeneous main effect. The substitution is a structural property of how competition and relationships interact in this market, operating with the same intensity for every requester type, currency tier, and market state. The size of the overall execution-cost reduction from inviting an additional quote, in contrast, depends on the requester's bargaining position—which in turn is shaped by trade size and client type, the same dimensions that will surface in Section IV.C as the incidence of the bilateral premium.

The cumulative evidence in this section establishes substitution at multiple levels: in non-parametric cell averages (Section III.A), in pooled regressions (Section III.B), within the same request (Section III.C), against a placebo benchmark (Section III.D), and uniformly across sub-groups (Section III.E).

IV. The Choice Between Relationship and Competition

While Section III examines how competition and relationships interact at the margin among multi-dealer trades ($D \geq 2$), we now turn to bilateral trades ($D = 1$). We first establish that bilateral and competitive trades form structurally separate kinds of executions and that the choice between them is sticky within requester. We then ask three questions: who pays the higher cost for executing bilaterally, who concentrates their flow in competitive trades, and how do the two differ within the same requester?

IV.A. Are Dealers Specialised in Relationship or Competition?

To establish that the relationship and competition channels are structurally distinct rather than two ends of a single continuum, we examine how individual requester-provider relationships allocate their flow between them. For each relationship, we compute the *relationship-channel share*: the fraction of trades executed bilaterally ($D = 1$).[4] A share near 0 or 1 means the relationship concentrates its flow in one channel; a share near the middle means it routinely uses both.

[INSERT FIGURE 9 HERE]

[4]The sample is restricted to Mixed RFQ requesters' one-sided trades within relationships that have at least ten such trades, so the share is estimated precisely enough to classify the relationship.

In Figure 9 we plot the distribution of this share at two levels. Panel A computes the share across requester-provider relationships. The distribution is strongly bimodal: 66% of relationships sit at the extremes, with a share below 10% or above 90%, while only 13% fall in the 30%–70% middle. Within a stable requester-provider match, the two channels do not mix: a given client routes almost all trades with one dealer through the relationship channel and almost all trades with another dealer through the competition channel, with little overlap between the two sets of dealers. Panel B computes the share at the provider level, pooling each provider's trades across all the requesters they serve. This distribution is not bimodal. Most providers fall in the middle of the distribution and handle both relationship channel and competition channel trades, just with different clients.

The contrast between the two panels pins down that the two-channel structure is a property of the specific requester-provider match, not of the provider. The same dealer simultaneously serves some clients through the relationship channel and other clients through the competition channel. The data rule out a reading in which small or specialised dealers concentrate in the relationship channel while larger, more diversified dealers run competitive auctions; the two-channel structure is encoded in the (requester, provider) match, not in the dealer's business model.

IV.B. Predicting the Next Move

The bimodal distribution at the relationship level could in principle reflect a single sharp draw at the start of a requester-provider match, after which the channel becomes mechanical. A more telling test of structural separation is whether the channel choice is sticky over time, in the sense that a requester's most recent channel use predicts the next one. We estimate the following regression:

$$\mathbf{1}\{D_{r,p,c,t} = 1\} = \beta\, L_{r,t} + \gamma' X_{r,p,c,t} + \alpha_{c,m} + \alpha_r + \alpha_p + \varepsilon_{r,p,c,t}, \tag{6}$$

where $L_{r,t}$ is an indicator equal to one if requester r's most recent one-sided trade prior to time t went through the relationship channel ($D = 1$), regardless of which provider executed it. $\alpha_{c,m}$ is a pair $\times$ year-month fixed effect, and α_r and α_p are requester and provider fixed effects. The sample is one-sided trades by Mixed RFQ requesters.

[INSERT TABLE IX HERE]

Table IX reports the results. In the specification (1) $\hat{\beta}$ is 0.78, i.e., a requester whose most recent one-sided trade went through the relationship channel is 78 percentage points more likely to

do the same on the next one-sided trade than a requester whose previous trade went through the competition channel. Adding pair × month fixed effects in specification (2) leaves the coefficient unchanged, while adding requester fixed effects in specification (3) absorbs average across-requester differences in channel propensity and $\hat{\beta}$ drops to 0.42. Provider fixed effects (specification (4) leave it at 0.40 and with trade-level controls (specification (5)) $\hat{\beta}$ just marginally drops to 0.38 but remains highly significant.

Moving across specifications, the drop from 0.78 to 0.38 means that roughly half of the unconditional persistence reflects cross-requester heterogeneity: some requesters routinely use the relationship channel and others routinely use the competition channel, and that cross-sectional pattern alone accounts for about 40 percentage points of the raw coefficient. The residual 0.38 is the within-requester persistence: the same requester is 38 percentage points more likely to use the relationship channel after using it last than after using the competition channel. Channel choice is sticky in two ways—stable across requesters, and persistent from one trade to the next within each requester.

The controls in specification (5) provide a first look at what shapes the channel choice beyond pure inertia. Deeper pair-specific dealer history with the executing provider raises the probability of relationship-channel routing (0.019 per log-unit of cumulative depth), as does more recent contact with that provider (-0.002 per log-minute since the last trade). Most strikingly, providers with any prior $D \geq 2$ history with this requester are 14.5 percentage points *less* likely to receive the next relationship-channel trade. The relationship-channel provider is not the provider quoting in this client's competitive auctions—a contrast we examine directly in the next subsection.

IV.C. Who Pays in Relationship Requests?

Above, we establish that relationship and competition trades are structurally distinct kinds of executions, and that the choice between them persists strongly within each requester. We now ask whether a relationship trade costs more than a competition trade for the same requester, and whether the answer differs across requesters.

We return to the tercile classification of requesters by pair-residualised mean transaction cost introduced in Section II: T1 comprises the tercile of lowest-cost requesters while T3 comprises the high transaction cost clients. Within each tercile we estimate the following regression:

$$TC_{i,p,c,t} = \beta \cdot \mathbf{1}\{D_{i,p,c,t} = 1\} + \gamma' X_{i,p,c,t} + \alpha_{r,c,m} + \varepsilon_{i,p,c,t}, \tag{7}$$

where $\mathbf{1}\{D = 1\}$ equals one if the trade is a relationship trade and zero if it is a competition trade at $D \geq 2$; X collects log notional and the quoted bid–ask spread; and $\alpha_{r,c,m}$ is a requester $\times$ pair $\times$ month fixed effect. The estimation sample restricts to Mixed RFQ requesters with at least five trades of each kind. The coefficient β is identified from variation across trades by the same requester in the same pair and the same month, and measures the within-cell premium of a relationship trade over a competition trade.

[INSERT TABLE X HERE]

Panel A of Table X reports the result. For T1 requesters the premium is small and statistically indistinguishable from zero ($\hat{\beta} = 0.04$). It rises with tercile, reaching 0.38 in T2 and 0.81 in T3. The highest-cost requesters pay roughly four-fifths of a basis point more for a relationship trade than for a competition trade in the same pair-month; the lowest-cost requesters pay nothing extra.

Panels B and C confirm the gradient under two robustness exercises. Panel B addresses the concern that classifying requesters on the regression outcome could induce selection bias. For each requester we randomly split trades 50/50, assign terciles using the classification half, and estimate equation (7) on the held-out half. The premium is essentially unchanged: 0.05 for T1, 0.37 for T2 and 0.73 for T3, respectively. In Panel C we replace the transaction cost-based terciles with total U.S. dollar notional traded, a structural characteristic of the requester independent of the regression outcome. Because larger requesters face systematically lower transaction costs, the volume tercile approximately inverts the TC terciles: tercile V3 (largest) is expected to behave like T1 and V1 is like T3. This is what we find in the data: the estimated coefficient is statistically highly significant and ranges from 0.16 for V3 to 0.67 for V2 and 0.58 for V1.

The extra cost of trading bilaterally is therefore not borne equally by all requesters. It falls on the highest-cost tercile: those who, conditional on currency pair, trade at systematically wider spreads. Next, we ask what distinguishes these requesters from the rest of the sample.

IV.D. Who Concentrates in Relationship Requests?

Section IV.B establishes that the choice between relationship and competition trades persists strongly within requester from one trade to the next. It does not, however, identify *which* requesters concentrate their flow at $D = 1$ and which route most of their trades to a multi-dealer auction. We now turn to that question.

Again, we use the tercile classification of requesters to fit a logistic classifier predicting T1 membership from a set of standardised requester-level features. Let X_r collect those features for

requester r, all rescaled to mean zero and unit standard deviation. The model is

$$\Pr(\mathbf{1}\{r \in T1\} = 1 \mid X_r) = \Lambda(X_r'\beta)\,, \tag{8}$$

where $\Lambda(\cdot)$ is the logistic CDF. Variable selection is performed by a logistic LASSO whose ℓ_1 penalty is chosen by 5-fold cross-validation; we then refit an unpenalized logistic regression on the LASSO-selected features and report the post-LASSO coefficients together with t-statistics. Inference is valid conditional on the LASSO having identified the correct sparse model (Belloni, Chernozhukov, and Hansen, 2014). Specification (1) in Table XI reports $\hat{\beta}$ for the contrast T1 vs. T2 ∪ T3; in specification (2) we drop T2 and report the sharper T1 vs. T3 contrast. The predictors span trading activity, network breadth, execution mode, quoted liquidity, and requester type.

[INSERT TABLE XI HERE]

Two predictors dominate both columns. T1 requesters use materially more providers than the rest of the sample ($\hat{\beta}_{\log N_{\text{prov}}}$ is 0.84 in specification (1) and 1.50 in specification (2), the largest magnitude in the model), and route a markedly smaller share of their trades through $D = 1$ ($\hat{\beta}_{\text{share } D=1}$ is -1.32 and −1.83, respectively). Beyond these, T1 requesters execute in tighter quoted bid–ask spreads (−0.20 and −0.41), trade smaller average clip sizes (−0.44 and −0.25), spread their flow more broadly across the pair universe rather than concentrating in G10 (−0.16 and −0.36), and are more often institutional as opposed to corporate clients (0.14 and 0.38). Total notional traded is dropped by the LASSO in both specifications: low-cost requesters are not distinguished by their cumulative trade volume.

The low-cost requester is therefore the large institutional client running a broad multi-dealer programme: many providers, dispersed across the pair universe, executing in tight quoted markets at modest clip sizes, and routing most of their flow to multi-dealer auctions. The high-cost requester displays the opposite profile, a narrow set of providers, concentrated in G10, with a much larger fraction of flow concentrated at $D = 1$.

IV.E. Competition vs. Relationship Trades

In Section IV.D we make some headway towards characterizing the low-cost requester in the cross-section. Now, we ask how the relationship and competition trades of the *same* requester differ from one another within each transaction cost tercile. For each requester-pair combination with

at least five trades of each kind we compute the average of a battery of trade-level and requester characteristics both for $D = 1$ as well as for $D \geq 2$ trades. The within-requester-pair difference

$$\delta^X_{r,c} = \bar{X}^{D=1}_{r,c} - \bar{X}^{D\geq 2}_{r,c} \tag{9}$$

is the object of interest. We demean trade and requester characteristics by currency-pair × year-month before aggregation, so the reported differences are net of common pair-month composition. We estimate $\bar{\delta}^X_T = \mathbb{E}[\delta^X_{r,c} \mid r \in T]$, the average difference across all requester-pair combinations within TC tercile T, and report the p-value from a paired t-test against zero.

[INSERT TABLE XII HERE]

Table XII organizes the ten features into three panels. Panel A collects the four whose gap scales monotonically with the TC tercile, with T1 as the weakest-sorting baseline and T3 as the strongest. For T1 requesters, the relationship and competition trades are statistically indistinguishable in size, and the chosen provider's market share in the pair is also statistically indistinguishable across the two kinds of trade. The cumulative pair-specific depth gap is positive but modest (0.10), and the transaction-cost gap is small (0.15 bps). For T1 requesters, relationship and competition trades look largely interchangeable within their own portfolio.

The sorting sharpens monotonically as we move from T1 to T3. The notional gap, insignificant in T1, grows to −0.24 in T2 and −0.33 in T3. That is, relationship trades in T3 are roughly 28% smaller than their competition trades for the same requester and pair. The depth gap rises from 0.10 in T1 to 0.19 in T3: the provider executing a relationship trade is materially more familiar than the provider executing a competition trade as we move up the tercile ladder. The chosen-provider pair-share gap, statistically zero in the T1 tercile, takes the opposite sign in T3 (0.002), with the sign flipping monotonically across terciles. The transaction-cost gap follows the same pattern, rising from 0.15 bps in T1 to 0.99 bps in T3. High-cost requesters reserve their relationship trades for smaller orders routed to a more familiar provider with a larger share of their business in the pair, and pay a larger premium for doing so.

Panel B reports four features whose within-requester-pair gap is highly significant in every tercile but does not vary across them. Relationship trades go to a provider with whom the requester traded more recently and more often than the competition-trade provider, and to one with materially *less* history of being invited to the requester's competitive auctions. These are universal mechanics of how a relationship trade is routed; the tercile gradient is confined to Panel A.

In Panel C we ask whether T3 requesters choose relationship trades specifically during periods when they have recently been concentrating their flow among fewer providers. Both measures we consider (the number of providers used in the past 180 days and the share with the most-used provider) show no within-requester-pair gap in T3. The captivity of high-cost requesters is structural, not driven by transient narrowing of their dealer set.

V. Conclusion

This paper provides the first comprehensive analysis of FX spot trading under the request-for-quote protocol. Using quote-level data from a major multi-dealer platform that covers the full life cycle of more than fourteen million requests across twenty-eight currency pairs, we document five results that are difficult to reconcile with standard models of OTC trading.

First, transaction costs are sharply non-monotonic in the number of direction-aligned quotes a requester solicits. Bilateral requests obtain tighter prices than requests sent to two providers, and the spike at $D = 2$ is roughly twice the level observed at $D = 1$ or $D = 3$–4. The pattern is robust across currency pairs, trade-size terciles, sub-periods, and the top decile of requesters by volume, and it is present at every within-request quote rank rather than driven by the executed quote alone.

Second, the $D = 2$ premium reflects client selection in the full sample, but not uniformly across client types. Requester fixed effects largely eliminate the average premium, yet low-cost, high-activity requesters continue to display a pronounced increase in costs when moving from one to two dealers. For these clients, the pattern points to a within-client trading strategy in which bilateral relationships and broader auctions are used selectively.

Third, competition and relationships are substitutes at the margin. The marginal value of an additional direction-aligned quote declines as the cumulative requester–provider relationship deepens, with a positive interaction coefficient that survives currency-pair, requester, and provider fixed effects. The substitution is reproduced in non-parametric cell means, in pooled regressions, within the same request, against a placebo benchmark that reshuffles depth within currency-month cells, and at near-identical intensity across G10 and emerging-market pairs, trade-size terciles, corporate and institutional requesters, and liquidity states. Within-request estimates further show that the depth premium reflects which (requester, provider) matches form, not within-match learning over time: once a stable (requester, provider) pair is fixed, the depth coefficient is indistinguishable from zero.

Fourth, bilateral and auction trades are structurally separate channels. The within-relationship share of bilateral trades is sharply bimodal at the requester–provider level but unimodal at the provider level: the same provider serves some requesters bilaterally and others through competitive auctions, while individual relationships concentrate flow in one channel. Channel choice is sticky from one trade to the next within a requester, and bilateral providers are not the requester's auction winners or near-winners. The two channels operate on different (requester, provider) matches drawn from different parts of the requester's counterparty roster.

Finally, we show that the cost of bilateral trading is highly uneven across clients. High-cost requesters pay a substantial premium for bilateral execution relative to competitive trades within the same requester, currency pair, and month. These clients are more likely to be corporate requesters with narrow provider networks and concentrated trading relationships. Low-cost requesters, by contrast, pay no comparable bilateral premium.

Taken together, our findings show that electronic RFQ platforms do not simply transform OTC trading into a standard competitive auction. Competition lowers costs only when it is sufficiently broad and when clients have credible outside options; otherwise, execution costs remain shaped by the structure of the client–dealer match. Relationships can deliver tight prices, but they can also generate rents when clients are captive to a narrow dealer network. The boundary between bilateral and competitive trading is therefore not merely a matter of how many dealers are contacted, but reflects how clients manage privacy, bargaining power, and relationship capital across trades. These results point to a richer view of OTC market design, in which search, information leakage, inventory management, and long-term relationships jointly determine the prices clients pay.

VI. Figures

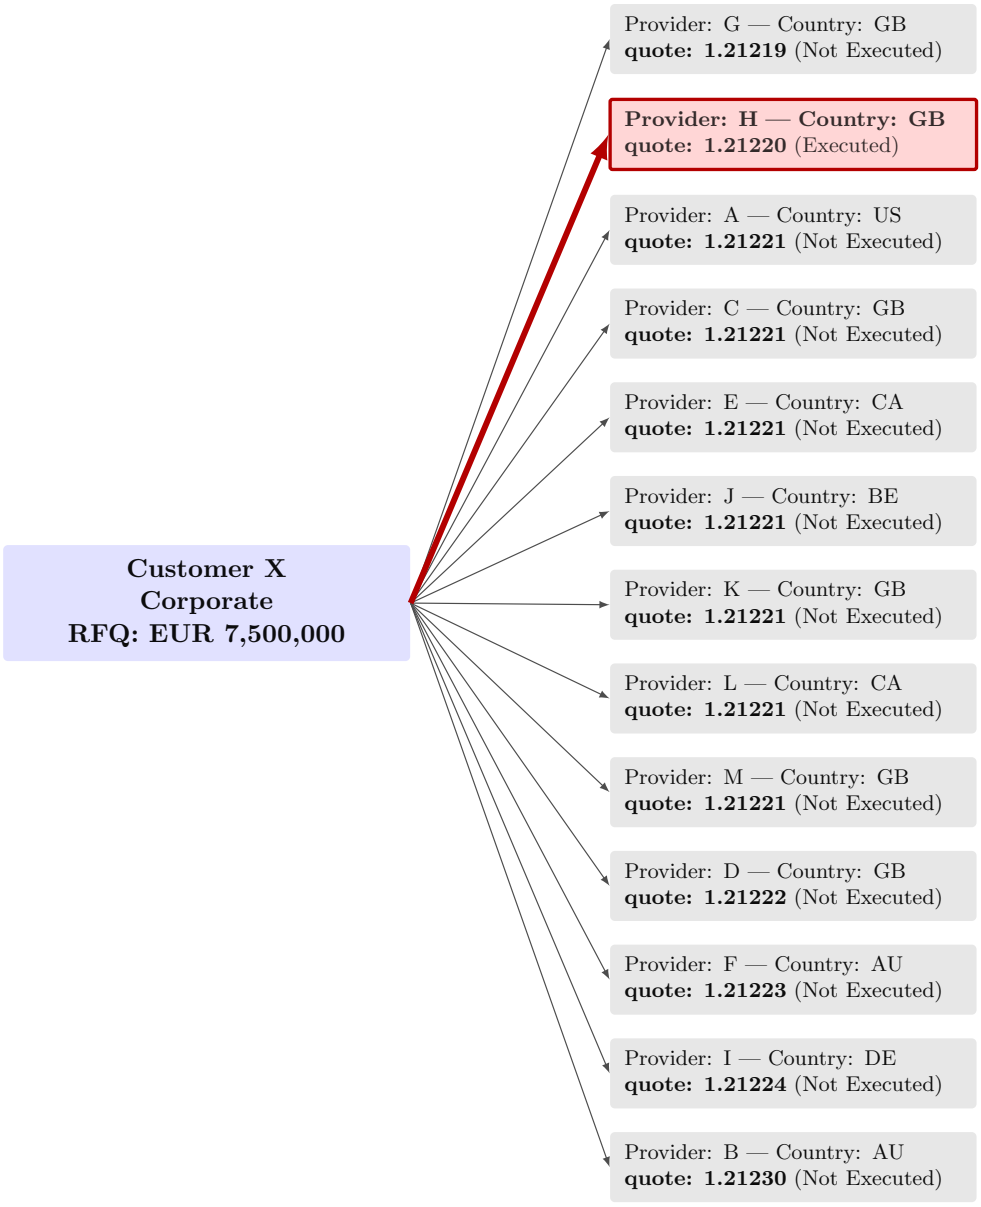


Figure 1. Example of RFQ trade in the FX spot market. This figure illustrates a representative example of RFQ dynamics involving a corporate customer seeking to purchase EUR 7,500,000 in the spot FX market. The (blue) node on the left represents the customer, who solicits quotes from multiple liquidity providers (gray nodes). Each provider responds with a quote, depicted as a line between the customer and the provider node. Only one quote results in execution (thick red line), while all others are non-executed responses (gray lines). Provider nodes are annotated with identifier, country, and execution status; the executed trade also includes the timestamp. All values and identifiers have been modified to preserve the anonymity of data from the trading platform.

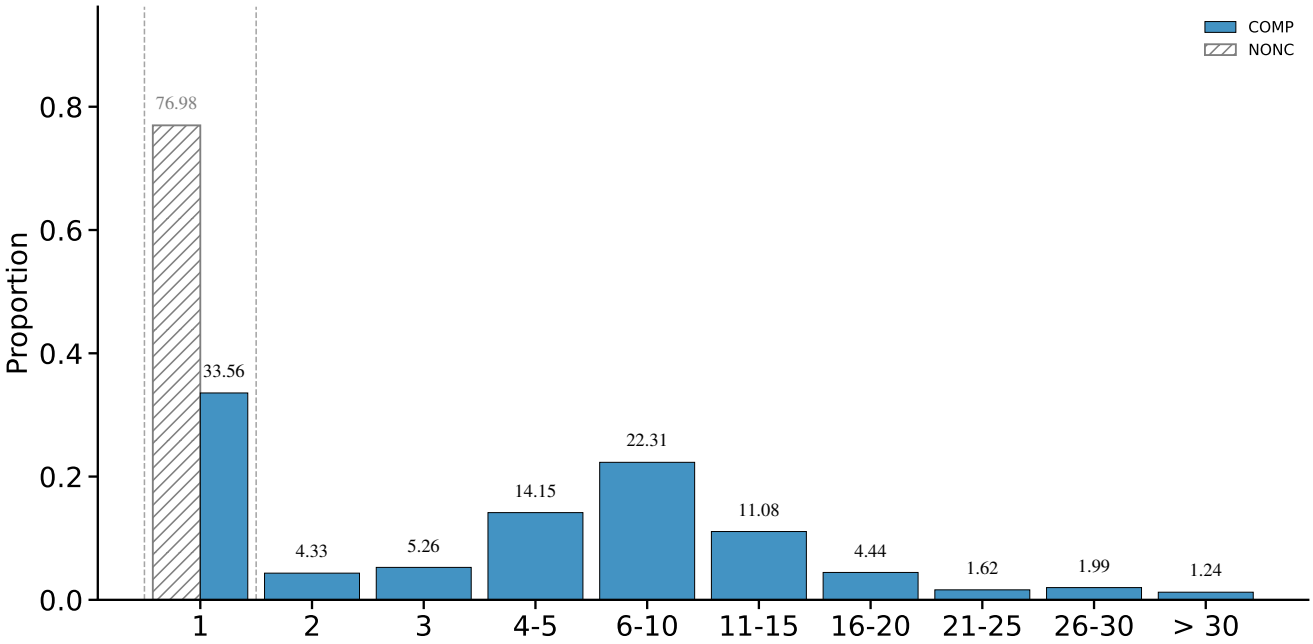


Figure 2. Distribution of quote requests. This figure shows the distribution of RFQs by the number of dealers requested per trade. Single-dealer requests (first bin) are split into two segments. The grey dashed segment denotes the share of non-competitive (NONC) trades relative to the full set of executed trades; the value displayed above the grey segment reports this share. Blue segments (including the first bin) denote competitive (COMP) trades and are expressed as shares of all COMP trades. The values displayed above the blue bars indicate the percentage of COMP trades in which the corresponding number of dealers is approached (e.g., 4.33% of all COMP trades involve two dealers). The sample period is January 2018 to July 2022.

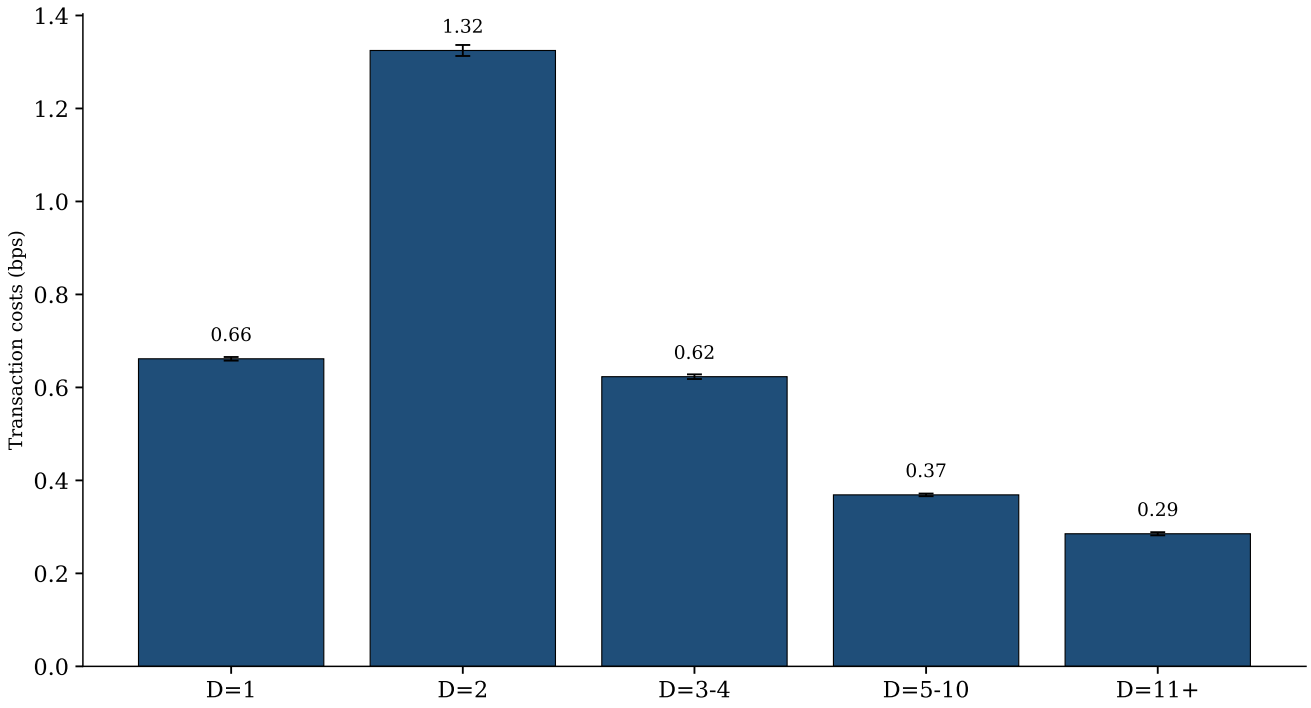


Figure 3. Transaction cost by the number of quotes per request. This figure shows the mean transaction cost (in basis points) for executed trades by the number of quotes per request, D. Transaction costs are computed as the signed log return between the traded price and the contemporaneous mid-price, expressed in basis points, and winsorized at the 1% tails. The sample period is January 2018 to July 2022.

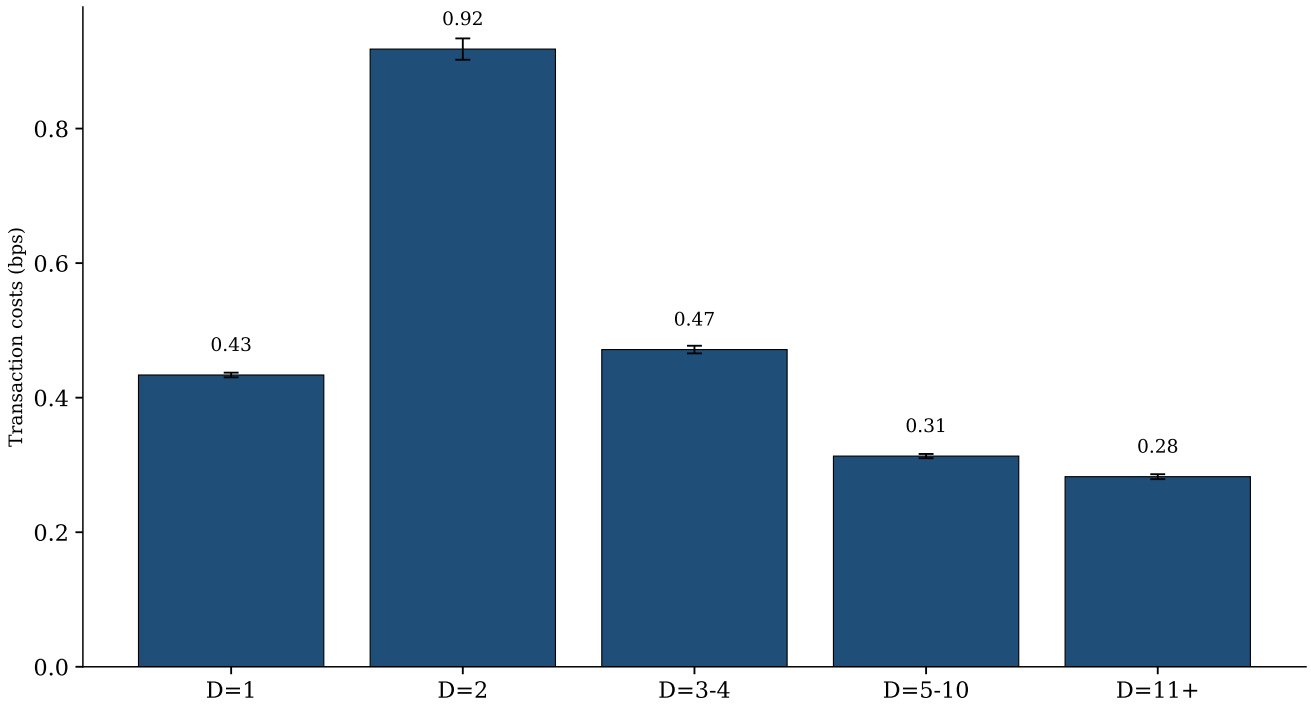


Figure 4. Transaction cost by the number of quotes per request: top-decile clients. This figure shows the mean transaction cost (in basis points) by the number of direction-aligned quotes per request, D, restricted to the top decile of requesters by total volume traded. The 90th-percentile volume cutoff is computed within each sub-period across all (sub-period, requester) pairs in the cleaned sample; a trade is kept if the requester is above the cutoff in the sub-period of the trade. Bars show the cell mean and error bars represent 95% confidence intervals based on heteroskedasticity-robust standard errors. The sample period is January 2018 to July 2022.

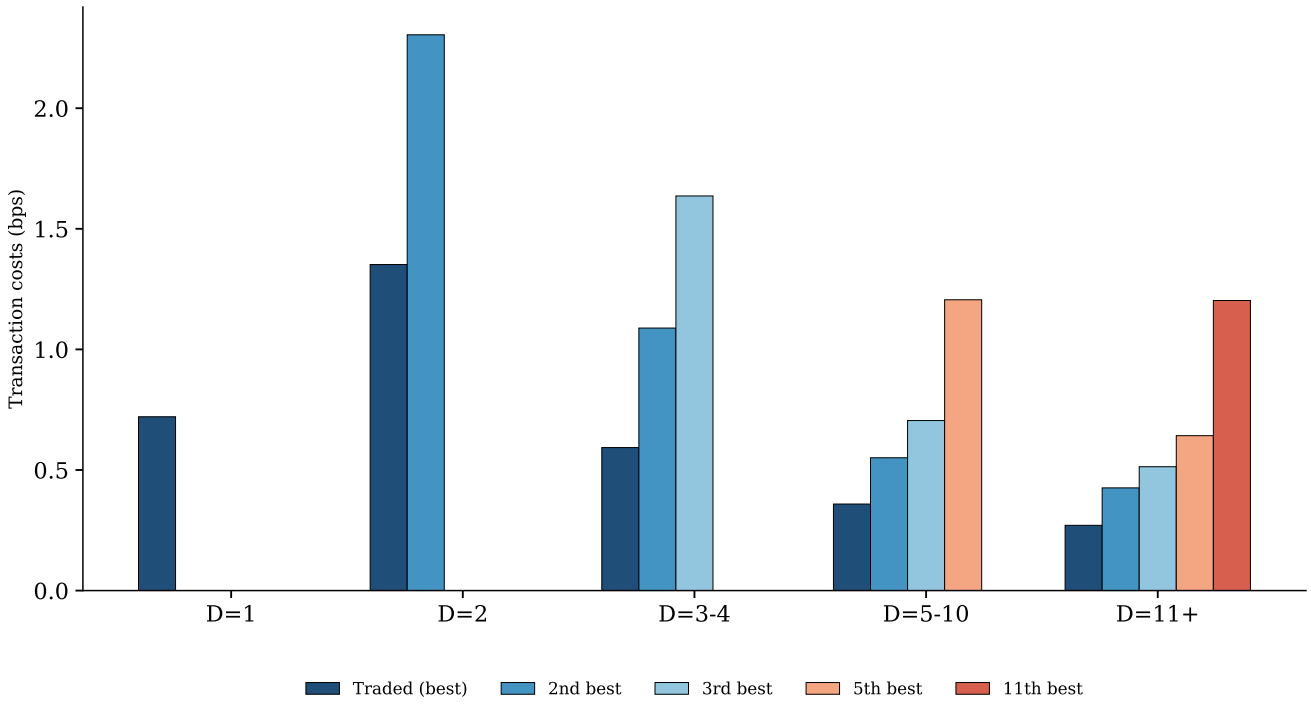


Figure 5. Transaction cost by the number of quotes per request and within-request quote rank. This figure shows the mean implied transaction cost (in basis points) by the number of approached dealers per request, D, and the rank of the quote within the same request. Within each request, we rank all direction-aligned quotes from lowest to highest implied transaction cost. The 1st-best quote is the cheapest offer the requester received on that request, and—by construction—is the executed ("traded") quote. The 2nd-best quote is the next-cheapest offer made on the same request but not selected; the 3rd-, 5th-, and 11th-best quotes are defined analogously and exist only for requests with at least that many direction-aligned quotes. The implied transaction cost of any quote—executed or not—is the signed deviation of that quote from the prevailing mid-price, expressed in basis points (equation 2). Within each D bin, bars report the cross-request mean of the k-th best quote for $k \in \{1, 2, 3, 5, 11\}$, restricted to requests with at least k direction-aligned quotes. Bars are mechanically absent where a rank cannot exist (e.g., the 3rd-best at $D = 2$, since only two direction-aligned quotes are received in that bin). The sample period is January 2018 to July 2022.

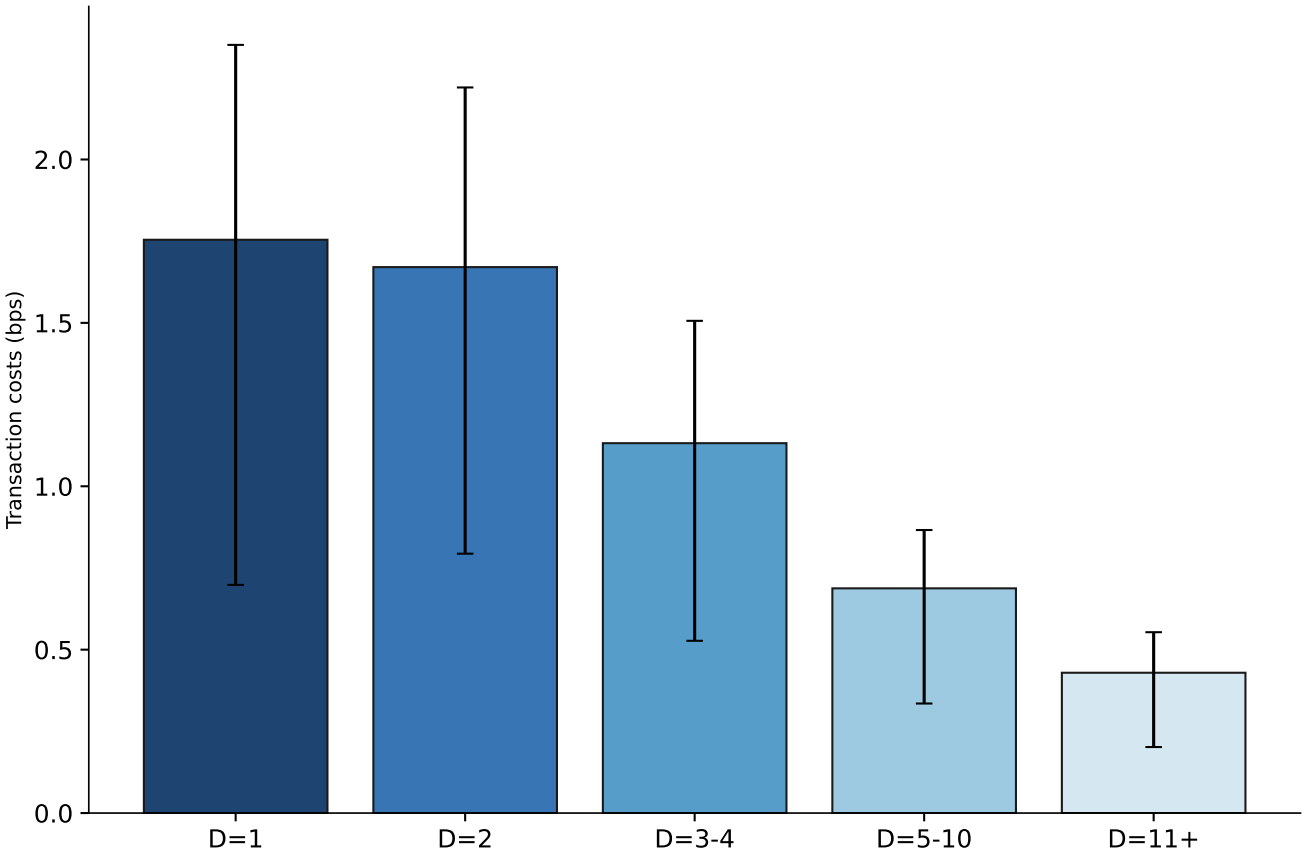


Figure 6. Requester-mean transaction cost by quotes. This figure shows the mean transaction cost (in basis points) by the number of quotes per request, D, computed as the equal-weighted average across requester means. For each requester, the mean transaction cost is first computed within each D-bin; the figure displays the simple average of these requester-level means across all requesters present in a given D-bin. The sample period is January 2018 to July 2022.

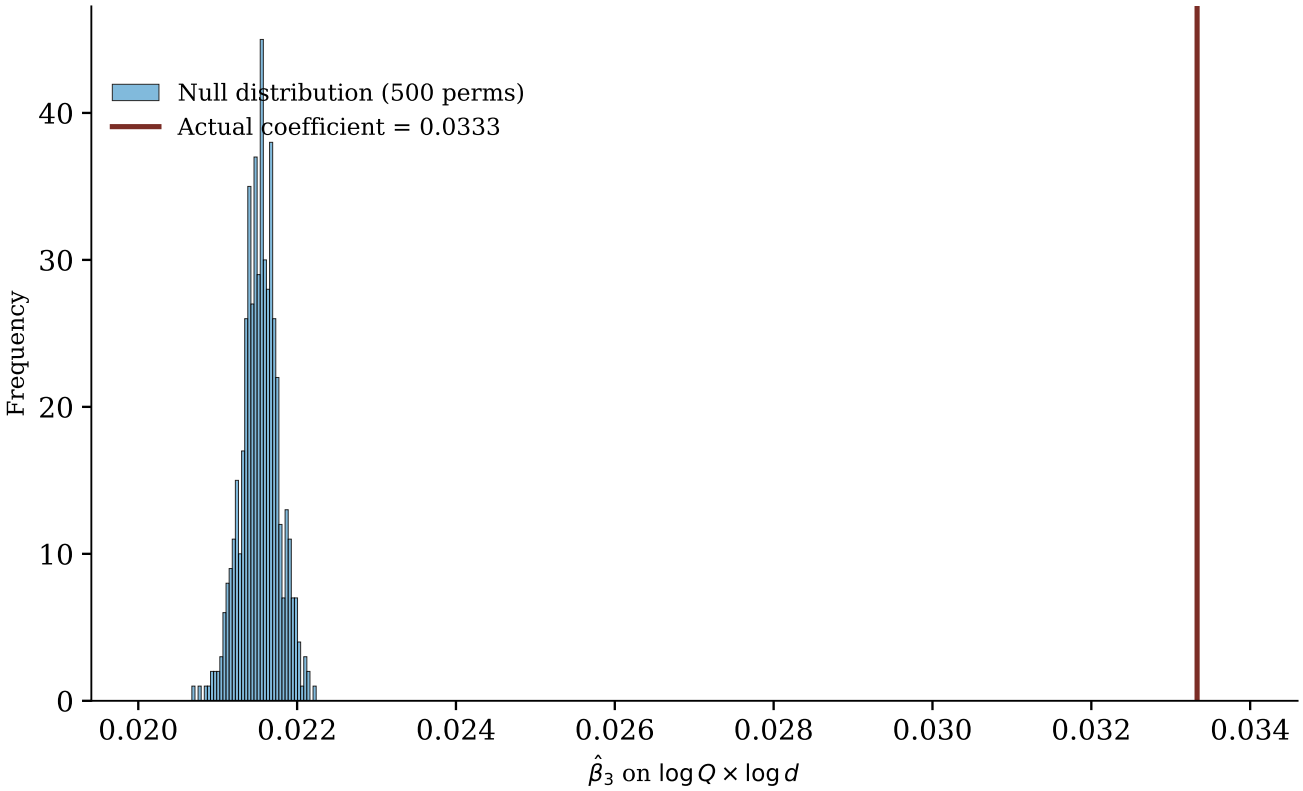


Figure 7. Placebo distribution. In this figure we randomly reshuffle $\log(1+d_{r,p,c,t})$ across trades within (pair, year-month) cells, re-estimate equation (4), and repeat the procedure 500 times. The figure shows the placebo distribution for the 500 draws (blue bars) and the unshuffled point estimate $\hat{\beta}_3 = 0.033$ (red bar).

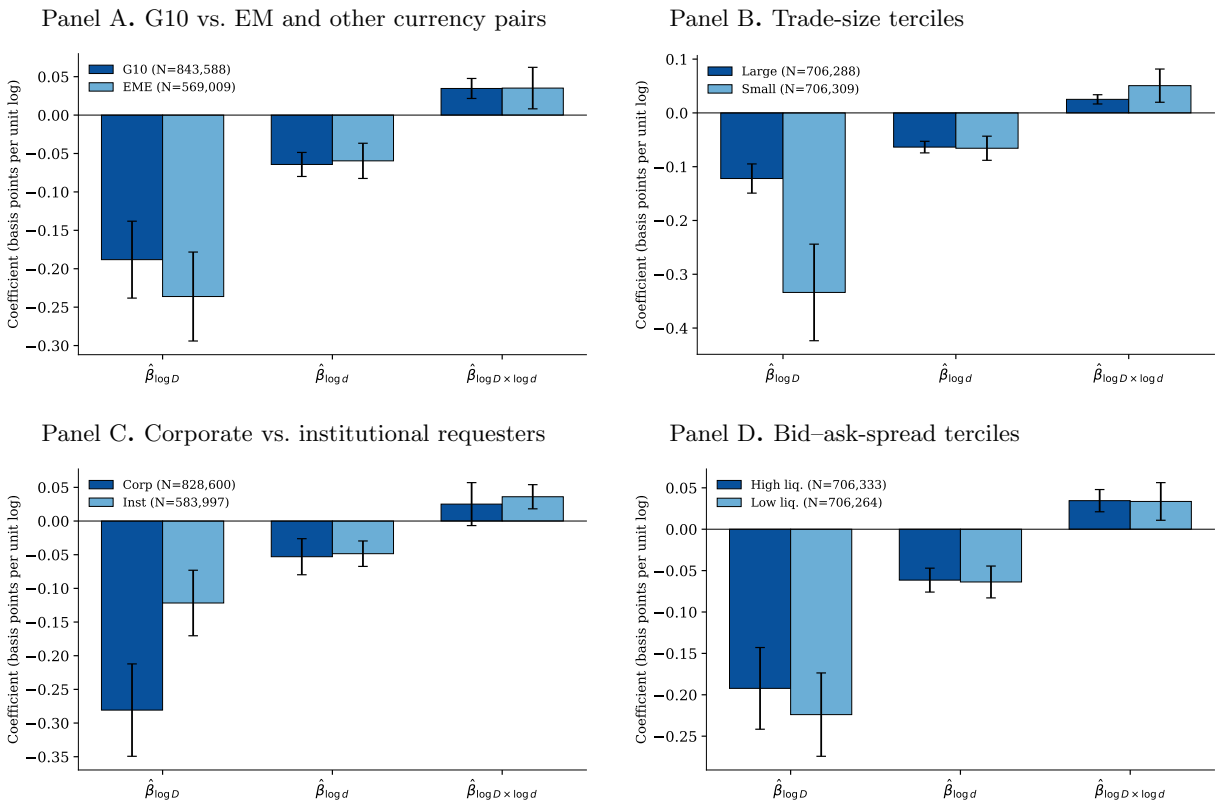


Figure 8. Heterogeneity of the competition effect across subgroups. This figure re-estimates equation (4) within four splits of the sample and reports the competition coefficient $\hat{\beta}_{\log D}$ for each subgroup, with 95% confidence-interval error bars constructed from standard errors two-way clustered by requester and provider. Panel A splits trades into G10 currency pairs versus EM and other pairs. Panel B splits trades into terciles of trade notional in U.S. dollars. Panel C splits requesters into corporate ("Corp") and institutional ("Inst") types. Panel D splits trades into terciles of the prevailing quoted bid–ask spread, capturing high, medium, and low contemporaneous liquidity. In each panel, the horizontal axis indexes the subgroup and the vertical axis reports $\hat{\beta}_{\log D}$. The sample period is January 2018 to July 2022.

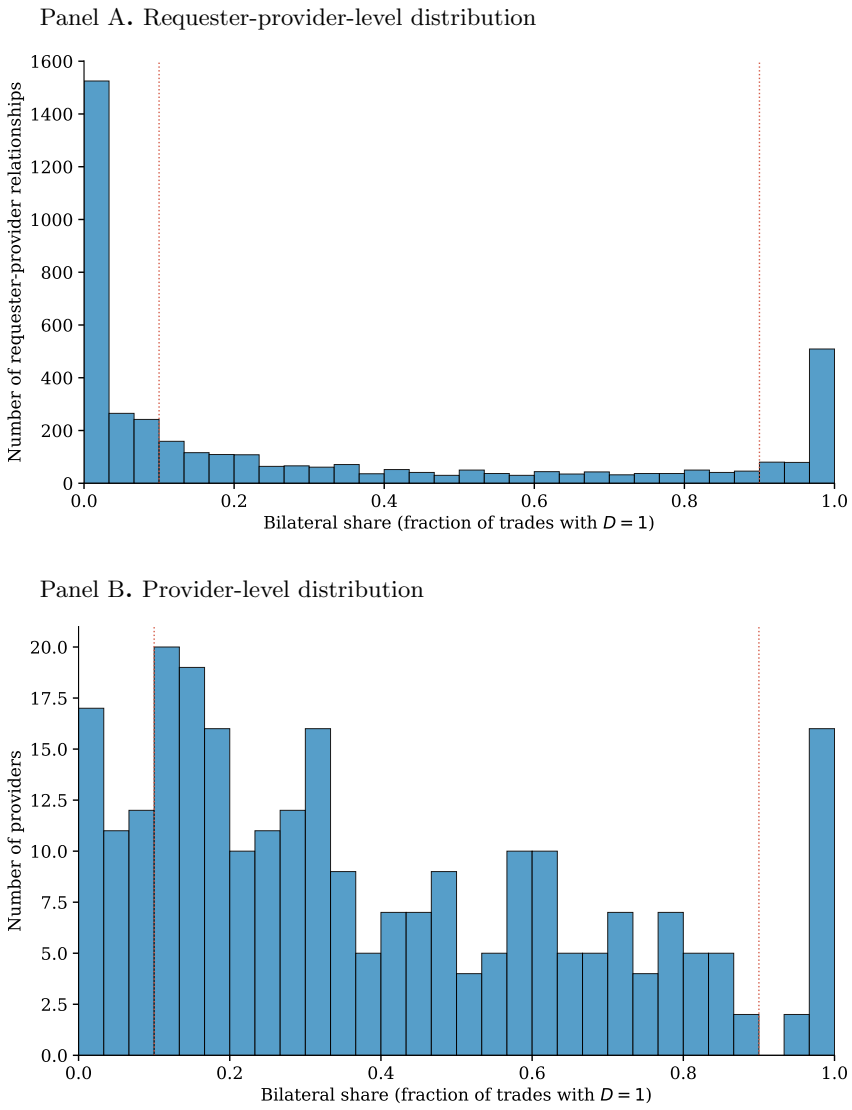


Figure 9. Distribution of bilateral shares at the requester-provider relationship level and at the provider level. This figure shows the distribution of bilateral shares across executed one-sided trades by Mixed RFQ requesters. The bilateral share for a requester-provider relationship is the fraction of one-sided trades in that relationship for which only one direction-aligned provider is invited ($D = 1$). Panel A plots the distribution across all requester-provider relationships with at least ten one-sided trades. Panel B plots the analogous distribution computed at the provider level, pooling each provider's one-sided trades across all the requesters they serve. The horizontal axis is the bilateral share, ranging from 0 (all trades go through the competition channel, $D \geq 2$) to 1 (all trades are bilateral, $D = 1$). Dotted vertical lines at 0.1 and 0.9 mark the bilateral-share extremes. The sample period is January 2018 to July 2022.

VII. Tables

Table I

Summary Statistics: RFQ

This table reports summary statistics for the sample of foreign exchange (FX) spot request-for-quote (RFQ) trades from a major multi-dealer platform. For each currency pair, we report the total number of observations (#Obs), executed trades (#T), and non-executed quotes (#Q), all in millions. We also report the average number of unique liquidity providers (#Prov.), the number of provider countries (#C), the number of unique requesters (#Req.), the percentage of financial institutions (%FI) and corporate clients (%CC), the number of unique provider–requester links (#Links), and the average executed trade size in millions of U.S. dollars (Notl.). Provider, requester, and link counts are computed separately within each of three subperiods and then averaged. The left panel reports USD-denominated pairs; the right panel reports EUR-denominated pairs. The sample period is January 2018 to July 2022.

Panel A: Full sample										
	#Obs	#T	#Q	#Prov.	#C	#Req.	%FI	%CC	#Links	Notl.
All	14.12	1.41	12.70	218	48	959	41	59	6810	2.46
Panel B: USD pairs										
	#Obs	#T	#Q	#Prov.	#C	#Req.	%FI	%CC	#Links	Notl.
AUD	0.39	0.06	0.33	111	38	152	46	54	906	1.59
CAD	0.73	0.07	0.66	120	37	190	63	37	1156	1.52
CHF	0.49	0.08	0.41	117	38	177	30	70	1082	2.61
DKK	0.08	0.01	0.07	72	31	69	36	64	370	1.28
EUR	3.90	0.35	3.54	184	47	641	33	67	4065	3.80
GBP	1.30	0.16	1.14	140	42	319	58	42	1974	1.91
HKD	0.16	0.02	0.13	76	28	83	73	27	476	2.79
HUF	0.10	0.01	0.09	63	25	52	35	65	288	1.04
JPY	0.35	0.06	0.29	111	35	144	72	28	950	2.61
MXN	0.49	0.05	0.44	109	27	135	62	38	852	1.31
NOK	0.12	0.01	0.11	90	29	86	44	56	533	2.28
NZD	0.10	0.02	0.08	77	33	67	61	39	417	0.87
PLN	0.17	0.01	0.15	86	27	88	34	66	513	1.70
SEK	0.24	0.02	0.22	105	33	138	41	59	778	1.77
ZAR	0.23	0.03	0.20	96	32	95	64	36	581	0.93
Panel C: EUR pairs										
	#Obs	#T	#Q	#Prov.	#C	#Req.	%FI	%CC	#Links	Notl.
AUD	0.26	0.03	0.23	109	34	163	28	72	945	1.34
CAD	0.26	0.02	0.24	107	31	134	28	72	835	2.89
CHF	0.80	0.07	0.73	130	39	287	39	61	1664	2.72
DKK	0.15	0.01	0.14	86	31	141	32	68	729	1.41
GBP	1.46	0.14	1.32	146	40	420	33	67	2570	2.36
HKD	0.08	0.01	0.07	77	26	78	50	50	422	2.14
HUF	0.27	0.02	0.25	91	29	169	5	95	1012	1.72
JPY	0.43	0.04	0.39	105	32	161	37	63	1010	1.97
MXN	0.19	0.01	0.18	74	23	81	11	89	466	1.83
NOK	0.25	0.02	0.23	107	32	161	37	63	991	1.83
PLN	0.54	0.03	0.50	105	30	233	7	93	1485	2.18
SEK	0.50	0.04	0.46	117	34	239	24	76	1395	2.10
ZAR	0.09	0.01	0.08	82	27	96	14	86	513	1.21

Table II

Summary Statistics: Client Groups

This table reports summary statistics for the three client groups defined by the per-pair classification (Single RFQ, Mixed RFQ, Multi RFQ). A client's classification is computed for each (requester, currency pair, subperiod) combination based on the requester's number of unique trade providers n_p, the minimum $\min(D)$, and the maximum $\max(QD)$ number of providers invited per request. A requester is Single RFQ in a given pair if she invites only one provider per request but uses multiple providers overall; Multi RFQ if she always invites multiple providers; and Mixed RFQ if she switches between modes. Panel A reports size and scale; Panel B reports composition and execution cost. N is the number of (requester, currency pair, subperiod) classifications in the group; all three groups are active in each of the 28 currency pairs of the working sample. *%Corp* and *%Inst* are the share of trades from corporate and institutional requesters. *Avg. Pairs/Req.* is the average number of currency pairs in which each requester is classified into the group within a subperiod. *%RFM* is the share of trades requested as two-way quotes. The sample period is January 2018 to July 2022.

Panel A: Size and scale

Group	N	Req.	Prov.	Trades	%Trades	Notl. ($B)	%Notl.
Single RFQ	928	383	322	186,303	13.2%	248.5	7.1%
Mixed RFQ	6,767	2,044	596	986,417	69.8%	2,462.7	70.7%
Multi RFQ	6,696	1,848	498	239,878	17.0%	770.4	22.1%
Total	14,391	2,878	654	1,412,598	100.0%	3,481.6	100.0%

Panel B: Composition and execution cost

Group	Med. Notl. (USD)	Avg. TC (bps)	Med. TC (bps)	%Corp	%Inst	Avg. Pairs/Req.	%RFM
Single RFQ	$767,519	0.693	0.431	23.7%	76.3%	2.42	9.4%
Mixed RFQ	$674,040	0.496	0.303	60.8%	39.2%	3.31	41.0%
Multi RFQ	$550,000	0.507	0.349	77.0%	23.0%	3.62	44.4%
Total	$676,400	0.524	0.325	58.7%	41.3%	5.00	37.4%

Table III

Transaction Costs by Quote Requests

This table reports average transaction costs (in basis points) by the number of approached dealers (D) for different subsamples. Transaction costs are computed as the signed difference between the traded price and the contemporaneous mid price, scaled by the mid price and expressed in basis points. Columns (1)–(5) report the mean transaction cost for each D-bin. Column (6) reports the difference between the $D = 5$–10 bin and the $D = 2$ bin along with the p-value from a Welch t-test. Column (7) reports the analogous difference for $D = 11+$ minus $D = 2$. The sample is restricted to non-competitive (COMP $= 0$) executed trades with $|\text{Tradetime} - \text{Quotetime}| \leq 60$ seconds. The sample period is January 2018 to July 2022.

Sample	(1) D=1	(2) D=2	(3) D =3–4	(4) D =5–10	(5) D=11+	(6) $\Delta_{5-10\|2}$ (p)	(7) $\Delta_{11+\|2}$ (p)
All (28 pairs)	0.66	1.33	0.62	0.37	0.29	−0.96 (< .001)	−1.04 (< .001)
EURUSD	0.48	0.80	0.43	0.19	0.07	−0.61 (< .001)	−0.73 (< .001)
USD-base (15)	0.60	1.36	0.57	0.31	0.24	−1.04 (< .001)	−1.12 (< .001)
EUR-base (13)	0.80	1.26	0.76	0.50	0.36	−0.76 (< .001)	−0.90 (< .001)
G10 (9)	0.53	1.29	0.49	0.23	0.15	−1.06 (< .001)	−1.14 (< .001)
EM/other (19)	0.89	1.37	0.83	0.59	0.45	−0.79 (< .001)	−0.93 (< .001)

Table IV

Transaction Costs and Client Selection

This table reports estimates from regressions of transaction costs (in basis points) on quote-bin indicator variables, progressively adding controls and fixed effects. D_k denotes an indicator for quote bin k, with $Q = 1$ as the omitted category. We estimate the following regression:

$$TC_{r,p,c,t} = \alpha + \sum_{k \neq 1} \delta_k D_{k,t} + \gamma' X_{r,p,c,t} + \text{FE} + \varepsilon_{r,p,c,t}$$

where r indexes the requester, p the provider, c the currency pair, t the trade, and $X_{r,p,c,t}$ includes the log of trade notional and the bid-ask spread. Columns progressively add currency-pair fixed effects (α_c), controls, provider fixed effects (α_p), requester fixed effects (α_r), and requester × calendar-week fixed effects ($\alpha_{r,w}$). Standard errors are two-way clustered by requester and provider; t-statistics are reported in parentheses. The sample period is January 2018 to July 2022.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$D = 1$	0.66***							
	(9.30)							
$D = 2$	1.32***	0.59***	0.04	0.41***	0.05	0.15***	0.11***	
	(11.22)	(4.88)	(0.87)	(4.92)	(1.03)	(3.69)	(3.02)	
$D = 3$–4	0.62***	-0.05	-0.24***	-0.11*	-0.22***	-0.04	-0.08*	
	(16.71)	(-0.62)	(-4.62)	(-1.82)	(-4.43)	(-0.82)	(-1.88)	
$D = 5$–10	0.37***	-0.33***	-0.31***	-0.36***	-0.29***	-0.12***	-0.17***	
	(18.13)	(-4.74)	(-6.60)	(-6.35)	(-6.42)	(-3.03)	(-4.27)	
$D = 11+$	0.29***	-0.46***	-0.29***	-0.46***	-0.28***	-0.13***	-0.19***	
	(10.16)	(-7.20)	(-7.22)	(-8.30)	(-7.21)	(-2.93)	(-5.77)	
$\log(D)$								-0.11***
								(-6.93)
log notional (USD)		0.01	0.05***	0.01**	0.05***	0.03***	0.05***	0.05***
		(0.62)	(7.75)	(2.02)	(8.36)	(4.16)	(9.26)	(8.06)
BA spread (bps)		0.03***	0.04***	0.04***	0.04***	0.08***	0.03***	0.04***
		(9.81)	(11.76)	(10.90)	(11.80)	(17.18)	(11.71)	(11.89)
Pair FE		✓	✓	✓	✓		✓	✓
Requester FE			✓		✓			✓
Provider FE				✓	✓			✓
Requester×Provider FE						✓	✓	
Obs. (mn)	1.41	1.41	1.41	1.41	1.41	1.41	1.41	1.41
R^2	0.20	0.15	0.34	0.22	0.35	0.36	0.41	0.35
Within R^2		0.05	0.02	0.04	0.02	0.04	0.01	0.02

Table V

Transaction Costs by Terciles

This table reports estimates of transaction costs (in basis points) on indicator variables for the number of direction-aligned quotes received per request, estimated separately within each TC tercile. For each tercile we estimate

$$TC_{i,p,c,t} = \alpha + \sum_{k \neq 1} \delta_k \mathbf{1}\{D_{i,p,c,t} = k\} + \gamma' X_{i,p,c,t} + \text{FE} + \varepsilon_{i,p,c,t},$$

where $TC_{i,p,c,t}$ is the realised transaction cost in basis points, $\mathbf{1}\{D_{i,p,c,t} = k\}$ is an indicator for the request falling in D-bin $k \in \{2, 3\text{-}4, 5\text{-}10, 11+\}$ with $D = 1$ as the omitted baseline, and $X_{i,p,c,t}$ comprises the log of trade notional and the prevailing quoted bid–ask spread.

We classify requesters into TC terciles based on their average pair-residualized transaction cost across the sample. For each trade we compute the residual $\widetilde{TC}_{i,p,c,t} = TC_{i,p,c,t} - \overline{TC}_c$, where $\overline{TC}_c$ is the mean transaction cost in currency pair c. We then take the requester-level mean of these residuals and split requesters into three equally-sized groups by quantile. T1 ("Low TC") contains the bottom third by client mean residualized cost, T3 ("High TC") contains the top third, and T2 contains the middle. The 959 unique requesters in the cleaned sample are split into approximately 320 requesters per tercile. The estimation sample is further restricted to Mixed RFQ requesters who use both channels, defined as requesters with at least five one-sided $D = 1$ trades and at least five $D \geq 2$ trades.

Within each tercile, three specifications are estimated. The (1) column reports the unconditional cell-mean differences (no fixed effects, no controls); the $D = 1$ row reports the level at the omitted baseline and the remaining rows report the difference from that baseline. The "(2)" column adds currency-pair × year-month fixed effects, requester fixed effects, and the controls $X_{i,p,c,t}$. The "(3)" column replaces the requester fixed effect with a provider fixed effect, retaining the pair × year-month fixed effects and the controls. Standard errors are two-way clustered by requester and provider and reported in parentheses below the coefficient estimates. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The sample period is January 2018 to July 2022.

	T1 (Low TC)			**T2**			**T3 (High TC)**		
	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
$D = 1$	0.29***	ref.	ref.	0.63***	ref.	ref.	2.19***	ref.	ref.
	(0.03)			(0.07)			(0.11)		
$D = 2$	0.29***	0.24***	0.29***	0.24**	−0.08	0.17***	−0.02	−0.54***	−0.29***
	(0.06)	(0.06)	(0.04)	(0.11)	(0.07)	(0.05)	(0.14)	(0.09)	(0.08)
$D = 3$–4	0.08*	0.05	0.11***	−0.01	−0.32***	−0.01	−0.59***	−1.18***	−0.86***
	(0.04)	(0.03)	(0.03)	(0.07)	(0.06)	(0.04)	(0.14)	(0.09)	(0.10)
$D = 5$–10	0.00	−0.08***	0.03	−0.03	−0.49***	−0.22***	−1.05***	−1.53***	−1.27***
	(0.03)	(0.03)	(0.02)	(0.07)	(0.05)	(0.04)	(0.13)	(0.11)	(0.10)
$D = 11+$	−0.05	−0.13***	−0.10***	−0.10	−0.50***	−0.30***	−1.31***	−1.95***	−1.37***
	(0.04)	(0.03)	(0.02)	(0.10)	(0.08)	(0.04)	(0.19)	(0.22)	(0.12)
log notional (USD)		0.04***	0.03***		0.09***	0.06***		0.03*	−0.01
BA spread (bps)		0.02***	0.02***		0.03***	0.03***		0.02***	0.02***
Pair×Month FE		✓	✓		✓	✓		✓	✓
Requester FE		✓			✓			✓	
Provider FE			✓			✓			✓
Obs. (mn)	0.96	0.96	0.96	0.32	0.32	0.32	0.14	0.14	0.14
R^2	0.00	0.14	0.14	0.00	0.25	0.26	0.05	0.47	0.36

Table VI

Transaction Costs and Relationship Depth

This table reports mean transaction costs (in basis points) in cells defined jointly by the number of approached dealers for the request and the cumulative depth of the requester–provider relationship. Columns partition trades by $D_{r,p,c,t}$, the number of direction-aligned quotes received for the underlying request; one-sided trades ($D = 1$) are excluded to isolate the within-competition margin. Rows partition trades by $d_{r,p,c,t}$, the cumulative count of prior trades between requester r and provider p in currency pair c strictly before the focal trade (no look-ahead). The "0 (fresh)" row denotes the first observed trade between the requester–provider–pair triple. Each cell reports the trade-weighted mean transaction cost across all trades falling in the corresponding (depth, competition) bin. The top-left cell (fresh relationship, lowest competition) and the bottom-right cell (deepest relationship, highest competition) are shaded to mark the diagonal corners of the grid. The sample period is January 2018 to July 2022.

Depth	D=2	D=3−4	D=5−10	D=11+
0 (fresh)	2.64	2.21	1.80	1.53
1–5	2.33	1.75	1.36	1.09
6–20	2.04	1.54	1.16	1.00
21–100	1.77	1.29	0.91	0.94
100+	1.53	0.87	0.72	0.83

Table VII

Competition and Relationship Trading

This table reports estimates from the following regression:

$$TC_{r,p,c,t} = \beta_1 \log D_{r,p,c,t} + \beta_2 \log(1 + d_{r,p,c,t}) + \beta_3 \log D_{r,p,c,t} \times \log(1 + d_{r,p,c,t}) + \gamma' X_{r,p,c,t} + \text{FE} + \varepsilon_{r,p,c,t}.$$

The dependent variable $TC_{r,p,c,t}$ is the realized transaction cost in basis points for the trade executed by requester r with provider p in currency pair c at time t. The variable $D_{r,p,c,t}$ denotes the number of direction-aligned quotes received by the requester for the underlying request, and $d_{r,p,c,t}$ is the cumulative count of prior trades between requester r and provider p in currency pair c up to (but not including) the focal trade. Controls $X_{r,p,c,t}$ comprise the log of trade notional and the prevailing quoted bid–ask spread. Columns progressively add fixed effects: currency-pair fixed effects α_c, requester fixed effects α_r, and provider fixed effects α_p. Standard errors are two-way clustered by requester and provider; t-statistics are reported in parentheses below the coefficient estimates. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The sample period is January 2018 to July 2022.

	(1) Pair FE	(2) +Interaction	(3) +Req FE	(4) +Prov FE	(5) Breadth	(6) Loyalty
$\log(D)$	-0.21^{***}	-0.42^{***}	-0.21^{***}	-0.21^{***}	-0.59^{***}	-0.04
	(-8.23)	(-10.94)	(-10.12)	(-10.42)	(-4.76)	(-0.88)
log(depth)	-0.06^{***}	-0.12^{***}	-0.04^{***}	-0.03^{***}		
	(-8.34)	(-7.57)	(-4.83)	(-4.38)		
$\log(D) \times \log(\text{depth})$		0.04^{***}	0.02^{***}	0.02^{***}		
		(5.22)	(4.86)	(5.38)		
log(breadth)					-0.48^{***}	
					(-6.65)	
$\log(D) \times \log(\text{breadth})$					0.16^{***}	
					(4.12)	
loyalty						1.35^{***}
						(3.57)
$\log(D) \times$ loyalty						-0.39^{*}
						(-1.94)
log notional (USD)	-0.01	-0.01	0.05^{***}	0.05^{***}	0.03^{***}	0.03^{***}
	(-1.13)	(-0.79)	(7.33)	(8.03)	(4.02)	(3.14)
BA spread (bps)	0.03^{***}	0.03^{***}	0.04^{***}	0.04^{***}	0.03^{***}	0.03^{***}
	(9.93)	(10.09)	(11.93)	(11.92)	(8.75)	(8.42)
R^2	0.14	0.15	0.34	0.35	0.17	0.16
Within R^2	0.05	0.05	0.02	0.02	0.08	0.07
Pair FE	✓	✓	✓	✓	✓	✓
Requester FE			✓	✓		
Provider FE				✓		

Table VIII

Within-Request Estimates of the Relationship Depth Effect

This table reports estimates of the following within-request specification:

$$QC_{i,p} = \beta \log(1 + d_{i,p}) + \gamma' X_{i,p} + \alpha_i + \varepsilon_{i,p}.$$

The dependent variable $QC_{i,p}$ is the "quoted " transaction cost (in basis points) implied by provider p's direction-aligned quote on request i, measured as the signed deviation of the quote from the prevailing mid-price. The variable $d_{i,p}$ is the cumulative count of prior executed trades between the requester on request i and provider p in the relevant currency pair, measured as of the time of request i (no look-ahead). Controls $X_{i,p}$ comprise the log of trade notional and the prevailing quoted bid–ask spread. Columns partition competitive requests ($D \geq 2$) by the number of direction-aligned quotes received: $D = 2$, $D = 3$–4, $D = 5$–10, and $D = 11+$. Panel A reports the baseline specification with request fixed effects α_i only, which absorb every request-level confound including currency pair, notional, time of day, requester identity, and prevailing market state. Panel B replaces the request fixed effect with currency-pair and provider fixed effects (α_c, α_p) and adds the controls $X_{i,p}$. Panel C uses requester–provider fixed effects ($\alpha_{r,p}$) with the same controls. The bottom row reports the number of provider-quote observations (in millions) underlying each column. Standard errors are two-way clustered by requester and provider; they are reported in parentheses below the coefficient estimates. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The sample period is January 2018 to July 2022.

	D=2	D=3−4	D=5−10	D=11+
Panel A: Request FE				
log(1 + depth)	−0.377***	−0.241***	−0.206***	−0.171***
	(0.045)	(0.075)	(0.019)	(0.022)
Within R^2	0.126	0.074	0.063	0.045
Panel B: Pair FE + Provider FE (plus log notional, BA spread)				
log(1 + depth)	−0.126***	−0.174***	−0.075***	−0.075***
	(0.023)	(0.029)	(0.010)	(0.012)
Within R^2	0.022	0.052	0.027	0.043
Panel C: Requester×Provider FE (plus log notional, BA spread)				
log(1 + depth)	−0.012	−0.021	−0.015	−0.005
	(0.024)	(0.015)	(0.011)	(0.018)
Within R^2	0.063	0.064	0.093	0.115
N (mn)	0.12	0.66	2.97	5.05

Table IX

Path Dependence of the Channel Choice

This table reports estimates of behavioural persistence in the relationship-channel choice. The dependent variable is an indicator equal to one if trade i is executed through the relationship channel ($D_i = 1$) and zero otherwise ($D_i \geq 2$). We estimate

$$\mathbf{1}\{D_i = 1\} = \beta\, L_{r,t} + \gamma' X_i + \alpha_{c,m} + \alpha_r + \alpha_p + \varepsilon_i,$$

where i indexes a trade with associated requester r, provider p, currency pair c, and year-month m. The focal regressor $L_{r,t}$ ("*Prev.* $D = 1$") equals one if requester r's most recent one-sided trade prior to i – regardless of which provider executed it – was itself bilateral ($D = 1$), and zero otherwise. Controls X_i comprise: *Log notional*, the log of trade notional in USD; *BA spread*, the prevailing quoted bid–ask spread in basis points; *Log depth*, the log of cumulative pair-specific depth between the requester and the executing provider; *Log recency*, the log of minutes since the requester's last trade with the executing provider; and *Prior* $D \geq 2$, an indicator for whether the executing provider has any prior $D \geq 2$ history with this requester. The sample is one-sided trades by Mixed RFQ requesters. Columns (1) to (5) progressively add fixed effects and controls. Standard errors are two-way clustered by requester and provider and reported in parentheses below the coefficient estimates. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The sample period is January 2018 to July 2022.

	(1)	(2)	(3)	(4)	(5)
Prev. $D = 1$	0.783***	0.730***	0.418***	0.396***	0.384***
	(0.025)	(0.026)	(0.030)	(0.030)	(0.029)
Log notional					−0.007***
					(0.002)
BA spread					0.001**
					(0.000)
Log depth					0.019***
					(0.004)
Log recency					−0.002***
					(0.001)
Prior $D \geq 2$					−0.145***
					(0.026)
N (mn)	0.58	0.58	0.58	0.58	0.58
R^2	0.599	0.615	0.713	0.733	0.742
Within R^2		0.521	0.190	0.181	0.200
Pair×Month FE		✓	✓	✓	✓
Requester FE			✓	✓	✓
Provider FE				✓	✓

Table X

Cost of Relationship Trading

This table reports estimates of the within-cell cost of executing through the relationship channel, estimated separately within each client tercile. For each tercile we estimate

$$TC_{i,p,c,t} = \beta \cdot \mathbf{1}\{D_{i,p,c,t} = 1\} + \gamma' X_{i,p,c,t} + \alpha_{r,c,m} + \varepsilon_{i,p,c,t},$$

where $\mathbf{1}\{D = 1\} = 1$ if the trade is a one-sided request to a single direction-aligned provider (the relationship channel) and 0 if it is executed at $D \geq 2$ (the competition channel); X comprises log trade notional and the quoted bid–ask spread; $\alpha_{r,c,m}$ is a requester × pair × month fixed effect.
Panel A classifies requesters into TC terciles by their average pair-residualised transaction cost across the cleaned sample. In Panel B we randomly split trades 50/50 for each requster, and assign terciles using the classification half, and estimate equation (7) on the held-out half. Panel C classifies requesters by total USD notional traded, V3 (largest). The estimation sample is Mixed RFQ requesters with at least five one-sided $D = 1$ and five $D \geq 2$ trades. Standard errors are two-way clustered by requester and provider. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***. The sample period is January 2018 to July 2022.

	(1)	(2)	(3)
Panel A: TC tercile classification (full sample)			
	T1 (Low TC)	T2	T3 (High TC)
$\mathbf{1}\{D = 1\}$	0.042	0.376***	0.812***
	(0.037)	(0.084)	(0.140)
Requester × Pair × Month FE	✓	✓	✓
N trades	649,782	171,494	65,619
Within R^2	0.005	0.024	0.042
Panel B: TC tercile classification, split sample			
	T1 (Low TC)	T2	T3 (High TC)
$\mathbf{1}\{D = 1\}$	0.050	0.371***	0.725***
	(0.040)	(0.092)	(0.143)
Requester × Pair × Month FE	✓	✓	✓
Estimation sample	Half-2	Half-2	Half-2
N trades	314,480	85,489	30,683
Within R^2	0.005	0.025	0.033
Panel C: Volume tercile classification (total USD notional)			
	V3 (Largest)	V2 (Mid)	V1 (Smallest)
$\mathbf{1}\{D = 1\}$	0.162***	0.671***	0.582***
	(0.044)	(0.161)	(0.202)
Requester×Pair×Month FE	✓	✓	✓
N trades	850,300	31,201	5,394
Within R^2	0.008	0.035	0.029

Table XI

Low and High Transaction Cost Requesters

This table reports coefficients from two classifiers predicting low-TC tercile (T1) membership at the requester level. Variable selection is performed by a logistic LASSO with ℓ_1 penalty chosen by 5-fold cross-validation across 30 candidates, scored by negative log loss; reported coefficients are from a post-LASSO unpenalized logistic regression on the selected features. Predictors are standardised to mean zero and unit standard deviation prior to fitting, so coefficient magnitudes are directly comparable within a column. Specification (1) classifies T1 against the pooled higher-tercile group (T2 $\cup$ T3); specification (2) drops T2 and classifies T1 against T3 only. t-statistics are reported in parentheses. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***. A dash indicates the LASSO has not selected the variable. Post-LASSO inference is conditional on the LASSO having identified the correct sparse model (Belloni, Chernozhukov, and Hansen (2014)). Requesters are restricted to those with at least 10 trades in the cleaned sample. The sample period is January 2018 to July 2022.

	(1) T1 vs. rest	(2) T1 vs. T3
Activity		
Log N trades	0.296**	0.429**
	(2.29)	(2.02)
Log total notional (USD)	–	–
Log mean notional (USD)	−0.443***	−0.246***
	(−7.76)	(−3.49)
Log active trading days	−0.224*	−0.423**
	(−1.83)	(−2.44)
Network breadth		
Log N unique providers	0.839***	1.503***
	(8.01)	(9.94)
Log N unique pairs	–	0.034
		(0.32)
Provider HHI	−0.106	0.137
	(−1.07)	(1.11)
Channel behaviour		
Share trades at $D = 1$	−1.319***	−1.828***
	(−7.62)	(−8.44)
Share one-sided $D = 1$	0.736***	0.890***
	(4.38)	(4.16)
Share two-way (RFM)	0.217***	0.212***
	(3.68)	(2.67)
Share trades in G10 pairs	−0.162**	−0.357***
	(−1.99)	(−3.41)
Quoted liquidity		
Mean BA spread (bps)	−0.201**	−0.405***
	(−2.46)	(−3.93)
Type		
Institutional (= 1)	0.136**	0.383***
	(2.48)	(4.52)
N requesters	2,570	1,714
$\Pr(T_1)$	0.333	0.500
Cross-validated C	0.386	0.728

Table XII

Relationship vs. Competition

This table reports within-(requester × pair) paired differences between the relationship channel ($D = 1$) and the competition channel ($D \geq 2$), estimated separately within each TC tercile. For each (requester, pair) cell with at least five trades in each channel we compute the cell mean of each feature on $D = 1$ trades and on $D \geq 2$ trades; the table reports the across-cell mean of those within-cell differences and the p-value from a paired t-test against zero. All features are demeaned by currency-pair × year-month before cell aggregation, so the reported gaps are net of pair-month composition and comparable in magnitude to within-FE regression coefficients.

Panel A collects the four features whose within-cell gap scales monotonically with TC tercile. Panel B reports four features whose gap is statistically significant in all three terciles but does not differentiate them: these are universal channel mechanics. Panel C reports portfolio-state features computed at the moment of the trade (N unique dealers used and a top-dealer-share HHI proxy, both windowed over the past 180 days); the absence of a significant gap in N unique dealers and the small magnitude of the top-dealer-share gap indicate that T3's captivity is structural rather than driven by time-varying portfolio breadth. The estimation sample is Mixed RFQ requesters with at least five $D = 1$ and five $D \geq 2$ trades, and trades that are either one-sided $D = 1$ or $D \geq 2$; RFM trades at $D = 1$ are excluded. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***. The sample period is January 2018 to July 2022.

	T1 (Low TC) $N_{\text{cells}} = 618$	T2 $N_{\text{cells}} = 412$	T3 (High TC) $N_{\text{cells}} = 371$
Panel A: Differentiating gradient (T1 → T3)			
Log notional (USD)	−0.020	−0.244***	−0.330***
Log cum pair depth (chosen dealer)	0.104***	0.146***	0.194***
Log chosen-dealer pair share	−0.001	−0.002**	0.002***
Transaction cost (bps)	0.149***	0.501***	0.993***
Panel B: Universal channel mechanics			
Log minutes since last with chosen dealer	−0.527***	−0.703***	−0.596***
Previous trade was one-sided $D = 1$	0.333***	0.235***	0.254***
Any $D = 1$ in last 5 trades with this dealer	0.428***	0.443***	0.408***
Chosen dealer: any prior $D \geq 2$ history	−0.074***	−0.139***	−0.126***
Panel C: Portfolio-state nulls (180-day window)			
N unique dealers, last 180 days	−0.032	0.137	0.014
Top-dealer share, last 180 days	0.009***	−0.001	0.009**

References

Baldauf, M., and J. Mollner (2024): "Competition and Information Leakage," *Journal of Political Economy*, 132(5), 1603–1641.

Baldauf, M., J. Mollner, and B. Williams (2025): "Competition and Privacy in Off-Market Trading," Working Paper, UBC Sauder School of Business.

Belloni, A., V. Chernozhukov, and C. Hansen (2014): "Inference on Treatment Effects after Selection among High-Dimensional Controls," *Review of Economic Studies*, 81, 608–650.

Duffie, D., P. Dworczak, and H. Zhu (2017): "Benchmarks in Search Markets," *The Journal of Finance*, 72(5), 1983–2044.

Duffie, D., N. Gârleanu, and H. Pedersen (2005): "Over-the-Counter-Markets," *Econometrica*, 73, 1815–1847.

Evans, M. D. D., and R. K. Lyons (2002): "Order Flow and Exchange Rate Dynamics," *Journal of Political Economy*, 110(1), 170–180.

Hau, H., P. Hoffmann, S. Langfield, and Y. Timmer (2021): "Discriminatory Pricing of Over-the-Counter Derivatives," *Management Science*, 67(11), 6660–6677.

Hendershott, T., and A. Madhavan (2015): "Click or Call? Auction versus Search in the Over-the-Counter Market," *The Journal of Finance*, 70(1), 419–447.

Hendershott, T. J., D. Li, D. Livdan, N. Schürhoff, and K. Venkataraman (2025): "Quote Competition in Corporate Bonds," *forthcoming, The Journal of Finance*.

Kargar, M., B. R. Lester, S. Plante, and P.-O. Weill (2025): "Sequential Search for Corporate Bonds," Working Paper, UCLA.

Lyons, R. K. (1995): "Tests of Microstructural Hypotheses in the Foreign Exchange Market," *Journal of Financial Economics*, 39(2–3), 321–351.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2016): "Information Flows in Foreign Exchange Markets: Dissecting Customer Currency Trades," *Journal of Finance*, 71(2), 601–634.

Ranaldo, A., and F. Somogyi (2021): "Asymmetric Information Risk in FX Markets," *Journal of Financial Economics*, 140(2), 391–411.

Riggs, L., E. Onur, D. Reiffen, and H. Zhu (2020): "Swap Trading After Dodd-Frank: Evidence From Index CDS," *Journal of Financial Economics*, 137(3), 857–886.

Wang, C. (2023): "The Limits of Multi-Dealer Platforms," *Journal of Financial Economics*, 149(3), 434–450.

Yueshen, B. Z., and J. Zou (2022): "Less is More," Working Paper, INSEAD.

A-I. Appendix: Robustness

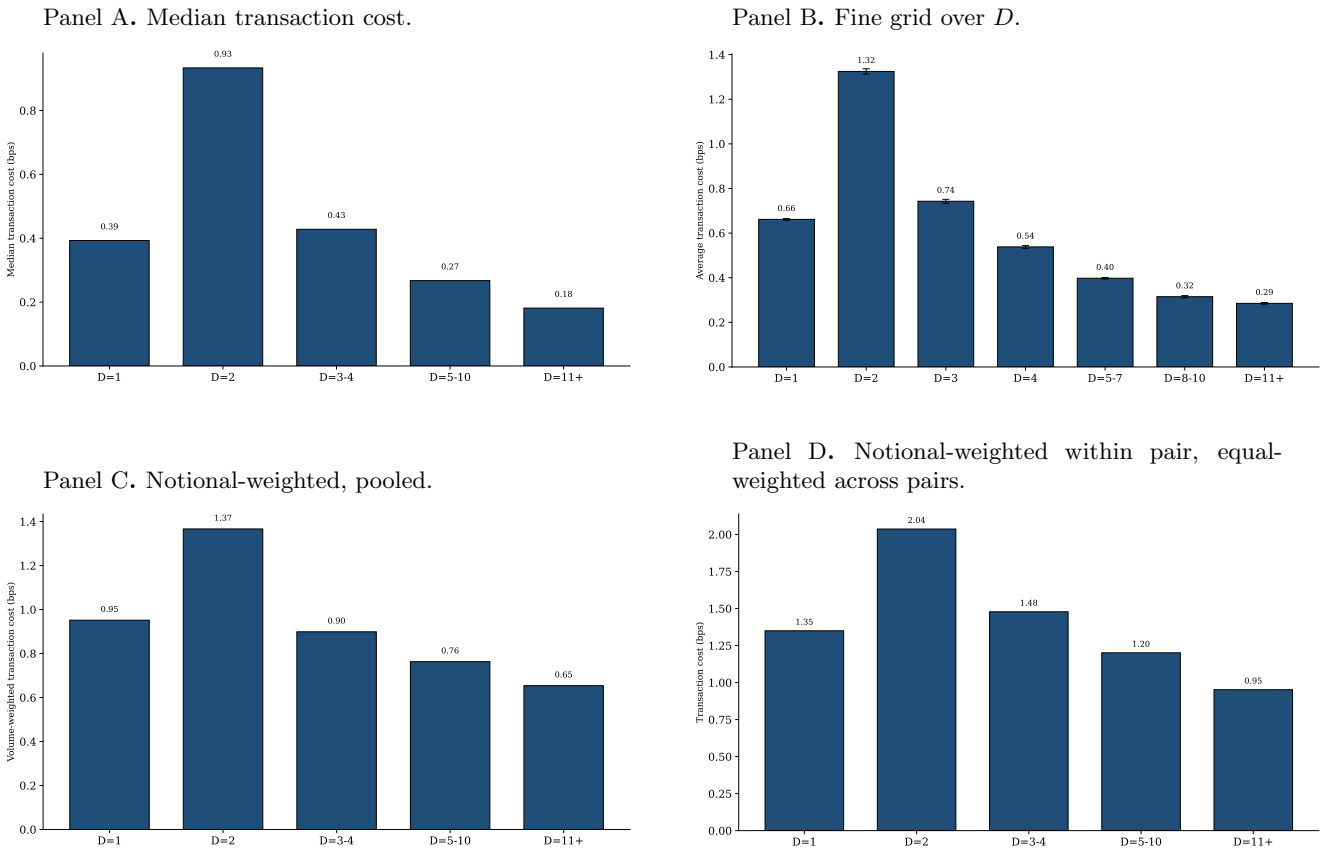


Figure A-1. Robustness I: Transaction costs by the number of dealers approached. This figure shows the transaction cost (in basis points) for executed trades by the number of dealers approached per request, D, under alternative summary statistics and weightings. Panel A reports the median transaction cost in place of the mean. Panel B uses a finer grid over D with bins $D = 1$, 2, 3, 4, 5–7, 8–10, 11+. Panel C reports the notional-USD-weighted mean pooled across all trades. Panel D computes the notional-USD-weighted mean within each currency pair and then averages equal-weighted across pairs. Whiskers on the mean-based panels denote 95% confidence intervals. The sample period is January 2018 to July 2022.

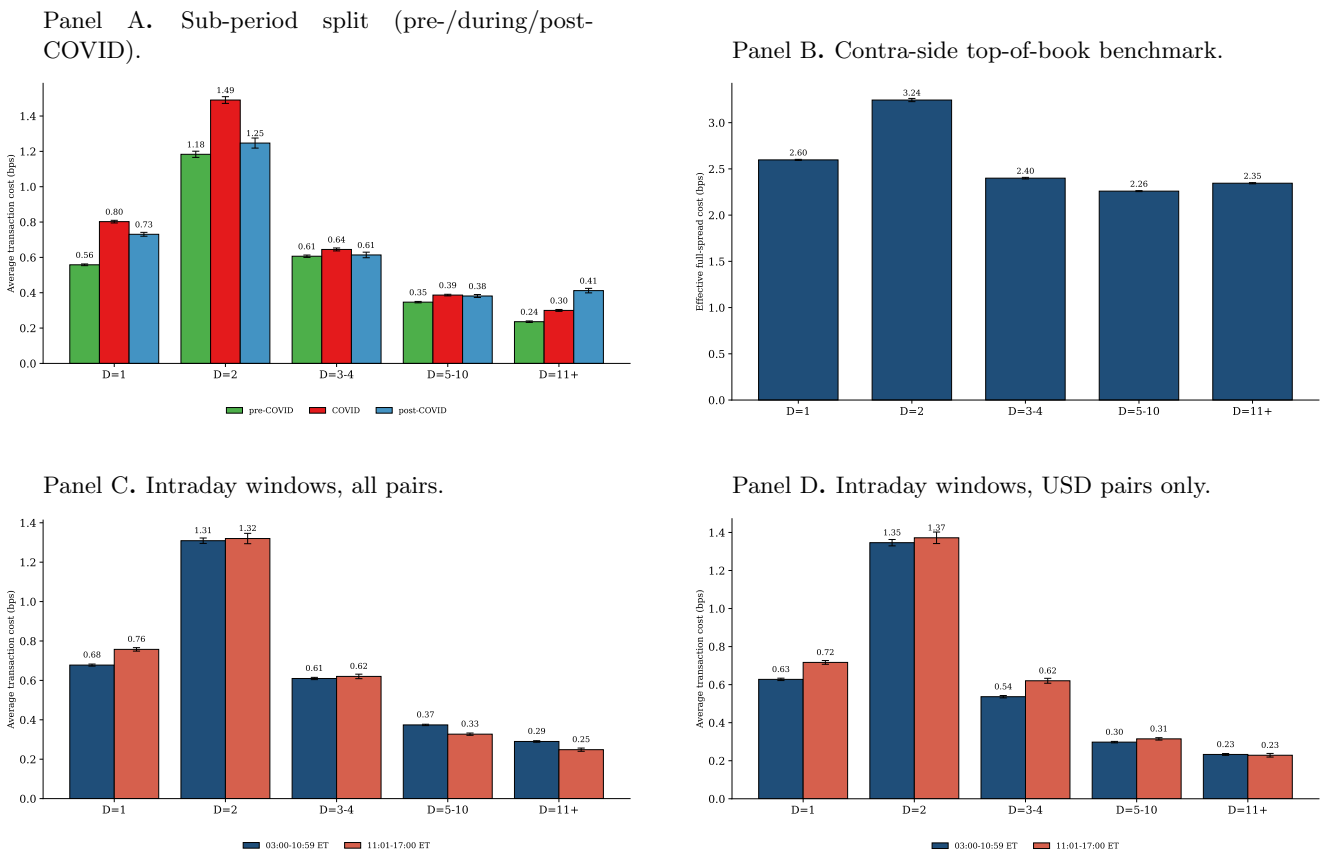


Figure A-2. Robustness II: Transaction costs by the number of dealers approached. This figure shows the mean transaction cost (in basis points) for executed trades by the number of dealers approached per request, D, under alternative sample splits and benchmarks. Panel A splits the sample into pre-COVID, COVID, and post-COVID sub-periods. Panel B replaces the mid-price benchmark with the contra-side composite top-of-book quote, so the transaction cost is measured as $(q_{i,p_i^*} - \text{Bid}_i)/\text{Bid}_i$ on buys and $(\text{Ask}_i - q_{i,p_i^*})/\text{Ask}_i$ on sells, both scaled by 10,000. Panels C and D split the sample by intraday window in US Eastern time (into a morning window associated with persistent USD-buying pressure (03:00–10:59 ET) and an afternoon window associated with persistent USD-selling pressure (11:01–17:00 ET); Panel D restricts Panel C to currency pairs containing USD. The sample period is January 2018 to July 2022.

Table A-I

Transaction Cost by Number of Quotes, Per Currency Pair

This table reports the mean transaction cost (in basis points) by the number of dealers approached per request, D, for each of the 28 currency pairs in the cleaned sample. The columns $\Delta_{5-10|2}$ and $\Delta_{11+|2}$ report the difference in mean transaction cost between $D \in \{5-10\}$ (resp. $D \geq 11$) and $D = 2$, with the p-value of a two-sided t-test of equal means in parentheses. Panel A (Panel B) lists results for USD pairs (EUR pairs). The sample period is January 2018 to July 2022.

	D=1	D=2	D=3-4	D=5-10	D=11+	$\Delta_{5-10\|2}$	(p)	$\Delta_{11+\|2}$	(p)
Panel A: USD pairs									
AUD	0.52	0.97	0.43	0.29	0.34	−0.68	<0.001	−0.63	<0.001
CAD	0.43	1.24	0.46	0.12	0.17	−1.11	<0.001	−1.07	<0.001
CHF	0.44	1.55	0.90	0.30	0.43	−1.26	<0.001	−1.12	<0.001
DKK	1.26	0.69	0.36	0.22	0.05	−0.47	<0.001	−0.64	<0.001
EUR	0.48	0.80	0.43	0.19	0.07	−0.61	<0.001	−0.73	<0.001
GBP	0.71	1.85	0.39	0.19	0.15	−1.65	<0.001	−1.70	<0.001
HKD	0.35	0.16	0.14	0.08	0.04	−0.08	<0.001	−0.12	<0.001
HUF	3.73	3.84	2.62	2.17	1.37	−1.67	<0.001	−2.48	<0.001
JPY	0.27	0.51	0.71	0.14	0.13	−0.37	<0.001	−0.38	<0.001
MXN	0.85	1.59	1.32	0.67	0.64	−0.92	<0.001	−0.95	<0.001
NOK	1.77	1.74	1.30	1.12	0.65	−0.62	<0.001	−1.09	<0.001
NZD	0.82	1.11	0.87	0.33	0.15	−0.78	<0.001	−0.96	<0.001
PLN	2.84	2.72	1.56	1.50	1.30	−1.22	<0.001	−1.42	<0.001
SEK	1.59	1.80	1.15	0.80	0.62	−1.01	<0.001	−1.18	<0.001
ZAR	1.74	2.38	1.78	1.42	1.09	−0.95	<0.001	−1.29	<0.001
Panel B: EUR pairs									
AUD	0.94	1.33	0.67	0.39	0.26	−0.94	<0.001	−1.07	<0.001
CAD	1.17	0.96	0.60	0.37	0.27	−0.59	<0.001	−0.69	<0.001
CHF	0.65	1.04	0.68	0.39	0.26	−0.66	<0.001	−0.78	<0.001
DKK	0.68	0.39	0.20	0.12	0.04	−0.27	<0.001	−0.35	<0.001
GBP	0.49	0.85	0.45	0.28	0.20	−0.57	<0.001	−0.65	<0.001
HKD	1.25	0.55	0.34	0.23	0.10	−0.32	<0.001	−0.45	<0.001
HUF	2.56	3.15	2.24	1.66	1.09	−1.49	<0.001	−2.06	<0.001
JPY	0.35	0.57	0.46	0.21	0.13	−0.36	<0.001	−0.44	<0.001
MXN	0.78	1.68	1.43	1.03	0.63	−0.65	<0.001	−1.05	<0.001
NOK	1.33	1.92	1.32	0.87	0.47	−1.05	<0.001	−1.45	<0.001
PLN	1.21	2.17	1.55	1.08	0.64	−1.09	<0.001	−1.53	<0.001
SEK	1.04	1.56	1.04	0.67	0.45	−0.89	<0.001	−1.11	<0.001
ZAR	3.69	2.49	2.25	1.36	0.88	−1.13	<0.001	−1.62	<0.001