Request for Quote Trading in the Foreign Exchange Market*

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

This version: February 2026

Abstract

We provide the first comprehensive analysis of FX spot trading via request for quote (RFQ) protocols. Execution costs are sharply non-monotonic in competition: bilateral requests obtain the tightest prices, while requesting quotes from two dealers is most expensive. This effect is driven by large clients and financial institutions. Concealing trade direction provides only limited benefits and only when dealers compete. Large clients receive markedly better prices, yet large trades remain costly. The biggest dealers quote most aggressively but still capture most revenues. The evidence challenges standard OTC trading theories and points to a framework in which competition, information design, and market power jointly determine execution prices.

JEL classification: F31, G14, G15, G20

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

*First version: 18 November 2025. 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. 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.

The foreign exchange (FX) market is the largest financial market in the world and underpins global trade, investment, funding, and risk management. However, despite its central role in finance, a substantial fraction of FX trading still occurs in opaque over-the-counter (OTC) environments where quotes are private and execution takes place bilaterally. The increasing use of electronic multi-dealer platforms has introduced a new hybrid trading architecture in which clients request quotes simultaneously from several dealers through a request for quote (RFQ) protocol, but prices remain private and non-binding until execution. This combination of bilateral relationships and real-time competition represents a significant departure from both classical dealer markets and fully transparent limit-order-book structures. Yet, we know surprisingly little about how clients strategically choose dealers, how dealers respond to competition, how informational asymmetries are managed, and how these frictions shape transaction costs at the moment of execution.

This paper provides the first comprehensive empirical analysis of RFQ dynamics in the FX market. Using proprietary data from one of the largest global multi-dealer platforms, we reconstruct the entire RFQ life cycle, starting with the client request, to the panel of dealer responses, ending with the final executed quote. For each RFQ, we observe the identity of the client and all dealers contacted, the direction and notional size of the trade, the full set of quotes submitted, and the quote ultimately accepted. Importantly, we match every RFQ to a high-frequency benchmark mid-price, allowing precise measurement of execution costs at the tick level. The dataset covers more than 22 million RFQs and nearly seven million executed trades between January 2018 and July 2022 across nineteen major currency pairs. To the best of our knowledge, this is the most detailed dataset ever studied on OTC FX trading, and it enables a unique exploration of price formation, liquidity provision, and strategic interaction in modern OTC markets.

First, we show that a higher number of banks competing for a trade does not necessarily lead to lower execution costs. Second, while large clients obtain systematically better pricing on average, individual large trades are not cheaper and instead exhibit higher costs within clients. Third, large dealers consistently offer better quotes but also generate most of the platform's total revenue. Fourth, hiding the intended trade direction through two-sided (or "request-for-market") quotes does not guarantee lower costs and, in some situations, revealing the direction in a single quote request can improve pricing. Fifth, competition is not only relevant at the RFQ stage but also matters in terms of client-dealer relationships: clients with less concentrated dealer relationships typically obtain lower transaction costs.

On our RFQ platform, traders can choose to request quotes from one, two or multiple competing

banks. Intuitively, competition for order flow by multiple banks should result in lower execution costs for the customer, proxied by an executed exchange rate closer to the prevailing mid rate in the interbank market. Our first stylised fact is that this does not hold unconditionally, challenging conventional wisdom that more competition lowers execution costs. That is, requests sent to just one bank obtain *better* execution prices on average than requests sent to multiple banks. However, the standard intuition obtains for requests going to multiple banks and transaction costs are decreasing in competition as measured by the number of banks contacted for a quote once more than one dealer is involved.

Our second finding highlights the difference between trade size on the one hand and client size on the other hand. Across clients, transaction costs are monotonically decreasing in the *total volume* a client trades over the full sample period. That is, large and high-volume clients benefit from execution costs that are only about a third of the cost that the smallest traders pay. In contrast, within large clients, the largest transactions tend to carry the highest transactions costs. Hence, for large clients, the cost-volume relationship is downward sloping in client volume, yet upward sloping in individual trade volume. This asymmetry points to the importance of client market power in FX trading: large clients can negotiate better overall terms, even though each individually large ticket remains relatively expensive to execute. Interestingly, transaction costs are decreasing with trade size for the group of smaller clients.

Third, repeating a similar analysis across dealers instead of clients, we find that large dealers (as measured by trading volume on the platform) offer markedly lower execution costs. Clients trading with the least active dealers face average costs that are around three times higher compared to clients trading with the most active dealers. As the largest 10% of dealers account for over 70% of the total volume on the platform, they still generate most of the overall gross revenue despite earning lower per unit transaction costs.

Fourth, one feature that is special in our setup is the option to request two-sided quotes. This distinctly separates the FX market we study from the only other RFQ market that has received significant attention in the literature so far, the market for corporate bonds (e.g. O'Hara and Zhou, 2021). Two-sided quote requests allow clients to mask their intended trade direction, which should limit strategic shading by dealers. However, we find that concealment only has a limited effect, likely because dealers infer trade direction from client history, flow patterns, or broader RFQ activity visible to them. Moreover, the standard intuition is reversed for clients that submit a quote request to only a *single* bank as they obtain lower transaction costs when this single request

is one-sided and the direction of trade is completely revealed to the dealer. That is, transparency is rewarded in a bilateral trade.

Fifth, we study client-dealer networks and the resulting competition effects on transaction costs. Measured at the client-level, more competition always leads to lower transaction costs. That is, single-quote request clients that trade with different dealers over time obtain better terms than clients that only ever trade with the same dealer. More generally, clients with more diverse and less concentrated dealer relationships experience lower transaction costs.

Our results highlight the interplay between client-dealer relationships, inventory management, information frictions, and bargaining power. The non-monotonic pricing pattern with respect to the number of requested dealers can for example be rationalised by differences in hedging incentives and internalisation capacity. When only one dealer is contacted, the dealer may internalise flow and optimise inventory management, resulting in tighter pricing—especially if the direction of trade is revealed via a one-sided quote request. Moreover, a bilateral trade also allows to avoid information leakage to other market participants as discussed in Hendershott and Madhavan (2015) in the context of RFQ in the corporate bond market or Riggs, Onur, Reiffen, and Zhu (2020) in the swap market. Introducing a second dealer weakens the relational advantage yet does not generate enough competitive pressure to fully discipline quotes, as multiple dealers may pre-hedge simultaneously. Only once a sufficiently large set of dealers is contacted do competitive mechanisms dominate and execution costs decline in line with standard intuition.

The distinction between client-level volume and transaction-level size also illuminates the role of bargaining power. Clients with sustained order flow benefit from better overall pricing, consistent with relationship capital and strategic value as recurring counterparties. Yet, larger individual trades face systematically higher costs within clients, reflecting inventory risk and balance sheet constraints at the trade level. A symmetric phenomenon exists on the supply side: large dealers provide the tightest quotes but still capture the majority of gross revenue, suggesting that liquidity provision benefits from scale economies and netting efficiencies that smaller dealers cannot replicate. Market concentration and cost efficiency thus appear positively related, challenging the notion that dealer dominance necessarily harms clients in OTC settings.

Taken together, the novel facts we uncover suggest that RFQ trading in modern FX markets does not conform to simple search-and-bargaining models. Instead, execution costs reflect an interaction between relationship capital, competitive search, inventory risk, funding constraints, and market structure. 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 risk, funding frictions, and long-term relationships on both the dealer and client side.

Literature Review. Our paper relates to several strands of the literature. First, it contributes to the emerging literature on request for quote (RFQ) dynamics in over-the-counter (OTC) markets. Kargar, Lester, Plante, and Weill (2025) and Hendershott, Li, Livdan, Schürhoff, and Venkataraman (2025) study search costs and quote competition in corporate bond markets, documenting how dealer responses shape execution quality and transaction costs. We extend the evidence in this literature to the FX market, where trading is more frequent, liquidity deeper, and the RFQ design enables strategic concealment of trade direction—features absent in corporate bonds. To our knowledge, this is the first study to analyse the complete quote and trade data for FX RFQs.

On the theory side, Wang (2023) builds a model of trading on multi-dealer platforms with RFQ features where dealers strategically choose how to respond to requests. Yueshen and Zou (2022) show theoretically that OTC market customers in equilibrium contact few dealers and submit fewer but larger orders, in order to incentivize maximum liquidity provision. We provide empirical evidence that allows for testing the theoretical mechanisms described in this literature.

Second, our work relates to the broader literature on market microstructure in FX markets. Prior research has examined price discovery, liquidity provision, and execution quality in dealer-to-dealer and dealer-to-client segments (e.g., Lyons, 1995; Evans and Lyons, 2002). These papers exclusively focus on executed trade prices, whereas we observe both ex-ante quotes and ex-post execution, allowing us to capture the competitive process itself and study the timing, dispersion, and aggressiveness of dealer responses in the spot market.

Third, we build on research that investigates price discrimination and heterogeneous pricing strategies across clients and dealers. Studies such as Hau, Hoffmann, Langfield, and Timmer (2021) highlight how dealers adjust quotes depending on client sophistication and their connection with other dealers, while Ranaldo and Somogyi (2021) and Menkhoff, Sarno, Schmeling, and Schrimpf (2016) document how trading behavior and execution outcomes differ systematically across customer types, including corporates, asset managers, and other financial institutions. Our dataset allows us to extend this literature by linking detailed RFQ-level competition measures to client type, thereby providing new evidence on how dealer pricing strategies and competitive

behavior vary across heterogeneous customers in the FX market.

More broadly, we contribute to theoretical and empirical work on OTC market structure and bargaining (e.g., Duffie, Gârleanu, and Pedersen, 2005; Duffie, Dworczak, and Zhu, 2017). Our results support recent arguments that platform design, search costs, and inventory management jointly shape pricing outcomes. The evidence we present—particularly on scale effects, information frictions, and the U.S. dollar funding premium—helps bridge FX microstructure and the broader literature on liquidity, funding constraints, and intermediation capacity in global financial markets.

I. Data

Request for Quotes Data. We obtain information on the complete request-for-quote (RFQ) trading dynamics from one of the largest secondary trading venues in the foreign exchange market. 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. For each RFQ, we observe an anonymous client identifier, the set of approached dealers—identifiable through unique dealer codes—the geographic location of each dealer, the trade direction (buy or sell), the currency pair requested, the notional volume, and the full sequence of dealer responses. Crucially, we also observe which quote the client ultimately accepts. This level of granularity enables us to reconstruct the complete life cycle of each RFQ, from initial request through competitive quoting to final execution. The data allow for detailed measurement of execution quality, assessment of price improvement, and analysis of liquidity provision and dealer competition at high frequency.

Our sample spans five years from January 2018 to July 2022 and covers 19 exchange rates vis-á-vis the U.S. dollar or the euro: Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), Danish krone (DKK), Euro (EUR), British pound sterling (GBP), Hong Kong dollar (HKD), Hungarian forint (HUF), Israeli new shekel (ILS), Japanese yen (JPY), Mexican peso (MXN), Norwegian krone (NOK), New Zealand dollar (NZD), Polish zloty (PLN), Romanian leu (RON), Swedish krona (SEK), Singapore dollar (SGD), Thai baht (THB), South African rand (ZAR). Thus, overall sample completes 37 exchanges rates - 18 denominared in each major currency - and the euro/U.S. dollar rate itself. To the best of our knowledge, this is the first and most comprehensive empirical study of RFQ dynamics in the FX market using transaction-level and quote data.

Figure 1 illustrates a hypothetical representative example of a client request: A customer seeks

to purchase EUR 7,500,000 in spot. She solicits quotes from multiple dealers. Each dealer responds with a quote as represented by the line between the client and the dealer node. Only one quote results in execution (shown with a thicker red line), while all others are non-executed responses (gray lines). In addition to the quoted price that a dealer provides, we observe an anonymous unique identifier for dealer, her country of location and the executed trade also includes the timestamp.

[INSERT FIGURE 1 HERE]

Reference FX Prices. Lastly, 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 in the previous 60 seconds, ensuring precise time alignment between trade and benchmark prices. Thus transaction costs (TC), expressed in basis points, are defined as

$$\text{TC}_i \equiv b_i \cdot \left(\frac{P_i^{\text{trade}} - P_i^{\text{mid}}}{P_i^{\text{mid}}} \right) \times 10{,}000 \tag{1}$$

where the indicator variable b_i 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. Allowing benchmark quotes to precede the trade by up to one minute and trimming outliers ensures that the benchmark reflects prevailing market conditions while remaining robust to spurious price spikes and stale quotes.[1]

Summary Statistics. Table I documents the richness and granularity of the RFQ database. Across all currency pairs, the dataset comprises 22.67 million observations, of which 6.81 million correspond to executed trades that can be linked to nearly 16 million associated quotes. These quotes are requested by 5,108 distinct clients from 1,156 dealers spanning 50 countries, resulting in more than 31,000 unique dealer–client relationships. The average notional value per trade is USD 2.09 million.

[INSERT TABLE I HERE]

[1] 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.

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 USD 3.28 million. This is followed by other major pairs such as GBP/USD, CAD/USD, AUD/USD, and USD/JPY, 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., 70.4% 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 AND TABLE II HERE]

The remaining 30.6% 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.

While over one third of competitive requests are still directed to only one dealer at a time, thus indicating limited dealer competition for a subset of the "competitive" quotes, the remaining distribution is highly skewed toward larger dealer sets. The mode occurs in the 6 to 10 dealers

range, which accounts for roughly one-fifth of all competitive RFQs. This suggests substantial heterogeneity in client behavior and a non-trivial portion of RFQ activity involves meaningful multi-dealer engagement.

II. Competition and Transaction Costs

A. Number of Quotes

We start by investigating the implications of the key feature in the setup of our FX platform—the ability to request quotes from multiple dealers. Figure 2 visualises the distribution of the various quotes in our sample. As discussed in Section I, over 60% of all competitive RFQs are addressed to more than one dealer, thus allowing us to directly study a first dimension of dealer competition, the relationship between execution costs at the transaction level and the number of dealers approached for a quote. Since our focus is on the effect of competition on pricing, we restrict the analysis in this section to competitive (COMP) trades.

Figure 3 plots the average execution cost—measured as the distance of the execute price relative to the contemporaneous benchmark mid-price from streaming quotes—conditional on the number of quote requests sent to dealers. The main result is obvious and striking: a higher number of competing dealers does not necessarily translate into lower execution costs. In fact, execution costs are lowest when only a single quote is requested, rise sharply when two dealers are approached, and then decline as the number of requested quotes increases, in line with standard intuition.

[INSERT FIGURE 3 AND TABLE III HERE]

The behaviour of transaction costs highlighted in Figure 3 also obtains across currency groups as documented in Table III: we observe the same pattern for pairs quoted against the U.S. dollar and the euro, as well as for G9 and emerging market currencies.[2] Across all currency pairs, costs increase by two thirds when moving from a single-dealer quote request to a quote requested from two dealers, reaching almost 1.5 bps. As quotes are requested from 11 or more dealers, they drop to almost a quarter of a basis point and can be as low as an eighth of a basis point for G9 currency transactions against the U.S. dollar. Importantly, the average cost differentials across buckets (Δ) are economically meaningful and statistically significant, with magnitudes increasing in the

[2]Moreover, the results in Figure 3 and Table III are robust to calculating (i) value-weighted averages across all trades as well as (ii) value-weighted averages normalised within each currency pair. These complementary measures of aggregation ensure that the results are not driven by disproportionately large transactions (captured by VW) or by differences in liquidity across currency pairs (captured by the pair-normalized measure).

distance between buckets—from 0.09 bps between adjacent bins to 1.3 bps when comparing the extremes. Monotonicity in this relationship is formally confirmed by the Patton and Timmermann (2010) test statistic, for which the associated p-value is effectively zero, rejecting the null of no ordered relationship. In Figure 4 we repeat the exercise for the top decile of clients in terms of volume and find that the results remain virtually unchanged—the effect is a systematic feature of FX RFQ market.

[INSERT FIGURE 4 HERE]

The pattern suggests the coexistence of two pricing regimes. Single-dealer requests appear to reflect strong client–dealer relationships, where the dealer potentially internalises flow, manages inventory efficiently, and therefore offers tight pricing. Introducing a second dealer weakens the bilateral relationship advantage, yet the competitive pressure is still too limited to fully discipline pricing, resulting in temporarily higher costs. Once clients request quotes from a sufficiently large group of dealers, classic competitive forces dominate and prices improve almost monotonically.

[INSERT FIGURE 5 AND TABLE IV HERE]

In addition to the transaction costs implied by the best quote, we show in Figure 5 hypothetical costs for additional quotes provided as part of the request. In particular, for every subsequent bin that contains at least k dealers, we add the transaction costs for the k^{th} best requested quote. That is, for the two-dealer bin, we show the costs implied by the first and second quote, for the bin with requests sent to three or four dealers, we plot bars representing the first, second and third best quote, and so on. Interestingly, the costs for the k^{th} best quote are monotonically decreasing across bins. That is, the higher level of competition that drives transaction costs once at least two quotes are requests also seems to improve the implied costs of the entire set of quotes. Moreover, additional quotes do not only lead to better prices but also to a larger price dispersion as measured by the interquartile range of prices within each request. In Table IV we show that regardless of the specification, additional requests lead to a significant increase in the dispersion of quotes within a request.

B. Determinants of Transaction Costs

Returning to the relationship between the number of quote requests and execution costs, we regress the deal-level transaction costs for all competitive trades on the logarithm of the number

of requested quotes and a set of control variables, using different combinations of fixed effects in Table V. Overall, increased competition by including additional dealers in a RFQ robustly reduces transaction costs. At the same time, the size of the trade does not significantly affect costs while poor liquidity conditions in the interbank market as measured by the size of the contemporaneous bid-ask spread are passed on and significantly increase costs. That is, additional costs incurred for managing inventory seem to be passed on to clients.

[INSERT TABLES V AND VI HERE]

As indicated in Figure 2, over 70% of transactions in our sample are executed non-competitively. In Table VI we directly assess the effect of executing trades competitively by restricting our sample to trades that include only a single dealer (either in a competitive trade with a single request or a non-competitive trade) but to clients that execute both types of trades in our sample. For both corporate clients as well as financial institutions, executing a competitive trade results in transaction costs that are around a third lower compared to trades at are executed non-competitively.

In sum, our results suggest that competition matters strongly for spot FX transactions. Clients that executed both competitive as well as non-competitive quotes enjoy much lower transaction costs under the competitive protocol. Moreover, competition within RFQs as measured by the number of contacted dealers lowers transaction costs. However, the fact that single-quote requests overall receive very favorable terms suggest the presence of an important relationship dimension when it comes to costs when trading spot FX. While we explore this angle further in Section V, we first study the effects of client and dealer size on trading terms in the next section.

III. Client and Dealer Size

In this section, we examine how heterogeneity across clients and dealers shapes execution outcomes in the FX RFQ market. The rich panel structure of our data allows us to characterize systematic differences in transaction costs across the distribution of trading activity on both sides of the market. We present three sets of results. First, we document a strong and monotonically decreasing relationship between client size and execution costs. Second, we show that execution costs are largely increasing in trade size. Third, we turn to the supply side and show that liquidity provision is highly concentrated, with the largest dealers offering the tightest prices and generating the bulk of total revenue.

A. *Client Size and Execution Costs*

We begin by grouping all clients into deciles based on the total notional volume they execute via competitive trades. Overall, those trades account for nearly 30% of the total volume of roughly USD 14 trillion over the sample period. Table VII summarizes the properties of each group, and Figure 6 plots average transaction costs across client-size deciles. Several patterns emerge.

First, trading activity is extremely dispersed. The 380 clients in the bottom decile execute less than USD 1 billion in total notional volume, whereas clients in the top decile trade roughly USD 3.5 trillion the same period from January 2018 to July 2022. Second, larger clients not only trade larger overall notional amounts, the average trade size as well as the average number of trades per client are both also monotonically increasing with client size. The average number of quote requests is also increasing with client size although the increase remains very moderate for the lower deciles, increasing by a little bit more than one when moving from decile 1 to decile 5. Third, the fraction of two-sided quote requests on the other remains fairly stable, dropping only by 7 percentage points when moving from decile 1 to decile 10. For most deciles apart from the extreme ones, the fraction of corporate clients is between 91% and 94%, with 23% of clients in the top decile being financial institutions.

[INSERT TABLE VII AND FIGURE 6 HERE]

Finally, and most importantly, execution costs decline steeply with client size. Equal-weighted costs fall from more than two basis points in the lowest decile to two thirds of a basis point for the largest clients, a drop of over 1.7 bps. This is economically very large in light of the fact that average transaction costs in our data are below one basis point. Moreover, the decline is monotonic, regardless of whether we consider equally- or volume-weighted averages, thus suggesting that large clients benefit not only from economies of scale but from strictly superior trading terms. The tighter execution conditions are likely not driven by lower information leakage as the largest clients also make significantly less use of two-sided quote requests. Such a behaviour is consistent with a reduced concern about information leakage when bargaining power is high.

B. *Trade Size and Client Characteristics*

Client-level averages mask important within-client variation. Tables V to **??** hint at higher transaction costs for larger trades while Table VII indicates that larger clients not only enjoy better conditions but also are the group that systematically executes the largest individual trades. In

Figure 7 we illustrate the regression results by plotting the average execution costs by trade deciles, whereby we allocate the same total volume and total number of trades per decile in Panels A and B, respectively. The figure largely confirms the finding that transaction costs are increasing in the trade size. However, the differences in Panels A and B point to the fact that the increase in transaction costs is largely driven by a relatively small number of relatively large trades. In terms of number of trades deciles, transaction costs very much U-shaped and only increase from decile 7 to decile 10, whereas in terms of total volume, transaction costs monotonically increase starting with decile 2.

[INSERT FIGURE 7 HERE]

To disentangle these competing forces, we provide a double sort on client and trade size in Table VIII. As in Table VII, we first sort on client size in deciles. Then, we consider quartiles for within-decile trade sizes.

[INSERT TABLE VIII HERE]

Across the various trade size quartiles, transaction costs are largely decreasing with client size, mirroring the main result from Table VII. Within client size deciles, we observe two types of patterns. For the largest clients (i.e., deciles 8 to 10), we observe a U-shaped pattern whereby the largest trades carry higher execution costs compared to those in the second quartile. Thereafter, however, transaction costs are increasing with smaller trade size. For most deciles with smaller clients, we observe decreasing transaction costs with individual trade size. In sum, the regression results in Tables V to **??** do not fully carry over and seem to be driven mainly by the fact that for large clients, the large trades are those that are most expensive. This is consistent with inventory risk and pre-hedging concerns that scale with trade size. For smaller clients this is apparently not a concern as their "large" trades are still very small overall.

[INSERT TABLE IX HERE]

Finally, we explore the fact that we can distinguish between corporate and financial institution clients in Table IX. While the clients in our sample are predominantly corporates (which is a special feature of our data that distinguishes our paper from other contributions), we do still observe a sizable fraction of financial institutions that trade on the platform as clients, especially among the largest clients. One might expect that these clients are more sophisticated compared to

corporates, and, thus should be able to negotiate better conditions for their trades. We find this to be mostly the case: financial institutions obtain significantly lower transaction costs compared to corporate clients for quote requests sent to up to 10 dealers. For the 11 to 20-dealer bin, however, the situation is reversed. That said, even though financial institutions account for less than 25% of the overall volume, they account for almost 60% of single-request trades. On the other hand, corporates account for a larger share of requests sent to many dealers.

C. Dealer Size, Pricing, and Revenue Concentration

We now turn to the supply side of the market. Again we group dealers into deciles based on their total notional traded volume (or the volume they "win" through their participation in the platform by providing quotes in response to requests) on the platform. Analogous to Table VII, we show dealer size and trade characteristics in Table X. Liquidity provision is highly concentrated although slightly less so compared to client volume. The top decile of dealers intermediates nearly 70% of total trading volume, reflecting the "core–periphery" structure of global FX markets. Dealers in the bottom half of the distribution account for only a negligible share of executed volume, despite responding to a sizable number of RFQs. On average, they can execute less than half of the quotes they provide. For larger dealers, the percentage of won trades drops to under 25%, implying that even though large dealers quote aggressively, they can execute less than one in four quotes they provide as a response to an RFQ.

[INSERT TABLE X AND FIGURE 8 HERE]

In Figure 8 we visualise the results from Table X and show transaction costs as a function of the dealer size in Panel A and total dealer revenue in Panel B. Execution costs decline (almost) monotonically with dealer size. Clients obtain the tightest prices from the most active dealers, whose costs are roughly one-fifth to one-third of those offered by dealers in the lowest decile (considering equal- and volume-weighted transaction costs, respectively). Despite offering the lowest costs, dealer revenue is almost exponentially increasing in dealer size: the volume of the business the large dealers win offsets the lower costs they charge. That is, the large fraction of the revenue the large dealers generate is not due to wide spreads, but instead they seem to be better able to internalize order flow, hedge more efficiently, and operate at scale. Smaller dealers, by contrast, win relatively few trades and must quote wider spreads to compensate for higher inventory and funding costs.

Taken together, the evidence shows that both sides of the FX RFQ market exhibit strong size-related heterogeneity. Large clients benefit from relationship capital and bargaining power, leading to uniformly tighter pricing. Large dealers, in turn, dominate liquidity provision and use their scale and superior risk management capabilities to offer better execution while still capturing the bulk of total revenue. At the same time, the higher costs of individual large trades demonstrate that inventory risk and pre-hedging concerns remain relevant in this market, even for the largest and most sophisticated participants.

IV. Transparency vs. Secrecy in Quote Requests

So far, we have focused on quotes that align with the ultimate direction of the trade—that is, quotes requested in the same direction as the eventual trade, regardless of whether they are provided as part of a one- or a two-sided quote request. However, our data allows to further differentiate as nearly one-third of all requests are requests-for-market (or RFMs), i.e., they include both bid and ask quotes. A request-for-market means that the client can conceal the true trading intention and limit potential information leakage. In addition, it allows to maintain greater bargaining flexibility and preserve anonymity in the face of potential strategic pricing by dealers. Overall, the prevalence of such behavior underscores a key tension in dealer–client trading: while clients seek competitive pricing through transparency, they must simultaneously manage the information revealed in the process. By focusing on the differences between one- and two-sided quote requests, we can study information about the direction of trade affects transaction costs.

Panel A in Figure 9 suggests that there is not much to gain from requesting a RFM from one or more dealers as the unconditional mean of transaction costs is virtually the same in both cases. Panel B compares the within effect once we control for currency pair and, most importantly, client fixed effects. There is a slight drop in transaction costs once the client requests two-sided quotes (significant at the 10% level), suggesting that there might be some value to keeping the intended direction of trade confidential at the quote request stage.

[INSERT FIGURES 9 AND 10 HERE]

In Figure 10 we repeat the analysis from Figure 3 but split up into average transaction costs for one-sided and RFM quote requests. Doing so reveals that the low average transaction costs for single-quote requests are entirely driven by one-sided requests whereas single-dealer RFM quotes are associated with much higher transaction costs. At the same time, a two-sided request

results in lower transaction costs for all except the most extreme RFQ bins. This finding further suggests that the low average transaction costs for single-dealer requests are really driven by relationship trading where the client doesn't see a need to hide the intended direction of trade. For trades where clients are not willing to be transparent with the dealer, competition works as expected and transaction costs are monotonically decreasing in dealer competition as measured by the number of dealers approached in a RFQ. Figure 11 shows that the negative relationship between transaction costs and dealer competition for RFMs holds for both corporate clients as well as financial institutions. On the other hand, also observe the increase in transaction costs for single-sided requests as we move from a single-dealer relationship trade to a two-dealer quote request.

[INSERT FIGURE 11 AND TABLE XI HERE]

In Table XI we formalise this insight in a regression where we restrict the sample to single- and two-dealer quote requests only. Furthermore, we allow for dummies for request-for-market trades as well as for two-dealer requests and we estimate the following specification:

$$\text{tx_cost}_i = \beta_0 + \beta_1 \,\text{TwoQuoteRequest}_i + \beta_2 \,\text{RFM}_i + \beta_3 \,(\text{TwoQuoteRequest}_i \times \text{RFM}_i) + \gamma X_i + \varepsilon_i, \quad (2)$$

where TwoQuoteRequest = 1 if the RFQ contains two aligned quotes (and 0 for a single quote) and RFM = 1 for a request-for-market. β_1 then measures the difference between two-quote and single-quote requests among non-RFM trades; β_2 measures the difference between RFM and non-RFM trades among single-dealer quote requests; and β_3 captures how the effect of requesting an additional quote differs in RFM trades.

In Column (1), which includes no fixed effects, the raw correlations suggest that both two-dealer requests and RFM trades are more expensive with β_1 and β_2 coefficient estimates equal to 1.23 and 0.7, respectively. However, the coefficient on the interaction term is -1.13, hence the multi-dealer quote request premium largely disappears for RFM trades: the net difference between two- and one-dealer quote requests within RFM trades is only $1.23 - 1.13 = 0.10$ basis points.

Once we introduce fixed effects in Columns (2)–(6), the interpretation changes substantially. After absorbing pair×day conditions and various requester and provider fixed effects, the coefficient on the two-dealer quote request dummy becomes negative (around -1.1 to -1.5 bps across specifications). This implies that, holding market conditions and client/dealer characteristics constant, requesting a second quote reduces markups by roughly one basis point. In contrast, the

coefficient on RFM becomes small and statistically insignificant, indicating that RFM trades are not systematically more expensive once selection is controlled for. The interaction term turns positive (around +0.5 to +0.7 basis points), implying that although competition lowers markups, the magnitude of that reduction is smaller for RFM trades. In other words, the competitive benefit of a second quote is attenuated—but not eliminated—when the trade involves multiple directions.

Overall, the comparison between Column (1) and Columns (2)–(6) reveals a clear selection pattern. In the raw data, two-dealer quote requests and RFM trades appear more expensive. After conditioning on detailed fixed effects, however, requesting an additional quote is associated with lower markups, and RFM per se does not raise costs. The positive association in Column (1) therefore reflects composition effects rather than a causal pricing premium.

In sum, the results in this section suggest that hiding the intended direction of the trade may help in certain situations, namely when requesting quotes from multiple dealers. However, a RFM request does not help with transaction costs for single-dealer quote requests, i.e., in situations where a tight client-dealer relationship and the implied transparency when asking for a one-sided quote seems to outweigh the potential benefit of hiding the direction of the trade. The evidence thus points to a coexistence of relationship-based and competition-based pricing regimes within the electronic RFQ market. While we study the pricing effects of quote competition in Section II we now proceed to further investigate the effects of trading relationships in Section V.

V. Dealer Relationships

In Figure 3 we show that single-quote requests result in lower transaction costs than quote requests from two dealers. Figure 10 highlights that this effect is entirely driven by one-sided requests, i.e. requests where the client is not only asking only a single dealer but where the client is also completely transparent about the intended direction of the trade. These results suggest that relationships are key to understanding the pricing in OTC markets. In this section, we explore this further.

A. Single Quotes and Multiple Dealers

Figure 12 shows that just over half of the clients that request *only* single quotes trade with a *single* dealer (this fraction is above 70% for institutional clients). The remaining 47% of clients that request single-quotes do so from more than one dealer over time. That is, over time, they still trade with more than one dealer, reflecting an additional dimension of competition that is

not revealed if single-quote requests are not dissected further. Finally, single-quote requests can also be made by clients that only occasionally ask a single dealer for a quote.

[INSERT FIGURE 12 AND TABLE XII HERE]

In Table XII we split up transactions costs for single-quote requests exactly among these three categories: (i) clients that always direct single-quote requests to the same dealer, (ii) clients that only request single quotes but from multiple dealers over time and (iii) clients that engage in single- as well as in multi-quote requests. Furthermore, we split the sample by client type (i.e. corporate clients and financial institutions) and whether the requests are one-sided or RFM. Regardless of the dimension of comparison, transaction costs are almost always monotonically decreasing when moving from single-quote request/single-dealer clients to single-quote request/multiple dealer clients and finally to the single-quote requests of clients that also use multiple quote requests. These results hold even if we relax the definition of the "single-dealer" relationship in Panel B where we allow the relationship to reset every year and in Panel C where we lower the threshold to 95% of all client transactions to be executed with a single dealer to qualify for the "single-dealer" category.

On average, transaction costs are more than 50% lower when single-quote requests are directed at multiple dealers and the differences are highly statistically significant for all comparisons. The biggest transaction cost improvements are realised for corporates issuing one-sided requests: transaction costs drop by over 60% when moving from a single to multiple dealers and by almost 80% for single-quote requests by clients that switch between single- and multiple-quote requests. The analysis in the table and highlighted in Figure 13 further allows to refine the insights from Figures 10 and 11: one-sided requests of corporate clients that only ever trade with a single dealer are considerably higher than transaction costs for other one-sided single-quote requests.

[INSERT FIGURE 13 AND TABLE XIII HERE]

In Table XIII we separate transaction costs within the single-quote segment based on transaction size as well as based on whether clients approach just one or multiple dealers over time. Within both categories, we recover the results from Table VIII where transaction costs are decreasing with the transaction size except for the largest quartile of trades.

B. Request for Quotes and Trade Executions

The previous section reveals a second layer of competition that does not only operate at the RFQ level but before requests are even sent out. That is, even though clients may send a particular request only to a single dealer, there is ex ante competition in the fact that for a sizable fraction of transactions, multiple dealers are considered for that single-quote request. We extend the analysis of the relationship between quote requests and subsequent trade executions by constructing a client-level measure of dealer participation across the full RFQ sample. For every client that doesn't only submit single-quote requests (i.e., all clients not included in Figure 12), we compute (i) the number of unique dealers from whom quotes are requested and (ii) the number of unique dealers with whom trades are ultimately executed. Figure 14 displays the resulting heatmap that measures the level of concentration in execution compared to how many dealers are contacted at the RFQ level.

[INSERT FIGURE 14 HERE]

Overall, execution is substantially more concentrated than quote solicitation. Overall, the modal outcome is on the diagonal—clients most often trade with as many dealers as they request quotes from over time. Excluding the single-quote request clients, the percentages vary between 34% and 46%. However, that also means that in up to two thirds of the cases, clients approach more dealers for quotes than they ultimately transact with. For example for the group of clients that trade with a total of only two dealers, 58% request quotes from more than two dealers over time and almost 20% will interact with five or more dealers. Overall, the patterns indicate that ex post dealer concentration does not imply limited ex ante competition. Furthermore, while looking at the number of quote requests is a natural choice to gauge competition, it may also be informative to consider how the sample of dealers at which the quote requests are direct changes over time and, possibly, from request to request.

C. Transaction Costs at the Client and Trade Level

Next we plot separate transaction costs for the various RFQ bins for the two client segments (corporate clients and financial institutions) and small vs. large clients in Figure 15 and small vs. large trades in Figure 16. The increase in transaction costs when moving from single- to multiple-quote requests is due to financial institutions as well as large corporate clients. However, the effect obtains for both small as well as large trades.

[INSERT FIGURES 15 AND 16 HERE]

Finally, we revisit Figure 3 at the client level, measuring competition by the number of dealers clients request quotes from on average. In Figure 17 we plot the average transaction costs based on the median number of requested quotes in a RFQ for a given client. Again, the single-quote request clients are special and we distinguish between those that transact with only one dealer and those that transact with multiple dealers. For all other clients, we sort them according to the median number of dealers they request quotes from over all their transactions.

[INSERT FIGURE 17 AND TABLE XIV HERE]

Measured at the client level, transaction costs no longer exhibit a hump-shaped pattern but are monotonically decreasing in the number of quotes requested. Table XIV provides additional information for the same sorting procedure based on the median number of dealers in a RFQ. Single-quote request clients that transact with more than a single dealer work with an average of just over four dealers (the median is three), the same as for clients that approach two dealers in a RFQ. There is a monotonic relationship between the median number of dealers approached in a RFQ and the average and median dealers transacted with. Both average as well as median concentration (as measured by the Herfindahl index) are monotonically decreasing in the number of dealers approached. Finally, the same holds for transaction costs, which decline sharply from 1.8 bps in the least competitive bin of multi-quote request clients to roughly 0.2 to 0.4 bps in the most competitive bins.

D. Client-Dealer Networks

Finally, we further explore the effects of client-dealer networks and the dealer concentration across clients and transactions. Figure 18 shows average dealer concentration at the client level across different trading-network categories, measured by the volume-weighted Herfindahl–Hirschman Index (HHI) computed from executed trades. Clients are classified into bilateral relationship clients (D=1), single-quote request clients who trade with multiple dealers over time (D>1), and multi-dealer RFQ clients grouped by the median number of aligned quotes in executed RFQs. The figure reports mean HHI within each category, with error bars indicating the interquartile range. By construction, dealer concentration is highest for relationship clients and naturally remains high for the other single-quote request clients, indicating persistent reliance on a small set of dealers even when clients rotate counterparties across trades. In contrast, concentration declines sharply and

monotonically as the number of aligned quotes increases, implying that the greater effective competition in executed RFQs is associated with substantially more diversified client-dealer trading networks.

[INSERT FIGURE 18 AND TABLE XV HERE]

In Table XV we report client-level summary statistics by quintiles of dealer concentration, measured by the volume-weighted Herfindahl–Hirschman Index (HHI). The table reveals a sharp and monotonic relationship between dealer concentration and trading outcomes. Clients in the least concentrated quintile (Q1) account for the vast majority of aggregate trading volume and trade with a large and diversified set of dealers, resulting in the lowest transaction costs. As concentration increases across quintiles, both the breadth of dealer relationships and trading frequency decline markedly, while reliance on a single dealer rises sharply. This shift is accompanied by a steady increase in transaction costs, with clients in the most concentrated quintile (Q5) paying more than three times the average costs of clients in the first quintile. Notably, differences in trade size alone cannot account for this pattern: while trade sizes vary across quintiles, the dominant change is in dealer dependence, as reflected in the top-dealer volume share approaching unity in the highest quintile. Taken together, the results suggest that greater diversification of client–dealer trading networks is associated with substantially lower execution costs, consistent with competition and outside-option effects playing a central role in pricing.

Finally, we conduct a double sort of average transaction costs based on (i) a client's dealer-network concentration and (ii) the median number of dealers approached in a RFQ at the client level. That is, clients are first sorted into terciles based on their volume-weighted Herfindahl–Hirschman Index (HHI) across executed trades, where "Low HHI" denotes a diversified trading network and "High HHI" indicates strong dealer dependence (high concentration). Within each HHI group, clients are then binned by the average number of dealers involved in their RFQs, separating out single-quote request clients trading with only one dealer (where HHI=1 by construction), single-quote request clients trading with multiple dealers and multi-quote RFQ clients by the median number of dealers approached in a RFQ. Each entry for transaction costs reports the average value in basis points, i.e., within cells clients are assigned equal weights.

[INSERT TABLE XVI HERE]

The results reveal two clear patterns. First, holding dealer dependence fixed (within an HHI tercile), transaction costs decline sharply as RFQs involve more dealers, consistent with stronger

competitive pressure improving execution prices: for example, within the High-HHI group average costs fall from 2.8 bps for D=1 to 2 bps for RFQs sent to two dealer and to around 0.4 bps for RFQs sent to 11 or more dealers. Second, holding RFQ competition fixed (within a dealer-count bin), costs are higher for clients with more concentrated trading networks, indicating that dealer dependence is associated with systematically worse pricing. This is most pronounced in for the low-competition bins (e.g., D>1 rises from 1.8 bps for Low HHI to around 2.4 bps for Medium/High HHI).

VI. Conclusion

This paper provides the first comprehensive examination of request-for-quote dynamics in the FX spot market, using a proprietary dataset that records every quote and execution from one of the largest multi-dealer trading platforms. We show that execution costs are jointly shaped by competition, information management, and the heterogeneous characteristics of both clients and dealers. A central finding is the non-monotonic relationship between the number of requested quotes and transaction costs: bilateral relationship trades obtain the tightest pricing, costs rise sharply when a second dealer is introduced, and competitive forces dominate only once a sufficiently large dealer set is engaged. This pattern highlights the coexistence—within the same electronic venue—of relationship-driven bilateral execution and competitive multi-dealer search.

We also show that clients strategically manage information by choosing whether to reveal or conceal trade direction. Concealing direction yields modest pricing benefits in competitive settings, consistent with concerns about information leakage and strategic quote shading. Yet the same behavior is counterproductive in bilateral RFQs, where transparency allows dealers to internalize flow, reduce hedging needs, and provide better prices. These results demonstrate that information design is an active margin of choice in FX markets, and that its value depends crucially on the surrounding competitive structure.

Client and dealer size introduce a further layer of heterogeneity. Large clients, who deliver steady and substantial order flow, obtain markedly lower transaction costs, reflecting both bargaining power and the economic value of durable trading relationships. On the supply side, a small group of global dealers intermediates the majority of RFQ activity. These institutions quote the tightest prices, achieve high hit ratios, and generate substantial revenue through scale economies and effective internalization. Smaller dealers, by contrast, win fewer trades and quote wider spreads, illustrating the importance of scale for liquidity provision in OTC markets.

Finally, we show that clients with more diverse and less concentrated dealer relationships experience lower transaction costs. In particular, single-quote request clients that trade with multiple dealers over time obtain significantly better terms than clients that always trade with the same dealer.

Taken together, our results reveal an FX RFQ ecosystem in which relationship capital, competitive search, information frictions, and balance-sheet constraints jointly determine transaction costs. These findings challenge canonical OTC trading models that rely on either pure search or pure bilateral bargaining, and instead point to the need for theoretical frameworks that integrate competition, information leakage, inventory and funding constraints, as well as economies of scale.More broadly, the results shed new light on the microstructure foundations of global currency markets and open several avenues for future research on the interaction of market design, dealer behavior, and international funding conditions.

VII. 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.

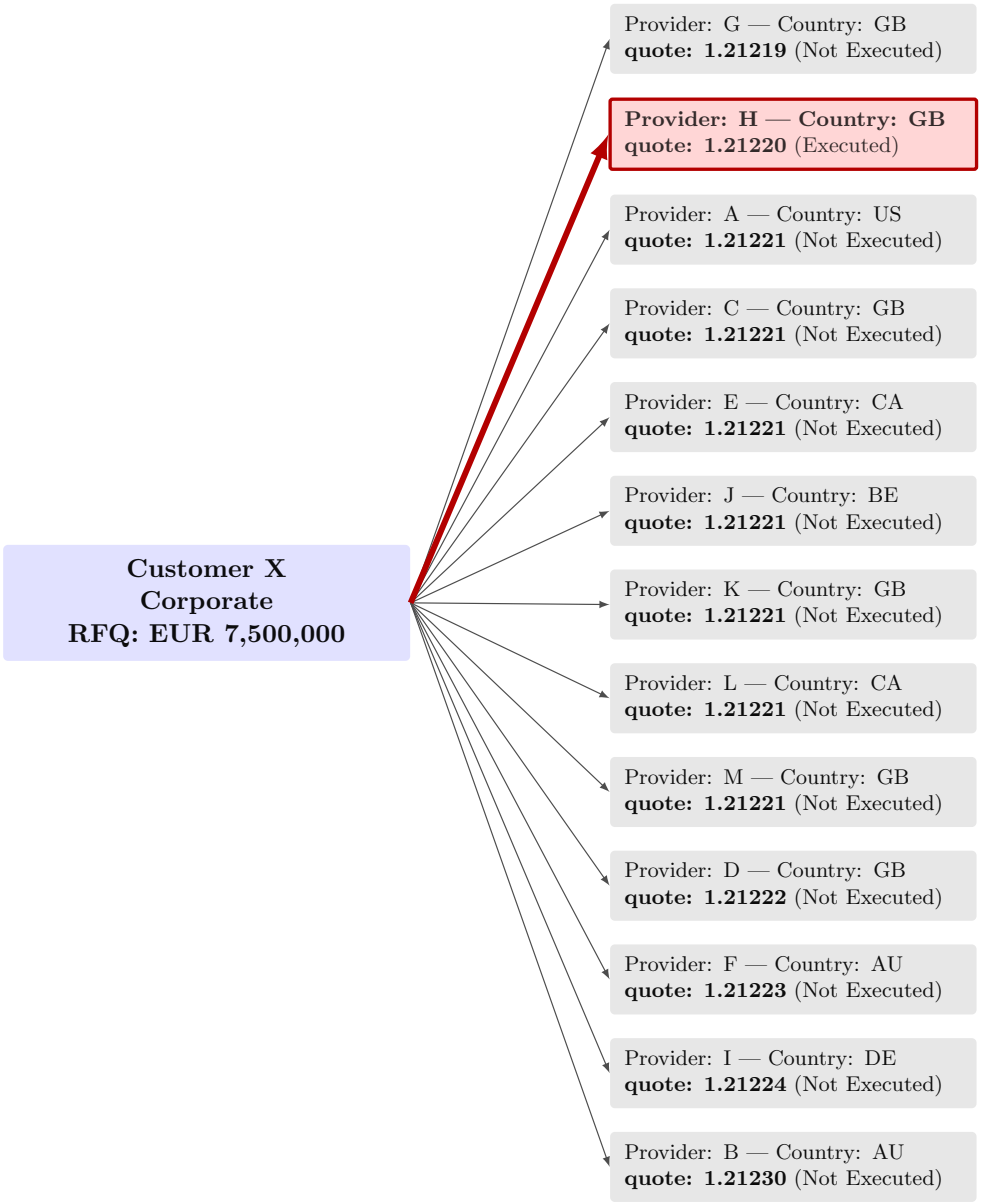


Figure 2. Distribution of quote requests.

This figure shows the distribution of RFQs by the number of dealers requested per trade. Single-dealer requests (bin = 1) 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 bin = 1) 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., 3.87

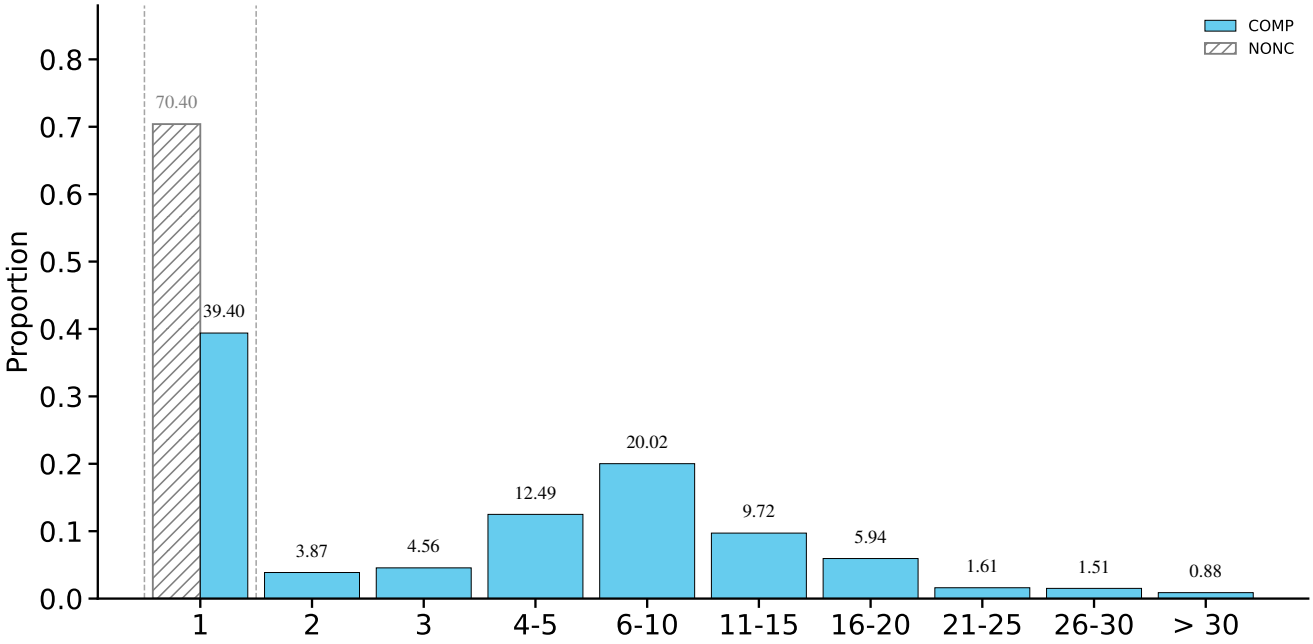


Figure 3. Quote requests and transaction costs.

This figure shows how average transaction costs (in basis points) vary with the number of dealer quotes requested in an RFQ, grouped into bins based on the total number of quotes. Bars report equally weighted (EW) average transaction costs within each bin. The sample period is January 2018 to July 2022.

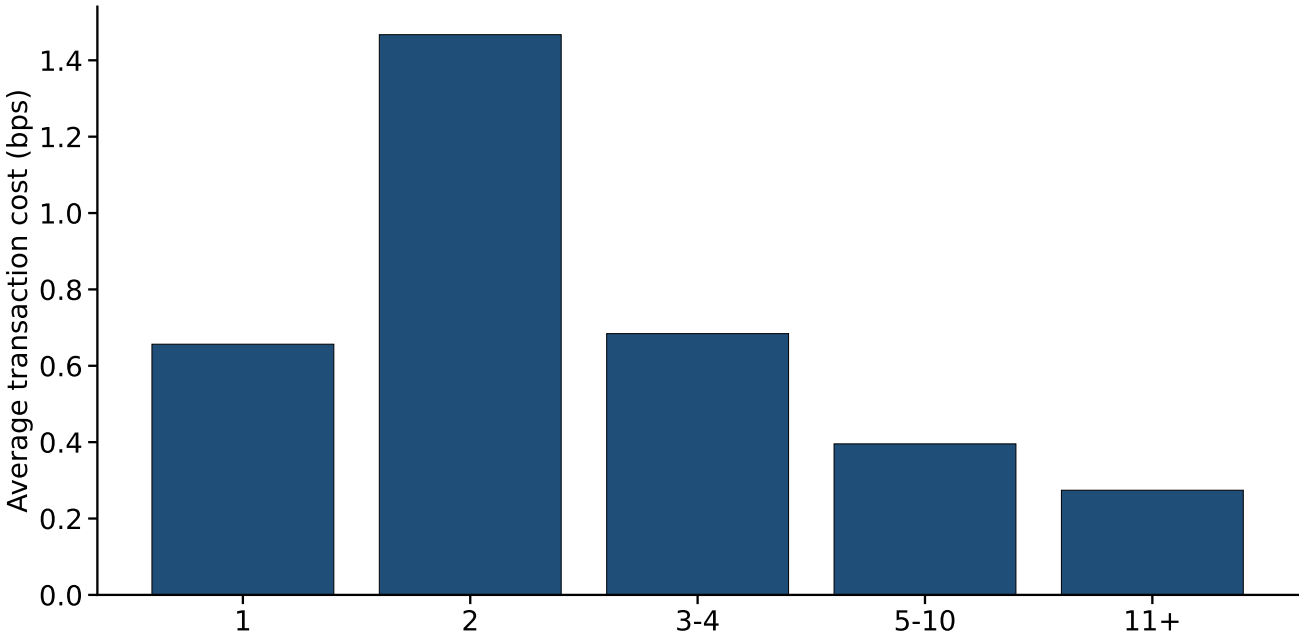


Figure 4. Quote requests and transaction costs: Top-decile clients

The figure shows average transaction costs (in basis points) by the number of dealers participating in the RFQ, grouped into bins, for the top decile of clients. Clients are ranked by traded notional within each subperiod. Bars report equally weighted (EW) average transaction costs within each bin. The sample period is January 2018 to July 2022.

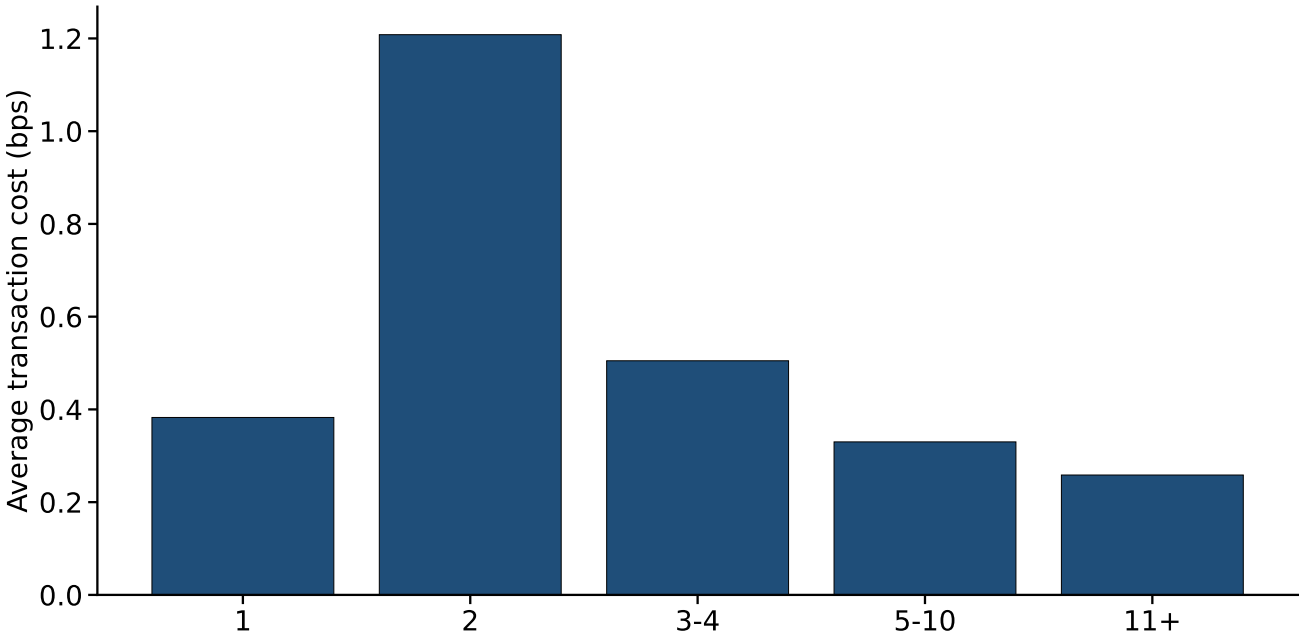


Figure 5. Transaction costs and the depth of dealer quotes in RFQs.
This figure shows how average transaction costs (in basis points) vary with the number of dealer quotes requested in an RFQ, grouped into bins based on the total number of participating dealers (1, 2, 3–4, 5–10, and 11+). For each bin, bars report equally weighted (EW) average transaction costs computed relative to the traded price ($k = 1$) and relative to the k-th best directional quote in the RFQ (k=2, 3, 5, and 11). The sample period is January 2018 to July 2022.

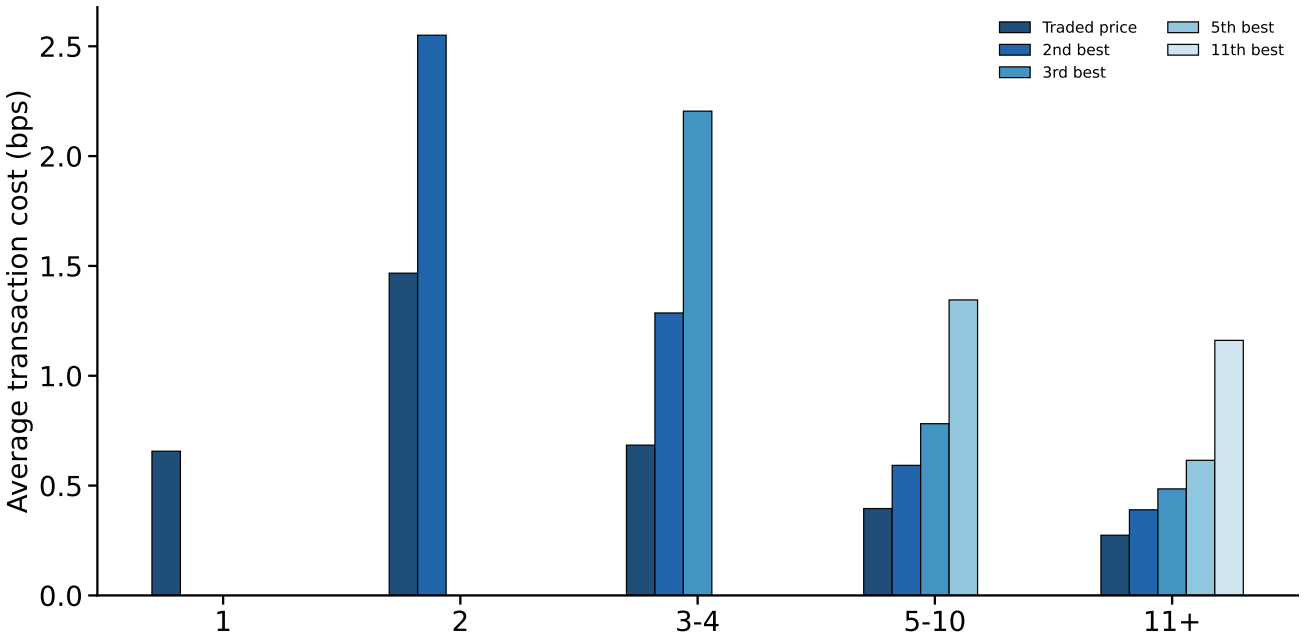


Figure 6. Transaction costs and client size.

This figure shows the relationship between transaction costs (in basis points) and client size, measured by total trading volume. Clients are sorted into deciles within each subperiod based on their traded notional amount. Decile-level statistics are averaged across regimes to account for ID reassignments. The solid lines report equally-weighted (EW) and volume-weighted (VW) average transaction costs by client decile. The bars show the share of total trading volume accounted for by each decile. The sample period is January 2018 to July 2022.

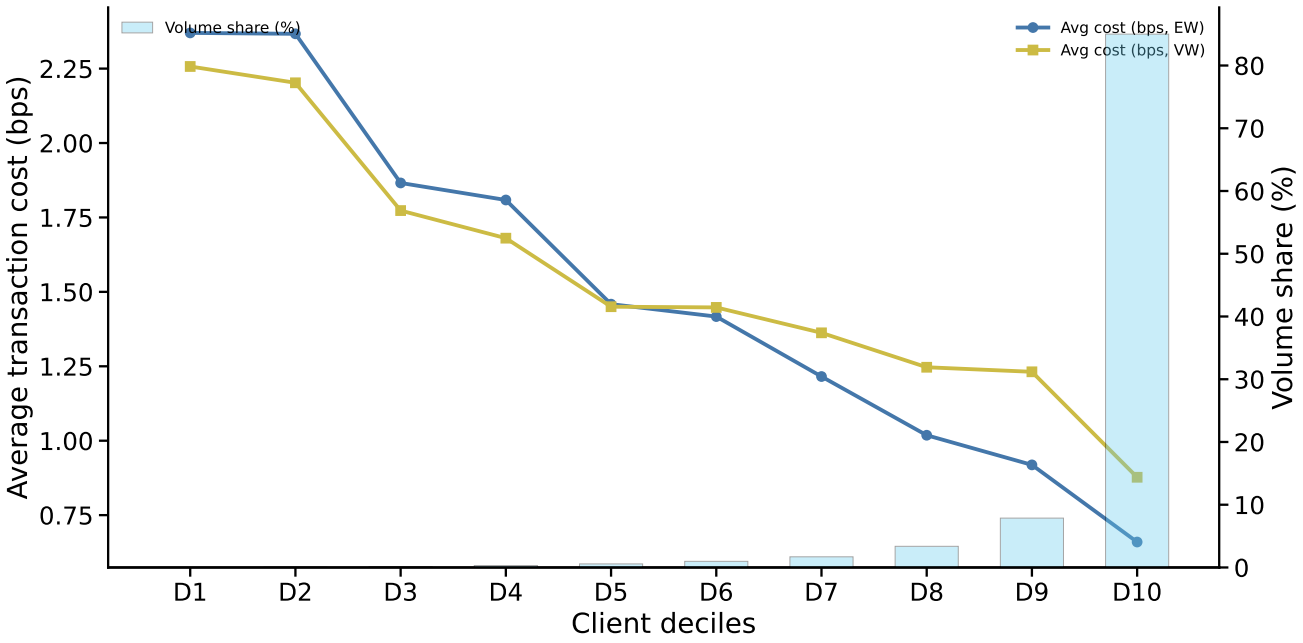
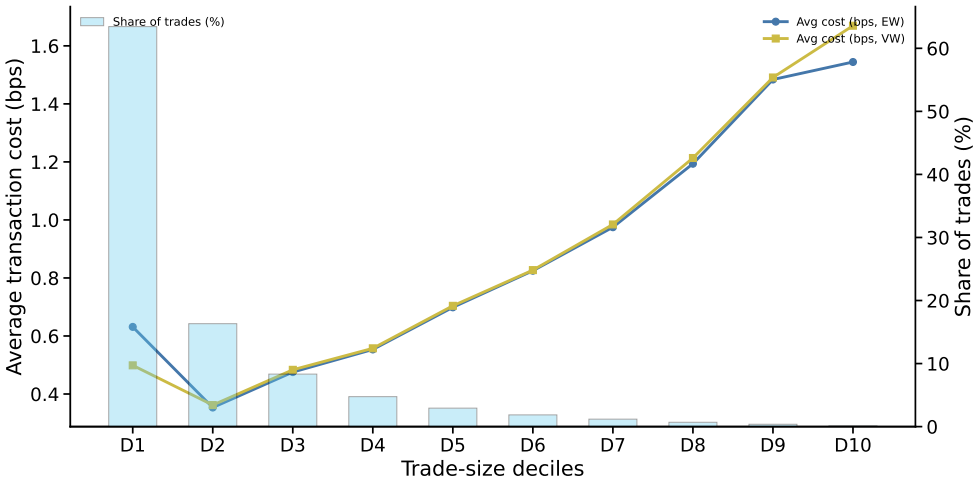


Figure 7. Costs and transaction size.

This figure shows average transaction costs across the trade-size distribution under two alternative binning approaches. Panel A reports results for equal-volume trade-size deciles, where each bin contains approximately the same share of total notional volume but a different number of trades. The bars show each bin's share of number of trades (in percent), while the lines display equal-weighted and volume-weighted average transaction costs in basis points. Panel B reports results for equal-count trade-size deciles, where each bin contains an equal number of trades. The bars represent each decile's share of total notional volume, and the lines show equal-weighted and volume-weighted average transaction costs. The sample period is January 2018 to July 2022.

Panel A. Equal Volume per Decile



Panel B. Equal Number of Trades per Decile

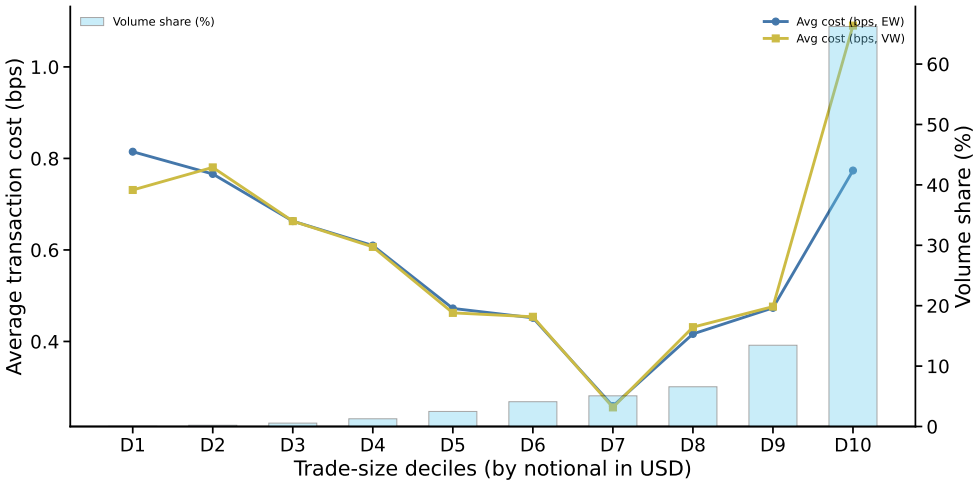
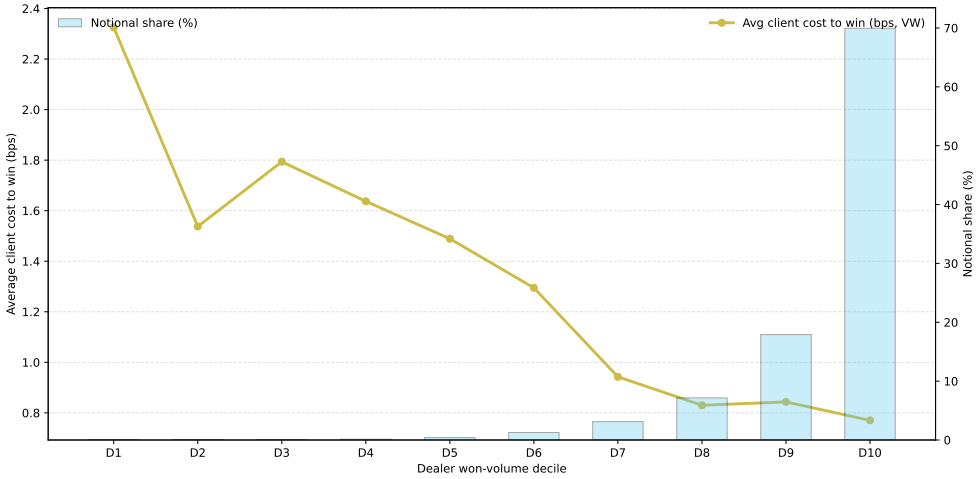


Figure 8. Dealer size, transaction costs, and dealer revenue.
This figure shows how transaction costs and dealer revenue vary with dealer activity. Panel A plots average transaction costs (in basis points relative to the mid-price) across dealer won-volume deciles. Within each subperiods, dealers are sorted into deciles based on executed notional volume; decile-level statistics are computed within subperiods and then averaged across regimes. The solid line reports the volume-weighted mean transaction cost for executed trades, while the bars show each decile's share of total executed notional volume. Panel B plots total gross dealer revenue (USD millions) across the same dealer activity deciles. Gross revenue is measured as the transaction-cost-implied markup relative to the mid-price multiplied by executed notional. The solid line reports total gross revenue, while the bars report each decile's share of total gross revenue. The sample is January 2018 to July 2022.

Panel A. Transaction Costs



Panel B. Revenues

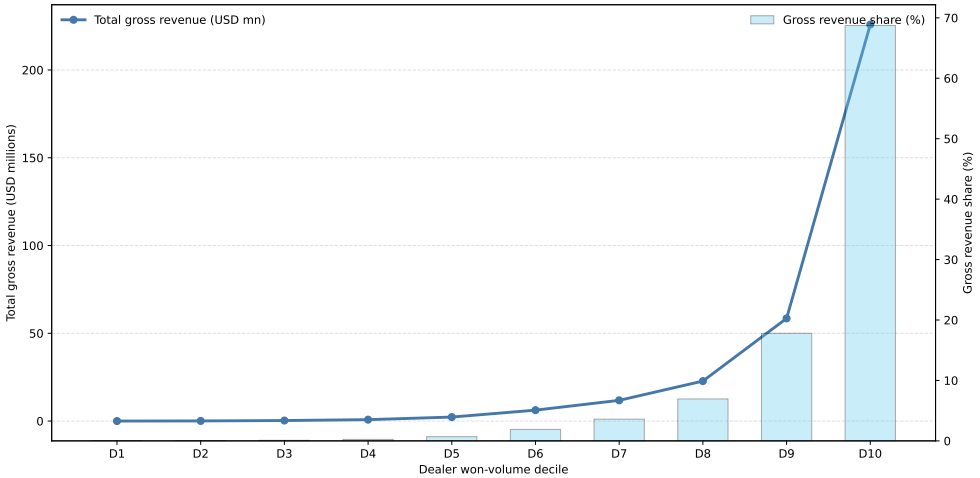
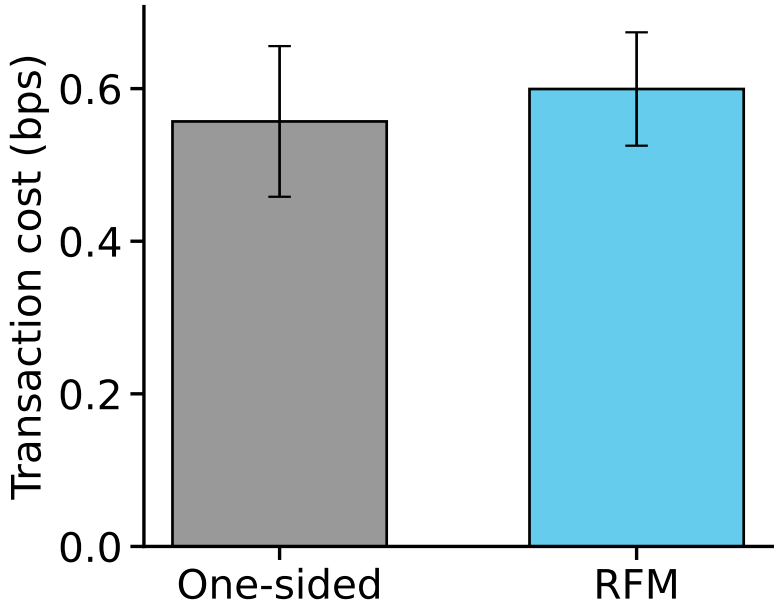


Figure 9. Average transaction costs for one-sided vs. RFM requests.
This figure reports regression-adjusted average transaction costs (in basis points) for one-sided and Request-for-Market (RFM) quote requests. Panel A reports marginal means from a regression of transaction costs on an indicator for RFM requests, controlling for trade size, trade direction, and contemporaneous bid-ask spreads. Panel B augments this specification with currency-pair-by-month, requester-by-month, and provider-by-month fixed effects. Standard errors are clustered at the requester level. The sample period is January 2018 to July 2022.

Panel A. Across Requests



Panel B. Within Requests

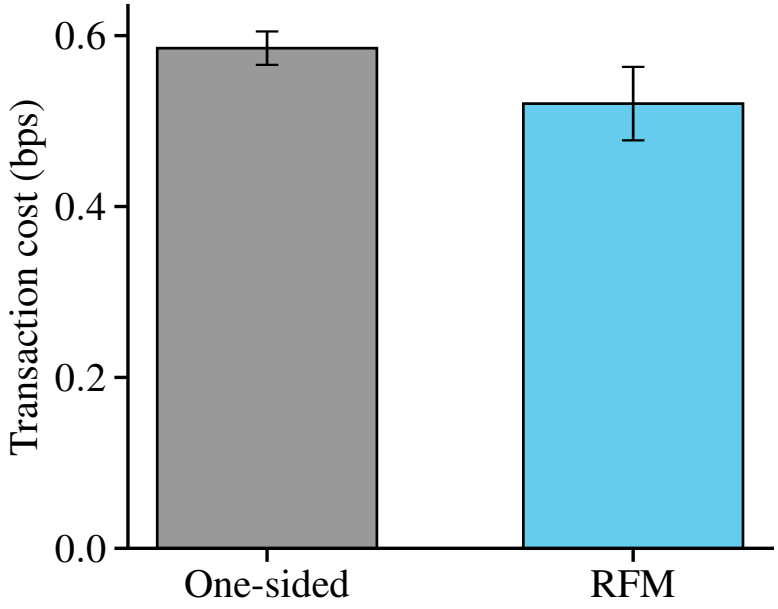


Figure 10. One-sided vs RFM quote requests across RFQs.
This figure shows how average transaction costs (in basis points) vary with the number of dealer quotes requested, separately for one-sided RFQs and request-for-market (RFM) RFQs. The x-axis groups trades into bins based on the total number of quotes in the RFQ. Bars report equally weighted (EW) average transaction costs within each bin. The sample period is January 2018 to July 2022.

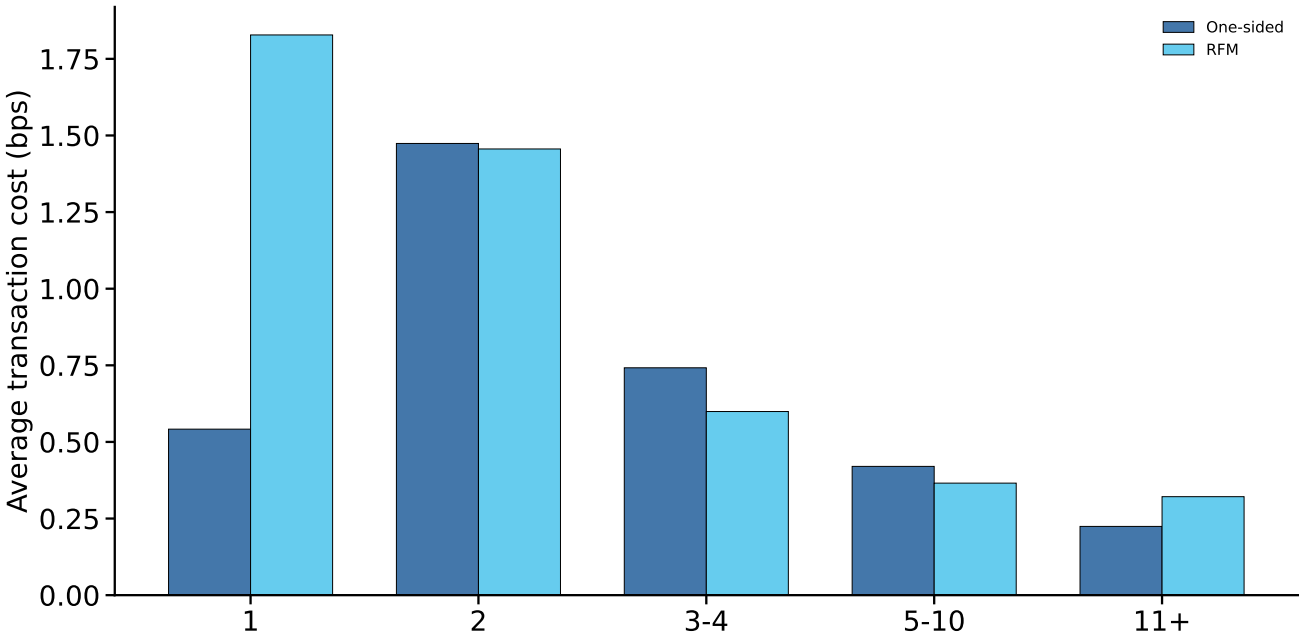


Figure 11. Transaction costs and RFQ competition by client type.
This figure shows how average transaction costs (in basis points relative to the mid-price) vary with the number of dealer quotes requested. RFQs are grouped into bins based on the total number of quotes. Bars report equally weighted (EW) average transaction costs for the full sample, corporate clients, and institutional clients. Panel A shows results for one-sided RFQs, in which the client requests quotes on a single side of the market. Panel B shows results for two-sided (RFM) RFQs, in which the client simultaneously requests buy and sell quotes. The sample period is January 2018 to July 2022.

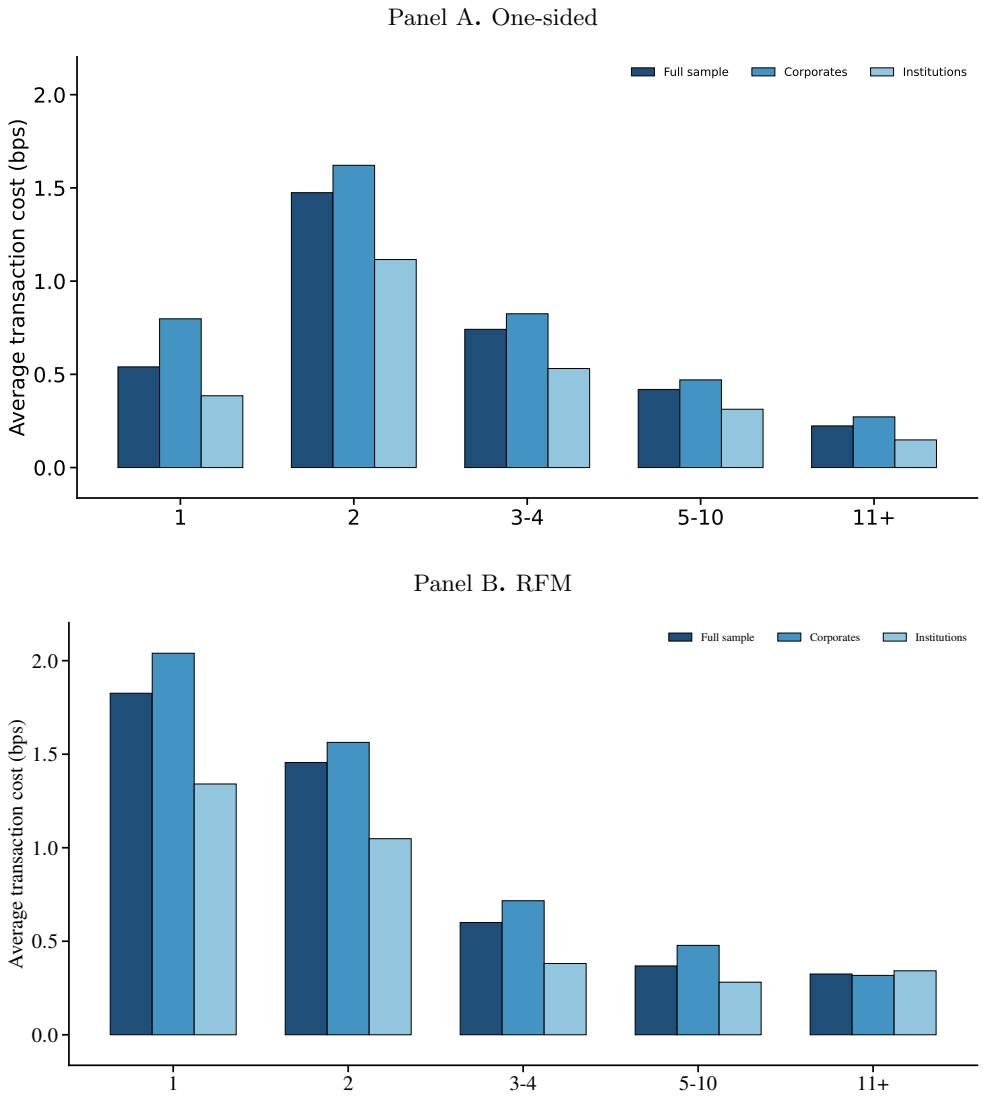


Figure 12. Requested dealers in single-quote requests.

This figure shows the distribution of requested dealer connections for single-quote request clients. The bars represent the density of clients across different categories of requested dealers. The last category (10+) consolidates all clients with 10 or more requested dealers. The figure differentiates between the full sample, corporate clients (Corp), and institutional clients (Inst), with each group shown in distinct colors: dark blue for the full sample, light blue for corporate clients, and light green for institutional clients. The sample period is January 2018 to July 2022.

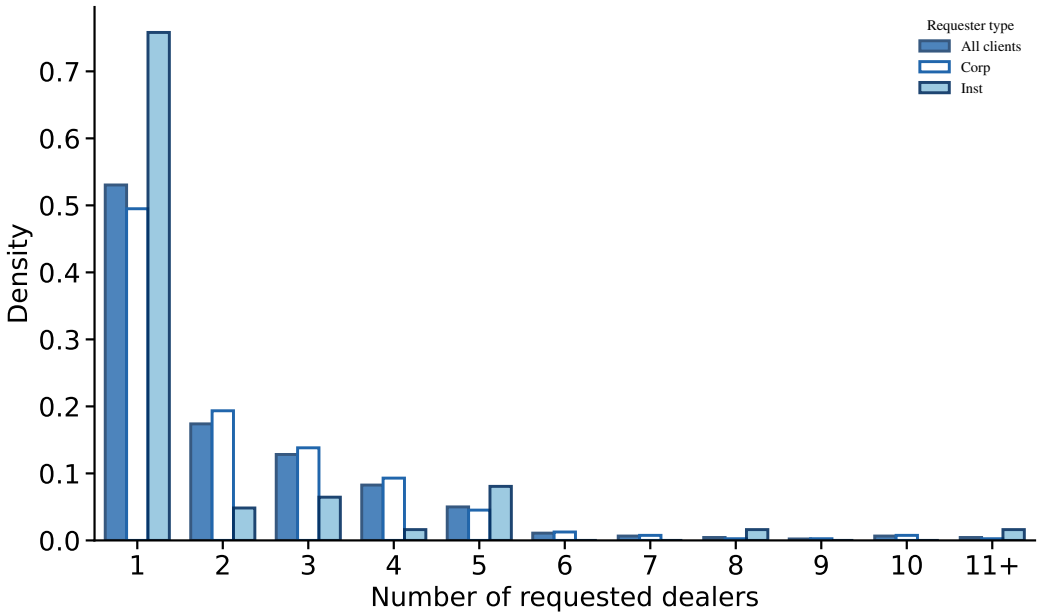


Figure 13.

Transaction costs and dealer relationships across RFQ types and client segments. This figure shows how average transaction costs (in basis points) vary with the number of dealer quotes requested in an RFQ and with the client's dealer-relationship structure. The x-axis groups RFQs into bins based on the total number of quotes requested. For the single-quote bin (1), bars split RFQs by the client–currency-pair dealer-relationship category computed within subperiods: $D = 1$ (the client trades with only one dealer), $D > 1, S$ (multiple relationships but RFQs are always single-dealer requests), and $D > 1, M$ (multiple relationships and the client sometimes requests multiple dealers). For all bins each bar reports the equally weighted mean transaction cost within the bin. Panels A, C, and E report one-sided RFQs, while Panels B, D, and F report request for market (RFM) RFQs. Panels A–B use the full sample, Panels C–D show results for corporate clients (CC), and Panels E–F show results for financial institutions (FI). The sample period is January 2018 to July 2022.

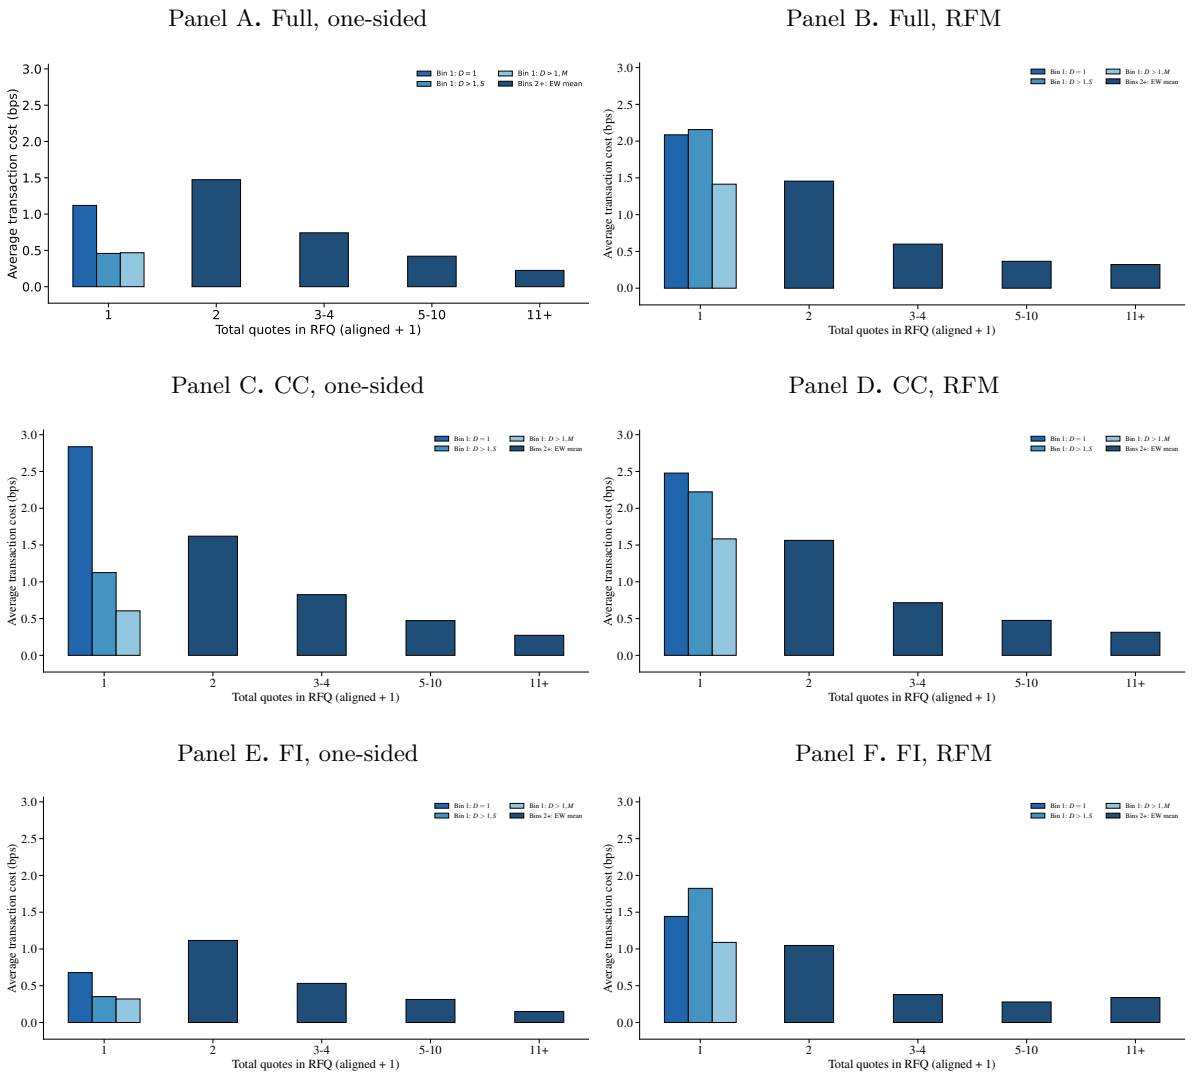


Figure 14. Requests and transactions.

The figure shows the joint distribution of clients' requested and trading dealer connections. Each cell reports the percentage of clients falling into a given combination of the number of dealers they requested quotes from (horizontal axis) and the number of dealers with whom they executed trades (vertical axis). Values of 11 or more dealers are grouped into the category "11+." Percentages are normalized row-wise, so each row sums to 100%. The bottom row is special as it contains only clients that submit only single-quote requests: the percentages are calculated based on how many dealers are approached for a quote over time. See Figure 12 for the description. The sample period is January 2018 to July 2022.

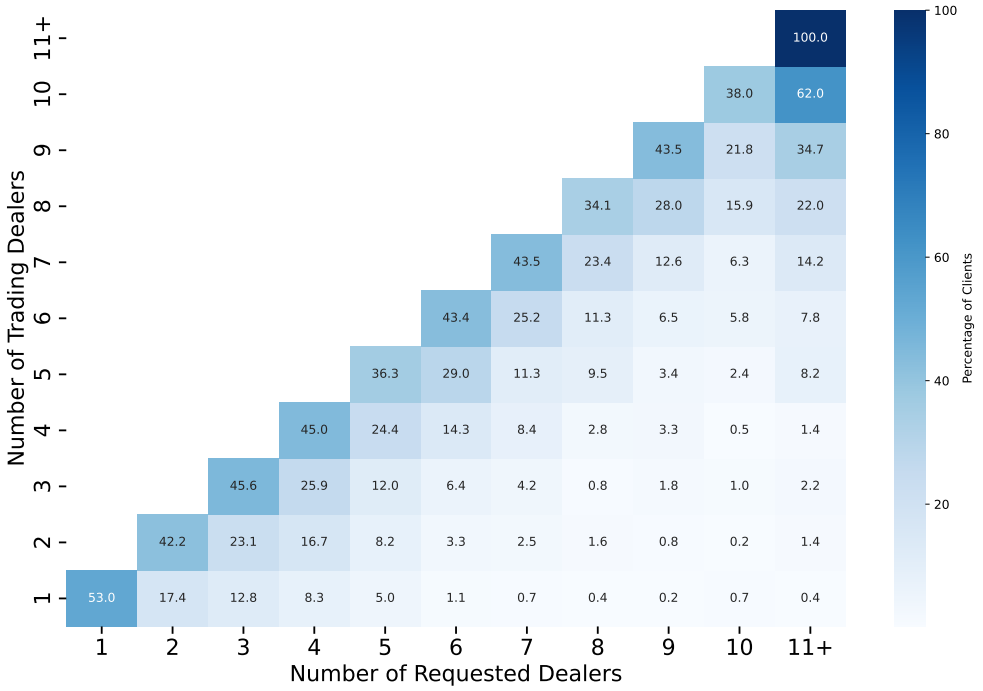


Figure 15. Transaction costs, RFQs, and client size.

This figure shows how average transaction costs (in basis points) vary with the total number of dealer quotes requested in an RFQ, separately for one-sided RFQs and request-for-market (RFM) RFQs, and for small versus large clients. In each panel, bars report equal-weighted mean transaction costs within quote-count bins, shown separately for the full sample, corporate clients, and financial institutions. Client size is defined based on total executed trading volume (USD notional) within each subperiod. Specifically, for each subperiod, clients are ranked by their total USD notional traded, and classified as a small client (SC) if their volume is below or equal to the median within that regime, and as a large client (LC) if their volume is above the median. The sample period is January 2018 to July 2022.

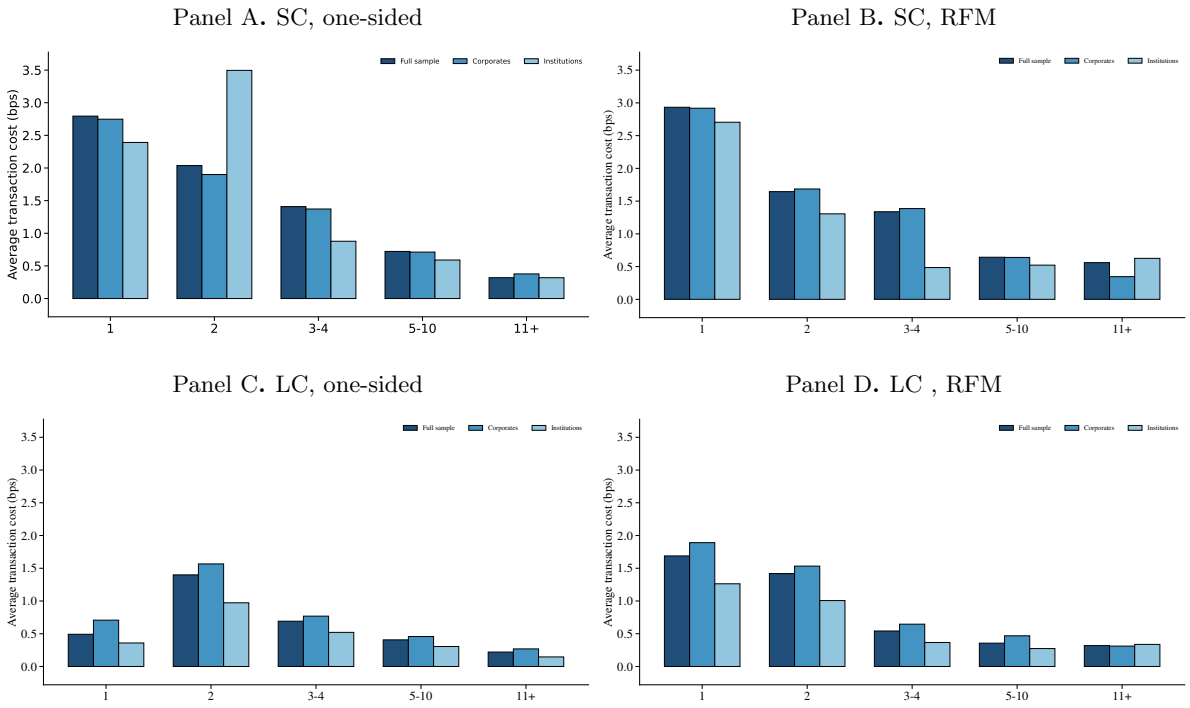


Figure 16. Transaction costs, RFQs, and trade size.

This table shows how average transaction costs (in basis points) vary with the total number of dealer quotes requested in an RFQ, separately for one-sided RFQs and Request for Market (RFM) RFQs, and for small versus large trades. Reported values are equal-weighted mean transaction costs within quote-count bins shown separately for the full sample, corporates, and financial institutions. Trade size is defined based on the USD notional of each executed transaction. Within each subperiod, trades are classified as a small trade (ST) if the USD notional is below or equal to the median notional in that subperiod, and as a large trade (LT) if the USD notional is above the median. The sample period is January 2018 to July 2022.

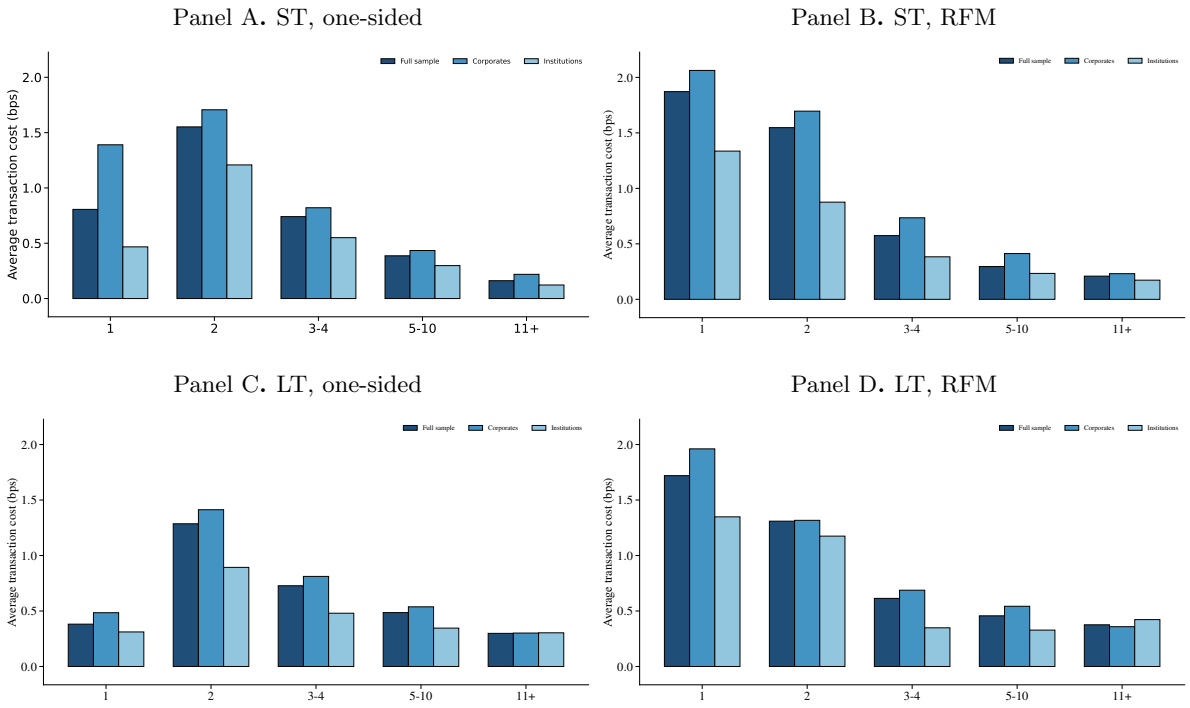


Figure 17. Transaction costs by trading relationships.
The figure shows the average transaction costs based on actual trading relationships at the client level. For single-quote request clients we distinguish between clients that always trade with the same dealer (D=1) and those that trade with multiple dealers (D>1). For clients that submit multiple-quote requests, we group them according to the median number of quotes requested. Error bars denote the interquartile range (25th–75th percentiles). The sample period is January 2018 to July 2022.

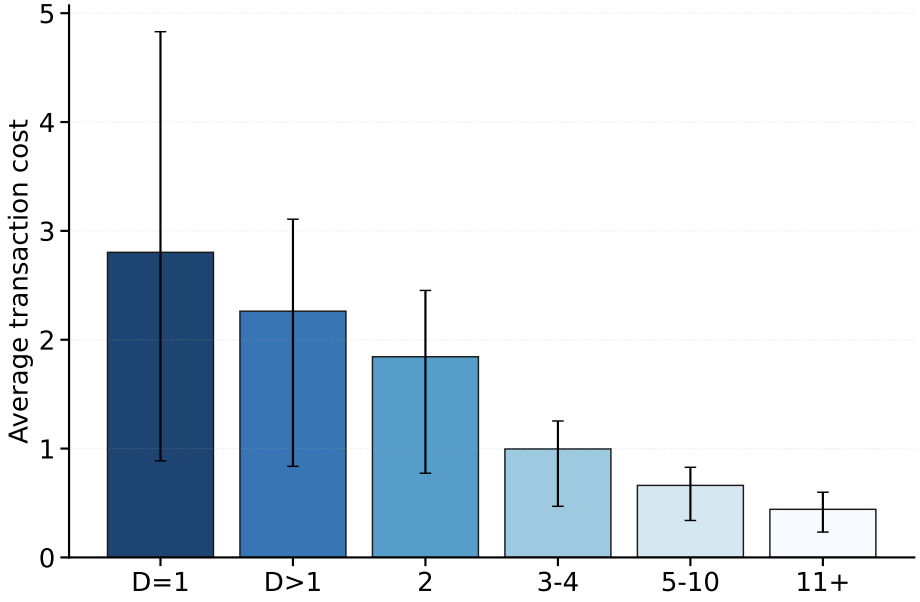
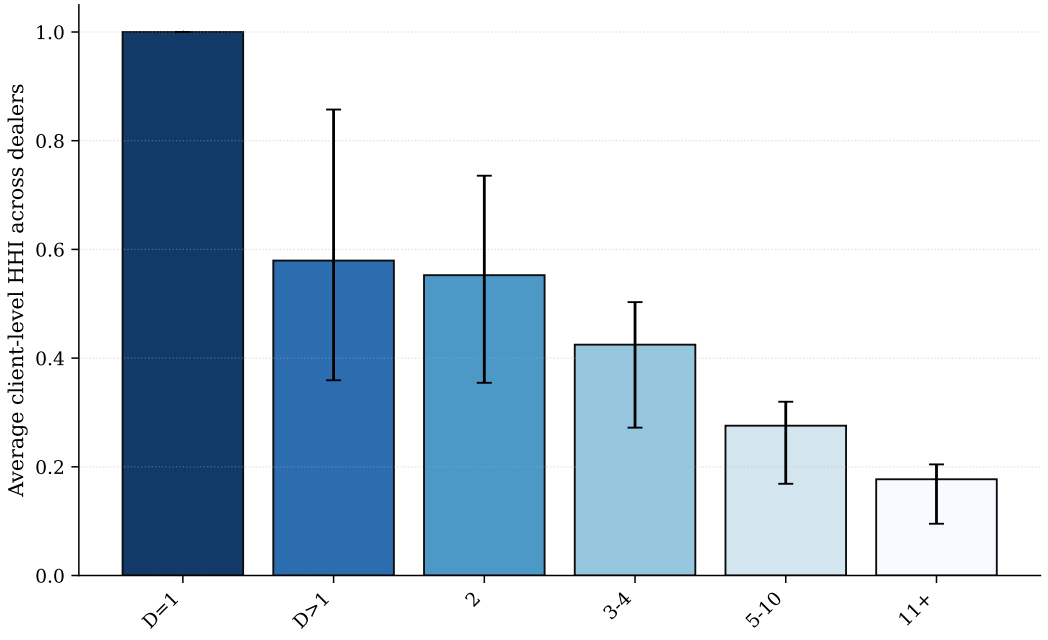


Figure 18. Average dealer concentration.

The figure shows average client-level dealer concentration, measured by the volume-weighted HHI across dealers, for seven client categories: bilateral relationship clients (D=1), single-quote multi-dealer clients (D>1), and multi-dealer RFQ clients grouped by the typical number of aligned quotes (2, 3–4, 5–10, 11+). Error bars report the interquartile range of HHI within each category. The sample period is January 2018 to July 2022.



VIII. Tables

Table I
Summary Statistics

The table reports summary statistics at the currency-pair level. For each currency pair, we report the total number of observations (in millions, #Obs), the number of executed trades (in millions, #T), and the number of non-executed quotes (in millions, #Q). We also report the average number of unique providers (#Prov.), unique requesters (#Req.), and unique provider–requester links (#Links). These counts are computed separately within each of the three subperiods (R1–R3) and then averaged across subperiods. In addition, the table reports the number of unique countries in which providers are located (#C). Finally, we report the average executed trade size in U.S. dollars (Not., in USD millions), as well as the average share of financial institutions (%FI) and corporate clients (%CC). The sample period spans January 2018 to July 2022.

	#Obs	#T	#Q	#Prov.	#C	#Req	%FI	%CC	#Links	Not.	#Obs	#T	#Q	#Prov.	#C	#Req	%FI	%CC	#Links	Not.
All	17.75	2.01	15.75	263	43	1298	9	91	9855	2.05										
	USD-Pairs										EUR-Pairs									
AUD	0.55	0.08	0.47	161	34	294	16	84	2311	1.31	0.29	0.03	0.26	151	29	342	10	90	2751	1.29
CAD	0.85	0.09	0.75	160	32	358	14	86	2792	1.42	0.28	0.02	0.26	147	26	316	11	89	2585	2.54
CHF	0.67	0.14	0.54	167	35	324	18	82	2583	2.13	0.87	0.10	0.77	173	34	528	10	90	4072	2.45
DKK	0.13	0.01	0.12	131	27	144	16	84	1157	1.03	0.18	0.02	0.16	135	24	324	9	91	2529	1.44
EUR	4.34	0.47	3.87	229	37	1001	9	91	7416	3.20										
GBP	1.63	0.20	1.43	182	34	554	12	88	4172	1.72	1.59	0.16	1.42	189	35	751	7	93	5599	2.16
HKD	0.32	0.04	0.28	129	26	163	19	81	1327	2.00	0.09	0.01	0.08	117	23	181	13	87	1630	2.05
HUF	0.11	0.01	0.10	106	23	93	19	81	840	0.91	0.30	0.02	0.28	125	25	300	6	94	2603	1.43
ILS	0.05	0.01	0.05	80	21	61	25	75	486	1.25	0.02	0.00	0.02	71	19	70	12	88	557	0.97
JPY	0.58	0.09	0.49	159	31	265	20	80	2144	2.10	0.45	0.04	0.41	154	29	340	12	88	2697	1.78
MXN	0.53	0.06	0.47	138	25	238	17	83	1994	1.21	0.23	0.02	0.22	106	21	176	10	90	1507	1.25
NOK	0.20	0.02	0.18	135	26	172	21	79	1354	1.60	0.27	0.03	0.24	144	26	329	11	89	2678	1.56
NZD	0.11	0.02	0.09	126	31	128	24	76	1060	0.83	0.06	0.01	0.05	118	24	136	14	86	1212	0.82
PLN	0.22	0.02	0.20	124	25	167	15	85	1460	1.27	0.63	0.05	0.59	146	27	434	7	93	3462	1.64
RON	0.03	0.00	0.03	85	19	73	9	91	584	1.39	0.46	0.02	0.44	95	19	232	4	96	1957	1.45
SEK	0.32	0.03	0.30	147	30	250	18	82	1876	1.50	0.55	0.05	0.50	156	28	450	9	91	3447	1.82
SGD	0.27	0.05	0.22	139	30	232	13	87	1860	1.14	0.14	0.01	0.12	125	26	223	9	91	1864	1.15
THB	0.02	0.00	0.02	92	22	83	14	86	613	1.38	0.02	0.00	0.02	82	22	104	5	95	712	1.32
ZAR	0.29	0.06	0.23	137	30	182	19	81	1401	0.75	0.10	0.01	0.09	125	25	211	9	91	1758	1.06

Table II
Decomposition of Requester Quoting Behaviour

The table reports summary statistics on client quoting behaviour. Clients are classified as: (i) *Single-dealer only*—RFQs sent exclusively to one dealer; (ii) *Mixed*—RFQs sent to one or to multiple dealers across requests; and (iii) *Multi-dealer only*—RFQs always sent to multiple dealers. *Clients* is the number of distinct clients in each category; $\bar{T}$ is the average number of executed trades per client; and $\bar{V}$ is the average notional per trade (USD millions). The sample period is January 2018 to July 2022.

	USD-Pairs									EUR-Pairs								
	Single Only			Mixed			Multiple Only			Single Only			Mixed			Multiple Only		
	Clients	$\bar{T}$	$\bar{V}$	Clients	$\bar{T}$	$\bar{V}$	Clients	$\bar{T}$	$\bar{V}$	Clients	$\bar{T}$	$\bar{V}$	Clients	$\bar{T}$	$\bar{V}$	Clients	$\bar{T}$	$\bar{V}$
AUD	53	197	1.66	76	161	3.87	165	25	3.03	62	27	0.91	74	73	2.13	206	13	2.50
CAD	58	90	1.71	95	219	3.24	205	24	3.51	57	16	1.80	68	65	2.04	191	14	2.80
CHF	56	141	1.98	86	342	2.92	181	27	3.87	79	68	2.14	151	133	2.73	298	23	3.42
DKK	19	18	1.05	26	82	2.83	99	17	1.87	52	16	1.62	56	43	1.91	215	11	2.60
EUR	135	128	2.77	400	279	5.79	466	46	5.69									
GBP	78	88	2.38	164	265	5.63	311	45	4.45	115	58	1.88	212	166	3.74	424	26	3.53
HKD	29	34	0.88	36	182	2.95	99	44	3.68	27	14	0.76	30	52	2.13	124	11	2.06
HUF	11	31	0.62	19	39	1.01	63	23	1.27	38	16	0.66	62	39	1.11	200	19	1.66
ILS	11	46	1.30	16	80	5.02	33	15	3.06	13	26	0.57	10	14	1.48	46	9	1.55
JPY	44	225	2.10	71	163	3.64	151	39	3.18	63	49	2.29	74	110	2.22	203	13	2.57
MXN	31	87	2.57	66	207	2.33	141	27	2.34	27	5	2.81	36	77	2.03	113	24	1.99
NOK	33	51	1.35	33	59	2.79	106	30	2.70	53	68	1.30	66	52	1.62	210	13	2.11
NZD	25	124	1.29	28	117	1.78	75	16	1.72	22	19	0.67	22	29	0.63	92	9	1.31
PLN	20	19	0.58	30	51	2.06	116	28	1.82	65	18	1.00	97	100	2.37	273	18	2.53
RON	11	2	0.93	13	17	1.12	50	12	1.25	30	6	0.99	46	71	1.68	155	22	1.83
SEK	34	49	0.94	56	65	1.91	159	26	1.96	69	46	2.38	111	69	1.81	270	16	2.40
SGD	36	50	1.49	53	212	2.53	144	18	2.23	35	9	2.55	45	54	1.83	143	11	1.83
THB	17	11	1.52	17	30	1.52	49	9	2.42	20	5	1.69	18	21	1.67	65	7	1.07
ZAR	32	209	0.97	47	199	1.22	103	19	1.58	41	17	0.58	44	29	1.54	125	11	1.35

Table III

Transaction Costs Across Currencies

The table reports average transaction costs conditional on the number of quotes received per request, ranging from 1 quote (1Q) to 10 or more quotes (11+). Averages reflect equal-weighted averages. The columns "5–10Q – 2Q" ("11+Q – 2Q") report the average difference (Δ) in transaction costs between the respective quote buckets, along with the p-value (p-val) from a difference-in-means test. Column PT displays the p-value from a Patton–Timmermann monotonicity test. The sample period is January 2018 to July 2022.

	Number of Quotes					5 to 10Q − 2Q		11+ − 2Q		PT
	1	2	3–4	5–10	11+	Δ	p-val	Δ	p-val	p-val
All	0.66	1.47	0.68	0.40	0.27	1.07	0.00	1.19	0.00	0.00
USD pairs	0.55	1.48	0.61	0.34	0.21	1.14	0.00	1.27	0.00	0.00
EUR pairs	0.98	1.45	0.84	0.49	0.40	0.95	0.00	1.04	0.00	0.00
USD G9	0.46	1.38	0.49	0.22	0.13	1.16	0.00	1.24	0.00	0.00
USD EME	1.01	1.83	1.11	0.88	0.55	0.95	0.00	1.28	0.00	0.00
EUR G9	0.77	1.21	0.65	0.37	0.26	0.84	0.00	0.95	0.00	0.00
EUR EME	2.04	1.99	1.35	0.77	0.67	1.22	0.00	1.32	0.00	0.00

Table IV
Quote Competition and Price Dispersion

This table reports RFQ-level regressions results of

$$\text{Disp}_i = \alpha + \beta \log(\#\text{Quotes}_i) + \gamma' \mathbf{X}_i + \delta_{p,d} + \delta_{p,d,h} + \delta_{r,m} + \delta_{r,p,m} + \varepsilon_i, \tag{3}$$

where Disp_i refers to the interquartile range of prices within each request, $\mathbf{X}_i$ includes the log notional trade amount and the bid–ask spread. $\delta_{p,d}$, $\delta_{p,d,h}$, $\delta_{r,m}$, and $\delta_{r,p,m}$ denote, respectively, pair-by-day, pair-by-day-by-hour, requester-by-month, and requester-by-pair-by-month fixed effects. Different columns include different subsets of these fixed effects, as indicated at the bottom of the table. The sample is restricted to competitive trades, aligned quotes within 60 seconds of the trade, and economically relevant quote sides (bid quotes for sell RFQs and ask quotes for buy RFQs). Standard errors are clustered at the requester level. The sample period is January 2018 to July 2022.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log(Quotes)	0.23	0.15	0.15	0.22	0.20	0.19	0.17	0.22
	(4.39)	(4.68)	(3.41)	(4.69)	(3.65)	(3.74)	(2.88)	(5.02)
log(Notional)			-0.23	-0.09	-0.08	-0.08	-0.08	-0.09
			(-2.35)	(-3.63)	(-2.72)	(-3.17)	(-2.73)	(-3.60)
BA-Spread			0.03	0.03	0.03	0.01	0.00	0.05
			(3.28)	(4.40)	(3.94)	(1.88)	(1.43)	(5.45)
Constant	0.45	0.55	3.35	1.44	1.38	1.45	1.44	1.38
	(3.17)	(4.63)	(2.46)	(4.99)	(3.82)	(4.70)	(4.02)	(4.78)
N(mn)	2.01	2.00	2.00	1.99	1.89	1.85	1.74	2.00
R^2	0.00	0.09	0.10	0.50	0.62	0.62	0.69	0.47
Adj. R^2	0.00	0.07	0.08	0.48	0.58	0.56	0.62	0.46
Within R^2	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00
Pair × Day	–	✓	✓	✓	✓	–	–	–
Pair × Day × Hr	–	–	–	–	–	✓	✓	–
Req × Mo	–	–	–	✓	–	✓	–	✓
Req × Pair × Mo	–	–	–	–	✓	–	✓	–

t statistics in parentheses

Table V
Determinants of Transaction Costs

This table reports regressions of transaction costs on a measure of intensity for competitively requested RFQ. The estimated specification is

$$\mathrm{TC}_{i,t} = \alpha + \beta \log(\mathrm{Quotes}_{i,t}) + X'_{i,t}\theta + \mathrm{FE} + \varepsilon_{i,t}, \tag{4}$$

where $\mathrm{TC}_{i,t}$ denotes execution costs in basis points relative to the prevailing mid-quote. The competition measure is the logarithm of the number of aligned dealer quotes received in the RFQ. The control vector $X_{i,t}$ includes the log notional trade amount, and the contemporaneous bid–ask spread. All specifications include high-dimensional fixed effects as reported in the table. Standard errors are clustered two-ways by requester and provider. The sample period is January 2018 to July 2022.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log(Quotes)	-0.27	-0.23	-0.24	-0.21	-0.20	-0.24	-0.27	-0.21
	(-6.74)	(-5.34)	(-6.74)	(-6.02)	(-5.69)	(-6.73)	(-6.74)	(-5.58)
log(Notional)	0.00	0.00	0.01	0.01	0.01	0.01	0.00	0.01
	(0.31)	(0.22)	(1.29)	(1.50)	(1.22)	(0.73)	(0.31)	(0.91)
BA-Spread	0.05	0.05	0.05	0.05	0.05	0.04	0.05	0.05
	(3.99)	(3.83)	(3.77)	(3.91)	(3.88)	(3.64)	(3.99)	(3.64)
Constant	0.89	0.80	0.75	0.71	0.72	0.77	0.89	0.70
	(6.63)	(5.63)	(6.57)	(6.11)	(6.14)	(6.32)	(6.63)	(5.93)
N(mn)	2.17	2.06	2.14	2.16	2.11	2.00	2.17	2.03
R^2	0.53	0.61	0.60	0.57	0.60	0.66	0.53	0.65
Adj. R^2	0.51	0.58	0.56	0.55	0.57	0.60	0.51	0.61
Within R^2	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Pair × Day	✓	✓	✓	✓	✓	–	✓	✓
Pair × Day × Hr	–	–	–	–	–	✓	–	–
Req × Mo	✓	–	✓	✓	–	✓	✓	–
Req × Pair × Mo	–	✓	–	–	–	–	–	✓
Prov × Mo	✓	✓	–	✓	–	✓	✓	–
Prov × Pair × Mo	–	–	✓	–	–	–	–	✓
Req × Prov	–	–	–	✓	–	–	–	–
Req × Prov × Mo	–	–	–	–	✓	–	–	–

t statistics in parentheses

Table VI
Choosing Competition: Within-Client Variation

This table reports estimates of the effect of the competitive protocol (COMP) on client execution costs. The dependent variable is the transaction cost per RFQ, TC_r, measured as the difference between the executed price and the prevailing mid-quote, expressed in basis points. The estimated specification is

$$\mathrm{TC}_r = \alpha + \beta\,\mathrm{COMP}_r + \boldsymbol{\gamma}'\mathbf{X}_r + \mathrm{FE} + \varepsilon_r,$$

where COMP_r equals one for competitively requested RFQs and zero otherwise. The vector $\mathbf{X}_r$ includes the log notional trade amount and the contemporaneous bid–ask spread. Columns (1)–(3) report results for corporate clients, while columns (4)–(6) report results for financial institutions. All specifications include these controls. Columns (1) and (4) include currency-pair-by-day fixed effects; columns (2) and (5) additionally include requester-by-day fixed effects; and columns (3) and (6) instead include requester-by-month fixed effects. Standard errors are two-way clustered at the requester and provider levels. The sample consists of clients who use both competitive and non-competitive protocols and spans January 2018 to July 2022.

	(1)	(2)	(3)	(4)	(5)	(6)
	Corporate Clients			Financial Institutions		
COMP	-1.60	-0.75	-0.64	-0.91	-0.46	-0.44
	(-5.37)	(-20.02)	(-9.45)	(-1.93)	(-9.17)	(-7.29)
log(Notional)	-0.24	-0.06	-0.06	-0.19	-0.00	-0.00
	(-5.43)	(-1.13)	(-1.30)	(-1.92)	(-0.31)	(-0.30)
BA-Spread	0.05	0.06	0.05	0.02	0.02	0.02
	(2.94)	(2.26)	(2.35)	(2.20)	(2.06)	(2.11)
Constant	5.16	2.81	2.80	3.04	0.80	0.80
	(6.66)	(4.74)	(5.14)	(2.25)	(8.34)	(7.97)
N (mn)	0.16	0.15	0.16	0.31	0.31	0.31
R^2	0.521	0.753	0.734	0.500	0.701	0.697
Adj. R^2	0.439	0.679	0.684	0.485	0.689	0.687
Within R^2	0.218	0.019	0.017	0.143	0.006	0.006
Request FE	Pair ×Day	Req ×Day Pair ×Day	Req ×Mo Pair ×Day	Pair ×Day	Req ×Day Pair ×Day	Req ×Mo Pair ×Day

t statistics in parentheses

Table VII

Trade Characteristics by Client Size

This table reports client characteristics and trading-cost measures by client volume decile. *Decile* indexes clients sorted into ten groups by total traded volume (1 = lowest, 10 = highest). *Clients* denotes the number of clients in each decile, and *Vol* is the total traded volume (in USD millions). *TC (Avg)* and *TC (Med)* report the equal-weighted and median transaction costs (in basis points), respectively. $\bar{Q}$ is the average number of dealer quotes per RFQ. *Multi* is the average share of two-sided requests (i.e., requests with quotes on both buy and sell sides). $\bar{T}$ denotes the average number of trades per client within the decile, *Avg Trade Size* is the mean trade size (in USD millions), and % *Corp* is the share of corporate clients in each decile. The sample period is January 2018 to July 2022.

Decile	Clients	Vol	TC (Avg)	TC (Med)	$\bar{Q}$	Multi	$\bar{T}$	Avg Trade Size	% Corp
1	380	0	2.36	2.27	2.60	0.46	5.80	0.37	0.89
2	380	2	2.32	2.14	3.07	0.47	15.82	1.04	0.94
3	379	6	1.92	1.82	3.02	0.47	29.98	1.65	0.94
4	380	13	1.82	1.67	3.54	0.43	47.78	2.16	0.96
5	379	23	1.47	1.48	3.70	0.45	65.73	2.98	0.93
6	380	39	1.33	1.39	4.30	0.43	87.02	3.39	0.94
7	379	69	1.21	1.33	4.85	0.44	131.64	3.82	0.92
8	380	137	1.07	1.30	5.64	0.41	240.44	5.07	0.91
9	379	323	0.93	1.22	6.03	0.38	410.89	8.11	0.91
10	380	3495	0.65	0.89	7.61	0.39	4243.85	7.93	0.77

Table VIII
Client and Transaction Size

The table reports average transaction costs, measured in basis points, sorted by client size (rows) and notional mount (columns). Differences between the largest and smallest groups (Q4 to Q1 in columns) and between the highest and lowest quote bins (D10 to D1 in rows) are shown together with corresponding p-values from Welch two-sample tests (weighted where appropriate). The sample period is January 2018 to July 2022.

	Large	Q3	Q2	Q1	Δ	p-val
D10	0.49	0.28	0.34	0.43	0.06	0.00
D9	0.89	0.67	0.70	1.25	-0.35	0.00
D8	0.96	0.92	1.49	2.05	-1.09	0.00
D7	1.09	1.11	1.14	1.41	-0.32	0.00
D6	1.19	1.33	1.38	1.51	-0.32	0.00
D5	1.29	1.39	1.71	1.60	-0.31	0.00
D4	1.56	1.78	1.84	1.51	0.05	0.19
D3	1.69	2.03	2.05	2.16	-0.47	0.00
D2	2.03	2.37	3.24	3.15	-1.12	0.00
D1	1.99	2.28	3.59	2.75	-0.76	0.00
D10-D1	-1.51	-2.00	-3.25	-2.32	0.81	0.00
p-val	0.00	0.00	0.00	0.00		

Table IX
Quote Competition Across Client Types

The table reports average transaction costs, measured in basis points, across different dimensions of market activity. Transaction costs are calculated as the difference between the executed price and the prevailing mid-quote at the time of trade, expressed in basis points. Reported averages are equally weighted. Within each quote bucket, transaction costs are separated out between financial institution and corporate clients. Differences between the client types (columns) and between the highest and lowest quote bins (rows) are shown together with corresponding p-values from Welch two-sample tests. The sample period is January 2018 to July 2022.

	Institutions	Corporates	Δ	p-val	N(Institutions)	N(Corporates)
1	0.43	0.99	-0.56	0.00	471,467	319,357
2	1.09	1.60	-0.50	0.00	20,119	57,494
3-4	0.46	0.78	-0.32	0.00	67,359	150,537
5-10	0.29	0.47	-0.18	0.00	225,697	298,697
11+	0.23	0.30	-0.06	0.00	135,974	258,618
11+ -1	-0.20	-0.69	0.50	0.00		
p(11+ - 1)	0.00	0.00				
n(11+)	135,974	258,618				
n(1)	471,467	319,357				

Table X

Trade Characteristics by Dealer Size

This table reports dealer activity, pricing, and market share by deciles of won volume. Dealers are ranked by their total *won* notional and split into deciles (D1 = least active, D10 = most active). Columns: **Dealer** is the number of distinct dealers in the decile; Vol is the aggregate notional of client trades won by those dealers (billions); Rev is a proxy for gross revenue, computed as the product of notional and trading costs (expressed in millions); TC (EW) and TC (VW) are equal- and volume-weighted average client transaction costs (bps) across wins; Win is the average RFQ share of won requests; Vol (Tot) is the decile's share of total executed notional in the sample. The sample period is January 2018 to July 2022.

Decile	Dealer	Vol	Rev	TC (EW)	TC (VW)	%Win	%Vol (Tot)
1	71	0.04	0.01	2.60	2.65	75.71	0.00
2	71	0.45	0.08	1.89	1.80	48.06	0.02
3	70	1.72	0.32	2.03	1.96	40.29	0.09
4	71	5.09	1.00	1.75	1.89	45.12	0.28
5	71	15.33	2.37	1.54	1.53	45.40	0.66
6	70	48.00	6.39	1.20	1.36	33.61	1.79
7	71	123.50	12.46	0.87	1.02	39.17	3.48
8	70	281.63	26.22	0.81	0.94	36.16	7.33
9	71	719.30	63.47	0.74	0.89	33.15	17.75
10	71	2912.44	245.35	0.53	0.83	23.89	68.60

Table XI

Hiding Direction: Single vs. Multi Quote Requests

We estimate the following specification:

$$\text{tx_cost}_i = \beta_0 + \beta_1 \,\text{TwoQuoteRequest}_i + \beta_2 \,\text{RFM}_i + \beta_3 \,(\text{TwoQuoteRequest}_i \times \text{RFM}_i) + \gamma X_i + \varepsilon_i, \quad (5)$$

where TwoQuoteRequest = 1 if the RFQ contains two aligned quotes (and 0 for a single quote), RFM = 1 for a request-for-market, and X_i includes the log notional trade amount and the contemporaneous bid-ask spread. The omitted category is therefore single-quote, single-direction trades. Columns differ in their fixed-effects structure. Standard errors are two-way clustered by requester and provider. The sample period is January 2018 to July 2022 and is restricted to single-dealer and two-dealer quote requests only.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Constant	3.27	1.11	1.10	1.14	1.11	1.05	
	(6.96)	(5.72)	(5.09)	(5.82)	(5.64)	(4.97)	
Two-dealer quote request	1.23	-1.34	-1.53	-1.17	-1.12	-1.13	
	(4.77)	(-3.84)	(-3.47)	(-3.24)	(-2.93)	(-3.47)	
RFM	0.70	-0.10	0.06	-0.00	0.02	0.15	
	(3.48)	(-0.43)	(0.25)	(-0.01)	(0.09)	(0.52)	
Two-dealer quote request × RFM	-1.13	0.66	0.74	0.51	0.47	0.47	
	(-3.79)	(2.12)	(1.95)	(1.67)	(1.47)	(1.46)	
log(Notional)	-0.23	-0.00	-0.01	-0.01	-0.01	-0.01	
	(-7.41)	(-0.29)	(-0.43)	(-0.62)	(-0.63)	(-0.72)	
BA-Spread	0.14	0.06	0.06	0.06	0.06	0.05	
	(8.17)	(3.21)	(3.22)	(3.17)	(3.23)	(2.47)	
N (mn)	0.96	0.94	0.92	0.94	0.92	0.83	
R^2	0.084	0.634	0.686	0.668	0.682	0.734	
Adj. R^2	0.084	0.608	0.651	0.641	0.656	0.679	
Within R^2	.	0.010	0.012	0.009	0.010	0.008	
Pair×Day	–	✓	✓	✓	✓	–	
Pair×Day×Hr	–	–	–	–	–	✓	
Req×Mo	–	✓	✓	✓	–	✓	
Req×Pair×Mo	–	–	–	–	–	–	
Prov×Mo	–	✓	–	✓	–	✓	
Prov×Pair×Mo	–	–	✓	–	–	–	
Req×Prov	–	–	–	✓	–	–	
Req×Prov×Mo	–	–	–	–	✓	–	

t-statistics in parentheses

Table XII
Transaction Costs for Single-Quote Requests

The table reports summary statistics of transaction costs for executed single-dealer requests for quotes (RFQs), defined as trades in which only one dealer participated in the request. N denotes the total number of single-dealer RFQs in each subsample. Clients are classified according to their dealer relationship within a given definition. $D = 1$ identifies client–pair relationships that involve only one dealer, with $N^{D=1}$ and $TC^{D=1}$ reporting the number of observations and mean transaction costs (in basis points), respectively. $D > 1$ refers to client–pair relationships involving multiple dealers across requests. Within this group, we distinguish between two disjoint subgroups. S ("single-request") denotes clients who always submit single-dealer RFQs within the relevant regime, with $N^{D>1,S}$ and $TC^{D>1,S}$ reporting the corresponding counts and mean transaction costs. M ("multi-dealer") denotes clients who sometimes submit multi-dealer RFQs within the same regime; $N^{D>1,M}$ and $TC^{D>1,M}$ are computed using only their single-dealer RFQs. N are reported in thousands. Transaction costs (TC) are measured relative to the prevailing mid-quote and expressed in basis points. p-val^S reports the p-value from a Welch difference-in-means test comparing $TC^{D=1}$ and $TC^{D>1,S}$. p-val^M reports the corresponding p-value comparing $TC^{D=1}$ and $TC^{D>1,M}$. Single-dealer relationships are defined either across the full sample (Panel A), at the yearly level (Panel B), or based on a 95% execution threshold with a primary dealer (Panel C). Results are shown for the full sample and separately for corporate clients (CC) and financial institutions (FI). The sample period is January 2018 to July 2022

	N				TC					
	All	$D=1$	$D>1,S$	$D>1,M$	All	$D=1$	$D>1,S$	$D>1,M$	p-val^S	p-val^M
Panel A: Single-Dealer Relationships, Full Sample										
Full sample	791	101	237	452	0.68	1.38	0.64	0.55	0.00	0.00
CC	319	29	52	238	1.04	2.85	1.52	0.71	0.00	0.00
FI	471	72	185	214	0.45	0.79	0.40	0.37	0.00	0.00
CC, one-sided	270	16	34	220	0.84	3.05	1.13	0.63	0.00	0.00
CC, RFM	49	13	18	18	2.13	2.59	2.25	1.66	0.00	0.00
FI, one-sided	450	64	181	205	0.40	0.70	0.37	0.34	0.00	0.00
FI, RFM	22	8	4	9	1.36	1.47	1.72	1.11	0.00	0.00
Panel B: Single-Dealer Relationships, by Year										
Full sample	791	120	227	444	0.68	1.51	0.56	0.52	0.00	0.00
CC	319	42	44	233	1.04	2.79	1.28	0.67	0.00	0.00
FI	471	78	183	211	0.45	0.81	0.39	0.36	0.00	0.00
CC, one-sided	270	22	31	217	0.84	3.04	0.92	0.61	0.00	0.00
CC, RFM	49	20	13	16	2.13	2.53	2.15	1.61	0.00	0.00
FI, one-sided	450	68	179	203	0.40	0.72	0.36	0.33	0.00	0.00
FI, RFM	22	10	4	8	1.36	1.44	1.77	1.07	0.00	0.00
Panel C: Single-Dealer Relationships, 95% of the trades										
Full sample	791	133	215	442	0.68	1.35	0.55	0.55	0.00	0.00
CC	319	42	42	235	1.04	2.74	1.24	0.70	0.00	0.00
FI	471	91	173	207	0.45	0.72	0.38	0.38	0.00	0.00
CC, one-sided	270	21	31	219	0.84	2.98	0.92	0.62	0.00	0.00
CC, RFM	49	21	12	16	2.13	2.50	2.08	1.68	0.00	0.00
FI, one-sided	450	83	169	198	0.40	0.64	0.35	0.35	0.00	0.00
FI, RFM	22	9	4	9	1.36	1.47	1.72	1.10	0.00	0.00

Table XIII
Single-Dealer Requests and Transaction Size

The table reports average transaction costs for single-dealer requests, sorting requests into different quartiles based on the individudal transaction size. Q4 (Q1) refers to the largest (smallest) transactions. D= 1 refers to single-quote requests that always trade with the same dealer, while D>1 refers to single-quote requests, where clients trade with multiple dealers. N^{RFQ} is the number of requests, $\overline{TC}$ is the average transaction cost in bps, and $\overline{V}$ is the average traded volume per transaction. The sample period is January 2018 to July 2022.

	Category	N^{RFQ}	$\overline{TC}$	$\overline{V}$		Category	N^{RFQ}	$\overline{TC}$	$\overline{V}$
Q4	D=1	10,438	1.31	3,686,443	Q4	D>1	18,7008	0.54	5,064,509
Q3	D=1	6,479	0.56	1,099,639	Q3	D>1	97,149	0.10	1,078,343
Q2	D=1	36,331	0.94	714,906	Q2	D>1	255,189	0.38	793,099
Q1	D=1	47,717	1.71	102,055	Q1	D>1	149,817	1.19	98,141

Table XIV
Client-Dealer Relationships

The table reports summary statistics on client-dealer relationships at the client level. $N^{Clients}$ refers to the number of unique clients in each category, $\sum V$ is the total notaionl volume traded (expressed in billions USD), $\frac{V_i}{\sum V}$ is the share of the group's notional relative to the aggregated traded volume, $\overline{HHI}$ ($\widetilde{HHI}$ is the average (median) Herfindahl index level, $\overline{TC}$ ($\widetilde{TC}$) are the average (median) transaction costs, $\overline{Vol}$ ($\widetilde{Vol}$) is the average (median) trade size, $\overline{N^{Dealers}}$ ($\widetilde{N^{Dealers}}$) is the average (median) number of dealers a client is transacting with. The sample period is January 2018 to July 2022.

	$N^{Clients}$	$\sum V$ (bn)	$\frac{V_i}{\sum V}$	$\overline{HHI}$	$\widetilde{HHI}$	$\overline{TC}$	$\widetilde{TC}$	$\overline{Vol}$ (m)	$\widetilde{Vol}$ (m)	$\overline{N^{Dealers}}$	$\widetilde{N^{Dealers}}$
D=1	244	40,626	0.01	1.00	1.00	2.80	2.19	7.06	0.60	1.00	1
D>1	142	160,540	0.04	0.58	0.56	2.26	1.73	2.28	0.88	4.13	3
2	1,520	1,329,469	0.32	0.55	0.51	1.84	1.41	2.83	1.13	4.44	3
3-4	803	532,641	0.13	0.42	0.37	1.00	0.77	3.37	1.66	5.86	5
5-10	836	970,123	0.24	0.28	0.23	0.66	0.55	4.01	2.16	8.75	8
11-20	227	850,987	0.21	0.18	0.14	0.46	0.38	5.89	3.20	16.09	16
21+	24	223,073	0.05	0.15	0.08	0.24	0.20	4.90	2.93	26.29	28.5

Table XV
Client-Level Trading Networks and Dealer Concentration

The table reports summary statistics at the client level by quintiles of dealer concentration, measured by the volume-weighted Herfindahl–Hirschman Index (HHI) computed from executed trades. $N^{Clients}$ denotes the number of unique clients in each quintile; $\sum V$ is the total executed notional volume (in billions of USD), and $\frac{V_i}{\sum V}$ is the group's share of aggregate volume. $\overline{TC}$ and $\widetilde{TC}$ report the mean and median transaction costs (in basis points), while $\overline{HHI}$ and $\widetilde{HHI}$ report the mean and median dealer concentration. $\overline{Vol}$ and $\widetilde{Vol}$ denote the mean and median trade size (in USD millions). $\overline{N^D}$ and $\widetilde{N^D}$ measure the average and median number of distinct dealers per client. $\overline{T}$ and $\widetilde{T}$ denote the mean and median number of executed trades per client. Finally, $\overline{Dep}$ and $\widetilde{Dep}$ report the mean and median share of a client's total volume executed with its largest dealer. The sample period is January 2018 to July 2022.

	$N^{Clients}$	$\sum V$	$\frac{V_i}{\sum V}$	$\overline{TC}$	$\widetilde{TC}$	$\overline{HHI}$	$\widetilde{HHI}$	$\overline{Vol}$	$\widetilde{Vol}$	$\overline{N^D}$	$\widetilde{N^D}$	$\overline{T}$	$\widetilde{T}$	$\overline{Dep}$	$\widetilde{Dep}$
Q1	760	3187	0.78	0.66	0.36	0.14	0.15	5	1	15	13	1657	272	0.24	0.24
Q2	759	552	0.13	1.07	0.57	0.26	0.26	4	1	7	7	534	81	0.38	0.38
Q3	759	144	0.03	1.35	0.79	0.40	0.39	3	1	5	4	145	42	0.53	0.53
Q4	759	94	0.02	1.68	1.01	0.59	0.57	3	1	3	3	100	20	0.71	0.71
Q5	759	131	0.03	2.19	1.36	0.95	1.00	4	0	2	1	202	6	0.97	1.00

Table XVI
Requests, Dealer Networks, and Trading Costs

The table reports average transaction costs, distinguishing between the number of dealers approached by clients and their Herfindahl-Hirschmann Index (HHI). The sample period is January 2018 to July 2022.

	D=1	D>1	2	3-4	5-10	11-20	21 +
Low HHI		1.75	1.71	1.05	0.65	0.45	0.20
Medium HHI		2.36	2.02	1.26	0.76	0.35	0.11
High HHI	2.79	2.41	2.04	1.46	0.76	0.44	0.77

References

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.

O'Hara, M., and X. Zhou (2021): "The Electronic Evolution of Corporate Bond Dealers," *Journal of Financial Economics*, 140(2), 368–390.

Patton, A. J., and A. Timmermann (2010): "Monotonicity in Asset Returns: New Tests With Applications to the Term Structure, the CAPM, and Portfolio Sorts," *Journal of Financial Economics*, 98(3), 605–625.

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.

AI. Appendix: Data

Data Cleaning. We restrict the sample to spot FX request-for-quote (RFQ) records and focus on the more liquid currency pairs for which (i) both USD- and EUR-denominated orientations are present in the RFQ data and (ii) reference mid-quotes are available from LSEG TickHistory. The sample includes 19 exchange rates vis-à-vis the U.S. dollar or the euro: Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), Danish krone (DKK), euro (EUR), British pound sterling (GBP), Hong Kong dollar (HKD), Hungarian forint (HUF), Israeli new shekel (ILS), Japanese yen (JPY), Mexican peso (MXN), Norwegian krone (NOK), New Zealand dollar (NZD), Polish zloty (PLN), Romanian leu (RON), Swedish krona (SEK), Singapore dollar (SGD), Thai baht (THB), and South African rand (ZAR). This yields 37 exchange rates in total: 18 USD-quoted pairs, 18 EUR-quoted crosses, and the EURUSD exchange rate itself.

We standardize currency orientation across USD-quoted pairs (e.g., EURUSD, AUDUSD) and EUR-quoted crosses (e.g., EURJPY, EURSEK). Whenever an observation is recorded in the reciprocal orientation, we invert the price, flip trade direction (Buy/Sell), and switch the bid/ask side to express all quotes in the target orientation (e.g., USDAUD observations are converted to AUDUSD).

We then apply a sequence of data cleaning steps. First, we remove invalid RFQs by dropping entire requests in which an executed *Buy* appears on the bid side or an executed *Sell* appears on the ask side. Second, we exclude a small set of trades flagged as non-competitive executions with ambiguous status indicators, by removing the entire corresponding request. Third, we de-duplicate quotes: within each *Request ID* we keep only the first executed quote (by trade time); among non-executed quotes we collapse duplicates defined by *Request ID*, client–dealer link, notional in EUR, side, and quote value. Fourth, we filter extreme price outliers using a conservative year-wise median-absolute-deviation rule applied to log quote values and to their reciprocals. If an executed quote is flagged as an outlier, we drop the entire request; otherwise, only non-executed outlier quotes are removed.

Finally, we add a *No Trade* indicator for requests without any execution and merge TickHistory reference bid/ask/mid quotes to executed trades using an as-of backward match (i.e., the most recent quote at or before the execution time). For EUR-denominated crosses, we construct the reference price using an EURUSD anchor and the USD leg of the quote currency. We then recompute *Notional in USD* from the reported notional currency and the matched reference mid-

quotes (using the EURUSD anchor where needed).