Bank of England

One dollar, many prices: dealer-specific pricing of synthetic dollar funding

Staff Working Paper No. 1,207
September 2026

Marco Grotteria and Alex Kontoghiorghes

Bank of England

Staff Working Paper No. 1,207

One dollar, many prices: dealer-specific pricing of synthetic dollar funding

Marco Grotteria[(1)] and Alex Kontoghiorghes[(2)]

Abstract

Does the price of synthetic dollar funding depend on which dealer intermediates the trade? We answer this question by comparing FX forwards in the same dollar currency pair, half-hour, and maturity bucket, which removes the common forward curve and isolates dealer-specific pricing. These dealer-specific prices vary substantially: their standard deviation is about 4 basis points per year, and the large dispersion is present both across dealers at a point in time and over time for a given dealer. Dealers charge 2.5 basis points more for buying dollars forward than for selling them, and the gap widens to 11.8 basis points at maturities under one month. The asymmetry survives controls for client-dealer relationships and trade characteristics. Differences in funding costs account for little of the observed price dispersion. Synthetic dollar funding therefore has many prices, reflecting dealer clientele and pricing power more than underlying funding conditions.

Key words: FX forwards, dealer market power, covered interest parity, synthetic dollar funding, funding constraints, EMIR.

JEL classification: E43, E52, G12, G15, G21.

---

(1) London Business School. Email: mgrotteria@london.edu
(2) Bank of England. Email: alex.kontoghiorghes@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees. We thank Svetlana Bryzgalova, Joao Cocco, Pasquale Della Corte, Rob Czech, Gerardo Ferrara, Iryna Kaminska, Daniel Ostry, Anna Pavlova, Andreas Schrimpf, Stephen Schaefer, Jonathan Wallen, the Bank of England Financial Market Infrastructure Data Team, and seminar participants at the Bank of England for helpful comments. We also thank Miruna-Daniela Ivan for their help obtaining some of the data.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Email: enquiries@bankofengland.co.uk

ISSN 1749-9135 (on-line)

I Introduction

Dollar funding outside the United States is intermediated on a vast scale through foreign-exchange derivatives. Insurers, pension funds, asset managers, and banks borrow in local currency and use FX forwards and swaps to obtain dollars, making the executed forward rate a transaction price of synthetic dollar funding. Yet the cross-currency basis—the standard measure of this funding cost—is an aggregate market price. It cannot reveal whether two institutions seeking otherwise comparable dollar funding pay the same price when they trade through different dealers. With roughly $80 trillion in FX forwards and swaps outstanding (Borio, McCauley, and McGuire, 2022), even modest dealer-level price differences can imply economically significant differences in the cost of dollar credit.

This paper asks whether the price of synthetic dollar funding depends on the dealer that intermediates the transaction. Using confidential transaction-level reports for outright deliverable FX forwards accessed at the Bank of England, we compare trades in the same currency pair, executed within the same half-hour and at similar maturities, but intermediated by different dealer groups. We find that synthetic dollars do not have a single transaction price. Near-identical trades differ by about 7 basis points per year between the 10th and 90th percentiles of dealer-month prices. Dealers transact at prices about 2.5 basis points higher when buying dollars forward than when selling them, and the gap rises to 11.8 basis points at maturities within one month. These differences survive controls for client-dealer relationships and observable transaction characteristics. Dealer prices also covary with secured dollar funding rates and scheduled dollar cashflows, and dealer balance-sheet conditions predict heterogeneous responses to common monetary-policy news.

Our data cover January 2021 through October 2025 and report executed forward rates, counterparty legal entity identifiers, trade direction, maturity dates, notionals, execution timestamps, and execution characteristics. These fields allow us to consolidate dealer legal entities to parent groups, identify the two sides of each dealer's forward book, and compare prices within client-dealer relationships. We focus on six dollar pairs—EUR, GBP, JPY, CHF, AUD, and CAD against the dollar—and on observed G15 dealer parents. Outright forwards are themselves an important segment of synthetic dollar funding: in April 2025 they accounted for about 20% of global daily FX turnover, and their turnover grew faster between 2022 and 2025 than that of FX swaps (Bank for International Settlements, 2025).[1]

Our empirical design isolates dealer-associated price differences from the common forward curve. We regress the log executed forward rate on dealer-by-direction indicators interacted with time to maturity, absorbing execution half-hour × currency pair × ten-day maturity-

[1]One advantage of using outright FX forward contracts is that they are priced all-in.

bucket fixed effects. Within these narrow comparison groups, the spot rate, interest-rate differentials, aggregate covered interest parity (CIP) deviations, term premia, and market-wide liquidity are common to the trades being compared. Identification therefore comes from differences in executed prices across transactions in the same dollar currency pair, at nearly the same time and maturity, but intermediated by different observed dealer parents on one or both sides of the contract. We express the resulting dealer-specific estimates in annualized log points and call them *dealer wedges*. Each wedge is measured relative to the omitted category on the same side of the trade, which contains counterparties outside the observed G15 dealer-parent set, including non-dealer clients and non-G15 dealers. The design therefore identifies how the price of synthetic dollars varies across intermediaries around the contemporaneous market price, not the aggregate cost of synthetic dollar funding.

Direction matters because the two sides of an FX forward expose a dealer to different funding terms. Throughout the paper, buy and sell are defined in foreign-currency (FC) terms. A dealer that sells FC forward receives—and thus buys—dollars at maturity; its hedge loads on dollar borrowing and foreign-currency lending spreads. A dealer that buys FC forward delivers dollars at maturity; its hedge instead loads on dollar lending and foreign-currency borrowing spreads. Accordingly, the difference between the sell-FC and buy-FC wedges loads on borrowing-minus-lending spreads within both currencies, whereas their sum nets the dealer's dollar spreads against its foreign-currency spreads and therefore speaks to differential dollar conditions across dealers. We use both objects to distinguish directional pricing from cross-dealer differences in dollar conditions.

The dealer wedges we estimate are equilibrium transaction-price differences that can reflect both intermediation costs and residual demand. Funding costs, balance-sheet capacity, inventories, and risk-bearing capacity can differ across dealers, while client relationships, operational constraints, switching costs, bargaining, and heterogeneous customer demand can prevent order flow from moving costlessly to the cheapest intermediary.[2] We therefore interpret a wedge as combining an intermediation-cost component with a residual-demand component.

To isolate the funding cost component of the dealer wedges and address the possibility that different dealers may serve different customers and intermediate different types of trades, we include more granular fixed effects and controls. First, we add client $\times$ dealer-parent fixed effects and, more stringently, client $\times$ dealer-parent $\times$ currency-pair fixed effects. These specifications absorb persistent pricing differences within client-dealer relationships,

[2]Collateral arrangements, execution protocols, relationship pricing, search frictions, and bargaining positions are among the channels through which limited customer substitution and dealer market power can enter transaction prices.

including relationship-specific differences across currencies. Second, we control for observable transaction characteristics associated with negotiated spreads, including trade size, execution venue, and currency-pair liquidity.

Our first result is substantial cross-dealer price dispersion. Monthly wedges have standard deviations of 4.4 basis points on the buy-FC side and 3.9 basis points on the sell-FC side. Removing both month and dealer fixed effects still leaves standard deviations of 3.1–3.4 basis points, so the dispersion is neither simply a common time-series component nor a permanent ranking of expensive and cheap dealers. Notional weighting leaves the level and time-series behavior of average wedges essentially unchanged, so small-volume dealer-months do not drive the result. For comparison, average bid-ask spreads in these markets are roughly 1 basis point, which highlights the economic significance of our estimates.

Second, prices differ systematically across the two sides of a dealer's book. The sell-FC wedge—the side on which the dealer buys dollars—averages 2.1 basis points per year, while the buy-FC wedge is close to zero. The sell-FC minus buy-FC difference is 2.52 basis points per year, with a standard error of 0.30. Its funding component loads on the dealer's borrowing-minus-lending spread in dollars and in foreign currency, capturing the cost of borrowing rather than lending in the two currencies. The gap remains between 2.2 and 2.7 basis points across specifications controlling for client-dealer relationships and observable transaction characteristics. It rises sharply at short maturities, from 2.07 basis points beyond three months to 11.82 basis points within one month.

Third, the sum of the two directional wedges measures dollar conditions across dealers (i.e., each dealer's dollar borrowing and lending spreads against its corresponding foreign-currency spreads). The sum is 1.78 basis points per year in the baseline specification, but falls to 0.20–0.95 basis points and becomes statistically indistinguishable from zero once we absorb client relationships and observable transaction characteristics. The cross section therefore suggests that differences in dollar funding conditions across major dealers explain little of the substantial dealer-price dispersion and directional asymmetry observed in the market.

Next, we show that the two sides of a dealer's book share a common dealer component. In the dealer-month panel, the slope of sell-FC wedges on buy-FC wedges, controlling for month effects, is 0.532 and remains between 0.52 and 0.56 after controlling for client relationships and transaction characteristics. The persistence of this relationship after absorbing dealer-client relationships is important: dealers that price one side of their book relatively high also tend to price the other side relatively high, even when the comparison is no longer driven by demand factors, e.g., the customers they serve.

We next measure dealer dollar conditions directly. Using confidential UK repo reports, we construct dealer-level secured dollar borrowing and lending rates. After absorbing dealer $\times$ currency-pair effects, a one-percentage-point increase in a buy-FC dealer's dollar repo lending rate is associated with a 1.12 basis point higher residual forward price; the analogous sell-FC borrowing-rate relationship is not statistically distinguishable from zero. Scheduled dollar cashflows from dealers' forward books provide a second measure of predetermined dollar conditions. Although the average wedges show little cross-dealer dispersion in relative dollar conditions after controls, within-dealer changes in observed dollar conditions covary with transaction prices.

Dealer conditions also shape the transmission of monetary-policy news. Around FOMC announcements, we compare transactions before and after the same announcement and ask whether the response to the common policy surprise varies with dealers' predetermined dollar positions. A \$1 billion increase in scheduled dollar inflows over the following seven days changes the response of the forward price to a one-unit statement surprise by approximately -0.08 to -0.10 basis points, depending on the dealer side; scheduled dollar outflows have no measurable differential effect. Dealers entering announcements with higher secured dollar rates likewise exhibit different pass-through.

Our evidence changes how synthetic dollar funding should be measured and interpreted. The aggregate cross-currency basis masks a distribution of transaction prices that varies across dealers, sides of the market, maturities, and dealer balance-sheet conditions. Two customers seeking otherwise comparable dollar funding can therefore face different prices simply because they trade through different dealers. Customers raising synthetic dollars transact on the sell-FC side, which is priced 2.5 basis points above the buy-FC side on average and 11.8 basis points above it within one month. The price a customer actually pays therefore depends on the dealer and the side it occupies, by most at the horizon where dollar funding stress binds.

Moreover, because balance-sheet requirements bind at the banking-group level, a constraint that tightens at a few dealers and one that tightens modestly across all dealers can move the average basis by a similar amount while generating very different patterns across dealers. As a result, the dispersion and direction of dealer prices are informative about intermediation capacity in a way that the level of the basis is not. This is particularly useful because these requirements are written on gross exposures, regardless of which currency a dealer delivers, whereas the pricing effects are asymmetric across the two sides of the market.

Related literature. Our results connect the literature on covered interest parity to the microeconomics of dealer intermediation. The classic transaction-cost literature replaces frictionless covered interest parity with side-specific no-arbitrage bounds (Frenkel and Levich, 1975, 1977; Clinton, 1988; Deardorff, 1979; Callier, 1981). We study the transaction prices selected within those bounds. The key empirical fact is that the selected price depends systematically on the intermediating dealer and on which side of the dealer's dollar book the transaction occupies. The relevant departure from a single CIP price is therefore not only a market-wide basis, but a cross section of dealer- and side-specific prices.

A large subsequent literature explains why an aggregate cross-currency basis can emerge and persist. Crisis-era evidence links dislocations in FX swap and cross-currency swap markets to dollar funding shortages at non-U.S. financial institutions (Baba, Packer, and Nagano, 2008). Post-crisis work relates CIP deviations to costly balance sheets, dollar scarcity, hedging demand, heterogeneous funding costs, and limits to scalable arbitrage (Borio, McCauley, McGuire, and Sushko, 2016; Du, Tepper, and Verdelhan, 2018; Ivashina, Scharfstein, and Stein, 2015; Avdjiev, Du, Koch, and Shin, 2019; Rime, Schrimpf, and Syrstad, 2022; Kloks, Mattille, and Ranaldo, 2026). Related work connects the basis to credit-market segmentation, FX-swap order flow, safe-dollar demand, Treasury-market intermediation, and the separation between dollar and Treasury convenience yields (Liao, 2020; Syrstad and Viswanath-Natraj, 2022; Jiang, Krishnamurthy, and Lustig, 2024; Du, Hébert, and Li, 2023; Du, Keerati, and Schreger, 2026); dollar funding needs also appear in the broader dash for dollars" during stress episodes (Cesa-Bianchi, Czech, and Eguren-Martin, 2021). This literature explains why the aggregate price of synthetic dollar funding can depart from CIP. Our object is different: the fixed effects absorb this aggregate price, and we study the dispersion of transaction prices around it across dealers and across the two sides of their books. This distinction matters for what we learn about dollar scarcity. Once we absorb client-dealer relationships and observable transaction characteristics, the sum of our directional wedges—the combination that captures differential dollar funding conditions—is statistically indistinguishable from zero. Thus, although aggregate dollar funding can be scarce, we find little evidence that major dealers systematically face different dollar funding conditions from one another.

Recent transaction-level studies provide the closest empirical benchmarks. Hacioğlu-Hoke, Ostry, Rey, Rousset Planat, Stavrakeva, and Tang (2024) map the UK FX derivatives market and document substantial differences across client sectors in the use of FX derivatives. Cenedese, Della Corte, and Wang (2021) exploit a UK leverage-ratio shock to show that affected dealers transact synthetic dollar funding at different prices from unaffected dealers. Their event identifies how a particular balance-sheet constraint shifts the prices of the dealers on which it binds relative to unaffected dealers. We ask instead how transaction prices are distributed across dealers in ordinary trading. Our half-hour $\times$ currency-pair $\times$ maturity-bucket fixed effects absorb the contemporaneous market price, so each dealer wedge measures a price relative to otherwise comparable transactions rather than the level of the CIP deviation. We estimate this cross-sectional distribution continuously over 58 months and separately on the two sides of the dealer book. The two approaches therefore isolate different margins: Cenedese et al. (2021) show that a discrete constraint can move affected dealers relative to unaffected dealers; we show that, in ordinary trading, substantial dealer-level dispersion remains around the common market price even though, after accounting for clientele and transaction characteristics, it does not appear as persistent cross-dealer dispersion in dollar funding conditions.

Our pricing framework is also closely related to Kloks, Mattille, and Ranaldo (2026), who show how balance-sheet constraints and segmentation generate a synthetic dollar premium when non-U.S. banks substitute between repo and off-balance-sheet FX swaps. Their analysis characterizes the shadow cost at which a borrower becomes indifferent across sources of dollar funding. Our dealer-level approach instead asks how prices are distributed across the intermediaries that supply those transactions. Near-identical trades can differ by roughly 7 basis points per year across dealer-months, and the dealer buy-dollar side is priced about 2.5 basis points above the dealer sell-dollar side on average. Yet, after absorbing client relationships and transaction characteristics, we find no corresponding cross-sectional dollar wedge across dealers. The aggregate synthetic dollar premium therefore coexists with substantial dealer-level price dispersion without requiring persistent dispersion in major dealers' dollar funding conditions.

This heterogeneity also connects our results to intermediary asset pricing. Limited arbitrage capital, funding liquidity, margins, and intermediary net worth can affect asset prices and risk premia (Shleifer and Vishny, 1997; Brunnermeier and Pedersen, 2009; Gârleanu and Pedersen, 2011; He and Krishnamurthy, 2013; Adrian, Etula, and Muir, 2014; He, Kelly, and Manela, 2017). Much of this literature emphasizes a marginal intermediary, a common intermediary state variable, or an aggregate balance-sheet constraint. Our transaction-level evidence identifies a complementary margin: the common intermediary component explains only part of the variation in FX-forward prices. Dealer-specific and side-specific conditions remain economically important even though we find little evidence of persistent cross-dealer differences in dollar funding conditions once clientele and transaction characteristics are absorbed.

Finally, the paper links the macroeconomics of dollar funding to the microstructure of OTC markets. Search-and-bargaining models imply that OTC transaction prices depend on outside options, trading networks, and access to alternative counterparties (Duffie, Gârleanu, and Pedersen, 2005). Empirically, dealer networks, repeated relationships, inventories, and price discrimination shape transaction prices in OTC markets (Li and Schürhoff, 2019; Hendershott, Li, Livdan, and Schürhoff, 2020; Di Maggio, Kermani, and Song, 2017; Hau, Hoffmann, Langfield, and Timmer, 2021; Colliard, Foucault, and Hoffmann, 2021).[3] In FX, order flow carries price-relevant information and customer groups differ in both the information contained in their trades and the prices they obtain (Evans and Lyons, 2002; Menkhoff, Sarno, Schmeling, and Schrimpf, 2016; Ranaldo and Somogyi, 2021). Our client-by-dealer fixed effects absorb persistent pricing differences within client-dealer relationships, shifting the comparison from customer types toward differences across dealers. The fact that these controls eliminate the cross-dealer dollar wedge while leaving substantial price dispersion and the directional gap largely intact points to a distinction between what dollars cost dealers and the prices at which dealers intermediate them.

Within FX derivatives, Wallen (2022) shows that the aggregate CIP basis contains markups as well as intermediation costs and that these markups rise when competition among synthetic-dollar suppliers falls. Our analysis studies a different margin of imperfect competition. Rather than asking how market power affects the level of the CIP basis, our specifications absorb the contemporaneous market-wide basis and we ask how transaction prices are distributed across dealers around that common price. This distinction lets us separate two sources of heterogeneity that are observationally combined in an aggregate basis: differences in dealers' dollar funding conditions and differences in the prices at which they intermediate their clients. We find substantial dispersion in the latter but no corresponding cross-dealer dollar-funding-cost differential.

Ferrara and Hall (2026) provide complementary evidence on the role of dealer relationships. They show that clients exposed to Credit Suisse's collapse substituted toward existing alternative dealers without differential cost increases. Their experiment therefore asks how customers reallocate across established relationships following a discrete dealer shock. We instead study the prices attached to those relationships in ordinary trading. Customers can have alternative dealers available when a major shock forces substitution and nevertheless face systematically different prices across those dealers in normal times. This distinction helps explain why dealer relationships matter for equilibrium pricing even when they do not prevent substitution following a large shock. In synthetic dollar funding, the residual demand facing each dealer therefore becomes part of the pricing object, linking dealer intermediation to demand-system approaches to asset pricing (Koijen and Yogo, 2019).

Synthetic dollar funding, in short, has many transaction prices. The aggregate basis captures the common price of synthetic dollars; our evidence uncovers the dealer-level distribution around it. That distribution varies substantially across dealers and across the two

[3] Hau et al. (2021) document dispersion across clients of a given dealer, or price discrimination along the customer dimension. Our first fact concerns dispersion across dealers for near-identical trades. The two margins are complementary: a dealer can transact with different customers at different prices, and customers can face different prices across dealers.

sides of their books even when persistent cross-dealer differences in dollar funding conditions do not, revealing intermediation and monetary-transmission margins that an aggregate CIP deviation cannot capture.

II Empirical Framework

This section defines the dealer wedge and the comparisons that identify it. The empirical object is the dealer-associated component of an executed FX-forward price after removing the forward-curve components common to transactions in the same currency pair, at nearly the same time and maturity. We estimate this component separately for the two sides of each dealer's forward book because buying and selling dollars expose the dealer to different funding terms. The resulting *dealer wedge* is an annualized log-price difference: it measures where a dealer transacts relative to otherwise comparable transactions.

Let S_t denote the spot exchange rate and let $F_t^{\text{bench}}(T)$ denote the benchmark forward rate for delivery at date T as implied by CIP. We quote each currency pair as dollars per unit of foreign currency and write $\tau = T - t$ for time to maturity in years. With continuously compounded benchmark yields,

$$\log F_t^{\text{bench}}(T) = \log S_t + \tau(r_{\text{USD},t} - r_{f,t}), \tag{1}$$

where $r_{\text{USD},t}$ and $r_{f,t}$ are the dollar and foreign-currency benchmark rates over the life of the contract. For transactions in the same currency pair, execution time, and maturity, this benchmark captures the spot exchange rate, the interest-rate differential, and the market-wide component of the forward curve. Rather than constructing $F_t^{\text{bench}}(T)$ from a particular set of benchmark rates, the empirical specification absorbs the common forward curve with granular fixed effects and uses the remaining price variation within narrowly defined comparison groups to estimate dealer-associated differences.

In fact, executed forward rates need not coincide across dealers. FX forwards are mostly bilateral over-the-counter contracts, and dealers differ in funding costs, secured market access, balance-sheet capacity, internal risk limits, and inventories. We refer to these differences as the *intermediation-cost* component of the dealer wedge (c). Customers may also value existing relationships, face operational constraints, or find it costly to redirect order flow, giving each dealer a different residual demand for its services. We refer to the resulting pricing component as *residual demand* (μ).[4] The transaction price reflects both components,

[4]Neither component must be positive: funding can be relatively cheap, a trade can release rather than consume balance-sheet capacity, and a dealer can concede price to attract flow or manage its inventory.

so

$$\omega = \underbrace{c}_{\text{intermediation cost}} + \underbrace{\mu}_{\text{residual demand}}. \tag{2}$$

Because these differences can depend on which currency the dealer receives at maturity, the direction convention is central. All dealer roles are defined in foreign-currency terms. For trade i, executed at t_i and maturing at T_i, let dealer parent $d(i)$ occupy side $j(i) \in \{\text{buy}, \text{sell}\}$. If $j(i) = \text{buy}$, dealer $d(i)$ *buys foreign currency forward*: it receives foreign currency and delivers dollars at maturity. We call this the buy-FC, or dollar-delivery, side. If $j(i) = \text{sell}$, the dealer *sells foreign currency forward*: it delivers foreign currency and receives dollars at maturity. We call this the sell-FC, or buy-dollar, side. A dealer that buys foreign currency forward is long the exchange rate S.

For a single dealer-side observation, write the executed transaction price as

$$\log F_i = \log F_{t_i}^{\text{bench}}(T_i) + \tau_i \, \omega_{d(i),j(i),t_i} + u_i. \tag{3}$$

The term ω_{djt} is the dealer wedge in annualized log points. Because ω_{djt} enters the transaction price multiplied by time to maturity, its natural unit is an annualized log-price difference.

Equation (3) decomposes an executed price into three parts: the common forward curve, a dealer- and side-specific price wedge, and a transaction-specific residual. The wedge is applied over the life of the contract, so its scaling by τ_i mirrors equation (1). There, the interest differential $r_{\text{USD},t} - r_{f,t}$ enters the log forward rate multiplied by maturity; here, a dealer transacting at an implied rate a few basis points away from the benchmark shifts the log price by the same difference cumulated over the contract horizon. A one-year contract priced at a wedge of five basis points therefore has a log price about 5×10^{-4} above the benchmark, whereas a one-month contract at the same annualized wedge is displaced by one twelfth as much. Expressing wedges in annualized units makes them comparable across maturities.

As we show in Appendix A, the economic content of the intermediation-cost component of the wedges depends on how a dealer hedges the position. A participant choosing between the forward and the same-side synthetic transaction faces the one-way replication bound; a dealer accommodating customer flow instead offsets the position with the opposite-side synthetic. Under the assumption that borrowing rates weakly exceed lending rates in each currency, the hedge bound is the tighter of the two on each side and is the one we use to interpret the data.[5] Under this bound, dollar borrowing and foreign-currency lending determine the funding component on the sell-FC side, while dollar lending and foreign-currency borrowing determine it on the buy-FC side. Thus, the same dealer can face different funding terms depending on which side of the forward it intermediates.

Therefore, dealers' funding components, $c^{\text{buy}}_{d,t}$ and $c^{\text{sell}}_{d,t}$ for the buy-FC and sell-FC sides respectively, are

$$c^{\text{buy}}_{d,t} = s^{L}_{d,\text{USD},t} - s^{B}_{d,\text{FC},t}, \qquad c^{\text{sell}}_{d,t} = s^{B}_{d,\text{USD},t} - s^{L}_{d,\text{FC},t}. \tag{4}$$

Here s^B and s^L denote borrowing and lending spreads relative to the benchmark rates. A sell-FC dealer receives dollars at maturity, so its hedge loads on its dollar borrowing spread and foreign-currency lending spread. A buy-FC dealer delivers dollars at maturity, so its hedge instead loads on its dollar lending spread and foreign-currency borrowing spread. Dollar funding conditions can therefore affect the two sides of a dealer's book differently.

This decomposition motivates two combinations of the observed directional wedges. For each dealer-month, define the *level* and the *difference*,

$$L_{d,t} \equiv \omega^{\text{sell}}_{d,t} + \omega^{\text{buy}}_{d,t}, \qquad \Delta_{d,t} \equiv \omega^{\text{sell}}_{d,t} - \omega^{\text{buy}}_{d,t}, \tag{5}$$

implying

$$L_{d,t} = (c^{\text{sell}}_{d,t} + c^{\text{buy}}_{d,t}) + (\mu^{\text{sell}}_{d,t} + \mu^{\text{buy}}_{d,t}), \tag{6}$$

and

$$\Delta_{d,t} = (c^{\text{sell}}_{d,t} - c^{\text{buy}}_{d,t}) + (\mu^{\text{sell}}_{d,t} - \mu^{\text{buy}}_{d,t}). \tag{7}$$

The funding components of the level and difference are

$$c^{\text{sell}}_{d,t}+c^{\text{buy}}_{d,t} = (s^{B}_{d,\text{USD},t}+s^{L}_{d,\text{USD},t})-(s^{B}_{d,\text{FC},t}+s^{L}_{d,\text{FC},t}), \qquad c^{\text{sell}}_{d,t}-c^{\text{buy}}_{d,t} = (s^{B}_{d,\text{USD},t}-s^{L}_{d,\text{USD},t})+(s^{B}_{d,\text{FC},t}-s^{L}_{d,\text{FC},t}). \tag{8}$$

The *level* nets the dealer's dollar borrowing and lending spreads against its foreign-currency borrowing and lending spreads. Because the wedges are measured relative to transactions facing the same contemporaneous forward curve, its funding component captures whether a dealer's dollar funding conditions are more expensive relative to its foreign-currency conditions than those of the dealers with which it is compared. The *difference* combines borrowing-minus-lending spreads within dollars and within foreign currency. It is therefore a round-trip

[5]That ordering is institutionalized in major currency blocs. For instance, the Federal Reserve's standing repo facility rate lies above the rate on its overnight reverse repurchase agreements, and the ECB's marginal lending facility rate lies above its deposit facility rate, so a bank that raises cash pays above the policy benchmark while a bank that places cash earns below it.

pricing object: under $s_k^B \geq s_k^L$, its funding component reflects the cost of borrowing rather than lending in the two currencies and is not specific to dollars. We use both combinations throughout because they answer different economic questions: $L_{d,t}$ captures relative dollar funding conditions, whereas $\Delta_{d,t}$ measures the asymmetry between the two sides of dealer intermediation.

The first empirical challenge is to separate the dealer-associated component ω from movements in the common forward curve. We do so by comparing transactions in narrowly defined comparison groups: the same currency pair, executed within the same 30-minute window, and with maturities in the same ten-day bucket. Let $c(i)$ denote the dollar currency pair, $w(i)$ the calendar 30-minute execution window, and $\ell(i)$ a ten-day maturity bucket. Let $b(i)$ denote the buy-FC dealer parent, which delivers dollars at maturity, and $s(i)$ the sell-FC dealer parent, which receives dollars. We retain trades in which at least one counterparty is an observed G15 dealer parent and drop intra-group transactions. Counterparties outside the observed G15 dealer-parent set—including non-dealer clients and non-G15 dealers—form the omitted category on their respective side of the trade. When both counterparties are observed G15 dealer parents, both dealer indicators enter the regression.

With $\mathcal{D}$ denoting the set of observed G15 dealer parents, our baseline specification is

$$\log F_i = \alpha_{w(i),c(i),\ell(i)} + \gamma_{c(i)}\tau_i + \sum_{d\in\mathcal{D}} \omega_d^{\text{buy}} \mathbb{1}\{b(i) = d\}\tau_i + \sum_{d\in\mathcal{D}} \omega_d^{\text{sell}} \mathbb{1}\{s(i) = d\}\tau_i + \varepsilon_i. \qquad (9)$$

The unit of observation i is a single executed forward contract—one record per transaction—and $\log F_i$ is the log of its executed forward rate, quoted in dollars per unit of foreign currency. There is one observation per contract rather than one per counterparty: both sides of a trade enter the same row through the indicators $\mathbb{1}\{b(i) = d\}$ and $\mathbb{1}\{s(i) = d\}$. The variable τ_i is the contract's time to maturity in years.

The fixed effect $\alpha_{w,c,\ell}$ absorbs components common to transactions in the same comparison group—the same currency pair, calendar half-hour, and ten-day maturity bucket—including the spot rate, interest-rate differentials, aggregate CIP deviations, term premia, and market-wide liquidity at that level of aggregation. The pair-specific slope $\gamma_c\tau_i$ absorbs residual linear carry within each maturity bucket. Identification therefore comes from comparing executed forward rates for transactions that face essentially the same contemporaneous forward curve but are intermediated by different observed dealer parents. This within-group comparison identifies the dealer wedges.

We estimate equation (9) at two frequencies because the two estimations answer different empirical questions. Estimating it once on the full transaction sample produces one buy-FC wedge ω_d^{buy} and one sell-FC wedge ω_d^{sell} for each dealer parent. These estimates summarize persistent cross-dealer pricing differences over the sample. Estimating the same specification separately in each calendar month produces the dealer-by-role panel $\{\hat{\omega}^j_{d,t}\}$. The monthly panel is the main empirical object in Section IV: it allows dealer prices to move over time and underlies the dispersion, persistence, maturity, factor-structure, and linkage results.

Every dealer coefficient has the same economic interpretation. Because the dealer indicators in equation (9) are interacted with τ_i, the coefficients are annualized log-price wedges. A positive $\hat{\omega}^{\text{buy}}_d$ means that the executed forward rate is higher when dealer d occupies the buy-FC side than when that side is occupied by a counterparty in the omitted category, conditional on the common comparison group, the pair-specific maturity slope, and the counterparty on the opposite side. A positive $\hat{\omega}^{\text{sell}}_d$ has the analogous interpretation when dealer d occupies the sell-FC side. The coefficients thus locate each observed dealer relative to the reference category; their differences across dealers describe the cross section of dealer-associated transaction prices.

The principal threat to interpreting the wedges as dealer-associated prices is remaining heterogeneity within these comparison groups. Dealers can transact at different exact moments within a half-hour, with different clients, collateral arrangements, trade sizes, venues, and relationships. Some of these features—such as relationship pricing or bargaining—are part of the equilibrium dealer price we seek to measure.[6] Client sorting is different: if particular customers systematically trade with particular dealers and would receive different prices regardless of dealer identity, the estimated wedges could reflect customer composition rather than cross-dealer pricing.

As mentioned, we address this concern by tightening the comparison within client-dealer relationships. We augment equation (9) first with client legal entity identifier (LEI) $\times$ dealer-parent fixed effects and then with client LEI $\times$ dealer-parent $\times$ currency-pair fixed effects. The first specification absorbs persistent pricing differences within each client-dealer relationship; the second additionally allows those relationship-specific pricing levels to differ across currency pairs. We separately control for observable transaction characteristics associated with negotiated bid-ask spreads, i.e., trade size, execution protocol, and pair liquidity. Admittedly, these controls do not eliminate the *whole* residual-demand component of wedges, but they absorb important sources of it and allow us to recompute the level and difference after holding fixed more of the customer relationship and execution environment.

The empirical analysis then uses the estimated wedges to answer three related questions. First, do otherwise comparable transactions converge to a common dealer price, or is there economically meaningful dispersion across dealer-months? Persistence, notional

[6]Timing within a window will only widen standard errors, but will not bias the wedge estimates unless execution timing is systematically related to dealer identity, in which case we want to capture this effect.

weighting, and the client and transaction controls help distinguish systematic dealer pricing from sampling noise and composition. Second, are the two sides of dollar intermediation priced symmetrically? Comparing buy-FC and sell-FC wedges, and tracing their maturity profiles, reveals whether the dealer's dollar role matters for transaction prices. Third, how much of dealer pricing is common to both sides of a dealer's book? We distinguish a common dealer-level component from side-specific variation. Finally, dealer-level repo rates and scheduled dollar exposures provide observable measures with which to assess whether these price differences covary with dealers' dollar conditions.

To examine how the two sides of a dealer's book move together, we relate sell-FC wedges to buy-FC wedges in the pooled dealer-month panel with month fixed effects as follows

$$\hat{\omega}^{\text{sell}}_{d,t} = a_t + \beta \hat{\omega}^{\text{buy}}_{d,t} + u_{d,t}. \tag{10}$$

The slope β measures whether dealers that price one side of their book relatively high also price the other side relatively high, while R^2 measures how much of the cross-dealer variation in the two wedges is shared. A strong relationship points to a common dealer component; a weak relationship leaves more of the variation specific to one side of the dealer's book. We repeat these exercises for each set of wedge estimates: the baseline wedges, the wedges estimated after absorbing client-dealer relationships, and those estimated after additionally controlling for observable transaction characteristics. Comparing the linkage across these specifications shows whether the common dealer component survives as we absorb important sources of residual demand and move closer to the intermediation-cost component of the wedge. Section IV turns to the evidence next and applies the same sequence of specifications throughout the analysis.

III Data

Our empirical design requires three pieces of information for each FX forward: the executed price, the identities of both counterparties, and the direction of each counterparty's position. Public FX datasets do not jointly report these fields.[7] We therefore use confidential UK EMIR transaction reports accessed at the Bank of England. These records let us compare executed forward prices across dealers while holding the currency pair, execution time, and maturity approximately fixed.

UK EMIR requires regulated entities to report derivative transactions to a central reposi-

[7] CLS settlement data lack dealer identities. The BIS triennial survey reports volumes but not prices. Thomson Reuters Matching records interdealer trades but not dealer-to-client trades. The DTCC swap data repository in the United States reports prices but redacts counterparty identities from the public tape.

tory. For each deliverable FX forward, the reports record the execution timestamp, currency pair, maturity date, notional amounts in both currencies, executed forward rate, legal entity identifiers (LEIs) for both counterparties, the reporting counterparty's buy/sell indicator, and the execution venue. We restrict the sample to outright forward contracts. Our sample covers January 2021 through October 2025, giving 58 monthly estimation windows. After reconstruction and cleaning, the unit of observation is an individual forward contract.

We measure dealer identity at the banking-group level because treasury, capital, funding, and balance-sheet constraints are managed within parent groups. We therefore map counterparty LEIs to ultimate parent groups using the Global LEI Foundation (GLEIF) hierarchy.[8] We focus on six core dollar pairs—AUD, CAD, CHF, EUR, GBP, and JPY against the dollar—and retain transactions in which at least one counterparty belongs to the G15 group of FX dealers. Fourteen G15 parent groups transact in these pairs during our sample, so dealer-level variables are defined for 14 parent groups. The resulting monthly dealer panel is balanced, with $14 \times 58 = 812$ dealer-month observations for each role. We exclude intra-group transactions because internal transfers need not reflect arm's-length prices.

Constructing one observation per economic transaction requires reconstructing each contract's reporting lifecycle. Raw EMIR records can contain the initial trade, subsequent amendments, valuation updates, cancellations, and daily state reports for the same contract. We use end-of-day state files to reconstruct the transaction, assign amendments to the originally executed price, discard cancelled contracts, and retain one record per forward. We then assign the buy-FC and sell-FC roles defined in Section II to the relevant counterparties.

A second construction step puts every transaction into a common economic convention. The same contract can be reported, for example, as USD/JPY or JPY/USD. Reversing that convention changes both the reported price and the economic direction of the trade. We therefore express all dollar pairs as dollars per unit of foreign currency. Given a reported forward rate F_t and end-of-day spot rate S_t, we compare $|F_t - S_t|$ with $|1/F_t - S_t|$. When the reciprocal is closer by more than a relative tolerance of $\varepsilon = 10^{-4}$, we invert the forward rate, exchange the two notional legs, and reverse the buy/sell indicator.[9]

We next remove observations likely to reflect reporting errors rather than economically extreme forward prices. Both filters operate within currency-pair-by-month groups and use the log forward price net of the same-day log spot rate. The first trims observations below the 1st percentile or above the 99th percentile of this quantity. Netting out spot before

[8] www.gleif.org

[9] This criterion is uninformative when S_t is close to one, which occurs for USD/CHF over parts of 2022 and 2023. For those observations, we determine orientation from the ratio of the two reported notional legs, which does not depend on the quote convention.

trimming prevents an unusually high or low spot exchange rate from being treated as an anomalous forward transaction. The second filter drops observations more than ten median-absolute-deviation units from the group median.

Table B.1 shows how these choices determine the final sample. The raw UK EMIR reports contain 178,473,017 FX forward transactions between January 2021 and October 2025. Restricting the data to the six core dollar pairs leaves 46,670,265 transactions, or 26.1 percent of the raw reports. The orientation and price filters described above remove another 2.1 percent, leaving 45,671,536 transactions. Requiring at least one G15 counterparty reduces the sample to 38,441,722 transactions. Finally, excluding intra-group transactions leaves 27,291,457 forwards, or 15.3 percent of the raw reports.

The intra-group restriction accounts for the largest reduction after the G15 filter. It removes 29.0 percent of transactions involving a G15 counterparty. Internal transfers between legal entities within a banking group are common in the EMIR data, but they need not represent arm's-length prices and therefore do not provide the dealer comparison required by our design.

The final sample contains 27,291,457 forwards over 58 months, or 470,542 transactions per month on average. Across the 14 dealer parents, each dealer accounts on average for 1,198,150 dealer-side transactions in the sell-FC role and 1,156,300 in the buy-FC role over the sample, corresponding to roughly 20,700 and 19,900 dealer-side transactions per dealer-role-month. The monthly wedges in Section IV are therefore estimated from many executed prices within each dealer-month rather than from a small number of trades.

The dependent variable in the transaction-level regressions is $\log F_i$, where F_i is the executed forward rate in dollars per unit of foreign currency. For a contract executed on date t_i and maturing on date T_i, we measure time to maturity in years as

$$\tau_i = \frac{\lfloor T_i - t_i \rfloor + 1}{365}, \tag{11}$$

so that a contract maturing on its execution date has one day of remaining maturity. We divide each trading day into 48 half-hour execution windows and assign contracts to ten-day maturity buckets,

$$\ell_i = \mathrm{ceil}(\tau_i/(10/365)) \,. \tag{12}$$

Equation (9) then compares transactions within half-hour execution-window × currency-pair × ten-day maturity-bucket comparison groups. These comparison groups are sufficiently narrow to absorb the common forward curve while leaving variation in the dealers intermediating otherwise similar trades.

Figure B.1 documents the variation underlying this comparison. Panel A counts transactions within each comparison group. Roughly one-quarter of comparison groups contain exactly two transactions, while about 15 percent contain at least 17. Thus many comparison groups contain several executed prices that can be compared after holding fixed currency pair, execution window, and maturity bucket. Panel B shows the maturity variation within execution-window $\times$ currency-pair groups: the median group spans five ten-day maturity buckets. Panels C and D ask whether those comparison groups also contain variation in dealer identity. On the buy-FC side, about 26 percent of comparison groups contain one G15 dealer parent and about 31 percent contain two; the sell-FC side has a similar distribution. The data therefore contain the within-cell transaction and dealer variation required to estimate equation (9).

Finally, for the funding tests in Section V, we supplement the forward transactions with confidential UK SFTR reports accessed at the Bank of England. These reports provide dollar repo transactions from which we construct dealer-level USD borrowing and lending rates. We aggregate the rates to the same parent groups used in the FX data and merge them to forward transactions by dealer and date. The resulting rates provide dealer-date proxies for the secured dollar borrowing and lending terms that enter the hedge portfolios in equation (4). Thus the EMIR data identify where each dealer transacts in the forward market, while the SFTR data provide an observable measure of the dealer's contemporaneous secured dollar conditions.

IV Results

The central empirical question is whether otherwise comparable FX forwards converge to a common price across dealers. They do not. Dealer wedges are dispersed across dealer-months, and the dispersion is asymmetric: the sell-FC side—on which the dealer buys dollars—is priced above the buy-FC side. These differences persist over time, survive controls for client relationships and transaction characteristics, and become substantially larger at short maturities. We decompose each dealer-month wedge into the sum $L_{d,t}$ and the difference $\Delta_{d,t}$ defined in equation (5) and examine both throughout this section: $\Delta_{d,t}$ is statistically and economically significant in every specification we consider, and $L_{d,t}$, though significant unconditionally, is shown below to reflect markup rather than differential dollar funding conditions across dealers.

We establish these facts using the two implementations of equation (9) introduced in Section II. The full-sample estimation produces one buy-FC and one sell-FC wedge for each dealer parent and is useful for displaying persistent cross-dealer differences. The monthly estimation produces the panel $\{\hat{\omega}^j_{d,t}\}$ used to measure dispersion and time variation. Throughout this section, we report $10{,}000\,\hat{\omega}$, so wedges are annualized basis points.

Dealer dispersion and the direction gap. Dealer prices differ substantially even after the common forward curve has been absorbed. Table 1 summarizes the monthly wedges. On the buy-FC side, the 10th and 90th percentiles are -4.34 and 3.54 basis points per year; on the sell-FC side they are -1.26 and 5.44. The corresponding standard deviations are 4.39 and 3.93 basis points. Forward contracts therefore exhibit a dispersed distribution of prices across dealer-months.

Removing month fixed effects leaves standard deviations of 3.74 basis points for buy-FC wedges and 3.22 for sell-FC wedges, so substantial cross-dealer variation remains after common monthly conditions are removed. Conversely, removing dealer fixed effects leaves standard deviations of 4.10 and 3.86 basis points, showing substantial movement within dealers over time. So, the dispersion is not simply a ranking of permanently expensive and cheap dealers. Removing both month and dealer effects still leaves standard deviations of 3.39 and 3.14 basis points. Consistent with these patterns, month effects account for 27 percent of buy-FC variance and 33 percent of sell-FC variance, whereas dealer effects account for 13 percent and 3 percent. Most observed dispersion therefore occurs at the dealer-month level rather than through a fixed ranking of dealer types.

The first question is whether the two sides of a dealer's forward book are priced differently and whether their average level indicates differential dollar funding conditions across dealers. The average sell-FC wedge—the side on which the dealer buys dollars—is 2.15 basis points per year, compared with -0.38 basis points on the buy-FC side. We summarize these two means using the difference and level defined in Section II. To estimate both objects, we pool the monthly dealer-role wedges and estimate

$$\hat{\omega}^{j}_{d,t} = \alpha + \beta\,\mathbb{1}\{j = \text{sell}\} + \varepsilon_{djt}. \tag{13}$$

The indicator equals one for a sell-FC observation and zero for a buy-FC observation. Thus α is the average buy-FC wedge and $\alpha + \beta$ is the average sell-FC wedge. It follows that $\hat{\beta}$ estimates the average difference, $\Delta_{d,t} = \omega^{\text{sell}}_{d,t} - \omega^{\text{buy}}_{d,t}$, while $2\hat{\alpha} + \hat{\beta}$ estimates the average level, $L_{d,t} = \omega^{\text{sell}}_{d,t} + \omega^{\text{buy}}_{d,t}$. Standard errors are two-way clustered by dealer parent and month, allowing observations for the same dealer to be correlated over time and observations in the same month to be correlated across dealers.

The two objects are both positive in the baseline specification, but they have different economic interpretations. The average difference is 2.52 basis points per year, with a standard error of 0.30. Dealers therefore transact at systematically higher prices on the buy-dollar side of their books. As discussed in Section II, the funding component of this difference combines borrowing-minus-lending spreads in dollars and foreign currency, so it captures the round-trip asymmetry between the two sides of intermediation rather than dollar funding conditions alone.

The average level is 1.78 basis points per year, with a standard error of 0.81. Unlike the difference, the funding component of the level nets the dealer's dollar borrowing and lending spreads against its foreign-currency spreads. A positive level is therefore consistent with dealers facing relatively more expensive dollar conditions than the counterparties against which their wedges are measured. At this stage, however, both the difference and the level are computed from the baseline wedges and therefore also contain the residual-demand component in equation (2).

The wedges entering equation (13) are themselves estimated from transactions, which raises a natural question: could estimation noise generate either the positive difference or the positive level? Under a standard measurement-error benchmark, it cannot. Write an estimated monthly wedge as

$$\hat{\omega}^{j}_{d,t} = \omega^{j}_{d,t} + e^{j}_{d,t}, \qquad \mathbb{E}\big[e^{j}_{d,t} \mid \omega^{j}_{d,t}\big] = 0, \tag{14}$$

where $\omega^{j}_{d,t}$ is the underlying dealer wedge and $e^{j}_{d,t}$ is estimation error. Mean-zero error makes the estimated wedges noisier but does not systematically shift their means. It therefore cannot generate either a positive sell-FC minus buy-FC difference or a positive sell-FC plus buy-FC level in expectation. The estimated difference of 2.52 basis points and level of 1.78 basis points are consequently not mechanical consequences of estimating the monthly wedges rather than observing them directly. Estimation noise affects the precision of both objects, but not their expected values.

Estimation error matters differently when we ask whether dealer wedges persist over time. Panel B of Table 1 relates each monthly wedge to the same dealer's wedge in the previous month. The pooled AR(1) coefficient is 0.411; estimated separately by role, it is 0.435 for buy-FC wedges and 0.388 for sell-FC wedges. Thus dealers with relatively high wedges in one month tend to have relatively high wedges in the next.

Because the monthly wedges are estimated from different sets of transactions, sampling noise should not persist mechanically from one month to the next. To make this argument explicit, suppose that the estimation errors in equation (14) are uncorrelated across adjacent monthly samples,

$$\mathbb{E}\big[e^{j}_{d,t} e^{j}_{d,t-1}\big] = 0. \tag{15}$$

Under this benchmark, estimation error increases the variance of the measured wedge but does not generate covariance between consecutive months. As a result, measurement error attenuates the estimated AR(1) coefficient toward zero. The observed persistence therefore provides evidence of persistence in the underlying dealer wedges rather than persistence generated by sampling noise.

This attenuation result also gives a conservative way to assess how much of the observed dispersion can be signal rather than estimation error. Let

$$\lambda = \frac{\mathrm{Var}(\omega)}{\mathrm{Var}(\hat{\omega})}$$

denote the share of the variance of estimated wedges attributable to variation in the underlying wedges. Under equations (14) and (15), the observed AR(1) coefficient satisfies

$$\hat{\rho} \approx \rho_\omega \lambda, \tag{16}$$

where ρ_ω is the true persistence of the underlying wedge. Since persistence cannot exceed one, $\rho_\omega \leq 1$, equation (16) implies

$$\lambda \geq \hat{\rho}. \tag{17}$$

Using the pooled estimate $\hat{\rho} = 0.411$, at least 41.1 percent of the observed variance is therefore attributable to variation in the underlying wedges under this benchmark. This is deliberately conservative: the bound reaches 41.1 percent only if the underlying wedge is perfectly persistent from one month to the next. Any persistence below one implies a larger signal share.

We can translate this variance bound into a bound on economically significant price dispersion. Panel A reports raw standard deviations of 4.39 basis points per year for buy-FC wedges and 3.93 for sell-FC wedges. Because standard deviations are the square roots of variances, a signal share of at least 0.411 implies that the standard deviation of the underlying wedges is at least $\sqrt{0.411} = 0.641$ times the observed standard deviation. The resulting lower bounds are $0.641 \times 4.39 = 2.81$ basis points per year for buy-FC wedges and $0.641 \times 3.93 = 2.52$ basis points per year for sell-FC wedges. Thus even an adjustment designed to attribute as much of the observed dispersion as possible to estimation noise leaves substantial dealer-level price dispersion. We return to this bound below once client relationships are held fixed.

The directional difference is also visible within individual dealer parents. Figure 1 plots the full-sample buy-FC wedge on the horizontal axis against the sell-FC wedge on the vertical axis. Every dealer lies above the 45-degree line, and most confidence intervals exclude that line. Thus the aggregate direction gap is not generated by one or two dealers: for the

Table 1

Monthly Dealer-Role Wedges: Direction Gap, Dispersion, and Persistence

Panel A: Distribution of monthly dealer-role wedge estimates

Role	N	Mean	Median	P10	P90	SD	SD net month FE	SD net dealer FE	SD net month and dealer FE
Buy FC	812	-0.376	-0.018	-4.34	3.54	4.385	3.735	4.095	3.388
Sell FC	812	2.147	2.165	-1.26	5.44	3.931	3.221	3.864	3.139

Sell-FC minus buy-FC difference, equation (13)

Specification	Sell − Buy	SE	N
Baseline	2.524***	0.296	1,624

Sell-FC plus buy-FC sum, equation (13)

Specification	Sell + Buy	SE	N
Baseline	1.782**	0.814	1,624

Panel B: AR(1) persistence of monthly dealer-role wedges

Statistic	Pooled	Buy FC	Sell FC
ρ	0.411***	0.435***	0.388***
Std. error	(0.098)	(0.134)	(0.072)
Observations	1,596	798	798

Note.– The unit of observation in Panel A and in the sell-FC minus/plus buy-FC blocks is a dealer parent by calendar month by role. Panel A reports the distribution of monthly dealer-by-role wedge estimates from equation (9), estimated separately by month. Units are annualized basis points. Buy FC denotes the dealer that buys foreign currency forward, sells dollars forward, and delivers dollars at maturity. Sell FC denotes the dealer that sells foreign currency forward, buys dollars forward, and receives dollars at maturity. Percentile columns report the cross-sectional distribution of monthly dealer-role wedge estimates. "SD net month FE" is the standard deviation of the residual after removing calendar-month fixed effects from the dealer-role wedge panel. "SD net dealer FE" is defined analogously after removing dealer-parent fixed effects. "SD net month and dealer FE" removes both sets of fixed effects. The standard deviations are not adjusted for sampling error in the monthly wedge estimates and are therefore upper bounds on the dispersion of the underlying wedges. The sell-FC minus buy-FC block reports the coefficient β on a sell-FC indicator from equation (13), equal to $\Delta_{d,t}$ in equation (5). The sell-FC plus buy-FC block reports $2\hat{\alpha} + \hat{\beta}$ from the same equation, equal to $L_{d,t}$. Standard errors in both blocks are two-way clustered by dealer parent and month. Panel B reports AR(1) regressions of monthly dealer-role wedges on their one-month lag. Standard errors in Panel B are clustered by dealer parent. *** 1%, ** 5%, * 10%.

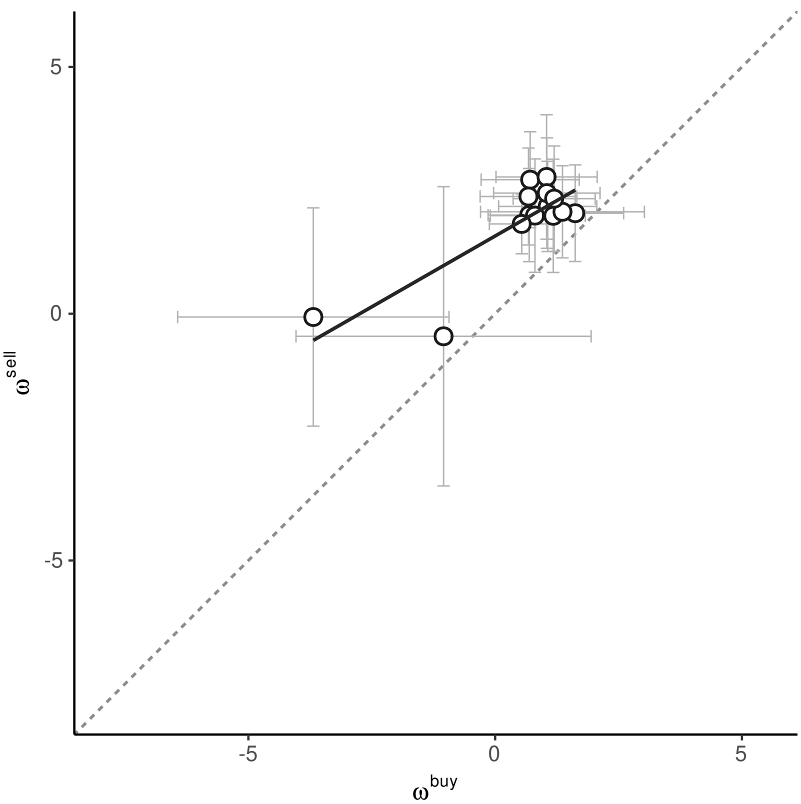

Figure 1. Direction Gap in Full-Sample Dealer Wedge Estimates. *Notes.* Each point is one dealer parent. The horizontal axis reports the full-sample buy-FC wedge $\hat{\omega}_d^{\text{buy}}$, and the vertical axis reports the full-sample sell-FC wedge $\hat{\omega}_d^{\text{sell}}$. Both wedges are estimated from equation (9) on the full transaction sample and are reported in annualized basis points. Standard errors are two-way clustered by buy-FC dealer parent and sell-FC dealer parent. Horizontal and vertical bars are 95% confidence intervals. The dashed line is the 45-degree line, where buy-FC and sell-FC wedges are equal. Points above the line have higher sell-FC wedges than buy-FC wedges. The solid line is the fitted regression of the full-sample sell-FC wedge on the full-sample buy-FC wedge across dealer parents.

observed dealer parents, the buy-dollar side is consistently priced above the sell-dollar side. The figure also shows why a single dealer intercept cannot capture the pattern. A dealer can have a positive sell-FC wedge and a near-zero or negative buy-FC wedge. Dealer identity and dealer direction therefore both matter for executed prices.

Client composition and transaction costs. A natural alternative explanation is that dealer wedges capture differences in customers or execution conditions rather than differences associated with the dealers themselves. Dealers serve different clients, and persistent client characteristics could therefore be attributed to dealer identity. Dealers also intermediate different transactions: negotiated spreads tend to be narrower for larger tickets, electronic execution, and more liquid currency pairs. We address these two concerns separately and then jointly.

To absorb persistent client composition, we first add client LEI × dealer-parent fixed effects to equation (9). This specification compares a dealer with its baseline only after allowing every client-dealer relationship to have its own average pricing level. We then tighten the comparison further with client LEI × dealer-parent × currency-pair fixed effects, allowing the pricing level of a given relationship to differ across currencies. To address observable transaction conditions, we add indicators for notionals above the pair-month median, electronic execution, and the three most liquid dollar pairs in the sample. Appendix C describes these controls and the corresponding specifications.

For each dealer d, role j, and month t, let $\hat{\omega}_{d,t}^{j,\text{base}}$ denote the baseline wedge and $\hat{\omega}_{d,t}^{j,\text{alt}}$ the wedge estimated under an alternative specification. We define

$$\Delta_{d,t}^{j} = \hat{\omega}_{d,t}^{j,\text{alt}} - \hat{\omega}_{d,t}^{j,\text{base}}. \tag{18}$$

A negative $\Delta_{d,t}^{j}$ means that the additional controls lower the estimated wedge. We estimate its average using

$$\Delta_{d,t}^{j} = \alpha + \varepsilon_{d,t}^{j}, \tag{19}$$

with standard errors two-way clustered by dealer parent and month. Every row of Table 2 uses the same 1,624 dealer-role-month observations, so differences across specifications are not driven by changes in sample composition. The controls affect the level and the difference very differently. The level, $L_{d,t} = \omega_{d,t}^{\text{sell}} + \omega_{d,t}^{\text{buy}}$, falls substantially as we absorb transaction characteristics and client-dealer relationships, while the difference, $\Delta_{d,t} = \omega_{d,t}^{\text{sell}} - \omega_{d,t}^{\text{buy}}$, remains close to its baseline value.

Transaction characteristics alone reduce the average dealer-role wedge by 0.42 basis points per year; client LEI × dealer-parent fixed effects reduce it by 0.48 basis points, and allowing those relationship effects to vary by currency pair reduces it by 0.54 basis points. Combining the transaction controls with the two client specifications reduces it by 0.74 and 0.79 basis points, respectively. The controls therefore remove a substantial component of the average level without materially changing the directional difference.[10]

Panel B shows how little the directional difference changes. The baseline sell-FC minus buy-FC gap is 2.52 basis points per year. Transaction controls raise it to 2.66, while the two client-relationship specifications yield gaps of 2.23 and 2.31 basis points. Adding transaction controls to those specifications yields 2.39 and 2.44 basis points. Every alternative estimate therefore remains within 0.31 basis points of the baseline. Persistent client-dealer compo-

[10] Because wedges are measured relative to an omitted category, a shift common to both roles changes their level relative to that category. The difference $\Delta_{d,t}$ is invariant to such a common shift.

sition and observable transaction characteristics do not account for the higher price on the dealer buy-dollar side.

Panel C shows what happens to the level. The baseline sell-FC plus buy-FC wedge is 1.78 basis points per year. Transaction-cost controls reduce it to 0.95 basis points, while client LEI $\times$ dealer fixed effects and client LEI $\times$ dealer $\times$ currency-pair fixed effects reduce it to 0.82 and 0.70 basis points, respectively. Combining the client controls with transaction characteristics leaves the level at 0.30 and 0.20 basis points. The level therefore becomes statistically indistinguishable from zero once we absorb observable characteristics associated with residual demand. This contrast between Panels B and C motivates our interpretation: the common component of dealer pricing is closely related to clientele and transaction conditions, whereas the directional difference survives these controls.

The economic interpretation of the level follows from equation (8). Its funding component compares a dealer's dollar funding conditions with its foreign-currency funding conditions and, because wedges are measured relative to otherwise comparable transactions, with those of the counterparties against which the dealer is estimated. The level therefore does not measure whether dollar funding is expensive in aggregate—the common cost of synthetic dollars is absorbed by the fixed effects. Instead, it asks whether dollar funding conditions differ systematically across dealers. Once we absorb client relationships and observable transaction characteristics associated with residual demand, the level is close to zero and statistically insignificant. The static cross section therefore provides little evidence of persistent differences in relative dollar funding conditions across major dealers.

Table 1 reports a pooled AR(1) coefficient of 0.411 in the baseline specification, which does not hold client relationships fixed or include transaction-cost controls. Part of that persistence could therefore reflect a dealer serving the same clients from month to month rather than a factor specific to the dealer itself. Table 3 re-estimates the same AR(1) regression on the wedges from each specification in Table 2. Absorbing client-relationship fixed effects lowers the pooled coefficient from 0.411 to 0.300, and it stays close to that value, between 0.29 and 0.36, once transaction-cost controls or currency-pair-specific relationship effects are added on top. About a quarter of the raw persistence in the baseline wedge is therefore attributable to persistent client-dealer relationships; a coefficient of roughly 0.30, still statistically significant, survives once relationships are held fixed.

A common dealer component across the two sides. The preceding results show that client relationships explain part of the level and persistence of dealer wedges but leave substantial dealer-level variation. We next ask whether this remaining variation is shared across the two sides of the same dealer's book. If the co-movement simply reflects the

Table 2

CLIENT COMPOSITION, TRANSACTION-COST CONTROLS, AND THE DIRECTION GAP

Panel A: Average change in monthly dealer-role wedges relative to baseline

Specification	Mean Δ	SE	N
Transaction-cost controls	−0.415***	(0.083)	1,624
Client LEI × dealer FE	−0.481**	(0.189)	1,624
+ transaction-cost controls	−0.742***	(0.216)	1,624
Client LEI × dealer × pair FE	−0.542**	(0.220)	1,624
+ transaction-cost controls	−0.792***	(0.246)	1,624

Panel B: Sell-FC minus buy-FC wedge difference

Specification	Sell − Buy	SE	N
Baseline	2.524***	(0.296)	1,624
Transaction-cost controls	2.658***	(0.273)	1,624
Client LEI × dealer FE	2.225***	(0.266)	1,624
+ transaction-cost controls	2.385***	(0.243)	1,624
Client LEI × dealer × pair FE	2.307***	(0.277)	1,624
+ transaction-cost controls	2.443***	(0.254)	1,624

Panel C: Sell-FC plus buy-FC wedge sum

Specification	Sell + Buy	SE	N
Baseline	1.782**	(0.814)	1,624
Transaction-cost controls	0.951	(0.726)	1,624
Client LEI × dealer FE	0.819	(0.689)	1,624
+ transaction-cost controls	0.297	(0.610)	1,624
Client LEI × dealer × pair FE	0.698	(0.706)	1,624
+ transaction-cost controls	0.197	(0.634)	1,624

Note.– The unit of observation is a dealer parent by calendar month by role, and every specification is estimated on the same 1,624 observations. Panel A reports pooled estimates of equation (19): the dependent variable is the change $\Delta^j_{d,t}$ in the monthly dealer-role wedge relative to the baseline specification, as defined in equation (18), so negative values indicate that the listed specification estimates a lower wedge. Panels B and C report estimates of β and of $2\alpha + \beta$, respectively, from equation (13), with the wedges re-estimated under each specification; these equal $\Delta_{d,t}$ and $L_{d,t}$ from equation (5). Transaction-cost controls are three trade-level proxies for the conditions under which negotiated bid-ask spreads narrow: an indicator for notional above the pair-month median, an indicator for execution on venue rather than bilaterally, and an indicator for EUR/USD, GBP/USD and JPY/USD, the three most liquid currency pairs in our sample. Appendix C gives details on the construction of the transaction-cost controls and the estimation specification. Client LEI × dealer FE denotes fixed effects for the client-counterparty legal entity identifier interacted with the dealer parent; Client LEI × dealer × pair FE further allows the client-dealer effect to vary by currency pair. For a dealer-role observation the client counterparty is the opposite counterparty in the trade. Units are annualized basis points. Standard errors, in parentheses, are two-way clustered by dealer parent and month. *** 1%, ** 5%, * 10%.

Table 3

AR(1) PERSISTENCE OF DEALER-ROLE WEDGES ACROSS SPECIFICATIONS

Statistic	Pooled	Buy FC	Sell FC
Baseline			
ρ	0.411***	0.439***	0.383***
Std. error	(0.097)	(0.132)	(0.072)
Observations	1,596	798	798
Transaction-cost controls			
ρ	0.403***	0.461***	0.344***
Std. error	(0.113)	(0.140)	(0.089)
Observations	1,596	798	798
Client LEI × dealer FE			
ρ	0.300***	0.289**	0.310***
Std. error	(0.092)	(0.113)	(0.080)
Observations	1,596	798	798
+ *transaction-cost controls*			
ρ	0.304**	0.310**	0.298***
Std. error	(0.102)	(0.122)	(0.090)
Observations	1,596	798	798
Client LEI × dealer × pair FE			
ρ	0.291***	0.330**	0.249***
Std. error	(0.081)	(0.113)	(0.071)
Observations	1,596	798	798
+ *transaction-cost controls*			
ρ	0.299***	0.360**	0.231**
Std. error	(0.089)	(0.120)	(0.078)
Observations	1,596	798	798

Note.– Each block reports the AR(1) coefficient of monthly dealer-role wedges on their first lag, estimated under the corresponding specification from Table 2. Standard errors, in parentheses, are clustered by dealer parent. *** 1%, ** 5%, * 10%.

Table 4
Buy–Sell Linkage in Dealer Wedges Across Specifications

Specification	Slope $\hat{\beta}$	SE	N	Adjusted R^2
Baseline	0.532***	(0.111)	812	0.550
Transaction-cost controls	0.547***	(0.108)	812	0.501
Client LEI × dealer FE	0.544***	(0.136)	812	0.492
+ transaction-cost controls	0.559***	(0.128)	812	0.448
Client LEI × dealer × pair FE	0.519***	(0.119)	812	0.510
+ transaction-cost controls	0.536***	(0.110)	812	0.470

Note.– The table reports the pooled panel version of equation (10): a regression of sell-FC wedges on buy-FC wedges with month fixed effects, re-estimated on the wedges from each specification in Table 2. The unit of observation is a dealer parent by month. Both variables are monthly dealer-role wedge estimates, in annualized basis points. Standard errors are two-way clustered by dealer parent and month. *** 1%, ** 5%, * 10%.

customers a dealer serves or the transactions it intermediates, it should weaken once we absorb those characteristics. If instead a common dealer component contributes to both sides, the relationship should survive these adjustments.

We test this implication by estimating equation (10), relating sell-FC wedges to buy-FC wedges with month fixed effects. We repeat the regression for the baseline wedges and for each set of wedges estimated with the client-relationship fixed effects and transaction-cost controls in Table 2. Table 4 reports the results.

The two sides move together, and the relationship is remarkably stable across specifications. In the baseline, the slope is 0.532 with an adjusted R^2 of 0.550: within a month, dealers with relatively high buy-FC wedges also tend to have relatively high sell-FC wedges. The slope remains between 0.52 and 0.56 across all five alternative specifications. In particular, controlling for client-dealer relationships leaves the slope essentially unchanged. The co-movement between the two sides therefore does not appear to arise simply because the same dealers repeatedly serve customers with similar pricing characteristics. It points instead to a common dealer component in transaction prices.

Classical sampling error in the estimated buy-FC wedge attenuates the coefficient toward zero. Yet, the results provide strong evidence that the two sides of a dealer's book share substantial variation. As regards the exact magnitude of the slope, it should not be given a structural interpretation: we cannot conclude that a one-basis-point change on one side mechanically produces a 0.53-basis-point change on the other.

Dynamics, factor structure, and maturity of the baseline wedges. We next return to the *baseline* dealer wedges to characterize their time-series dynamics, factor structure, and maturity profile. These exercises describe the unconditional distribution of dealer transaction prices before absorbing client relationships or transaction characteristics. We first examine how average dealer prices move over time and how much of the cross-sectional variation is common across dealers, and then ask where along the maturity structure the directional asymmetry is concentrated.

We first trace the average baseline wedge through time. For each role $j \in \{\text{buy}, \text{sell}\}$, we estimate

$$\hat{\omega}^j_{d,t} = \sum_{\mu} \theta^j_\mu \, \mathbb{1}\{t = \mu\} + u^j_{d,t}, \tag{20}$$

where μ indexes calendar months. Because the regression contains month indicators but no dealer fixed effects, $\hat{\theta}^j_\mu$ is the equal-weighted average baseline wedge across dealer parents for role j in month μ. Figure 2 plots these averages and 95 percent confidence intervals, with standard errors clustered by dealer parent.

The time series reinforces the directional result. Sell-FC wedges generally lie above buy-FC wedges and display larger month-to-month movements, while buy-FC wedges remain closer to zero through much of the sample. More of the variation in average dealer pricing therefore occurs on the side on which dealers buy dollars.[11]

One concern is that equal-weighted averages give small dealer-months the same influence as large ones. Figure 3 therefore weights each dealer-role-month by its share of monthly notional on the relevant side of the market. The resulting series closely resembles the equal-weighted series: the sell-FC wedge remains elevated relative to the buy-FC wedge. The directional time-series pattern is therefore not driven by low-volume dealer-months.

The baseline dealer-month panel also contains substantial common variation across dealers. Figure 4 decomposes each role-specific panel into principal components. The first component explains 51 percent of the variance of buy-FC wedges and 46 percent of the variance of sell-FC wedges. No single component, however, accounts for the remaining variation. The baseline price distribution therefore contains an important common component alongside substantial dealer-specific variation.

The maturity structure places the directional asymmetry at the front end of the forward curve. Table 5 re-estimates the baseline monthly wedges separately for contracts with maturities up to one month, between one and three months, and at least three months. The average sell-FC wedge falls sharply with maturity, from 10.18 basis points per year at maturities up to one month to 3.65 between one and three months and 1.58 beyond three months. Buy-FC wedges are much closer to zero, ranging from −1.65 at short maturities to −0.49 at long maturities.

The difference between the two sides rises accordingly. The sell-FC minus buy-FC differ-

[11] Figures B.2 and B.3 report the monthly cross-sectional distribution of buy-FC and sell-FC wedges separately by currency pair. Dispersion is present in every core pair. It is widest in CHF/USD and JPY/USD and narrowest in GBP, and the sell-FC distribution sits above the buy-FC distribution in all six pairs.

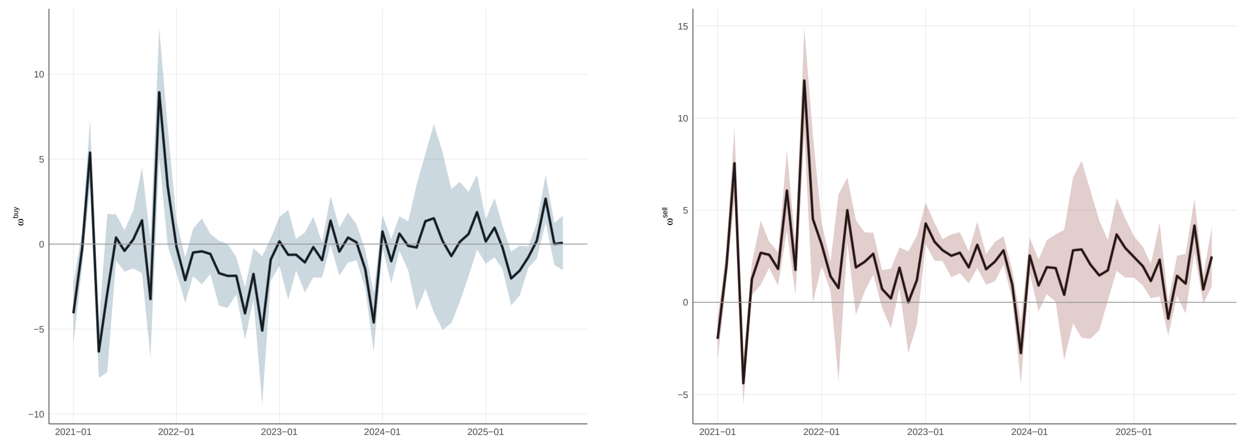

(A) Buy-FC wedge ω^{buy}

(B) Sell-FC wedge ω^{sell}

Figure 2. Monthly Dealer-Wedge Effects. *Notes.* The figure plots calendar-month fixed effects from separate role-specific regressions of monthly dealer-role wedges on month indicators, as in equation (20). The unit of observation is a dealer parent by calendar month by role. The regressions include no dealer fixed effects, so each point is the equal-weighted average wedge for the role in that month. The solid line is the estimated month effect. Shaded bands are 95% confidence intervals based on standard errors clustered at the dealer parent level. Units are annualized basis points.

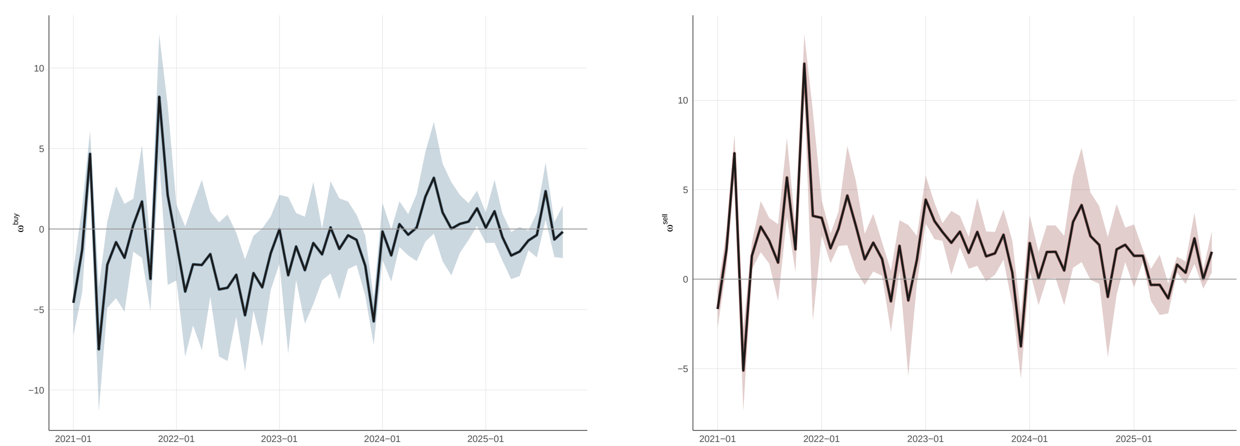

(A) Notional-weighted buy-FC wedge

(B) Notional-weighted sell-FC wedge

Figure 3. Notional-Weighted Monthly Dealer-Wedge Effects. *Notes.* The figure plots calendar-month fixed effects from separate role-specific weighted least-squares regressions of monthly dealer-role wedges on month indicators, as in equation (20). The unit of observation is a dealer parent by calendar month by role. Weights are each dealer parent's share of total monthly notional on the relevant side of the market. The regressions include no dealer fixed effects, so each point is the notional-weighted average wedge for the role in that month. The solid line is the estimated weighted month effect. Shaded bands are 95% confidence intervals based on standard errors clustered at the dealer parent level. Units are annualized basis points.

Table 5

Dealer Wedges by Maturity Bucket

Panel A: Monthly dealer-role wedges by maturity bucket

Role	Maturity bucket	N	Mean	Median	SD	SD net month FE	SD net dealer FE	SD net month and dealer FE
Buy FC	Short ($\leq$1m)	812	-1.646	-1.192	23.413	17.391	22.734	16.465
Buy FC	Medium (1–3m)	812	-0.443	-0.187	7.155	6.003	6.846	5.630
Buy FC	Long ($\geq$3m)	812	-0.489	-0.070	4.714	4.386	4.377	4.021
Sell FC	Short ($\leq$1m)	812	10.178	10.712	22.975	16.839	22.787	16.582
Sell FC	Medium (1–3m)	812	3.653	3.258	7.062	5.843	7.010	5.779
Sell FC	Long ($\geq$3m)	812	1.580	1.720	4.282	3.913	4.169	3.789

Panel B: Sell-FC minus buy-FC tests by maturity bucket

Maturity bucket	Sell − Buy	SE	N	Adjusted R^2
Short ($\leq$1m)	11.824***	2.292	1,624	0.060
Medium (1–3m)	4.096***	0.749	1,624	0.076
Long ($\geq$3m)	2.069***	0.270	1,624	0.050

Panel C: AR(1) persistence by maturity bucket

Maturity bucket	Statistic	Pooled	Buy FC	Sell FC
Short ($\leq$1m)	ρ	0.104	0.040	0.159***
	Std. error	(0.069)	(0.095)	(0.049)
	Observations	1,596	798	798
Medium (1–3m)	ρ	0.240***	0.227**	0.252***
	Std. error	(0.054)	(0.091)	(0.047)
	Observations	1,596	798	798
Long ($\geq$3m)	ρ	0.458***	0.499***	0.414***
	Std. error	(0.150)	(0.165)	(0.135)
	Observations	1,596	798	798

Note.– Panel A reports monthly dealer-by-role wedge estimates estimated separately by maturity bucket. The unit of observation is a dealer parent by calendar month by role by maturity bucket. Units are annualized basis points. Short maturities are contracts with time to maturity up to one month; medium maturities are contracts with time to maturity between one and three months; long maturities are contracts with time to maturity of at least three months. The residualized SD columns are defined as in Table 1 and are not adjusted for sampling error in the monthly wedge estimates. Panel B reports estimates of equation (13) estimated separately by maturity bucket, with standard errors two-way clustered by dealer parent and month. Panel C reports AR(1) regressions of monthly dealer wedges on their one-month lag within each maturity bucket. Standard errors in Panel C are clustered by dealer parent. *** 1%, ** 5%, * 10%.

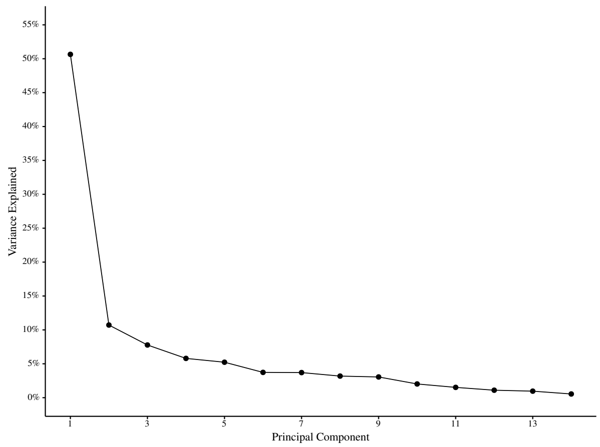

(A) Buy-FC wedges

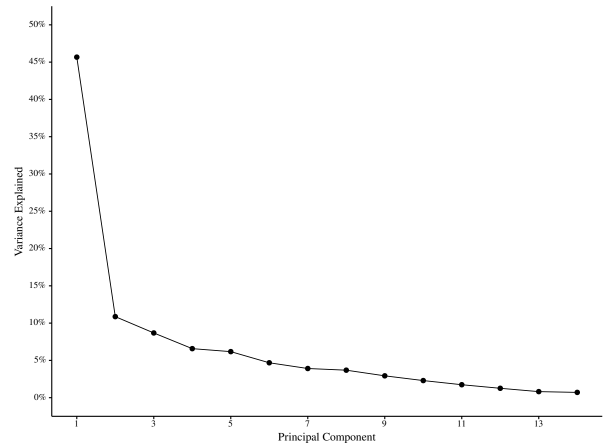

(B) Sell-FC wedges

Figure 4. Factor Structure of Monthly Dealer Wedges. *Notes.* The figure reports the share of variance explained by each principal component of the dealer-by-month wedge panel. Panel A uses buy-FC wedges, and Panel B uses sell-FC wedges. The horizontal axis indexes principal components ordered by explained variance. The vertical axis reports the share of total variance explained.

ence is 11.82 basis points per year within one month, 4.10 basis points between one and three months, and 2.07 basis points beyond three months. Because all wedges are annualized, this gradient is not mechanically generated by the shorter life of the contract: the annualized directional difference itself is more than five times as large within one month as beyond three months. The dealer buy-dollar premium is therefore concentrated at the front end.

The level also varies strongly with maturity. The sell-FC plus buy-FC wedge is 8.53 basis points per year within one month, 3.21 basis points between one and three months, and 1.09 basis points beyond three months. The maturity gradient in the baseline level should be interpreted with the same qualification as the unconditional level discussed above: it combines intermediation costs with residual-demand pricing. In particular, its larger value at short maturities does not by itself establish that differential dollar funding conditions are more important at the front end.

Finally, persistence varies systematically with maturity. The pooled AR(1) coefficient is 0.104 at short maturities, 0.240 at medium maturities, and 0.458 at long maturities. Persistence is therefore weakest where the directional gap is largest. This comparison requires caution because the monthly wedges need not be estimated with the same precision across maturity groups. Let λ denote the share of observed wedge variance attributable to the underlying wedge rather than estimation error. Under equation (16), $\hat{\rho} \approx \rho_\omega \lambda$, so greater estimation noise at short maturities lowers the measured AR(1) coefficient even if the persistence of the underlying dealer wedge is unchanged. The maturity pattern in Panel C should therefore be read as a pattern in the persistence of the estimated wedges, not necessarily as evidence that latent dealer prices become more persistent with maturity. In sum, we conclude the baseline wedges vary across time, across dealers, and sharply with maturity. Most importantly, the directional difference is largest at short maturities: the price gap between the sell-FC and buy-FC sides is more than five times as large within one month as beyond three months.

V Dollar Funding Shocks and Monetary Transmission

To study whether dealer-specific dollar conditions affect transaction prices, we use two direct measures of those conditions: secured dollar funding rates and scheduled dollar cashflows. These measures capture both the current cost of obtaining dollars and dealers' anticipated future dollar needs. We first ask whether dealers facing different dollar conditions price otherwise comparable forwards differently. We then examine whether dealer-specific dollar conditions shape the transmission of common funding shocks by studying heterogeneous responses to FOMC announcements, which provide a plausibly exogenous shock to the value of dollar funding. This approach allows us to test whether dealers with different funding needs exhibit different pass-through of the same monetary-policy surprise.

The evidence shows that dealers' own secured dollar rates covary with transaction prices, and both secured rates and scheduled dollar inflows predict heterogeneous responses around FOMC announcements. The announcement evidence is asymmetric: dealers scheduled to pay dollars out over the following week pass a tightening surprise through to forward prices more strongly than dealers scheduled to receive them. Measured against a dealer with no scheduled exposure, the difference comes from the inflow side—expected dollar inflows dampen pass-through, while expected outflows do not amplify it.

We study these relationships at the transaction level. Let $f_i = 10^4 \times \log F_i$ denote the log executed forward rate of trade i, expressed in basis points. Each dealer characteristic is measured at the dealer–day level and attached to the relevant side of the transaction. Because a trade can contain an observed dealer on either or both sides, we estimate the relationship separately by dealer role. Let Z_i denote the side-specific dealer characteristic used in a given specification.

We use two empirical comparisons. The first asks whether dealers facing different dollar conditions transact at different residual forward prices in the full sample:

$$f_i \;=\; \theta' Z_i \;+\; \mu_{w(i),c(i),m(i)} \;+\; \alpha_{s(i),c(i)} \;+\; \delta_{b(i),c(i)} \;+\; \varepsilon_i, \qquad (21)$$

where $w(i)$ is the execution window, $c(i)$ the currency pair, $m(i)$ the maturity bucket, $s(i)$

the sell-FC dealer, and $b(i)$ the buy-FC dealer. The execution-window $\times$ pair $\times$ maturity effects μ absorb the common forward curve. The coefficient θ is therefore identified from differences in dealer characteristics among transactions in the same currency pair, execution window, and maturity bucket rather than from aggregate movements in dollar rates. Our preferred specification additionally includes sell-FC $\times$ pair and buy-FC dealer $\times$ pair effects, α and δ. These fixed effects absorb persistent differences in where each dealer prices each currency pair, so θ is identified from changes in a dealer's characteristic relative to that dealer's own average for the pair.

The second comparison asks whether dealers with different pre-existing dollar conditions respond differently to the same monetary-policy news. We use the 39 FOMC statement dates in our sample. Let k index announcements, let STMT_k denote the high-frequency statement surprise of Acosta, Brennan, and Jacobson (2024), and let Post_i equal one for trades executed after the 2:00 p.m. New York release. For dealer characteristics $Z_{i,k-1}$ measured on the previous weekday, we estimate

$$\begin{aligned} f_i \;=\; & \beta' Z_{i,k-1}\,(\mathrm{STMT}_k \times \mathrm{Post}_i) \;+\; \gamma' Z_{i,k-1}\,\mathrm{Post}_i \\ & +\; \lambda\,(\mathrm{STMT}_k \times \mathrm{Post}_i) \;+\; \psi\,\mathrm{Post}_i \;+\; \alpha_{s(i),c(i),k} \;+\; \delta_{b(i),c(i),k} \;+\; \varepsilon_i. \end{aligned} \tag{22}$$

Equation (22) contains the full set of lower-order terms implied by the triple interaction. Dealer $\times$ pair $\times$ announcement-date fixed effects are included separately for the two sides of the transaction. These fixed effects absorb the levels of $Z_{i,k-1}$ and STMT_k, as well as their interaction. The coefficient of interest, β, measures whether the within-day price response to the same statement surprise differs across dealers according to their predetermined dollar conditions.

The FOMC sample contains contracts with no more than 30 days to maturity executed between 11:00 a.m. and 5:00 p.m. New York time on the 39 statement dates from January 2021 through October 2025. Because STMT_k is common to all trades on announcement date k, the identifying variation is at the dealer–event level. For a given dealer, currency pair, and announcement date, equation (22) compares forward prices before and after 2:00 p.m.; it then relates that pre–post difference to the dealer's predetermined exposure interacted with the announcement surprise. Standard errors are three-way clustered by date $\times$ pair, sell-FC dealer $\times$ pair $\times$ date, and buy-FC dealer $\times$ pair $\times$ date.

A. Repo Exposure

A dealer's secured dollar rate provides an observable proxy for the dollar terms it faces outside the forward market. We construct daily, volume-weighted dollar repo borrowing and lending rates from UK SFTR reports accessed at the Bank of England. All rates are expressed in percentage points. Following the hedge portfolios in Section II, we pair the sell-FC dealer—which receives dollars in the forward and must source dollars to hedge that position—with its repo borrowing rate. We pair the buy-FC dealer—which delivers dollars in the forward and places dollars in the hedge—with its repo lending rate.

The full-sample evidence is asymmetric. A dealer's repo lending rate is positively related to the forward price on the buy-FC side, and that relationship survives the dealer $\times$ pair controls. The analogous relationship between the sell-FC dealer's borrowing rate and the forward price disappears once the same controls are included.

Table 6 establishes this result by progressively tightening the comparison. Columns (1)–(3) relate the forward price to the sell-FC dealer's borrowing rate; columns (4)–(6) use the buy-FC dealer's lending rate. Columns (1) and (4) absorb execution-window $\times$ pair effects, columns (2) and (5) additionally absorb maturity, and columns (3) and (6) add dealer $\times$ pair effects on both sides of the trade. We treat columns (3) and (6) as the preferred specifications because they compare movements in a dealer's own repo rate with movements in its transaction prices after removing persistent dealer-pair pricing differences. The samples differ across the two sides because each specification requires the corresponding dealer's repo rate to be observed.

On the buy-FC side, the coefficient falls from 5.00 in column (4) to 3.17 in column (5) and 1.12 in column (6). Thus part of the raw relationship reflects maturity composition and persistent differences across dealers and currency pairs. Even after absorbing both, however, a one-percentage-point increase in the buy-FC dealer's repo lending rate is associated with a 1.12 basis point higher log forward price, with a standard error of 0.21. This sign matches equation (4): a buy-FC dealer that can place dollars at a higher rate is willing, other things equal, to pay more for foreign currency forward.

The sell-FC side does not show the corresponding within-dealer relationship. Its repo borrowing coefficient falls from 3.75 in column (1) to 0.36 after controlling for maturity and to -0.19 in column (3), with a standard error of 0.50. The latter estimate is close to zero and statistically insignificant. The positive raw association on this side therefore disappears once the comparison absorbs maturity and persistent dealer-pair pricing differences.[12]

We conclude that observable secured dollar conditions covary with residual forward prices, but the relationship that survives dealer-pair controls appears on the side whose hedge places

[12] Appendix A implies that dollar borrowing pressure enters through the sell-FC side, since hedging that side requires sourcing dollars. Two readings are consistent with the null result. The reported volume-weighted daily repo rate may be a poor proxy for the marginal cost of dollars at the moment of execution, since the dealer need not fund the hedge at the daily average. Alternatively, the sell-FC price may be set against residual customer demand rather than at the hedge boundary, in which case the funding term is dominated by the residual-demand component m_{djt} in equation (2).

Table 6

REPO RATES AND FORWARD PRICES

	Dependent variable: f_i, log forward rate (bps)					
	(1)	(2)	(3)	(4)	(5)	(6)
Sell-FC dealer repo borrowing rate$_t$	3.75***	0.36	−0.19			
	(1.20)	(0.50)	(0.50)			
Buy-FC dealer repo lending rate$_t$				5.00***	3.17***	1.12***
				(1.63)	(0.88)	(0.21)
Fixed effects						
Window × pair	Yes			Yes		
Window × pair × maturity		Yes	Yes		Yes	Yes
Sell-FC dealer × pair			Yes			Yes
Buy-FC dealer × pair			Yes			Yes
Observations	9,718,907	9,718,907	9,718,907	7,912,633	7,912,633	7,912,633

Note.– The table estimates equation (21), $f_i = \theta' Z_i + \mu_{w(i)c(i)m(i)} + \alpha_{s(i)c(i)} + \delta_{b(i)c(i)} + \varepsilon_i$. The unit of observation is an FX-forward transaction in core dollar pairs. The dependent variable f_i is the log executed forward rate in the canonical dollars-per-foreign-currency quote, in basis points. The repo rate is the dealer's volume-weighted average dollar repo rate on the trade date, in percentage points. Each side is paired with the rate that enters its hedged break-even quote: the borrowing rate is attached to the sell-FC dealer, which must source dollars, and the lending rate to the buy-FC dealer, which places them. Samples differ across columns because each requires the relevant dealer's rate to be observed. Standard errors, in parentheses, are three-way clustered by date × pair, sell-FC dealer × pair, and buy-FC dealer × pair. *** 1%, ** 5%, * 10%.

dollars rather than on the side whose hedge sources them.

The announcement design asks a different question: do dealers facing different secured dollar conditions transmit the same monetary-policy surprise differently? They do. Table 7 uses repo rates measured on the weekday preceding each announcement. Column (1) attaches the sell-FC dealer's borrowing rate; column (2) attaches the buy-FC dealer's lending rate. The corresponding samples contain 108,549 and 63,788 transactions.

The triple interactions are negative on both sides. A one-percentage-point higher pre-announcement repo borrowing rate at the sell-FC dealer changes the post-announcement response to a one-unit statement surprise by −90.4 basis points, with a standard error of 53.4. The corresponding coefficient for the buy-FC dealer's lending rate is −116.8, with a standard error of 39.5. The second estimate is considerably more precise; the first provides weaker evidence of the same sign. The two regressions show that the price response to a common monetary-policy surprise varies systematically with the secured dollar conditions faced by the dealer before the announcement.

The sign has a direct interpretation under the canonical quote convention. Holding spot and foreign-currency rates fixed, a higher forward rate corresponds to a higher implied dollar interest rate. The negative triple interactions therefore imply that, following a positive tightening surprise, dealers entering the announcement with higher own secured dollar rates

Table 7

Repo Rates and FOMC Pass-Through

	Dependent variable: f_i, log forward rate (bps)	
	(1) Sell-FC	(2) Buy-FC
Triple interaction: repo rate$_{k-1}$ × STMT × Post		
Sell-FC dealer borrowing rate	-90.4^{*}	
	(53.4)	
Buy-FC dealer lending rate		-116.8^{***}
		(39.5)
Lower-order terms		
Repo rate$_{k-1}$ × Post	0.09	-2.01^{*}
	(1.54)	(1.20)
STMT × Post	−183.8	−122.7
	(237.4)	(143.0)
Post	6.22	14.79^{***}
	(6.80)	(4.90)
Fixed effects		
Sell-FC dealer × pair × date	Yes	Yes
Buy-FC dealer × pair × date	Yes	Yes
Observations	108,549	63,788

Note.– The table estimates equation (22). The levels of the repo rate and of STMT, and their product, are absorbed by the fixed effects; all remaining terms are reported. Columns differ only in which repo rate enters $Z_{i,k-1}$, each side being paired with the rate that enters its hedged break-even quote: column (1) the sell-FC dealer's borrowing rate, since hedging that side requires sourcing dollars, and column (2) the buy-FC dealer's lending rate, since hedging that side places them. Samples differ because each column requires a different dealer's rate to be observed. The sample is FX-forward transactions in core dollar pairs with time to maturity of thirty days or less, executed between 11:00 a.m. and 5:00 p.m. New York time on FOMC statement dates. The dependent variable f_i is the log executed forward rate in the canonical dollars-per-foreign-currency quote, in basis points. Repo rates are volume-weighted average dollar rates in percentage points, measured on the weekday preceding the announcement. STMT is the high-frequency FOMC statement surprise; Post indicates trades executed after the 2:00 p.m. release. Standard errors, in parentheses, are three-way clustered by date × pair, sell-FC dealer × pair × date, and buy-FC dealer × pair × date. *** 1%, ** 5%, * 10%.

increase the forward price by less than dealers entering with lower rates. The evidence is about *differential* pass-through across dealers; it does not imply that a higher repo rate itself causes weaker monetary transmission.[13]

B. Dollar Exposure

Repo rates measure the terms on which dealers transact in secured dollar markets. We next use the forward book itself to measure whether dollars are scheduled to arrive at, or leave, a dealer in the near future. This measure asks whether a dealer's predetermined dollar position

[13] For the same reason we do not read the common STMT×Post coefficient as aggregate pass-through. The dealer × pair × date fixed effects are defined at daily frequency and therefore do not absorb the intraday response of the spot exchange rate, which enters the executed forward rate directly.

is related to its current forward price and, more sharply, whether that position predicts its response to a common monetary-policy shock.

For each forward, we sign the dollar leg from the dealer's perspective and assign the cashflow to the contract's maturity date. We then sum the dealer's scheduled net dollar cashflows over the following seven days:

$$X_{d,t} \; = \sum_{\tau \in [t,\, t+7]} \text{Cashflow}_{d,\tau}. \tag{23}$$

Positive $X_{d,t}$ denotes scheduled net dollar inflows, which we call long-dollar exposure; negative $X_{d,t}$ denotes scheduled net dollar outflows, or short-dollar exposure. Exposure is measured in billions of dollars and is not standardized. We use its value on the previous weekday in every regression, so the exposure variable is determined before the transaction and, in the FOMC specifications, before the announcement.

Net exposure imposes a strong symmetry restriction: an additional dollar scheduled to arrive and an additional dollar scheduled to leave must affect the outcome by equal amounts with opposite signs. We therefore also decompose net exposure into

$$X^{-}_{d,t} \; = \; \max(-X_{d,t},\, 0)\,, \qquad X^{+}_{d,t} \; = \; \max(X_{d,t},\, 0)\,, \tag{24}$$

where $X^{-}_{d,t}$ measures short-dollar exposure and $X^{+}_{d,t}$ long-dollar exposure. The net, absolute, and split specifications impose successively weaker restrictions on these two components. Because $X = X^{+} - X^{-}$, the net specification constrains the coefficient on X^{+} to equal minus the coefficient on X^{-}. Because $|X| = X^{+} + X^{-}$, the absolute specification constrains the two coefficients to be equal. The split specification estimates them separately and is therefore the preferred specification for determining which side of the dealer's scheduled dollar position drives the relationship.

The level regressions show that scheduled cashflows covary with forward prices, but not symmetrically across dealer roles or exposure directions. Table 8 estimates equation (21) using exposure measured on the previous weekday. All specifications absorb execution-window $\times$ pair $\times$ maturity effects and dealer $\times$ pair effects on both sides, so the coefficients compare transactions after removing the common forward curve and persistent dealer-pair pricing differences.

On the sell-FC side, a \$1 billion increase in net scheduled dollar inflows is associated with a 0.132 basis point lower forward price. Absolute exposure also enters negatively, at -0.168. The split specification reveals that these reduced-form summaries mask an asymmetric relationship: short-dollar exposure enters at -0.042 basis points per billion and long-dollar exposure at -0.170. Both estimates are negative, rather than equal and opposite as the net specification requires. The absolute specification's equality restriction also fails to describe this pattern. The sell-FC level relationship therefore cannot be summarized as a symmetric response to net dollar exposure.

The buy-FC side is weaker and different. Neither net exposure, at 0.065, nor absolute exposure, at 0.042, is distinguishable from zero. In the split specification, short-dollar exposure enters at -0.051 basis points per billion, with a standard error of 0.017, whereas the long-dollar coefficient is 0.036 with a standard error of 0.049. Thus scheduled cashflows are related to residual transaction prices, but the level estimates do not reveal a common exposure-price relationship across the two dealer roles.

This ambiguity makes the FOMC design particularly useful. Cross-sectional price levels can reflect dealer conditions correlated with scheduled cashflows even after the fixed effects. Around an announcement, by contrast, equation (22) asks whether dealers with different predetermined exposures respond differently to the same monetary-policy news. We therefore place greater interpretive weight on the announcement-window evidence than on the level correlations.

The FOMC evidence provides a much cleaner directional pattern: heterogeneous pass-through is concentrated among dealers expecting dollar inflows. Table 9 estimates equation (22) using the same net, absolute, and split exposure measures. Net exposure predicts weaker pass-through on both dealer sides, with triple-interaction coefficients of -0.071 for the sell-FC dealer and -0.084 for the buy-FC dealer. Absolute exposure is negatively associated with pass-through on the buy-FC side, at -0.095, but its sell-FC estimate of -0.063 is imprecise.

The split specification shows where the net-exposure result comes from. Short-dollar exposure has little measurable relationship with differential pass-through: the coefficients are 0.046 on the sell-FC side and 0.008 on the buy-FC side, with standard errors of 0.086 and 0.061. Long-dollar exposure, by contrast, enters negatively on both sides. A \$1 billion increase in scheduled dollar inflows changes the post-announcement response to a one-unit statement surprise by -0.080 basis points for the sell-FC dealer and -0.105 basis points for the buy-FC dealer, with standard errors of 0.039 and 0.025. The evidence therefore rejects a simple symmetric scarcity story in which dollar inflows and outflows produce mirror-image responses. The measurable heterogeneity is concentrated among dealers with dollars scheduled to arrive.

Under the assumed sign convention for STMT, the economic interpretation is that dealers with more near-term dollar inflows transmit a tightening surprise less strongly into forward prices. Because a higher canonical forward rate corresponds, holding the other CIP compo-

Table 8

LAGGED DOLLAR EXPOSURE AND FORWARD PRICES

	Dependent variable: f_i, log forward rate (bps)					
	Net		Absolute		Split	
	(1) Sell-FC	(2) Buy-FC	(3) Sell-FC	(4) Buy-FC	(5) Sell-FC	(6) Buy-FC
Net exposure X_{t-1}	-0.132^{***}	0.065				
	(0.028)	(0.048)				
Absolute exposure $\|X\|_{t-1}$			-0.168^{***}	0.042		
			(0.029)	(0.051)		
Short USD X^{-}_{t-1}					-0.042^{***}	-0.051^{***}
					(0.013)	(0.017)
Long USD X^{+}_{t-1}					-0.170^{***}	0.036
					(0.029)	(0.049)
Fixed effects						
Window × pair × maturity	Yes	Yes	Yes	Yes	Yes	Yes
Sell-FC dealer × pair	Yes	Yes	Yes	Yes	Yes	Yes
Buy-FC dealer × pair	Yes	Yes	Yes	Yes	Yes	Yes
Observations	17,588,298	16,613,247	17,588,298	16,613,247	17,588,298	16,613,247

Note.– The table estimates equation (21), $f_i = \theta' Z_i + \mu_{w(i)c(i)m(i)} + \alpha_{s(i)c(i)} + \delta_{b(i)c(i)} + \varepsilon_i$. The unit of observation is an FX-forward transaction in core dollar pairs. The dependent variable f_i is the log executed forward rate in the canonical dollars-per-foreign-currency quote, in basis points. X is the dealer's net scheduled USD cashflow over the seven-day horizon in equation (23), in billions of dollars and not standardized, with positive values denoting net dollar inflows; X^- and X^+ are the nonnegative components in equation (24). All exposure variables are measured on the previous weekday. The column heading gives the dealer side to which exposure is attached. Standard errors, in parentheses, are three-way clustered by date × pair, sell-FC dealer × pair, and buy-FC dealer × pair. *** 1%, ** 5%, * 10%.

nents fixed, to a higher implied dollar rate, a dealer with an additional \$1 billion of scheduled inflows raises the forward price by roughly 0.08–0.11 basis points less in response to a positive one-unit statement surprise. Dealers with scheduled outflows do not display the corresponding increase in pass-through. This pattern is consistent with near-term dollar availability affecting how dealers absorb a common monetary-policy shock, but the reduced-form design does not establish the mechanism.

Scheduled cashflows also predict post-announcement price movements independently of the surprise. The exposure $\times$ Post coefficients are negative and significant in every column, at -0.0021 and -0.0025 per \$1 billion of net exposure on the sell-FC and buy-FC sides, and -0.0024 and -0.0031 for the long-dollar component. Dealers with more scheduled dollar inflows lower their prices after the release whatever the statement contains. These coefficients are far smaller in raw magnitude than the triple interactions, but the two are not directly comparable: the triple interaction is measured per unit of STMT, and a realized surprise is a small fraction of one unit.

The two exercises lead to the same reduced-form conclusion. Dealer-level dollar conditions matter not only for the cross section of forward prices but also for how those prices respond to common monetary-policy news. Secured dollar rates covary with transaction prices and predict differential FOMC pass-through. Scheduled dollar cashflows likewise predict pass-through, with the measurable response concentrated among dealers expecting dollar inflows rather than outflows. Measured directly, then, differential dollar conditions do move dealer prices, even though they leave no separately identifiable trace in the level once fixed effects absorb markup alongside them (Section IV, Table 2, Panel C). The two findings are consistent: the earlier null bounds what the static cross section can recover, while these estimates identify the same margin from within-dealer variation in observable dollar conditions. Together they connect the dealer-specific price distribution documented in Section IV to observable dealer balance-sheet conditions.

VI Conclusion

Does the cost of synthetic dollar funding depend on the dealer through which a customer trades? The evidence says yes: otherwise comparable forward transactions carry different prices across dealers, and those differences depend systematically on which side of the dollar market the dealer occupies. Rather than a single price of dollar funding, the market contains a distribution of dealer-specific prices.

The evidence identifies two dimensions of this distribution. First, the average sell-FC wedge—the side on which the dealer buys dollars—exceeds the buy-FC wedge by 2.5 basis

Table 9

Dollar Exposure and FOMC Pass-Through

	Dependent variable: f_i, log forward rate (bps)					
	Net		Absolute		Split	
	(1) Sell-FC	(2) Buy-FC	(3) Sell-FC	(4) Buy-FC	(5) Sell-FC	(6) Buy-FC
Triple interactions: exposure$_{k-1}$ × STMT × Post						
Net exposure X	-0.071^{***}	-0.084^{***}				
	(0.024)	(0.016)				
Absolute exposure $\lvert X \rvert$			-0.063	-0.095^{***}		
			(0.043)	(0.028)		
Short USD X^-					0.046	0.008
					(0.086)	(0.061)
Long USD X^+					-0.080^{**}	-0.105^{***}
					(0.039)	(0.025)
Lower-order interactions: exposure$_{k-1}$ × Post						
Net exposure X	-0.0021^{***}	-0.0025^{***}				
	(0.0004)	(0.0003)				
Absolute exposure $\lvert X \rvert$			-0.0022^{***}	-0.0029^{***}		
			(0.0007)	(0.0005)		
Short USD X^-					-0.0023	-0.0083^{***}
					(0.0032)	(0.0020)
Long USD X^+					-0.0024^{***}	-0.0031^{***}
					(0.0007)	(0.0004)
STMT×Post	-449.8^{***}	-567.1^{***}	-438.5^{***}	-540.8^{***}	-459.5^{***}	-583.2^{***}
	(123.2)	(88.11)	(134.2)	(91.13)	(145.6)	(94.62)
Post	9.356^{**}	10.71^{***}	9.888^{**}	12.04^{***}	10.28^{**}	13.86^{***}
	(3.801)	(2.735)	(4.035)	(2.865)	(4.447)	(3.019)
Fixed effects						
Sell-FC dealer × pair × date	Yes	Yes	Yes	Yes	Yes	Yes
Buy-FC dealer × pair × date	Yes	Yes	Yes	Yes	Yes	Yes
Observations	111,738	75,009	111,738	75,009	111,738	75,009

Note.– The table estimates equation (22). The upper panel reports the triple interactions between dealer exposure, the FOMC statement surprise, and Post; the second panel reports the corresponding exposure × Post interactions. The table also reports the common STMT×Post term and Post. The levels of exposure and of STMT, and their product, are absorbed by the fixed effects. The common STMT×Post coefficient is not an estimate of aggregate pass-through: the dealer × pair × date fixed effects are defined at daily frequency and do not absorb the intraday spot-exchange-rate response, which enters the executed forward rate directly. The sample is FX-forward transactions in core dollar pairs with time to maturity of thirty days or less, executed between 11:00 a.m. and 5:00 p.m. New York time on FOMC statement dates. The dependent variable f_i is the log executed forward rate in the canonical dollars-per-foreign-currency quote, in basis points. X is the dealer's net scheduled USD cashflow over the seven-day horizon in equation (23), in billions of dollars and measured on the weekday preceding the announcement; X^- (short USD, net dollar outflows) and X^+ (long USD, net dollar inflows) are the nonnegative components in equation (24). The column heading gives the dealer side to which exposure is attached. STMT is the high-frequency FOMC statement surprise; Post indicates trades executed after the 2:00 p.m. release. Standard errors, in parentheses, are three-way clustered by date × pair, sell-FC dealer × pair × date, and buy-FC dealer × pair × date. *** 1%, ** 5%, * 10%.

points per year. This gap remains between 2.2 and 2.7 basis points across specifications that absorb client-dealer relationships and observable transaction characteristics. Second, the cross section is substantially wider than this average direction gap: the 10th-to-90th percentile distance across dealer-months is roughly 7 basis points per year, and economically significant dispersion remains after removing both common month effects and persistent dealer effects.

The level of the two directional wedges provides a different result. The sell-FC plus buy-FC wedge is 1.78 basis points per year in the baseline specification, but falls to 0.20–0.95 basis points and becomes statistically indistinguishable from zero once we absorb client relationships and observable transaction characteristics associated with currency forward demand. The distinction matters because the sum isolates relative dollar conditions across dealers, whereas the difference captures borrowing-versus-lending spreads within currencies. Our estimates therefore show substantial dealer-level price dispersion and a robust directional difference, but little evidence that differences in dollar funding conditions across major dealers explain either pattern.

The additional evidence links this price distribution to dealer balance-sheet conditions. Dealers' secured dollar rates are related to their forward prices, and around FOMC announcements the same policy surprise produces different price responses across dealers. Dealers entering an announcement with higher own secured dollar rates pass through the surprise less strongly. Dealers with larger scheduled dollar inflows over the following week likewise raise forward prices by less after a tightening surprise, while dealers with scheduled dollar outflows show no comparable differential response. These estimates show that dealer balance sheets shape both the level of synthetic dollar prices and their response to common monetary shocks.

Aggregate CIP basis is better understood as an average across heterogeneous intermediaries than as a uniform shadow price of dollar funding. The common component of the basis captures an aggregate state of dollar markets, but the price actually paid by a customer also depends on the dealer that intermediates the transaction, the side of the dealer's dollar book, and the dealer's near-term funding position. Synthetic dollar funding therefore has many prices: the transaction data reveal the dealer-level market structure beneath it.

References

Acosta, M., C. M. Brennan, and M. M. Jacobson (2024). Constructing high-frequency monetary policy surprises from sofr futures. *Economics Letters 242*, 111873.

Adrian, T., E. Etula, and T. Muir (2014). Financial intermediaries and the cross-section of asset returns. *The Journal of Finance 69*(6), 2557–2596.

Avdjiev, S., W. Du, C. M. Koch, and H. S. Shin (2019). The dollar, bank leverage, and deviations from covered interest parity. *American Economic Review: Insights 1*(2), 193–208.

Baba, N., F. Packer, and T. Nagano (2008). The spillover of money market turbulence to fx swap and cross-currency swap markets. *BIS Quarterly Review None*(None), None.

Bank for International Settlements (2025). OTC foreign exchange turnover in april 2025. Triennial central bank survey, Bank for International Settlements, Basel.

Borio, C., R. N. McCauley, and P. McGuire (2022). Dollar debt in fx swaps and forwards: huge, missing and growing. *BIS Quarterly Review None*(None), None.

Borio, C., R. N. McCauley, P. McGuire, and V. Sushko (2016). Covered interest parity lost: understanding the cross-currency basis. *BIS Quarterly Review*.

Brunnermeier, M. K. and L. H. Pedersen (2009). Market liquidity and funding liquidity. *Review of Financial Studies 22*(6), 2201–2238.

Callier, P. (1981). One way arbitrage, foreign exchange and securities markets: A note. *The Journal of Finance 36*(5), 1177–1186.

Cenedese, G., P. Della Corte, and T. Wang (2021). Currency mispricing and dealer balance sheets. *The Journal of Finance 76*(6), 2763–2803.

Cesa-Bianchi, A., R. Czech, and F. Eguren-Martin (2021). Dash for dollars. Bank of England Staff Working Paper series 932, Bank of England.

Clinton, K. (1988). Transactions costs and covered interest arbitrage: Theory and evidence. *Journal of Political Economy 96*(2), 358–370.

Colliard, J.-E., T. Foucault, and P. Hoffmann (2021). Inventory management, dealers' connections, and prices in over-the-counter markets. *The Journal of Finance 76*(5), 2199–2247.

Deardorff, A. V. (1979). One-way arbitrage and its implications for the foreign exchange markets. *Journal of Political Economy 87*(2), 351–364.

Di Maggio, M., A. Kermani, and Z. Song (2017). The value of trading relations in turbulent times. *Journal of Financial Economics 124*(2), 266–284.

Du, W., B. Hébert, and W. Li (2023). Intermediary balance sheets and the treasury yield curve. *Journal of Financial Economics 150*(3), 103722.

Du, W., R. Keerati, and J. Schreger (2026). Decoupling dollar and treasury privilege. Working Paper 35000, National Bureau of Economic Research.

Du, W., A. Tepper, and A. Verdelhan (2018). Deviations from covered interest rate parity. *The Journal of Finance 73*(3), 915–957.

Duffie, D., N. Gârleanu, and L. H. Pedersen (2005). Over-the-counter markets. *Econometrica 73*(6), 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.

Ferrara, G. and H. Hall (2026). Trading relationships in fx derivatives: Lessons from credit suisse's collapse. Staff Working Paper 1154, Bank of England. Updated version of the working paper originally published in 2025.

Frenkel, J. A. and R. M. Levich (1975). Covered interest arbitrage: Unexploited profits? *Journal of Political Economy 83*(2), 325–338.

Frenkel, J. A. and R. M. Levich (1977). Transaction costs and interest arbitrage: Tranquil versus turbulent periods. *Journal of Political Economy 85*(6), 1209–1226.

Gârleanu, N. and L. H. Pedersen (2011). Margin-based asset pricing and deviations from the law of one price. *Review of Financial Studies 24*(6), 1980–2022.

Hacioğlu-Hoke, S., D. Ostry, H. Rey, A. Rousset Planat, V. Stavrakeva, and J. Tang (2024). Topography of the fx derivatives market: A view from london. Staff Working Paper 1103, Bank of England.

Hau, H., P. Hoffmann, S. Langfield, and Y. Timmer (2021). Discriminatory pricing of over-the-counter derivatives. *Management Science 67*(11), 6660–6677.

He, Z., B. Kelly, and A. Manela (2017). Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics 126*(1), 1–35.

He, Z. and A. Krishnamurthy (2013). Intermediary asset pricing. *American Economic Review 103*(2), 732–770.

Hendershott, T., D. Li, D. Livdan, and N. Schürhoff (2020). Relationship trading in over-the-counter markets. *The Journal of Finance 75*(2), 683–734.

Ivashina, V., D. S. Scharfstein, and J. C. Stein (2015). Dollar funding and the lending behavior of global banks. *The Quarterly Journal of Economics 130*(3), 1241–1281.

Jiang, Z., A. Krishnamurthy, and H. Lustig (2024). Dollar safety and the global financial cycle. *The Review of Economic Studies 91*(5), 2878–2915.

Kloks, P., E. Mattille, and A. Ranaldo (2026). Segmented dollar funding. *Journal of Financial Economics 184*, 104348.

Koijen, R. S. J. and M. Yogo (2019). A demand system approach to asset pricing. *Journal of Political Economy 127*(4), 1475–1515.

Li, D. and N. Schürhoff (2019). Dealer networks. *The Journal of Finance 74*(1), 91–144.

Liao, G. Y. (2020). Credit migration and covered interest rate parity. *Journal of Financial Economics 138*(2), 504–525.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2016). Information flows in foreign exchange markets: Dissecting customer currency trades. *The Journal of Finance 71*(2), 601–634.

Ranaldo, A. and F. Somogyi (2021). Asymmetric information risk in FX markets. *Journal of Financial Economics 140*(2), 391–411.

Rime, D., A. Schrimpf, and O. Syrstad (2022). Covered interest parity arbitrage. *The Review of Financial Studies 35*(11), 5185–5227.

Shleifer, A. and R. W. Vishny (1997). The limits of arbitrage. *The Journal of Finance 52*(1), 35–55.

Syrstad, O. and G. Viswanath-Natraj (2022). Price-setting in the foreign exchange swap market: Evidence from order flow. *Journal of Financial Economics 146*(1), 119–142.

Wallen, J. (2022). Markups to financial intermediation in foreign exchange markets. Working paper.

Appendix

A Forward Sides, Funding Rates, and Dealer Wedges

This appendix derives the one-way bounds that link FX forward sides to dealer funding rates. The key point is that the spot and money-market hedge is self-financing: it costs zero at date t. As it is standard in the presence of transaction costs, the asset prices are only defined up to some inequalities (Deardorff, 1979; Callier, 1981). We would need a model of the demand side and customers' bargaining power in order to derive a solution for the forward rates.

We consider two possible strategies. A one-way arbitrage following Deardorff (1979), that is, a market participant with a desired terminal currency transaction compares the direct forward transaction with the synthetic transaction generated by spot and money-market trades, and chooses the cheaper way to obtain the same terminal currency position. A hedge strategy, that is, a dealer accepts a side only if the zero-initial-cost forward-plus-hedge position generates a weakly non-negative terminal surplus.

All exchange rates are quoted in the canonical convention: dollars per unit of foreign currency. Let S_t denote the spot exchange rate and let $F_t(T)$ denote the forward rate for delivery at date T. Thus one unit of foreign currency converts into S_t dollars, and one dollar converts into $1/S_t$ units of foreign currency. A higher forward rate means that foreign currency is more expensive in dollars. Write

$$\tau = T - t$$

for time to maturity. With continuously compounded benchmark rates, define the benchmark log forward rate

$$f_t(T) = \log S_t + \tau(r_{\$,t} - r_{f,t}),$$

where $r_{\$,t}$ and $r_{f,t}$ are the dollar and foreign-currency benchmark rates over the life of the contract.

There are two sides of the forward. If dealer d buys foreign currency forward, it receives one unit of foreign currency and delivers dollars at maturity. This is the buy-FC side, or the dealer sell-dollar side. If dealer d sells foreign currency forward, it delivers one unit of foreign currency and receives dollars at maturity. This is the sell-FC side, or the dealer buy-dollar side.

Let $R^B_{d,\mathrm{USD},t}$ and $R^L_{d,\mathrm{USD},t}$ denote the dealer's gross dollar borrowing and lending rates over $[t, T]$. Let $R^B_{d,\mathrm{FC},t}$ and $R^L_{d,\mathrm{FC},t}$ denote the corresponding foreign-currency borrowing and lending rates. Write

$$R^B_{d,\text{USD},t} = \exp\{(r_{\$,t} + s^B_{d,\text{USD},t})\tau\}, \qquad R^L_{d,\text{USD},t} = \exp\{(r_{\$,t} + s^L_{d,\text{USD},t})\tau\},$$

and

$$R^B_{d,\text{FC},t} = \exp\{(r_{f,t} + s^B_{d,\text{FC},t})\tau\}, \qquad R^L_{d,\text{FC},t} = \exp\{(r_{f,t} + s^L_{d,\text{FC},t})\tau\}.$$

The superscript B denotes the dealer's borrowing rate. The superscript L denotes the dealer's lending or investment rate. We will assume that

$$R^B_{d,k,t} \geq R^L_{d,k,t}, \qquad k \in \{\text{USD}, \text{FC}\}.$$

This assumption is needed to rule out the trivial same-currency arbitrage of borrowing and lending in the same currency at a positive spread. It also describes the money markets in which dealers actually fund: central bank standing facilities institutionalize the wedge in both currency blocs, since the ECB's marginal lending facility rate lies above its deposit facility rate and the Federal Reserve's standing repo facility rate lies above the rate on its overnight reverse repurchase agreements, so a bank that raises cash pays above the policy benchmark while a bank that places cash earns below it.

A terminal payoff vector is written as

$$(x, y),$$

where x is dollars and y is units of foreign currency at maturity. Under this convention, the sell-FC forward payoff to the dealer is

$$(F^{\text{sell}}_{dt}, -1),$$

and the buy-FC forward payoff to the dealer is

$$(-F^{\text{buy}}_{dt}, +1).$$

The forward has zero upfront payment.

Synthetic terminal transactions There are two zero-initial-cost synthetic terminal transactions. They are the money-market alternatives to direct forward trading.

First, consider the synthetic purchase of one unit of foreign currency for delivery at T.

To receive one unit of foreign currency at maturity, the dealer lends

$$\frac{1}{R^L_{d,\mathrm{FC},t}}$$

units of foreign currency today. Since the spot quote is dollars per unit of foreign currency, buying this amount of foreign currency costs

$$\frac{S_t}{R^L_{d,\mathrm{FC},t}}$$

dollars today. If the dealer borrows this dollar amount, the dollar repayment at maturity is

$$\frac{S_t R^B_{d,\mathrm{USD},t}}{R^L_{d,\mathrm{FC},t}}.$$

Thus the zero-cost synthetic transaction that delivers +1 unit of foreign currency at maturity has terminal payoff

$$\left(-\frac{S_t R^B_{d,\mathrm{USD},t}}{R^L_{d,\mathrm{FC},t}}, +1\right). \tag{25}$$

Define

$$F^+_{dt} \equiv \frac{S_t R^B_{d,\mathrm{USD},t}}{R^L_{d,\mathrm{FC},t}}. \tag{26}$$

Then (25) can be written as $(-F^+_{dt}, +1)$. This is the synthetic future purchase price of foreign currency.

Second, consider the synthetic sale of one unit of foreign currency for delivery at T. To owe exactly one unit of foreign currency at maturity, the dealer borrows

$$\frac{1}{R^B_{d,\mathrm{FC},t}}$$

units of foreign currency today. Converting this amount into dollars gives

$$\frac{S_t}{R^B_{d,\mathrm{FC},t}}$$

dollars today. Lending these dollars until maturity gives

$$\frac{S_t R^L_{d,\mathrm{USD},t}}{R^B_{d,\mathrm{FC},t}}$$

dollars at maturity. The foreign-currency borrowing requires repayment of one unit of foreign currency. Thus the zero-cost synthetic transaction that delivers -1 unit of foreign currency at maturity has terminal payoff

$$\left(\frac{S_t R^L_{d,\text{USD},t}}{R^B_{d,\text{FC},t}}, -1 \right). \tag{27}$$

Define

$$F^-_{dt} \equiv \frac{S_t R^L_{d,\text{USD},t}}{R^B_{d,\text{FC},t}}. \tag{28}$$

Then (27) can be written as $(F^-_{dt}, -1)$. This is the synthetic future sale price of foreign currency.

Replication bounds and hedge bounds This subsection separates two distinct one-way comparisons.

Define

$$F^-_{dt} \equiv \frac{S_t R^L_{d,\text{USD},t}}{R^B_{d,\text{FC},t}}, \qquad F^+_{dt} \equiv \frac{S_t R^B_{d,\text{USD},t}}{R^L_{d,\text{FC},t}}. \tag{29}$$

The first number, F^-_{dt}, is the dollar receipt generated by the zero-cost synthetic sale of one unit of foreign currency. The second number, F^+_{dt}, is the dollar cost generated by the zero-cost synthetic purchase of one unit of foreign currency. If dealer borrowing rates weakly exceed dealer lending rates in both currencies, then

$$F^-_{dt} \leq F^+_{dt}. \tag{30}$$

Indeed,

$$\frac{F^+_{dt}}{F^-_{dt}} = \frac{R^B_{d,\text{USD},t}}{R^L_{d,\text{USD},t}} \cdot \frac{R^B_{d,\text{FC},t}}{R^L_{d,\text{FC},t}} \geq 1.$$

The inequality is strict if at least one borrowing-lending spread is strict.

Consider a dealer that sells one unit of foreign currency forward. The dealer's forward payoff is

$$(F^{\text{sell}}_{dt}, -1).$$

There are two relevant one-way comparisons.

First, the dealer can compare the forward sale with the same-side synthetic sale $X^-_{dt} = (F^-_{dt}, -1)$. Both transactions deliver one unit of foreign currency at maturity. The forward sale is weakly better than the synthetic sale if and only if it gives at least as many dollars:

$$F^{\text{sell}}_{dt} \geq F^-_{dt}. \tag{31}$$

This is the same-side replication bound. It is a one-way dominance condition: if $F^{\text{sell}}_{dt} < F^{-}_{dt}$, then selling foreign currency forward is dominated by synthetically selling foreign currency through spot and money markets.

Second, the dealer can hedge the sell-FC forward position. To hedge the obligation to deliver one unit of foreign currency, the dealer uses the synthetic purchase $X^{+}_{dt} = (-F^{+}_{dt}, +1)$. Combining the forward and the hedge gives

$$(F^{\text{sell}}_{dt}, -1) + (-F^{+}_{dt}, +1) = (F^{\text{sell}}_{dt} - F^{+}_{dt}, 0).$$

This combined position costs zero at date t (both the forward and the hedge are zero-cost). Hence the dealer's zero-initial-cost forward-plus-hedge position generates a weakly non-negative terminal surplus if and only if

$$F^{\text{sell}}_{dt} \geq F^{+}_{dt}. \tag{32}$$

This is the hedge bound.

Because $F^{-}_{dt} \leq F^{+}_{dt}$, the hedge bound implies the replication bound:

$$F^{\text{sell}}_{dt} \geq F^{+}_{dt} \quad \Longrightarrow \quad F^{\text{sell}}_{dt} \geq F^{-}_{dt}.$$

Thus, on the sell-FC side, the hedge bound is the tighter bound. The zero-surplus boundary is

$$\bar{F}^{\text{sell}}_{dt} = F^{+}_{dt} = \frac{S_t R^{B}_{d,\text{USD},t}}{R^{L}_{d,\text{FC},t}}. \tag{33}$$

The hedge strategy only implies $F^{\text{sell}}_{dt} \geq \bar{F}^{\text{sell}}_{dt}$. Equality requires an additional zero-surplus, competitive-pricing, or customer-outside-option assumption.

Now consider a dealer that buys one unit of foreign currency forward. The dealer's forward payoff is

$$(-F^{\text{buy}}_{dt}, +1).$$

Again there are two relevant one-way comparisons.

First, the dealer can compare the forward purchase with the same-side synthetic purchase $X^{+}_{dt} = (-F^{+}_{dt}, +1)$. Both transactions deliver one unit of foreign currency at maturity. The forward purchase is weakly better than the synthetic purchase if and only if it requires no larger dollar payment:

$$F^{\text{buy}}_{dt} \leq F^{+}_{dt}. \tag{34}$$

This is the same-side replication bound. If $F^{\text{buy}}_{dt} > F^{+}_{dt}$, then buying foreign currency forward is dominated by synthetically buying foreign currency through spot and money markets.

Second, the dealer can hedge the buy-FC forward position. The dealer will receive one unit of foreign currency from the forward, so the hedge must create a terminal obligation of one unit of foreign currency. The hedge is the synthetic sale $X_{dt}^{-} = (F_{dt}^{-}, -1)$. Combining the forward and the hedge gives

$$(-F_{dt}^{\text{buy}}, +1) + (F_{dt}^{-}, -1) = (F_{dt}^{-} - F_{dt}^{\text{buy}}, 0).$$

This combined position costs zero at date t (both the forward and the hedge are zero-cost). Hence the dealer's zero-initial-cost forward-plus-hedge position generates a weakly non-negative terminal surplus if and only if

$$F_{dt}^{\text{buy}} \leq F_{dt}^{-}. \tag{35}$$

This is the hedge bound.

Because $F_{dt}^{-} \leq F_{dt}^{+}$, the hedge bound implies the replication bound:

$$F_{dt}^{\text{buy}} \leq F_{dt}^{-} \quad \Longrightarrow \quad F_{dt}^{\text{buy}} \leq F_{dt}^{+}.$$

Thus, on the buy-FC side, the hedge bound is the tighter bound. The zero-surplus boundary is

$$\bar{F}_{dt}^{\text{buy}} = F_{dt}^{-} = \frac{S_t R_{d,\text{USD},t}^{L}}{R_{d,\text{FC},t}^{B}}. \tag{36}$$

Again, the hedge only implies $F_{dt}^{\text{buy}} \leq \bar{F}_{dt}^{\text{buy}}$. Equality requires an additional zero-surplus, competitive-pricing, or customer-outside-option assumption.

For a dealer buy and sell quotes, the hedge participation constraints are

$$F_{dt}^{\text{buy}} \leq F_{dt}^{-}, \qquad F_{dt}^{\text{sell}} \geq F_{dt}^{+}.$$

Thus the dealer's maximum zero-surplus F^{buy} is F_{dt}^{-} and the dealer's minimum zero-surplus F^{sell} is F_{dt}^{+}. With no additional markup, the dealer sell minus buy spread is therefore

$$F_{dt}^{-} - F_{dt}^{+},$$

, which is non-negative by (30). With market power, inventory effects, balance-sheet costs, relationship pricing, or customer inelasticity, the actual F^{buy} and F^{sell} can be different from F_{dt}^{-} and F_{dt}^{+}, respectively.

If the customer has the same synthetic outside options and competition eliminates surplus, then customer participation gives the opposite inequalities. A customer selling foreign currency will not accept $F_{dt}^{\text{buy}} < F_{dt}^{-}$, and a customer buying foreign currency will not accept $F_{dt}^{\text{sell}} > F_{dt}^{+}$. Combining these customer constraints with the dealer's no-loss constraints selects the boundary quotes:

$$F_{dt}^{\text{buy}} = F_{dt}^{-}, \qquad F_{dt}^{\text{sell}} = F_{dt}^{+}.$$

Without those additional customer-access and zero-surplus assumptions, only the hedge inequalities are implied.

Dealer wedges The dealer wedges defined in the empirical sections are broader concepts than the boundary funding component. In fact, as mentioned, actual transaction prices can differ from the boundary quote because a forward position can use balance-sheet capacity, affect inventory risk, relax an existing exposure, or reflect residual demand from customers. For convenience,or $j \in \{\text{buy}, \text{sell}\}$ let c_{dt}^{j} capture the intermediation costs, which include the funding component and the shadow value of banace-sheet capacity usage, and let μ_{dt}^{j}, collect the nonfunding components and any markup or concession in annualized log points. Define dealer wedges by

$$\log F_{dt}^{j} = f_t(T) + \tau\omega_{dt}^{j}.$$

The decomposition is

$$\omega_{dt}^{\text{sell}} = c_{dt}^{\text{sell}}, + \mu_{dt}^{\text{sell}}, \tag{37}$$

$$\omega_{dt}^{\text{buy}} = c_{dt}^{\text{buy}}. + \mu_{dt}^{\text{buy}}. \tag{38}$$

These equalities define the residual terms m_{dt}^{j} without imposing zero surplus.

Dollar scarcity maps naturally into the sell-FC, dealer buy-dollar side because hedging that side requires borrowing dollars today, namely c_{dt}^{sell} reflects $s_{d,\text{USD},t}^{B} - s_{d,\text{FC},t}^{L}$. Conversely, a buy-FC dealer delivers dollars at maturity, but its hedge lends dollars today and borrows foreign currency, namely c_{dt}^{buy} reflects $s_{d,\text{USD},t}^{L} - s_{d,\text{FC},t}^{B}$. Dollar borrowing pressure therefore does not enter the buy-FC and sell-FC wedges in the same way.

The model does not imply that the estimated buy-FC wedge must be non-positive or that the estimated sell-FC wedge must be positive for every dealer. Such a sign result would require additional assumptions about competition, customer outside options, and the interpretation of the residual term. The empirical wedges in this paper are log-price differences relative to an omitted category after fixed effects absorb common forward-curve components. They can therefore be positive or negative on either side. The prediction is about relative pricing and comparative statics: dollar borrowing pressure should raise the sell-FC wedge relative to the buy-FC wedge, all else equal. If the nonfunding term on the buy-FC side is sufficiently larger than the nonfunding term on the sell-FC side, the sell-minus-buy wedge can be negative even when borrowing spreads exceed lending spreads in both currencies.

Two combinations of dealer wedges isolate different funding terms. Their *difference* is

$$\omega_{dt}^{\text{sell}} - \omega_{dt}^{\text{buy}} = \left(s_{d,\text{USD},t}^{B} - s_{d,\text{USD},t}^{L}\right) + \left(s_{d,\text{FC},t}^{B} - s_{d,\text{FC},t}^{L}\right) + \left(\mu_{dt}^{\text{sell}} - \mu_{dt}^{\text{buy}}\right). \tag{39}$$

The difference therefore loads on the dealer's borrowing-minus-lending spread in both dollars and foreign currency, together with any difference in nonfunding components across the two sides. It is a round-trip pricing object rather than a dollar-specific measure.

The *level*, defined as the sum of the two wedges, is

$$\omega_{dt}^{\text{sell}} + \omega_{dt}^{\text{buy}} = \left(s_{d,\text{USD},t}^{B} + s_{d,\text{USD},t}^{L}\right) - \left(s_{d,\text{FC},t}^{B} + s_{d,\text{FC},t}^{L}\right) + \left(\mu_{dt}^{\text{sell}} + \mu_{dt}^{\text{buy}}\right). \tag{40}$$

Its funding component instead nets the dealer's dollar borrowing and lending spreads against its foreign-currency borrowing and lending spreads. The level therefore captures relative dollar funding conditions, together with the components of dealer pricing that enter both sides of the book. We report both objects in the main text because the difference measures the directional asymmetry in dealer pricing, whereas the level speaks to differential dollar conditions across dealers. To reduce the contribution of nonfunding components, we re-estimate dealer wedges after absorbing client-dealer relationships and controlling for observable transaction characteristics associated with pricing. These controls reduce the scope for the level and difference to reflect observable differences in residual demand.

B Additional Results

Table B.1
Sample construction

	Transactions	% of previous	% of raw
Raw UK EMIR FX forward reports	178,473,017	—	100.0
Six core dollar pairs	46,670,265	26.1	26.1
Orientation and price filters	45,671,536	97.9	25.6
At least one G15 parent counterparty	38,441,722	84.2	21.5
Excluding intra-group transactions	27,291,457	71.0	15.3

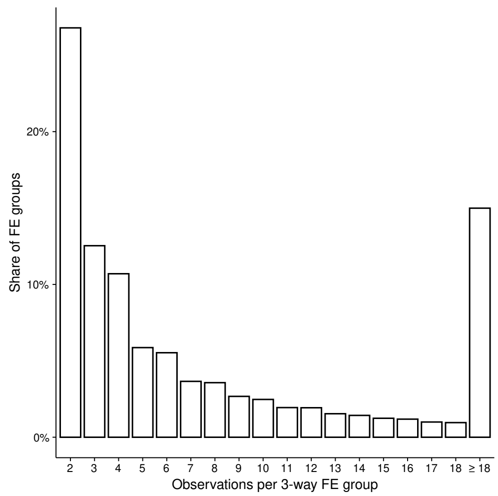

(A) Observations per group

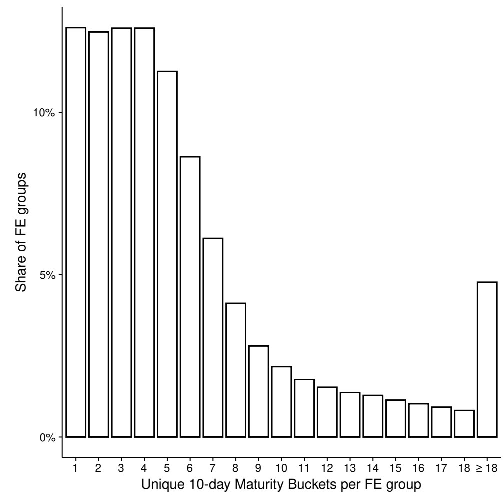

(B) Maturity buckets per window-pair group

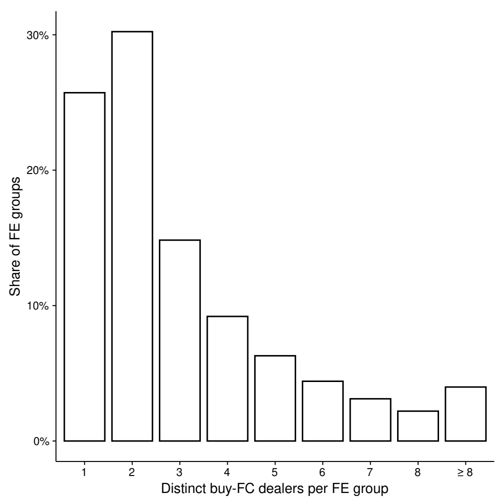

(C) Buy-FC dealers per group

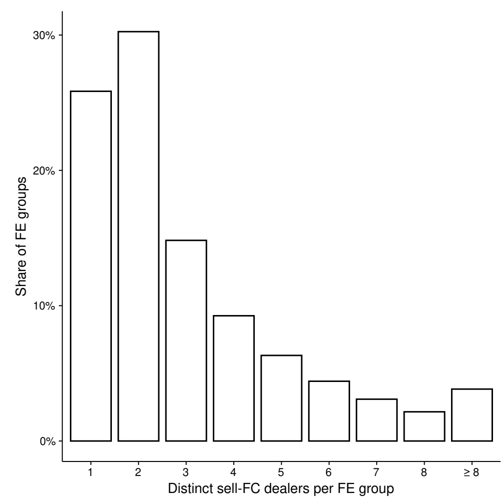

(D) Sell-FC dealers per group

Figure B.1. Fixed-Effect Group Composition. *Notes.* The figure summarizes the cells used to estimate equation (9). Panels A, C, and D use execution-window ×currency-pair ×ten-day maturity-bucket cells. Panel A plots the distribution of the number of transactions per cell. Panel B plots the number of ten-day maturity buckets observed within each execution-window ×currency-pair group. Panels C and D plot the number of distinct observed G15 dealer parents on the buy-FC and sell-FC sides of each pricing cell. Bar heights are shares of cells.

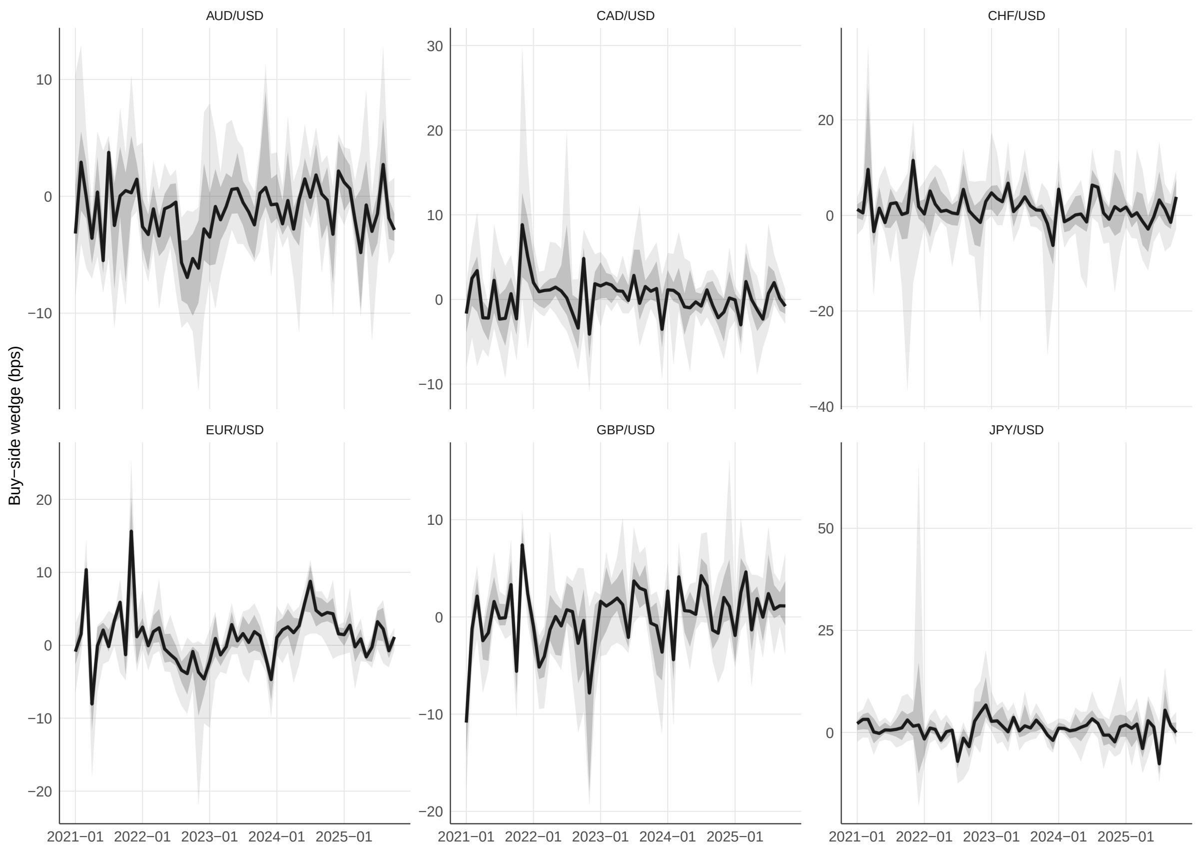

Figure B.2. Buy-foreign-currency Dealer Wedges by Currency Pair, Core Dollar Pairs. *Notes.* The figure plots the monthly cross-sectional distribution of buy-foreign-currency dealer wedges separately by core dollar currency pair. The solid line is the median; dark shading is the interquartile range; light shading is the 10th–90th percentile range. Units: annualized basis points.

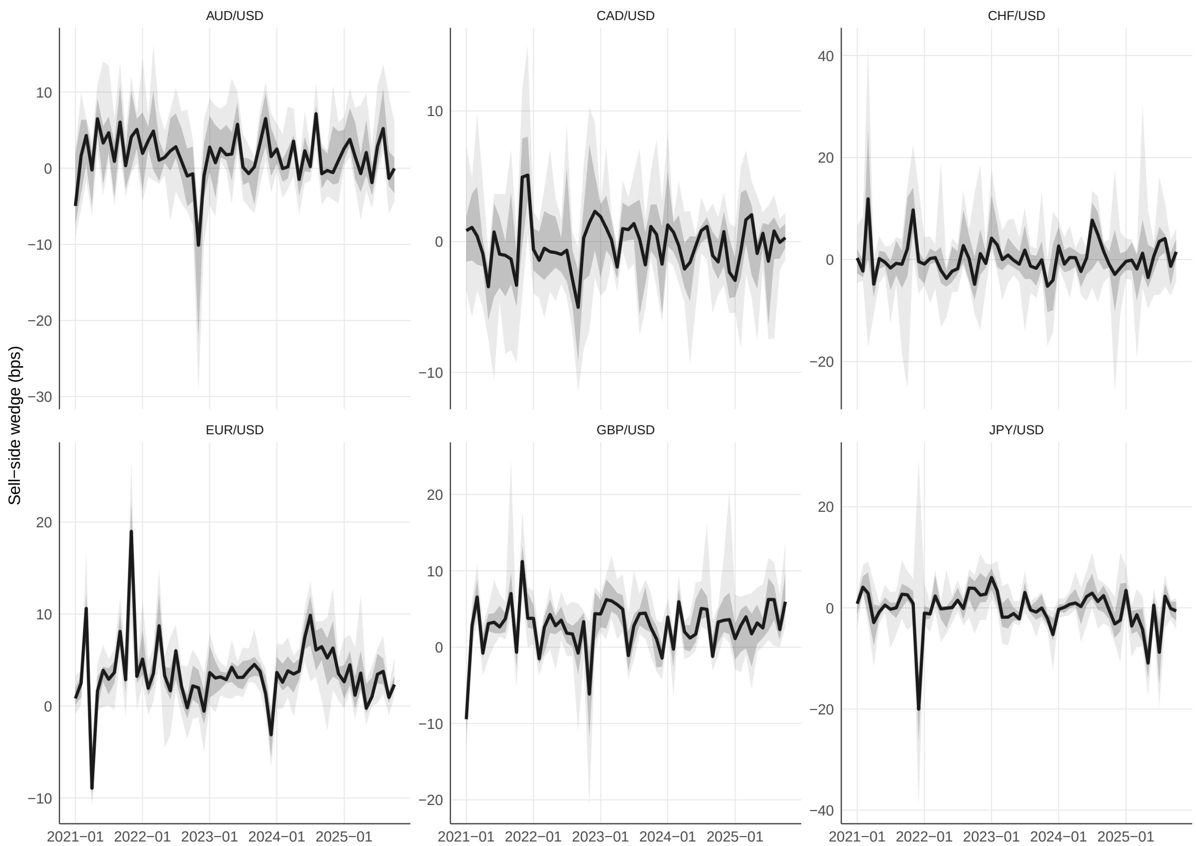

Figure B.3. Sell-foreign-currency Dealer Wedges by Currency Pair, Core Dollar Pairs. *Notes.* The figure plots the monthly cross-sectional distribution of sell-foreign-currency dealer wedges separately by core dollar currency pair. The solid line is the median; dark shading is the interquartile range; light shading is the 10th–90th percentile range. Units: annualized basis points.

C Trade-Characteristic Controls

This appendix describes the transaction-cost controls used in Table 2 and how they enter the estimation of the dealer wedges.

Controls. We construct three trade-level indicators from the baseline estimation sample. Each marks a condition under which a negotiated round-trip spread is narrower, because the customer has a better outside option or the trade is cheaper to intermediate.

Large trade equals one if the trade's dollar notional exceeds the median dollar notional within currency pair and month. Let N_i^{FC} and N_i^{USD} denote the notional amounts of the foreign-currency and dollar legs of trade i, taken after the orientation adjustment of Section III, so that N_i^{FC} is always the leg denominated in the non-dollar currency of the pair. The dollar size of the trade is

$$\text{USD Notional}_i = \begin{cases} N_i^{\mathrm{USD}}, & \text{if the dollar leg is reported,} \\ N_i^{\mathrm{FC}}\, F_i, & \text{otherwise,} \end{cases} \tag{41}$$

where F_i is the executed forward rate in dollars per unit of foreign currency. The second line reconstructs the dollar leg from the foreign-currency leg and the contract's own price, which is exact for a deliverable forward and retains observations that reliance on the reported dollar leg alone would discard. Because the quote convention is dollars per unit of foreign currency, the conversion multiplies by F_i rather than dividing by it; the orientation step of Section III ensures that convention holds before the calculation is made.

Electronic execution equals one for trades executed on venue and zero for bilateral transactions, constructed as the complement of the reported OTC execution flag.

Liquid currency pair equals one for EUR/USD, GBP/USD and JPY/USD.

All three are populated for the full baseline sample, so adding them leaves the estimation sample unchanged and the rows of Table 2 remain comparable.

Dealer presence and dealer direction. The controls are allowed to move the level of dealer pricing and the sell-minus-buy gap separately. For each trade i, let S_i and B_i equal one when a G15 dealer parent is on the sell-FC and buy-FC side of the trade respectively, and define

$$P_i = S_i + B_i, \qquad D_i = S_i - B_i. \tag{42}$$

A trade with an observed dealer on the sell side alone has $P_i = 1$ and $D_i = 1$; one with an observed dealer on the buy side alone has $P_i = 1$ and $D_i = -1$; one with observed dealers on both sides has $P_i = 2$ and $D_i = 0$. A term entering the price with loading $\lambda P_i + \delta D_i$ therefore shifts the sell-side price by $\lambda + \delta$, the buy-side price by $\lambda - \delta$, their average by λ, and their difference by 2δ. Interacting a control with P_i and with D_i thus lets it shift the level of pricing and the direction gap independently, rather than imposing that a characteristic which raises one also raises the other.

Specification. Let $\widetilde{X}_{ik} = X_{ik} - \bar{X}_k$ denote control k in deviation from its sample mean. We re-estimate equation (9) with six added interactions:

$$\begin{aligned} p_i = \Gamma_{w,c,\ell} + \alpha_c \tau_i + \sum_d \omega_d^{\text{sell}} \, \tau_i \, 1\{d = \text{seller}_i\} + \sum_d \omega_d^{\text{buy}} \, \tau_i \, 1\{d = \text{buyer}_i\} \\ + \sum_k \lambda_k \, P_i \tau_i \widetilde{X}_{ik} + \sum_k \delta_k \, D_i \tau_i \widetilde{X}_{ik} + \varepsilon_i, \end{aligned} \tag{43}$$

where p_i is the log executed forward rate, τ_i is time to maturity in years, and $\Gamma_{w,c,\ell}$ denotes the baseline execution-window $\times$ currency-pair $\times$ maturity-bucket fixed effects. Each interaction is multiplied by τ_i, as the dealer indicators are, so all coefficients are in annualized log points.

Demeaning the controls fixes the point at which the dealer coefficients are evaluated. With $\widetilde{X}_{ik}$ in place of X_{ik}, ω_d^{sell} and ω_d^{buy} are the wedges of a dealer transacting a trade with the sample mean of every control, and are therefore on the same scale as the baseline estimates. Without demeaning they would be the wedges at $X_{ik} = 0$, a small trade executed bilaterally in a less liquid pair, and would not be comparable across specifications. The monthly dealer-role wedges $\hat{\omega}_{d,t}^{j}$ from equation (43) replace the baseline wedges, and the resulting panel is carried through equations (18)–(13) exactly as in the baseline.