

Network Structure and Pricing in the FX Market.

January 9, 2020

Joel Hasbrouck
NYU Stern

Richard M. Levich
NYU Stern

Joel Hasbrouck, NYU Stern School of Business, 44 West 4th Street, New York, NY 10012.
E-mail: jhasbrou@stern.nyu.edu

Richard M. Levich, NYU Stern School of Business, 44 West 4th Street, New York, NY 10012. E-mail:
rlevich@stern.nyu.edu

This paper reflects the views of the authors and should not be interpreted as reflecting the views of CLS Bank International or New York University. We thank Rob Franolic, Dino Kos, and Irene Mustich for their assistance in obtaining data for this study. We are indebted to seminar audiences at the Securities and Exchange Commission, the Federal Reserve Bank of New York, the Cox School (SMU), and the Commodities Futures Trading Commission. We take responsibility for all errors.

Disclosures: One of us (Hasbrouck) teaches a course for a financial institution that engages in FX market making.

Network Structure and Pricing in the FX market

Abstract

Foreign exchange (FX) settlement data define a network, for which we may construct centrality measures and profit attributions. Our sample of settlement data from CLS Bank spans diverse currency pairs, participants and execution platforms over the Aprils of 2013 and 2016.

We assign settlement members to (five) groups ranked by unweighted degree centrality. We define an average centrality differential as the return to the more-central counterparty in the trade, and model this as a function of the two counterparties' centrality groups. Estimates of the average centrality differential are generally positive: the more-central counterparty receives a higher return. Additionally, the differential generally increases as the counterparties' centralities diverge. These two results are consistent with a pervasive centrality premium. The estimates are robust to the choice of pre- or post-settlement benchmarks, to inclusion of settlement size interactions, and to grouping on volume-weighted degree centrality. Across currency pairs the centrality profit varies considerably, and typically amounts to about one-third of bid-ask half-spread. The centrality premium is consistent with the hypothesis that central agents exercise bargaining power. We also find, however, evidence suggesting that the premium is partially offset by losses that central agents incur in supplying liquidity.

I. Introduction

An active area of research views the agents in an over-the-counter financial market and their trades as defining the nodes and connecting edges of a network. In such networks agents are typically heterogeneous in the number, proximity and strength of their connections to others. These attributes are formally summarized by an agent's centrality. A more central agent may have an advantage in intermediation, which may be passed on to trading counterparties in the form of lower markups, a centrality discount. On the other hand, centrality might be associated with bargaining power and higher markups, that is, a centrality premium. The existing evidence on centrality premia vs. discounts is mixed, but generally more supportive of centrality premia. This paper investigates centrality price effects in the global foreign exchange (FX) market, using settlement data from the CLS Bank. We find strong evidence in favor of a centrality premium, of a magnitude slightly smaller than the bid-ask spread in this market.

Centrality premia have been documented in: the US municipal bond market (Li and Schürhoff (2012, 2019a)); the US corporate bond market (Di Maggio, Kermani and Song (2017)); and European overnight bank lending (Gabrieli and Georg (2017)). Centrality discounts are found in US debt securitizations (Hollifield, Neklyudov and Spatt (2017)). Julliard, Liu, Seyedan, Todorov and Yuan (2019) find that for UK repo and reverse repo operations haircuts are lower when the bank counterparty is more central. The only FX study that has come to our attention is Hagströmer and Menkveld (2019). Although their focus is information diffusion, they examine the relation between bid-ask spreads and centrality for eight large dealers in the CHF/EUR pair on the EBS trading platform for three weeks in January 2015. (About 30.8% of the CHF deals in their sample are against the EUR.) They conclude that "empirical support for [a centrality premium] is weak, at best." This study provides more conclusive evidence.

The settlement data underlying the present study are supplied by CLS Bank. CLS is a major provider of FX settlement services, presently covering eighteen major currencies. Most settlements are for spot exchanges, and these generally correspond to spot trades. Our sample covers the Aprils of 2013, and 2016, and consists of all individual settlements in these months. Hasbrouck and Levich (2019) describe this sample in detail and analyze the implications for liquidity. Both sides of the

settlement have anonymized identification codes, which allow us to construct in the present paper the network mapping, compute centralities, and impute trading revenues.

Our data comprise an unusually broad and comprehensive sample of FX market activity, not limited to any single execution method. We cannot, however, identify the execution platform, the liquidity supplier, the liquidity demander, participant type, or indeed any characteristic of the participants aside from their settlements. The data underlying the above-cited studies often include one or more of these identifiers. Our settlement members are partitioned into five groups of increasing centrality, and most inferences are based on these groupings.

To build the network we impute a link between two nodes if they share at least one settlement. Our principal measure of node centrality is unweighted degree centrality, the number of links that originate from the node, or equivalently, the number of neighboring nodes. We also investigate volume-weighted degree centrality and find that our results are essentially unchanged. This is of interest because ranking by volume-weighted centrality is equivalent to ranking by volume share, arguably a simpler and more straightforward measure of importance.

For each settlement we impute a return to the buyer of the base currency relative to a proximate benchmark price (the midpoint of a near-contemporaneous bid and ask). We then model the dependence of this return on the centralities of the buyer and seller. We find that the profit is positive when the buyer's centrality exceeds the seller's, and negative when the seller's centrality is the larger. This is consistent with a positive return to centrality, a centrality premium. The centrality premium is pervasive and present for returns computed with pre- and post-settlement benchmarks. It is present in two widely separated sample months (the Aprils of 2013 and 2016), and in all currency pairs.

It is reasonable to suppose that dealers have high centrality and customers have low centrality. It might therefore be hypothesized that grouping on centrality is simply an alternative way to recover the customer-dealer distinction, and that our centrality premium is simply the usual dealer markup. We find, however, centrality premia in settlements between adjacent groups whose high centralities suggest that all group members are likely to be dealers, and in settlements between adjacent groups whose low centralities suggest that all members are customers.

Starting from the presumption that more central members have lower marginal costs of intermediation, a centrality premium might suggest that these agents possess a measure of relative bargaining power. We cannot rule this out. Our evidence on the centrality premium, however, is a cross-sectional finding, based on high- and low-centrality interactions relative to a point-in-time benchmark. Across time we find that when high-centrality members are buying the base currency, there is usually a contemporaneous decline in the exchange rate. This holds for most currency pairs and it is often statistically significant. It is less uniform than our cross-sectional results but it nevertheless suggests a cost that might partially account for the centrality premium.

The paper is organized as follows. The next section reviews the literature on centrality in trading networks and FX microstructure. Section III discusses the data. In Section IV we describe the properties of the settlement network and its participants. Our measures of centrality pricing and the econometric specification are developed in Section V. Section VI describes our centrality grouping procedures; estimation results are presented in Section VII. Section VIII explores time series variation in centrality-based order flows. Section IX concludes the paper.

II. Literature review

The present study draws on two strands of literature. The first involves the network perspective on over-the-counter (OTC) markets. Our remarks here aim to illuminate the connection between network centrality and transaction prices. The second line of analysis deals with FX market microstructure. Here we seek to establish the relevance of network models for our clearing network.

A. Network perspectives on OTC markets

An OTC market (like the FX market) generally lacks customer-facing centralized or consolidated trading mechanisms. With few opportunities for direct customer-to-customer exchange, most customer trades are dealer-to-customer. Dealer-to-dealer (interdealer) trade occurs via mechanisms or protocols to which customers have at best limited access. In this view all trade is bilateral. (There are, for example, no mechanisms like the opening and closing single-price double-sided auctions used in equity markets.) The economic effect of immediate interest is the

connection between pricing in these bilateral trades and the counterparty centralities within the network defined by these trades.

The theoretical motivation for the centrality-pricing connection arises most prominently in the models of OTC markets based on search and bargaining. Centrality in these models is generally a proxy variable for some agent attribute like search speed or bargaining power. The argument for these proxy relationships is strong. An agent who can search quickly will contact more potential counterparties, trade more frequently, and trade with a larger set of counterparties, increasing (at least in ex post measurement) her centrality. A larger set of potential counterparties also lowers the cost of walking away from a take-it-or-leave-it bid or ask quote, and therefore translates into higher bargaining power and better terms of trade. These remarks apply generally to any agent, customers as well as dealers. In the empirical studies of US securities trading, however, a market that exhibits a centrality premium is generally held to be one in which a more central dealer imposes higher spreads on their customers (or lower spreads, in the case of a centrality discount).

Although they do not explicitly refer to centrality, Duffie, Garleanu and Pedersen (2005) clearly model various aspects of the economic effects described above. In DGP investors (customers) and marketmakers (dealers) are distinct agent classes. Search-and-bargaining occurs in dealer-to-customer trade, and (in their baseline model) in customer-customer trade. Interdealer trade occurs without search or delay at a fixed price. Search intensity in the dealer-customer market, which is generally fixed (and equivalent to speed), is in equilibrium inversely related to the bid-ask spreads faced by customers. In a comparative-statics sense, therefore, the model suggests a centrality discount. DGP also allude to another centrality-related effect. Some variants of the model feature an investor class whose members are more sophisticated, “that is, have better access to other investors or to marketmakers who do not have total bargaining power ... [and] receive a tighter bid-ask spread,” (p. 1817). Better access is arguably equivalent to better connectivity and higher centrality. Note that in this context, the connectivity and centrality pertain to the customer. The assumption that investors executing smaller trades are unsophisticated accords with a general empirical finding that smaller trades receive worse prices. Finally, although the model differentiates between investors and marketmakers, DGP suggest a blurring of this distinction as,

“... sophisticated investors would, under certain trading conditions, profit from executing as many trades and possible, and would start acting like marketmakers,” (p. 1830).

Whereas DGP assume interdealer trade to be frictionless, Neklyudov (2013, 2019) models it as a search market, with search intensities that are randomly distributed across dealers. Maintaining the proxy relation between search intensity and centrality, the variation in search intensity implies cross-sectional variation in centrality. This is an important generalization because in many empirical studies, dealer centrality is measured with respect to interdealer trades. Neklyudov’s baseline equilibrium model generates a centrality discount (lower spreads for the customers of central dealers), but under some assumptions can lead to a premium.

Li and Schürhoff (2019b) model a segmented customer clientele wherein one class prefers intermediation speed. Central dealers are faster and can charge a premium to these fast-preference investors. This accords with the practice now common in many markets of charging extra for faster channels of information and market access. Hollifield, Neklyudov and Spatt (2017) also consider a customer clientele segmented on speed preferences. In contrast to LS, however, the fast-preference investors are also more sophisticated and have higher bargaining power. The bargaining power translates into better terms of trade against the central dealers who service them, a centrality discount.

The empirical evidence on whether and why centrality induces a discount or commands a premium is mixed. In US municipal bonds Li and Schürhoff find that dealers who are more central with respect to interdealer trading charge customers higher markups (a centrality premium). As noted above, they view this premium as compensation for a superior (faster) intermediation service. In US securitizations, Hollifield, Neklyudov and Spatt find that central dealers have lower markups. This centrality discount is explained by their sophistication. Di Maggio, Kermani and Song (2017) find a centrality premium in the US corporate bond market.

These studies generally analyze the relation between dealer-customer pricing and dealer centrality in the interdealer market. Two recent studies model dealer-customer pricing as a function of customer attributes, particularly the number of dealer counterparties. In FX forwards Hau, Hoffmann, Langfield and Timmer (2019) find that customers with few counterparties receive worse prices, but that this disadvantage is eliminated when they employ multi-dealer execution

systems. In the US corporate bond market Hendershott, Li, Livdan and Schürhoff (2019) examine customer-dealer pricing for one important class of customers, insurance companies. In their model, customers initially select the dealers with whom they will subsequently trade. This number is endogenous, and its relation to pricing is non-monotonic. Initially, as a customer begins to add dealer counterparties, she receives better terms of trade due to competition. The growth also, however, dilutes the value of the dealer-customer relationship. At some point the loss from the relationship dilution outweighs the gain from competition and the customer's terms of trade worsen. Both studies are connected to ours in that within the broader context of the overall trading networks a customer's number of dealer counterparties is her degree centrality.

Most of these studies use data that identify the broker/dealers and the customer sides of trades (although not the customer identities). The customer/dealer demarcation in these data is faithful to the distinction assumed in the economic models surveyed above. In other OTC settings, customer/dealer distinctions may be less meaningful. Julliard, Liu, Seyedan, Todorov and Yuan (2019) study the UK repo/reverse repo market. Their reporting entities are banks, but some counterparties are nonbanks. They find that for UK repo and reverse repo operations haircuts are lower when the bank counterparty is more central. For a repo this corresponds to a centrality premium; for a reverse repo, a discount. (A lower haircut favors the party that posts collateral. In a repo, this party is the bank; in a reverse repo, this party is the counterparty lender.) The European overnight bank loan network studied in Gabrieli and Georg (2017), also lacks customer/dealer identifiers and there is no institutional basis for making such a distinction. The extent to which certain banks function as intermediaries can be inferred only from their transactional patterns, and most importantly, from their centralities. In pricing, Gabrieli and George find a centrality premium. As in these two studies, our settlement data do not have customer/dealer identifiers.

This discussion has examined network centrality in the context of the relatively recent search and bargaining view of OTC markets. Over time and across markets, however, OTC trade has involved a wide range of bilateral, multilateral, and centralized mechanisms. Interdealer trade on the London Stock Exchange circa 1991 occurred via bilateral negotiation and multiple electronic limit order books (Reiss and Werner (2004); Reiss and Werner (1998)). In the US Nasdaq stock market during the same era, interdealer trade was facilitated by the SelectNet (originally bilateral)

and Instinet (centralized) systems. Interdealer platforms in the US Treasury market include voice brokers, and BrokerTec and eSpeed, both electronic limit order books (Adrian, Fleming and Vogt (2017); Fleming and Remolona (1997); Fleming, Mizrach and Nguyen (2018)). A recent SIFMA survey of the US bond market covers nineteen providers and forty-two trading protocols (SIFMA (2016)). In the FX market, bilateral trading has been facilitated by the Reuters Dealing 2000-1 system Evans and Lyons (2002).¹ Centralized systems include EBS and Reuters Matching (Chaboud, Chiquoine, Hjalmarsson and Vega (2014); Chaboud, Dao and Vega (2019); King, Osler and Rime (2012); Mancini, Ranaldo and Wrampelmeyer (2013)). Multilateral platforms include Currenex, Hotspot and FXall (Moore, Schrimpf and Sushko (2016)). According to Sinclair (2018), market participants wishing to trade FX have more than 75 different FX venues at their disposal.

Given the increasingly widespread use of trading mechanisms that differ from the search and bargaining protocols featured in many theoretical models, one might question the predictions on network centrality and pricing that arise from these models. There are several justifications. Firstly, as a practical computational matter, a trading network and its centrality metrics can be constructed from any record of bilateral trades, even if the execution platforms are centralized, broadly visible and accessible. Secondly, although the search and bargaining models presume successive random encounters with counterparties, search can still occur over trading venues, the dark pools, limit order books, and so forth that comprise modern fragmented markets. Thirdly, centrality is essentially a summary proxy, for speed, number of potential counterparties, or bargaining power, even in the search and bargaining models. These proxy relationships are plausible in other mechanisms as well.

B. Foreign Exchange Market Structure and Institutions

King, Osler and Rime (2012, 2013) and Evans and Rime (2019) broadly survey the FX market, its trading arrangements, and the effects of these arrangements on exchange rate

¹ “All trades on this system take the form of electronic bilateral conversations. A conversation is initiated when a dealer calls another dealer using the system to request a quote. ... A dealer who has been called is expected to provide a fast two-way quote with a tight spread. Quotes are take-it-or-leave-it, and if not dealt or declined quickly (i.e., within seconds), the quoting dealer retracts the quote, ending the conversation,” Evans and Lyons (2002). The sequential, bilateral request for quote (RFQ) is the protocol that most closely resembles the search and bargaining paradigm commonly used in network models.

determination. The present discussion more narrowly focuses on features important for intermediation and connecting our analysis to the papers discussed above.

Relative to the markets considered in other centrality studies, the FX market exhibits some similarities. Many of the customers (large global investors) and many of the dealers (large money center banks) are the same. There are also, though, some striking differences. With daily global turnover exceeding \$5T (US dollar equivalent) the FX market is by far the largest (Bank for International Settlements (2016)). Yet by number of “securities” traded, it is arguably the smallest. CLS Bank settles eighteen currencies, which account for most FX trading. LS’s sample of municipal bonds, by contrast, comprises over a million issues.

The small number of actively traded currencies carries implications for intermediation. Many of the studies discussed above find interdealer intermediation chains, successive transfers of positions consisting of very specific securities held in quantities that are not highly divisible. In the US municipal bond market, Li and Schürhoff find that in over twenty percent of the customer sales, the position passes through two or more dealers en route to the ultimate customer buyer. Hollifield, Neklyudov and Spatt (2017) find intermediation chains in the securitization market: of the 77,045 complete chains in their sample, 21,036 (27%) involve more than one dealer. The arrangement of these transfers is search-intensive. The FX market, in contrast, has many centralized and multilateral platforms that cover multiple participants (such as ICAP’s EBS or Reuters FXall). In these systems, a large initial quantity can be reallocated to multiple counterparties in smaller amounts, largely eliminating the role for search-based intermediation chains.²

FX market transparency is relatively low. Whereas many of the OTC markets discussed in the previous section have developed some form of trade reporting (TRACE for US Treasury and corporate bonds, EMMA for US municipal bonds, for example), this has not yet occurred in the spot FX market. Under MiFID II, FX forwards involving a European counterparty are subject to

² The portfolio shifts model of Evans and Lyons (2002) of intermediation in the FX market reflects the centralization in the interdealer market. There are three rounds of trading. In the first round, new information arrives, dealers quote prices to their customers, and customers trade. This market is segmented across dealers and since prices and trades are not reported, the information is not fully revealed. The second round of trading is interdealer, but follows the same protocol: each dealer quotes a price, and trades against the quoted prices of others. The third round of trading is dealer-to-customer, and results in further risk-sharing between the dealers and customers.

transaction reporting (Hau, Hoffmann, Langfield and Timmer (2019)), but this differs from real-time trade publication.

As a final point of distinction, most microstructure studies (both theoretical and empirical) involve securities bought and sold against a given numeraire. Currency exchanges are more symmetric. Every currency is a numeraire in its home country, but in an exchange of two currencies either can be considered the numeraire. Thus, although sophisticated in many respects, the FX market bears a formal resemblance to a barter network.⁴

III. Data

CLS Bank provided us with settlement records for the Aprils of 2013 and 2016 (corresponding to BIS triennial survey months). For each settlement, a record contains ISO codes for the two currencies exchanged and their amounts; the date and time (to the millisecond) when the first party submitted a settlement instruction; and various other fields. Participants have anonymized identifiers. These replace the Bank Interchange Codes (BICs) in the original data, but for simplicity we continue to refer to them as BICs. The BIC identifiers are unique within a given April/year, but do not persist over years.⁵

Hasbrouck and Levich (2019) compare CLS settlement volume to other measures of FX spot activity. After adjustment for reporting of prime-brokered trades, we find that CLS settlements account for about 36% of BIS turnover. This is substantially more comprehensive than the coverage of EBS (6%, in 2016) or Reuters FXall (6%), the two largest electronic execution platforms. The relative composition of CLS spot activity accords closely to BIS figures.

The most active pair is EUR/USD, but even excluding this pair, most settlements have the EUR or USD on one side. Settlement quantities are strongly clustered: about forty percent of all settlements are sized at one million units of the base currency. Settlement prices are also clustered. In 2010, settlements were essentially priced on a grid with a tick-size of 0.0001 (except for those

⁴ In a barter network, price discovery (often a focus of microstructure models) may be subject to surprising impediments. Feldman (1973) shows that without a universally held numeraire commodity (“money”), sequential bilateral barter need not converge to Pareto optimal outcomes.

⁵ CLS also provided us with April 2010 data. The counts and values in the April 2010 sample agree with other sources, but the identified member counts are low. Hasbrouck and Levich (2019) incorporate the 2010 data: the analyses in that paper do not rely on member identifiers. These are needed in the present analyses, and so we drop the 2010 data.

with the JPY as the quote currency). By 2016, this had shifted to 0.00001. (The JPY pairs were initially on a 0.01 grid, and subsequently, 0.001.)

CLS settlement data have been used in other analyses (see Hasbrouck and Levich). To our knowledge, however, only one other study has used counterparty information to construct the settlement network. For a 2011-2012 sample of daily settlement flows León-Janampa José (2017) builds the network and examines its stability and risk properties. Our sample is shorter, but it consists of individual settlement records. Our analysis is also aimed at pricing.

For market prices, we rely on the Olsen quotes for thirteen major pairs. Olsen collects streaming quotes from banks and multilateral platforms. These quotes are representative, but they are not comprehensive, nor are they visible to all market participants. Our data is a subset of the streamed quotes with a ten-second time granularity: in each ten-second window, Olsen records the first bid-ask pair submitted. This is not sufficiently precise to accurately match each settlement to a prevailing bid and ask, but it is close enough to fix a local benchmark price for allocating trading profits and losses. Like the settlement prices, the bids and asks are clustered. In 2010, most quotes use four decimal digits of precision, but by 2016, the fifth decimal place is widely used (except in the USD/MXN pair). Bid-ask spreads in all pairs narrow over time (Hasbrouck and Levich (2019)).

Most of our empirical analysis follows from the assumption that settlements correspond to trades. With respect to prices, quantities, and identities of trading parties, this correspondence is highly accurate: accepted settlement instructions result in irrevocable transfers of high value. The correspondence in time, however, is inexact. Submission of settlement instructions lags the trade confirmation by an amount that generally depends on the participant and the execution mechanism. Hasbrouck and Levich analyze the delays by matching them to recently posted quotes. They find that roughly 50% of spot settlements can be closely matched to quotes posted in the previous ten seconds, and that about 80% can be matched to quotes posted in the last minute.

IV. Settlement networks

A. Overall activity

The distribution of activity across settlement members is highly concentrated. Table 1, Panel A reports counts and distribution points. In 2016 there are 21,824,492 settlement sides

(twice the number of settlements) allocated across 16,163 active members (those who engaged in at least one settlement). To protect the anonymity of the largest members we report grouped data in the top tail of the distribution (here and throughout the paper). The top five members average 2,208,252 sides per member (accounting for about 11 million sides), but the median (50th percentile) is only 14 sides per member. Panel B reports distributions where the settlements are differentiated by instrument type. (Side and member counts reflect only the indicated instrument.) Activity is not uniform across instruments (spot settlements dominate), but activity is similarly concentrated. Panel C specializes to spot settlements by currency pair. The number of active members ranges from a low of 239 (AUD/JPY 2016) to a high of 5,699 (EUR/USD 2016).

Settlement volume is also concentrated. In Table 2 we compute for each member the value of their settlement sides (USD equivalent), rank in descending order, and report the cumulative proportion (relative to the total value of all settlement sides). In both years, the top ten members account for about half of the settlement volume. At all rankings at or below 200 members, the cumulative shares rise from 2013 to 2016, suggesting an increase in concentration between the two periods.⁶

B. Network construction

A network is defined by its nodes and edges (links). It is straightforward to associate the nodes with settlement members, which are identified within a given month by anonymized codes. Initially we impute an edge between two members if they engaged in at least one mutual settlement in any currency pair and any instrument. Given the large number of participants the usual graph visualizations are too dense to be illuminating. The number of edges and nodes can be managed by limiting the network to the most active members. The networks of the most-active members, however, uniformly resemble star networks, in which each member is connected to almost all other members.

⁶ We also examined concentration by currency. Across major currencies (such as the EUR and USD) the same members dominate the rankings. In the Scandinavian currencies (such as the Danish, Norwegian, and Swedish kroner) there is more variation: the set of top-ranked members includes members (presumably local banks) whose overall rankings are not among the top. Across instrument types the set of top ranked members is generally consistent, with the exception of option exercises: the sellers of over-the-counter FX options apparently constitute a narrower and more specialized segment of the market.

The networks are more usefully characterized in other ways. The importance of a node can be quantified by its centrality. Among alternative centrality measures, the simplest is degree centrality: the number of edges that meet at a node, or alternatively the number of a node's immediate neighbors. Our primary results are based on this measure, but we also consider an alternative, weighted degree centrality, where each edge is weighted by the volume (US dollar equivalent) that passes through the edge. The weighted degree centrality for a given node is proportional to the volume share of the node (relative to the total volume across the network).

Table 3 reports the distribution of degree centrality. In the overall statistics (Panel A) and in all the instrument type subsamples (Panel B), over ten percent of the members have only one counterparty. As a point of comparison, Hendershott, Li, Livdan and Schürhoff (2019) find that about thirty percent of their insurance companies have only one dealer counterparty. It is reasonable to conjecture that in the FX settlement institutional hierarchy, members connected to only one other (at least 10%) are customers; members connected to more than, say, 100 others are more likely dealers; in the middle are larger customers, correspondent banks, and smaller non-bank dealers.

We next consider the distribution of shortest-path length. The purpose of the settlement network is intermediation. Between two randomly selected members, the shortest path connecting them is arguably the most efficient settlement path. Over all pairs, we determine the length of the shortest path. Intuitively, a path length of one most likely corresponds to a settlement in which one side is a dealer (D) and the other side is a dealer or customer (C), that is, D-D or D-C. Table 4 reports the distributions of shortest path lengths. In 2016 only 0.1% of the paths fall into this category. A path of length two might be C-D-C (54.01%); length three, C-D₁-D₂-C (45.81%). Note that the shortest path distribution is defined over all possible network paths, not the sample distribution of intermediation chains.

Finally, we examine nodes' neighbors. Financial networks are often characterized as core-periphery, hypothetically consisting of a small number of highly interconnected dealers and a larger number of customers each of whom is linked to at most a few of the dealers. If we provisionally define the top-twenty-five members as the "core", we can then ask how many of the remaining members are neighbors of (one step away from) at least one core member. The remaining (non-

neighbor) members are at least two steps distance from a core member. From an institutional perspective, the non-core intermediary might be a correspondent bank. Table 5 reports the counts. In 2013, all but 9.49% of settlement members are one step removed from a core member; in 2016 this has dropped to 5.88%.

C. Dealers

Are the most active and central members of the settlement network “dealers”? Although our identifiers are anonymized, the list of direct (that is, non-third-party) CLS members posted on the CLS website includes many major money-center banks that manage large FX desks. It is therefore likely that our top twenty-five set includes many dealers. It is then sensible to ask whether analysis of the settlement flows can illuminate dealers’ positions, trading and profits.

In this section we analyze imputed positions. Intermediaries are presumed to balance the rates at which buy and sell orders flow into the market. This does not suffice, however, to ensure their viability. The positions of cash and securities will tend to diverge over time, hitting limits imposed by regulation or prudent risk management. Garman (1976) suggests that a dealer might set her bid and ask quotes to elicit an order imbalance that would restore or stabilize their position. This mechanism underlies the inventory control models of market making (Amihud and Mendelson (1980); Hasbrouck and Sofianos (1993); Ho and Stoll (1980); Ho and Stoll (1981); Madhavan and Smidt (1993); Stoll (1976), among others). In currency markets, both Lyons (1995) and Yao (1997) study samples of dealer positions and find strong mean reversion around zero.

When the data record does not report an agent’s position, it is common to impute a position at a given time by cumulating all the agent’s trades and transfers up to that time. There is still some indeterminacy arising from the unknown starting position, but this simply means that the paths of the true and imputed positions will differ by a constant. If the true position is stationary about zero, the imputed position will be stationary around the starting value. On the other hand, if the trades and transfers are only partially observed, the imputed position will incorporate a cumulative error that will generally behave as a random walk.

If central members are de facto dealers, and if the settlement data constitute a comprehensive record of their trades, their imputed positions (cumulative settlement flows) should

be stationary. As an illustration Figure 1 plots the implied positions for a representative currency (the Australian dollar) during April 2016 for the top twelve members, based on volume-weighted degree centrality in settlements for the currency. The time scale covers 22 trading days (with weekends removed). In the AUD the positions of members ranked 5, 7, and 11 are visually relatively flat, but all suggest divergence at multiday horizons. Most of the other AUD plots might suggest periods of local stationarity, but the dominant pattern is more consistent with a random walk or time-trend. Plots for other currencies are similar. More formally, across the top-ranked members in all currencies, unit root tests on cumulative flows do not generally reject the null hypothesis (of a unit root) or time trend.

In summary, neither visual examination nor formal testing suggests that imputed positions are stationary. This implies that the cumulated settlements are poor approximations of actual dealing positions. It is perhaps unsurprising that the settlement data do not capture all activity. Some transfers are settled directly (for example, over the SWIFT system). More importantly, though, most trades that a bank makes with its customers are settled “on us”, that is, on the bank’s books. Although these trades might cause large inflows and outflows to and from a bank’s currency positions, they are not associated with any interbank settlements. This leakage might be one-sided if customer needs are correspondingly unbalanced. Among a bank’s customers, for example, imports of goods and services might dominate exports, or capital flows might favor home or foreign markets. Such effects might explain the time-trends displayed by some of the cumulative member flows.⁷

These limitations are important for the imputations of profits described in the following section. With a complete record of a bank’s trades and transfers, it would be possible to accurately compute the bank’s trading profits. Our estimates of trading costs and revenues are less comprehensive, and only reflect components of profits.

⁷ The restriction that a member’s true position (net of non-CLS flows) is stationary (does not contain a unit root) can be used to partially characterize the properties of the non-CLS flows. In the long run the cumulation of these flows must offset the long-run variation in the accumulated CLS flows.

V. Pricing and profits

We seek to characterize the relationship between the settlement price and the centralities of two participants. Our approach is to impute a trading return as the difference between the (log) settlement price and a nearby benchmark market price, and then model the dependence of this return on buyer and seller centralities. The profit computations resemble those used for conventional transaction cost analysis (TCA) in equity markets. They are also similar to the cost imputations used in Ramadorai (2008) to study a 1994-2001 sample of customer FX trades supplied by a custody bank. We first describe the general framework, and then turn to the econometric specifications.

A. Imputation framework

In an FX trade one currency is exchanged for another; the relative quantities given and received imply an exchange rate; each currency is received by one side of the trade and paid out by the other side. There is a high degree of symmetry. In economic terms, either currency could be considered the numeraire or “home” currency, and either side could be considered the buyer.

Organized FX trading nevertheless follows certain operational conventions. For a given pair of currencies, one is conventionally designated the base currency, and the other is the quote currency. In the EUR/USD market, for example, the base currency (the EUR) is bought or sold in partially standardized quantities (usually in multiples of one million euros) by payment or receipt of the quote currency (the USD). The price (exchange rate) in the market is stated as the number of USD that is being exchanged for one EUR.

We impute a return to the buyer of the base currency as $\pi = m - p$ where p is the log settlement price and m is a log benchmark price that ideally represents an unbiased valuation free of trading frictions. In equity market studies the average (midpoint) of the bid and offer at a nearby time is a common choice.⁸ We follow this practice using the Olsen quotes: we set m equal to the log

⁸ The return imputation, $\pi = m - p$, is similar to a measure used in equity transaction cost analysis (TCA). In the equity setting, the two sides of the trade are differentiated as liquidity providers (“makers” or “dealers”) or liquidity demanders (“takers” or “customers”). For example, when a customer buyer lifts a dealer’s offer at time t , paying p_t , the quantity $m_{t+\ell} - p_t$ for some delay $\ell > 0$ is considered the (customer’s) realized cost, and its negative is considered the dealer’s revenue. Customer buys and sells are treated asymmetrically. When a customer seller hits a dealer’s bid (receiving p_t), the realized cost is $p_t - m_{t+\ell}$, and the dealer’s revenue is $m_{t+\ell} - p_t$. Our

midpoint of the bid and ask quote immediately (within a few seconds) after the settlement, or (in alternative analyses) immediately prior to the settlement. The return to the seller of the base currency is $-\pi$. Since the returns are zero-sum it suffices to analyze the return to the buyer.

We model π as a function of the buyer's and seller's centralities. The economic network models reviewed above do not suggest an exact functional form, so we use a grouping approach that balances flexibility and parsimony. We order settlement members by centrality and assign them to groups indexed from one (low centrality) to five (high). The details of this assignment are discussed in detail below.

The statistical model is a linear regression in which π is regressed against classification dummy variables corresponding to fixed effects. The most important fixed effects here involve the interactions of buyer and seller centralities. For settlement k , these fixed effects are denoted by:

$$\mu_{ij} = E[\pi_k | \text{buyer} \in \text{group } i \text{ and seller} \in \text{group } j] \text{ where } i, j = 1, \dots, 5 \quad (1)$$

Because these effects are driven by the centralities of the buyer and seller, we designate the μ_{ij} as centrality differentials. With some abuse of notation, we will write the regression as:

$$\pi_k = \mu_{ij} + \epsilon_k \quad (2)$$

where i and j depend on k , and ϵ_k is a mean-zero error.⁹

The centrality differential μ_{ij} is the mean return differential (to the base currency buyer) conditional on a buyer from group i and a seller from group j . Symmetry suggests that the buyer's differential μ_{ij} should equal the seller's differential if the centrality positions were reversed, that is, $\mu_{ij} = -\mu_{ji}$. By the same reasoning, the diagonal of μ contains differentials for buyers and sellers in equal centrality groups, which should in expectation equal zero. We generally impose the zero-

data do not identify customer and dealer, however, so our profit imputation is always computed for the base-currency buyer. SEC Rule 605 requires market centers to report average realized costs computed with $\ell = 5$ minutes. The rule also mandates reporting of statistics based on the quote midpoint prevailing when the customer order arrives at the market center, that is, a pre-trade benchmark.

⁹ To express this more formally as a linear regression, define z_k as the 5×5 matrix with $z_{kij} = 1$ if the buyer is in centrality group i and the seller is in group j and $z_{kij} = 0$ otherwise. Then $\pi_k = \text{vec}(z_k)' \text{vec}(\mu) + \epsilon_k$, where $\text{vec}(\cdot)$ is the vectorization operator that stacks the columns of the argument in a single column.

diagonal restriction in estimation. (In alternative unconstrained estimations, not reported here for the sake of brevity, diagonal entries are generally small.)

Across all random assignments of buyer and seller groups we expect $E\pi = 0$. A non-zero exchange rate drift, however, will enter as an offset to the profits of all base-currency buyers. We allow for this by including an intercept in the regression:

$$\pi_k = \alpha + \mu_{ij} + \epsilon_k. \quad (3)$$

With an intercept and an unconstrained μ matrix, the regressors are linearly dependent. Zeroing the diagonal interaction terms removes these dependencies. Additionally, we introduce in some specifications fixed effects for currency pair and settlement size classifications.

Although symmetry of centrality effects suggests that $\mu_{ij} = -\mu_{ji}$, we do not impose this restriction in estimation, preferring to let the data support or refute the hypothesis. Nevertheless, assuming that, given one settlement member in group i and the other in j , either member is equally likely to be the buyer, it is sensible to average over these assignments. Letting $i^* = \max(i, j)$ and $j^* = \min(i, j)$, we define the average centrality differential as $\bar{\mu}_{i^*j^*} = (\mu_{i^*j^*} - \mu_{j^*i^*})/2$. This represents a centrality differential defined from perspective of the agent with the higher centrality, whether that agent is the buyer or the seller. It is zero by construction when $i = j$. If the $\bar{\mu}_{ij}$ for $i > j$ are generally positive, we describe the network as exhibiting a centrality premium; if negative, a centrality discount.

B. Econometric specifications

To complete the description of our statistical models, we expand the notation to reflect the dimensions and attributes of our data. First let κ index the thirteen currency pairs, $\kappa \in \{AUD/JPY, \dots, USD/MXN\}$. In a sample month (April of 2013 or 2016) we have Olsen bid and ask quotes at intervals of approximately ten seconds. Let $t = 1, \dots, T_\kappa$ index these intervals. Following the earlier discussion, the durations between quote times are not exactly ten seconds, nor are the intervals exactly aligned across currency pairs. Finally let $k = 1, \dots, N_{\kappa t}$ index the settlements for currency pair κ in interval t . The return calculations in most of our specifications use a post-settlement benchmark:

$$\pi_{kt\kappa}^{Post} = m_{t+1,\kappa} - p_{kt\kappa} \quad (4)$$

where $m_{t+1,\kappa}$ is the log bid-ask midpoint at the start of interval $t + 1$, that is, the interval following settlement k . (Alternative specifications use pre-settlement midpoints.) With this expanded notation, the regression specification (with intercept) is:

$$\pi_{kt\kappa}^{Post} = \alpha_{\kappa} + \mu_{ij} + \epsilon_{kt\kappa} \quad (5)$$

In this specification the μ_{ij} estimates are pooled over all currency pairs. We also present separate estimates by currency pair.

The structure of the regression errors in (5) is as follows. For a given currency pair, the settlements contained in an interval will typically involve diverse buyers and seller. We assume that price variation not explained by centrality is independent across settlements. All settlement returns in the same interval, however, depend on the same benchmark, $m_{t+1,\kappa}$. Given the limitations of the quote data, it is prudent to view the bid-ask midpoints as noisy estimates of fair-value prices. This noise introduces a common error component. The regression error covariances are therefore modeled as:

$$Cov(\epsilon_{kt\kappa}, \epsilon_{k^*,t\kappa}) = \begin{cases} \sigma_{f\kappa}^2 + \sigma_{g\kappa}^2 & \text{if } k = k^* \\ \sigma_{f\kappa}^2 & \text{if } k \neq k^* \end{cases} \quad (6)$$

Here, $\sigma_{f\kappa}^2$ impounds common-factor variation in the benchmark, and $\sigma_{g\kappa}^2$ is settlement-specific. Errors are assumed to be independent across different intervals and across currency pairs. Estimates and standard errors of the μ fixed effects are GLS, using the estimates of $\sigma_{f\kappa}^2$ and $\sigma_{g\kappa}^2$.

VI. Group construction

Our empirical specifications rely on grouping settlement members by centrality. We construct our centrality groups with three objectives: maintaining (for a given centrality measure) a monotonic ordering; ensuring adequate settlement volume in all groups; and, ensuring anonymity for the settlement members in each group. To promote the last objective, we require that each group have at least five settlement members. Groups are formed from high to low. Within each year/currency pair we rank members in descending order of (unweighted) degree centrality. For each member we compute settled volume (USD equivalent) and their proportion relative to the total. (Since each settlement has two sides the total settled volume is twice that of the usual reported volume.) Starting from the member with highest centrality we accept members into the

highest (fifth) group until there are at least five members and the cumulative volume proportion reaches twenty percent. The volume proportion target for the remaining four groups is reset to equal one quarter of the unallocated remainder from the highest group. The second-highest group is constructed in similar fashion, the volume proportion target is recalculated, and so forth.

Table 6 Panel A reports (by currency pair) counts, average degree centrality, and volume proportion for each group. For example, in the 2016 AUD/JPY pair the top five members (by centrality) account for 56.7% of the volume. With 43.3% of volume to be allocated to the remaining four groups, the target volume share in each group becomes 10.8%. To construct the second-highest group we continue through the ranked members until we have at least five more members with at least 10.8% of the volume. These five members account for 21.5% of the volume. The cumulative allocated volume is now 78.2%, and the target for the remaining three groups is set to 7.3%. The third-highest group also contains five members and 8.7% of the volume. The process continues for the lower groups.

For all currency pairs volume is concentrated in the larger groups. In April 2016, for all pairs, the top two groups contain the minimum number of members (five). In seven of the thirteen pairs, the third group also contains five members. The lowest two groups are more numerous. The volume proportions for the top two groups generally increase between 2013 and 2016, suggesting that concentration has also increased.

Table 6 Panel B reports statistics for groups formed by weighted degree centrality. The volume weighting tends to increase even further the volume concentration in the upper groups. In the AUD/JPY pair, for example, the top group formed using unweighted centrality accounts for 56.5% of the volume, and the top group using weighted centrality has 64.6% of the volume.

Table 7 Panel A reports settlement volumes between base-currency buyer and seller groups. For example, in April 2016, the total settlement value between group-1 buyers and group-5 sellers is 366,376M USD. In these tabulations the unit of observation is the settlement: the volume is not doubly counted. Table 7 Panel B reports these values as percentages. (The percentages in each 5×5 annual block total to one hundred.) The blocks in both panels are approximately diagonal. We'd expect that in a settlement flow each party is equally likely to be the buyer, but no accounting identity ensures a strict symmetry.

It is not surprising that there are high activity within and between groups 4 and 5. Most, possibly all, members of these groups might be dealers. Active interdealer trading occurs in most OTC markets. It is worth commenting, however, on the activity within and between the lower centrality groups. In a pure dealer market, customer-to-customer trade is not possible. In the present sample, activity within and between groups 1 and 2 is low, but certainly non-zero. Finally, we note that, consistent with the increase in concentration, volume percentages involving the top two centrality groups increase between 2013 and 2016.

VII. Results

A. Centrality differentials, pooled estimates

Our baseline specification is equation (5). While we report below separate estimates for each year and currency pair, the broad features of the estimates are most distinctive in a pooled specification estimated in each year over all currency pairs. More specifically, the centrality differentials μ_{ij} are pooled (and constrained to zero for $i = j$). Intercepts and variance parameters, however, are specific to the currency pair.

Table 8, Panel A, reports the μ_{ij} estimates (2016 on the left; 2013 on the right). The estimated $\mu_{21} = 0.058$, for example, implies that when a group-2 buyer trades with a group-1 seller, the buyer realizes an average additional return of 0.058 basis points. Conversely when a group-1 buyer trades with a group-2 seller, the average incremental return to the buyer is -0.106 bp, that is, a loss to the buyer and a gain to the seller.

The 2016 estimates exhibit several distinctive patterns. Firstly, the lower-triangle entries in are positive (consistent with a buyer advantage when the buyer has the larger centrality). Secondly the upper-triangle entries are negative (consistent with a buyer disadvantage when the seller has the higher centrality). Thirdly, there is a general consistency between the orderings of gains and relative centrality advantage. The fifth row, for example, corresponds to settlements in which a group-5 (highest centrality) buyer settles against (left to right) sellers of increasing centrality. As the difference in centrality narrows, the buyers' differential shrinks. Similarly, the fifth column corresponds to trades in which a group-5 seller settles against buyers of increasing centrality (top to bottom). Here too, as the seller's relative centrality advantage shrinks, expected seller gains

(formally, buyer losses) decline. This monotonicity should in principle apply to all rows and columns. Although this is a general pattern, it is not uniform. In the first column, for example, when trading with a group-1 seller, a group-5 buyer has a smaller gain than a group-4 buyer (0.126 bp vs 0.129 bp).

The broad features of the centrality differentials are important for our conclusions. At first glance it might be conjectured that grouping agents by centrality is simply an alternative way to recover the customer-dealer distinction, that agents with high centrality are dealers, those with low centrality are customers, and that our centrality premium reflects nothing more than the dealer markup customarily charged to the customer. While we don't reject the logic of this assertion, we question whether it can account for all our results. The sub- and super-diagonal elements of the matrix of the μ estimates are centrality differentials between adjacent groups. If all members of groups 1 and 2 were otherwise undifferentiated customers, for example, the centrality differentials for settlements involving a group-1 buyer and group-2 seller would be zero. The 2016 pooled estimate is -0.106 (with a t-value of -12.09). Similarly, the estimated centrality differential for a group-2 buyer and group-1 seller is 0.058 (6.48). The argument also applies to adjacent upper groups. If all members of groups 4 and 5 were otherwise undifferentiated dealers, the centrality differentials μ_{54} and μ_{45} would also be zero; the estimates suggest otherwise.

The absence of a clear dealer/customer demarcation is consistent with broad trends in market structure. Commenting on the 2019 BIS survey results, Schimpf and Sushko (2019a, b) note an increased presence of nonbank liquidity suppliers, a decline in interbank volume, and a rise in prime-brokered customer access to platforms traditionally restricted to dealers.

Panel B summarizes other model parameters. In both years the α_κ intercepts are generally small (although statistically significant). In both years and in all currency pairs, the common component of the regression errors ($\sigma_{f\kappa}^2$) generally dominates the idiosyncratic component ($\sigma_{g\kappa}^2$).

Table 8 Panel C reports the $\bar{\mu}_{ij}$ average centrality differentials. For example, when a group-4 member trades with a group-1, the centrality differential (averaged over group-4 buyers and group-4 sellers) is 0.135 bp. Here, too, the gains to the centrality-dominant side are generally increasing in the difference in the counterparties' centralities, but not uniformly ($\bar{\mu}_{21} < \bar{\mu}_{31} <$

$\bar{\mu}_{41} \neq \bar{\mu}_{51}$, for example). The average centrality differentials between adjacent centrality groups are generally positive and significant: the centrality premium is pervasive.

The general features of the 2013 and 2016 estimates are similar. For the 2013 centrality differentials, the lower-triangle values are uniformly positive (implying a centrality advantage for buyers), but the upper-triangle values are not uniformly negative. The average centrality differentials are consistently monotonic, however.

Although highly statistically significant, the magnitude of the average centrality differential is small. Consider a representative trade in which one million British pounds are purchased for 1.4 million US dollars. If the less-central side is group-1 and the more-central is group-5, the estimate of $\bar{\mu}_{51} = 0.1250$ basis points implies a gain to the group-5 member of $0.125 \times 10^{-4} \times 1.4 \times 10^6 = 17.50$ US dollars. The high settlement volumes, on the other hand, imply larger aggregate amounts. By applying the centrality differentials (Table 8 Panel A) to the intergroup settlement volumes (Table 7 Panel A), we may estimate an aggregate value. This is approximately 42.4M USD (for 2016) and 34.0M USD (2013). Since there are approximately twenty trading days in 2016, this implies a daily gain for the more central settlement members of about two million dollars.

B. Centrality differentials by settlement size, pooled estimates

Hasbrouck and Levich (2019) discuss in detail the settlement sizes in the CLS samples. The modal size is one million units of the base currency. Approximately 40% of the settlements are for this quantity. To assess the relation between size and price, we define a classification variable $Size \in \{< 1M, \geq 1M\}$ and modify regression specification (5) to interact size with centrality:

$$\pi_{kt\kappa}^{Post} = \alpha_{\kappa} + \mu_{ij} \otimes Size + \epsilon_{kt\kappa} \quad (7)$$

The size interaction effectively doubles the number of μ_{ij} differentials, and for brevity we do not report the full set of estimates.

Table 9 Panel A presents the estimates of the average centrality differentials $\bar{\mu}_{ij}$ for $i > j$. The estimates for small settlements are on the top; those for large settlements are below; the large-minus-small differences are at the bottom. The $\bar{\mu}_{ij}$ estimates are generally positive for both small and large settlements, and the magnitudes are generally increasing with the difference between

centrality groups. Consistent with earlier results, this implies an advantage to the side with the relatively higher centrality, a centrality premium.

The $\bar{\mu}_{ij}$ estimates are generally higher for small settlements. The large-minus-small differences are mostly negative and statistically significant. This implies that the centrality advantage is lower for larger settlements. This is broadly consistent with the evidence from other OTC markets that smaller dealer-to-customer trades have higher markups. Representative studies include: Reiss and Werner (1996) and Bernhardt, Dvoracek, Hughson and Werner (2004) for UK equities; Green, Hollifield and Schurhoff (2007) and Harris and Piwowar (2006) for US municipal bonds; Bessembinder, Maxwell and Venkataraman (2006) and Edwards, Harris and Piwowar (2007) for US corporate bonds.

C. Centrality differentials with ex ante benchmarks, pooled estimates

The estimates in Tables 8 and 9 are based on the π^{Post} return given in equation (4), which benchmarks the price to a post-settlement bid-ask midpoint. To investigate differential returns based on a pre-settlement benchmark, we define:

$$\pi_{ktk}^{Pre} = m_{t,\kappa} - p_{ktk} \quad (8)$$

In equity transaction cost measurement a post-trade benchmark is generally viewed as impounding the price impact attributable to the trade. A pre-trade benchmark is in principle independent of all effects of the trade (except for “leakage” of the trader’s intentions).

In the earlier discussion of settlement characteristics, we noted that settlement instructions are submitted with a delay relative to the trade occurrence. The effect of this delay on the calculation of π_{ktk}^{Post} is likely to be minimal: a quote timestamped after the settlement time (like $m_{t+1,\kappa}$ in (4)) is almost certain to lag the actual trade time as well. Determination of π_{ktk}^{Pre} is more sensitive. Given submission latencies, a benchmark midpoint that appears to be pre-settlement (like $m_{t,\kappa}$ in (8)) might actually have been set after the occurrence of the trade that generated the settlement.

Table 10 reports estimates of regression specification (5) with π_{ktk}^{Pre} as the dependent variable. In lieu of the full set of centrality differentials, Panel A reports the average centrality

differentials $\bar{\mu}_{ij}$ for $i > j$. They are very close to the corresponding values computed from the π_{ktk}^{Post} specifications (Table 8, Panel B).

D. Centrality differentials with volume-weighted centrality groups, pooled estimates

The estimates of fixed effects and centrality profits in Table 11 are based on member groupings using volume-weighted centrality. The estimates differ from those based on unweighted centrality reported in Table 8, but they are broadly similar in sign and magnitude. This similarity is important due to the equivalence of ranking by volume-weighted centrality or by proportional volume share. While most centrality measures attempt to illuminate some distinctive property of the node in the context of the network, volume share is a straightforward alternative measure of relative importance.

E. Centrality differentials by currency pair

It is clear from Table 6 that the number and concentration of settlement members varies considerably across currency pairs. In 2016 the top five settlement members account for 56.6% of the settlement volume in the GBP/JPY pair, but only 35.6% in the USD/CAD. The estimates of centrality effects reported to this point are pooled over all pairs. The pooling illuminates the broader properties of the centrality effects but suppresses differences among pairs.

Most importantly, settlement patterns within and between groups vary considerably. Table 12 reports the settlement volume percentages between groups (analogous to Table 7 Panel B) by currency. Within each year-pair block, settlements within group 5 (the most central) are always high, but are often smaller than group 5 settlements against lower groups. This is consistent with the supposition that order flow involving lower groups is netted out before the imbalance is passed on in the interdealer market. The volume of settlements within and between lower groups is small, but nontrivial. In the EUR/USD pair, for example, 4.4% of settlement volume occurs within and between groups 1 and 2.

We next estimate specification (5) separately for each currency pair and year. Table 13 Panel A reports estimates of average centrality differentials $\bar{\mu}_{ij}$. Given the smaller size of these subsamples, the estimates are noisier and t-values are lower, relative to the pooled those derived from the full sample. The monotonicity patterns in $\bar{\mu}_{ij}$ across the centrality pairings are still

evident, but they are less consistent. Nevertheless, for settlements between group 5 and group 1, $\bar{\mu}_{51}$ is positive and statistically significant. It is usually (but not always) the largest entry. The $\bar{\mu}_{51}$ estimates vary considerably across currency pairs, ranging from 0.073 bp (for the EUR/USD) to 0.444 bp (NZD/USD).

To gauge the magnitudes of the average centrality differentials, we may provisionally assume that the most central settlement members (group 5) are dealers and the least central members (group 1) are customers. In this case, the centrality profit $\bar{\mu}_{51}$ should be approximately equal to one-half the bid-ask spread, the cost of liquidity for one leg of a round-trip trade. Table 14 reports the $\bar{\mu}_{51}$ (estimates from Table 12) and half-spread estimates from Hasbrouck and Levich (2019); Figure 2 depicts a log-log scatterplot. The correlation between the logs of the two series is 0.71. A simple log-log regression implies that the half-spreads are about three times the $\bar{\mu}_{51}$ value. Viewed as an estimate of the half-spread, therefore, $\bar{\mu}_{51}$ is certainly biased downwards. Considering that it is constructed without the benefit of explicit customer/dealer identification, however, it is a remarkably good proxy.

VIII. Centrality-based order flow

The analysis to this point establishes that in a trade between agents of unequal centrality the more central player realizes a larger share of the trading profit, that is, a centrality premium. The possible reasons for this premium include bargaining power, implicit compensation for some other services (like intermediation speed in Li and Schürhoff (2019a)), or bearing the risk of exposure to adverse price movements. In this section we investigate the latter, by examining the relations between flows originating from central vs. non-central participants and contemporaneous returns.

Evans and Rime (2019) survey the extensive literature on currency order flows and returns. Evidence confirming a connection comes from a wide range of time periods and data sources, but the one set of authoritative studies is based on the accurately time-stamped records of trades and quotes from the electronic interdealer platforms (Chaboud, Chiquoine, Hjalmarsson and Vega (2014); Hagströmer and Menkveld (2019); Mancini, Ranaldo and Wrampelmeyer (2013), among others). A second major group of studies examines order flows differentiated by participant

classification. Evans and Rime (2016) study order flow in the Norwegian kroner broken out into banks, non-bank financial and non-financial customers. Bjonnes, Osler and Rime (2009) examine a sample drawn from the trading records at a single bank, where the differentiation is in the bank's counterparties. Rinaldo and Somogyi (2018, RS) analyze the information content of hourly flows classified as banks, investment funds, non-bank/non-fund financials (such as insurance) and corporate (any non-financial).

The present analysis is most strongly connected to RS in that their flows are also constructed from CLS settlements. Despite this commonality of source, however, the two samples differ markedly. Most significantly, in RS the settlement sides are classified by type. Our anonymized identifiers are not mapped to type. The flows in RS are aggregated over type and over time (hourly); ours are disaggregated. The identifiers and disaggregation in the present study are essential to the centrality determinations discussed earlier. RS' sample covers a longer span than ours (2012 to 2017) and it is continuous. Most importantly, the flows in RS's data are also differentiated as maker vs. taker (liquidity supplier vs. demander). This last feature is key to RS's identification of the information content of these flows, and its absence precludes us from making similar attributions.

The analysis of our profit attributions explores cross-section variation in terms of trade, that is, as a function of the centralities of the participants. Here, we analyze time variation. We examine the correlation between flows differentiated by centrality and contemporaneous returns. The unit of observation in the cross-sectional regression is a settlement; the unit of observation here is a five-minute interval. We stress that we are only asserting correlation, not causation. The usual sequential trade models posit a causal link as the dealer revises her bid and ask quotes in the direction of the most recent customer trade. Momentum and some dynamic hedging strategies, however, generate order flow in reaction to returns.

We use the five centrality groups based on unweighted degree centrality adopted above, and determine for each settlement the centrality groups of the buyer of the base currency and the seller. For a given (April) year, currency pair, and approximate 5-minute interval t , let v_{ijt} denote the total volume (USD equivalent) of all settlements in the interval in which the base currency buyer's centrality group is i and the seller's group is j . Our net centrality flow is

$$NCF_t = \sum_{i=1}^5 \sum_{j=1}^5 \text{sign}(i - j) v_{ijt} \quad (9)$$

In this sum, settlement volume is signed positive when the buyer has the centrality advantage, and negative when the seller has the advantage. When buyer and seller are in the same group, the volume contributes nothing. The construction is neutral in that there is no presumption that more-central counterparties are relatively informed or uninformed, or more likely to be liquidity suppliers or demanders. Because order impact in most markets has been found to be a concave function of size, we construct a signed logarithm as:

$$\text{Log}NCF_t = \text{Sign}(NCF_t) \log(1 + |NCF_t|) \quad (10)$$

The estimated specification is then

$$r_t = a + b \text{Log}NCF_t + e_t \quad (11)$$

where r_t is the return on the base currency over the interval (in basis points, using bid-ask midpoints).

Table 15 reports the estimates and t-values. Most of the b estimates are negative and statistically significant. This is clearest in the specifications for 2016, where the b estimate is negative in eleven (and statistically significant in six) of the thirteen pairs. Estimates for 2013 are similar.

The implication is that the more central members are buying in falling markets. This behavior is broadly consistent with liquidity provision. If more-central members are buying, less-central members are selling, and acting as liquidity demanders. In this view, the trading gains that the more-central members realize against less-central members are offset by price declines in acquired positions. We earlier raised the question of whether the central members are dealers. Analysis of positions inferred from cumulative settlements did not support this conjecture, but the settlement record is only a partial one. The present analysis of central flows is more supportive. The signs of the b estimates do not suggest that more central players are informed or that they are momentum traders.

IX. Conclusions

The network perspective has provided a framework for analyzing numerous over-the-counter markets. Centrality, in the broad sense of connectivity and embeddedness, facilitates

intermediation. All else equal centrality lowers the costs for the better-connected agents. These cost reductions might be passed on to counterparties and end-users (a centrality discount).

Alternatively, centrality might be associated with network dominance, greater bargaining power, and higher charges (a centrality premium). The evidence on existing markets is mixed. This paper examines the question as it pertains to the FX market, which is, at least by value of traded volume, by far the largest and most economically important.

Our analysis is based on CLS settlement data for most major pairs in the Aprils of 2013 and 2016, a sample that is unusually comprehensive in scope. The institutions in these data carry anonymized identifiers, which allow us to construct the network and its measures. They include both small customers and major dealers. The former might trade with one counterparty; the latter might have thousands. Activity is very concentrated in a few members who are large and central.

This study presents strong evidence for a centrality premium in this market. As a representative figure, when a member in the highest centrality group trades against a seller in the lowest group (or vice versa), the former an average centrality differential of about 0.125 basis points (cf. Table 8, Panel C). Averaged across all currency pairs, this is approximately one third of the bid-ask (half) spread. This differential appears across a range of currency pairs, in pooled estimates, with ex ante and ex post benchmarks, using both volume-weighted and unweighted degree centrality measures, and in both 2013 and 2016 samples.

Although the evidence supporting a centrality premium in FX market is strong, the underlying contributing factors are less clear. The premium is consistent with central agents possessing stronger bargaining power. In discussing the centrality premium in the US municipal bond market, however, Li and Schürhoff (2019a) note that the more central brokers are providing faster intermediation, a differentiated service which may be valuable to their customers and might potentially explain the premium.

Li and Schürhoff's identification of intermediation speed is not meaningful in our setting, but we investigate another mechanism that might work in a similar fashion. Our characterization of the FX centrality premium examines cross-sectional differences in prices paid and received by central vs. non-central participants. Across time, however, there is also variation and persistence in the aggregate flows to and from the central participants. These are associated with exchange rate

movements. When central participants are (in the aggregate) buying the base currency, the less-central agents are selling. If the central agents were informed traders, their purchases would coincide with increases in the exchange rate. We find the opposite, that the central participants are buying into declines in the exchange rate. This behavior is more consistent with liquidity provision. In this view, the central premium may be in part compensation for supplying liquidity to aggregate order imbalances. We note, however, that the estimates in our flow/exchange rate regressions are not as uniform and statistically significant as those supporting the centrality premium. As such, we view our evidence of liquidity provision as suggestive rather than definitive.

Finally, although the CLS data are more comprehensive than most alternative sources, they do not cover trades, such as those occurring between a bank and the bank's customers, where settlement would simply occur on the books of the bank. These would not normally require settlement via a third-party intermediary. Our results, therefore, can at best suggest profit attributions and transfers occurring between agents who belong to the clearing systems. Participants that are peripheral within the network, and who may therefore have poor terms of trade within the network, may nevertheless be in a stronger position when trading against their own customers.

X. References

- Adrian, Tobias, Michael J. Fleming, and Erik Vogt, 2017, An Index of Treasury Market Liquidity: 1991-2017 (Federal Reserve Bank of New York).
- Amihud, Yakov, and Haim Mendelson, 1980, Dealership markets: Market-making with inventory, *Journal of Financial Economics* 8, 31-53.
- Bank for International Settlements, 2016, Triennial Central Bank Survey: Foreign Exchange Turnover in 2016.
- Bernhardt, Dan, Vladimir Dvoracek, Eric Hughson, and Ingrid M Werner, 2004, Why do larger orders receive discounts on the London Stock Exchange?, *The Review of Financial Studies* 18, 1343-1368.
- Bessembinder, Hendrik, William Maxwell, and Kumar Venkataraman, 2006, Market transparency, liquidity externalities, and institutional trading costs in corporate bonds, *Journal of Financial Economics* 82, 251-288.
- Bjonnes, Geir H., Carol L. Osler, and Dagfinn Rime, 2009, Asymmetric Information in the Interbank Foreign Exchange Market (SSRN), <http://ssrn.com/paper=1343833>.
- Chaboud, Alain, Ben Chiquoine, Erik Hjalmarsson, and Clara Vega, 2014, Rise of the machines: algorithmic trading in the foreign exchange market, *Journal of Finance* 60, 2045-2084.
- Chaboud, Alain, Avery Dao, and Clara Vega, 2019, What makes HFT's tick? (SSRN), <https://ssrn.com/abstract=3407970>, doi:<https://dx.doi.org/10.2139/ssrn.3407970>.
- Di Maggio, Marco, Amir Kermani, and Zhaogang Song, 2017, The value of trading relations in turbulent times, *Journal of Financial Economics* 124, 266-284.
- Duffie, D., N. Garleanu, and L. H. Pedersen, 2005, Over-the-counter markets, *Econometrica* 73, 1815-1847.
- Edwards, Amy K., Lawrence E. Harris, and Michael S. Piwowar, 2007, Corporate Bond Market Transaction Costs and Transparency, *The Journal of Finance* 62, 1421-1451.
- Evans, Martin D. D., and Richard K. Lyons, 2002, Order Flow and Exchange Rate Dynamics, *The Journal of Political Economy* 110, 170-180.
- Evans, Martin D. D., and Dagfinn Rime, 2016, Order flow information and spot rate dynamics, *Journal of International Money and Finance* 69, 45-68.
- Evans, Martin D.D., and Dagfinn Rime, 2019, Microstructure of Foreign Exchange Markets (SSRN), <https://ssrn.com/abstract=3345289>, doi:<https://dx.doi.org/10.2139/ssrn.3345289>.
- Feldman, Allan M., 1973, Bilateral trading processes, pairwise optimality, and Pareto optimality, *Review of Economic Studies* 50, 463-473.
- Fleming, Michael J, and Eli M Remolona, 1997, Price formation and liquidity in the US Treasury market: evidence from intraday patterns around announcements (SSRN), <https://ssrn.com/abstract=943513>.
- Fleming, Michael J., Bruce Mizraeh, and Giang Nguyen, 2018, The microstructure of a U.S. Treasury ECN: The BrokerTec platform, *Journal of Financial Markets* 40, 2-22.
- Gabrieli, Silvia, and Co-Pierre Georg, 2017, A network view on interbank liquidity (SSRN), <https://ssrn.com/abstract=2544603>.

- Garman, Mark, 1976, Market microstructure, *Journal of Financial Economics* 3, 257-275.
- Green, Richard C., Burton Hollifield, and Norman Schurhoff, 2007, Financial Intermediation and the Costs of Trading in an Opaque Market, *Rev. Financ. Stud.* 20, 275-314.
- Hagströmer, Björn, and Albert J Menkveld, 2019, Information revelation in decentralized markets, *Journal of Finance* forthcoming.
- Harris, L. E., and M. S. Piwowar, 2006, Secondary trading costs in the municipal bond market, *Journal of Finance* 61, 1361-1397.
- Hasbrouck, Joel, and Richard M. Levich, 2019, FX liquidity and market metrics: new results using CLS bank settlement data (Finance Dept., Stern School, New York University), <https://ssrn.com/abstract=2912976>, doi:<https://ssrn.com/abstract=2912976>.
- Hasbrouck, Joel, and George Sofianos, 1993, The trades of market makers: An empirical examination of NYSE specialists, *Journal of Finance* 48, 1565-1593.
- Hau, Harald, Peter Hoffmann, Sam Langfield, and Yannick Timmer, 2019, Discriminatory pricing of over-the-counter derivatives (International Monetary Fund).
- Hendershott, Terrence, Dan Li, Dmitry Livdan, and Norman Schürhoff, 2019, Relationship Trading in OTC Markets, *Journal of Finance* forthcoming.
- Ho, Thomas S. Y., and Hans R. Stoll, 1980, On Dealer Markets Under Competition, *Journal of Finance* 35, 259-267.
- Ho, Thomas, and Hans R. Stoll, 1981, Optimal dealer pricing under transactions and return uncertainty, *Journal of Financial Economics* 9, 47-73.
- Hollifield, Burton, Artem Neklyudov, and Chester Spatt, 2017, Bid-Ask Spreads, Trading Networks, and the Pricing of Securitizations, *The Review of Financial Studies* 30, 3048-3085.
- Julliard, Christian, Zijun Liu, Seyed E. Seyedan, Karamfil Todorov, and Kathy Yuan, 2019, What drives repo haircuts? evidence from the UK market (SSRN), <https://ssrn.com/abstract=3374969>.
- King, Michael R., Carol L. Osler, and Dagfinn Rime, 2012, Foreign Exchange Market Structure, Players, and Evolution, in Jessica James, Ian W. Marsh, and Lucio Sarno, eds.: *Handbook of Exchange Rates* (John Wiley & Sons).
- King, Michael R., Carol L. Osler, and Dagfinn Rime, 2013, The market microstructure approach to foreign exchange: Looking back and looking forward, *Journal of International Money and Finance* 38, 95-119.
- León-Janampa José, Henry, 2017, Network analysis and systemic FX settlement risk, *Statistics & Risk Modeling* 34, 91.
- Li, Dan, and Norman Schürhoff, 2012, Dealer networks (SSRN), <http://ssrn.com/paper=2023201>, doi:10.2139/ssrn.2023201.
- Li, Dan, and Norman Schürhoff, 2019a, Dealer networks, *The Journal of Finance* 74, 91-144.
- Li, Dan, and Norman Schürhoff, 2019b, Internet appendix for "Dealer Networks," *Journal of Finance*, doi:10.1111/jofi.12728.
- Lyons, Richard K., 1995, Tests of microstructural hypotheses in the foreign exchange market, *Journal of Financial Economics* 39, 321-351.

- Madhavan, Ananth, and Seymour Smidt, 1993, An analysis of changes in specialist inventories and quotations, *Journal of Finance* 48, 1595-1628.
- Mancini, Lorian, Angelo Ranaldo, and Jan Wrampelmeyer, 2013, Liquidity in the foreign exchange market: measurement, commonality, and risk premiums, *Journal of Finance* 68, 1805-1841.
- Moore, Michael, Andreas Schrimpf, and Vladyslav Sushko, 2016, Downsized FX markets: causes and implications, *BIS Quarterly Review* December, 35-51.
- Neklyudov, Artem, 2013, Bid-ask spreads and the decentralized interdealer markets: Core and peripheral dealers (University of Lausanne).
- Neklyudov, Artem, 2019, Bid-ask spreads and the over-the-counter interdealer markets: Core and peripheral dealers, *Review of Economic Dynamics*.
- Ramadorai, Tarun, 2008, What determines transaction costs in foreign exchange markets?, *International Journal of Finance & Economics* 13, 14-25.
- Ranaldo, Angelo, and Fabricius Somogyi, 2018, Heterogeneous Information Content of Global FX Trading (University of St.Gallen), <https://ssrn.com/abstract=3263279>, doi:10.2139/ssrn.3263279.
- Reiss, Peter C, and Ingrid M Werner, 2004, Anonymity, adverse selection, and the sorting of interdealer trades, *The Review of Financial Studies* 18, 599-636.
- Reiss, Peter C., and Ingrid M Werner, 1996, Transaction costs in dealer markets: evidence from the London Stock Exchange, in Andrew W Lo, ed.: *The industrial organization of the securities industry* (University of Chicago Press, Chicago).
- Reiss, Peter C., and Ingrid M. Werner, 1998, Adverse selection in Dealers' Choice of Interdealer Trading Systems (Graduate School of Business, Stanford University), <https://ssrn.com/abstract=168188>.
- Schrimpf, Andreas, and Vladyslav Sushko, 2019a, FX trade execution: complex and highly fragmented, *BIS Quarterly Review* 39-51.
- Schrimpf, Andreas, and Vladyslav Sushko, 2019b, Sizing up global foreign exchange markets, *BIS Quarterly Review* 21-38.
- SIFMA, 2016, SIFMA Electronic Bond Trading Report: US Corporate & Municipal Securities.
- Stoll, Hans, 1976, Dealer inventory behavior: An empirical examination of Nasdaq stocks, *Journal of Financial and Quantitative Analysis* 11, 359-380.
- Yao, Jian, 1997, Market making in the interbank foreign exchange market, PhD Dissertation (Stern School of Business, New York University, available through Proquest, Document ID 304422990).

XI. Tables

Table 1. Distribution of CLS activity across participants, by April/year

CLS settlements in the Aprils of 2013, and 2016. Member counts reflect only those members who engaged in at least one settlement. Each settlement has two sides: the number of sides is $2 \times$ the number of settlements.

Panel A. Distribution quantiles and mean of the top five refer to $n_i = \# \text{ settlement sides}$ for member i (all instruments, all currencies).

	No. sides	No. members	Distribution quantiles							Mean of top five
			5%	10%	25%	50%	75%	90%	95%	
2016	21,824,492	16,163	1	2	5	14	35	81	155	2,208,252
2013	27,212,924	11,651	1	2	6	15	39	109	237	2,490,092

Panel B. Distribution quantiles and mean of the top five refer to $n_{ij} = \# \text{ settlement sides}$ for member i in instrument type j (all currencies).

	Instrument Type	No. sides	No. members	Distribution quantiles							Mean of top five
				5%	10%	25%	50%	75%	90%	95%	
2016	Spot	19,877,612	10,737	1	1	2	5	15	42	84	2,098,130
	Forward	655,348	11,332	1	1	2	5	14	32	60	46,316
	Swap (Near)	490,516	7,270	1	1	2	6	15	30	51	29,314
	Swap (Far)	486,026	7,140	1	1	2	6	14	30	52	29,267
	Option	75,082	1,383	1	1	1	1	3	15	69	6,309
	Other	239,908	7,164	1	1	1	3	8	19	36	16,737
2013	Spot	24,980,670	7,262	1	1	2	5	15	41	428	2,359,721
	Forward	951,088	9,090	1	1	2	7	18	52	98	69,911
	Swap (Near)	472,234	5,638	1	1	3	6	16	41	143	27,055
	Swap (Far)	477,112	5,635	1	1	3	6	16	42	142	27,501
	Option	45,280	551	1	1	1	5	23	126	272	3,042
	Other	286,540	4,771	1	1	1	3	6	24	80	21,131

Panel C. Distribution quantiles and mean of the top five refer to $n_{i\kappa} = \# \text{ settlement sides for member } i \text{ in currency pair } \kappa$ (spot settlements only).

	Currency pair	No. sides	No. members	Distribution quantiles							Mean of top five
				5%	10%	25%	50%	75%	90%	95%	
2016	AUD/JPY	361,472	239	1	1	1	4	53	1,114	7,883	46,036
	AUD/USD	1,837,326	2,462	1	1	1	2	7	19	88	198,266
	EUR/CHF	251,474	803	1	1	1	2	8	93	339	28,056
	EUR/GBP	554,492	1,533	1	1	1	2	6	34	219	57,598
	EUR/JPY	512,270	683	1	1	1	2	15	194	811	61,743
	EUR/USD	4,243,332	5,699	1	1	1	3	10	21	63	451,291
	GBP/JPY	284,468	511	1	1	1	3	6	69	595	36,206
	GBP/USD	1,751,862	4,395	1	1	1	3	8	18	49	180,501
	NZD/USD	684,980	784	1	1	1	2	6	72	631	79,834
	USD/CAD	1,490,500	1,153	1	1	1	2	10	87	935	155,860
	USD/CHF	557,830	2,399	1	1	1	2	5	13	39	63,476
	USD/JPY	3,702,542	2,494	1	1	1	2	5	61	242	393,765
	USD/MXN	520,382	1,106	1	1	1	2	7	22	128	63,506
2013	AUD/JPY	533,456	326	1	1	1	9	79	1,682	4,143	66,339
	AUD/USD	1,999,050	1,336	1	1	1	2	13	245	1,942	182,630
	EUR/CHF	436,544	829	1	1	1	4	26	217	853	40,070
	EUR/GBP	563,310	1,414	1	1	1	3	9	81	364	53,578
	EUR/JPY	1,852,548	751	1	1	1	6	81	683	3,798	216,964
	EUR/USD	5,724,600	3,515	1	1	1	3	9	99	763	513,372
	GBP/JPY	324,614	455	1	1	1	2	16	192	2,060	38,637
	GBP/USD	1,786,302	2,713	1	1	1	2	7	44	199	165,451
	NZD/USD	553,414	542	1	1	1	5	38	757	2,375	54,785
	USD/CAD	1,085,714	747	1	1	1	3	31	482	3,452	100,239
	USD/CHF	693,162	1,346	1	1	1	2	6	71	397	68,521
	USD/JPY	6,240,280	1,884	1	1	1	3	12	230	1,920	655,056
	USD/MXN	546,588	877	1	1	1	3	7	64	695	54,571

Table 2. Settlement activity by value (all instruments, all currencies)

For each settlement member we compute the value of all their settlement sides (USD equivalent) over the April of the indicated year. We rank them in descending order and report the cumulative percentage (relative to the value of settlements sides across all members).

Rank by value of settlement sides	Cumulative % of value	
	2016	2013
5	33.3	31.6
10	52.3	46.5
20	68.9	60.1
50	87.8	76.6
100	93.4	87.4
200	96.0	93.9
500	97.8	98.0
1,000	98.7	99.1

Table 3. Distribution of degree centrality

CLS settlements in the Aprils of 2013 and 2016. For a given April/year, a participant is included as a network node if they engaged in at least one settlement in any pair, and a link between two participants is imputed if they were counterparties in at least one settlement. The degree centrality of a participant (node) is the number of links (edges) incident at that node.

Panel A. Member degree centrality, all currencies, all instruments.

	No. active members	Distribution quantiles									Mean of top five
		1%	5%	10%	25%	50%	75%	90%	95%	99%	
2016	16,163	1	1	1	2	4	7	11	13	22	5,689
2013	11,651	1	1	1	2	4	8	12	15	81	4,004

Panel B. Member degree centrality, by instrument, all currencies

		No. active members	Distribution quantiles									Mean of top five
			1%	5%	10%	25%	50%	75%	90%	95%	99%	
2016	Spot	10,737	1	1	1	2	5	15	42	84	3,503	2,098,130
	Forward	11,332	1	1	1	2	5	14	32	60	211	46,316
	Swap (Near)	7,270	1	1	1	2	6	15	30	51	646	29,314
	Swap (Far)	7,140	1	1	1	2	6	14	30	52	660	29,267
	Option	1,383	1	1	1	1	1	3	15	69	1,281	6,309
	Other	7,164	1	1	1	1	3	8	19	36	298	16,737
2013	Spot	7,262	1	1	1	2	5	15	41	428	20,246	2,359,721
	Forward	9,090	1	1	1	2	7	18	52	98	347	69,911
	Swap (Near)	5,638	1	1	1	3	6	16	41	143	1,156	27,055
	Swap (Far)	5,635	1	1	1	3	6	16	42	142	1,156	27,501
	Option	551	1	1	1	1	5	23	126	272	2,104	3,042
	Other	4,771	1	1	1	1	3	6	24	80	793	21,131

Panel C. Member degree centrality, spot settlements, Olsen currency pairs

		No. active members	Distribution quantiles							Mean of top five	
			1%	5%	10%	25%	50%	75%	90%		95%
2016	AUD/JPY	239	1	1	1	1	2	8	20	55	84
	AUD/USD	2,462	1	1	1	1	2	3	6	10	493
	EUR/CHF	803	1	1	1	1	2	4	11	23	193
	EUR/GBP	1,533	1	1	1	1	1	3	8	18	312
	EUR/JPY	683	1	1	1	1	1	5	15	28	175
	EUR/USD	5,699	1	1	1	1	2	4	8	10	1,442
	GBP/JPY	511	1	1	1	1	2	3	11	20	148
	GBP/USD	4,395	1	1	1	1	2	3	7	10	1,011
	NZD/USD	784	1	1	1	1	1	3	11	22	179
	USD/CAD	1,153	1	1	1	1	1	4	11	26	313
	USD/CHF	2,399	1	1	1	1	1	3	5	9	475
	USD/JPY	2,494	1	1	1	1	1	3	8	13	604
	USD/MXN	1,106	1	1	1	1	2	3	5	11	264
2013	AUD/JPY	326	1	1	1	1	2	8	17	32	111
	AUD/USD	1,336	1	1	1	1	2	5	19	55	367
	EUR/CHF	829	1	1	1	1	2	7	22	40	255
	EUR/GBP	1,414	1	1	1	1	2	4	12	21	369
	EUR/JPY	751	1	1	1	1	3	10	24	44	266
	EUR/USD	3,515	1	1	1	1	2	5	10	24	861
	GBP/JPY	455	1	1	1	1	1	5	11	21	122
	GBP/USD	2,713	1	1	1	1	2	4	9	17	607
	NZD/USD	542	1	1	1	1	3	9	34	52	195
	USD/CAD	747	1	1	1	1	2	8	31	53	238
	USD/CHF	1,346	1	1	1	1	2	3	11	24	285
	USD/JPY	1,884	1	1	1	1	2	5	16	47	552
	USD/MXN	877	1	1	1	1	2	4	11	29	196

Table 4. Distribution of shortest paths

Each pair of nodes in the settlement network we compute the length of the shortest path between them. (For example, if member A and member B both settle against member C, the length of the path is two. Settlements between two members are of length one. A path with length zero is included as a formalism to connect a member (node) with itself. The table reports the distribution of these shortest paths.

2016		2013	
Length of shortest path	Percent	Length of shortest path	Percent
0	0.01	0	0.01
1	0.06	1	0.10
2	54.01	2	52.81
3	45.81	3	41.74
4	0.12	4	5.21
5	0.00	5	0.12

Table 5. Core-connected members

A member is considered core-connected if it is a (one-step) neighbor of one or more of the top twenty-five (ranked by volume) settlement members.

	2016	2013
All members (nodes)	16,163	11,651
Neighbors of top 25 nodes	15,213	10,545
Not neighbors of top 25 nodes	950	1,106
	(5.88%)	(9.49%)

Table 6. Degree centrality groups: descriptive statistics

For each April (year)/currency pair, spot settlement members are ranked by descending degree centrality. Four groups are constructed to approximate balanced volume shares (20% in each group), subject to the constraint that each group contains no fewer than five members. Table entries report the number of members, average degree centrality, and volume share (USD equivalent) in each group.

Panel A. Groups formed on unweighted degree centrality.

		No. settlement members					Average unweighted degree centrality					Volume share (percent)					
		Group					Group					Group					
		All	1	2	3	4	5	1	2	3	4	5	1	2	3	4	5
2016	AUD/JPY	239	217	7	5	5	5	4.2	36.3	51.4	65.0	83.8	6.4	6.7	8.7	21.5	56.7
	AUD/USD	2,462	2,436	11	5	5	5	3.0	139.5	274.8	378.0	493.0	10.1	10.5	17.6	16.8	45.0
	EUR/CHF	803	765	22	6	5	5	3.1	39.5	99.3	158.0	193.2	11.3	11.4	11.5	20.8	45.0
	EUR/GBP	1,533	1,495	22	6	5	5	3.0	59.8	168.7	267.2	311.6	11.7	11.8	12.2	25.6	38.8
	EUR/JPY	683	655	12	6	5	5	4.0	53.8	108.3	149.0	176.8	8.4	8.8	10.4	25.4	47.0
	EUR/USD	5,699	5,657	25	7	5	5	3.3	140.5	526.0	1054	1442	11.2	11.3	12.7	20.8	44.0
	GBP/JPY	511	487	9	5	5	5	2.9	31.3	56.8	83.4	148.4	6.3	6.3	11.5	19.2	56.6
	GBP/USD	4,395	4,360	20	5	5	5	3.0	129.9	442.8	721.8	1011	11.2	11.3	11.4	26.7	39.4
	NZD/USD	784	761	8	5	5	5	3.3	64.3	100.6	133.4	179.0	9.1	9.6	12.3	24.6	44.4
	USD/CAD	1,153	1,127	11	5	5	5	4.0	84.2	148.8	214.6	313.2	14.4	16.1	15.5	18.4	35.6
	USD/CHF	2,399	2,356	27	6	5	5	2.4	52.1	185.8	379.4	474.6	9.2	10.9	11.3	20.2	48.4
	USD/JPY	2,494	2,463	14	7	5	5	3.1	90.6	235.1	389.8	603.8	6.9	11.1	11.5	20.8	49.8
	USD/MXN	1,106	1,082	9	5	5	5	2.7	52.4	101.4	171.6	263.8	7.3	7.8	12.7	24.5	47.7
2013	AUD/JPY	326	291	17	8	5	5	4.1	20.5	43.5	72.0	110.8	9.7	9.8	10.7	17.3	52.5
	AUD/USD	1,336	1,293	25	5	8	5	5.7	97.4	154.8	228.8	366.8	13.6	13.7	16.3	22.0	34.4
	EUR/CHF	829	791	20	8	5	5	5.4	60.6	125.9	185.6	255.0	14.3	15.1	17.8	20.1	32.7
	EUR/GBP	1,414	1,352	40	11	6	5	3.3	39.7	117.5	224.2	369.2	11.7	12.0	18.0	20.4	37.9
	EUR/JPY	751	704	32	5	5	5	5.8	56.3	146.2	188.2	266.4	10.9	11.0	12.2	24.6	41.2
	EUR/USD	3,515	3,461	35	6	8	5	4.8	158.1	317.0	493.0	860.8	15.4	15.7	15.7	17.2	36.0
	GBP/JPY	455	419	18	5	8	5	2.9	18.6	40.2	64.5	122.4	8.8	8.9	11.5	12.6	58.2
	GBP/USD	2,713	2,661	31	8	8	5	3.5	91.4	208.5	331.8	607.0	12.8	12.9	14.2	20.5	39.6
	NZD/USD	542	504	22	6	5	5	6.2	57.3	104.5	143.6	194.6	14.9	14.9	19.2	17.7	33.3
	USD/CAD	747	701	26	7	8	5	5.5	64.7	106.4	155.5	238.2	15.7	15.9	18.9	18.4	31.1
	USD/CHF	1,346	1,311	19	6	5	5	3.7	75.8	147.0	193.6	284.6	13.4	14.7	14.3	27.1	30.5
	USD/JPY	1,884	1,846	23	5	5	5	6.2	137.3	230.0	299.8	551.6	14.3	14.4	14.6	19.6	37.1
	USD/MXN	877	839	21	7	5	5	3.4	51.5	104.6	139.8	195.8	13.2	13.5	13.9	15.5	44.0

Panel B. Groups formed on volume-weighted degree centrality

		No. settlement members					Average unweighted degree centrality					Volume share (percent)					
		Group					Group					Group					
		All	1	2	3	4	5	1	2	3	4	5	1	2	3	4	5
2016	AUD/JPY	239	218	6	5	5	5	4.4	40.5	45.2	65.8	82.8	5.8	6.6	7.8	14.9	64.9
	AUD/USD	2,462	2,437	10	5	5	5	3.2	118.0	311.2	369.8	450.4	9.2	10.3	10.6	19.5	50.4
	EUR/CHF	803	773	14	6	5	5	3.5	43.4	101.8	144.2	182.2	10.3	10.3	10.9	16.1	52.4
	EUR/GBP	1,533	1,497	19	7	5	5	3.1	54.7	155.9	281.8	282.4	10.6	11.0	11.1	19.9	47.4
	EUR/JPY	683	657	11	5	5	5	4.1	61.3	105.0	150.0	168.4	8.5	9.0	9.3	16.6	56.6
	EUR/USD	5,699	5,662	21	6	5	5	3.6	135.5	592.2	971.0	1390	11.1	11.2	11.2	19.6	46.9
	GBP/JPY	511	486	10	5	5	5	3.0	32.3	52.0	81.0	142.6	5.4	5.8	8.5	17.4	62.9
	GBP/USD	4,395	4,359	21	5	5	5	3.0	128.1	460.0	758.4	901.4	9.9	10.4	10.3	21.9	47.4
	NZD/USD	784	762	7	5	5	5	3.4	68.4	107.0	129.2	172.4	8.7	9.2	10.9	16.9	54.4
	USD/CAD	1,153	1,126	11	6	5	5	4.1	87.2	122.2	223.0	283.0	12.4	13.8	13.9	18.3	41.5
	USD/CHF	2,399	2,366	17	6	5	5	2.6	76.8	160.0	348.4	440.8	9.4	9.7	10.1	18.8	51.9
	USD/JPY	2,494	2,472	7	5	5	5	3.5	156.4	209.4	459.0	504.6	8.7	8.9	10.2	21.1	51.2
	USD/MXN	1,106	1,084	7	5	5	5	2.8	78.0	78.6	156.2	253.2	6.9	6.9	9.8	18.5	58.0
2013	AUD/JPY	326	293	15	8	5	5	4.3	24.7	38.4	63.2	108.0	7.8	8.3	8.8	18.0	57.1
	AUD/USD	1,336	1,299	19	8	5	5	6.4	81.9	205.9	205.8	342.4	12.6	12.6	14.5	21.5	38.7
	EUR/CHF	829	794	18	7	5	5	5.8	63.1	141.0	189.2	203.8	12.8	13.1	13.6	17.7	42.8
	EUR/GBP	1,414	1,372	23	9	5	5	3.9	61.0	120.1	258.2	286.8	11.1	11.2	11.8	21.7	44.1
	EUR/JPY	751	712	20	9	5	5	6.6	55.3	95.8	191.2	248.4	8.7	9.0	9.1	20.5	52.7
	EUR/USD	3,515	3,469	27	9	5	5	5.6	128.4	311.8	539.2	746.6	12.7	13.1	14.0	20.5	39.7
	GBP/JPY	455	425	13	7	5	5	3.2	21.2	51.3	51.2	122.4	7.2	7.5	8.4	18.7	58.2
	GBP/USD	2,713	2,676	20	7	5	5	4.2	108.2	235.9	461.6	439.2	11.8	12.1	12.7	22.6	40.8
	NZD/USD	542	508	17	7	5	5	7.0	51.6	114.6	120.0	171.6	12.7	13.5	13.2	19.1	41.6
	USD/CAD	747	714	16	7	5	5	6.8	85.4	100.0	159.4	204.2	15.1	15.3	17.0	21.5	31.2
	USD/CHF	1,346	1,313	15	8	5	5	4.1	68.6	129.1	217.2	218.4	10.4	11.2	11.8	21.6	44.9
	USD/JPY	1,884	1,847	19	8	5	5	6.6	113.1	246.3	319.2	413.6	10.5	11.1	11.9	21.3	45.2
	USD/MXN	877	852	9	6	5	5	4.3	56.0	94.2	132.6	183.0	10.6	10.7	13.0	20.4	45.3

Table 7. Settlement flows between centrality groups

For each April (year)/currency pair we group settlement members according to (unweighted) degree centrality (5-high). Panel A reports total settlement volume (USD equivalent) between groups; buyer and seller directions refer to the base currency; M is million; B is billion. Panel B reports these volumes as percentages of the annual total.

Panel A. Settlement volume (USD equivalent)

		Buyer Centrality Group	Seller Centrality Group				
			1	2	3	4	5
2016	1	67,864M	98,680M	122,383M	213,179M	366,376M	
	2	103,118M	106,141M	122,364M	225,722M	452,082M	
	3	125,114M	117,775M	105,275M	214,665M	547,851M	
	4	215,179M	203,730M	214,044M	299,526M	922,582M	
	5	366,797M	451,591M	556,050M	934,071M	1,611B	
2013	1	184,198M	233,558M	254,670M	368,804M	719,470M	
	2	254,094M	194,079M	259,752M	374,965M	772,372M	
	3	281,069M	250,127M	240,717M	411,113M	812,779M	
	4	389,637M	360,539M	395,623M	353,344M	984,801M	
	5	751,006M	776,700M	793,693M	986,991M	1,395B	

Panel B. Percent of settlement volume (USD equivalent)

		Buyer Centrality Group	Seller Centrality Group				
			1	2	3	4	5
2016	1	0.8	1.1	1.4	2.4	4.2	
	2	1.2	1.2	1.4	2.6	5.2	
	3	1.4	1.3	1.2	2.4	6.3	
	4	2.5	2.3	2.4	3.4	10.5	
	5	4.2	5.2	6.3	10.7	18.4	
2013	1	1.4	1.8	2.0	2.9	5.6	
	2	2.0	1.5	2.0	2.9	6.0	
	3	2.2	2.0	1.9	3.2	6.4	
	4	3.0	2.8	3.1	2.8	7.7	
	5	5.9	6.1	6.2	7.7	10.9	

Table 8. Centrality differentials, pooled estimates

CLS spot settlements in the Aprils of 2013 and 2016, restricted to thirteen Olsen currency pairs. $\kappa \in \{AUD/JPY, \dots, USD/MXN\}$ indexes currency pairs; $t = 1, \dots, T_\kappa$ indexes intervals (of approximately ten seconds) defined by Olsen bid and ask quote records; $k = 1, \dots, N_{\kappa t}$ indexes settlements within each interval. The base currency buyer's return is $\pi_{kt\kappa}^{Post} = m_{t+1,\kappa} - p_{kt\kappa}$ where $p_{kt\kappa}$ is the settlement price (exchange rate) and $m_{t+1,\kappa}$ is the midpoint of the Olsen bid and ask at the start of the interval following the settlement. Settlement members are placed in centrality groups labeled from 1 (low) to 5 (high). The centrality differential is $\mu_{ij} = E[\pi_{kt\kappa}^{Post} | buy \in group i, seller \in group j]$ for $i, j = 1, \dots, 5$. The specification is

$$\pi_{kt\kappa}^{Post} = \alpha_\kappa + \mu_{ij} + \epsilon_{kt\kappa}$$

where α_κ is a pair-specific intercept. The error specification is

$$Cov(\epsilon_{kt\kappa}, \epsilon_{k^*t\kappa}) = \begin{cases} \sigma_{f\kappa}^2 + \sigma_{g\kappa}^2 & \text{if } k = k^* \\ \sigma_{f\kappa}^2 & \text{if } k \neq k^* \end{cases}$$

Panel A reports the μ_{ij} estimates; Panel B, the α_κ , $\sigma_{f\kappa}^2$, and $\sigma_{g\kappa}^2$ estimates. The μ_{ij} and α_κ are scaled to basis points; t-values are reported in parentheses. Panel C reports average centrality differentials (and associated t-values): $\bar{\mu}_{ij} = (\mu_{ij} - \mu_{ji})/2$ for $i > j$.

Panel A. Centrality differentials, μ_{ij}

Buyer centrality group	2016					2013				
	Seller centrality group					Seller centrality group				
	1	2	3	4	5	1	2	3	4	5
1		-0.106 (-12.09)	-0.099 (-14.38)	-0.141 (-27.69)	-0.125 (-30.56)		-0.034 (-8.10)	-0.031 (-8.20)	-0.043 (-13.64)	-0.072 (-29.56)
2	0.058 (6.48)		-0.108 (-17.78)	-0.084 (-18.54)	-0.113 (-36.60)	0.037 (9.11)		-0.035 (-10.13)	-0.043 (-14.64)	-0.030 (-13.62)
3	0.104 (15.27)	0.065 (10.75)		-0.014 (-3.38)	-0.046 (-17.04)	0.049 (13.16)	0.062 (17.68)		0.023 (8.32)	-0.002 (-1.18)
4	0.129 (25.30)	0.059 (12.94)	-0.013 (-3.20)		-0.039 (-18.18)	0.083 (26.52)	0.065 (21.85)	-0.032 (-11.60)		-0.003 (-1.74)
5	0.126 (31.17)	0.078 (25.31)	0.014 (5.30)	0.026 (12.43)		0.102 (42.37)	0.047 (21.67)	0.010 (4.71)	0.010 (4.96)	

Panel B. Currency pair fixed effects and variance parameters.

	2016				2013			
	Pair effects		Variance parameters		Pair effects		Variance parameters	
	Estimate	t-Value	$\sigma_{f\kappa}^2$	$\sigma_{g\kappa}^2$	Estimate	t-Value	$\sigma_{f\kappa}^2$	$\sigma_{g\kappa}^2$
AUD/JPY	-0.041	(-2.77)	11.427	7.005	0.068	(6.98)	6.552	4.276
AUD/USD	0.032	(3.80)	8.772	3.753	0.033	(6.77)	3.108	1.363
EUR/CHF	0.018	(2.33)	2.875	0.615	0.026	(4.20)	2.451	0.400
EUR/GBP	-0.001	(-0.14)	6.378	2.753	0.006	(0.80)	3.569	2.424
EUR/JPY	0.017	(1.95)	4.556	3.071	0.046	(7.19)	4.406	4.355
EUR/USD	0.023	(4.28)	3.933	2.053	0.002	(0.39)	2.217	1.815
GBP/JPY	0.026	(2.38)	4.240	3.748	0.095	(9.86)	3.890	3.777
GBP/USD	0.018	(3.02)	3.745	2.141	-0.002	(-0.53)	1.860	1.320
NZD/USD	0.047	(4.20)	10.535	4.639	0.040	(4.80)	4.727	2.940
USD/CAD	0.016	(2.39)	4.778	2.221	0.004	(0.70)	3.135	0.834
USD/CHF	0.019	(2.16)	5.879	1.381	0.015	(2.07)	3.496	2.480
USD/JPY	0.030	(5.31)	4.443	3.041	0.053	(9.98)	4.293	3.730
USD/MXN	0.007	(1.51)	0.664	0.563	-0.001	(-0.18)	0.906	0.798

Panel C. Average centrality differentials, $\bar{\mu}_{ij}$ for $i > j$

Higher centrality group	2016				2013			
	Lower centrality group				Lower centrality group			
	1	2	3	4	1	2	3	4
2	0.082				0.035			
	(13.34)				(12.67)			
3	0.102	0.087			0.040	0.048		
	(21.53)	(20.71)			(15.83)	(20.84)		
4	0.135	0.072	0.000		0.063	0.054	-0.027	
	(39.12)	(23.25)	(0.10)		(30.45)	(27.99)	(-15.46)	
5	0.125	0.096	0.030	0.033	0.087	0.039	0.006	0.007
	(46.39)	(48.07)	(17.97)	(26.60)	(57.20)	(28.83)	(4.94)	(5.70)

Table 9. Centrality differentials and settlement size

The sample is CLS spot settlements in the Aprils of 2013 and 2016, restricted to thirteen Olsen currency pairs. Then $\kappa \in \{AUD/JPY, \dots, USD/MXN\}$ indexes currency pairs; $t = 1, \dots, T_\kappa$ indexes intervals (of approximately ten seconds) defined by Olsen bid and ask quote records; $k = 1, \dots, N_{\kappa t}$ indexes settlements within each interval. The base currency buyer's return is $\pi_{kt\kappa}^{Post} = m_{t+1,\kappa} - p_{kt\kappa}$ where $p_{kt\kappa}$ is the settlement price (exchange rate) and $m_{t+1,\kappa}$ is the midpoint of the Olsen bid and ask at the start of the interval following the settlement. Settlement members are placed in centrality groups labeled from 1 (low) to 5 (high). The centrality differential is $\mu_{ij} = E[\pi_{kt\kappa}^{Post} | buy \in group i, seller \in group j]$ for $i, j = 1, \dots, 5$. The specification is

$$\pi_{kt\kappa}^{Post} = \alpha_\kappa + \mu_{ij} \otimes Size + \epsilon_{kt\kappa}$$

where $Size$ is a classifier, $Size \in \{< 1M, \geq 1M\}$ in units of the base currency; and, α_κ is a pair-specific intercept. The error specification is

$$Cov(\epsilon_{kt\kappa}, \epsilon_{k^*t\kappa}) = \begin{cases} \sigma_{f\kappa}^2 + \sigma_{g\kappa}^2 & \text{if } k = k^* \\ \sigma_{f\kappa}^2 & \text{if } k \neq k^* \end{cases}$$

Panel A reports estimates of average centrality differential (and associated t-values): $\bar{\mu}_{ij} = (\mu_{ij} - \mu_{ji})/2$ for $i > j$ (in basis points). Panel B contains the $\alpha_\kappa, \sigma_{f\kappa}^2$, and $\sigma_{g\kappa}^2$ estimates.

Panel A. $\bar{\mu}_{ij}$ estimates

Settlement size	Higher centrality group	2016				2013			
		Lower centrality group				Lower centrality group			
		1	2	3	4	1	2	3	4
<i>Small:</i> <i>< 1M units</i> <i>of the base</i> <i>currency</i>	2	0.151 (16.41)				0.064 (15.02)			
	3	0.148 (20.20)	0.125 (16.32)			0.056 (14.09)	0.044 (13.24)		
	4	0.144 (28.54)	0.097 (19.41)	0.017 (4.35)		0.075 (24.87)	0.041 (14.38)	-0.019 (-7.53)	
	5	0.151 (37.11)	0.117 (38.03)	0.034 (14.33)	0.013 (8.05)	0.085 (39.43)	0.009 (4.48)	0.006 (3.32)	-0.010 (-5.94)
<i>Large:</i> <i>≥ 1M units</i> <i>of the base</i> <i>currency</i>	2	0.026 (3.18)				0.014 (3.85)			
	3	0.069 (11.24)	0.070 (14.09)			0.030 (9.02)	0.053 (16.53)		
	4	0.126 (27.04)	0.056 (14.37)	-0.016 (-4.20)		0.053 (18.65)	0.065 (25.25)	-0.034 (-14.29)	
	5	0.105 (29.52)	0.081 (31.48)	0.026 (11.32)	0.058 (31.63)	0.090 (41.79)	0.068 (36.35)	0.007 (3.77)	0.022 (13.98)
<i>Difference:</i> <i>Large minus</i> <i>small</i>	2	-0.124 (-10.10)				-0.050 (-8.89)			
	3	-0.079 (-8.33)	-0.055 (-6.04)			-0.027 (-5.17)	0.008 (1.84)		
	4	-0.018 (-2.62)	-0.041 (-6.42)	-0.033 (-6.07)		-0.022 (-5.30)	0.025 (6.43)	-0.015 (-4.27)	
	5	-0.045 (-8.43)	-0.037 (-9.23)	-0.007 (-2.23)	0.045 (18.69)	0.005 (1.67)	0.059 (22.40)	0.001 (0.42)	0.032 (14.13)

Panel B. Estimates of α_{κ} , $\sigma_{f\kappa}^2$, and $\sigma_{g\kappa}^2$

	2016				2013			
	Pair effects, α_{κ}		Variance parameters		Pair effects		Variance parameters	
	Estimate	t-Value	$\sigma_{f\kappa}^2$	$\sigma_{g\kappa}^2$	Estimate	t-Value	$\sigma_{f\kappa}^2$	$\sigma_{g\kappa}^2$
AUD/JPY	-0.036	(-2.42)	11.426	7.005	0.064	(6.52)	6.550	4.277
AUD/USD	0.035	(4.21)	8.771	3.753	0.031	(6.23)	3.108	1.363
EUR/CHF	0.023	(2.93)	2.874	0.614	0.024	(3.87)	2.453	0.398
EUR/GBP	0.004	(0.36)	6.379	2.752	0.004	(0.54)	3.570	2.423
EUR/JPY	0.021	(2.46)	4.553	3.073	0.042	(6.52)	4.406	4.353
EUR/USD	0.027	(4.97)	3.933	2.052	-0.001	(-0.21)	2.216	1.815
GBP/JPY	0.032	(2.91)	4.237	3.750	0.092	(9.37)	3.890	3.777
GBP/USD	0.023	(3.76)	3.745	2.141	-0.005	(-1.11)	1.859	1.320
NZD/USD	0.051	(4.55)	10.534	4.638	0.037	(4.43)	4.728	2.940
USD/CAD	0.020	(2.95)	4.778	2.221	0.002	(0.36)	3.136	0.833
USD/CHF	0.023	(2.63)	5.877	1.381	0.012	(1.71)	3.497	2.479
USD/JPY	0.034	(5.86)	4.442	3.042	0.050	(9.29)	4.292	3.731
USD/MXN	0.010	(2.18)	0.663	0.564	-0.002	(-0.38)	0.905	0.798

Table 10. Average centrality differentials, ex ante benchmark

CLS spot settlements in the Aprils of 2013 and 2016, restricted to thirteen Olsen currency pairs. $\kappa \in \{AUD/JPY, \dots, USD/MXN\}$ indexes currency pairs; $t = 1, \dots, T_\kappa$ indexes intervals (of approximately ten seconds) defined by Olsen bid and ask quote records; $k = 1, \dots, N_{kt}$ indexes settlements within each interval. The base currency buyer's return is $\pi_{kt\kappa}^{Pre} = m_{t\kappa} - p_{kt\kappa}$ where $p_{kt\kappa}$ is the settlement price (exchange rate) and $m_{t\kappa}$ is the midpoint of the Olsen bid and ask at the start of interval k . Settlement members are placed in centrality groups labeled from 1 (low) to 5 (high). The centrality differential is $\mu_{ij} = E[\pi_{kt\kappa}^{Pre} | \text{buy} \in \text{group } i, \text{seller} \in \text{group } j]$ for $i, j = 1, \dots, 5$. The specification is

$$\pi_{kt\kappa}^{Pre} = \alpha_\kappa + \mu_{ij} + \epsilon_{kt\kappa}$$

where α_κ is a pair-specific intercept. The error specification is

$$\text{Cov}(\epsilon_{kt\kappa}, \epsilon_{k^*t\kappa}) = \begin{cases} \sigma_{f\kappa}^2 + \sigma_{g\kappa}^2 & \text{if } k = k^* \\ \sigma_{f\kappa}^2 & \text{if } k \neq k^* \end{cases}$$

Panel A reports estimates of centrality advantage (and associated t-values): $\bar{\mu}_{ij} = (\mu_{ij} - \mu_{ji})/2$ for $i > j$ (in basis points). Panel B contains the α_κ , $\sigma_{f\kappa}^2$, and $\sigma_{g\kappa}^2$ estimates.

Panel A. Estimates of $\bar{\mu}_{ij}$ for $i > j$

Higher centrality group	2016				2013			
	Lower centrality group				Lower centrality group			
	1	2	3	4	1	2	3	4
2	0.081 (12.99)				0.033 (11.16)			
3	0.102 (21.29)	0.088 (20.73)			0.042 (15.77)	0.048 (19.62)		
4	0.137 (39.29)	0.073 (23.39)	-0.001 (-0.45)		0.068 (31.51)	0.056 (27.66)	-0.026 (-14.36)	
5	0.128 (46.55)	0.098 (48.19)	0.029 (17.43)	0.033 (26.44)	0.089 (55.90)	0.038 (27.07)	0.005 (3.94)	0.005 (4.19)

Panel B. Estimates of α_κ , $\sigma_{f\kappa}^2$, and $\sigma_{g\kappa}^2$

	2016				2013			
	Pair effects		Variance parameters		Pair effects		Variance parameters	
	Estimate	t-Value	$\sigma_{f\kappa}^2$	$\sigma_{g\kappa}^2$	Estimate	t-Value	$\sigma_{f\kappa}^2$	$\sigma_{g\kappa}^2$
AUD/JPY	-0.031	(-2.12)	10.254	7.197	0.060	(6.52)	5.593	4.552
AUD/USD	0.031	(3.90)	7.661	3.920	0.030	(6.53)	2.661	1.520
EUR/CHF	0.019	(2.61)	2.404	0.863	0.026	(4.41)	2.196	0.540
EUR/GBP	0.003	(0.28)	5.765	2.724	0.005	(0.67)	3.070	2.580
EUR/JPY	0.019	(2.31)	3.966	3.135	0.043	(7.27)	3.569	4.494
EUR/USD	0.023	(4.57)	3.389	2.144	-0.001	(-0.34)	1.681	2.031
GBP/JPY	0.024	(2.34)	3.839	3.624	0.076	(8.24)	3.308	3.851
GBP/USD	0.016	(2.85)	3.259	2.164	-0.004	(-0.89)	1.558	1.404
NZD/USD	0.050	(4.72)	9.137	5.053	0.034	(4.27)	4.215	3.076
USD/CAD	0.018	(2.78)	4.096	2.300	0.003	(0.65)	2.841	0.919
USD/CHF	0.016	(1.87)	5.339	1.450	0.017	(2.52)	2.953	2.725
USD/JPY	0.031	(5.84)	3.906	3.086	0.049	(10.31)	3.361	4.084
USD/MXN	0.001	(0.30)	0.548	0.609	-0.002	(-0.37)	1.191	0.634

Table 11. Average centrality differentials, volume-weighted centrality grouping

The sample is CLS spot settlements in the Aprils of 2013 and 2016, restricted to thirteen Olsen currency pairs. Then $\kappa \in \{AUD/JPY, \dots, USD/MXN\}$ indexes currency pairs; $t = 1, \dots, T_\kappa$ indexes intervals (of approximately ten seconds) defined by Olsen bid and ask quote records; $k = 1, \dots, N_{\kappa t}$ indexes settlements within each interval. The base currency buyer's return is $\pi_{kt\kappa}^{Post} = m_{t+1,\kappa} - p_{kt\kappa}$ where $p_{kt\kappa}$ is the settlement price (exchange rate) and $m_{t+1,\kappa}$ is the midpoint of the Olsen bid and ask at the start of the interval following the settlement. Settlement members are grouped based on their volume-weighted degree centrality. The centrality differential is $\mu_{ij} = E[\pi_{kt\kappa}^{Post} | buy \in group i, seller \in group j]$ for $i, j = 1, \dots, 5$. The specification is

$$\pi_{kt\kappa}^{Post} = \alpha_\kappa + \mu_{ij} + \epsilon_{kt\kappa}$$

where α_κ is a pair-specific intercept. The error specification is

$$Cov(\epsilon_{kt\kappa}, \epsilon_{k^*t\kappa}) = \begin{cases} \sigma_{f\kappa}^2 + \sigma_{g\kappa}^2 & \text{if } k = k^* \\ \sigma_{f\kappa}^2 & \text{if } k \neq k^* \end{cases}$$

Panel A reports average centrality differentials (and associated t-values): $\bar{\mu}_{ij} = (\mu_{ij} - \mu_{ji})/2$ for $i > j$ (in basis points). Panel B contains the α_κ , $\sigma_{f\kappa}^2$, and $\sigma_{g\kappa}^2$ estimates.

Panel A. Estimates of $\bar{\mu}_{ij}$ for $i > j$

Higher centrality group	2016				2013			
	Lower centrality group				Lower centrality group			
	1	2	3	4	1	2	3	4
2	0.043 (7.29)				0.124 (31.67)			
3	0.093 (18.49)	0.051 (10.19)			0.078 (23.36)	0.014 (5.46)		
4	0.106 (31.14)	0.090 (24.62)	0.043 (12.58)		0.110 (44.71)	0.049 (23.96)	0.011 (5.16)	
5	0.112 (45.01)	0.089 (44.80)	0.102 (58.43)	0.093 (76.30)	0.123 (70.64)	0.083 (54.80)	0.051 (38.87)	0.013 (13.73)

Panel B. Estimates of α_κ , $\sigma_{f\kappa}^2$, and $\sigma_{g\kappa}^2$

	2016				2013			
	Pair effects		Variance parameters		Pair effects		Variance parameters	
	Estimate	t-Value	$\sigma_{f\kappa}^2$	$\sigma_{g\kappa}^2$	Estimate	t-Value	$\sigma_{f\kappa}^2$	$\sigma_{g\kappa}^2$
AUD/JPY	-0.046	(-3.06)	11.421	7.010	0.068	(7.06)	6.551	4.276
AUD/USD	0.026	(3.19)	8.772	3.752	0.034	(6.87)	3.113	1.357
EUR/CHF	0.013	(1.62)	2.880	0.612	0.027	(4.35)	2.444	0.401
EUR/GBP	-0.006	(-0.61)	6.385	2.749	0.006	(0.81)	3.570	2.420
EUR/JPY	0.013	(1.52)	4.561	3.069	0.047	(7.38)	4.404	4.357
EUR/USD	0.019	(3.44)	3.935	2.050	0.002	(0.56)	2.217	1.815
GBP/JPY	0.022	(1.99)	4.239	3.750	0.096	(9.90)	3.888	3.773
GBP/USD	0.014	(2.34)	3.745	2.142	-0.002	(-0.47)	1.862	1.317
NZD/USD	0.042	(3.74)	10.530	4.641	0.039	(4.76)	4.733	2.931
USD/CAD	0.010	(1.56)	4.780	2.220	0.004	(0.76)	3.136	0.831
USD/CHF	0.014	(1.64)	5.880	1.379	0.015	(2.14)	3.495	2.479
USD/JPY	0.025	(4.50)	4.444	3.040	0.053	(10.13)	4.292	3.730
USD/MXN	0.002	(0.42)	0.662	0.565	-0.002	(-0.28)	0.906	0.795

Table 12. Settlement volume between centrality groups, by currency pair

The sample is CLS spot settlements in the Aprils of 2013 and 2016, restricted to thirteen Olsen currency pairs. Settlement members are placed in centrality groups labeled from 1 (low) to 5 (high) based on their unweighted degree centrality. Table entries are intergroup settlement volumes as a percentage of the total. (Within each 5 × 5 block the values sum to one hundred.)

	High centrality group	2016					2013				
		Low centrality group					Low centrality group				
		1	2	3	4	5	1	2	3	4	5
AUD/JPY	1	0.3					1.1				
	2	0.9	0.4				1.3	0.5			
	3	1.5	0.9	0.4			2.9	2.5	1.0		
	4	3.2	2.6	3.0	3.4		2.7	3.8	5.0	1.0	
	5	6.6	8.2	11.2	27.4	30.0	10.4	11.0	9.0	21.1	26.8
AUD/USD	1	0.7					1.4				
	2	1.9	1.2				3.3	1.1			
	3	3.1	3.8	2.4			3.9	3.5	2.5		
	4	3.4	3.4	6.0	2.2		6.5	6.5	7.3	3.8	
	5	10.5	9.5	17.5	16.4	18.0	10.6	11.8	12.9	16.2	8.7
EUR/CHF	1	0.8					1.4				
	2	3.0	0.7				3.3	1.7			
	3	2.4	2.3	1.0			6.3	5.6	2.6		
	4	7.5	6.3	3.7	2.7		6.3	7.2	4.7	2.9	
	5	7.9	9.8	12.6	18.6	20.5	10.0	10.8	13.9	16.2	7.2
EUR/GBP	1	0.5					0.5				
	2	1.9	0.6				2.1	0.9			
	3	4.0	3.0	0.8			4.7	4.9	1.7		
	4	7.7	7.8	5.3	4.7		6.5	4.3	6.9	3.3	
	5	8.7	9.6	10.5	20.9	13.9	9.0	10.9	16.2	16.4	11.7
EUR/JPY	1	0.9					0.8				
	2	1.2	0.7				2.1	0.6			
	3	2.6	2.3	0.4			2.9	2.8	1.3		
	4	4.1	3.8	4.0	5.4		5.2	5.1	4.8	3.5	
	5	7.1	8.9	11.2	28.1	19.3	10.0	10.8	11.3	27.2	11.6
EUR/USD	1	0.8					1.6				
	2	2.4	1.2				4.7	1.8			
	3	3.6	2.4	1.1			3.9	4.7	2.0		
	4	5.8	5.2	4.6	2.7		6.1	5.2	6.5	1.6	
	5	9.1	10.2	12.5	20.7	17.8	13.0	13.2	12.3	13.5	9.9

	High centrality group	2016					2013				
		Low centrality group					Low centrality group				
		1	2	3	4	5	1	2	3	4	5
GBP/JPY	1	0.4					0.7				
	2	0.7	0.1				0.9	0.2			
	3	1.3	1.6	1.9			1.2	2.2	2.0		
	4	2.7	3.2	2.4	2.2		3.9	2.1	4.2	1.5	
	5	7.1	7.0	14.1	25.9	29.6	10.2	12.2	11.4	11.9	35.4
GBP/USD	1	0.9					1.2				
	2	2.9	1.0				3.1	1.0			
	3	2.2	2.1	1.0			3.4	4.4	1.3		
	4	6.5	6.5	5.7	5.9		5.4	4.2	6.0	3.9	
	5	8.9	9.1	10.8	23.0	13.5	11.5	12.1	11.9	17.8	12.9
NZD/USD	1	0.6					1.9				
	2	1.4	1.0				3.9	2.0			
	3	2.1	2.2	1.3			5.3	3.8	3.7		
	4	4.1	5.6	5.8	5.3		6.5	6.5	6.7	2.1	
	5	9.4	8.1	11.7	23.0	18.3	10.3	11.7	15.2	11.4	9.0
USD/CAD	1	2.0					1.5				
	2	6.1	2.7				3.7	2.2			
	3	3.7	4.1	1.5			9.4	5.0	3.5		
	4	6.2	6.0	6.5	2.5		5.3	6.2	6.2	3.2	
	5	8.6	10.6	13.7	13.1	12.6	10.2	12.5	10.2	12.7	8.2
USD/CHF	1	0.6					1.4				
	2	2.1	1.0				4.0	1.1			
	3	2.2	2.5	0.8			4.5	2.8	1.1		
	4	4.0	6.1	3.5	1.9		7.4	10.5	9.0	4.5	
	5	8.7	9.1	12.8	23.1	21.5	8.2	9.8	10.0	18.2	7.4
USD/JPY	1	0.4					1.7				
	2	1.1	1.1				4.1	1.6			
	3	2.0	2.9	1.1			3.4	3.2	1.2		
	4	3.1	3.7	4.2	3.5		6.1	6.6	6.4	3.0	
	5	6.8	12.2	11.8	23.6	22.6	11.7	11.5	13.7	14.1	11.6
USD/MXN	1	0.3					1.2				
	2	0.9	0.3				3.7	1.8			
	3	1.7	1.7	0.9			3.2	2.8	1.9		
	4	4.6	3.6	6.0	3.7		4.5	3.3	3.9	2.2	
	5	6.8	9.1	14.3	27.4	18.9	12.7	13.6	14.1	14.9	16.3

Table 13. Average centrality differentials by currency pair.

The sample is CLS spot settlements in the Aprils of 2013 and 2016, restricted to thirteen Olsen currency pairs. Then $\kappa \in \{AUD/JPY, \dots, USD/MXN\}$ indexes currency pairs; $t = 1, \dots, T_\kappa$ indexes intervals (of approximately ten seconds) defined by Olsen bid and ask quote records; $k = 1, \dots, N_{\kappa t}$ indexes settlements within each interval. The base currency buyer's return is $\pi_{kt\kappa}^{Post} = m_{t+1,\kappa} - p_{kt\kappa}$ where $p_{kt\kappa}$ is the settlement price (exchange rate) and $m_{t+1,\kappa}$ is the midpoint of the Olsen bid and ask at the start of the interval following the settlement. The centrality differential is $\mu_{ij} = E[\pi_{kt\kappa}^{Post} | buy \in group i, seller \in group j]$ for $i, j = 1, \dots, 5$. The specification is

$$\pi_{kt}^{Post} = \alpha + \mu_{ij} + \epsilon_{kt}$$

The error specification is

$$Cov(\epsilon_{kt}, \epsilon_{k^*t}) = \begin{cases} \sigma_f^2 + \sigma_g^2 & \text{if } k = k^* \\ \sigma_f^2 & \text{if } k \neq k^* \end{cases}$$

The specification is estimated separately for each (April) year and currency pair. Panel A reports estimates of average centrality differential (and associated t-values): $\bar{\mu}_{ij} = (\mu_{ij} - \mu_{ji})/2$ for $i > j$ (in basis points). Panel B contains the α , σ_f^2 , and σ_g^2 estimates.

Panel A. Estimates of $\bar{\mu}_{ij}$ for $i > j$

		2016				2013			
		Low group				Low group			
	High group	1	2	3	4	1	2	3	4
AUD/JPY	2	0.225 (1.86)				0.214 (2.68)			
	3	0.330 (4.18)	0.156 (1.88)			0.100 (2.99)	-0.075 (-1.89)		
	4	0.312 (6.01)	0.171 (3.52)	0.221 (4.54)		0.112 (2.49)	0.102 (3.53)	0.118 (4.88)	
	5	0.218 (5.38)	0.116 (3.87)	0.052 (2.03)	-0.026 (-1.55)	0.177 (9.79)	0.017 (0.88)	0.063 (3.43)	-0.087 (-8.40)
	2	0.196 (9.99)				-0.058 (-6.79)			
AUD/USD	3	0.208 (12.69)	0.027 (2.12)			0.042 (6.42)	0.031 (4.79)		
	4	0.213 (14.91)	0.013 (0.94)	-0.008 (-0.90)		0.066 (11.58)	0.047 (8.30)	-0.051 (-11.46)	
	5	0.228 (25.84)	0.112 (14.38)	0.035 (6.54)	0.053 (9.46)	0.092 (20.72)	0.060 (14.95)	0.020 (6.23)	0.031 (9.09)
	2	0.253 (9.16)				0.179 (14.50)			
	3	0.223 (9.65)	0.121 (4.94)			-0.021 (-2.67)	0.118 (15.03)		
EUR/CHF	4	0.106 (5.88)	0.066 (4.98)	-0.112 (-6.31)		0.042 (5.09)	0.043 (7.18)	0.083 (10.08)	
	5	0.143 (10.73)	0.079 (7.92)	0.038 (4.24)	0.012 (1.71)	0.099 (14.77)	0.145 (27.17)	0.094 (18.98)	-0.012 (-2.87)

		2016				2013			
		Low group				Low group			
	High group	1	2	3	4	1	2	3	4
EUR/GBP	2	0.204 (5.39)				0.257 (7.98)			
	3	0.132 (5.73)	0.148 (6.44)			-0.109 (-6.34)	0.048 (2.97)		
	4	0.176 (9.62)	0.090 (6.60)	-0.024 (-1.50)		0.164 (8.40)	0.186 (11.26)	0.064 (5.18)	
	5	0.083 (4.78)	0.088 (6.58)	0.009 (0.69)	0.013 (1.57)	0.105 (7.55)	0.083 (7.33)	0.075 (9.04)	-0.024 (-2.94)
	2	0.297 (4.85)				0.110 (4.84)			
EUR/JPY	3	0.130 (3.90)	0.185 (5.04)			-0.013 (-0.87)	0.044 (2.76)		
	4	0.183 (6.57)	0.066 (2.49)	-0.075 (-3.34)		0.122 (9.98)	0.119 (9.89)	0.000 (0.01)	
	5	0.127 (6.63)	0.011 (0.66)	0.022 (1.53)	0.039 (4.73)	0.115 (12.00)	0.056 (6.81)	-0.036 (-4.62)	-0.089 (-20.49)
	2	-0.003 (-0.30)				0.070 (14.81)			
	3	0.071 (9.82)	0.101 (13.56)			0.057 (12.40)	0.055 (13.59)		
EUR/USD	4	0.069 (12.64)	0.082 (14.52)	-0.001 (-0.29)		0.053 (14.50)	0.030 (7.85)	-0.076 (-24.65)	
	5	0.073 (15.45)	0.076 (22.03)	0.047 (16.22)	0.018 (8.82)	0.053 (20.08)	-0.021 (-8.87)	-0.022 (-9.86)	0.010 (4.63)
	2	0.437 (3.65)				-0.495 (-4.79)			
	3	0.233 (3.03)	-0.037 (-0.56)			0.206 (3.08)	-0.050 (-1.13)		
	4	0.137 (2.56)	-0.023 (-0.37)	0.112 (2.32)		0.023 (0.56)	0.002 (0.03)	-0.011 (-0.31)	
GBP/JPY	5	0.199 (5.50)	-0.071 (-2.51)	-0.062 (-3.03)	-0.123 (-9.05)	-0.026 (-1.09)	0.044 (2.47)	0.036 (1.87)	-0.077 (-4.06)
	2	0.086 (6.37)				0.046 (5.42)			
	3	0.050 (3.21)	0.103 (6.91)			0.078 (9.24)	0.059 (9.75)		
	4	0.142 (16.15)	0.074 (9.65)	-0.043 (-5.40)		0.063 (9.38)	0.041 (5.45)	0.004 (0.82)	
	5	0.109 (14.31)	0.106 (17.69)	-0.013 (-2.28)	0.061 (17.54)	0.117 (26.84)	-0.006 (-1.47)	0.011 (2.95)	0.037 (10.85)
GBP/USD	2	0.260 (5.07)				0.154 (4.93)			
	3	0.360 (7.24)	-0.055 (-1.40)			0.035 (1.95)	0.094 (4.57)		
	4	0.452 (17.06)	0.076 (3.63)	-0.014 (-0.74)		0.127 (7.70)	0.124 (6.42)	0.056 (4.04)	
	5	0.444 (22.60)	0.056 (3.27)	0.046 (3.42)	0.008 (0.91)	0.240 (17.96)	0.139 (10.33)	0.105 (10.96)	-0.038 (-3.36)
	NZD/USD	2	0.260 (5.07)				0.154 (4.93)		
3		0.360 (7.24)	-0.055 (-1.40)			0.035 (1.95)	0.094 (4.57)		
4		0.452 (17.06)	0.076 (3.63)	-0.014 (-0.74)		0.127 (7.70)	0.124 (6.42)	0.056 (4.04)	
5		0.444 (22.60)	0.056 (3.27)	0.046 (3.42)	0.008 (0.91)	0.240 (17.96)	0.139 (10.33)	0.105 (10.96)	-0.038 (-3.36)

		2016				2013			
		Low group				Low group			
	High group	1	2	3	4	1	2	3	4
USD/CAD	2	0.118 (4.99)				-0.028 (-3.44)			
	3	0.146 (9.54)	0.109 (9.40)			-0.145 (-15.67)	-0.069 (-6.53)		
	4	0.165 (15.17)	0.113 (10.95)	0.030 (4.08)		0.033 (5.71)	0.078 (17.50)	0.061 (7.60)	
	5	0.161 (19.26)	0.104 (17.00)	0.004 (0.88)	0.061 (12.59)	0.096 (21.84)	0.123 (32.86)	0.186 (31.45)	0.024 (7.32)
	2	0.109 (3.46)				0.093 (5.04)			
USD/CHF	3	0.131 (3.77)	-0.039 (-1.52)			0.102 (6.23)	0.045 (2.58)		
	4	0.282 (14.93)	0.089 (5.72)	0.047 (3.16)		0.171 (14.00)	0.113 (9.63)	0.115 (11.98)	
	5	0.163 (11.44)	0.135 (13.25)	0.057 (7.47)	-0.010 (-1.75)	0.187 (14.62)	0.048 (4.53)	0.055 (5.75)	-0.045 (-6.20)
	2	0.031 (1.71)				-0.014 (-2.09)			
	3	0.045 (3.46)	0.078 (8.89)			0.124 (18.73)	0.041 (6.77)		
USD/JPY	4	0.131 (12.73)	0.064 (8.35)	0.020 (3.02)		0.080 (14.56)	0.046 (9.79)	-0.041 (-9.90)	
	5	0.088 (12.59)	0.125 (28.62)	0.034 (8.94)	0.041 (14.55)	0.082 (20.19)	0.065 (18.12)	-0.036 (-12.75)	0.033 (10.41)
	2	0.118 (2.96)				0.005 (0.31)			
	3	0.250 (7.19)	0.103 (3.83)			0.103 (6.21)	-0.024 (-1.69)		
	4	0.300 (16.29)	-0.019 (-1.34)	0.026 (2.08)		0.006 (0.48)	0.076 (5.27)	-0.020 (-1.60)	
USD/MXN	5	0.225 (17.44)	0.084 (7.98)	-0.001 (-0.06)	0.062 (13.23)	0.149 (17.82)	0.008 (1.10)	-0.032 (-5.94)	-0.017 (-2.96)

Panel B. Estimates of α , σ_f^2 , and σ_g^2

	2016				2013			
	Intercept		Variance parameters		Intercept		Variance parameters	
	Estimate	t-Value	σ_{fk}^2	σ_{gk}^2	Estimate	t-Value	σ_{fk}^2	σ_{gk}^2
AUD/JPY	-0.005	(-0.25)	11.433	6.996	0.058	(4.68)	6.540	4.277
AUD/USD	0.023	(2.53)	8.771	3.753	0.041	(7.37)	3.107	1.364
EUR/CHF	0.011	(1.17)	2.862	0.624	0.080	(10.99)	2.427	0.414
EUR/GBP	-0.027	(-2.20)	6.374	2.755	0.016	(1.54)	3.569	2.421
EUR/JPY	0.019	(1.65)	4.541	3.083	0.031	(3.81)	4.404	4.353
EUR/USD	0.013	(2.36)	3.933	2.053	-0.003	(-0.64)	2.217	1.814
GBP/JPY	0.008	(0.51)	4.240	3.738	0.076	(6.16)	3.884	3.779
GBP/USD	0.004	(0.63)	3.745	2.141	-0.007	(-1.34)	1.858	1.321
NZD/USD	0.067	(4.95)	10.533	4.636	0.018	(1.56)	4.714	2.947
USD/CAD	0.040	(5.25)	4.777	2.221	0.012	(1.91)	3.136	0.830
USD/CHF	0.001	(0.10)	5.874	1.384	0.040	(3.85)	3.496	2.476
USD/JPY	0.051	(8.55)	4.443	3.041	0.034	(6.02)	4.293	3.729
USD/MXN	0.004	(0.59)	0.663	0.562	0.007	(1.00)	0.904	0.798

Table 14. Average centrality differentials and bid-ask spreads by currency pair

The half-spread (in basis points) is from Hasbrouck and Levich (2019, Table 5). The average centrality differentials for group 5 vs. group 1 (also in basis points) are from Table 13. The last row reports unweighted averages across the thirteen currency pairs.

	Half-spread	Average centrality differential, $\bar{\mu}_{5,1}$
AUD/JPY	0.598	0.243
AUD/USD	0.463	0.219
EUR/CHF	0.550	0.114
EUR/GBP	0.443	0.112
EUR/JPY	0.366	0.099
EUR/USD	0.222	0.069
GBP/JPY	0.549	0.040
GBP/USD	0.349	0.121
NZD/USD	0.730	0.313
USD/CAD	0.427	0.176
USD/CHF	0.520	0.177
USD/JPY	0.270	0.115
USD/MXN	1.590	0.159
Average	0.544	0.174

Table 15. Centrality flows and returns

For each April (year) and currency pair, Olsen quotes are grouped into intervals of (approximately) five minutes. Settlement members are partitioned in groups numbered from 1 (low centrality) to 5 (high). For interval t , let v_{ijt} denote the total volume (USD equivalent) of all settlements in the interval in which the base currency buyer's centrality group is i and the seller's group is j . The net centrality flow is

$$NCF_t = \sum_{i=1}^5 \sum_{j=1}^5 \text{sign}(i - j) v_{ijt}$$

The signed logarithm of the flow is $\text{Log}NCF_t = \text{Sign}(NCF_t) \log(1 + |NCF_t|)$ where AmtUSD_{ij} is the amount of the settlement (in millions USD). The signed logarithm of the net centrality flow is $\text{Log}NCF_i = \text{Sign}(NCF_i) \log(|NCF_i|)$. Table reports coefficient estimates for the regression

$$r_t = a + b \text{Log}NCF_t + e_t$$

where r_t is the return on the base currency over the interval (in basis points, using bid-ask midpoints). t -values use Newey-West (heteroscedasticity- and autocorrelation-consistent) standard errors.

	2016		2013	
	<i>a</i>	<i>b</i>	<i>a</i>	<i>b</i>
AUD/JPY	-0.070 (-0.84)	-0.409 (-6.07)	0.029 (0.37)	-0.188 (-3.82)
AUD/USD	0.012 (0.19)	-0.061 (-2.76)	-0.004 (-0.09)	-0.035 (-2.84)
EUR/CHF	0.007 (0.29)	-0.042 (-2.24)	0.010 (0.57)	-0.019 (-1.86)
EUR/GBP	-0.027 (-0.63)	0.026 (1.10)	0.014 (0.39)	-0.043 (-2.16)
EUR/JPY	-0.068 (-1.19)	0.063 (1.25)	0.061 (0.74)	0.124 (3.56)
EUR/USD	0.027 (0.73)	-0.030 (-2.63)	0.021 (0.53)	0.016 (1.59)
GBP/JPY	-0.065 (-0.88)	-0.029 (-0.58)	0.055 (0.69)	0.098 (1.59)
GBP/USD	0.036 (0.79)	-0.031 (-1.92)	0.025 (0.73)	-0.023 (-2.15)
NZD/USD	0.036 (0.61)	-0.148 (-4.19)	0.056 (1.23)	-0.125 (-5.63)
USD/CAD	-0.076 (-1.54)	-0.064 (-3.28)	-0.017 (-0.64)	-0.030 (-3.02)
USD/CHF	0.002 (0.06)	-0.015 (-0.79)	-0.018 (-0.43)	-0.023 (-1.33)
USD/JPY	-0.067 (-1.14)	-0.058 (-3.05)	0.057 (0.91)	-0.044 (-2.52)
USD/MXN	0.001 (0.01)	-0.037 (-1.11)	-0.023 (-0.41)	-0.058 (-2.26)

Figure 1. Implied positions in the AUD, April 2016, top twelve settlement members

For a given settlement member, let f_t denote the net inflow to the member in the t^{th} ten-minute interval in April of 2016 (weekends excepted). Ignoring the initial holding, the member's position evolves as the cumulative inflow $x_t = \sum_{s=1}^t f_s$. The figures plot x_t for April 2016, by currency and by settlement member for the top-twelve members, ranked on volume-weighted degree centrality.

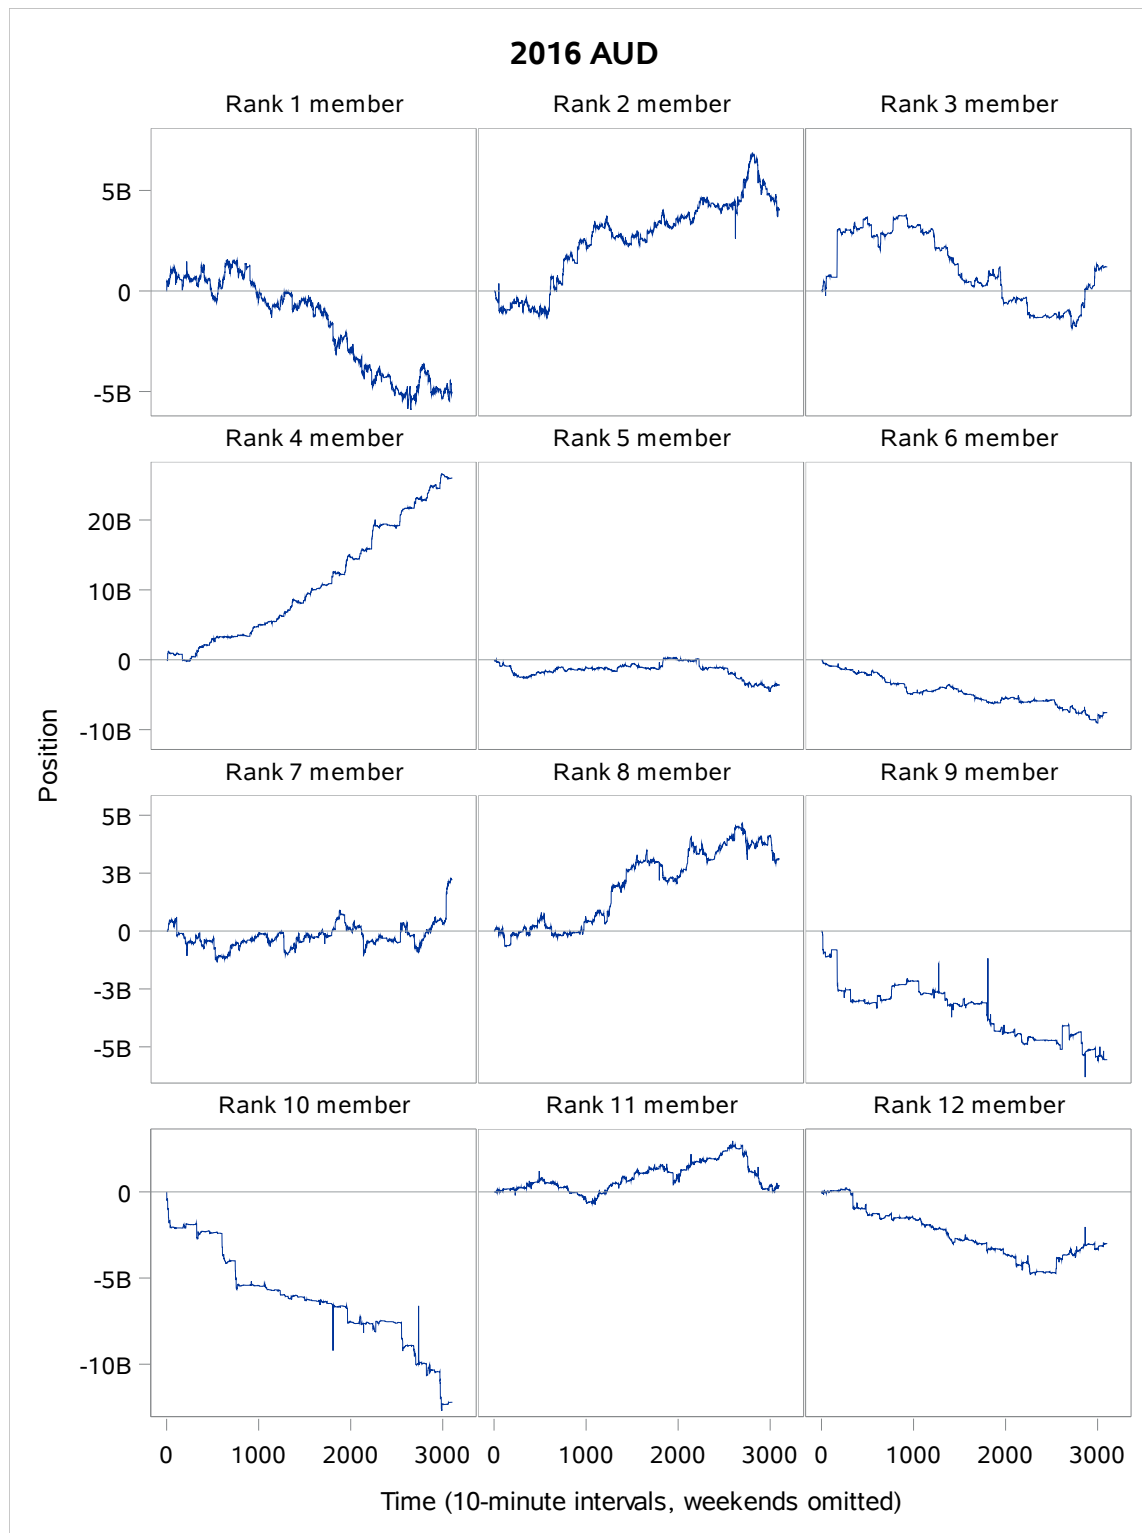


Figure 2. Average centrality differentials and half-spreads

Log-log scatterplot of $\bar{\mu}_{51}$ (the group-5 vs. group-1 average centrality differential) vs. one-half the mean bid-ask spread. The half-spread estimates are computed from the values in Hasbrouck and Levich (2019). The $\bar{\mu}_{51}$ estimates are from Table 13.

