FX Liquidity and Market Metrics: New Results Using CLS Bank Settlement Data

Joel Hasbrouck
NYU Stern

Richard M. Levich
NYU Stern

September 30, 2019

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, discussing institutional background, and comments. We are also indebted to Alex Ferreira, Yalin Gündüz, Carol Osler, Angelo Ranaldo, Andreas Schrimpf, and audiences at the 2019 CEBRA / Federal Reserve Bank of New York Conference, the 2019 INFINITI Conference, the International Conference on High Frequency Exchange Rate Dynamics (National Graduate Institute for Policy Studies, Tokyo), Cubist Systematic Strategies, Vanderbilt University, Università della Svizzera Italiana, the University of Utah, the Carey School at Johns Hopkins, the Stockholm School of Business, and Lund University for comments on earlier drafts. We take responsibility for all remaining errors.

Online Appendix 1 presents supplemental tables and figures. Online Appendix 2 discusses the reconciliation of activity measures constructed from CLS settlements BIS survey, and EBS.

This paper supersedes an earlier manuscript entitled “FX Market Metrics: New Results Using CLS Bank Settlement Data.”

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

Abstract

Using a new and comprehensive sample of foreign currency settlement instructions submitted to the CLS Bank, we investigate activity and liquidity in the foreign exchange market. The settlement data are observed at high frequency and span a wide range of currencies, participants, and trading mechanisms. With respect to overall turnover, they are substantially more comprehensive than activity on the EBS and Reuters electronic execution platforms. The relative settlement activities across currency pairs accord closely to BIS survey estimates and are more consistent with BIS estimates than EBS volumes. We estimate price impact coefficients using three alternative approaches. The estimated coefficients generally decline from April 2010 to April 2013 and rise from April 2013 to April 2016. This suggests that market liquidity rises and then falls for larger orders that would be broken up and executed, but the net change between 2010 and 2016 cannot be clearly signed. In contrast, Olsen bid-ask spreads generally decline, suggesting an ongoing improvement in liquidity for smaller orders. Additionally, we find that from 2010 to 2016 median settlement sizes and price clustering decline, which is consistent with a broad shift to algorithmic trading.

Keywords: Foreign exchange, CLS Bank, market microstructure, liquidity, algorithmic trading.

JEL Classification: F31, G12, G15, G23

1. Introduction and Motivation

With global trading volume in excess of $5 trillion per day, foreign exchange (FX) is often acknowledged as the world's largest financial market. However, because the FX market is geographically dispersed, lightly regulated and without centralized record keeping empirical studies of the market have been handicapped. Past studies have relied on small slices of the market such as specialized electronic platforms (e.g. EBS or Reuters), individual bank data, or screen capture sources (e.g. Olsen and Associates) which collect samples of indicative prices. These studies have developed many valuable empirical regularities about the market, but they have limitations. There are few reliable indicators of deal characteristics and market liquidity. In this paper, we use a large database of FX settlement instructions from the CLS Bank to address these questions and illuminate earlier findings.

The importance of settlement data is suggested by its close connection to trading processes. In a hypothetical trade, once counterparties have agreed on terms (perhaps via an electronic platform), they transmit instructions to CLS bank, which acts as an intermediary to facilitate final clearing and settlement. Because these instructions initiate irrevocable transfers, the essential details – the amounts to be exchanged, the price, and identifiers for each counterparty – are accurate and authoritative.

By reconciling our CLS sample to BIS survey figures, we estimate that CLS handles roughly 37% of spot volume. The CLS settlements span a wide range of quantities, as small as a few pennies to several billion US dollars. The currencies eligible for CLS settlement account for more than 90% of all global FX trading. The parties eligible to use CLS include 70 settlement members, but also (as of 2019) over 25,000 third-party members. Importantly for our purposes, members may funnel their settlement instructions through CLS regardless of whether these trades were arranged on ECNs, via direct dealing, or by voice brokers. This broad spectrum of transactions allows us to compute metrics that are representative of a larger portion of the market.

Our sample comprises all CLS settlement instructions submitted in the Aprils of 2010, 2013, and 2016. These periods correspond to BIS survey months, and we determine that the settlement and BIS data agree in most respects. Summary estimates formed in each of these months reveal new details about FX liquidity. To proxy for price impact, we construct illiquidity ratios for thirteen currency pairs over fixed volume intervals (Barardehi, Bernhardt and Davies (2018)), illiquidity ratios over fixed time intervals (Amihud (2002)), and impact estimates based on bulk volume classification (Easley, Lopez de Prado and O'Hara (2016)). Across currency pairs and time, these proxies are highly and positively correlated with each other, and they are negatively correlated with turnover. Between the most- and least-actively traded pair (the EUR/USD and the AUD/JPY) there is approximately a ten-fold difference in liquidity (based on the 2016 fixed-volume interval estimates).

Across the three samples, for most currency pairs, these price impact proxies generally decline between 2010 and 2013, but increase between 2013 and 2016, suggesting that market liquidity first improved, but then worsened. Bid-ask spreads estimated from Olsen data, however, generally declined over both periods, implying ongoing improvement in liquidity. This is not logically inconsistent because the spread and illiquidity ratio reflect different dimensions of liquidity.

We find that variation in impact proxies across currency pairs is more internally consistent than that of EBS-based price-impact measures. We attribute this to EBS' low volume in currencies for which it is not the dominant market. CLS settlement volumes better reflect overall patterns of FX market turnover as found in the BIS survey data.

Turning to other features, we find that settlement sizes are strongly clustered at one million units of the base currency. Despite this, however, median settlement sizes have declined, from about $1M in 2010 to $750K in 2016. We also find a change in price clustering. In 2010, settlement prices fall on a grid defined by the traditional pip size (0.0001 for most quote currencies; 0.01 for the Japanese yen). By 2016, the price grid has become finer by a factor of ten. We observe a similar change in clustering for bid and ask quotes. The smaller settlement sizes and finer price grids are consistent with a broad shift toward algorithmic trading.

For a trade settled through CLS, the amounts of each currency exchanged (and implicitly the price) are essentially exact because they are confirmed by both sides of the trade. Settlement instructions are submitted, however, subsequent to the trade. In many cases submission may be virtually instantaneous, using straight-through-processing that is linked with the firm's trading systems. In some cases, though, counterparties may separate the confirmation of an FX trade (e.g. via a platform) from the back-office processing of settlement instructions. The former is time sensitive while the latter is not. Despite the possibility of delays in reporting to CLS Bank, we find that most settlements are consistent with market prices observed within the sixty-second window prior to the submission. We argue that measures, which are constructed using time- or volume-aggregated returns and volumes are reliable price impact proxies when imputed transaction times are subject to this degree of measurement error.

This paper proceeds as follows. In the following section we summarize the relevant literature and establish the context for our study. We then describe the CLS Bank and its operations in Section 3. Section 4 presents summary features of the settlement data. In Section 5 we discuss the BIS survey methodology, compare the coverage of BIS and CLS samples, and perform a similar analysis for EBS data. In Section 6 we augment the settlement data with bids and asks (collected at a ten-second frequency) and discuss the correspondence between settlements and market transactions. Section 7 discusses the price impact methodology, and Section 8 presents the estimates. A summary concludes the paper in Section 9.

2. Literature review.

Although we discuss many properties of the settlement data, this study ultimately focuses on liquidity. In line with most preexisting studies, we rely on bid-ask spread and the price impact measures of trading cost. We begin, however, with the observation that the study of FX liquidity has developed jointly with that of price discovery. While liquidity broadly relates to trading cost, price discovery refers to the sources of information and the incorporation of information into prices. The two concepts are economically connected in that the costs of price discovery (due, for example, to asymmetric information) are presumed to be passed through as trading costs. The correspondence is not exact, however. O'Hara (2003) emphasizes the disconnect, noting that a market with perfect agreement on security values (and thus no need for price discovery) may nevertheless exhibit costly trading (illiquidity).

The connections (and the distinctions) between liquidity and price discovery are particularly important in the FX market. The traditional and still-dominant view is that this market is essentially a classic dealer market with two segments. Dealers trade against customers over-the-counter, and dealers trade against each other in the interdealer market.

Most studies of FX liquidity concentrate on the interdealer market. This may be due to data availability, but the interdealer market is also important because it is primary for price discovery. A dealer brings to the interdealer market flow originating from the dealer's customers, which suggests that the interdealer market aggregates the order flows and information of all participants. Numerous studies (discussed below) support this view, and the present paper does not suggest otherwise. With respect to liquidity, however, we will argue that measures based on the interdealer market are incomplete, and that settlement data offer a broader and more accurate characterization.

Empirical studies of price discovery employ various approaches and data samples. Evans and Lyons (2002a, 2008) analyze price discovery using a 1996 sample of DM/$

activity from an interdealer bilateral negotiation platform (the Reuters D2000-1 system). They find that order flow is the main determinant of exchange rate movements and that public news is transmitted through the interdealer order flow. Although bilateral negotiation may still occur in the interdealer market, the two limit order markets, EBS/ICAP and Reuters Matching (which is distinct from the D2000-1 system), are currently more prominent. Of these two systems, EBS has been more closely studied. Several early analyses use a sample of one-minute time-aggregated volume and price data for the EUR/USD and USD/JPY spot exchange market from January 1999 through February 2004: Chaboud, Chernenko and Wright (2008) find volume surges subsequent to scheduled US macro news announcements; Berger, Chaboud, Chernenko, Howorka and Wright (2008) correlate volume and price responses to macroeconomic news announcements. Chaboud, Chiquoine, Hjalmarsson and Vega (2014) find that algorithmic trading is associated with (and indeed causes) fewer triangular arbitrage opportunities, smaller price-change autocorrelations, and reduced volatility, which implies improvement in efficiency and price discovery. Breedon, Rime and Vitale (2016) examine the relation between EBS order flow and Reuters survey forecasts of future currency values. Hagströmer and Menkveld (2019) study a sample of EBS data at 100-millisecond intervals to characterize patterns of price discovery around the January 15, 2015 revaluation of the Swiss franc. They find that EBS participants continued to quote, while updates from alternative venues essentially ceased, consistent with the position of EBS as the dominant venue for price discovery.

The interdealer market's importance for price discovery has motivated studies of its liquidity. Using the 1996 Reuters D-2000-1 DM/$ sample, Evans and Lyons (2002b) assess liquidity using order impact coefficients. As with price discovery, however, most studies have focused on EBS. Breedon and Ranaldo (2013) estimate bid-ask spreads in six EBS currency pairs over a 1997-2007 sample and document intraday patterns. Mancini, Ranaldo and Wrampelmeyer (2013) (MRW) utilize one-second time-aggregated data for nine currency pairs. Their sample period 2007-09 allows the authors to analyze liquidity before and after the Lehman Brothers bankruptcy, and also to identify a common factor that explains co-movement in liquidity across the nine currency pairs as well as liquidity in other principle financial markets. Karnaukh, Ranaldo and Söderlind (2015) (KRS) use a longer sample of EBS data covering January 2007 – May 2012 to estimate and explain time varying and cross-sectional variation in FX liquidity. Like MRW, KRS find evidence of "significantly stronger commonality in periods of market stress—as indicated by high FX and stock market volatility, tight funding constraints (high TED spread), and losses of carry trade portfolios."[1]

Although the EBS studies constitute the largest group of related papers, other analyses use data that provide, at least along some dimensions, coverage similar to the CLS data considered here. Banti, Phylaktis and Sarno (2012) construct liquidity measures based in part on the daily flows of investment institutions reported by a major custodian bank. Evans and Rime (2016) employ daily data on order flows by participant class for the Norwegian kroner.

While these studies reflect rigorous analysis and offer many useful insights, the specificity of the samples limits the generality of the conclusions. Sarno and Taylor (2001) state, "One important consequence of decentralization in the foreign-exchange market is a degree of fragmentation; because not all dealer quotes are observable, transactions may occur at the same time at different prices," (p. 5). To this we would add that since no trading platform accounts for a large share of volume across all currencies and across all classes of market participants, liquidity measures based on any single platform may not be representative of the market as whole.

[1] KRS also analyze low frequency (LF) daily data over the same sample period and find that liquidity measured with LF data co-moves with HF measures. See King, Osler and Rime (2013) for a critical survey of the FX market microstructure literature and Berger, Chaboud and Hjalmarsson (2009) for a review of the literature on FX volatility.

Existing liquidity studies are overwhelmingly EBS-based. Although EBS claims to be "the recognized primary source of global transactional spot FX market data," it is not uniformly preeminent. King, Osler and Rime (2012) note that, "EBS has long dominated interbank trading for the EUR, JPY, and CHF, while Reuters [Matching] dominates the GBP, AUD, CAD, and the Scandinavian currencies."[2] An EBS-based liquidity estimate may not, therefore, be representative of interdealer executions on Reuters Matching, let alone interdealer trades arranged via direct negotiation or using a voice broker. Moreover, many major players in the FX market (such as non-dealing banks and non-bank financial institutions) have not, at least until recently, possessed the ability to trade in the interdealer market.

Accumulating evidence suggests recent shifts in the patterns of FX trading. The BIS Triennial Survey began tracking execution data in 2013. Moore, Schrimpf and Sushko (2016, MSS) offer a detailed analysis of the 2016 figures. Among their various findings about the changing nature of how and by whom FX transactions are executed, we highlight the following:

- "The number of dealer banks willing to warehouse risks has declined, while non-bank market-makers have gained a stronger footing as liquidity providers, even trading directly with end users."
- "The structure of the market may be slowly shifting towards a more relationship-based form of trading, albeit in a variety of electronic forms."

[2] London FX Ltd. (2017) lists currency pairs with an indication of "primary liquidity source" (EBS or Reuters). According to this tabulation: the seven EBS-dominant pairs are EUR/USD, USD/JPY, EUR/JPY, USD/CHF, EUR/CHF, AUD/JPY, and GBP/JPY; neither system dominates trading in AUD/NZD; and the twenty-five remaining active pairs (notably, the majority) are Reuters-dominant. We are not aware of any comprehensive studies of the Reuters Matching platform.

MSS further note that non-bank electronic market-makers have grown to become "top liquidity providers in FX markets." These non-bank participants include XTX Markets, Virtu Financial, Citadel Securities, GTS and Jump Trading. MSS report that these non-bank market-makers are active on multilateral trading platforms (such as Currenex, Hotspot and FXall) where they "provide prices to bank's e-trading desks, retail aggregators, hedge funds and institutional clients" thus acting as liquidity providers. Collectively, these electronic communication networks and dark pools account for about 10% of global FX turnover.[3]

At the same time, top-tier dealer banks have become large scale "internalizers" – meaning that they seek to match offsetting customer orders on their own books rather than immediately hedging them in the inter-dealer market. In 2016, single-bank platforms accounted for 25% of global FX turnover, up from 16% in 2013. Because of these changes, MSS report a decline in turnover arranged through EBS and Reuters combined to only 13% in 2016 compared with nearly 16% in 2013.

Importantly, despite the apparent fragmentation of FX trading and liquidity across many new players and venues, MSS (2016) note that "electronic venues such as EBS and Reuters Matching play a key role in price discovery." Indeed, citing "market sources" MSS conclude that "EBS and Reuters Matching have remained the primary reference sources for benchmark pricing of major currency pairs."

Our own market sources (engaged in currency hedge fund and currency overlay/asset management) acknowledge that they typically inspect prices on EBS and Reuters to gain a sense of current valuations. For execution, however, they turn to their relationship bank or one of the multi-bank platforms that uses a Request for Quote (RFQ) protocol. Our sources believe that this strategy typically results in better execution prices

[3] The BIS Survey Reporting Guidelines describe dark pools as "Private platforms for trading securities (especially for large trade sizes), where access is restricted and quotes are not revealed," and note that, "They are operated by some of the main FX dealing banks, as well as broker-dealers (e.g. BGC) and platform providers. Examples [include] BGC [and] Hotspot QT." (Bank for International Settlements (2015, p. 15)).

and/or additional information about market conditions and trends which they value. Our source accounts conform to the trend, described in MSS (2016), of single-bank and multi-bank platforms gaining popularity apparently at the expense of EBS and Reuters.

Thus, while EBS and Reuters Matching may be considered to act as the primary reference sources for benchmark pricing, it seems more challenging, given the fragmentation of the market, to accept that EBS or Reuters Matching can well serve as the sole venue for estimating market liquidity. Indeed, MSS (2016, p. 35) seem to agree when they conclude "Such changes in the composition of market participants and their trading patterns may have significant implications for market functioning and *FX market liquidity resilience* going forward." (our emphasis)

EBS/Reuters and CLS are likely to have substantial overlap in their clienteles. Both systems originally limited direct participation to major banks, but then expanded to allow indirect participation. EBS and Reuters are now open to smaller bank and non-bank institutions through prime brokerage arrangements.[4] Analogously, CLS participation was initially limited to direct settlement members, but now access is more broadly available through third-party clearing arrangements. Thus, the available evidence suggests that neither EBS/Reuters nor CLS can be considered purely interdealer institutions, but the relative proportions of dealer and customer activities in these systems are not known.

Other relevant CLS-based studies include: Fischer and Ranaldo (2011); Gargano, Riddiough and Sarno (2018); Ranaldo and Somogyi (2018). All use CLS data, but the samples differ in range, grouping, and aggregation. Fischer and Ranaldo, using daily total

[4] King, Osler and Rime (2012) note that in 2010, 30% of London spot FX volume was executed via prime brokerage (PB). In the 2016 BIS survey, total EBS and Reuters daily spot turnover ("electronic indirect") was $372,983M, while PB transactions comprised $564,007M. The PB numbers include all execution methods and platforms, but the relative magnitudes suggest that a substantial portion of EBS and Reuters activity is occurring through PB.

settlement volume document a 5% increase in trading volume on US Federal Open Market Committee announcement days. The other two studies use settlement flows inferred from the settlement data and differentiated by participant.[5] Over the 2012-2017 sample period, Ranaldo and Somogyi find that hourly net buy order flow originating from non-market-making banks, investment funds, and non-bank financial firms generally predicts a permanent increase in the spot rate. Conversely the net buy flow from nonfinancial corporate participants generally predicts a decline. They furthermore identify a profitable trading rule based on these flows. Gargano, Riddiough and Sarno establish that daily volume (over 2012 to 2017) is predictive in a manner that suggests the existence of asymmetric information.

Our CLS sample does not cover a continuous time period, and participants are not grouped by type. It is, however, more detailed in key respects. There is a record for each of nearly 30 million spot settlement agreements, with anonymized identifiers and millisecond timestamps. This facilitates a well-informed characterization of settlement sizes, prices, timing features, and leads to our modified illiquidity ratios.

3. CLS Bank Operations

A foreign exchange settlement is the last stage of the trade process and comprises irrevocable transfers (in opposing directions) of the two currencies. It is initiated when the two parties to the transaction separately submit instructions that name each other as the counterparty and specify the terms of the settlement (the amounts of the two currencies

[5] In each currency pair CLS identifies market-making banks (based on previous trading activity). All other participants are presumed to be takers. In the typical trade, a taker executes an order against a bid or ask quote supplied by a market-maker. If the trade occurs at the bid, the taker is selling, and (if at the ask) buying. CLS then constructs an hourly series of net taker buy volume (or, equivalently, net market-maker sell volume). Additionally, CLS reports net taker buy volume by participant class: corporate, investment funds, non-market-maker banks, and non-bank financial institutions. See Quandl (2018).

being exchanged, who is receiving which currency, and when the settlement is scheduled to occur). When the details match, the transfers proceed.

A settlement is generally distinct from what might be considered, in other contexts, a trade. An execution on an electronic platform, for example, would typically report price, quantity, and a time stamp. The settlement instructions would also include counterparty identifications but would not identify the platform or any other attribute of the execution process (such as the time stamp). There is another important distinction. Because many execution mechanisms, such as electronic limit order books or voice brokers, provide pre-trade anonymity, the resulting trades are presumed to be arms-length transactions at market prices. Settlements, however, are bilateral transfers, and the terms of the exchange aren't necessarily close to current market prices. An exercise of an FX option, for example, involves a transfer at the exercise price, which is likely (conditional on the exercise) to differ from the current market price.

CLS Bank operates the largest FX settlement service. Developed and owned by a consortium of major banks, it began operations in 2002, and is generally considered to be, "the sole multi-currency settlement system of its kind, offering both liquidity savings and settlement risk mitigation across all the major currencies, and the only one that operates on a global basis across all the major currencies," (Financial Stability Oversight Council (2012)).[6] It was originally formed to address Herstatt risk, a reference to a 1974 incident of settlement failures in the US dollar/Deutschemark market that involved transfers between entities in different time zones.[7] Herstatt risk is pernicious not simply because of

[6] The initials "CLS" denote "continuous linked settlement," but the settlement procedure is now generally characterized as "payment versus payment" (PVP). See CLS Group (2013).
[7] On June 26, 1974, Herstatt Bank received Deutschemark settlement payments at its offices in Cologne Germany, but was later that day closed down and forced to cease operations by German banking regulators. It was thus unable to deliver US dollars to its counterparties once US banks opened for business.

the loss of principal (which in the global FX market could be substantial), but also because of a systemic cascade effect should dealers withdraw from the market and be unwilling to quote and trade with their normal counterparties.

The CLS settlement process is payment-versus-payment (PVP). While the details of the entire system are complex, the general PVP principle is straightforward.[8] Both counterparties independently submit to CLS Bank detailed settlement instructions ("submissions"), which CLS then matches. On the agreed-upon settlement date, during the settlement cycle window, CLS Bank receives currency A from one counterparty and currency B from the other counterparty. Once both amounts have been received and CLS has verified that all details match, CLS releases the funds and pays out both counterparties.[9] Once settlement has been concluded, it is final and irrevocable. If counterparty B fails to provide adequate funding, CLS suspends the failing counterparty and takes remedial action to protect the full amount of counterparty A's principal, which avoids settlement risk.[10] The transaction between A and B is left to settle in some other manner. Kahn and Roberds (2001) and Lacker (2001) discuss netting, risk mitigation, and incentives for monitoring in the CLS system.

---

[8] For example, of transactions submitted to CLS, only those that are matched and not rescinded will be settled, subject to satisfying certain risk tests. More detail on how CLS works is available here: https://www.cls-group.com/About/CG/Pages/CorePrinciples.aspx

[9] The FX market generally works on a "T+2" settlement schedule (or "T+1," if both parties are in North America). That is, when a spot trade occurs on day "T", settlement instructions are submitted to CLS contemporaneously, but these instructions specify that the transfer should actually occur two days later. Forwards and far legs of swaps, of course, will have varied settlement dates, and so will depart from this convention. The date-time stamps on our data refer to the submission (of the settlement instructions).

[10] The Allsopp Report, an influential document that prefigured CLS, refers to a "guaranteed refund system," wherein "counterparties are guaranteed that any settlement payment they make will be cancelled or returned if their counterparties fail to pay what they owe," (Bank for International Settlements (1996)). This contrasts with the "guaranteed delivery system" used in regulated futures and options markets, where counterparties post collateral and a clearinghouse guarantees delivery.

CLS settlement operations are contingent on real-time gross settlement domestic payments systems, countries' acceptance of the legality of a foreign entity (CLS Bank) to deem a transaction final and irrevocable, and CLS Bank's acceptance of counterparty risk. CLS settlement is therefore only available for a restricted set of eligible currencies, eligible products, and eligible counterparties or members.

In April 2016, there were 18 CLS-eligible currencies including the major G-10 currencies plus the Korean won, South African rand, and others. Collectively, these 18 currencies accounted for 92.8% of global turnover in the 2016 BIS survey although this overstates the potential reach of CLS because both currencies as well as both counterparties must be CLS-eligible to settle in CLS. Levich and Packer (2017) estimate that 2013 turnover among all pairs of the then-17 CLS currencies measures 90.46% of global turnover.

In April 2016, CLS settlement was available for spot FX trades, outright forwards, FX swaps, and currency swaps. Collectively, these three products accounted for 95.0% of global turnover in the 2016 BIS survey. FX options (representing the final 5.0% in the BIS survey) are a special case. The initial payment of an FX option premium does not settle through CLS (the premium is simply a one-way payment from the buyer to the seller). However, an option exercise is CLS eligible and appears as a spot settlement when exercised.

Direct participation in CLS is limited to settlement member financial institutions (currently 70 in number). In addition, though, settlement members can grant indirect access to other institutions ("third parties"). The settlement member, designated in this capacity as a Third Party Service Provider (TPSP), acts as a gatekeeper to CLS, assumes the risks of dealing with their third-party clients, and charges these clients for their services. The arrangement is distinct from a prime brokerage relationship, but obviously exhibits certain similarities.[11] The number of third-party members is large and growing: Levich and Packer (2017) report 11,000 in 2014; CLS' current website claims over 25,000.

Third party institutions can be commercial banks, central banks, non-bank financial institutions, corporations and investment funds. These non-bank institutions are important in that they do not fit cleanly into the customer/dealer dichotomy. They must be sufficiently large and sophisticated to prefer settlement of their transactions through CLS but would not typically be acting as an FX dealer. They may also be subsidiaries, affiliates, or other sub-units of settlement members. This is significant because it precludes identifying any given member as a distinct and independent economic agent. Although members and third parties have the right to submit eligible transactions for settlement in CLS, they are under no obligation to do so. Bilateral settling (the accepted practice prior to the start of CLS) is still an available option.[12]

4. Data and summary statistics.

Our data sample consists of all submissions to CLS during the Aprils of 2010, 2013, and 2016. This sample was chosen to correspond to the BIS triennial surveys. Each data record reflects submissions by both sides and corresponds to one settlement. By

[11] In a prime brokerage relationship, the client trades using the credit and authority of the named broker. In the third party settlement relationship, "Third party service providers interface with CLS on behalf of their third parties and take legal obligation for their payments. Third party service providers handle all instructions and funding on behalf of their third parties," CLS Bank (2013). Thus, while a prime broker arrangement is a sponsorship that provides access to a trading platform (like EBS), the third party arrangement provides access to the settlement mechanism.

[12] A CLS survey of their own settlement members reported that bilateral netting was used to settle 25.8% of turnover even for trades involving CLS-eligible currencies (CLS Group (2014)). Members may also elect other settlement methods such as on-us (when the counterparty holds an account at the member's financial institution), bilateral netting, or other PVP systems. Given the risk mitigation advantages associated with using CLS and the large number of member counterparties, it is unclear why counterparties select bilateral settlement. See Kos and Levich (2016) for further discussion.

convention, the first member to submit settlement instructions is designated as the "trading party," and the currency they are receiving as the "buy currency"; the other (subsequently arriving) submitter is considered the counterparty, and their received currency is the "sell currency". There is no economic content to these designations, however, as the essentials of the settlement would be identical if the designations were to be reversed. The data are time-stamped with one-millisecond precision.[13] These times impound, however, a random delay relative to the original trade. We discuss this at greater length in Section 6.2.

Table 1 reports total sample counts and settlement values, categorized by CLS' classification of instrument type. In all years spot settlements are the most numerous, comprising over 90% of all settlements. Their proportion by value is much smaller, however, only about 20% to 30%. The reverse holds for FX swap settlements. The near and far legs taken together represent less than 5% of the total settlement counts, but 50% to 70% by value.[14] Settlements in the options category reflect exercises (not sales). Over time (across the three April samples) the spot value proportion falls, and the swap value proportion rises.

Most of our analyses focus on spot settlements. Figure 1 depicts the histogram and sample CDF of spot settlement sizes for April 2016, in units of the base currency.

---

[13] Each record reports the time when CLS accepted the submission of the trading party, the accept time of the counterparty's submission, and the time when the instructions were matched. Because the trading party is designated when the first submission is processed, it is the earliest. It is therefore closest in time to the trade or similar event that motivated the settlement. References to "time" in this paper accordingly denote the trading party accept times. Submissions are generally processed continuously, on arrival. In our sample, however, each day generally contains one interval of three or four minutes where submissions are queued and the accept times are batched. We estimate that this affects about 0.4% of the observations. Data from 2016 onwards do not contain queued intervals.

[14] FX swaps by definition involve paired near and far legs. The near- and far-leg counts in Table 1 are very close, but not exactly equal. We believe that the difference is due to minor timing discrepancies in the database extraction of our sample.

(Corresponding figures for other years are reported in Online Appendix 1.) The distribution reflects numerous small trades. The traditional minimum size on the interdealer trading platforms is one million units of the base currency (Chaboud, Chernenko and Wright (2008), for example). Roughly 25% of the spot settlements are smaller, suggesting that the settlement data capture at least some non-interdealer activity.

The distribution of spot settlement sizes also exhibits a strong clustering. Histogram peaks fall on "natural" multiples, such as 1,000, 2,000, 5,000, 10,000, 20,000, 50,000, and so on. The distribution is strongly concentrated at one million, with well-defined subsidiary peaks at 500,000, 100,000, and 10,000. We stress that the units in the figure are not restated to a common numeraire: a settlement for one million Euros in the EUR/USD pair lies at the same horizontal location as a settlement of one million USD in the USD/JPY pair. When measured in units of the quote currency settlement sizes are also clustered but the peaks are not as sharply defined, nor do they occur on natural multiples (supplementary figures in Online Appendix 1).

This pattern suggests that clustering arises from trading conventions in the base currency, and furthermore that, given the relative stability of exchange rates over this period, the clustering in quote currency amounts is mostly reflective and derivative of the clustering in the base currency. For FX microstructure analysis, the practical importance of this observation is that apparent variation in settlement (and, presumably, trade) values may be mostly a function of the numeraire currency and/or variation in the exchange rate. As the numeraire is often set (in the present paper and elsewhere) to USD equivalent, the clustering of trade size will be obvious only in pairs for which the USD is the base currency.

Over time there is a distinct trend toward smaller spot settlements. The median size, for example, goes from 998,973 (USD equivalent) in 2010 to 797,073 in 2013, and then to 758,073 in 2016 (Online Appendix 1). This is consistent with recent trends in trade sizes in many other markets, where the drop is commonly attributed to a technology-related decline in fixed (per trade) costs and the rise in algorithmic trading.[15] Another contributing factor may be growth in the number of CLS third-party members.

Table 2 reports shares of spot turnover by currency. The dominant currencies are the USD, EUR, and JPY. The percentages total to two hundred because each settlement has two sides/currencies. The percentages are very close to the corresponding figures from the corresponding BIS surveys.[16] These results pertain to shares; Section 5 discusses reconciliation of the total amounts.

Table 3 reports for each currency the relative shares (by value) of the contra currencies, that is, the currencies on the other side of the settlements. For brevity, the table reports only 2016 values. Values for 2010 and 2013 are similar and are reported in Online Appendix 1. Percentages sum to one hundred across each row. For example, the first row corresponds to the AUD: relative to the total value of all settlements that have the AUD on one side, 1.4% (by value) have the CAD on the other side. These shares suggest that while the USD is usually the dominant contra currency, there are some notable exceptions. The USD share is relatively small for the Nordic currencies (DKK, NOK, SEK), each of which is much more likely to be exchanged for EUR. In many currencies (CAD, HKD, ILS, KRW, SGD, ZAR) the entry in the USD column exceeds ninety percent. This may reflect the use of the USD as a vehicle currency: if a currency pair has no established market, each may be

---

[15] For example, in the Aprils of 2010, 2013, and 2016, the average trade sizes for NYSE-listed equities are 320, 239, and 200 shares, respectively.

[16] For the Canadian dollar (CAD), the CLS shares are markedly higher than the BIS shares. In 2010, a Bank of Canada assessment of risks in the FX market recommended (as a first priority), "Establish same-day USDCAD settlement in CLS," and (as a second priority), "Increase use of CLS for FX transactions …," Bank of Canada (2010). In 2011, FX Week reported, "Speaking at the FX Invest North America congress in Toronto, Donna Howard, chief of the financial markets department at the BoC, … said the priority for the Canadian FX market in particular is to establish same-day USD/CAD settlement in CLS – targeted for 2011. Howard said that, while the self-regulatory nature of the FX industry is a strength, the central bank has used 'moral suasion' to ensure all five major Canadian banks are live on CLS by end-2010," FX Week (2011). Thus, the high relative use for CLS in CAD settlements may reflect regulatory pressures.

converted to/from USD as an intermediate step. The KRW is an extreme case in that *all* CLS spot settlements involve the USD on the other side.

Although settlement instructions may be submitted at any time, the submissions are not uniform in time. Figures detailing spot settlement activity are presented in Online Appendix 1, but the main results can be summarized as follows. Intraday turnover is elevated in three periods corresponding to business hours in Tokyo, London, and New York. Major currencies tend to follow this pattern, but activity in other currencies is more concentrated in local business hours. There are also regularities at shorter periods. Activity plots in 2013 and 2016 (but not 2010) exhibit hourly peaks.[17] We discuss other features of settlement timing in Section 6.2.

We noted above the clustering in the distribution of settlement quantities. Clustering is also a feature of settlement prices. Traditionally, most exchange rates were quoted to the fourth decimal place, implying a tick size ("pip") of 0.0001 (USD per EUR, for example). Yen exchange rates were traditionally quoted to the second decimal place, 0.01 (JPY per USD, for example). A reasonable null hypothesis is that given sufficient mixing in exchange rates the values in the pip digit would be uniformly distributed (10% on each digit). Table 4 summarizes the actual sample distribution. For brevity, digits "0" and "5" are reported separately, and the remaining eight digits are summarized as a single group ("Other"). The null hypothesis of 10% probability on each digit implies 10% on "0" and "5" and 80% on "Other". In all years the frequencies of "0" and "5" are similar and slightly above 10%. Across settlement type, clustering is strongest is option settlements. Since these reflect option exercises, it is reasonable to assume that exercise prices in the FX

---

[17] The reasons for the hourly peaks are unclear. Between 2010 and 2013 there were two institutional changes: CLS introduced an aggregation service; and the number of third-party settlement members increased substantially. Although it would be reasonable to conjecture that the hourly peaks arose from aggregation, we find hourly peaks in non-aggregated as well as aggregated settlements. Nor does it seem obvious why hourly settlement might be preferred by third-party members.

market (like those in the equity market) are fixed at natural multiples of the price increment. Near and far swap settlement rates also exhibit modest clustering.

Table 4 also tabulates occurrence frequencies for the next finer digit, also known as the micro pip or pipette: 0.00001 for most currency pairs, and 0.001 for the JPY pairs. Clustering here is more extreme. In all years the full range of digits 0-9 is used, but “0” dominates. There is also a stronger time trend: overall the “0” frequencies are 55.3%, 33.1%, and 27.5% in 2010, 2013, and 2016 respectively.

Price clustering is common in securities markets. In equities markets it is generally attributed to negotiation costs (Harris (1991)) or collusion by quote setters (Christie and Schultz (1994)). In FX markets, clustering in indicative quotes has been studied by Hasbrouck (1999) and Sopranzetti and Datar (2002). Osler (2003) analyses the effects of clustering in the trigger prices for stop-loss and take-profit orders. She finds that trigger-price clustering can account for price dynamics associated with technical trading rules, and various aspects of extreme price movements (Osler (2005) and Osler and Savaser (2011)). Chaboud, Dao and Vega (2019, CDV) examine the effects of changes in tick size on the EBS trading platform. We are not aware of any other studies that examine clustering in FX trades or settlements.

On the EBS trading platform, CDV note that prior to March 4, 2011 prices were quoted to four decimal places (one pip). This regime would include our April 2010 sample. On March 7, 2011, the fifth decimal place became available. In 2012, use of the fifth place was restricted to “0” and “5”, and this regime was presumably in place during our 2013 and 2016 samples. We find that the decline in settlement price clustering is slightly stronger for EBS pairs, but it is clearly evident in non-EBS pairs as well (Supplemental Table S2 in Online Appendix 1): the increased usage of finer increments is not limited to EBS.

The decline in price clustering over the 2010-2016 period containing our data is consistent with both negotiation cost and collusion hypotheses. During this period, US and UK authorities investigated and brought charges in matters related to FX benchmark price setting, and these enforcement actions would have generally discouraged further collusion. With respect to negotiation costs, if these are viewed broadly as including costs of monitoring the market, repricing orders, and so forth, it is logical to associate the decline in clustering with the rise in algorithmic trading.

5. Reconciliation and comparison of the CLS settlement data with other sources

Because CLS data have not been extensively used in research studies, it is useful to examine its coverage and consistency with other sources. Specifically, we reconcile and compare the composition of CLS data with BIS survey figures and data collected from Reuters and EBS electronic platforms used in earlier studies. We summarize the results below; details are provided in Online Appendix 2.

We believe that data from the BIS Triennial Survey offers the most comprehensive picture of the global FX market. The BIS survey constructs global aggregates from figures supplied by participating central banks, adjusting for double counting within and across national boundaries. To facilitate comparability with our settlement data (and with Reuters and EBS) we also adjust the BIS data for multiple counting of prime-brokered trades. The BIS does not classify prime-brokered trades in 2010, so we focus primarily on 2013 and 2010. Net of our prime-brokerage adjustment, CLS spot settlement volume in 2013 and 2016 accounts for 36.1% and 36.9% of BIS spot turnover. For those years, EBS and Reuters report the spot volume executed on their platforms. Their combined volume accounts for 15.0% and 13.1% of adjusted BIS spot turnover. Thus, the CLS settlement data are substantially more comprehensive.

We also examine activity across currency pairs. It was noted above that most prior FX liquidity studies are based on EBS data. One comprehensive study, Mancini, Ranaldo and Wrampelmeyer (2013) reports activity for a 2008-2009 sample. We compare their EBS trade and volume figures with 2010 BIS survey and 2010 CLS settlement quantities. Across currency pairs, BIS and CLS values substantially agree. Relative to these sources, though, EBS activity estimates slightly over-weight the pairs for which EBS is the dominant platform, but substantially under-weight the Reuters-dominant pairs. We will show below that this underweighting strongly affects liquidity estimates based on EBS data.

6. Settlements and market prices.

Because all trades end in settlement, it might be supposed that there is close agreement between settlement prices and market quotes. Settlements can also arise, however, from transfers (such as option exercises) that do not represent arms-length transactions. Discrepancies might also stem from price changes occurring over the delay between trade and the submission of settlement instructions.

To investigate the correspondence between market and settlement prices, we supplement the CLS data with Olsen quotes. Olsen Financial Technologies, a commercial data provider (olsendata.com), has compiled historical bid and ask data for major currency pairs. The quotes are streamed by consolidators and major banks. Our data are constructed over ten-second intervals, and within each interval Olsen supplies the first new bid-ask pair. In practice, these observations are close to the start of the interval.

We view these as indicative prices. They are not necessarily firm (available for immediate execution) nor are they necessarily the best bid and offer available to any participant. Our Olsen sample consists of thirteen major pairs, which altogether contain approximately ninety percent of our CLS spot settlements.

6.1. Spreads and clustering in bids and asks

For each Olsen record we compute the bid-ask midpoint, the absolute spread (the ask less the bid) and the proportional spread (the absolute spread divided by the midpoint). Table 5 reports the medians of absolute spreads (ask less the bid), bid-ask midpoints, and relative spreads (absolute divided by the midpoint) by currency pair and year. Both absolute and relative spreads decline over the 2010-2016 period. Figure 2 depicts the median relative spreads by year and pair. These exhibit a clear downward trend.

Like the settlement prices, bids and asks are clustered. We examined the frequency of digits in the decimal place corresponding to the pip (the second place for the JPY pairs, the fourth place for all others) and in the next (micro pip) decimal place. In all years there is little discernible clustering in the pip place. The micro pip frequency, though, exhibits both clustering and a trend. In 2010 the frequency of a "0" digit across bids and asks in all pairs is 74.9% (implying that the other digits are not generally used). In 2013 and 2016 the "0" frequencies drop to 26.7% and 19.5%, consistent with the trend in settlement prices.

6.2. Bids, asks, and settlement prices.

We now examine the joint behavior of CLS settlement prices and Olsen bid-ask quotes. As an illustration, Figure 3 plots CLS settlements and Olsen bid-ask midpoints for the EUR/USD pair on April 17, 2013. (The vertical scale is set to show relevant detail, and so a small number of outlier settlements lie beyond the displayed range.) The figure shows that while the line defined by the bid-ask midpoint is sharply defined, the settlements are visually blurred. That is, the settlement exchange rates appear to exhibit high local variation. In addition, there are clear hourly effects, on-the-hour concentrations of settlements at away-from-the market rates (notably, at 9:00, 11:00, 12:00, and 13:00). There is directional variation in the peaks: the rates are sometimes above and sometimes below the market.

If the settlement prices simply matched the prevailing quote midpoint, submitted with a fixed delay, the two series would be identical, up to a horizontal shift corresponding to the delay. Instead, the midpoint line tends to define, along the time axis, the leading edge of a broad cloud. This is most clearly visible around 16:00. This pattern suggests random delays in the submission of settlement instructions.[18]

We aim to estimate delays in settlement submissions by using the Olsen quotes as benchmark prices. This has two purposes. First, the distributions of estimated delays are interesting because they may reflect differences in settlement procedures associated with pair, settlement size, execution methods, and so forth. Second, the estimated delays may be used to correct the timing of settlements, leading to more accurate estimates of liquidity.

For each settlement, our approach involves looking backwards from the submission time until we find an Olsen price that is acceptably close to the settlement price. We consider a range of acceptance criteria. The most stringent acceptance criterion is that the settlement price must lie at an Olsen bid or ask, or inside of this range. That is, given a settlement initiated at time t and price p_t, we look backwards until we find an Olsen bid and ask such that $bid_{t-s} \leq p_t \leq ask_{t-s}$, implying a delay of s.

Our Olsen data are not, however, comprehensive: we observe quotes roughly every ten seconds. We therefore replace bid_{t-s} and ask_{t-s} in this rule with the backwards running minimum bid and maximum ask, defining acceptance with delay s if

$$\min_{\tau \in [t-s,t]} bid_\tau \leq p_t \leq \max_{\tau \in [t-s,t]} ask_\tau \tag{1}$$

Table 6 reports the distribution of imputed delays based on this criterion for each sample period in the rows labeled "bid and ask". In April 2010, for example, 20.1% of the spot settlement prices can be contained in an Olsen bid-ask found within five seconds, and an additional 18.0% are contained in an Olsen bid-ask found five-to-ten sections prior to the settlement. Above ten seconds, the acceptance tails off, and 6.6% of the settlements can't be matched to an acceptable quote within ten minutes. Results for 2013 are similar, but in

[18] As noted above, CLS submissions are not directly generated by the execution platform, and the execution platform is not identified to CLS. For settlement purposes, it is essential that the parties to the trade agree on buyer and seller identities, the price and the quantity, but submission time and method of execution are of lesser importance.

2016 the match rates in the first and second intervals drop markedly (to 9.6% and 9.7%, respectively).

The restriction to an Olsen bid-ask interval, however, may well be too restrictive. In contrast to the National Best Bid and Offer (NBBO) widely used in equity market studies, the Olsen quotes do not constitute a continuous record of actionable prices nor are they available (or even visible) to all market participants. Given these considerations, we also investigate less restrictive acceptance criteria. In the first alternative we redefine the acceptance range as $[bid - spread, ask + spread]$, where $spread$ denotes the median bid-ask spread for the pair in the given sample year. The second alternative acceptance range is $[bid - 5 \times spread, ask + 5 \times spread]$. These are denoted in the table as $\pm spread$ and $\pm 5 \times spread$, respectively. Not surprisingly, the broader acceptance ranges shift the distribution in favor of shorter imputed delays.

In each case the imputed delays appear to worsen markedly between 2013 and 2016. This can be attributed, however, to the post-2013 tightening of bid-ask spreads, and the implied shrinking of our acceptance regions. We therefore also consider acceptance regions based on the traditional tick size (pip, 0.01 for the JPY pairs, and 0.0001 for all others). Corresponding to the spread-based intervals, we investigate ranges of $\pm pip, \pm 2 \times pip, and \pm 10 \times pip$. These results are reported in the last three groups of Table 6. For 2010 and 2013, the results for pip-based and $spread$-based acceptance ranges are similar. In 2016, however, the longer delays found in the $spread$-based imputations are not present in the pip-based imputations. To arrive at representative figures, we consider the $\pm 2 \times pip$ match region. After one minute, the percentages of unmatched settlements are 17.6%, 10.8%, and 15.4% in the Aprils of 2010, 2013 and 2016, respectively. That is, 82.4%, 89.2%, and 84.6% of settlements *can* be matched within the minute.

Estimates reported in Online Appendix 1 illustrate other aspects of variation in imputed delays. Across currencies, delays are somewhat longer for the JPY pairs (AUD/JPY, EUR/JPY, GBP/JPY, and USD/JPY). The distributions also vary with settlement size.

Relatively large and relatively small settlements have longer imputed delays. Strikingly, the shortest delays are found for settlements of exactly one million units of the base currency. Consistent with the clustering at this size noted earlier, it seems likely that such standard sizes would be more likely to have routine and automated generation of settlement instructions, perhaps because they arise from trades executed on electronic markets.

7. Estimating price impact

The preceding section establishes that bid-ask spreads in the FX market dropped markedly over the 2010-2016 sample period, consistent with improved liquidity. The bid-ask spread is most meaningful for small orders that can be executed in one trade. Larger orders that are split over time, however, also incur price impact costs because earlier trades in the sequence adversely move the price for later trades. In this section we consider three impact estimates: the classic Amihud illiquidity ratio formed over intervals of fixed time (Amihud (2002)); the illiquidity ratio formed over intervals containing a given traded volume (Barardehi, Bernhardt and Davies (2018)); and a regression estimate based on the bulk-volume classifier (BVC) suggested in Easley, Lopez de Prado and O'Hara (2016, ELO).

To motivate these measures, we start with a simple linear model of price change:

$$\Delta p_j = \lambda x_j + u_j \tag{2}$$

where x_j is the quantity of the j^{th} incoming active order, signed positive when the trader is buying and negative when the trader is selling; $\Delta p_j = p_j - p_{j-1}$ is the first difference of the price; and, u_j is an innovation attributed to non-trade information. The impact coefficient, $\lambda > 0$, is the parameter of interest. This specification can be motivated from an asymmetric information model (following Kyle (1985), Glosten and Milgrom (1985), or Easley and O'Hara (1987)) augmented with public non-trade information that enters the price through the disturbance, u_j.

With signed orders and reliable time stamps, a specification like (2) can be estimated directly (as in MRW, for example). Our settlement flows, however, are unsigned:

we cannot tell which side initiated the trade. Furthermore, the analysis of the previous section establishes that while the settlement time stamps may correspond approximately to actual trade times, the sequencing of settlements can't be assumed to accurately reflect the sequencing of the original trades. These difficulties can be partially mitigated by working with volumes and price changes aggregated over intervals that include multiple trades.

Specification (2) can be reworked in terms of volume by forming an illiquidity ratio for the j^{th} order as:

$$I_j = \frac{|\Delta p_j|}{Volume_j} = \frac{|\lambda x_j + u_j|}{|x_j|} \tag{3}$$

where $Volume_j = |x_j|$ is the unsigned order magnitude. Over a sample of orders, illiquidity can be summarized by the mean or median of I_j. If $u_j \approx 0$, then $I_j \approx \lambda$, which suggests that ignorance of order signs might not necessarily significantly impair estimation of λ. Missequencing of settlements, however, and misalignment with prices might lead to more serious problems.

Although these sequencing and alignment concerns might seem specific to our setting, O'Hara (2015) and Easley, Lopez de Prado and O'Hara (2016) suggest that they are endemic in high-speed fragmented markets, due to random intermarket reporting delays. To deal with both effects, they advocate aggregating price changes and volumes over multiple trades. The rationale is that the total volume and the end-to-end price change over an interval are relatively insensitive to the ordering of trades within the interval.

Adapting the notation for interval k, $Volume_k = \sum_j |x_j|$, the sum of the individual volumes contained in the interval. The price change over the interval is $\Delta p_k = \sum_j \Delta p_j$, and the illiquidity ratio for the interval is $I_k = |\Delta p_k| / Volume_k$. Note that the denominator is not the absolute value of the interval's net order flow $\left(Volume_k \neq \left|\sum_j x_j\right|\right)$. As a result, even if $u_k = \sum_j u_j \approx 0$, we would not expect $I_k \approx \lambda$. Aggregation therefore presents a trade-off, introducing a new source of error as it mitigates others.

In equity applications, illiquidity ratios are generally formed from returns and volumes measured daily, that is, over intervals of fixed time. Our fixed time illiquidity ratios, denoted I_k^{Time}, are defined for two-minute intervals with non-zero settlement volume. Alternatively, Barardehi, Bernhardt and Davies (2018, BBD) investigate equity illiquidity ratios (and other statistics) formed over intervals defined by traded volume. Analogously, our settlement volume illiquidity ratios denoted I_k^{Volume}, are computed over intervals with \$100M USD cumulative settlement volume.

For each currency pair, we compute I_k^{Time} and I_k^{Volume} for all intervals (two-minute or \$100M USD settlements) in the Aprils of 2010, 2013 and 2016. In each month we winsorize the sample values at 95%, and report means of the winsorized samples. We do this because the monthly distributions of the individual I_k^{Time} values exhibit large outliers associated with intervals that have large price changes on low volumes. The I_k^{Volume}, which by design are not formed for low-volume intervals, are much less prone to this problem. We nevertheless winsorize the monthly I_k^{Volume} samples to be consistent with our treatment of the I_k^{Time}.

In equity markets, BBD find that illiquidity ratios computed over volume-based intervals capture institutional trading costs better than those computed over fixed-time intervals. In view of this finding and the I_k^{Time} outliers, the volume-based interval estimates are our preferred measures.

The illiquidity ratio, it will be recalled, is used in situations where net order flow is not observed and specification (2) cannot therefore be estimated directly. ELO suggest a method for imputing net order flow in aggregated data. Their technique, bulk-volume classification (BVC), for volume-based intervals (bars, in their terminology) can be outlined as follows. Let x_k^{Buy} and x_k^{Sell} denote the total buy and sell volume over interval k $\left(x_k^{Buy},\ x_k^{Sell} > 0\right)$. The total interval volume is $Volume_k = x_k^{Buy} + x_k^{Sell}$, and the net signed order flow is $x_k = x_k^{Buy} - x_k^{Sell}$. Define the standardized price change over the interval as $z_k = \Delta p_k / \sqrt{\sigma_{\Delta p}^2}$ where $\sigma_{\Delta p}^2$ is estimated price-change variance. The BVC imputation of buy volume is $\hat{x}_k^{Buy} = Volume_k \Phi(z_k)$, and sell volume is $\hat{x}_k^{Sell} = Volume_k\left(1 - \Phi(z_k)\right)$ where Φ is the standard normal distribution function. A positive price change attributes more volume to buys; a negative price change, to sales. The BVC-imputed net order flow is

$$\hat{x}_k = \hat{x}_k^{Buy} - \hat{x}_k^{Sell} = Volume_k(2\Phi(z_k) - 1). \tag{4}$$

BVC classification is often used in settings where signed orders and reliable timestamps allow for comparison and validation. In US futures data, ELO find that BVC classifications outperform traditional high frequency tick-rule (TR) and quote-based (Lee-Ready) classifications as predictors of informed trading (Lee and Ready (1991, LR)). In US equities data, Chakrabarty, Pascual and Shkilko (2015, CPS) find that TR and LR classifiers provide more accurate measures of net order imbalance than BVC, but also that imbalances estimated by TR, LR, and BVC classifiers possess comparable power in explaining returns, liquidity and trading costs. Their Table 6 reports regression estimates of returns against BVC-based net order imbalances. We estimate similar regressions, though to a different purpose: CPS are concerned primarily with overall explanatory power of the regression; we are interested in the regression coefficient as a measure of liquidity.

Specifically, we use $\hat{x}_k$ in lieu of x_k in the interval analog to specification (2):

$$\Delta p_k = \lambda^{BVC} \hat{x}_k + u_k \tag{5}$$

Then we estimate the impact coefficient, λ^{BVC} by OLS. Of course (and as CPS point out), since the contemporaneous Δp_k is used to construct $\hat{x}_k$ (via the z_k term), the measurement error $\hat{x}_k - x_k$ is correlated with the residual u_k. Thus, while the estimated λ^{BVC} measures the association between returns and net order flow, it cannot be interpreted as a measure of causal impact.

8. Results

For each currency pair and three one-month samples we compute the three price impact estimates described above. Summary statistics for settlement-volume illiquidity

I^{Volume}, fixed-time illiquidity I^{Time}, and OLS estimates for λ^{BVC} are reported in Tables 7, 8, and 9, respectively. The first three columns of each table report for each year the means (or, for λ^{BVC}, the OLS estimate) and standard errors; the remaining columns indicate the direction and significance of year-vs-year differences. Alternatively, the estimates and 95% confidence bounds are depicted in Figures 4, 5, and 6.

All impact estimates are scaled to have dimensions of basis points per $1M US traded. In principle (for a given currency pair and year) they should agree, but generally $I^{Volume} < \lambda^{BVC} < I^{Time}$. The high estimates for I^{Time} reflect the outliers discussed in the last section. Despite the differences in overall scale, the measures are highly correlated. Across the currency pair/sample year panel, Table 10, Panel A, reports the pairwise correlations among these measures and log($BIS\ Turnover$); Panel B reports the partial correlations among the illiquidity measures, after controlling for log($BIS\ Turnover$). All measures are negatively correlated with turnover. The correlations and partial correlations between I^{Volume} and I^{BVC} are the strongest. All correlations in Table 10 are statistically significant with p-values below 0.0001.

In equity market settings, a price impact coefficient can be used to map a hypothetical trade onto an expected price change. Our proxies can support similar calculations, but in our context the mapping is one of association, not causation. Subject to this caveat, we consider the variation in I^{Volume} across currency pairs (Table 7) for 2016. The smallest estimate (highest liquidity) is in the EUR/USD pair ($0.0117\ bp/\$1M$), which is also the pair with the highest turnover. The largest estimate (lowest liquidity) is in the AUD/JPY pair (0.1120). Thus, the illiquidity estimates for the highest and lowest pairs diverge by approximately ten-fold. Indeed, even in moving from the EUR/USD to the GBP/USD (the pair with the highest turnover to the third-highest) the illiquidity ratio approximately doubles, from 0.0117 to 0.0227. Across all currency pairs, the relation between turnover and illiquidity is generally negative (consistent with Table 10), but we note one outlier. In the USD/CHF pair illiquidity is low ($I^{Volume} = 0.0274$) while turnover is also among the lowest. The standard errors of the means in Table 7 are small. Using standard multiple comparison procedures, the differences between 2016 means are statistically significant at the $\alpha = 0.05$ level, except for USD/CHF vs. EUR/GBP and USD/CAD vs. EUR/CHF. These patterns are also generally found for the I^{Time} and λ^{BVC} illiquidity measures (Tables 8 and 9, respectively). Supplementary Table S4 in the Online Appendix provides further details.

These inferences, though, are subject to considerations that don't arise in equity markets. The Kyle model invoked to motivate equation (2) comprises one security, one numeraire currency, one signal, and one market; the foreign exchange market has multiplicities on all of these dimensions.[19] For example, the AUD/USD 2016 estimate for I^{Volume} is 0.0350 $bp/\$1M$. That is, the AUD appears to be about three times more liquid when the USD is the contra-currency than when the JPY is on the other side.

The distribution of trading across currency pairs is a feature of the foreign exchange market that has no exact parallels in equity trading. The AUD can be bought and sold against many other currencies, and order flows in an AUD pair can plausibly reflect information in either or both the AUD and the contra currency. Order flows in the AUD/USD pair, for example, might reflect AUD-specific information (such as developments in the markets for mineral exports) or USD-specific fundamentals (such as a US election result). In the former case, one would also expect trading to be distributed over other AUD currency pairs; in the latter, over other USD pairs. The distribution of settlement volume across currencies can help reconcile the liquidity differences. From Table 3, the AUD is more actively traded against the USD (72.2%) than against the JPY (11.4%).

---

[19] Although a US stock trades on multiple exchanges, these exchanges are tightly linked, and the reported consolidated volume is comprehensive. In a given currency pair, by contrast, even with the settlement data, we can infer only a portion of volume. With a given price change attributed to a smaller portion of volume, the apparent impact will be larger. This affects all the estimates.

The cross-sectional variation in liquidity differs from that suggested by EBS estimates. MRW estimate price impact coefficients using precisely time-stamped quote and trade data from EBS. Like the three estimators considered here, the MRW price impact coefficient can be viewed as a proxy for Kyle's lambda. Their specification (their equation (1)) is a generalization of (2) that allows for additional lags. Figure 7 is a log-log scatterplot of the MRW mean price impact coefficients and 2010 I^{Volume} estimates for the nine currency pairs common to both studies. Visually there appears to be a weak positive correlation between the two measures, but the dependence becomes stronger if we exclude the Reuters-dominant currency pairs (denoted by triangles). More formally, if we include all nine pairs, the correlation between the MRW impact coefficients and I^{Volume} estimates is 0.635 with a p-value of 0.067. If we restrict the sample to the five EBS-dominant pairs, the estimated correlation is 0.952 with a p-value of 0.009. Given the small sample sizes, however, we view these results as suggestive rather than definitive.

Both the MRW price impact and illiquidity measures are ratios with volume in the denominator. When a pair is Reuters-dominant, EBS accounts for a relatively small share of volume. A price impact measure based on EBS volume will therefore be biased upwards, which is consistent with Figure 7.

We now turn to impact variation over time. For all three measures and most currency pairs, impact generally declines from 2010 to 2013 and then rebounds between 2013 and 2016. The 2010 to 2013 year-on-year changes are predominantly negative and significant in Tables 8, 9, and 10; the 2013 to 2016 changes are predominantly positive and significant. Judging from price impact, then, liquidity improved and then worsened. The pervasive declines in the bid-ask spreads, on the other hand, suggest ongoing improvement in liquidity. These two results are not contradictory: the spread measures liquidity for smaller trades; the impact estimate, for larger trades that must be broken up.

For the impact estimates it is difficult to generalize about the net change between 2010 and 2016. For I^{Volume} the change from 2010 to 2016 is significantly positive more often than significantly negative (seven of the thirteen pairs, including the EUR/USD, vs. four, last column of Table 7). The λ^{BVC} estimates are similar (Table 9), but for the I^{Time} estimates (Table 8) decreases dominate.

9. Conclusions

This paper provides a first look at CLS Bank FX settlement data. Like EBS and Reuters quotes and trades analyzed by others, the submitted settlement instructions constitute an ongoing record of FX transactions. Compared with EBS and Reuters data, however, the CLS settlement data have several distinctive advantages. Firstly, they are more comprehensive across trading platforms. Whereas a given currency pair will most likely concentrate on one or the other platform, the CLS data can cover both, and other mechanisms besides (such as bilateral negotiation and voice-brokered trading). Secondly, the CLS data are more comprehensive in terms of volume coverage. In 2016 we estimate that EBS and Reuters together account for about 13.1% of BIS spot turnover, after adjustment for prime brokerage. The corresponding figure for CLS is 36.9%.

The comprehensiveness of the CLS settlement flows offer insights into patterns of exchange. Settlement amounts are highly clustered in size: the modal quantity is one million units of the base currency, with additional clustering on natural multiples (two million, five million) or sub-multiples (such as 500,000 or 100,000). Against most other currencies, the USD is the dominant contra currency: most settlements have the USD on one side of the trade. The Scandinavian currencies (DKK, NOK, SEK) are the exceptions, with the EUR being the dominant contra currency.

Our CLS settlement instructions have been submitted and accepted by both sides of the trade. They are highly accurate, therefore, with respect to prices, quantities, identities of trading parties, and similar terms. The settlement submission and acceptance processes induce delay, however, relative to actual trade times. Our estimates suggest that for approximately 80%-90% of the settlements, the price is within the bounds set by market bids and asks ($\pm$2 pips) in the minute prior to the settlement acceptance time. This suggests a strong correspondence between trades and settlements over minute (and longer) intervals, but not at the second or millisecond accuracy that might be inferred from the precision of the timestamps.

The fact that timestamps and sequencing in settlement instructions do not precisely correspond to those of market transactions means that the many liquidity estimates that are contingent on accurate timing are not available. These problems are not unique to FX markets, however. O'Hara (2015) and others have noted that similar timing concerns arise in equity markets. As algorithms break large parent orders into numerous smaller child orders, the pricing sequence on these orders becomes less informative about market liquidity. In this respect the trading environments in FX and equities have become more similar.

We construct three alternative measures of FX order impact: a standard Amihud illiquidity ratio using fixed time intervals; an illiquidity ratio using intervals fixed in traded volume; and an impact coefficient estimated using bulk-volume classification. All approaches use aggregate data and so should be less sensitive to incorrect sequencing. We calculate point estimates of the mean and standard deviation for each of these estimators, for 13 currency pairs, and three sample months. The correlations between our estimators generally exceed 0.9, suggesting close agreement. The correlations of each estimator with turnover (both BIS and CLS measures) are strongly negative, which while not unexpected, is reassuring. Across currency pairs, the differences in illiquidity measures are economically meaningful and generally statistically significant.

Across currency pairs, our illiquidity measures are positively correlated with the price impact estimates computed by MRW based on EBS data. This association is strong, however, only for the EBS-dominant currency pairs. For Reuters-dominant pairs, the EBS-based impact estimates are positive outliers: our estimates imply that these markets are in fact more liquid than the EBS-based impact measures would suggest.

Our results illuminate changes in liquidity over time. Over the Aprils of 2010, 2013, and 2016, bid-ask spreads exhibit a strong downward trend. As a measure of trading cost, the spread is most meaningful for smaller orders that can be completed in a single execution. Our illiquidity estimates correspond to price impact coefficients, which affect the cost of large orders that must be worked over time. Our estimates generally suggest that impact declined between 2010 and 2013 (an improvement in liquidity) and increased between 2013 and 2016 (a decline in liquidity). The net changes between 2010 and 2016, however, are mixed in direction.

The settlement data allow us to establish additional results about FX trading. In many cases, our results conform with prior beliefs about a market well-known for its opaqueness, light regulation and fragmented trading. Bid-ask spreads are smaller and liquidity better on highly traded currency pairs compared to others. Liquidity varies considerably over the 24-hour trading days with recognizable patterns for many currency pairs. But some results come with a little surprise. The typical trade size of FX swaps is many times as large as a typical spot transaction. Trade size clustering and price clustering is evident, although the latter appears to be on the decline. The average spot settlement size declines consistently from 2010 through 2016.

Our paper does not address the factors that might explain these cross-sectional and time series patterns. Some part of the explanations may be found in the CLS data itself, such as the number of active counterparties for each currency pair and the nature of their network structure. Other factors such as changes in bank regulation, capital requirements and restrictions on market-making may also play a role. As larger samples of settlement data become available, longer series of the liquidity measures can be formed, and the relations between liquidity and FX pricing can be studied over longer horizons and during specific macroeconomic and political events. We see these as a fertile area for future research.

References

Amihud, Yakov, 2002, Illiquidity and stock returns: Cross-section and time-series effects, *Journal of Financial Markets* 5, 31-56.

Bank for International Settlements, 1996, Settlement risk in foreign exchange transactions http://www.bis.org/cpmi/publ/d17.pdf.

Bank for International Settlements, 2015, Trienniel Central Bank Survey of Foreign Exchange and OTC Derivatives Market Activity: Reporting Guidelines for Turnover in April 2016.

Bank of Canada, 2010, The Canadian foreign exchange market: developments and opportunities, in Canadian Foreign Exchange Committee, ed.

Banti, Chiara, Kate Phylaktis, and Lucio Sarno, 2012, Global liquidity risk in the foreign exchange market, *Journal of International Money and Finance* 31, 267-291.

Barardehi, Yashar H, Dan Bernhardt, and Ryan J Davies, 2018, Trade-time measures of liquidity, *The Review of Financial Studies* 32, 126-179.

Berger, David, Alain Chaboud, and Erik Hjalmarsson, 2009, What drives volatility persistence in the foreign exchange market?, *Journal of Financial Economics* 94, 192-213.

Berger, David W., Alain P. Chaboud, Sergey V. Chernenko, Edward Howorka, and Jonathan H. Wright, 2008, Order flow and exchange rate dynamics in electronic brokerage system data, *Journal of International Economics* 75, 93-109.

Breedon, Francis, and Angelo Ranaldo, 2013, Intraday Patterns in FX Returns and Order Flow, *Journal of Money, Credit & Banking (Wiley-Blackwell)* 45, 953-965.

Breedon, Francis, Dagfinn Rime, and Paolo Vitale, 2016, Carry trades, order flow, and the forward bias puzzle, *Journal of Money, Credit and Banking* 48.

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.

Chaboud, Alain P., Sergey V. Chernenko, and Jonathan H. Wright, 2008, Trading Activity and Macroeconomic Announcements in High-Frequency Exchange Rate Data, *Journal of the European Economic Association* 6, 589-596.

Chakrabarty, Bidisha, Roberto Pascual, and Andriy Shkilko, 2015, Evaluating trade classification algorithms: Bulk volume classification versus the tick rule and the Lee-Ready algorithm, *Journal of Financial Markets* 25, 52-79.

Christie, William G., and Paul H. Schultz, 1994, Why do NASDAQ market makers avoid odd-eighth quotes?, *Journal of Finance* 49, 1813-1840.

CLS Bank, 2013, Best practices for third parties and third party service providers, (CLS Bank International).

CLS Group, 2014, CLS Settlement Methods Survey: Summary Report.

Easley, David, Marcos Lopez de Prado, and Maureen O'Hara, 2016, Discerning information from trade data, *Journal of Financial Economics* 120, 269-285.

Easley, David, and Maureen O'Hara, 1987, Price, trade size, and information in securities markets, *Journal of Financial Economics* 19, 69-90.

Evans, Martin D. D., and Richard K. Lyons, 2002a, Order Flow and Exchange Rate Dynamics, *The Journal of Political Economy* 110, 170-180.

Evans, Martin D. D., and Richard K. Lyons, 2002b, Time-varying liquidity in foreign exchange, *Journal of Monetary Economics* 49, 1025-1051.

Evans, Martin D. D., and Richard K. Lyons, 2008, How is macro news transmitted to exchange rates?, *Journal of Financial Economics* 88, 26-50.

Evans, Martin D. D., and Dagfinn Rime, 2016, Order flow information and spot rate dynamics, *Journal of International Money and Finance* 69, 45-68.

Financial Stability Oversight Council, 2012, 2012 Annual Report, (U.S. Department of the Treasury).

Fischer, Andreas M., and Angelo Ranaldo, 2011, Does FOMC news increase global FX trading?, *Journal of Banking & Finance* 35, 2965-2973.

FX Week, 2011, EC Suggests Corporate Threshold, CLS Gets Bank of Canada Plug, FX Week.

Gargano, Antonio, Steven J. Riddiough, and Lucio Sarno, 2018, The value of volume in foreign exchange (SSRN), https://ssrn.com/abstract=3019870.

Glosten, Lawrence R., and Paul R. Milgrom, 1985, Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71-100.

Hagströmer, Björn, and Albert J Menkveld, 2019, Information revelation in decentralized markets, *Journal of Finance* forthcoming.

Harris, Lawrence E., 1991, Stock price clustering and discreteness, *Review of Financial Studies* 4, 389-415.

Hasbrouck, Joel, 1999, Security bid/ask dynamics with discreteness and clustering: Simple strategies for modeling and estimation, *Journal of Financial Markets* 2, 1-28.

Kahn, Charles M., and William Roberds, 2001, The CLS bank: a solution to the risks of international payments settlement?, *Carnegie-Rochester Conference Series on Public Policy* 54, 191-226.

Karnaukh, Nina, Angelo Ranaldo, and Paul Söderlind, 2015, Understanding FX Liquidity, *Review of Financial Studies* 28, 3073-3108.

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.

Kos, Dino, and Richard M. Levich, 2016, Settlement risk in the global fx market: how much remains?

Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315-1336.

Lacker, Jeffrey M., 2001, The CLS bank: a solution to the risks of international payments settlement? A comment, *Carnegie-Rochester Conference Series on Public Policy* 54, 227-233.

Lee, Charles M. C., and Mark J. Ready, 1991, Inferring trade direction from intraday data, *Journal of Finance* 46, 733-746.

Levich, Richard M., and Frank Packer, 2017, Development and functioning of FX markets in Asia and the Pacific, *Financial Markets, Institutions & Instruments* 26, 3-58.

London FX Ltd., 2017, Electronic Brokers Web page at http://www.londonfx.co.uk/autobrok.html, Accessed on October 13, 2017.

Mancini, Loriano, 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.

O'Hara, Maureen, 2003, Presidential address: Liqudity and price discovery, *Journal of Finance* 58, 1335-1354.

O'Hara, Maureen, 2015, High frequency market microstructure, *Journal of Financial Economics* 116, 257-270.

Osler, Carol L, 2003, Currency orders and exchange rate dynamics: an explanation for the predictive success of technical analysis, *The Journal of Finance* 58, 1791-1819.

Osler, Carol L., 2005, Stop-loss orders and price cascades in currency markets, *Journal of International Money and Finance* 24, 219-241.

Osler, Carol, and Tanseli Savaser, 2011, Extreme returns: The case of currencies, *Journal of Banking & Finance* 35, 2868-2880.

Quandl, 2018, CLSRF|CLS FX Flow Data|Quandl Web page at https://www.quandl.com/databases/CLSRF/documentation/features, Accessed on December 31, 2018.

Ranaldo, Angelo, and Fabricius Somogyi, 2018, Heterogeneous Information Content of Global FX Trading (University of St.Gallen), https://ssrn.com/abstract=3263279.

Sarno, Lucio, and Mark P. Taylor, 2001, The microstructure of the foreign-exchange market: a selective survey of the literature (Princeton University).

Sopranzetti, Ben J, and Vinay Datar, 2002, Price clustering in foreign exchange spot markets, *Journal of Financial Markets* 5, 411-417.

Table 1. Counts and values of CLS settlements

The sample comprises all CLS settlements for which instructions were submitted within April of the indicated year. USD values are computed using average spot exchange rates over the April/year.

		N	%	USD amount	%
2010	Spot	7,265,894	91.8%	$13,574.272B	30.8%
	Outright Forward	258,584	3.3%	$2,804.908B	6.4%
	Near Leg FX Swap	135,910	1.7%	$10,857.259B	24.6%
	Far Leg FX Swap	135,508	1.7%	$10,817.973B	24.6%
	Other	84,399	1.1%	$5,655.755B	12.8%
	Option	38,499	0.5%	$346.309B	0.8%
	All	7,918,794	100.0%	$44,056.475B	100.0%
2013	Spot	12,490,361	91.8%	$14,195.135B	25.6%
	Outright Forward	475,567	3.5%	$1,961.552B	3.5%
	Near Leg FX Swap	236,143	1.7%	$16,814.548B	30.4%
	Far Leg FX Swap	238,582	1.8%	$16,808.103B	30.4%
	Other	138,208	1.0%	$5,090.657B	9.2%
	Option	22,642	0.2%	$476.827B	0.9%
	All	13,601,503	100.0%	$55,346.823B	100.0%
2016	Spot	9,937,550	91.2%	$10,388.432B	20.1%
	Outright Forward	327,590	3.0%	$1,630.189B	3.2%
	Near Leg FX Swap	245,209	2.2%	$17,827.466B	34.5%
	Far Leg FX Swap	242,964	2.2%	$17,755.223B	34.4%
	Other	110,939	1.0%	$3,618.706B	7.0%
	Option	37,541	0.3%	$388.837B	0.8%
	All	10,901,793	100.0%	$51,608.852B	100.0%

Table 2. Market share by currency, spot settlements

The sample is all spot settlements initiated on CLS during the Aprils of 2010, 2013, and 2016. For each settlement, both sides are valued in US dollars using the average exchange rates over the month. Currency totals therefore sum to twice the total turnover. The BIS shares exclude ILS, which the survey does not separately report.

	2010		2013		2016	
	CLS	BIS	CLS	BIS	CLS	BIS
USD	85.4%	84.4%	83.8%	87.0%	85.3%	90.5%
EUR	45.9%	49.1%	38.5%	38.8%	34.9%	33.9%
JPY	19.2%	21.3%	29.6%	31.5%	25.6%	25.8%
GBP	14.3%	15.1%	11.9%	11.7%	14.3%	13.8%
AUD	9.2%	7.9%	11.0%	10.1%	9.7%	9.3%
CAD	7.7%	5.5%	8.3%	4.8%	9.5%	6.8%
CHF	6.5%	6.5%	4.6%	4.3%	4.0%	3.7%
NZD	2.1%	1.5%	2.2%	2.0%	3.1%	2.6%
SEK	1.5%	1.3%	1.5%	1.4%	2.5%	2.2%
SGD	1.4%	1.1%	1.1%	1.1%	2.2%	1.8%
KRW	1.4%	1.5%	1.2%	1.0%	1.5%	1.9%
HKD	1.4%	1.3%	0.8%	1.1%	1.1%	1.5%
NOK	1.2%	0.9%	1.3%	1.1%	2.0%	1.9%
MXN	1.1%	1.3%	2.4%	2.9%	2.4%	2.8%
DKK	0.8%	0.3%	0.4%	0.4%	0.5%	0.5%
ZAR	0.7%	0.6%	1.1%	1.0%	1.2%	1.0%
ILS	0.1%		0.2%		0.2%	
All	200.0%	200.0%	200.0%	200.0%	200.0%	200.0%

Table 3. CLS settlement contra currencies

The sample is all CLS spot settlements in April 2016. A row summarizes all settlements in which the row-currency is involved in the exchange. Percentage entries in the row reflect the total USD equivalent value of the settlement, broken out by the other currency in the exchange. For example, of the total dollar value of all settlements involving the AUD, 1.4% occurred in the AUD/CAD pair.

	AUD	CAD	CHF	DKK	EUR	GBP	HKD	ILS	JPY	KRW	MXN	NOK	NZD	SEK	SGD	USD	ZAR
AUD		1.4	0.2		6.2	1.9			11.4			0.0	6.4	0.0	0.2	72.2	
CAD	1.4		0.2		4.0	1.3	0.0		1.9		0.1	0.0	0.3	0.0		90.7	
CHF	0.4	0.5			30.8	2.8			1.5			0.1	0.1	0.3		63.4	
DKK					76.0	0.7								0.0		23.3	
EUR	1.7	1.1	3.5	1.1		7.5	0.0	0.0	6.4		0.1	3.8	0.2	4.7	0.1	69.7	0.1
GBP	1.3	0.9	0.8	0.0	18.3		0.0		7.8		0.0	0.1	0.3	0.1	0.0	70.2	0.0
HKD		0.1			0.8	0.5			0.2							98.4	
ILS					3.0											97.0	
JPY	4.3	0.7	0.2		8.7	4.4	0.0					0.0	0.5	0.0	0.0	81.1	0.0
KRW																100.0	
MXN		0.3			1.5	0.0										98.2	
NOK	0.0	0.2	0.2		66.9	0.9			0.1				0.0	4.4		27.3	
NZD	20.3	0.9	0.1		2.4	1.6			4.1			0.0		0.0	0.1	70.3	
SEK	0.0	0.0	0.4	0.0	67.0	0.8			0.0			3.5	0.0			28.2	
SGD	1.0				1.2	0.3			0.5				0.2			96.9	
USD	8.2	10.1	2.9	0.1	28.5	11.8	1.3	0.2	24.3	1.8	2.8	0.6	2.5	0.8	2.5		1.4
ZAR					1.7	0.1			0.6							97.5	

Table 4. Clustering in settlement rates

The sample is CLS settlements for Aprils of 2010, 2013, and 2016. A pip is the traditional tick size in a currency pair, in units of the quote currency: 0.01 for the JPY, 0.0001 for all others (the second decimal place for JPY, the fourth for all others). The micro pip digit is the third place for the JPY and the fifth place for all others. 0 and 5 are tabulated separately; the remaining eight digits are summarized in "Other". Under the null hypothesis of equal probability of each digit, the frequencies on 0/5/Other would be 10%/10%/80%.

			Percent frequency of pip digit			Percent frequency of micro pip digit		
Instrument		N	0	5	Other	0	5	Other
All	2010	7,918,794	12.1	11.4	76.5	55.3	7.0	37.7
	2013	13,601,503	11.6	10.8	77.6	33.1	11.7	55.2
	2016	10,901,793	11.9	11.0	77.1	27.5	10.0	62.5
Spot	2010	7,265,894	11.8	11.3	76.8	56.4	7.0	36.6
	2013	12,490,361	11.5	10.8	77.7	33.3	11.9	54.7
	2016	9,937,550	11.8	11.0	77.2	27.7	10.1	62.2
Outright Forward	2010	258,584	11.2	10.8	78.0	27.0	8.7	64.3
	2013	475,567	10.7	10.3	79.0	18.3	9.7	72.0
	2016	327,590	10.8	10.4	78.8	15.3	9.6	75.1
Near Leg FX Swap	2010	135,910	16.6	12.2	71.2	54.8	5.1	40.1
	2013	236,143	15.1	11.5	73.4	41.2	7.9	50.9
	2016	245,209	14.0	11.2	74.8	32.0	7.7	60.3
Far Leg FX Swap	2010	135,508	13.8	11.9	74.3	45.9	7.1	47.0
	2013	238,582	11.6	10.6	77.8	31.0	9.5	59.4
	2016	242,964	12.0	10.8	77.2	25.7	9.3	65.1
Other	2010	84,399	14.5	12.1	73.4	57.7	5.1	37.2
	2013	138,208	11.6	10.8	77.6	41.3	8.6	50.1
	2016	110,939	12.0	11.1	76.9	25.3	8.2	66.5
Option	2010	38,499	42.9	21.5	35.6	77.6	3.5	18.9
	2013	22,642	54.1	24.6	21.3	85.8	1.5	12.6
	2016	37,541	38.8	18.6	42.6	63.8	4.1	32.1

Table 5. Summary statistics on Olsen quotes

The sample is Olsen bids and asks in April of the indicated year. Within each ten-second window Olsen reports the first bid and ask. For each such observation, the spread is the ask less the bid (in units of the quote currency, scaled by 10^4), the bid-ask midpoint is the average of the bid and ask (in units of the quote currency), and the relative spread is the spread divided by the midpoint (scaled to basis points). Table entries are medians.

	Spread × 10^4			Bid-ask midpoint			Relative Spread × 10^4 (bp)		
	2010	2013	2016	2010	2013	2016	2010	2013	2016
AUD/JPY	300.000	210.000	100.000	86.7400	101.6040	84.1300	3.547	2.046	1.196
AUD/USD	2.600	1.700	0.700	0.9268	1.0372	0.7660	2.833	1.628	0.925
EUR/CHF	3.000	2.200	1.200	1.4337	1.2175	1.0920	2.087	1.805	1.099
EUR/GBP	2.000	1.600	0.700	0.8767	0.8519	0.7946	2.260	1.878	0.886
EUR/JPY	300.000	220.000	90.000	125.4750	128.4305	123.8005	2.357	1.708	0.732
EUR/USD	2.000	1.400	0.500	1.3424	1.3046	1.1351	1.474	1.073	0.444
GBP/JPY	430.000	320.000	170.000	143.3450	151.2900	156.5985	2.984	2.165	1.098
GBP/USD	2.500	2.100	1.000	1.5339	1.5295	1.4270	1.624	1.360	0.697
NZD/USD	4.000	2.300	1.000	0.7114	0.8459	0.6886	5.587	2.730	1.459
USD/CAD	3.500	2.000	1.100	1.0033	1.0171	1.2813	3.450	1.949	0.854
USD/CHF	2.800	1.800	1.000	1.0684	0.9349	0.9634	2.638	1.931	1.040
USD/JPY	200.000	130.000	60.000	93.4150	98.2725	109.3655	2.125	1.309	0.540
USD/MXN	54.000	55.000	56.000	12.2210	12.1937	17.4647	4.418	4.531	3.180

Table 6. Imputed submission delays

The sample is CLS spot settlements merged with Olsen quotes during the month of April of the indicated year. Table entries indicate the distribution of imputed delays in settlement times relative to Olsen bid and ask quotes. For each spot settlement, we look backwards until we find a match region (variously defined) that contains the settlement price. Using the "bid and ask" criterion, a settlement at time t priced at p_t is considered to be matched with delay s if

$$\min_{\tau\in[t-s,t]} bid_\tau \leq p_t \leq \max_{\tau\in[t-s,t]} ask_\tau$$

Under the second definition, the settlement is matched at time s if

$$\min_{\tau\in[t-s,t]} bid_\tau - spread \leq p_t \leq \max_{\tau\in[t-s,t]} ask_\tau + spread$$

Where spread is the median spread for the currency pair estimated over the month. Under the third definition, a match is inferred if

$$\min_{\tau\in[t-s,t]} bid_\tau - 5\times spread \leq p_t \leq \max_{\tau\in[t-s,t]} ask_\tau + 5\times spread$$

The last three match regions are defined analogously but using the pip (tick size) instead of the spread. The pip size is 0.0001 for all quote currencies except JPY (for which the pip size is 0.01.) "NM" denotes "not matched (within ten minutes).

Match region		0-5s	5s-10s	10s-20s	20s-30s	30s-1m	1m-10m	NM
Bid and ask	2010	20.1%	18.0%	13.6%	8.0%	12.4%	21.2%	6.6%
	2013	20.6%	19.4%	15.7%	8.6%	12.9%	16.7%	6.2%
	2016	9.6%	9.7%	14.2%	11.2%	21.9%	25.2%	8.2%
$\pm spread$	2010	36.4%	31.8%	7.3%	3.1%	5.0%	11.7%	4.8%
	2013	36.0%	35.1%	7.7%	3.6%	5.6%	8.5%	3.5%
	2016	22.0%	21.8%	11.3%	8.1%	14.1%	16.8%	6.0%
$\pm 5\times spread$	2010	48.4%	42.1%	3.5%	0.4%	0.7%	2.7%	2.1%
	2013	46.3%	45.3%	3.0%	0.7%	1.1%	2.2%	1.4%
	2016	40.5%	39.2%	4.3%	2.1%	3.5%	6.7%	3.7%
$\pm pip$	2010	29.3%	25.8%	10.1%	5.1%	8.0%	16.0%	5.6%
	2013	32.4%	31.4%	9.5%	4.8%	7.3%	10.6%	4.1%
	2016	24.7%	24.4%	10.5%	7.2%	12.3%	15.2%	5.6%
$\pm 2\times pip$	2010	34.9%	30.5%	7.9%	3.5%	5.6%	12.6%	5.0%
	2013	37.8%	37.0%	6.7%	3.0%	4.7%	7.6%	3.2%
	2016	33.0%	32.3%	7.4%	4.4%	7.4%	10.7%	4.7%
$\pm 10\times pip$	2010	47.5%	41.3%	3.7%	0.6%	1.0%	3.4%	2.5%
	2013	46.8%	45.8%	2.8%	0.5%	0.9%	1.8%	1.3%
	2016	46.7%	44.9%	2.2%	0.6%	1.1%	2.4%	2.0%

Table 7. Volume-based illiquidity ratios, I^{Volume}

Intervals are constructed to contain clearing sequences of \$100M US equivalent. For sequence k, the illiquidity ratio is $I_k = |\Delta p_k| / (Settlement\ Volume)_k$, where Δp_k is the log price change over the interval, and $(Settlement\ Volume)_k$ is the cumulative settlement volume over the interval, scaled to $bp/\$1M$. I_k is computed for all sequences in the Aprils of 2010, 2013, and 2016, and the observations are winsorized at 95%. The first three columns contain means, referred to as I^{Volume} in the text, and (in parentheses) standard errors. The remaining columns indicate the direction and statistical significance of changes between the indicated years (based on a t-test for differences in means): −−/++ denote decreases/increases with a p-value of 0.01 or better; −/+, with a p-value of 0.05 or better; NS, no significant change.

	Levels			Changes		
	2010	2013	2016	2010-2013	2013-2016	2010-2016
AUD/JPY	0.1386	0.1119	0.1120	−−	NS	−−
	(0.0048)	(0.0028)	(0.0028)			
AUD/USD	0.0280	0.0169	0.0350	−−	++	++
	(0.0002)	(0.0001)	(0.0003)			
EUR/CHF	0.0137	0.0165	0.0274	++	++	++
	(0.0002)	(0.0002)	(0.0006)			
EUR/GBP	0.0283	0.0269	0.0374	−	++	++
	(0.0004)	(0.0004)	(0.0006)			
EUR/JPY	0.0398	0.0465	0.0503	++	++	++
	(0.0004)	(0.0004)	(0.0009)			
EUR/USD	0.0103	0.0089	0.0117	−−	++	++
	(<0.0001)	(<0.0001)	(<0.0001)			
GBP/JPY	0.1097	0.1412	0.0909	++	−−	−−
	(0.0031)	(0.0057)	(0.0025)			
GBP/USD	0.0224	0.0148	0.0227	−−	++	NS
	(0.0002)	(0.0001)	(0.0002)			
NZD/USD	0.0737	0.0480	0.0639	−−	++	−−
	(0.0015)	(0.0009)	(0.0011)			
USD/CAD	0.0285	0.0154	0.0275	−−	++	−
	(0.0003)	(0.0001)	(0.0002)			
USD/CHF	0.0352	0.0319	0.0400	−−	++	++
	(0.0005)	(0.0005)	(0.0007)			
USD/JPY	0.0181	0.0168	0.0156	−−	−−	−−
	(0.0001)	(<0.0001)	(0.0001)			
USD/MXN	0.0470	0.0324	0.0536	−−	++	++
	(0.0011)	(0.0005)	(0.0009)			

Table 8. Illiquidity ratios based on fixed-time intervals, I^{Time}

Fixed interval illiquidity ratios based on two-minute windows, in units of bp/$1M (USD), winsorized at 95%. Means, denoted I^{Time} in the text, and (in parentheses) standard errors are reported for the ratios in the Aprils of 2010, 2013, and 2016. The remaining columns indicate the direction and statistical significance of changes between the indicated years (based on a t-test for differences in means): −−/++ denote decreases/increases with a p-value of 0.01 or better; −/+, with a p-value of 0.05 or better; NS, no significant change.

	Levels			Changes		
	2010	2013	2016	2010-2013	2013-2016	2010-2016
AUD/JPY	2.6932	1.6495	1.7379	−−	++	−−
	(0.0151)	(0.0113)	(0.0139)			
AUD/USD	0.1364	0.0472	0.1058	−−	++	−−
	(0.0014)	(0.0004)	(0.0009)			
EUR/CHF	0.2629	0.2142	0.4830	−−	++	++
	(0.0020)	(0.0016)	(0.0031)			
EUR/GBP	0.3721	0.3586	0.4756	−−	++	++
	(0.0031)	(0.0027)	(0.0038)			
EUR/JPY	0.2606	0.2389	0.4750	−−	++	++
	(0.0024)	(0.0022)	(0.0040)			
EUR/USD	0.0271	0.0186	0.0260	−−	++	−−
	(0.0003)	(0.0001)	(0.0002)			
GBP/JPY	1.8205	2.0576	1.4834	++	−−	−−
	(0.0094)	(0.0095)	(0.0089)			
GBP/USD	0.1132	0.0651	0.0930	−−	++	−−
	(0.0011)	(0.0006)	(0.0009)			
NZD/USD	1.1234	0.4325	0.4949	−−	++	−−
	(0.0093)	(0.0037)	(0.0045)			
USD/CAD	0.2974	0.0990	0.1225	−−	++	−−
	(0.0029)	(0.0009)	(0.0012)			
USD/CHF	0.3922	0.3006	0.4052	−−	++	++
	(0.0035)	(0.0025)	(0.0032)			
USD/JPY	0.0521	0.0256	0.0293	−−	++	−−
	(0.0005)	(0.0002)	(0.0002)			
USD/MXN	0.8396	0.4869	0.8619	−−	++	++
	(0.0044)	(0.0035)	(0.0064)	−−	++	−−

Table 9. Impact coefficients from volume-based intervals and bulk-volume classification

Intervals are constructed to contain settlement sequences of $100M US. The regression specification is $\Delta p_k = \lambda^{BVC}\hat{x}_k + u_k$ where Δp_k is the log price change over interval k; x_k is the signed order flow imputed by bulk volume classification: $\hat{x}_k = (settlement\ volume)_k \times \left(2\Phi\left(\Delta p_k/\sigma_{\Delta p}\right) - 1\right)$. λ^{BVC} is estimated via OLS, in units of $bp/\$1M$ US. Estimates and standard errors are reported for the Aprils of 2010, 2013, and 2016. The remaining columns indicate the direction and statistical significance of changes between the indicated years (based on a t-test for differences in means): −−/++ denote decreases/increases with a p-value of 0.01 or better; −/+, with a p-value of 0.05 or better; NS, no significant change.

	Levels			Changes		
	2010	2013	2016	2010-2013	2013-2016	2010-2016
AUD/JPY	0.3096	0.2566	0.2840	−−	++	−−
	(0.0046)	(0.0034)	(0.0036)			
AUD/USD	0.0667	0.0415	0.0881	−−	++	++
	(0.0003)	(0.0002)	(0.0003)			
EUR/CHF	0.0451	0.0391	0.0586	−−	++	++
	(0.0004)	(0.0004)	(0.0006)			
EUR/GBP	0.0715	0.0666	0.0893	−−	++	++
	(0.0004)	(0.0005)	(0.0005)			
EUR/JPY	0.0967	0.1107	0.1329	++	++	++
	(0.0006)	(0.0005)	(0.0010)			
EUR/USD	0.0288	0.0243	0.0303	−−	++	++
	(<0.0001)	(<0.0001)	(<0.0001)			
GBP/JPY	0.2437	0.3911	0.2470	++	−−	NS
	(0.0048)	(0.0065)	(0.0051)			
GBP/USD	0.0563	0.0367	0.0559	−−	++	NS
	(0.0002)	(0.0002)	(0.0002)			
NZD/USD	0.1671	0.1105	0.1527	−−	++	−−
	(0.0013)	(0.0011)	(0.0011)			
USD/CAD	0.0703	0.0388	0.0686	−−	++	−−
	(0.0003)	(0.0003)	(0.0003)			
USD/CHF	0.0886	0.0780	0.0933	−−	++	++
	(0.0005)	(0.0006)	(0.0007)			
USD/JPY	0.0461	0.0437	0.0456	−−	++	NS
	(0.0001)	(0.0001)	(0.0001)			
USD/MXN	0.1256	0.0851	0.1344	−−	++	++
	(0.0012)	(0.0008)	(0.0009)			

Table 10. Correlations among liquidity estimates

I^{Volume} is the illiquidity ratio based on settlement-volume intervals; I^{Time} is illiquidity based on fixed two-minute intervals; λ^{BVC} is the OLS estimate of the impact coefficient in a regression of price changes against net order flow imputed with a bulk volume classification rule. Measures are estimated for thirteen currency pairs and three one-month samples (the Aprils of 2010, 2013 and 2016). BIS turnover is measured in $B US per day. The p-values of all tests against the null (equal to zero) are below 0.0001.

Panel A. Pearson Correlations

	$\log(I^{SV})$	$\log(I^{Fixed})$	$\log(\lambda^{BVC})$	$\log(BIS\ turnover)$
$\log(I^{SV})$	1.000	0.909	0.995	-0.783
$\log(I^{Fixed})$	0.909	1.000	0.905	-0.946
$\log(\lambda^{BVC})$	0.995	0.905	1.000	-0.778
log(Volume)	-0.783	-0.946	-0.778	1.000

Panel B. Pearson correlations, partial, controlling on $\log(BIS\ turnover)$

	$\log(I^{SV})$	$\log(I^{Fixed})$	$\log(\lambda^{BVC})$
$\log(I^{SV})$	1.000	0.706	0.986
$\log(I^{Fixed})$	0.706	1.000	0.709
$\log(\lambda^{BVC})$	0.986	0.709	1.000

Figure 1. CLS spot settlement quantities, April 2016

Histogram and CDF of settlement quantities measured in units of base currency.

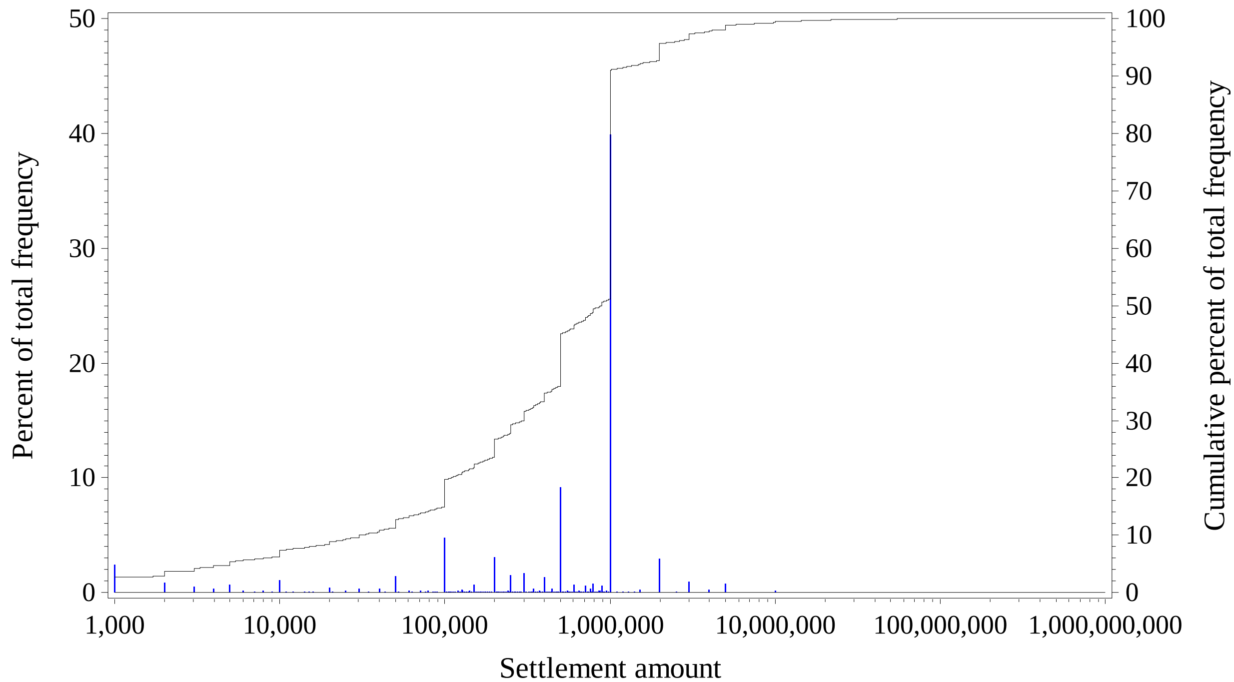

Figure 2. Median relative bid-ask spreads by year and currency pair.

Olsen bids and asks are collected at ten-second intervals. For each observed bid-ask pair, the relative spread in basis points is $\frac{ask-bid}{(ask+bid)/2} \times 10{,}000$. Figures depict medians for the indicated year and currency pair.

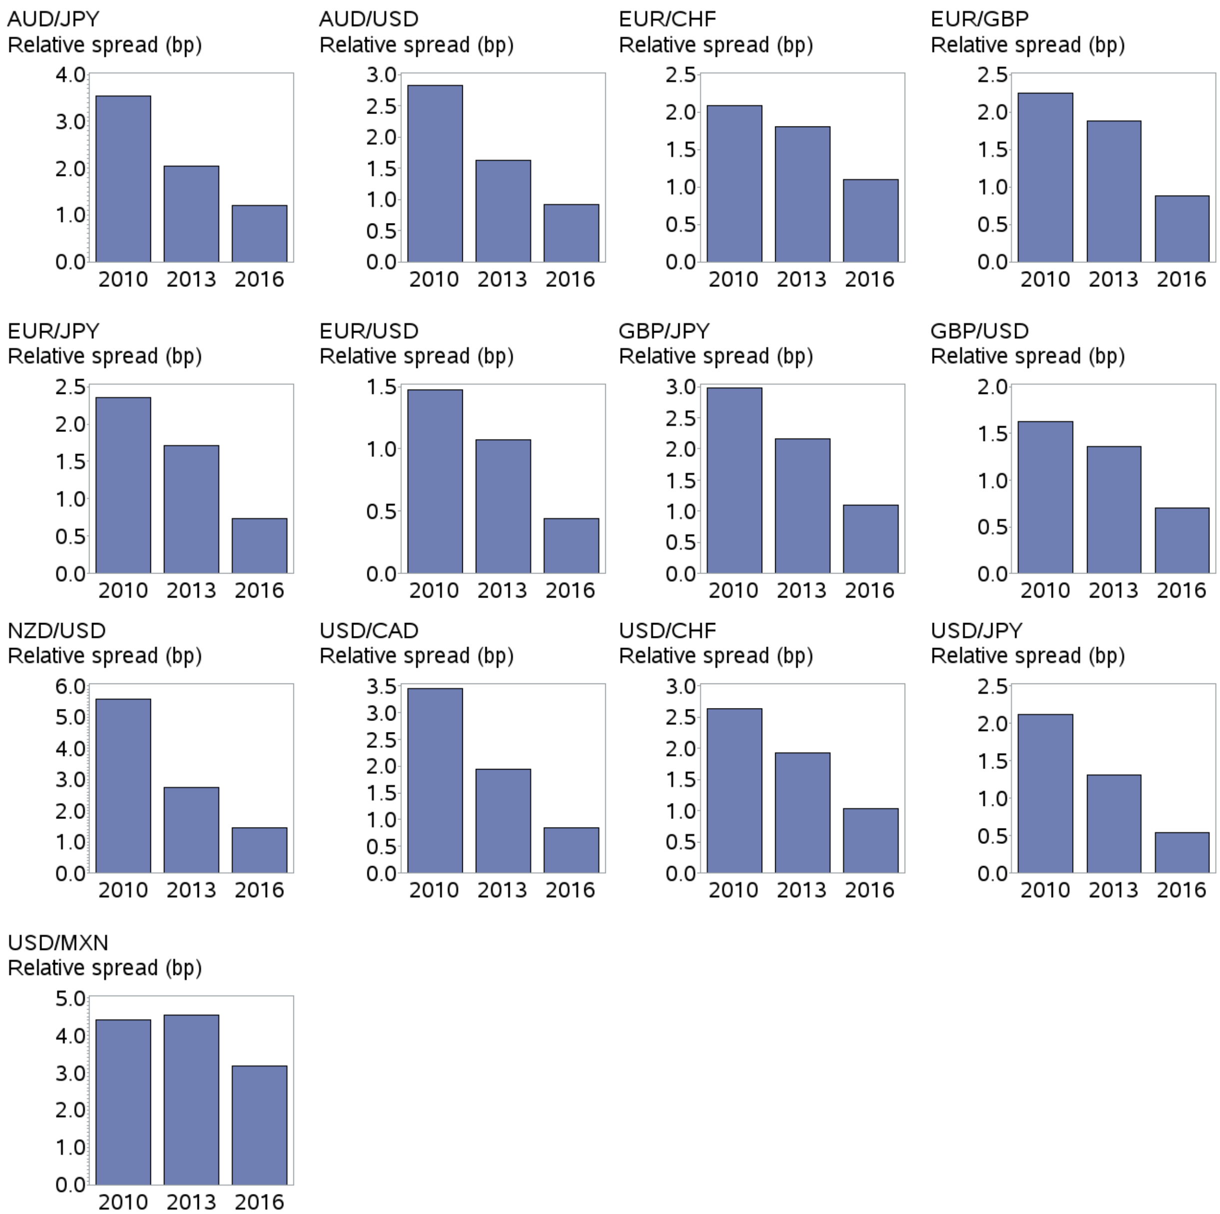

Figure 3. Settlement rates and market rates, EUR/USD, April 17, 2013

CLS settlement accept times and reported rates are shown as gray dots. Bid-ask midpoints of Olsen quotes are shown as a black line.

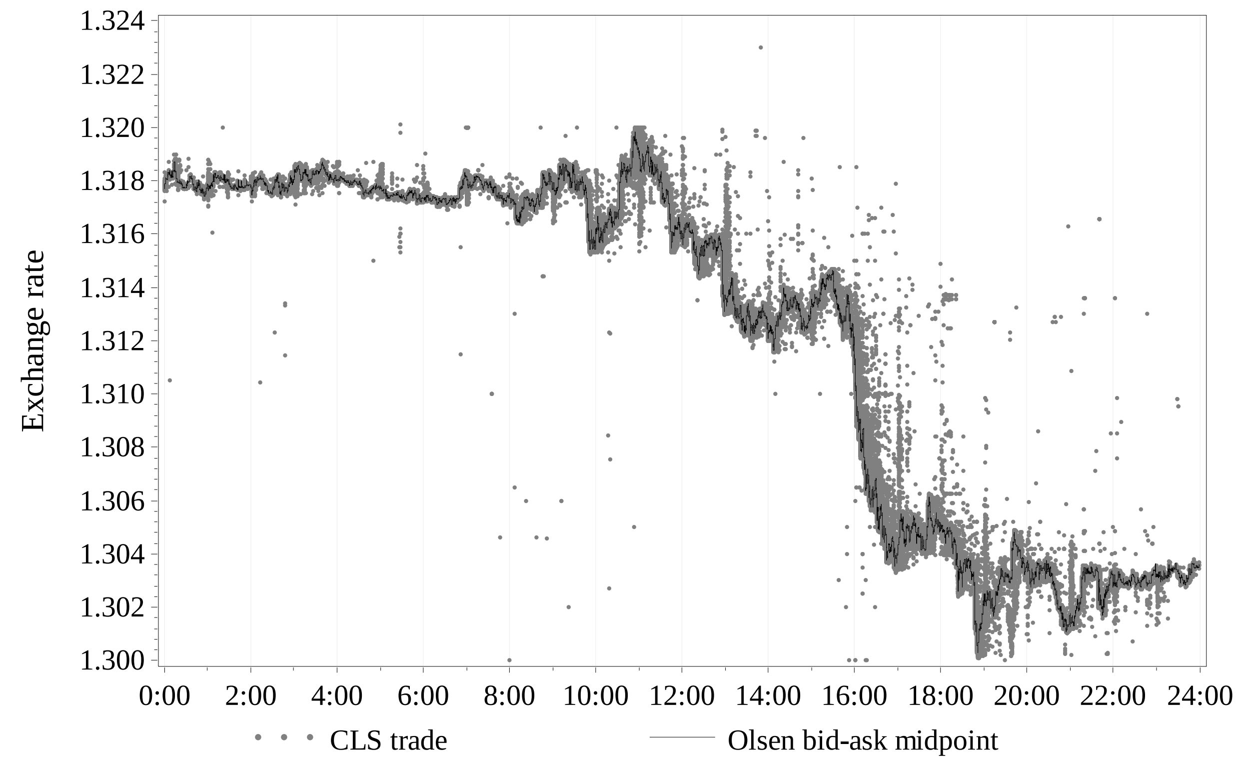

Figure 4. Settlement volume illiquidity ratios

Time-scaled illiquidity ratios based on $100M US sequences, in units of bp/$1M (USD). Means and 5% confidence intervals (cf. Table 11).

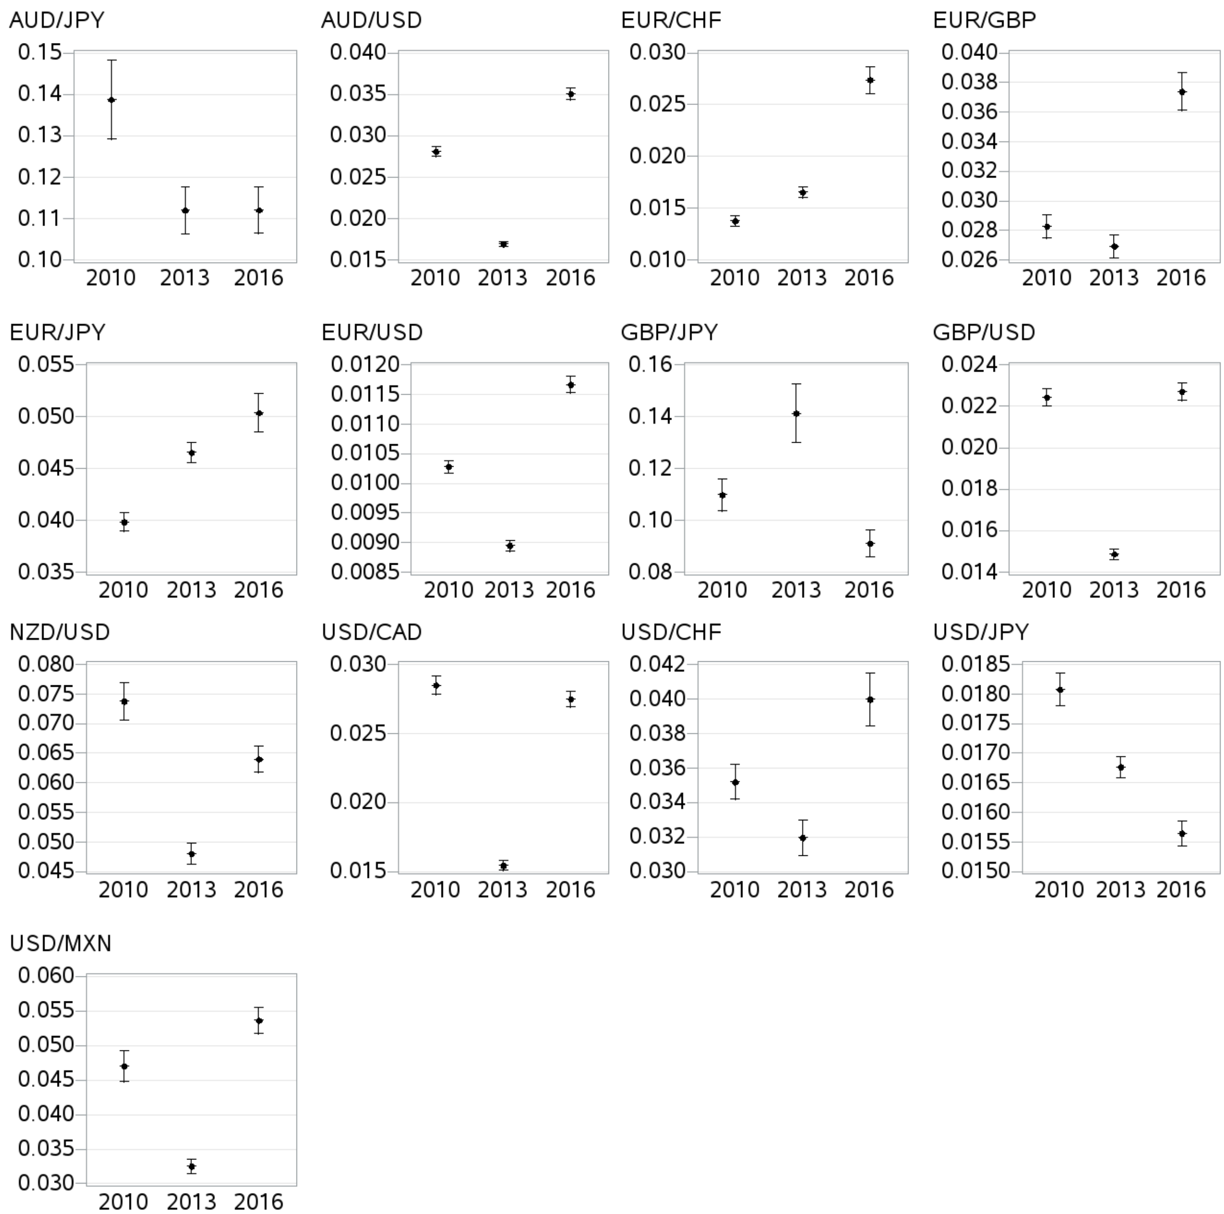

Figure 5. Fixed-interval illiquidity ratios

Illiquidity ratios based on two-minute intervals, in units of bp/$100M (USD). Means and 5% confidence intervals (cf. Table 12).

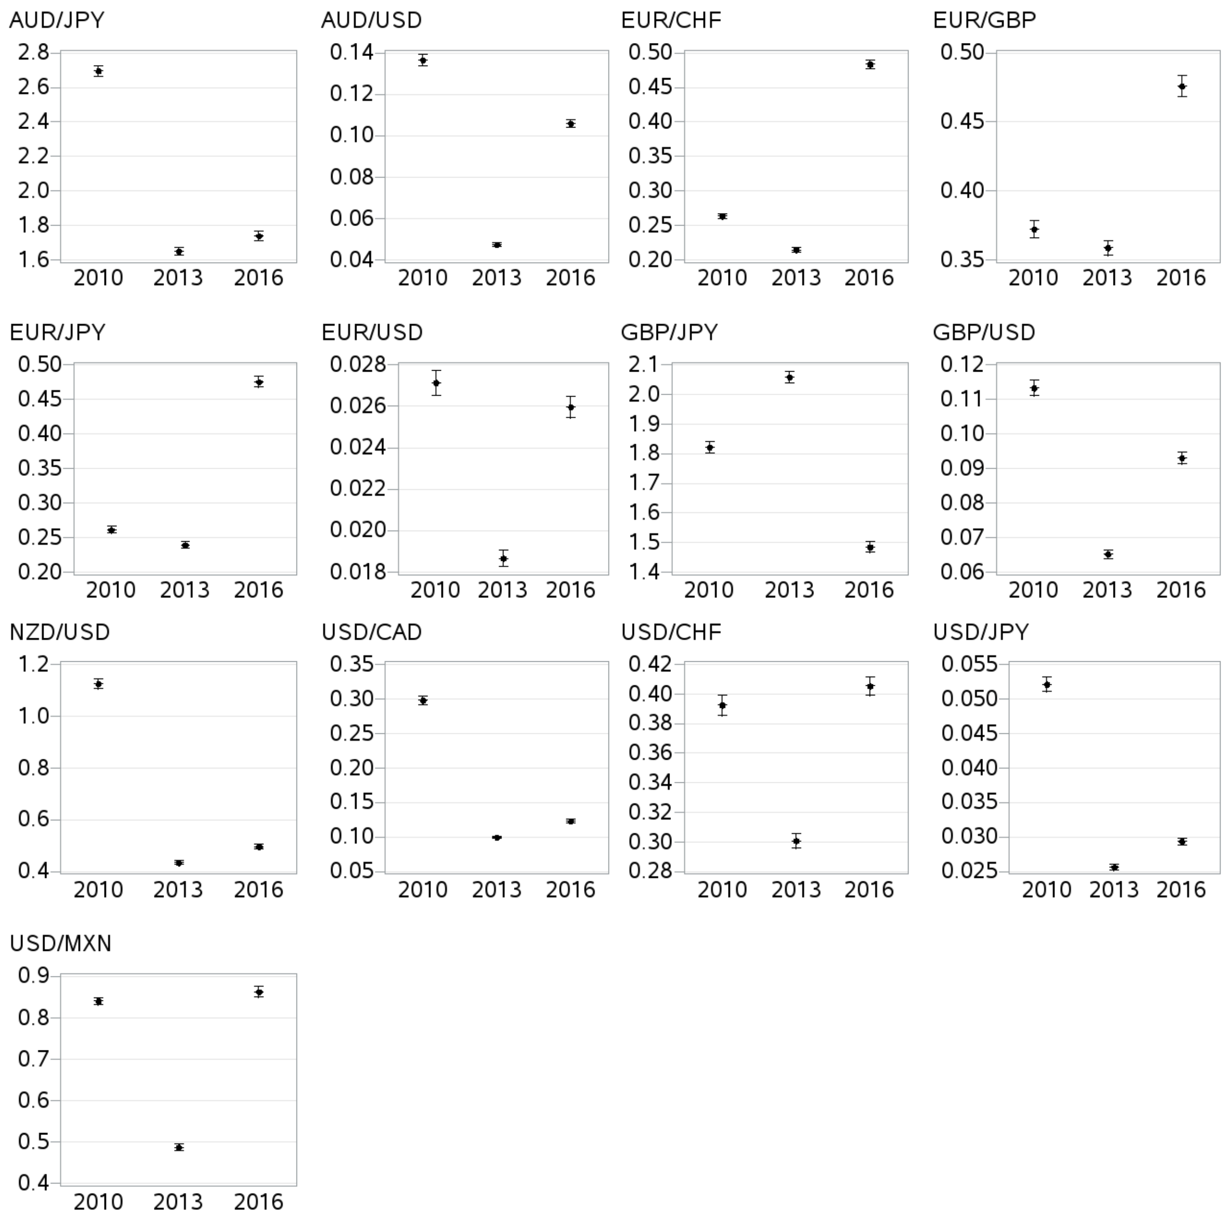

Figure 6. Bulk-volume classification impact coefficients

Intervals are constructed to contain settlement sequences of \$100M US. The regression specification is $\Delta p_k = \lambda^{BVC}\hat{x}_k + u_k$ where Δp_k is the log price change over interval k; x_k is the signed order flow imputed by bulk volume classification: $\hat{x}_k = (settlement\ volume)_k \times \left(2\Phi\left(\Delta p_k/\sigma_{\Delta p}\right) - 1\right)$. λ^{BVC} is estimated via OLS, in units of $bp\ per$ \$100M US. Vertical bars demarcate lower and upper confidence intervals (at $p = 0.05$).

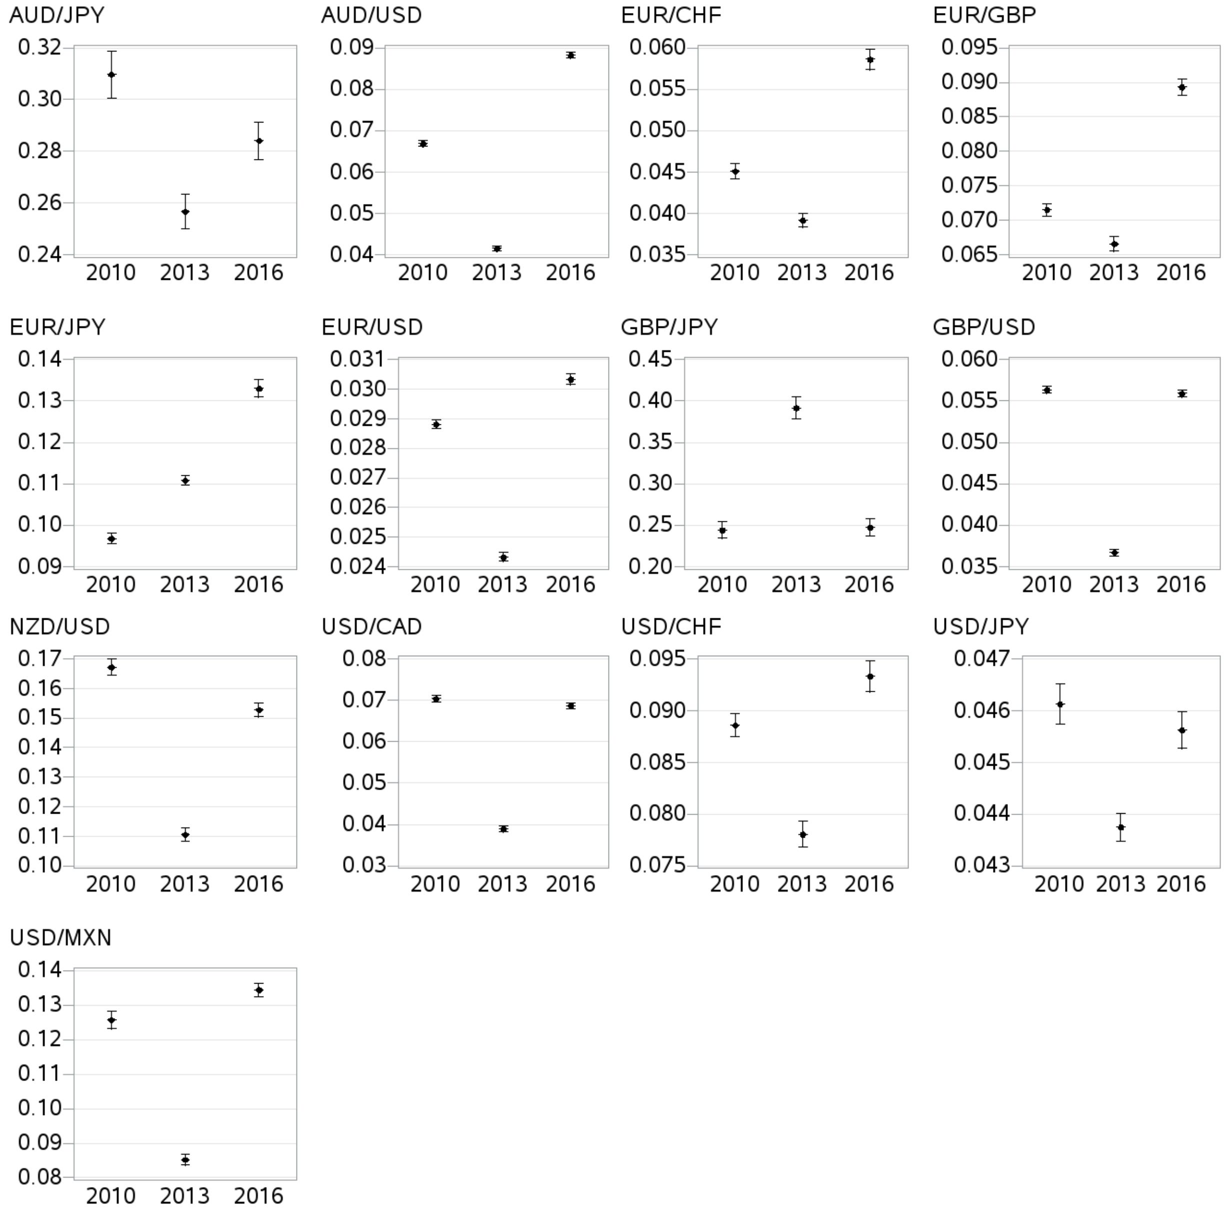

Figure 7. Comparison of settlement-based and EBS-based impact estimates

Log-log scatterplot of EBS estimated price impact coefficients (on the vertical axis) vs. I^{Volume} impact estimates (on the horizontal axis). The I^{Volume} impact estimates are means based on CLS settlement data for April 2010; the EBS price impact coefficients are the means reported in Mancini, Ranaldo and Wrampelmeyer (2013), internet appendix IA.III. EBS-dominant pairs are plotted as circles; Reuters-dominant are plotted as triangles.