Foreign Exchange Fixings and Returns Around the Clock

Ingomar Krohn[*] Philippe Mueller[†] Paul Whelan[‡]

Link to the latest version

ABSTRACT

This paper documents a new stylised fact in foreign exchange markets: intraday currency returns display prolonged reversals around the major benchmark fixings, characterised by an appreciation of the U.S. dollar pre-fixing and a depreciation thereafter. Tracing returns around the clock, the major fixing during Asian trading hours (Tokyo) and two major fixings during European and U.S. hours (Frankfurt and London) generate a distinct '*W*' shaped return pattern over the 24-hour trading day. On either side of the reversal, price drifts persist for hours; moreover, they are a systematic feature of the data being present every day of the week, month of the year, and during each of the 20 years in our sample. We argue these findings require two ingredients (*i*) a structural demand for dollar immediacy at local currency fixing times; and (*ii*) pre-fix hedging risk management practices by financial intermediaries. Consistent with this conjecture, we show our findings are amplified in states of high anticipated volatility, low liquidity, and that 'arbitrageurs' can exploit these patterns after taking transaction costs into account.

Keywords: foreign-exchange, intraday and overnight returns, high-frequency data, corporates, pension funds, insurance companies, intermediation.

This Version: December 2019.

We would like to thank Daniel Andrei, Patrick Augustin, Nick Baltas, Nina Boyarchenko, Gino Cenedese, Pasquale Della Corte, Michael Fleming, Patrick Green, Farouk Jivjav, Christian Julliard, Nina Karnaukh, Sven Klingler, Paul Meggyesi, Albert Menkveld, Tarun Ramadorai, Dagfinn Rime, Elvira Sojli, Gyuri Venter, Tony Zhang, and seminar and conference participants at various universities and financial institutions. The authors gratefully acknowledge financial support from the Canadian Derivatives Institute. Paul Whelan gratefully acknowledges financial support from the Danish Council for Independent Research grant no. 9037-00105B. Ingomar Krohn gratefully acknowledges financial support from the Economic and Social Research Council, grant no. 1500668. The usual disclaimer applies.

[*]Copenhagen Business School, Solbjerg Pl. 3, 2000 Frederiksberg, DK, Email: ik.fi@cbs.dk

[†]Warwick Business School, Coventry, CV4 7AL, UK, Email: Philippe.Mueller@wbs.ac.uk

[‡]Copenhagen Business School, Solbjerg Pl. 3, 2000 Frederiksberg, DK, Email: pawh.fi@cbs.dk

Empirical work on exchange rates and currency markets is typically based on daily (or monthly) currency returns that are measured at the London fix (i.e., at 4:00 p.m. London time).[1] However, the foreign exchange market trades continuously on a 24-hour decentralised basis between participants spread across the globe and in different time zones, from 5:00 a.m. Sydney time on Monday morning until 5:00 p.m. New York time on Friday afternoon. This means that the FX market continues to trade through the London fix even though the rates measured at the London fix are often used as "closing prices." Of the global estimated \$6.6 trillion daily turnover, around \$1.2 trillion is traded during U.S. hours, roughly \$2.4 trillion during London hours, and the remaining volume of \$3.0 trillion is distributed across a large number of local markets (see BIS (2016)).

In this paper, we study returns for the G9 currencies at high-frequency intervals and around the clock with a particular focus on the major currency fixes around the world and we document a set of novel and robust stylised facts with respect to the intraday return patterns. First, we show that the U.S. dollar systematically appreciates against all currencies ahead of the three major currency fixes in Tokyo, Frankfurt and London, while displaying a strong reversal pattern thereafter. This means that throughout the day the price of the U.S. dollar reaches a local maximum at the Tokyo, ECB and London fix, respectively. If we take the view of an investor that goes long foreign currencies, the value of the foreign currency portfolio exhibits a significant V-shape around the currency fixes that take place at 10:00 a.m. Tokyo time, 2:15 p.m. Frankfurt time and 4:00 p.m. London time. Aggregated over the day and across currencies, the foreign exchange portfolio exhibits a distinct W-shaped pattern over a 24 hour period.[2]

Second, we document that the pattern is very robust over time, across currencies as well as across different data sets. Furthermore, the pattern is not driven by day of the week or month of the year effects and is present in all years of our sample, i.e., it is a pervasive feature of the high frequency data. At the same time, we show that the pattern is statistical and economically stronger during periods of high ex-ante volatility as measured by the implied volatility extracted from currency options or the equity VIX. Moreover, in periods of low liquidity (as proxied by higher than usual bid-ask spreads) the patterns are also amplified. Finally, we show that the large

[1] See, e.g., Thomson Reuters (2017).

[2] The ECB fix and the London fix are less than three hours apart, while there is a depreciation of the U.S. dollar after the ECB fix and a subsequent appreciation ahead of the London fix the overall pattern is dominated by the pre-ECB fix appreciation of the U.S. dollar and the post-London fix depreciation.

intraday swings are not easy to exploit once transaction costs are accounted for, although we argue that large players in the market that can trade at tight spreads should be able to take advantage of the predictable patterns. Overall, the evidence suggests that dealer immediacy models and limited risk bearing capacity are a potential explanation for our findings.

To establish the empirical facts we construct a panel of 5-minute spot returns around the clock using high-frequency data on a set of nine currencies vis-à-vis the U.S. dollar: the Australian dollar (AUD), the Canadian dollar (CAD), the euro (EUR), the British pound (GBP), the Japanese yen (JPY), the New Zealand dollar (NZD), the Norwegian krone (NOK), the Swedish krona (SEK), and the Swiss franc (CHF). Our sample period spans January 1999 to December 2018 during which these pairs cover approximately 75% of the total daily turnover in the foreign exchange market (see BIS (2019)).

Using the intraday panel of spot FX returns, we define different windows based on the three major fixes as well as the market opening time in Europe. As the end (and beginning) of the day we define 5:00 p.m. New York time—this is the time when bid-ask spreads in the FX market are the highest and trading volume is the lowest. At the same time, it is the time of day when major banks settle their positions and calculate their risk positions for the end of the day. However, our choice of time to measure the end of the trading day deviates from the London fixing time that is generally used to calculate closing prices in foreign exchange markets. Measured in local time, the major fixes happen in the morning in Tokyo, just after lunch in Frankfurt and towards the end of the trading day in London. As the dollar reaches a local maximum at each of the fixes, it is obvious that the patterns we detect are not a result of a particular time of day but are directly related to trading activity at and around the fixes. We also show that in order to explain the patterns over a 24 hour period it is necessary to move beyond a story focusing on intraday vs. overnight returns as is regularly done in the existing literature—it is simply not possible to consistently tie the W-shaped return pattern to market opening and closing times alone. For example, the U.S. dollar appreciates until 10:00 a.m. local time (9:00 p.m. ET) before the Tokyo fix whereas before the ECB fix the U.S. dollar appreciates until 2:15 p.m. local time or 8:15 a.m. ET.

The intraday swings are both statistically highly significant and economically large. A portfolio that goes long all G9 currencies exhibits daily average swings of around 2 basis points (or over 5% annualized). This may seem small but given the very tight bid-ask spreads in the FX market and the extremely high trading volume this translates into significant swings in dollar terms. The high statistical significance and the remarkable persistence also indicates that the swings are systematic and not driven by a number of outliers. There is a large literature highlighting the importance of the Tokyo and London fixes and there is a body of evidence pointing to unsophisticated hedgers and speculators trading immense quantities of currencies at the fixing price. Moreover, there is the general notion that hedgers (i.e., corporates) usually sell U.S. dollars at the fix and buy local (foreign) currencies. As such it may seem counterintuitive that they appear to be able to transact exactly at the local peak, i.e., when the U.S. dollar is highest. This suggests that the hedgers' supply of U.S. dollars at the fix must be outstripped by a demand for U.S. dollars, also emanating from seemingly unsophisticated investors who prefer to trade at the fix. We argue that this demand is driven by liability-driven investors that purchase U.S. dollar denominated securities such as Treasuries on an ongoing basis.

We link our results to an existing literature on financial intermediation and limited risk-bearing capacity in financial markets. In particular, the currency patterns we document appear to be the high-frequency analogue to the price reversals observed around Treasury auctions over the course of multiple days and first highlighted by Lou, Yan, and Zhang (2013). Liquidity providers in FX markets are regularly faced with excess dollar demand at the fixes although the exact magnitude is unknown. Thus, they have to trade taking uncertainty about the net demand for U.S. dollars into account and they face a trade-off between arbitraging the difference between the pre-fix price and the expected price of the U.S. dollar at the fix and hedging the uncertainty about the expected price at the fix. Consistent with this line of argument we show that the reversal patterns are magnified during periods of high volatility and low liquidity.

Finally, we assess the profitability of simple intraday trading strategies that require rebalancing of positions around the fix and in line with the windows we define in order to exploit the predictable patterns we observe. Following existing research (e.g., Menkhoff, Sarno, Schmeling, and Schrimpf (2012), Lustig, Roussanov, and Verdelhan (2011)) we use the bid-ask spread as a proxy for transaction costs and we consider different magnitudes of the spread that have been argued to accurately reflect trading costs faced by dealers in the FX market (Gilmore and Hayashi (2011), Gargano, Riddiough, and Sarno (2018)). Even though intraday trading strategies require a very high turnover, we document that for some currencies (i.e., the EUR, GBP and JPY among others) it is possible to generate positive net returns when "trading the W" assuming that traders get slightly tighter (and more realistic) spreads compared to the quoted prices in our benchmark data base. Our results suggest that traders may be able to achieve annualised Sharpe Ratios between 0.5 and 0.8 from trading the longer windows we define. However, it seems almost impossible to reap short-term profits from "trading the fix", i.e., from taking positions for a short period of time around the fix. Despite the fact that reversals around the fix are quite large, transaction costs eat away potential profits. Or put differently, around the fix dealers set bid-ask spreads at levels that make it impossible to profitably trade but once traders are willing to take on more risk and hold the currency positions for the full window they are adequately compensated and are able to earn excess returns.

The paper is organized as follows: In Section I we describe the data and in Section II we discuss the relevance of the currency fixes. In Section III we present the main empirical stylised facts while Section IV presents some robustness tests and additional results. In Section V we discuss some implications of an explanation based on dealer's limited risk-bearing capacity. Moreover, we show to what extent traders may be able to profit from trading both the W as well as the fix when taking transaction costs into account. Section VI concludes.

Literature Review

Our paper is primarily related to the literature on financial intermediation and limited risk-bearing capacity in financial markets. The Duffie (2010) presidential address surveys the literature, reviews a host of empirical regularities and discusses competing structural explanations. For equities, Chen, Noronha, and Singal (2004) document that the prices of stocks delated from the S&P 500 index drop after deletion dates and rebound in the following months. For bonds, Lou, Yan, and Zhang (2013) show that in the run up to U.S Treasury auctions yields (prices) gradually rise (decline) in anticipation of the auction date and thereafter revert to pre-auction levels. And in option markets, Nagel (2012) argues that short-term reversal strategies can be interpreted as compensation for liquidity provision by dealers.

Our paper is also related to a long standing puzzle in foreign exchange markets documented by Cornett, Schwarz, and Szakmary (1995), Ranaldo (2009) and Breedon and Ranaldo (2013) who argue that foreign currencies depreciate in local trading hours. More recently, Jiang (2017) argues a depreciation of the U.S dollar in U.S. trading hours can be rationalised within the context of a long run risk framework à la Bansal and Yaron (2004). With respect to these papers, our contribution is two fold: Firstly, our granular dissection allows identification of price reversals around major currency fixes as opposed to drifts within market opening and closing times. Second, we provide an explanation for this pattern in terms of dollar demand shocks in the presence of limited risk bearing capacity and study the determinants of price impediments to trade.

Finally, our paper is related to a literature in market microstructure studying foreign exchange benchmarks. For the London fix, Evans (2018) assesses price dynamics in tight windows around the fix while Evans, O'Neill, Rime, and Saakvitne (2018) show differences in trading behavior across investor types; for the Tokyo fix, Ito and Yamada (2016) provides evidence of short lived price spikes. Our paper distinguishes itself from these studies along important dimensions. First, instead of focusing on dynamics in a short window around one particular fix, we assess return patterns over the entire trading day taking into account all three major currency fixes. Second, we document that persistent intraday pricing patterns have low frequency asset pricing implications. Lastly, our evidence is based on one of the largest FX intraday datasets analysed to date that allows us to confirm that the stylised facts we uncover are a systematic and robust feature of the data.

I. Data

The empirical analysis is based on one of the most comprehensive high-frequency foreign exchange data sets analysed to date and is constructed from multiple high-quality data sources. Our full sample starts in January 1999 and ends in December 2018, covering 20 years of high-frequency tick-by-tick data for the G10 currencies, including the Australian dollar (AUD), the Canadian dollar (CAD), the euro (EUR), the Japanese yen (JPY), the New Zealand dollar (NZD), the Norwegian krone (NOK), the Swedish krona (SEK), the Swiss franc (CHF), and the British pound (GBP), vis-à-vis the U.S. dollar. These currencies are consistently among the most liquid currencies over the sample period and together they account for close to 75% of the total daily turnover in the foreign exchange market according to the latest triannual BIS survey (see BIS (2019)).[3]

[3]Note that the total sum of market shares equals 200% as two currencies are involved in each foreign exchange transaction. The sum of market shares for our sample is over 172%. Assuming that the other currencies are only traded against the G10 gives a lower bound for the market share of 146% or almost three quarters of 200%.

A. Data Sources

We compile our data from multiple sources including *Thomson Reuters*, *CME* and *Datastream.*

Thomson Reuters data. For the full sample period we have high-frequency bid and ask quotes from *Thomson Reuters Tick History* (TRTH) that provides tick by tick indicative quotes that were available to market participants in real-time. Starting in April 2006 we also have data from the Thomson Reuters dealing platform ("Matching Prices with Volumes Daily (D5)") that provides real-time data on traded prices as volumes.[4] The platform provides an anonymous central limit order book for spot, forward and NDF trading for the interbank FX community in over 80 currency pairs and with over 1,100 subscribers. Together with *Electronic Broking Services* (EBS), *Matching* is the leading inter-dealer platform for foreign exchange trading with a daily volume for spot transactions exceeding 100 billion U.S. dollars (compared to around 76 billion U.S. dollars traded on EBS).[5] While not all currency pairs are equally liquid on both platforms (*Matching* for example is clearly the leading platform for Commonwealth currencies), Breedon and Vitale (2010) show that returns from both platforms for a given currency pair are highly correlated. The two primary electronic communication networks (ECN) account for under 10% of total spot transactions and the proportion of the two venues is further declining. However, BIS (2018) document that they remain crucial for price discovery in the foreign exchange market, leading for examples price changes in futures markets.

The main difference between the D5 and the TRTH data is that we have information on traded prices from *Matching*, while we source quoted prices from TRTH. The TRTH data on the other hand collects indicative quotes from individual banks from a platform that allows market participants to trade directly with banks (i.e., "bank-to-client" trading) and, as such acts more as an aggregator of quotes. In addition, the D5 data includes information on number of trades and volumes and an identifier for the initiator of the trade, allowing us to calculate various accurate measures related to foreign exchange order flow.

CME data. In addition to the Reuters data we also collect data on foreign exchange futures from the *Chicago Mercantile Exchange* (CME) that were introduced after the breakdown of the

[4] *Thomson Reuters Tick History* is now part of *Refinitiv*, the re-branded *Financial & Risk* business of Thomson Reuters.

[5] EBS is now part of CME, offering an OTC platform alongside the foreign exchange futures and options traded on the exchange.

post WWII Bretton Woods agreement in 1972 for the period from June 2006 to December 2018.[6] Unlike most foreign exchange instruments, FX futures are exchange-traded and marked-to-market on a daily basis. While the overall daily transaction volume is magnitudes smaller than the volume in the OTC market with only around 127 billion U.S. dollars per day (compared to 6.6 trillion U.S. dollars global daily turnover), the futures are very liquid and the average volume is around one million contracts per day. From the CME database we collect high-frequency bid and ask prices as well as transaction volumes. With respect to OTC indicative spot quotes, futures prices and quotes have two attractive properties: (*i*) quotes are real-time executable; and (*ii*) futures prices incorporate the cost of carry which in the currency is an implied interest rate differential.

Options data. Finally, to measure implied volatilities, we use information on over-the-counter foreign exchange options at the daily frequency from two different sources. For the period January 1999 to February 2013 we use one-month plain-vanilla European call and put options obtained through JP Morgan for all currencies against the U.S. dollar. We extract implied volatilities for at-the-money (ATM), 10-delta, 25-delta calls, as well as 10-delta and 25-delta puts. To extend the sample period until December 2018, we supplement our sample using options data from Bloomberg from March 2013 onwards. In contrast to the data from JP Morgan, the implied volatilities on Bloomberg are quoted as at-the-money straddles, risk-reversals and butterflies. We follow Carr and Wu (2009) to back out plain-vanilla implied volatilities for the same 10-delta and 25-delta strikes as the JP Morgan data.[7] Equipped with the cross-section of implied volatilities we follow Mueller, Stathopoulos, and Vedolin (2016) to construct daily measures of model-free implied volatility for the respective currency pairs.

B. Data Sample and Currency Returns

From the TRTH and CME data sets we obtain the best bid and ask quote recorded to the nearest even second. After applying a number of filters to correct the data for outliers, the price at each five-minute tick is obtained by linearly interpolating from the average of the bid and ask quotes for the two closest ticks. If no quote was submitted during a specific interval, we fill the gap with the most recent available price. The quotes are then used to construct the mid prices as

[6] See `www.cmegroup.com/education/files/understanding-fx-futures.pdf`.

[7] For most currency pairs Bloomberg provides foreign exchange options data from January 2009 onwards.

well as the currencies' net returns at five minute intervals. In addition, the bid-ask prices also allow to calculate returns net of transaction costs. Following previous studies (e.g. Andersen, Bollerslev, Diebold, and Vega (2003)) we exclude quotes that are submitted on days that are associated with low trading activity. We remove all quotes on weekends between Friday 5:00 p.m. and Sunday 5:05 p.m. (Eastern Standard Time, EST). Similarly, we drop quotes around fixed holidays, i.e., Christmas (24 to 26 December), New Year (31 December to 2 January), and 4 July, and around flexible holidays, such as Good Friday, Easter Monday, Memorial Day, Labor Day, and Thanksgiving (including the day after). We express all spot rates in U.S. dollar per foreign currency. Hence, an increase of the (log) exchange rate s_t can be interpreted as an appreciation of the U.S. dollar vis-à-vis the foreign currency.

II. Currency Fixings

According the to the latest BIS (2019) survey, daily foreign exchange turnover reached around 6.6 trillion U.S. dollars in 2018; most of the trade is in FX swaps (3,202, numbers in billion U.S. dollars), spot transactions (1,987), currency forwards (999), and options (294). In contrast, exchange traded derivatives such as futures and options only reach a daily turnover of a measly 127 billion U.S. dollars. Thus, the overwhelming majority of foreign exchange transactions happen in the over-the-counter (OTC) market.

There are essentially two types of execution models possible in the OTC FX market. On the one hand, banks may act as a broker on behalf of their clients, helping them trade and obtain liquidity without taking on any inventory risk. On the other hand, banks may act as principals that actively take positions when trading with clients and that manage the risks that result from holding FX inventory.

Banco Santander for example clarifies the nature of the trading relationship between the client and the bank in an "FX Disclosure Notice."[8] In the notice, the bank clearly states that it acts as a principal in the wholesale FX market and lays out a menu of execution types that clients can choose from based on their needs and sophistication. For example, an *at best* order will be executed using the bank's 'discretion and expertise to achieve the best price,' while a *stop* orders will be

[8]See `https://www.santander.com/en/landing-pages/foreign-exchange-disclosure-notice` for the full text.

executed at the 'nearest possible level given prevailing market conditions.' A further option is to place *fix* order, where the bank executes the 'transaction at the published *fix rate* after applying a bid offer spread around that mutually agreed upon price.'

The notice also states that the bank will engage in pre-hedging, i.e., the bank will manage the inventory risks associated with the anticipated execution of client orders. This means that the bank will actively execute transactions for hedging purposes either preceding or following client orders.[9] In particular, the bank cautions that these transactions may have and 'unintended effect' by 'impacting the benchmark fixing or related markets.' This means that there could be unintended and, more importantly, negative consequences for clients that choose to trade at the fix, begging the question why they should choose to do so in the first place over choosing any of the other execution types.

A currency fixing is a pre-set time of day when bids and offers are aggregated and a reference price is published. The most popular fixings are the London, ECB and Tokyo fixings.[10] Figure 1 depicts these fixings visually in Eastern Time (ET, the time in New York) 'around the clock'. The coloured blocks in Figure 1 show the regular trading hours in the futures markets of each location. The figure begins at 5:00 p.m. ET which is the end of the trading day in New York and roughly the beginning of the trading day in Australasia. The first major currency fixing that occurs is Tokyo at 10:00 a.m. local time which is 9:00 p.m. ET (or 8:00 p.m. depending on DST). The red, green and yellow blocks overlap, meaning that as Japanese trading is closing European markets are opening. The beginning of the trading day in New York (we assume 8:00 a.m. for currencies) happens close to the 'ECB fix' at 8:15 a.m. ET (2:15 p.m. local time) but the timing is clearly not exactly aligned. Moreover, the ECB fix is also not aligned with the usual release time of macro announcements at 8:30 a.m. ET. As we argue later, the distinction in timing is important when considering intraday price movements in exchange rates. The final and most important fix of the day is the London fix at 4:00 p.m. local time (or 11:00 a.m. ET).

[9] Risk management activities prior to trade acceptance—which are distinct from pre-hedging activities—are collectively referred to as 'last look' and are required because spot rates fluctuate continuously which may lead to differences between quoted spot rates and executed spot rates. The Global Foreign Exchange Committee has recently called for a revision of rules regarding the use of last look which potentially allow the banks to exploit customer information: `www.globalfxc.org/docsthe_role_of_cover_and_deal_arrangements.pdf`.

[10] As of 2018 there are now currency fixings published every 30 minutes.

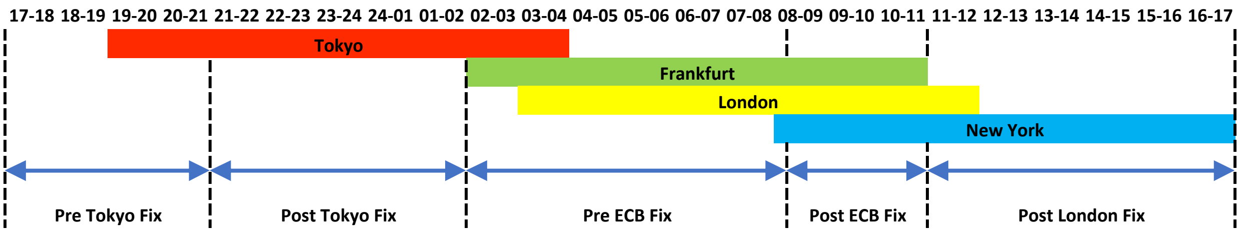

Figure 1. Currency Fixes across Time Zones

While all have an impact foreign exchange markets, they differ from each other with respect to institutional characteristics, publication time of reference rates, as well as the methodologies to compute fixing rates. In what follows, we provide a summary of the institutional characteristics of the three major fixes in currency markets, while a detailed comparison of FX benchmarks can for example be found in FSB (2014).

First, the Tokyo fix rates are published at 10:00 a.m. local time whereby each bank determines its own individual fixing rate for their customers. This is a major difference compared to the ECB and London fixes where only one reference rate is published. The rates of the Tokyo fix are based on transacted prices, which bank sample from its own customer transactions at 9:55:00 a.m. Further, the fixing rate applies not only to pre-fixing but also to post-fixing customer orders submitted after 10:00 a.m. The Tokyo fix, therefore, has far reaching consequences for banks for the rest of the trading day (see, e.g., Ito and Yamada (2016)).

Second, reference rates from the ECB fix are based on a daily teleconference between eurozone central banks at 2:15 p.m. CET. The reference rates are the average of quoted bid and offer prices against the euro, which means that the ECB reference rate is not based on actual transactions. Since 1st July 2016 onward, the ECB has delayed publication its reference rates until 16:00 CET, while the methodology to compute the reference rates has remained the same (ECB (2019)). The ECB reference rates are particularly used by non-financial corporations in the euro-area that use forward contracts for hedging purposes (FSB (2014)).

Lastly, the fixing rate in London is set at 4:00 p.m. BST and published by WM/Reuters. In contrast to the Tokyo Fix, the London fix rate applies to all banks and pre-fixing order that arrive before 4:00 p.m. The fixing rate is computed based on trades, and quotes for less liquid currency pairs, in a window around 4:00 p.m. In a five-minute interval around the fix (3:57:30 p.m. to 4:02:30 p.m.) traded rates are sourced every second from major FX platforms and after validation checks, a median trade based bid and offer rate is calculated from the pooled sample of trades.[11] Finally, the mid-rate is based on the median bid and offer trade rates, adjusted for a minimum spread, and validated before published to market participants. The London fix is prominently used by various groups of market participants that value their international portfolio positions (Melvin and Prins (2015)).

Customers in foreign exchange markets rely heavily on benchmark rates which generates a concentration of trading around the fixes. To make this point clear, from D5 we compute the volume in every 1 minute interval normalised by the total volume that occurred on that day. We do this for all days and compute 1-minute averages across the sample period from 1 January 2006 to 31 December 2018. Figure 2 plots this statistic for EUR/USD, GBP/USD and JPY/USD pairs. A number above '1' means there is more volume at this point compared to the average across all points in the day and vice-versa for a number below '1'. Volumes spike dramatically at each of the fixings with the spike in JPY trade being largest at the Tokyo fix (6 times the daily average) and the spike in GBP being largest at the London fix (17 times the daily average). Interestingly, while there is a spike in volumes at the ECB fix there is a larger spike in volumes 15 minutes later at 8:30 a.m., coinciding with the time of macro releases (5 to 8 times the daily average). We also note that volumes for all three pairs are larger than average when European markets are open and that during European and U.S. hours volumes spike on the hour driven by the lesser known hourly benchmark fixings. Finally, Figure 2 shows that volumes continue to be significant for a considerable amount of time after the London fix providing clear evidence against the idea that the fix marks the end of the foreign exchange trading day.

[INSERT FIGURE 2 HERE]

III. Currency Returns Around the Clock

In this section, we provide novel evidence on the intraday behaviour of foreign exchange and, in particular, document the following stylised fact: *exchange rate returns display a predictable intraday seasonality such that the U.S. dollar appreciates in the run up to foreign exchange fixings*

[11] Before 14th December 2014, the length of the window to calculate the fixing rate was limited to a one-minute interval from 3:59:30 p.m. to 4:00:30 p.m.

and depreciates thereafter. This pattern is robust across currency exchanges and sample periods. Furthermore, it is present for all days of the week, weeks of the month and months of the year, and it is not concentrated on special days when macro information is released or when the central bank makes an announcement.

A. Dissecting Currency Returns

Unless otherwise noted, we take the perspective of a New York based investor and work with Eastern Time (ET) so that exchange rates are quoted in units of foreign currency per USD.[12] We define daily close-to-close log spot returns (Δs_d^{CTC}) as the percent change in the mid-quote between 5:00 p.m. on day d and 5:00 p.m. on day $d-1$, i.e.,[13]

$$\Delta s_d^{CTC} = s_d^{5:00p.m.} - s_{d-1}^{5:00p.m.}.$$

Next, we split the day into different periods guided by the timing of the three main currency fixes across the globe: (i) the Tokyo fix at 10:00 a.m. local time; (ii) the ECB fix at 2:15 p.m. local time; and (iii) the London fix at 4:00 p.m. local time. In ET, within these splits log spot returns are given by

$$\begin{aligned}
\Delta s_d^{\text{Pre-T}} &= s_{d-1}^{9:00p.m.} - s_{d-1}^{5:00p.m.} && \text{pre-Tokyo fix (pre-T)} \\
\Delta s_d^{Post-T} &= s_{d}^{2:00a.m.} - s_{d-1}^{9:00p.m.} && \text{post-Tokyo fix (post-T)}^{14} \\
\Delta s_d^{Pre-E} &= s_{d}^{8:15a.m.} - s_{d}^{2:00a.m.} && \text{pre-ECB fix (pre-E)} \\
\Delta s_d^{E-L} &= s_{d}^{11:00a.m.} - s_{d}^{8:15a.m.} && \text{ECB fix to London fix (E-L)} \\
\Delta s_d^{Post-L} &= s_{d}^{17:00p.m.} - s_{d}^{11:00a.m.} && \text{post-London fix (post-L)}
\end{aligned}$$

In order to be able to distinguish between the post-Tokyo and the pre-ECB fix periods we use 8:00 a.m. Frankfurt time (or 2:00 a.m. ET), i.e., the beginning of the FX trading day in Europe.

[12]Eastern Time denotes the generalised time zone in New York. During summer and winter the time zones in use in New York are Eastern Daylight Time (EDT) and Eastern Standard Time (EST), respectively. Note that DST is in operation for about two thirds of the year.

[13]Note that our choice of closing time differs from the London fix time at 4:00 p.m. that is normally used to define daily currency returns and that underlies the standard WM/Reuters FX data available on *Datastream*, for example.

Similarly, we could define the start of the FX trading day in New York as 8:00 a.m. However, since the ECB fix is scheduled at 8:15 a.m. ET we use the window between the ECB fix and 5:00 p.m. ET to define the intraday period in NY. In addition, the start of our intraday period is almost aligned with the start of the open outcry hours on CME, where various currency futures and options are actively traded.[15] Consequently, the overnight period is defined as the window between 5:00 p.m. ET and 8:15 a.m. the next day. As the major U.S. macro announcements come out at 8:30 a.m. ET (i.e., before the NYSE opens), our definition of the U.S. FX trading day means these points occur during the intraday window.

It is worth noting that our approach ensures that the sum of the intraday and overnight return components on a daily level exactly add up to the close-to-close return (i.e. $\Delta s_d^{ID} + \Delta s_d^{ON} = \Delta s_d^{CTC}$). By construction, this also holds at the monthly frequency even though slight differences arise as we drop weekends and holidays from our sample and, thereby, we ignore returns generated on these days.

B. Currency Returns Around the Clock

B.1. Reuters TRTH Indicative Quotes

We begin our analysis by plotting the annualized average cumulative 5-minute log returns between 5:00 p.m. ET on day t and 5:00 p.m. ET on day $t+1$, for the sample period 1 January 1999 to 31 December 2018 and the G10 currencies. Figure 3 plots the average annualised returns to the EUR, JPY and GBP pairs while Figure 4 shows both hourly and cumulative returns for each of the nine currencies in separate panels.

All currencies show a distinct pattern of depreciation ahead of the Tokyo fix at 9:00 p.m. ET followed by a reversal thereafter. Once European markets open at 2:00 a.m. ET, all currencies depreciate against the USD ahead of the ECB fix. This drop is much stronger for the European currencies and more muted for the AUD, NZD and CAD. The period between the ECB and London fix doesn't show a clear pattern in the cross-section aside from the EUR and JPY who appreciate until one hour before the London fix. After the London fix, all currencies show a strong appreciation versus the USD until the end of the business day in the U.S. at 5:00 p.m. ET with

[15] An overview of currency futures trading hours can be found at `http://www.cmegroup.com/trading-hours.html#fx`.

the JPY being the sole exception and moving in the opposite direction. Overall, all currencies except the JPY appreciate during the U.S. intraday period and depreciate overnight.

[INSERT FIGURES 3 AND 4 HERE]

Figure 5 plots the average cumulative returns over a 24-hour period of the unconditional dollar portfolio that goes long all foreign currencies in equal weights along with the average hourly returns of the dollar portfolio.[16] Aggregating across currencies, the consistent depreciation of foreign currencies before the Tokyo fix and after European markets open combined with the depreciation of the U.S. dollar during the intraday period leads to a distinctive *'W'*-shaped pattern of the cumulative returns measured over a full day. Overall, there is a significant appreciation of the U.S. dollar during the overnight period of around 4.5% per year followed by a reversal during the day of just over 5%. Over the full sample period, the U.S. dollar has depreciated against the basket of currencies at a rate of roughly 0.6% per year. Given the size of the FX spot market, this translates into very large sums. Using daily turnover numbers from the 2016 Triannual BIS survey, the pattern we detect implies potential daily swings between 430 and 870 million U.S. dollars depending on whether we calculate the potential impact based on daily U.S. dollar turnover or based on the daily turnover of the currencies in our sample.

[INSERT FIGURE 5 AND TABLE 1 HERE]

Table 1 summarizes Figures 3 to 5 formally by reporting average FX log returns (i.e., exchange rate changes) along with t-statistics for the various intraday sub-periods as defined above.[17]

As discussed above, all foreign currencies depreciate against the U.S. dollar after trading in New York ceases and in anticipation of the Tokyo fix. The Australian and New Zealand dollar (-7.70% and -9.24%, respectively) show the most negative average returns, while the Japanese yen and the Swiss franc depreciate the least compared to other currency pairs. It is worth highlighting that irrespective of the magnitude of the returns, average annualized returns of all G10 currencies

[16] We follow Lustig, Roussanov, and Verdelhan (2011) in constructing the dollar portfolio using the G10 currencies from our sample.

[17] Note that at this stage we explicitly take Daylight Savings Time into account by calculating pre- and post-Tokyo fix returns using windows that line up around 10:00 a.m. Tokyo time. During the winter months when New York follows EST, this means 8:00 p.m. ET and during the summer months when New York follows EDT this means 9:00 p.m. ET. All figures are plotted using ET.

are different from zero at the 1% level of significance. The reversal after the Tokyo fix is equally statistically significant for all currencies in our sample. The magnitude by the Norwegian krone and the Japanese yen with 8.22% and 7.72% per annum, respectively. Not very surprisingly, the dollar portfolio exhibits a very strong and significant reversal pattern as well, dropping around 5% before the Tokyo fix and recouping the losses thereafter.

Leading up to the ECB fix, the European currencies and the Japanese yen depreciate against the U.S. dollar. The point estimates are large in both statistical and economic terms. The highest drops are posted by the euro, and the Swedish krona with -7.12% and -5.92% measured on an annual basis, respectively. Between the ECB and the London fix, currencies don't move as consistently in the cross-section as during other windows. However, in general, the dollar appreciates in advance of the London Fix with the largest drop observed for the major European pairs. The t-statistics for these pairs are significant at the 1% level, which shows that, even in this short period, a pre-fix dollar appreciation is a robust feature of the data.

After the London fix, the pattern is striking: with the exception of the JPY, all currencies appreciate strongly (i.e., between 3.79% for the New Zealand dollar and 6.59% for the euro) during the period between the London fix and the close of markets in the U.S., whereas the Japanese yen depreciates by 2.59%. Overall, the dollar portfolio appreciates by almost $\sim$ 4% and movements for all currencies are strongly statistically significant.

The last column in Table 1 makes clear that the pattern we document is an intraday seasonality (i.e., a predictable component) that does not carry over to close-to-close returns. In fact, with the exception of the Swiss franc (average annual appreciation of 2.47%) and the Australian dollar (average appreciation of 1.16%) none of the currencies in our sample moves by more than a percent on average over the whole sample period we consider and none of the close-to-close returns are statistically significant. The dollar portfolio for example appreciates on average by only 90 basis points per year.

B.2. Alternative Data Sets

The reversal of the dollar around the major currency fixings is both economically large and highly statistically significant in annualised terms. Moreover, the pattern is remarkably robust over time. That said, on a daily basis the movements are on the order of a few basis points, begging the question whether the pattern is an artefact of using TRTH indicative quotes to calculate log spot changes. To this end, we consider additional data sources to verify the stylised facts we report based on the TRTH data. For the sample period April 2006 to December 2018 we obtain data from the Thomson Reuters dealing platform (D5) that provides a historical archive of transacted prices alongside associated volumes. This allows us to compute volume-weighted-average-prices (VWAPs) for arbitrary intervals.

[INSERT FIGURES 6 HERE]

Finally, we also exploit the third source of data and plot the cumulative returns of CME futures for the three currency pairs since June 2006 over a 24-hour period in the bottom panel of Figure 6. As with the other data sets, currencies follow a W-shaped pattern around the clock. The visual conclusion is confirmed when we calculate average annualized returns for the different windows as in Table 1.[18]

[INSERT TABLE 2 HERE]

To illustrate how robust and consistent the W-shaped pattern is in currency returns, we show in Table 2 the average returns during the respective windows for the three pairs EUR, GBP and JPY using the TRTH, Reuters D5 Matching and CME data sets for the common (and post-crisis) sample from January 2009 to December 2018. This provides both a robustness check with respect to the sample period as well as the data source.

IV. Further Empirical Results

In this section we dissect the dollar return reversal pattern around the Tokyo, Frankfurt and London fixes along a number of dimensions, day of the week, month of the week, year by year analysis, and also study the extent to which foreign exchange reversals and currency benchmark fixings are unique to pairs bilateral to the U.S. dollar.

[18]See the Online Appendix for the detailed numbers.

A. Calendar Effects

A.1. Day of the Week Effects

Table 3 reports average FX log returns for the Dollar portfolio along with t-statistics for the various intraday sub-periods defined in the previous section. Considering first the Tokyo reversal, the pre/post fix dollar appreciation/depreciation are consistent across days of the week both in terms of magnitude and statistical significance. Dollar appreciations are highest on Tuesdays (7.14%) and Fridays (7.51%) and lowest on Thursdays (2.70%). Subsequent depreciations are remarkably symmetric across all days. The pre ECB / pre LND (pre-E + E-L) dollar depreciation is less consistent across the week, and is significant on Mondays, Wednesdays and Fridays, while the post LND dollar depreciation is large in magnitude and significantly positive on Wednesdays, Thursdays Fridays. Aside from differences in magnitudes, each day of the week displays the same W-shaped intraday return pattern where the sequence of reversals has consistent signs around the fixes. Close-to-close returns are statistically different than zero on any day, however due to the large pre-E dollar appreciations on Mondays and Fridays are at least economically significantly negative, while the particularly large dollar depreciations post London on Wednesdays and Thursday generate sizeable negative close to close returns.

In summary, the W-shaped dollar portfolio return pattern is present with consistent signs in each day of the week with varying levels of statistical significance. For the Tokyo fix, returns display a systematic reversal that is symmetric in economic magnitude and strongly statistically significant. For the London fix, the pattern is mostly present although its magnitude and statistical significance display a greater degree of heterogeneity.

[INSERT TABLE 3 HERE]

A.2. Month of the Year Effects

Table 4 displays summary statistics for the Dollar portfolio returns defined within our fix windows for each month of the year, averaged across all trading days in the sample. As expected, returns for a given currency may display larger variation in their intraday patterns and statistical significance is somewhat lowers. However, the reversal pattern for Tokyo is again remarkable consistent, symmetric in magnitude, and in almost all cases statistically significant at the 1% level. Pre-ECB returns are negative in 10/12 months and pre-LND (pre-E + E-L) returns are negative in 9/12 months although rarely significantly so on a month-by-month basis. Post London fix returns are positive in 11/12 months, the exception being October, and significant at close to the 5% level in 6/12 months. Taking a step back and considering the broader patterns a W-shaped return pattern is clearly visible in each month of the year. Point estimates for each month almost always go in the direction of a W-shaped pattern and so aside from the reduced power of the test, the Tokyo, Frankfurt, London reversal is a pervasive feature throughout the calendar year.

[INSERT TABLE 4]

B. Persistence over Time

The results in Table 1 are highly statistically significant, implying that a strategy going long and short the respective currencies around the fixes should result in very high returns before transaction costs. In Figure 7 we show how returns associated with the V-shaped reversals around the fixes develop over time. Ignoring transaction costs, we construct total return indices for the three major currency pairs (EUR, GBP, JPY) against the U.S. dollar as well as for the dollar portfolio. Before the Tokyo and ECB fix we take a short position that is reversed post-Tokyo and post-London fix respectively.[19] Panels (a) and (b) plot the evolution of a portfolio with an initial value of one dollar and a trading strategy around the Tokyo and ECB/London fix, respectively, while panels (c) and (d) display the annual returns for the dollar portfolio around the respective fixes.

[INSERT FIGURE 7 HERE]

First, the total return indices highlight the strong persistence of the V-return pattern over the entire sample period. All portfolios (with the exception of the Japanese yen around the ECB/London fix) accrue steadily over the whole period but with a stronger appreciation around the local fixes. An investment of one U.S. dollar in the yen trading strategy around the Tokyo fix climbs to almost 13 U.S. dollars by the end of 2018. The same strategy for the euro, pound and dollar portfolio only results in a final portfolio value of 10, 5.3 and 9.2 U.S. dollars respectively. In contrast, the same strategy implemented around the ECB/London fix results in portfolio values

[19] Note that we ignore the period between the ECB and the London fix for this exercise.

of 15.3, 10.7 and 5 U.S. dollars for the euro, pound and dollar portfolio, respectively, while the yen portfolio actually loses about 6% of value.

Second, the plots confirm that return reversals are particularly strong during times of financial turmoil and financial crises. While the returns are relatively low in 2007 (still between 5% and 10% for the dollar portfolio) they reach 10%-20% per year in 2008 and 2009. Furthermore, panels (c) and (d) also show particularly high returns to the dollar portfolio during 2001 when the dot-com bubble burst.

In summary, not only are fix patterns a feature of the data in each day of the week and month of the year but the pattern has persisted over a long sample period. This particular point highlights that high frequency variations can have low frequency impacts is they occurs in the same direction systematically over extended sample periods.

C. Re-basing and Dollar Specialness

Finally, we investigate whether the W-shaped intraday pattern we document in Section III is a unique feature of the U.S. dollar or whether it is shared by alternative base currencies. To this end, we compute base currency portfolios, i.e., the analogue to the dollar portfolio, for the euro, pound and yen, respectively. This can be done very easily using the fact that cross-pairs can be expressed as the ratio of two dollar based currency pairs, i.e.,

$$S^{\frac{GBP}{EUR}} = \frac{S^{\frac{GBP}{USD}}}{S^{\frac{EUR}{USD}}}$$

such that the exchange between the GBP and EUR ($S^{\frac{GBP}{EUR}}$) equals the ratio of the pound-dollar ($S^{\frac{GBP}{USD}}$) and the euro-dollar ($S^{\frac{EUR}{USD}}$) exchange rate. Re-basing all G10 rates vis-à-vis the three currency pairs we can construct a euro, pound and yen base portfolio as described in Section III, respectively.

In Figure 8 we plot the cumulative 5-minute intraday spot returns for the three alternative base portfolios. As per our usual convention, a positive return implies an appreciation of the foreign currencies and a depreciation of the respective base currency.

[INSERT FIGURE 8 HERE]

It is obvious that all the three base portfolios in Figure 8 exhibit significant intraday patterns. However, none of the patterns for the euro, pound or yen are similar to the W-shape displayed by the dollar portfolio in Figure 5. The euro is rather flat during the overnight period in Europe and then depreciates in the build-up to the ECB fix, reaching a low at the fix. Thereafter, the currency gradually appreciates until the end of the day in New York. The pound on the other hand reaches a low early on in the trading day in London (around 10 a.m. London time) before gradually appreciating through both the ECB and the London fix until the end of the day in New York. Finally, the yen exhibits a gradual appreciation through the intraday period in Asia with a small peak at the Tokyo fix. There is a second local peak about an hour before the ECB fix at 7 a.m. ET and, finally, the yen reaches the intraday peak at the London fix before sharply depreciating until the end of the day in New York. For all base portfolios, the patterns are not driven by the U.S. dollar as is evident when comparing the cumulative returns from the base portfolio with all currencies with those of a version where we exclude the dollar.

While it is of independent interest to investigate the patterns for the different base currencies, we merely use Figure 8 to highlight that the main stylised fact we document in Section III is specific to the U.S. dollar and leave further exploration of the respective base currency patterns for future research. In summary, none of the other base portfolios has systematic reversals around all the major fixes as is the case for the U.S. dollar, an indication of the importance of the U.S. dollar that is used either as a base or quote currency in 88% of all foreign currency transactions. Given the importance of the fixes for currency trading, only the U.S. dollar seems to be used systematically in transactions at the fix around the trading day. Moreover, only for the U.S. dollar must we infer a systematic excess demand on average at the respective fixing times. The euro base portfolio for example only has local through at the ECB fix mirroring the pattern of the euro against the U.S. dollar apparent in Figure 3. Similarly, the yen base portfolio appreciates after the London fix, in line with the depreciation of the yen against the U.S. dollar during the same window.

V. FX Liquidity and Risk Bearing Capacity

In this section we briefly discuss how the main patterns we present in Section III fit existing evidence and models for the Treasury auctions market. Furthermore, we also present results taking trading costs into account, thus shedding light on the question whether it is possible for traders to profit from the predictable return patterns we document for intraday currency returns.

A. Intermediation, Inventory Risk and Risk Bearing Capacity

Grossman and Miller (1988) emphasize one particular role of dealers and specialists in asset markets, namely to offer immediacy to investors who arrive and trade asynchronously. Intermediaries in turn then generate profits by absorbing the resulting order imbalances and subsequently trading them away. In his presidential address, Duffie (2010) discusses how trading opportunities coupled with slow movement of investment capital can cause price reversals of the kind we have documented around currency fixings. One example he highlights is the price pattern of Treasury securities around (pre-scheduled) auction dates studied in Lou, Yan, and Zhang (2013) where prices of Treasury securities gradually decline in anticipation of the auction date before recovering thereafter.

Apart from the frequency of the reversal, the price patterns of Treasury securities around auctions strongly resemble the patterns we document around the global currency fixing times. Borrowing from the framework of Vayanos and Wang (2009), Sigaux (2018) formalises the intuition about the mechanism in two-period model with liquidity traders and liquidity providers where the traders are infinitely risk averse. At the intermediate date there is an uncertain net supply of the risky asset in the market and liquidity providers have to take this uncertainty into account. Consequently, the demand of liquidity providers has both an arbitrage and a hedging component.

The interpretation of the model in Sigaux (2018) can be adapted (re-labelled) for currency markets. Instead of having a net supply of Treasuries as the difference between the amount of securities issued and the amount of securities absorbed by natural buyers we have a net demand for U.S. dollars (or, equivalently, a net supply of foreign currency) as the difference between the U.S. dollar supply by corporates and the U.S. dollar demand by liability-driven investors. Liquidity providers have to trade taking uncertainty about the net demand for U.S. dollars into account and they face a trade-off between arbitraging the difference between the pre-fix price and the expected price of the U.S. dollar at the fix and hedging the uncertainty about the expected price at the fix.

B. Currency Return Patterns, Volatility and Liquidity

One implication of the of the framework laid out above is that the magnitude of the V-shaped reversals around the fixes should depend on risk and liquidity in the FX market. During periods of elevated risk levels dealers should earn a higher risk premium for providing liquidity around the fix and, similarly, the provision of liquidity should be compensated better during periods of low liquidity. To this end we use bid-ask spreads to proxy for time variation in liquidity and implied volatilities extracted from either FX or equity options to proxy for risk in the foreign exchange market.

We expect that when volatilities are high or when liquidity is low (or, equivalently, spreads are high) the magnitude of the patterns we document are higher as well. To test this conjecture we compare average returns to trading the W during days when volatility is high (or liquidity low) to the average returns during the normal days (i.e., when volatilities and spreads are low). To this end we sort all days in the sample based on the previous day's volatility or normalised spread into quintiles. The high volatility and high spread days are the top quintile over the whole sample period whereas we designate the bottom three quintiles (or the bottom 60% of days) as the "normal" days.

FX volatilities are the implied volatilities we extract from currency option prices for the respective currencies. For the dollar portfolio we use the average implied volatility for all nine currencies in our sample. Alternatively, we also use the VIX taken from CBOE to sort the days in our sample. This means that in the case of equity implied volatility the sorting variable is the same for all currencies. For the spreads we sort based on the respective window, i.e., around the Tokyo fix or around the European fixes to allow for different liquidity stances in the Asian as well as the European or North American markets.[20]

In Figure 9 we plot the cumulative FX returns for the dollar portfolio as well as the EUR,

[20] We also sort according to the total intraday average spread. The results are qualitatively the same. Furthermore, there is no substantial difference to sorting with a one-day lag as opposed to sorting contemporaneously. Finally, the results are also largely robust to comparing top vs. bottom quintiles or quartiles as opposed to top vs. the bottom three quintiles.

GBP and JPY for the high and low spread days (top vs. bottom three quintiles), respectively. As is evident from the figure, the magnitudes of the currency movements are much more pronounced for the low liquidity (high spread) days as opposed to the high liquidity (low spread) periods. To formalise this insight we first calculate returns earned from trading the W, i.e., we take a short position in the foreign currency before the fix and we revert to a long position after the fix. The "Tokyo" window reverses the position at the Tokyo fix whereas the "Europe" goes short the foreign currencies before the ECB fix and long after the London. Thus, the period between the two fixes is ignored in this analysis—as is evident from Table 1 there is no clear pattern for the interfix period. All returns are based on midquotes and, thus, ignore transaction costs.

Table 5 displays the corresponding results for sorts on spreads as well as volatility for either the full sample going back to January 1999 or a shorter sample period starting in January 2009. For all currencies (including the dollar portfolio) and both samples the returns during low liquidity (high spread) periods are higher than during normal periods. For the dollar portfolio the annualised average differences are 8.4% and 6.6% for the full sample period and even higher when measured over the last ten years only. The differences are highly statistically significant for the EUR, GBP and dollar portfolio and marginally significant for the JPY during the Tokyo window. Overall, the results are remarkably consistent across currencies, windows and samples.

The table also shows the result for the volatility sorts. Again, results are quite robust and very consistent. Only the JPY exhibits a (non-significant) negative difference during the Europe window and using the VIX as a sorting variable. All other sorts result in positive return differences implying a larger magnitude of the pattern we report on high volatility days. Moreover, all differences with the exception of the JPY are statistically significant during the full sample. During the late sample the GBP and the dollar portfolio differences for the Europe window remain positive but are no longer statistically significant.

[INSERT TABLE 5 AND FIGURES 9 AND 10 HERE]

Finally, in Figure 10 we plot the cumulative FX returns for the dollar portfolio as well as the EUR, GBP and JPY sorted on either the FX implied volatilities or the VIX. While there are some differences between FX IV and VIX sorts, the high and low volatility days often seem to coincide, leading to average cumulative returns that are highly correlated between across the two sorting variables. In summary, we find that the magnitudes of the V-shaped reversals around the fix are higher on high volatility days as well as on low liquidity days as measured using bid-ask spreads.

C. *Intraday Profitability and Transaction Costs*

The stylised facts regarding the intraday dynamics of currencies around the major fixes presented in Sections III and B are largely based on indicative high-frequency mid-quotes and, hence, do not account for transaction costs. In this subsection we examine whether the trading strategies implied by the return patterns are profitable in a practical setting that explicitly takes bid-ask spreads into account. Indeed, the patterns presented in Table 1 suggest that a trader would have to moved aggressively in and out of positions up to four times over the course of a 24-hour period to exploit the systematic exchange rate movements we document. Is this evidence of market inefficiency, or, consistent with a story of financial intermediation in which dealers are setting spreads in order to offset intraday swings?

Using the quoted high-frequency bid and ask prices from our benchmark TRTH data set as a proxy for the effective spread we calculate OTC returns net of transaction costs. However, there is evidence that spreads reported to standard databases are substantially wider compared to the effective spreads based on firm quotes and executed trades (see, e.g., Gilmore and Hayashi (2011)), leading to measures of net returns that are too conservative compared to what professional traders that move large volumes would obtain. Gargano, Riddiough, and Sarno (2018) compare the bid-ask spreads from *Datastream* with quoted prices from other data providers in the years after the financial crisis and they suggest decreasing indicative spreads by up to 75% in order to obtain a more realistic proxy of the transaction costs that big traders in the over-the-counter FX market face. When considering the profitability of the trading strategies based on the over-the-counter rates, we take an agnostic approach and report results for different spread adjustments ranging from zero to 75% in line with the existing literature.[21]

We start our analysis by exploring the price changes of the three currency pairs EUR, GBP and JPY around the fixes focusing on the full windows of depreciation and appreciation around

[21] More specifically, we follow the approach in Gargano, Riddiough, and Sarno (2018) and verify that the daily bid-ask spreads from the TRTH database closely resemble those from the publicly available indicative quotes obtained through *Datastream*. Moreover, we confirm that they are roughly comparable in size over the full sample period. This means that the arguments supporting the use of smaller bid-ask spreads to obtain more realistic results are broadly applicable in our context as well.

the respective fixes.[22] We define net returns to a strategy where an investor enters a short position in a foreign currency before the fix and closes the position at the fix on day t as $\Delta s_t^{pre} = -s_t^a + s_t^b$. The superscript a (b) refer to the ask (bid) price in FX spot (futures) markets when TRTH (CME) data is employed. Analogously, net returns to a strategy that buys foreign currencies at the fix and closes the position at some point in the hours after the fix, are defined as $\Delta s_t^{post} = s_t^b - s_t^a$. Starting with the EUR we go short at the open of market in Europe at 2:00 a.m. ET, reverse the position at the ECB fix and then hold the EUR until the end of the U.S. trading day at 5:00 p.m. For the GBP we use the same window with the exception that we revert the position at the London fix. Finally, for the JPY we take a short position at 5:00 p.m. ET, revert at the Tokyo fix and hold the JPY until 2:00 a.m. Panel A in Table 6 presents the results for holding the three currency pairs during the respective windows around the local fix. As with the shorter window around the fix, the annualised returns for trading the three currencies around the fixes are very high as long as transaction costs are ignored. Shorting the currencies before the fix and reverting the position afterwards yields annualised average returns of 13.6%, 11.2% and 12.9% for the EUR, GBP and JPY, respectively. Incorporating the full transaction costs as implied by the indicative quotes again largely reverses the picture and average returns for most windows turn negative, reflecting the huge turnover that the trading strategy implies. However, reducing the bid-ask spread by 50% now yields positive returns for most windows and all currencies amounting to 6.6%, 4.2% and 3.4% for the EUR, GBP and JPY, respectively. Using T-bill rates we also calculate Sharpe ratios for trading the W-shaped intraday patterns that are reported in Panel B of Table 6. In line with the results for the gross returns, Sharpe ratios are negative using the full transaction costs, and positive and very high if transaction costs are ignored. Market participants that are able to trade at better bid-ask spreads may earn Sharpe ratios ranging between 0.5 and 0.7 when exploiting the predictable intraday patterns. As a comparison, in a last step we repeat the same exercise and implemented the trading strategy in FX futures markets. As the reported prices refer to firm quotes, we only consider the case with full transaction costs that are derived from bid and ask prices recorded on the CME platform. Results are reported in the row $\text{BA}_{CME}^{100\%}$. As shown, even after accounting for the full spread as a proxy for transaction costs returns are positive for half of the reported intraday periods. Generating Sharpe ratios of 0.99 and 0.52, in

[22] The results for fixed windows of n-hours pre and n-hours post fixings are reported in the online appendix.

particular the systematic trends during the pre-E and post-T periods can be profitably exploited by investors.

[INSERT TABLE 6 HERE]

In summary, while there is strong intraday predictability around the fixes, it is not obvious that this can be exploited by the average trader. First, returns from trading a relatively small window around the fix are usually more than offset by transaction costs. Second, holding the currency positions for a longer window that allows to exploit the persistent drift patterns throughout the day may lead to positive excess returns, at least for traders that are able to get reasonably good conditions to trade.

VI. Conclusion

Benchmark currency fixes around the globe mark key times in the day when non-specialised participants in the foreign exchange market see their orders executed, causing tremendous spikes in volumes. While it may be argued that it is possible to obtain better prices throughout the day, it seems that the transparency provided by the benchmark rates is valued more highly than any potential benefit from having an order executed at a slightly more favourable price. Specialised participants in the market such as dealers and brokers can thus anticipate the exact timing when liquidity shocks in the foreign exchange market may regularly hit. Moreover, it seems to be the case that the two main categories of non-specialised participants in foreign exchange markets are foreign corporates (i.e., exporters) who regularly supply U.S. dollars to buy back local currencies and liability driven investors (such as Japanese pension funds for example) who regularly purchase U.S. securities such as Treasuries and, therefore, have a large demand of U.S. dollars.

This paper studies demand shocks for U.S dollar in high frequency around currency fixings and documents a new empirical fact: the U.S. dollar systematically appreciates ahead of the three major currency fixings and depreciates thereafter. That is, the U.S. dollar reaches an intraday peak at the Tokyo, ECB and London fix, respectively. The geographically distinct nature of foreign exchange fixings results in a distinct W-shaped return pattern over a 24 hour period. This pattern is robust over time, is present in all G10 currencies, in different data sets, is not driven by calendar effects, and is persistent throughout all years in our sample.

We argue this finding is consistent with the requirement of liquidity providers to meet the demand for U.S. dollars at the fix, which in turn generates inventory risk that must be managed by discounting foreign currency valuations relative to the dollar. More generally, we argue that our findings are consistent with early contributions on intermediation such Grossman and Miller (1988) as well as with set of theoretical predictions neatly summarised by the Duffie (2010) AFA presidential address. Lifting predictions from this literature, we show that intraday swings and bid-ask spreads are increasing in fundamental volatility and dealers' limited risk bearing capacity. Moreover, price reversals around fixes are only present in foreign exchange returns with U.S dollar base risk. Finally, consistent with limited capital arguments, we show that arbitrageurs with deep enough pockets can likely exploit the predictable return patterns even after taking into account transaction costs.

VII. Appendix: Tables

Table 1. Intraday Returns and the Tokyo, ECB and London Fix

This table reports annualized average returns for different intraday periods around the Tokyo, ECB and London fix using the TRTH data base. Positive values imply the foreign currency appreciates versus the U.S. dollar. The dollar portfolio "DOL" is an equal weighted average of all nine currencies in our sample. The windows are defined as follows: the pre-Tokyo fix window starts at 5:00 p.m. ET until the Tokyo fix at 10:00 a.m. local time ("pre-T"), followed by the post-Tokyo window ("post-T") that runs until 2:00 a.m. ET (when European markets open). The pre-ECB fix window ("pre-E") spans the period between European opening hours until the ECB fix at 8:15 a.m. ET. The "interfix" window ("E-L") covers the period between the ECB and the London fix at 11:00 a.m. ET. The final window spans the period after the London fix ("post-L") starting at 11:00 a.m. ET and ending at 5:00 p.m. ET. Thus, the intraday period is the sum of the "E-L" and the "post-L" windows whereas the overnight period covers the "pre-T", "post-T" and "pre-E" windows. All times are measured in Eastern Time, taking into account daylight savings time. Data is daily and covers the sample period from January 1999 to December 2018 (5009 observations). Returns are measured as the average log changes in the mid quote for the respective currency. All numbers are annualised and *t*-statistics are in parentheses.

	pre-T	post-T	pre-E	E-L	post-L	CTC
AUD	-7.70	5.19	-0.75	-0.36	4.79	1.16
	(-8.50)	(4.26)	(-0.54)	(-0.29)	(3.53)	(0.41)
CAD	-4.32	4.24	-1.68	-1.56	3.98	0.66
	(-9.66)	(7.71)	(-1.64)	(-1.32)	(3.74)	(0.33)
CHF	-3.15	3.64	-5.46	2.07	5.37	2.47
	(-5.46)	(5.51)	(-3.42)	(1.59)	(4.66)	(0.98)
EUR	-5.10	6.48	-7.12	0.09	6.59	0.94
	(-9.55)	(9.83)	(-5.59)	(0.07)	(6.13)	(0.42)
GBP	-4.89	3.53	-5.67	0.71	6.27	-0.06
	(-9.22)	(5.29)	(-4.41)	(0.70)	(6.75)	(-0.03)
JPY	-4.68	8.22	-2.27	1.25	-2.59	-0.08
	(-6.53)	(8.17)	(-1.99)	(1.12)	(-2.59)	(-0.04)
NOK	-5.03	7.72	-3.93	-4.28	4.92	-0.59
	(-7.48)	(9.67)	(-2.38)	(-3.10)	(3.89)	(-0.22)
NZD	-9.26	6.06	-0.12	1.54	3.79	1.00
	(-8.76)	(5.13)	(-0.08)	(1.18)	(2.63)	(0.69)
SEK	-5.21	6.05	-5.92	-1.53	6.32	-0.29
	(-7.21)	(7.61)	(-3.66)	(-1.12)	(4.88)	(-0.11)
DOL	-5.42	5.55	-3.53	-0.08	4.38	0.91
	(-12.24)	(9.40)	(-3.66)	(-0.09)	(4.96)	(0.51)

Table 2. Intraday Returns for EUR, GBP and JPY Across Different Data Sets

This table reports annualized average returns for different intraday periods around the Tokyo, ECB and London fix for the EUR, GBP and JPY using three different data sets, TRTH, D5 and CME, respectively. Positive values imply the foreign currency appreciates versus the U.S. dollar. The windows are defined as follows: the pre-Tokyo fix window starts at 5:00 p.m. ET until the Tokyo fix at 10:00 a.m. local time ("pre-T"), followed by the post-Tokyo window ("post-T") that runs until 2:00 a.m. ET (when European markets open). The pre-ECB fix window ("pre-E") spans the period between European opening hours until the ECB fix at 8:15 a.m. ET. The "interfix" window ("E-L") covers the period between the ECB and the London fix at 11:00 a.m. ET. The final window spans the period after the London fix ("post-L") starting at 11:00 a.m. ET and ending at 5:00 p.m. ET. All times are measured in Eastern Time, taking into account daylight savings time. Data is daily and covers a sample period from January 2009 to December 2018 (2515 observations) where we have quotes from all three data sets. Returns are measured as the average log changes in the mid quote (for the TRTH and CME samples) or the value-weighted average price (or VWAP, for the D5 sample) for the respective currency. All numbers are annualised and t-statistics are in parentheses.

		pre-T	post-T	pre-E	E-L	post-L	CTC
EUR	TRTH	-4.35	4.74	-6.63	0.13	5.01	-1.10
		(-6.04)	(4.91)	(-3.83)	(0.08)	(3.42)	(-0.36)
	D5 VWAP	-3.35	3.94	-7.78	0.64	5.05	-1.51
		(-4.19)	(3.85)	(-4.47)	(0.39)	(3.48)	(-0.50)
	CME	-5.09	4.26	-7.98	0.68	5.89	-2.24
		(-4.60)	(4.51)	(-4.55)	(0.41)	(4.02)	(-0.72)
GBP	TRTH	-3.93	2.46	-3.44	-0.74	6.40	0.75
		(-5.00)	(2.28)	(-1.78)	(-0.52)	(4.88)	(0.25)
	D5 VWAP	-3.62	2.35	-3.57	-0.57	6.21	0.80
		(-4.61)	(2.23)	(-1.85)	(-0.40)	(4.85)	(0.27)
	CME	-4.94	1.99	-3.22	-0.78	6.06	-0.89
		(-3.09)	(1.90)	(-1.66)	(-0.53)	(4.50)	(-0.27)
JPY	TRTH	-4.83	7.92	-3.71	0.72	-2.40	-2.30
		(-4.89)	(5.50)	(-2.50)	(0.47)	(-1.73)	(-0.75)
	D5 VWAP	-5.79	6.83	-2.74	0.91	-1.09	-1.87
		(-5.81)	(4.56)	(-1.75)	(0.58)	(-0.75)	(-0.59)
	CME	-3.16	5.95	-4.52	2.26	-3.07	-2.53
		(-2.69)	(4.11)	(-3.03)	(1.43)	(-2.14)	(-0.80)

Table 3. Day of the Week Effects

This table reports annualized average returns for different intraday periods around the Tokyo, ECB and London fix for the dollar portfolio "DOL" which is an equal weighted average all nine currencies in our sample. The sample is split according to the various days of the week, i.e., Monday through Friday. Positive values imply the foreign currency portfolio appreciates versus the U.S. dollar. The windows are defined as follows: the pre-Tokyo fix window starts at 5:00 p.m. ET until the Tokyo fix at 10:00 a.m. local time ("pre-T"), followed by the post-Tokyo window ("post-T") that runs until 2:00 a.m. ET (when European markets open). The pre-ECB fix window ("pre-E") spans the period between European opening hours until the ECB fix at 8:15 a.m. ET. The "interfix" window ("E-L") covers the period between the ECB and the London fix at 11:00 a.m. ET. The final window spans the period after the London fix ("post-L") starting at 11:00 a.m. ET and ending at 5:00 p.m. ET. All times are measured in Eastern Time, taking into account daylight savings time. Data is daily and covers the sample period from January 1999 to December 2018. The number of observations ranges between 934 and 1032. Returns are measured as the average log changes in the mid quote for the respective currency. All numbers are annualised and t-statistics are in parentheses.

	pre-T	post-T	pre-E	E-L	post-L	CTC
Monday	-5.27	5.11	-6.35	2.14	1.50	-2.87
Obs: 934	(-4.54)	(3.77)	(-2.78)	(1.20)	(0.89)	(-0.77)
Tuesday	-7.14	6.35	-1.58	0.25	2.65	0.53
Obs: 1027	(-7.65)	(4.96)	(-0.71)	(0.13)	(1.33)	(0.13)
Wednesday	-4.65	5.97	-1.99	-1.05	5.16	3.45
Obs: 1031	(-5.44)	(4.72)	(-0.96)	(-0.54)	(2.16)	(0.84)
Thursday	-2.70	5.64	-0.90	-3.27	6.93	5.70
Obs: 1011	(-2.58)	(4.28)	(-0.40)	(-1.54)	(3.57)	(1.40)
Friday	-7.51	5.39	-7.27	1.29	5.62	-2.48
Obs: 1006	(-7.66)	(3.80)	(-3.57)	(0.50)	(3.15)	(-0.60)

Table 4. Month of the Year Effects

This table reports annualized average returns for different intraday periods around the Tokyo, ECB and London fix for the dollar portfolio “DOL” which is an equal weighted average all nine currencies in our sample. The sample is split according to the various months of the year, i.e., January through December. Positive values imply the foreign currency portfolio appreciates versus the U.S. dollar. The windows are defined as follows: the pre-Tokyo fix window starts at 5:00 p.m. ET until the Tokyo fix at 10:00 a.m. local time (“pre-T”), followed by the post-Tokyo window (“post-T”) that runs until 2:00 a.m. ET (when European markets open). The pre-ECB fix window (“pre-E”) spans the period between European opening hours until the ECB fix at 8:15 a.m. ET. The “interfix” window (”E-L”) covers the period between the ECB and the London fix at 11:00 a.m. ET. The final window spans the period after the London fix (“post-L”) starting at 11:00 a.m. ET and ending at 5:00 p.m. ET. All times are measured in Eastern Time, taking into account daylight savings time. Data is daily and covers the sample period from January 1999 to December 2018. The number of observations ranges between 384 and 438. Returns are measured as the average log changes in the mid quote for the respective currency. All numbers are annualised and t-statistics are in parentheses.

	pre-T	post-T	pre-E	E-L	post-L	CTC
January	-7.39	7.10	-6.94	-1.24	6.99	-1.48
Obs: 402	(-4.70)	(3.17)	(-1.74)	(-0.36)	(2.15)	(-0.22)
February	-3.70	7.88	-8.92	-2.86	7.57	-0.03
Obs: 384	(-2.70)	(3.37)	(-2.75)	(-0.83)	(2.37)	(-0.00)
March	-3.23	1.94	-4.20	2.87	4.63	2.01
Obs: 438	(-2.00)	(1.01)	(-1.32)	(0.95)	(1.31)	(0.32)
April	-4.81	3.81	-1.92	6.15	8.11	11.35
Obs: 412	(-3.00)	(2.17)	(-0.60)	(2.00)	(3.13)	(1.98)
May	-6.75	5.66	-4.23	-3.78	1.28	-7.81
Obs: 422	(-4.45)	(2.93)	(-1.24)	(-1.18)	(0.42)	(-1.25)
June	-4.81	6.46	-5.19	1.96	7.25	5.66
Obs: 426	(-3.26)	(2.71)	(-1.53)	(0.57)	(2.30)	(0.87)
July	-6.87	6.75	-4.16	0.51	8.85	5.08
Obs: 420	(-4.73)	(3.80)	(-1.40)	(0.15)	(3.42)	(0.86)
August	-7.30	1.80	-1.71	1.97	0.52	-4.71
Obs: 440	(-5.25)	(1.00)	(-0.56)	(0.62)	(0.20)	(-0.86)
September	-7.08	2.80	1.90	3.68	1.03	2.34
Obs: 403	(-4.15)	(1.36)	(0.52)	(1.10)	(0.32)	(0.37)
October	-6.65	5.49	-3.68	0.85	-1.31	-5.29
Obs: 438	(-3.99)	(2.40)	(-1.16)	(0.25)	(-0.43)	(-0.82)
November	-4.71	7.91	-4.29	-6.68	2.14	-5.63
Obs: 408	(-3.12)	(3.69)	(-1.10)	(-2.03)	(0.69)	(-0.85)
December	-2.08	11.42	0.25	-5.97	6.55	10.17
Obs: 416	(-1.33)	(5.99)	(0.08)	(-1.92)	(1.91)	(1.62)

Table 5. Volatility, Liquidity and the *W*

This table reports annualized average return differences for different intraday periods around the Tokyo as well as the ECB and London fix for the EUR, GBP, JPY and the dollar portfolio (DOL) using midquotes from the TRTH data base. The returns are calculated as the average cumulative returns for the reversal around the respective fixes where a short position in the foreign currency is taken before the fix and a long position after the fix. For the Tokyo fix the short position is held from 5:00 p.m. ET to the Tokyo fix at 10:00 a.m. local time (taking DST into account), and the long position is held from the Tokyo fix until 2:00 a.m. ET. For the "Europe" window the short position is held from 2:00 a.m. ET until the ECB fix at 2:15 p.m. local time, and the long position is held starting with the London fix at 4:00 p.m. local time until 5:00 p.m. ET. The three-hour period between the two fixes is dropped. All sample days are sorted according to the previous day's average relative bid-ask spread during the respective window (either Tokyo or Europe), the implied volatility extracted from FX options or the CBOE VIX, respectively. As volatilities and liquidity have a skewed distribution, the differences are calculated as the average returns during the top quintile of volatility and spreads days versus the average returns for the bottom three quintiles. Returns are in percent and annualised; t-statistics are in parentheses. Time is Eastern Time (ET) and the sample periods start in January 1999 and January 2009 for Panels A and B, respectively.

	Panel A: January 1999–December 2018					
	Spreads		FX IV		VIX	
	Tokyo	Europe	Tokyo	Europe	Tokyo	Europe
EUR	13.09	10.25	11.82	11.70	9.81	8.90
	(6.89)	(2.82)	(6.71)	(3.42)	(5.21)	(2.44)
GBP	9.64	8.09	7.76	8.64	5.44	4.80
	(5.13)	(2.13)	(4.44)	(2.46)	(2.79)	(1.25)
JPY	5.02	1.16	10.08	3.62	6.03	-0.26
	(1.76)	(0.32)	(3.76)	(1.11)	(2.26)	(-0.07)
DOL	8.41	6.61	7.20	7.64	5.21	3.77
	(5.13)	(2.24)	(4.63)	(2.80)	(3.11)	(1.28)
	Panel B: January 2009–December 2018					
	Spreads		FX IV		VIX	
	Tokyo	Europe	Tokyo	Europe	Tokyo	Europe
EUR	14.35	12.56	13.55	12.39	11.63	12.62
	(5.40)	(2.53)	(5.26)	(2.63)	(4.45)	(2.45)
GBP	9.63	6.36	8.91	6.79	7.38	1.94
	(3.54)	(1.13)	(3.36)	(1.33)	(2.65)	(0.34)
JPY	9.96	8.16	12.86	9.49	3.86	2.28
	(2.57)	(1.68)	(3.55)	(2.24)	(1.03)	(0.47)
DOL	10.73	6.07	8.89	5.58	7.41	2.77
	(4.47)	(1.45)	(3.86)	(1.44)	(3.02)	(0.65)

Table 6. Trading the *W* for EUR, GBP and JPY Including Transaction Costs
This table reports annualized average returns for different intraday periods around the Tokyo, ECB and London fix for the EUR, GBP and JPY taking transaction costs into account. This means, foreign currencies are bought at the bid and sold at the ask. Returns already take the direction of trading into account, i.e. before the fixes we take a short position that is reversed afterwards. We only report results for selected windows: The window for the EUR and the GBP around the ECB and London fix, respectively, starts at 2:00 a.m. ET and ends at 5:00 p.m. ET; the window for the JPY around the Tokyo fix starts at 5:00 p.m. and ends at 2:00 a.m. Data is daily and covers a sample period from January 1999 to December 2018 (5009 observations). Returns are annualised and take into account daily trading, and transaction costs are measured using the indicative bid-ask spreads. In addition to the results based on the indicative quotes we also report returns and Sharpe Ratios using tighter bid-ask spreads (reduced by 50%) as well as results ignoring transaction costs. Data from CME ($BA^{100\%}_{CME}$) is daily and covers the sample period from January 2009 to December 2018 (2515 observations). As quotes are firm and tradeable, full transaction costs are accounted for.

Panel A: Gross Returns

	EUR			GBP			JPY		
	pre-E	post-E	pre/post-E	pre-L	post-L	pre/post-L	pre-T	post-T	pre/post-T
$BA^{100\%}$	0.04	-0.59	-0.55	-1.92	-0.82	-2.74	-4.91	-1.22	-6.13
$BA^{50\%}$	3.56	3.04	6.60	1.52	2.73	4.24	-0.12	3.49	3.37
$BA^{0\%}$	7.08	6.68	13.75	4.96	6.27	11.23	4.67	8.21	12.87
$BA^{100\%}_{CME}$	5.53	0.58	6.11	0.06	-4.89	-4.83	-11.23	2.41	-8.82

Panel B: Sharpe Ratios

	EUR			GBP			JPY		
	pre-E	post-E	pre/post	pre-L	post-L	pre/post	pre-T	post-T	pre/post
$BA^{100\%}$	0.00	-0.09	-0.06	-0.27	-0.21	-0.33	-1.55	-0.28	-1.12
$BA^{50\%}$	0.62	0.41	0.70	0.20	0.65	0.50	-0.06	0.77	0.60
$BA^{0\%}$	1.24	0.90	1.47	0.67	1.50	1.32	1.44	1.82	2.33
$BA^{100\%}_{CME}$	0.99	0.08	0.65	0.01	-0.99	-0.51	-2.25	0.52	-1.27

VIII. Appendix: Figures

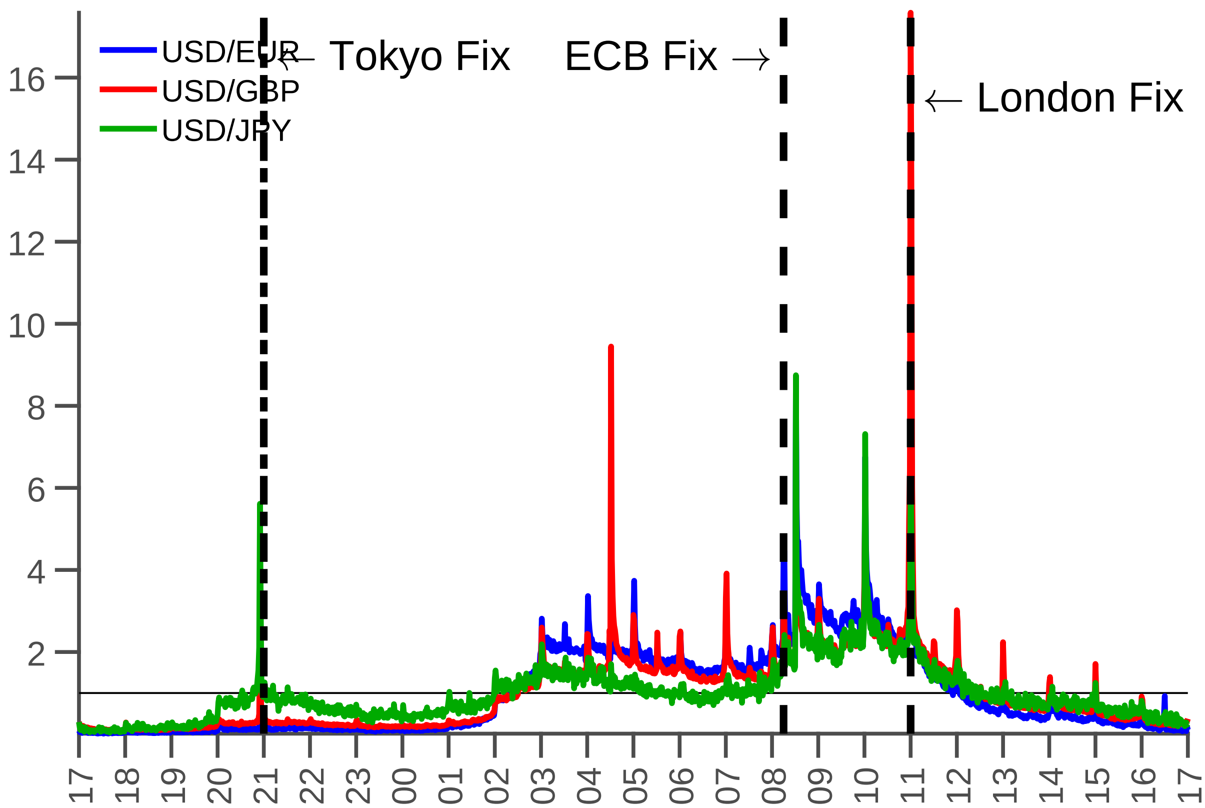

Figure 2. Normalised Volumes for EUR, GBP and JPY
This figure shows normalised volume over the course of the trading data, based on traded volume recorded on Thomson Reuters *Matching* interdealer trading platform (D5). The volume is computed for every one-minute interval normalised by the total volume on the day. The plot shows the average normalised volume over the sample period from January 2006 to December 2018.

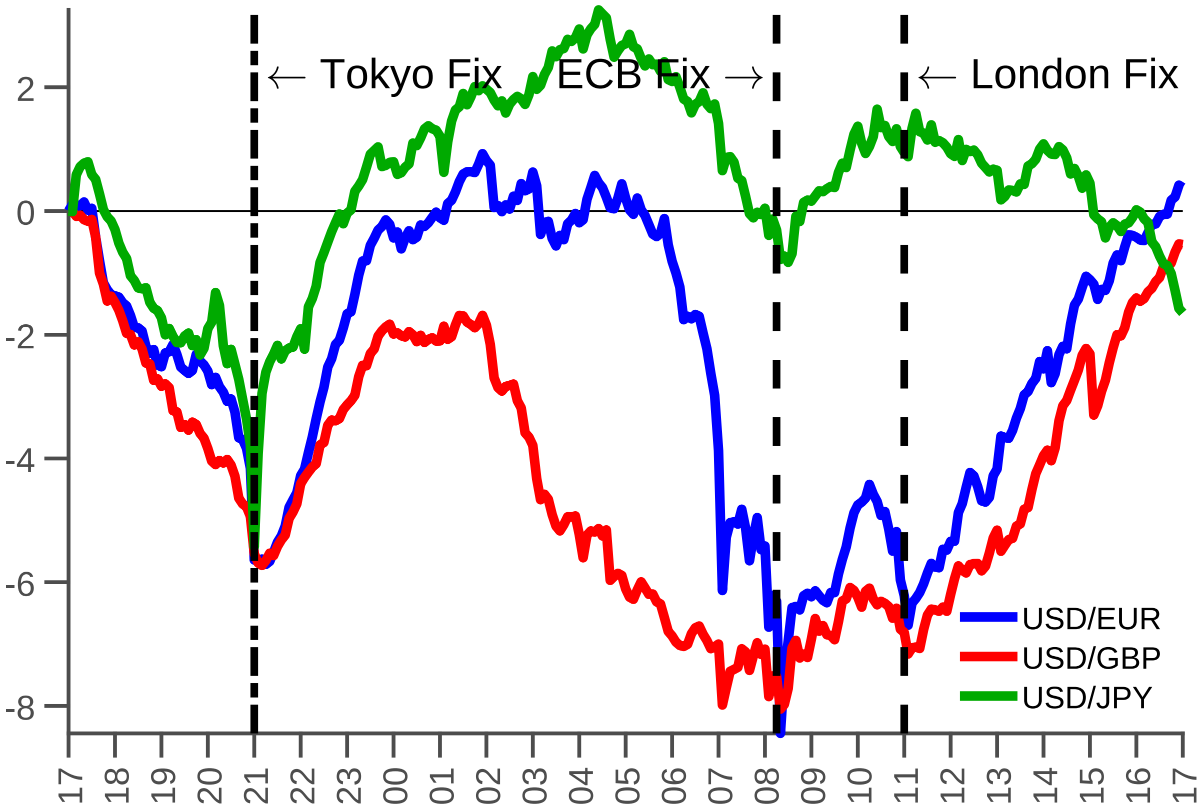

Figure 3. Cumulative 5-min Returns for EUR, GBP and JPY

This figure displays cumulative average annualised 5-min returns ($-\Delta s$) over the course of a trading day for the USD/EUR (blue), USD/GBP (red), and USD/JPY (green) pairs, respectively. An increase means the foreign currency appreciates against the U.S. dollar. The three black dashed lines at 9:00 p.m., 8:15 a.m., and 11:00 a.m. refer to the Tokyo fix, the ECB fix and the London fix, respectively. All returns are annualised and expressed in percent. The time is measured in Eastern Time (ET). The sample period is January 1999 to December 2018.

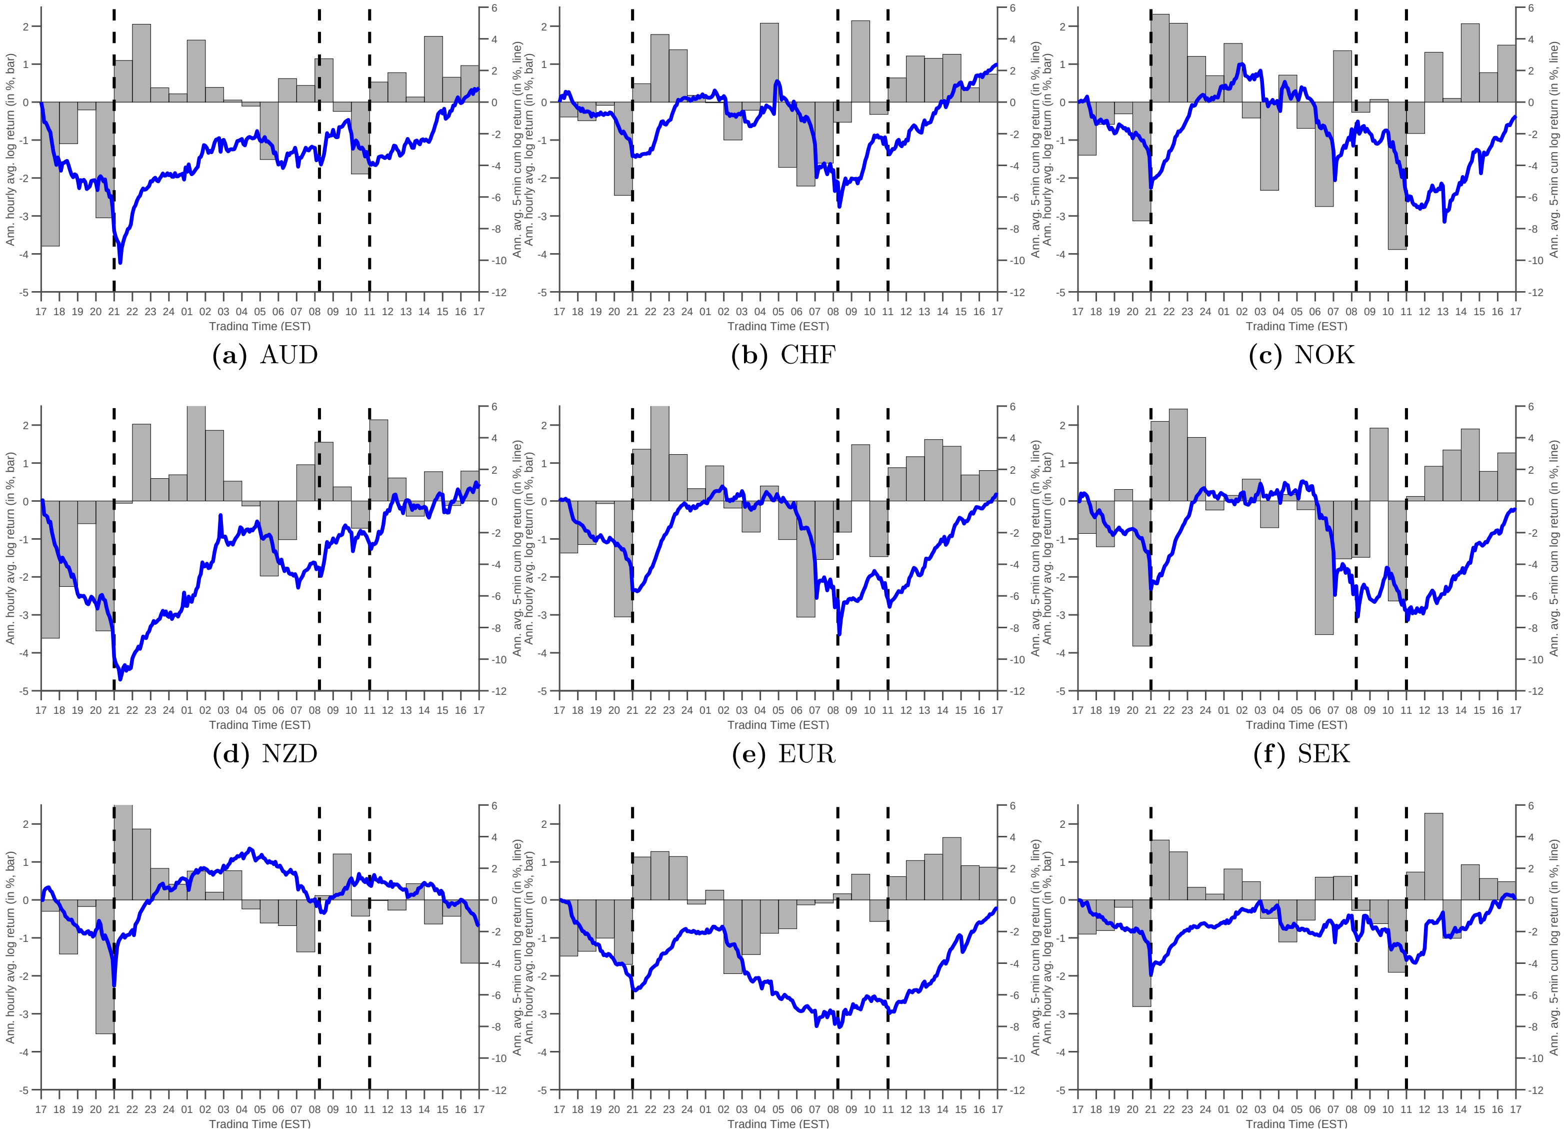

(a) AUD **(b)** CHF **(c)** NOK

(d) NZD **(e)** EUR **(f)** SEK

(g) JPY **(h)** GBP **(i)** CAD

Figure 4. Hourly and Cumulative 5-min Returns for G9 Currencies

This figure displays cumulative average annualized 5-min returns ($-\Delta s$) over the course of a trading day. An increase means the foreign currency appreciates against the U.S. dollar. The three black dashed lines at 9:00 p.m., 8:15 a.m., and 11:00 a.m. refer to the Tokyo fix, the ECB fix and the London fix, respectively. All returns are annualised and expressed in percent. The time is measured in Eastern Time (ET). The sample period is January 1999 to December 2018.

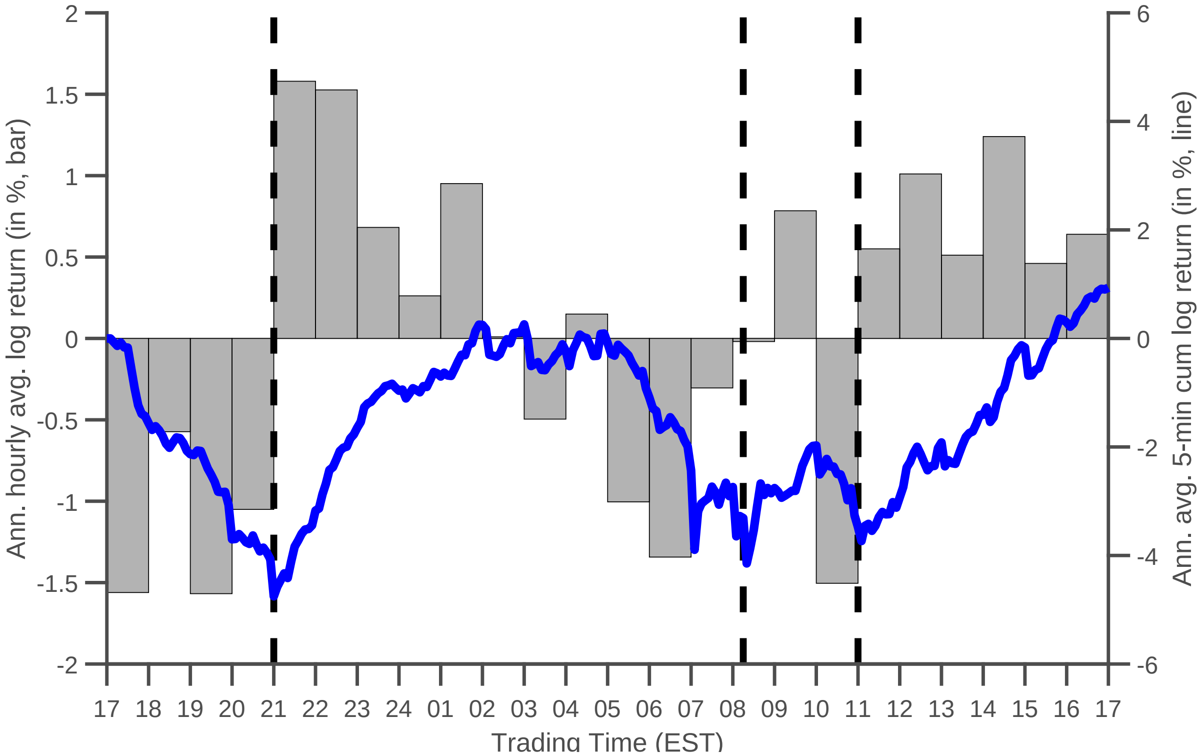

Figure 5. Hourly and Cumulative 5-min Returns for the Dollar Portfolio (DOL)
This figure displays cumulative average annualised 5-min returns ($-\Delta s$) over the course of a trading day. An increase means the foreign currency appreciates against the U.S. dollar. The three black dashed lines at 9:00 p.m., 8:15 a.m., and 11:00 a.m. refer to the Tokyo fix, the ECB fix and the London fix, respectively. All returns are annualised and expressed in percent. The time is measured in Eastern Time (ET). The sample period is January 1999 to December 2018.

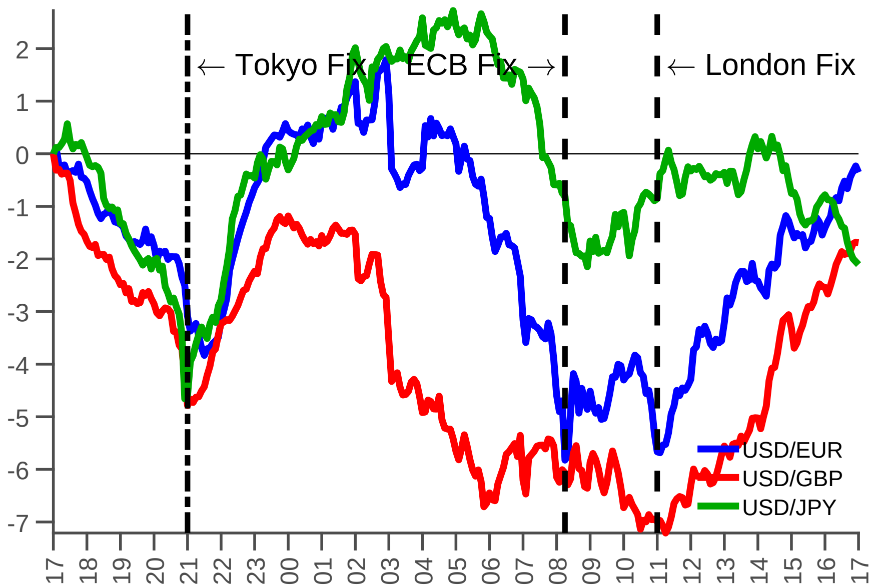

(a) D5

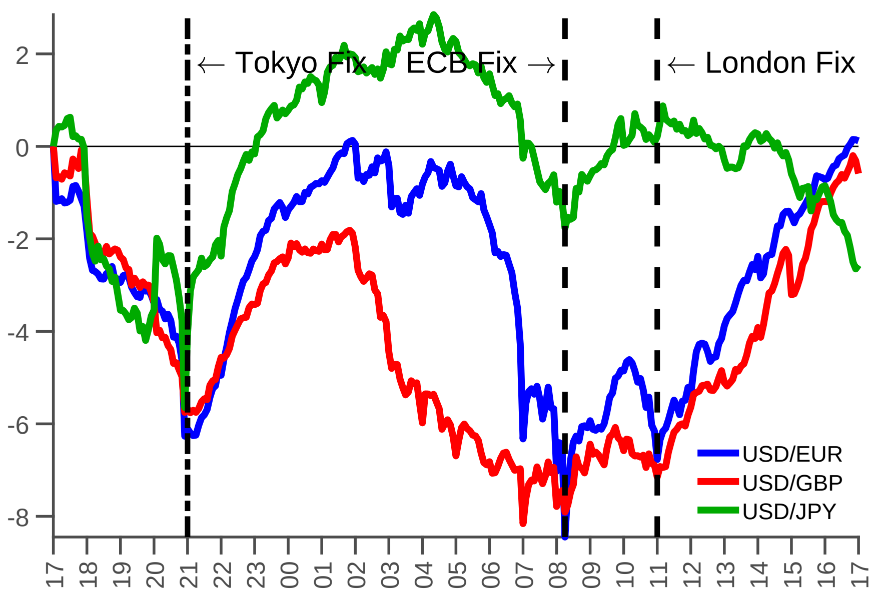

(b) CME

Figure 6.

Cumulative 5-min Returns computed from VWAP (D5) and from FX Futures Quotes (CME)
This figure displays cumulative average annualized 5-min returns ($-\Delta s$) over the course of a trading day calculated using value-weighted average prices from D5 (Panel (a)) and firm FX futures quotes from CME (Panel (b)), respectively. An increase means the foreign currency appreciates against the U.S. dollar. The three black dashed lines at 9:00 p.m., 8:15 a.m., and 11:00 a.m. refer to the Tokyo fix, the ECB fix and the London fix, respectively. All returns are annualised and expressed in percent. The time is measured in Eastern Time (ET). The sample period is January 1999 to December 2018.

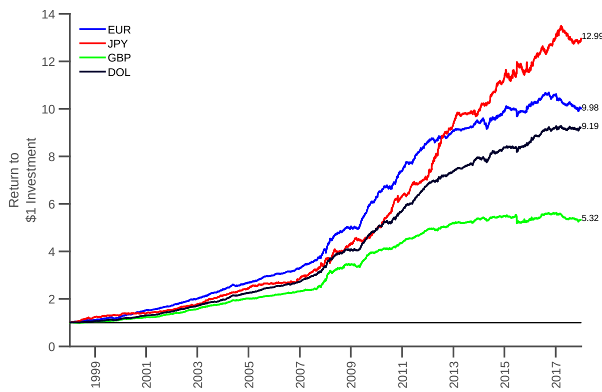

(a) Total Return Index: Tokyo Fix

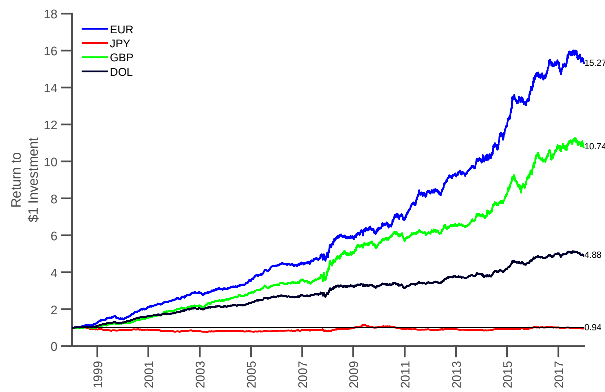

(b) Total Return Index: ECB/London Fix

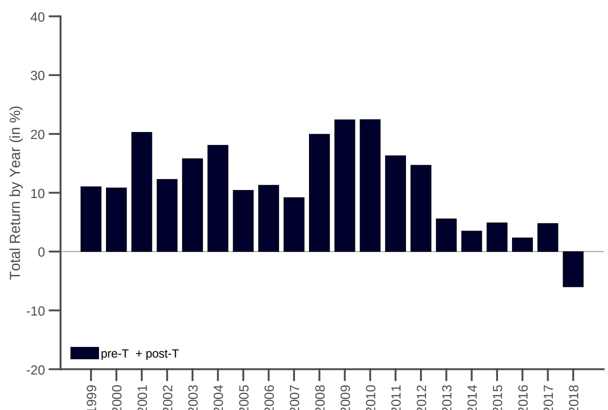

(c) DOL around Tokyo Fix

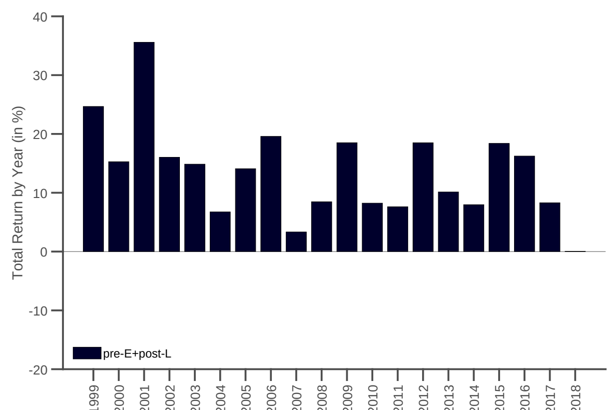

(d) DOL around ECB/London Fix

Figure 7. Total Return Indices and Year-By-Year Performance: Trading the *W*
The figures show the performance of a trading strategy around the Tokyo and ECB/London fix, where an investor takes a short-position during the pre-fix and a long-position during the post-fix, respectively. The top panel shows the total return indices for the three major currencies (EUR, GBP, JPY) and the dollar portfolio (DOL) with an initial investment of one U.S. dollar. The bottom panel shows the total returns split year by year for the dollar portfolio. The sample period is January 1999 to December 2018.

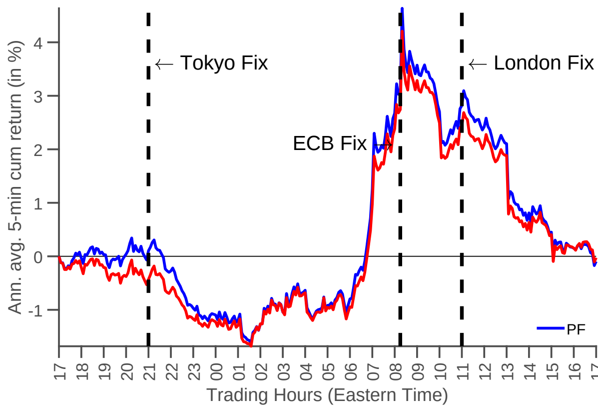

(a) EUR portfolio

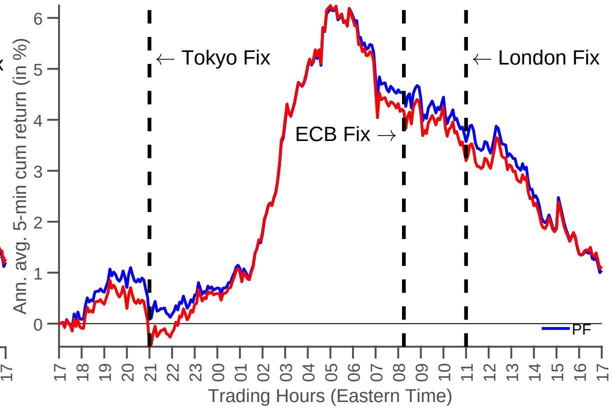

(b) GBP portfolio

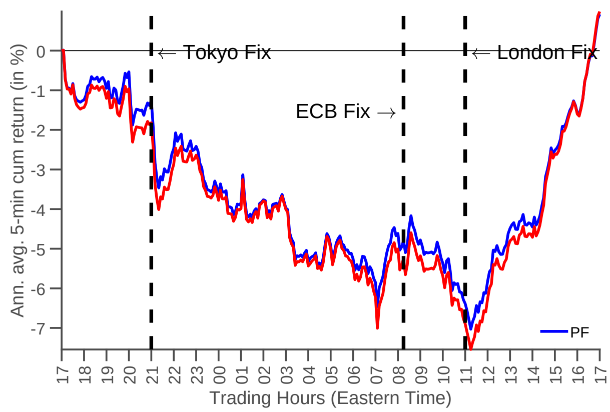

(c) JPY portfolio

Figure 8. Cumulative 5-minute returns: Alternative Base Currency Portfolios
These figures display cumulative average annualized 5-min returns ($-\Delta s$) over the course of a trading day for base currency portfolios, whereby the euro (EUR), British pound (GBP), and Japanese yen (JPY) serve as base currency, respectively. Base currency portfolios are computed by calculating the unconditional average of returns denominated in a non-U.S. dollar currency (i.e. EUR or GBP or JPY). A positive return implies foreign currencies appreciate vis-à-vis the non-U.S. dollar base currency. The time is measured in Eastern Time (ET). The sample period is January 1999 to December 2018.

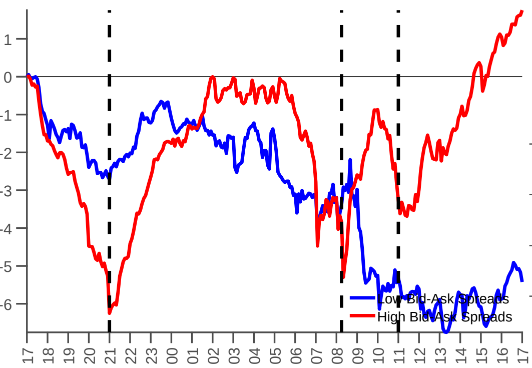

(a) DOL

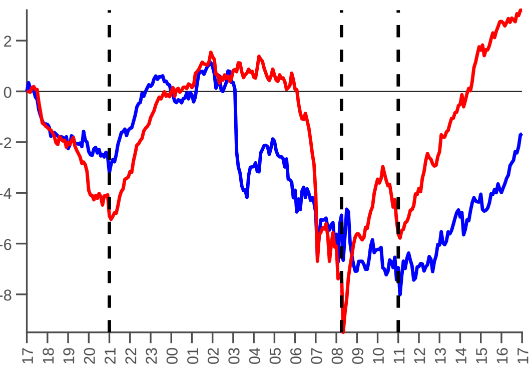

(b) EUR

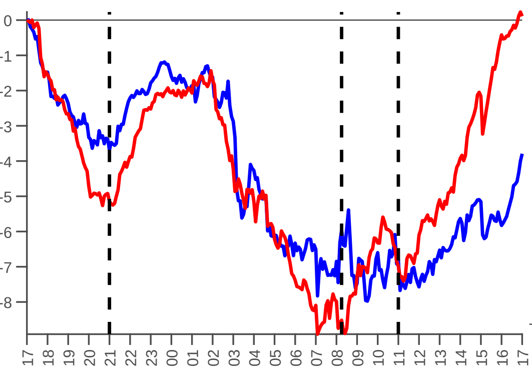

(c) GBP

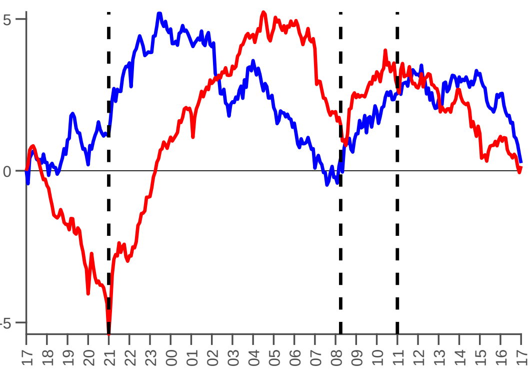

(d) JPY

Figure 9. Cumulative Returns and Average Bid-Ask Spreads
The figure displays cumulative average annualized 5-min returns ($-\Delta s$) over the course of a trading day for returns sorted on the previous day's average relative bid-ask spreads. The sample is split into quintiles and the figures plot the intraday cumulative returns for the highest and the three lowest quintiles, respectively. An increase means the foreign currency appreciates against the U.S. dollar. The three black dashed lines at 9:00 p.m., 8:15 a.m., and 11:00 a.m. refer to the Tokyo fix, the ECB fix and the London fix, respectively. Time is measured in Eastern Time (ET) and the sample period is January 1999 to December 2018.

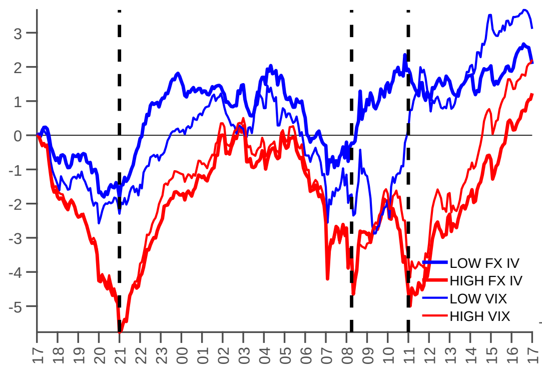

(a) DOL

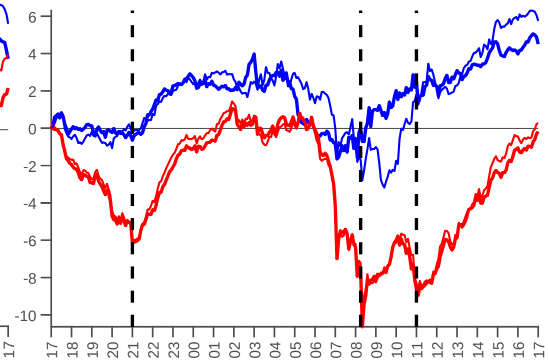

(b) EUR

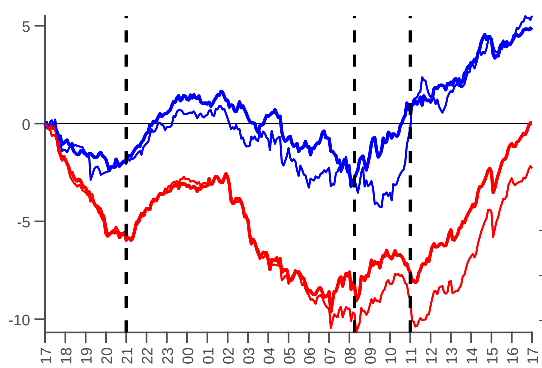

(c) GBP

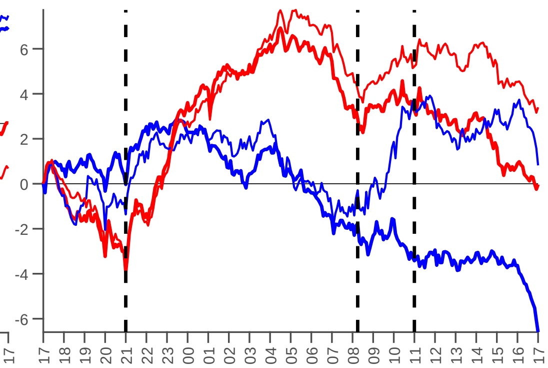

(d) JPY

Figure 10. Cumulative Returns and Time-Varying Volatility

The figure displays cumulative average annualized 5-min returns ($-\Delta s$) over the course of a trading day for returns sorted either on the previous day's model free implied volatility extracted from currency option prices (thick lines) or the previous day's CBOE VIX (extracted from options on the S&P 500 index). The sample is split into quintiles and the figures plot the intraday cumulative returns for the highest and the three lowest quintiles, respectively. An increase means the foreign currency appreciates against the U.S. dollar. The three black dashed lines at 9:00 p.m., 8:15 a.m., and 11:00 a.m. refer to the Tokyo fix, the ECB fix and the London fix, respectively. Time is measured in Eastern Time (ET) and the sample period is January 1999 to December 2018.

References

Andersen, T. G., T. Bollerslev, F. X. Diebold, and C. Vega (2003): "Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange," *American Economic Review*, 93(1), 38–62.

Bansal, R., and A. Yaron (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59, 1481–1509.

BIS (2016): "Triennial Central Bank Survey–Foreign exchange and OTC derivatives markets in 2016," Discussion paper, Bank for International Settlements.

——— (2018): "Monitoring of Fast-Paced Electronic Markets," Discussion paper, Bank for International Settlements.

——— (2019): "Triennial Central Bank Survey–Foreign Exchange Turnover in April 2019," Discussion paper, Bank for International Settlements.

Breedon, F., and A. Ranaldo (2013): "Intraday Patterns in FX Returns and Order Flow," *Journal of Money, Credit and Banking*, 45(5), 953–965.

Breedon, F., and P. Vitale (2010): "An Empirical Study of Portfolio-Balance and Information Effects of Order Flow on Exchange Rates," *Journal of International Money and Finance*, 29(3), 504–524.

Carr, P., and L. Wu (2009): "Variance Risk Premiums," *Review of Financial Studies*, 22(3), 1311–1341.

Chen, H., G. Noronha, and V. Singal (2004): "The price response to S&P 500 index additions and deletions: Evidence of asymmetry and a new explanation," *The Journal of Finance*, 59(4), 1901–1930.

Cornett, M. M., T. V. Schwarz, and A. C. Szakmary (1995): "Seasonalities and intraday return patterns in the foreign currency futures market," *Journal of Banking and Finance*, 19(5), 843–869.

Duffie, D. (2010): "Presidential Address: Asset Price Dynamics with Slow-Moving Capital," *The Journal of Finance*, 65(4), 1237–1267.

ECB (2019): "Framework for the euro foreign exchange reference rates," Discussion paper, European Central Bank.

Evans, M. (2018): "Forex trading and the WMR Fix," *Journal of Banking and Finance*, 87, 233–247.

Evans, M. D., P. O'Neill, D. Rime, and J. Saakvitne (2018): "Fixing the Fix? Assessing the Effectiveness of the 4pm Fix Benchmark," Working Paper.

FSB (2014): "Foreign Exchange Benchmarks," Discussion paper, Financial Stability Board, Financial Stability Board, Final Report.

Gargano, A., S. J. Riddiough, and L. Sarno (2018): "The Value of Volume in Foreign Exchange," Working Paper.

Gilmore, S., and F. Hayashi (2011): "Emerging Market Currency Excess Returns.," *American Economic Journal: Macroeconomics*, 3(4), 85 – 111.

Grossman, S. J., and M. H. Miller (1988): "Liquidity and market structure," *the Journal of Finance*, 43(3), 617–633.

Ito, T., and M. Yamada (2016): "Puzzles in the Forex Tokyo "Fixing": Order Imbalances and Biased Pricing by Banks," Discussion paper, Center on Japanese Economy and Business, No. 352.

Jiang, Z. (2017): "Currency Returns in Different Time Zones," Working Paper, Northwestern University.

Lou, D., H. Yan, and J. Zhang (2013): "Anticipated and Repeated Shocks in Liquid Markets," *The Review of Financial Studies*, 26(8), 1891–1912.

Lustig, H., N. Roussanov, and A. Verdelhan (2011): "Common Risk Factors in Currency Markets," *Review of Financial Studies*, 24(11), 3731–3777.

Melvin, M., and J. Prins (2015): "Equity hedging and exchange rates at the London 4 p.m. fix," *Journal of Financial Markets*, 22, 50–72.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2012): "Currency Momentum Strategies," *Journal of Financial Economics*, 106, 660–684.

Mueller, P., A. Stathopoulos, and A. Vedolin (2016): "International Correlation Risk," *forthcoming, Journal of Financial Economics.*

Nagel, S. (2012): "Evaporating liquidity," *The Review of Financial Studies*, 25(7), 2005–2039.

Ranaldo, A. (2009): "Segmentation and time-of-day patterns in foreign exchange markets," *Journal of Banking and Finance*, 33, 2199–2206.

Sigaux, J.-D. (2018): "Trading ahead of Treasury auctions," ECB Working Paper Series, No 2208.

Thomson Reuters (2017): "WM/Reuters FX Benchmarks," Technical Report, Thomson Reuters.

Vayanos, D., and J. Wang (2009): "Liquidity and asset prices: A unified framework," Discussion paper, National Bureau of Economic Research.