# WORKING PAPER

## Microstructure of foreign exchange markets

NORGES BANK

# Microstructure of Foreign Exchange Markets[*]

Martin D.D. Evans[†]        Dagfinn Rime[‡]

February 2019

## Abstract

This article presents an overview of research on the Microstructure of Foreign Exchange Markets. We begin by summarizing the institutional features of FX trading and describe how they have evolved since the 1980s. We then explain how these features are represented in microstructure models of FX trading. Next, we describe the links between microstructure and traditional macro exchange-rate models and summarize how these links have been explored in recent empirical research. Finally, we provide a microstructure perspective on two recent areas of interest in exchange-rate economics: the behavior of returns on currency portfolios, and questions of competition and regulation.

---

# 1 Introduction

Research on the Microstructure of Foreign Exchange Markets examines the behavior of foreign exchange (FX) instruments (such as spot and forward) in an environment that replicates the key features of trading in the actual market. Traditional macro exchange-rate models pay little attention to how trading in the FX market takes place. The implicit assumption in these models is that the details of trading (i.e., who quotes prices and how trade takes place) are unimportant for the behavior of FX prices over months, quarters or longer. Microstructure models, by contrast, examine how information relevant to the trading decisions of market participants becomes embedded in the prices of FX instruments via the trading process. According to this perspective, trading is not an ancillary market activity that can be ignored when considering how the prices of FX instruments behave. Rather, trading is an integral part of the process through which the prices of FX instruments are determined and evolve through time.

The past two decades have witnessed rapid growth in the volume of FX microstructure research. Early papers in the literature focused on partial equilibrium models that captured the key features of FX spot trading. These models provided a new and rich array of empirical predictions that were confirmed in subsequent empirical research. This research provided a new perspective on the proximate drivers of exchange rates over short horizons, ranging from a few minutes to a few weeks. More recent research has moved away from the traditional partial equilibrium domain of microstructure models to focus on the links between FX trading and macroeconomic conditions. The goal of this research is to provide a set of micro-foundations for the exchange-rate dynamics that respect the institutional setting of the market which have been missing from general equilibrium macro models.

In this article we first provide a broad overview of the FX market, noting key features that are incorporated into microstructure models. This discussion updates earlier reviews in Osler (2009), Evans (2011) and Evans and Rime (2012). Next, we consider the links between microstructure and macro exchange-rate models. Here we examine how research in FX microstructure complements rather than competes with research based on traditional models. Finally, we provide a microstructure perspective on two recent areas of interest: (i) returns on currency portfolios, and (ii) FX trading, competition and regulation.

# 2 Microstructure Models and Market Structure

Microstructure models of the foreign exchange market attempt to incorporate the features of trading that are essential to understanding the economic processes generating movements in FX prices and trading patterns. In particular, the models allow us to examine how information is transmitted from one participant to another as trading takes place, and to study how this information transmission process ultimately leads to the determination of spot exchange rates. The fact that the models describe, in detail, how trading takes place between participants does not mean that the researchers using these models are only interested in trading. Their focus remains on understanding exchange-rate dynamics, but they are using models that make detailed predictions about trading activity as well.

## 2.1 Market Structure

No model can incorporate all the institutional features of trading in the FX market – it is far too complex. Furthermore, over time there have been considerable changes in the ways that trading takes place, both in terms of market participants and the venues in which they trade. The key and ongoing challenge facing microstructure researchers is to "see through" these institutional changes so as to focus on a small number of features that are essential for understanding the main economic mechanisms at work. Here we provide a broad overview of trading in the market and describe how key aspects of trading have changed.

Over time trading in the FX market has involved several different types of participants. The FX dealers working at major banks have always been the central intermediaries. More specifically, they play the role of liquidity suppliers who are willing to be the counter-parties to trades initiated by others. End-users comprise the other principal type of participant. This group comprises the financial and non-financial customers of dealer-banks, as well as other (non-dealer) entities that trade on electronic systems. Trading by financial end-users is generated by allocative, speculative and risk management motives. Trading by non-financial end-users is generated by goods trade, cross-border investment and risk management.

Changes in the composition of trading by market participants are shown in Figure 1. Here we plot total spot trading volume and the volume shares for different groups of participants reported in the Triennial Survey of Global FX Markets coordinated by the Bank for International Settlements (BIS). The figure shows that global volume has grown from USD 430 bn in 1992, to a peak of USD 2000 bn in 2013. The share involving non-financial end-users was relatively stable at around 15 percent, before falling to around eight percent in the last few years. Dealer-banks participated in 70

Figure 1: Spot FX. Total volume (USD bn, rhs) and counterparties share of volume (percent, lhs)
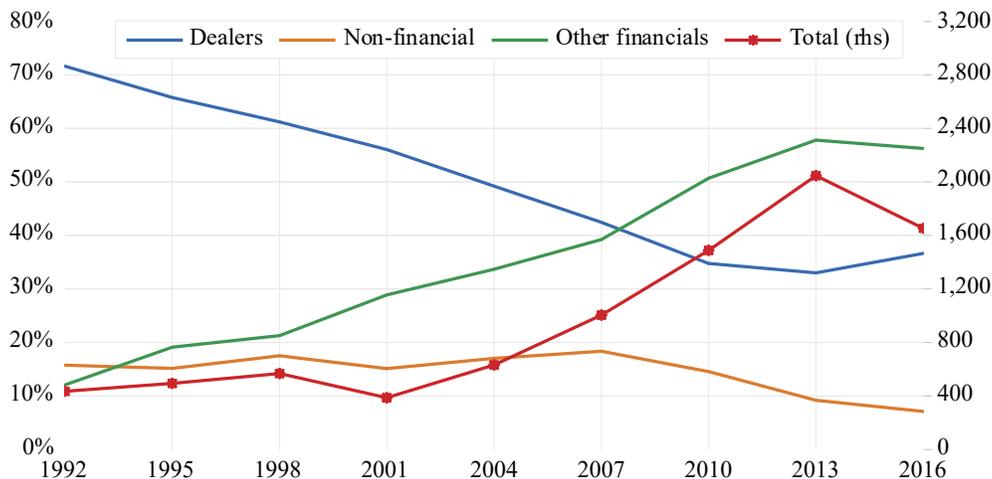


Note: Graph shows total global FX spot volume, in USD billions, in red-dot marked line, on right hand side. Other lines show the market share of three key counter-party groups. Source: BIS Triennial FX survey of FX dealing banks.

percent of all volume in 1992, but their share has been steadily trending downward until recently, leveling off at around 35 percent. The decline of dealer-banks' participation has been accompanied by an increase in the share of trading by financial end-users; which has risen from 12 percent in 1992 to 56 percent in 2016. This shift in the trading shares of dealer-banks and financial end-users reflects the increasing importance of FX as a separate asset class and the growth in electronic trading.

Figure 2 provides an overview of how the institutional structure of FX trading has changed between the 1980s and the present day. Until the 1980s FX trading took place in a two-tier market comprising the interbank and retail tiers: Within the interbank tier, trades took place directly between the dealers working at the trading desks of major banks (shown by link 1), and indirectly via voice brokers that matched dealer counterparties (shown by link 2).[1] Trades between dealers and their customers took place in the retail tier of the market (3). Trades in both tiers were either conducted by phone calls or electronically (via a form of email) between counterparties located all around the world (i.e., (5)), but most trading in the interbank tier took place between dealers located in a few financial centers: Tokyo, Singapore, Frankfurt, New York, and particularly London. Trading took place 24 hours a day, but activity was heavily concentrated during the daytime hours

---

[1]Hereafter, we simply refer to link $x$ in Figure 2 by $(x)$.

of the main financial centers.

Since the mid-1990s most interbank trade has taken the indirect form in which dealers submit market and limit orders to buy and sell currencies to electronic brokerages (4). The main brokerages, run by EBS and Reuters, were introduced in 1992. These systems prioritize limit orders so that those with the best prices are matched first with incoming market orders. Electronic brokers made inter-dealer risk sharing more efficient, requiring fewer trades to efficiently allocate a given volume. They also allowed dealers to more easily manage the risks associated with filling large customer orders within the constraints they face on both the duration and size of their FX positions.[2]

Panel (c) in Figure 2 shows how the market structure became more complex in the last decade. The trading structure has become more fragmented as in other financial markets (e.g. equities). End-users can now trade on Multi-bank platforms (6), through Prime Brokerage accounts on the EBS and Reuters electronic limit order books (7,8), and on platforms that stream/aggregate prices from other platforms, so-called Retail Aggregators (10). The major banks have responded to these developments by developing their own electronic platforms (i.e., Single-bank platforms (9)) in order to capture end-users' business, and by providing liquidity to other venues, like the Multi-bank platforms.

Figure 3 plots data published by the Federal Reserve Bank of New York on FX execution methods over the past two decades. The plots clearly show that Single-bank and Multi-bank platforms have become the dominant trading venues, while the electronic brokers run by EBS and Reuters have lost market share. These developments have reduced trading transparency because far less information is available across the market in real time from Single-bank platforms than can be obtained from the EBS and Reuters systems. These changes have also had an impact on market liquidity, as we discuss below.

Despite the introduction of new trading venues that are accessible to end-users, dealer-banks remain the key liquidity providers in FX trading. Table 1 shows the distribution of spot trading for a major bank in 2012. The rightmost column shows that trading on Single-bank and Multi-bank platforms together make up more than 50 percent of this bank's trading. The two left columns, covering Major and Medium-sized banks, shows that other banks are active both at Single-bank platforms, Multi-bank platforms, in addition to traditional interbank trading.

The continued importance of dealer-banks as liquidity suppliers reflects several factors. Although end-users now have access to many electronic trading platforms that in principle allow them to by-pass dealer-banks and trade end-user to end-user, the liquidity available on these plat-

---

[2]For example, dealers typically face strict limits on the size of their overnight positions.

Figure 2: Evolution of FX market structure

(a) 1980s



(b) Early and mid-1990s

(c) 2010
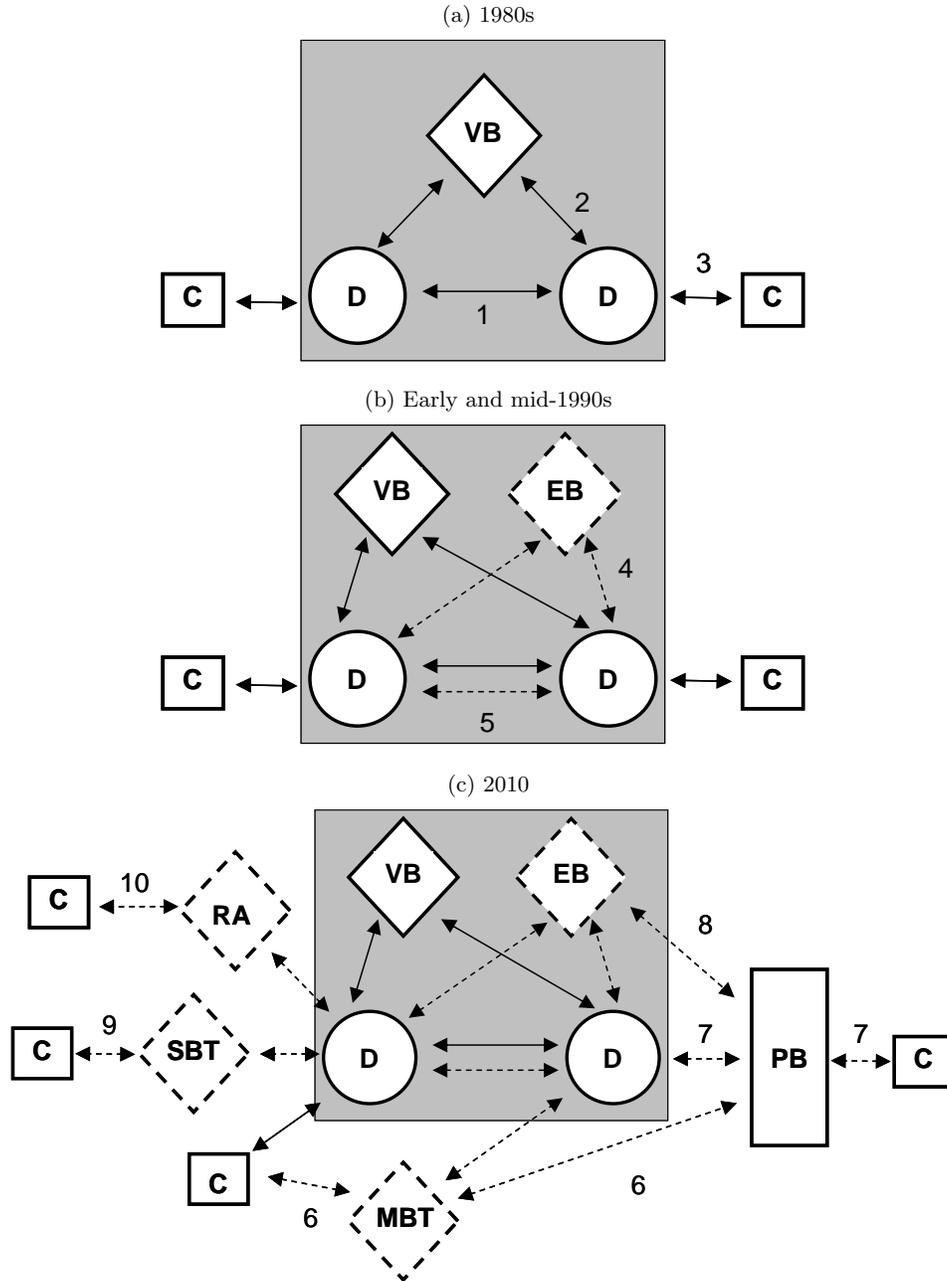
Notes: D=dealer, C=client, VB=voice broker, EB=electronic broker, PB=prime broker, MBT = multi-bank trading system, SBT=single-bank trading system, RA=retail aggregator. Solid lines represent voice execution methods. Dashed lines represent electronic execution methods. Source: King et al. (2012)

Figure 3: Spot FX. Share of execution method for U.S. market
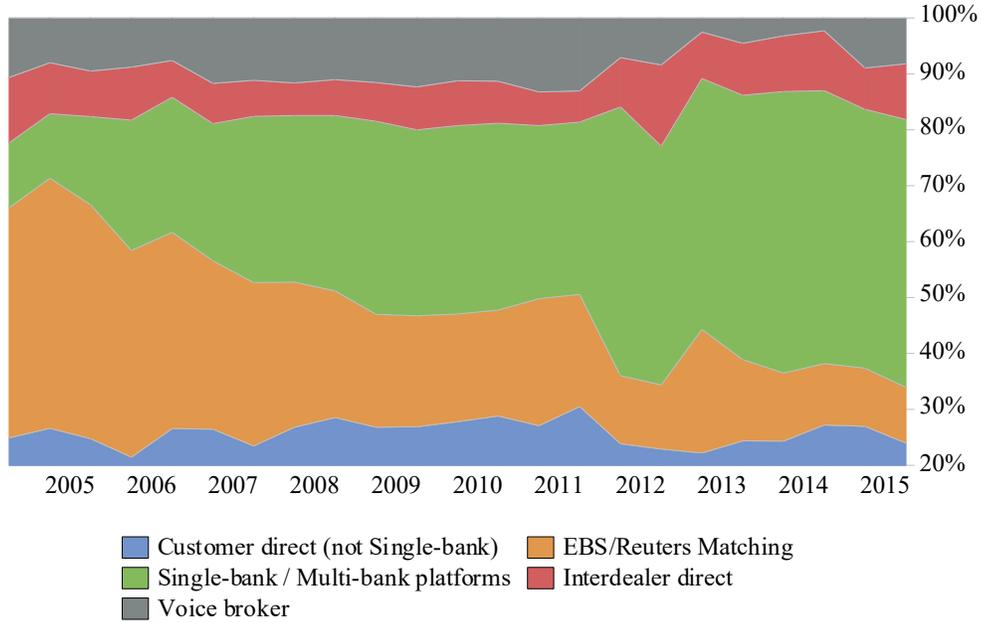


Note: Graph show distribution of FX spot volumes in U.S. market on different execution methods. Source: Bi-annual survey by New York Fed FX Committee.

forms is limited. This can be seen in Figure 4. Here we plot the size distribution for voice and electronic trades in 2013 across major banks. The figure shows that almost 40 percent of the electronic trades are at the USD 5 million limit imposed by the platforms. Trades above USD 5 million are executed by voice between end-users and dealer-banks. Notice, also, that roughly 30 percent of voice trades are for sizes of at least USD 50 million. These data emphasize the fact that dealer-banks are still the dominant providers of liquidity to end-users who want to purchases or sell large volumes of foreign currencies. They only face competition in liquidity provision from Multi-bank platforms and Retail Aggregators for small trades. Thus, an end-user wanting to purchase or sell large amounts of FX must still trade with a dealer-bank in much the same way as they did in the 1980s.

## Table 1: Execution Method and End-User Groups

|  | Major banks | Medium banks | Real Money | Leveraged Investors | Large Corporates | Small Corporates | Brokers | Small banks | Total |
|---|---|---|---|---|---|---|---|---|---|
| Direct | 0 | 0 | 14 | 2 | 25 | 69 | 0 | 2 | 1 |
| SBT | 5 | 22 | 33 | 1 | 18 | 30 | 14 | 65 | 12 |
| MBT | 51 | 21 | 53 | 46 | 57 | 1 | 25 | 14 | 39 |
| Prime brokered | 2 | 30 | 0 | 50 | 0 | 1 | 60 | 3 | 21 |
| Interbank | 42 | 27 | 0 | 0 | 1 | 0 | 0 | 16 | 26 |
| Total | 53 | 10 | 1 | 5 | 1 | 0 | 24 | 5 |  |

Note: Table shows the distribution of trading volume for a large bank in 2012 across venues and customer type. Columns shows how the different customer types have traded across venues, while the Total at bottom shows the relative importance of each counter-party group with this bank. The rightmost column marked total split the bank's total volume across the venues. SBT and MBT is short for Single-bank and Multi-bank trading platforms. Source: (Bjønnes and Kathiziotis, 2016).

## Figure 4: Spot Trade Size Distribution (Major Currencies)
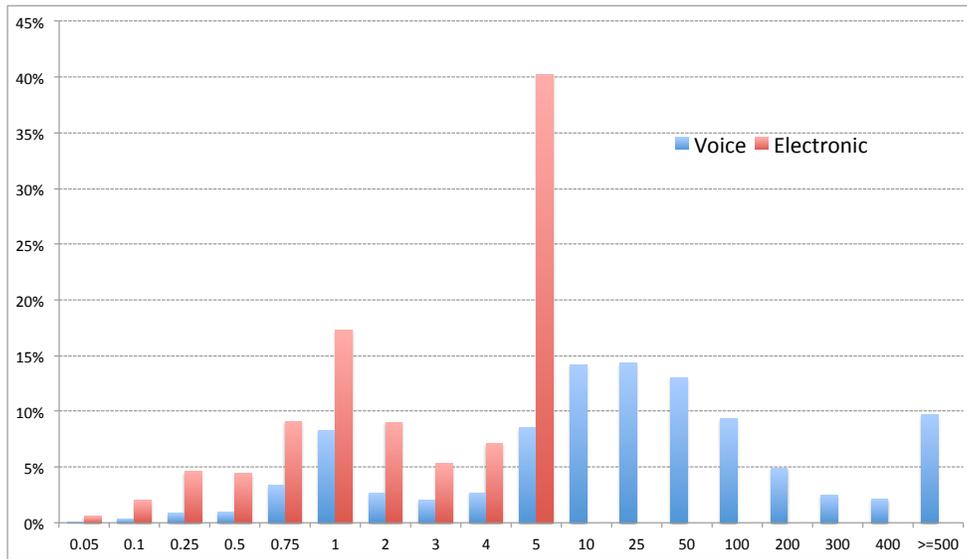


Notes: The figure plots the faction of spot voice and electronic trading volume in five major currency pairs (EURUSD, GBPUSD, USDJPY, USDCAD and USDAUD) across different trade sizes, ranging from USD 0.05 million, to at least USD 500 million.

Dealer-banks are only willing to act as liquidity suppliers to end-users wishing to execute large trades if they can profitably manage the associated risks. Historically, these risks were managed by combining the information in the bank's customer order flow (i.e., the flow of orders to buy and sell FX from the banks' customers) with information on prices and trades from the interbank tier of the

market. This information processing problem became more challenging as end-users gained access to Multi-bank platforms and Retail Aggregators, because these platforms diverted customer flows from dealer-banks. The introduction of Single-bank platforms allowed individual banks to counter this trend, but in so doing they also introduced competitors to the EBS and Reuters platforms. As we noted above, the net result of these competitive pressures has been a reduction in market-wide transparency. Dealer-banks with successful Single-bank platforms have managed to retain customer order flows that provide them with the necessary information to manage risk effectively, but their actions have reduced how representative EBS and Reuters are for market-wide trading conditions. With larger volumes netted internally within the major banks, it has become harder for dealer-banks without successful Single-bank platforms to profitably manage the risks of supplying liquidity across the market.

## 2.2 Microstructure Models

The first models of FX trading considered the situation facing a single dealer in a partial equilibrium setting (see, e.g., Lyons, 1995). Subsequent models examine the trading decisions of a large number of dealers in a competitive market setting. These models incorporate the institutional features of the trading in the 1980s: they distinguish between trades between dealers in the interbank tier, and trades between dealers and end-users in the retail tier (see, e.g., Evans and Lyons 2002). Even in this comparatively simple setting, it is not straightforward to characterize the optimal trading strategies of individual dealers and end-users in the equilibrium of the model. This technical challenge has hampered the development of trading models that incorporate the richer institutional details of the modern FX market. For example, there are no models that have dealers competing with electronic platforms for end-users' order flows, nor are there models that allow dealers to manage risk through the use of both market and limit orders. In short, there appears to be something of a disconnect between the complex institutional structure of modern FX trading, and the highly simplified environment found in even the newest microstructure trading models.

This disconnect is less important when we consider the economic mechanisms at work in the market. Despite their institutional simplicity, existing microstructure models describe a complex process through which information is first conveyed to dealers by end-user order flows, and then used by dealers in their trades to manage risk. This process aggregates the information in end-user flows across the market into a form that dealers find relevant in the determination of the prices they quote for further trades. In short, the models describe the process by which the optimal trading decisions of end-users and dealers lead to the incorporation of price-relevant information

8

into exchange rates.

The institutional developments described above have changed the details of this process but not its essential characteristics. End-user flows still convey information to dealers which they use to manage the risks of supplying liquidity. The main difference is that end-user flows now arrive at banks electronically via multiple platforms as well as by the traditional voice channel. In addition, dealer-banks now have many more venues and ways in which to execute their risk-management trades. Nevertheless, as before, in aggregate these dealer trades convey information across the market which is eventually incorporated into exchange rates. Undoubtedly, the fragmentation of trading in the modern FX market makes this aggregation process more complex. However, recent evidence in Evans (2018), covering the EBS platform until 2015, confirms that order flow continues to account for a very high share, more than 90 percent, of FX price variation. Furthermore, by the end of daily trading, FX prices must still induce end-users to absorb dealer-banks' net intraday positions, just as in the 1980s.

# 3   Linking microstructure and macro exchange-rate models

Models of FX microstructure focus on how the information driving the decisions of market participants becomes embedded in exchange rates via the trading process. In contrast, macro models link exchange-rate dynamics to changing expectations concerning the future paths of interest rates and other variables, collectively referred to as "exchange-rate fundamentals". Here we show how the information driving the trading decisions in microstructure models is linked to the changing expectations about fundamentals. This discussion makes use of the framework found in Evans and Rime (2016).

We start with the definition of the expected log excess return on holding FX between the periods $t$ and $t + 1$:

$$\delta_t = \mathbb{E}_t^{\mathrm{M}} s_{t+1} - s_t + \hat{r}_t - r_t, \tag{1}$$

where $s_t$ is the log exchange rate, measured as the dollar price of FX. We refer to $\delta_t$ as the foreign exchange risk premium.[3] Here $r_t$ and $\hat{r}_t$ are the logs of the U.S. and foreign one-period nominal interest rates and $\mathbb{E}_t^{\mathrm{M}}$ denotes expectations conditioned on common information known to market participants at the start of period $t$, $\Omega_t^{\mathrm{M}}$. These participants comprise dealer-banks and the end-users that supply liquidity via limit orders on electronic trading platforms. Their expectations

---

[3]We use this terminology purely for pedagogical convenience. In the analysis that follows, we make no assumptions about why $\delta_t$ differs from zero at any point in time. It could represent equilibrium compensation for risk, or some kind of market friction.

are important because the quoting of prices by dealer-banks and the submission of limit orders literally establish the prices at which FX transactions take place, which in turn determines the spot exchange rate.[4]

Next, we rewrite (1) as a difference equation in $s_t$ and solve forward $H$ periods. Applying the Law of Iterated Expectations to the resulting expression produces

$$s_t = \mathbb{E}_t^{\text{M}} \sum_{i=0}^{H-1} (\hat{r}_{t+i} - r_{t+i}) - \mathbb{E}_t^{\text{M}} \sum_{i=0}^{H-1} \delta_{t+i} + \bar{s}_t, \tag{2}$$

where $\bar{s}_t \equiv \mathbb{E}_t^{\text{M}} s_{t+H}$ is the expected long-run exchange rate. We consider the implications of (2) for the $h$-period depreciation rate, $\Delta^h s_{t+h} = s_{t+h} - s_t$, where $h < H$. By definition, this rate equals the sum of expected depreciation rate, $\mathbb{E}_t^{\text{M}} \Delta^h s_{t+h}$, and the $h-$period-ahead forecast error $s_{t+h} - \mathbb{E}_t^{\text{M}} s_{t+h}$; components that can be directly computed from (2) as

$$\mathbb{E}_t^{\text{M}} \Delta^h s_{t+h} = \mathbb{E}_t^{\text{M}} \sum_{i=0}^{h-1} (r_{t+i} - \hat{r}_{t+i}) + \mathbb{E}_t^{\text{M}} \sum_{i=0}^{h-1} \delta_{t+i} \qquad \text{and} \tag{3a}$$

$$s_{t+h} - \mathbb{E}_t^{\text{M}} s_{t+h} = - \left( \mathbb{E}_{t+h}^{\text{M}} - \mathbb{E}_t^{\text{M}} \right) \sum_{i=h}^{H+h-1} (r_{t+i} - \hat{r}_{t+i}) - \left( \mathbb{E}_{t+h}^{\text{M}} - \mathbb{E}_t^{\text{M}} \right) \sum_{i=h}^{H+h-1} \delta_{t+i} + \bar{s}_{t+h} - \mathbb{E}_t^{\text{M}} \bar{s}_{t+h}. \tag{3b}$$

Substituting these expressions into the identity $\Delta^h s_{t+h} = \mathbb{E}_t^m \Delta^h s_{t+h} + s_{t+h} - \mathbb{E}_t^m s_{t+h}$ produces

$$\Delta^h s_{t+h} = \mathbb{E}_t^{\text{M}} \sum_{i=0}^{h-1} (r_{t+i} - \hat{r}_{t+i}) + \mathbb{E}_t^{\text{M}} \sum_{i=0}^{h-1} \delta_{t+i}$$
$$- \left( \mathbb{E}_{t+h}^{\text{M}} - \mathbb{E}_t^{\text{M}} \right) \sum_{i=h}^{H+h-1} (r_{t+i} - \hat{r}_{t+i}) - \left( \mathbb{E}_{t+h}^{\text{M}} - \mathbb{E}_t^{\text{M}} \right) \sum_{i=h}^{H+h-1} \delta_{t+i}$$
$$+ \bar{s}_{t+h}, - \mathbb{E}_t^{\text{M}} \bar{s}_{t+h}. \tag{4}$$

Equation (4) identifies all the proximate factors that can drive the $h$-period depreciation rate. Importantly, this expression follows simply from the Law of Iterated Expectations and the definition of the risk premium in (1). It contains no assumptions about FX trading, the behavior of interest rates or the expected long-run exchange rate. Notice, also, that the equation holds equally for

---

[4]In reality all FX transactions take place at either a bid or ask price, but we can safely ignore this complication here. For concreteness, we can simply think of the exchange rate as the mid-point of the bid and ask prices for FX trades at a point in time.

time periods of any duration (i.e., from seconds to quarters). Consequently, (4) provides a general framework to explore the links between microstructure and macro models of exchange-rate behavior across all frequencies.

Equation (4) shows that the depreciation rate can be decomposed into two sets of components. The first set comprises market participants' current period$-t$ expectations concerning the immediate path of interest rates and the risk premium. The second set comprises the revision in market participants' expectations, between periods $t$ and $t + h$, concerning the path for interest rates and the risk premium further into the future, and the long-horizon level of the exchange rate. Importantly, a variable can *only* drive the depreciation rate if they contribute to participants' period$-t$ expectations, or to the revision in their expectations between $t$ and $t + h$.

It is well-established that interest differentials and other macro variables have very little forecasting power for short- and medium-term depreciation rates. Microstructure models provide a perspective on this lack of forecastability by focusing on the first two terms on the right-hand side of the equation: $\mathbb{E}_t^{\mathrm{M}} \sum_{i=0}^{h-1} (r_{t+i} - \hat{r}_{t+i})$ and $\mathbb{E}_t^{\mathrm{M}} \sum_{i=0}^{h-1} \delta_{t+i}$. Since both terms involve the expectations of market participants, a macro variable can only have forecasting power insofar as is affects these expectations. This rules out variables that are not contemporaneously known. For example, since there are reporting lags for many macroeconomic variables, such as GDP, there is no way for the information contained in these variables to be directly and contemporaneously incorporated into market participants' expectations. In contrast, variables that are contemporaneously known to market participants, like order flows, are candidate forecasting variables.

Two strands of empirical research explore this implication. The first considers the forecasting power of order flows for future depreciation rates. Papers by Bjønnes et al. (2005), Berger et al. (2008), Evans and Lyons (2005), King et al. (2010), Chinn and Moore (2011), Evans and Lyons (2013), Evans and Rime (2016) and others all show that the flow of FX orders received by dealer-banks from end-users have significant forecasting power over horizons of days and weeks. Equation (4) tells us that this forecasting power must reflect a correlation between these flows and expectations concerning the future path of interest rates and the risk premia. This implication is confirmed by the second strand of the recent literature. For example, Evans and Lyons (2013) find that order flows have forecasting power for future inflation, GDP and monetary growth in the U.S. and Germany. These variables should be correlated with the path of interest rates that are relevant determinants of the EURUSD depreciation rate. In contrast, Evans and Rime (2016) show that the flows in the EURNOK market are correlated primarily with expectations concerning the risk premia.

Empirically, most of the variation in depreciation rates appear unforecastable. According to equation (4), these variations originate from revisions in participants' expectations concerning future interest rates, risk premia, and the long-run exchange rate. In principle, these revisions can be induced by public news or by the flow of new information reaching participants as they trade.

Macro data release are one source of such news. Many papers have shown that exchange rates react to the data releases in the direction predicted by standard models. However, the movements in exchange rates immediately following releases account for a very small fraction of the variance in (daily) depreciation rates. Microstructure models provide a richer perspective on the exchange-rate effects of data releases. They show that releases can have indirect effects because it takes time for market participants to reach a consensus about the implications of the release. This consensus-building process takes place via trading between participants following the release and can be identified from the order flows. Evans and Lyons (2008) estimate that more than one-third of the total variance in daily spot-rate changes can be related to the direct and indirect effects of macro data releases and other news sources, with indirect effects contributing roughly 60 percent more than the direct effects. These estimates, and similar results in Love and Payne (2008), clearly indicate that the indirect effects of news operating via order flow are an important component of exchange-rate dynamics.

The flow of orders received by dealer-banks from end-users, flows that are unrelated to the data releases, also appear to be an important source of information to market participants. Models developed in Evans (2010) and Evans and Lyons (2013) show how non-public information about the current (unreported) state of the economy is carried in dispersed form by the end-user flows reaching individual dealer-banks, and is then aggregated via inter-dealer trading into a form that drives market-wide expectations.[5] In support of this idea, these papers show that end-user flows carry incremental (non-public) information about current and future macro variables. In a similar vein, Rime et al. (2010) show that order flows have short-term forecasting power for specific macro data releases, while Evans and Rime (2016) find that end-user flows in the EURNOK market produce revisions in expectations about future risk premia and the long-horizon exchange rate. Together, these findings provide a structural interpretation of the strong contemporaneous correlation between deprecation rates and order flows reported by Evans and Lyons (2002) and many others.

In macro models, the risk premium embedded in the exchange rate is determined by the covariance between currency returns and the stochastic discount factor (pricing kernel) that determines

---

[5]Albuquerque et al. (2008) show that order flow from equity markets may also carry relevant information for exchange rates.

the prices of all financial assets. Microstructure models have a different perspective on the determinants of the risk premium. In the equilibrium of a trading model (see, e.g., Evans and Lyons 2002), dealers quote FX prices to ensure that risks are shared efficiently. This means that the price quoted at the end of daily trading must embed a risk premium that provides end-users with the incentive to absorb any net imbalance in dealer positions. Intraday FX prices embed risk premia that are similarly determined by risk-sharing. One implication of this microstructure perspective is that order flows should be correlated with risk premia. This idea has been explored empirically in a recent paper by Breedon et al. (2016).

# 4    Recent Research

Research in FX microstructure brings added perspective to two recent areas of interest in exchange-rate economics. The first area focuses on the behavior of returns on portfolios of foreign currencies. The second considers questions of competition and regulation of FX trading.

## 4.1    Portfolio Returns and Liquidity

Interest in the behavior of returns on currency portfolios largely originates from work by Lustig and Verdelhan (2007). Historically, research on exchange-rate behavior had focused on the behavior of individual foreign currency returns. Drawing on the larger finance literature, Lustig and Verdelhan (2007) argued that it would be easier to identify the systematic factors driving currency returns by examining the behavior of portfolio returns than the returns on individual currencies. This idea has proved appealing to numerous researchers, and has resulted in a growing number of papers; including Burnside et al. (2011), Òscar Jordàa and Taylor (2012), Hassan (2013), Menkhoff et al. (2012a,b) and Colacito and Croce (2011), among others.

As a purely statistically matter, the rationale for studying returns on currency portfolios rather than individual currency returns is straightforward. By combining individual currency returns into portfolios we effectively apply a statistical filter that reduces the importance of so-called idiosyncratic factors affecting individual currency returns. That said, the economic rationale for the focus on portfolios is more complex. Indeed, perspectives from FX microstructure can provide guidance to avoid potential pitfalls.

The problem concerns the economic relevance of the computed returns. Available data makes it possible to construct returns on portfolios with many different currencies going back several decades. However, it is unclear whether these returns on many of the individual currencies are

representative of the returns an international investor would have received in real time based on actual trades. Although the spot FX market is often cited as the world's largest financial market, one characterized by huge trading volumes and very high liquidity, this characterization overlooks the substantial differences in trading volume and liquidity across different currency pairs. Investors wanting to exchange some currency pairs face significant trading costs, particularly for large trades. Drawing reliable inferences from the behavior of individual returns computed without regard to these costs is difficult, and is even a greater challenge when they contribute to portfolio returns.

To gain some appreciation of the scale of the problem, we begin by comparing the size of the bid-ask spreads for spot trades across a large number of currency pairs at a point in time. In reality, the spread between the bid and ask prices at which an investor can trade vary according to the trading venue, the size of the trade, and (in some cases) the identity of the investor. For the moment, we ignore these complications. Instead, we focus on the spreads between the best bid and ask limit prices on the EBS and Reuters platforms across 61 currency pairs.[6] These data represent the smallest spreads available to typical sophisticated investors who trade small amounts via Prime Brokerage accounts. Figure 5 plots the median spreads (measured in basis points) from April 2013. The spreads range from as little as one quarter of a basis point for pegged currency pairs like the USDHKD, to as much as sixteen basis points in the case of the USDHUF. Clearly, these are substantial differences. Moreover, small differences in spreads can be economically important to an investor. For instance, a 1 basis point increase in bid-ask spread reduces the monthly return by 0.12 percent measured at an annual rate $((1 + 0.01/100)^{12} - 1)$.

The differences in spreads depicted in Figure 5 also shows up in trading volumes. Table 2 shows the share of spot trading volume for the top 29 currency pairs in April 2013 together with the median spreads on the EBS and Reuters platforms. The share data come from the BIS 2013 Triennial Survey. As the table shows, spot trading is heavily concentrated in a few currencies. This is a persistent pattern in the data. All the earlier BIS surveys showed that trades between a few major currency pairs accounted for the lion's share of spot trading volume. Table 2 also shows that there is no simple monotonic relationship between the share of trading volume and the median spreads. Spreads generally tend to be higher for currency pairs with smaller shares, but some pairs with similar shares have very different spreads on these platforms (see, e.g., USDTRY and USDSEK).

---

[6]When a currency pair trades on both the EBS and Reuters ECNs, we plot the smaller of the two median spreads.

Figure 5: Cross-Currency Spread Distribution (basis point) April 2013
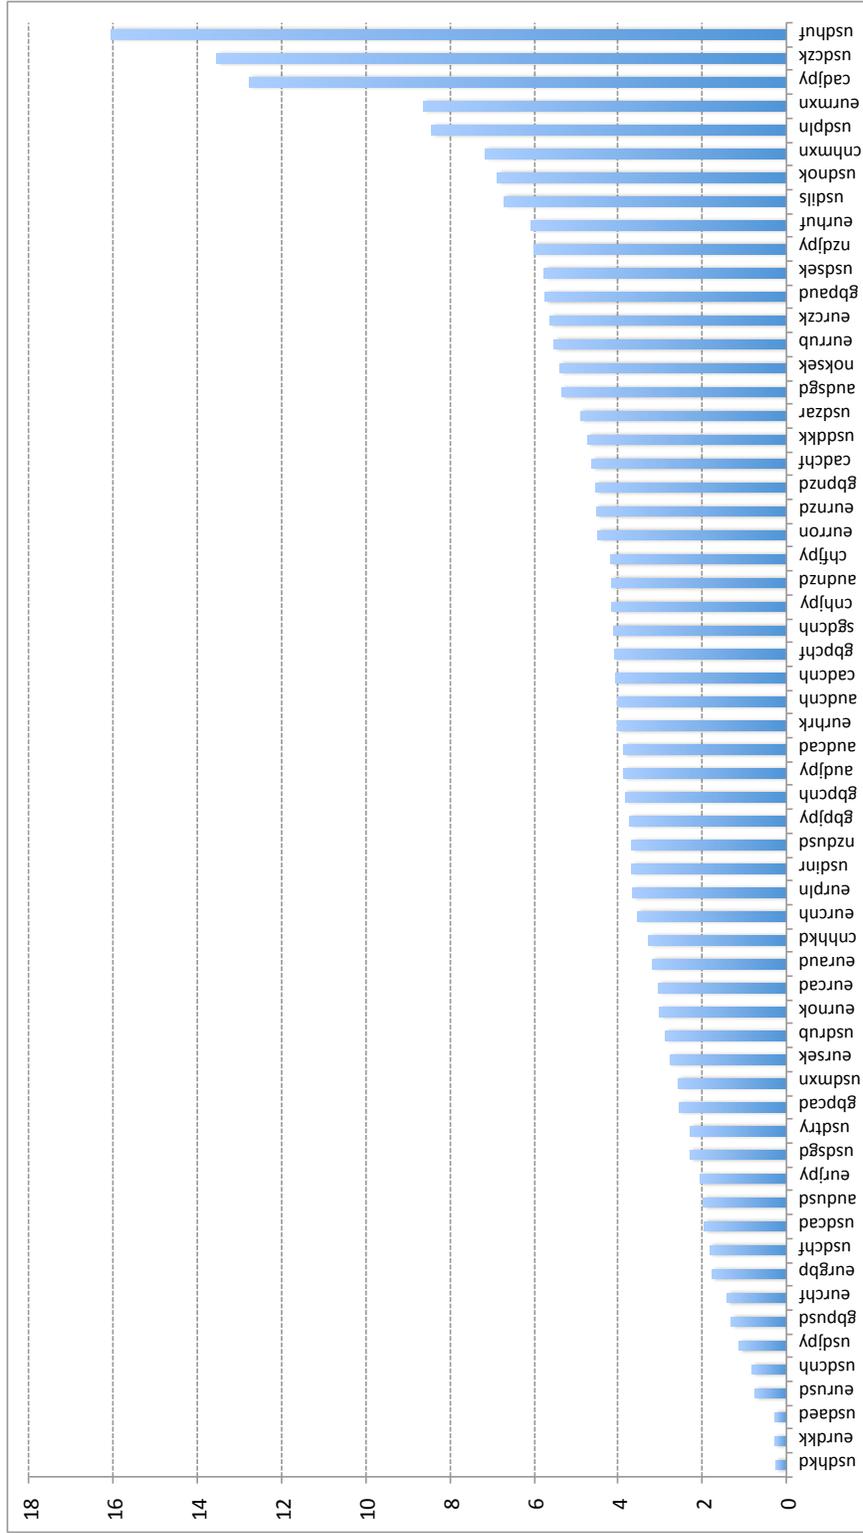


Notes: The plot shows median relative inside spread, in basis points, on the EBS/Reuters platforms during April of 2013

Table 2: Spreads and Trading Volume

| Rank | Currency Pair | Share | Spread | | Rank | Currency Pair | Share | Spread |
|------|---------------|-------|--------|---|------|---------------|-------|--------|
| 1 | EURUSD | 24.12 | 0.754 | | 16 | USDTRY | 1.17 | 2.296 |
| 2 | USDJPY | 18.30 | 1.132 | | 17 | USDSEK | 1.03 | 5.758 |
| 3 | GBPUSD | 8.82 | 1.308 | | 18 | USDZAR | 0.96 | 4.894 |
| 4 | AUDUSD | 6.80 | 1.979 | | 19 | USDINR | 0.93 | 3.675 |
| 5 | USDCAD | 3.73 | 1.965 | | 20 | USDNOK | 0.91 | 6.875 |
| 6 | USDCHF | 3.44 | 1.813 | | 21 | EURSEK | 0.52 | 2.773 |
| 7 | EURJPY | 2.77 | 2.051 | | 22 | USDPLN | 0.41 | 8.445 |
| 8 | USDMXN | 2.39 | 2.574 | | 23 | EURAUD | 0.39 | 3.189 |
| 9 | USDCNH | 2.10 | 0.822 | | 24 | EURNOK | 0.37 | 3.027 |
| 10 | EURGBP | 1.90 | 1.772 | | 25 | EURCAD | 0.28 | 3.046 |
| 11 | NZDUSD | 1.54 | 3.680 | | 26 | EURPLN | 0.26 | 3.664 |
| 12 | USDRUB | 1.48 | 2.879 | | 27 | EURDKK | 0.23 | 0.269 |
| 13 | EURCHF | 1.33 | 1.405 | | 28 | EURHUF | 0.18 | 6.063 |
| 14 | USDHKD | 1.28 | 0.257 | | 29 | EURCNH | 0.02 | 3.535 |
| 15 | USDSGD | 1.22 | 2.295 | | | | | |

Notes: The table reports the percentage share of trading volume for each currency pair reported in the BIS 2013 survey, and the median spread in basis points between the best limit bid and offer prices on the EBS/Reuters platforms during April 2013.

Figure 5 and Table 2 provide information on the degree of cross-currency heterogeneity in spreads and trading volumes at a particular point in time. This is only part of the picture. Figure 6 shows that there is also a significant degree of time series variation in bid-ask spreads. This figure plots spreads from three different sources for six representative currency pairs. The WM Company and Reuters have been a primary source for the exchange rate data found in the databases used by many researchers. The bid and ask prices they report come from trades in a one-minute window around 4:00 pm London (local) time, that are used to compute benchmark exchange rates, known as the WMR Fix. The second source is FXall, a Multi-bank platform that is accessible to financial and non-financial corporations as well as individuals. The third source is the Reuters D2000-2 platform. Access to this system was restricted to dealer-banks until the introduction of Prime Brokerage accounts in 2004.

Figure 6: Spread Comparisons
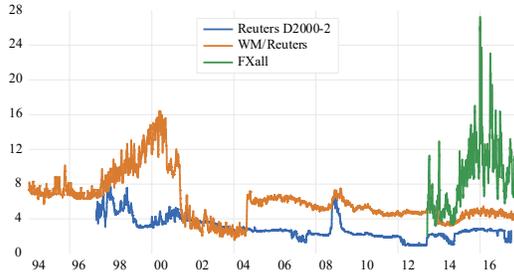
A: USDAUD
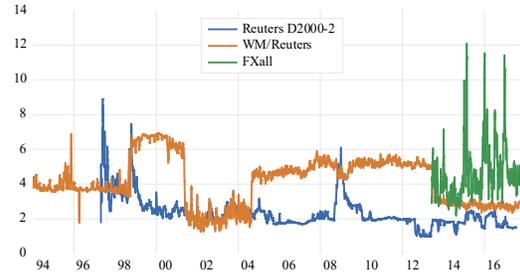
B: USDCAD
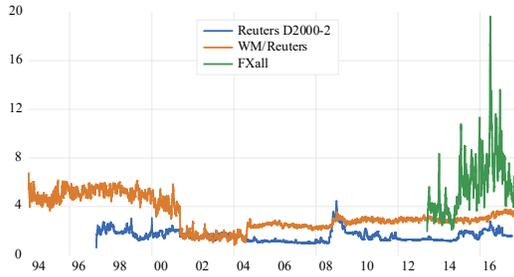
C: GBPUSD
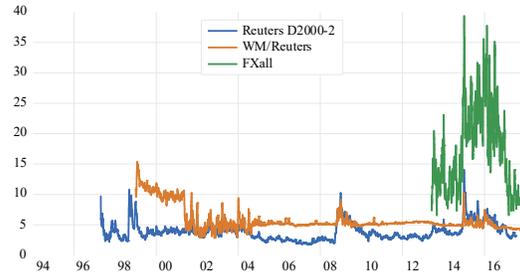
D: EURNOK
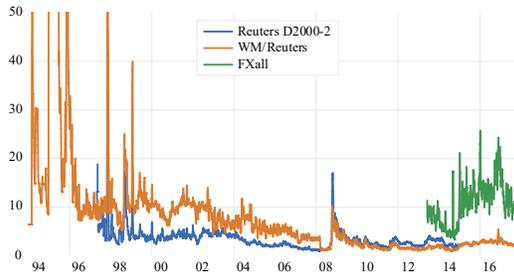
E: USDMXN
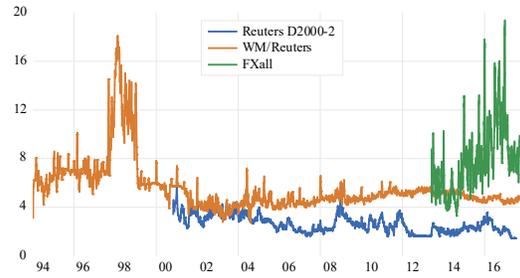
F: USDSGD

The plots in Figure 6 display two noteworthy features. First, there is considerable time-series volatility in the spreads from each individual source. This is particularly noticeable in the case of the FXall spreads, but it is also substantial in the other spreads for several currency pairs. Second, spreads can differ significantly from each other depending on their source. This is also clear from Table 3 which reports statistics for spreads sourced from the Fix and the Reuters D2000-2 platform. In particular, the volatility of spreads on Reuters D2000-2 are significantly larger than those from WMR. These differences are not surprising. Microstructure models emphasize the fact

Table 3: Descriptive statistics for spreads from Reuters Matching and WMR-fix. 2005-2014

|        |     | Std. Dev. | Max   | Min  | Skew  |
|--------|-----|-----------|-------|------|-------|
| AUD    | D2  | 0.40      | 3.74  | 0.43 | 2.06  |
|        | WMR | 0.19      | 1.68  | 0.60 | -0.07 |
| CAD    | D2  | 0.33      | 3.31  | 0.45 | 2.11  |
|        | WMR | 0.19      | 1.55  | 0.55 | -0.85 |
| EUR    | D2  | 1.92      | 49.48 | 0.25 | 9.97  |
|        | WMR | 0.17      | 1.41  | 0.59 | 0.35  |
| EURNOK | D2  | 0.42      | 6.35  | 0.37 | 3.36  |
|        | WMR | 0.15      | 2.46  | 0.80 | 5.30  |
| GBP    | D2  | 0.36      | 3.82  | 0.68 | 2.64  |
|        | WMR | 0.14      | 1.80  | 0.70 | 0.15  |
| MXN    | D2  | 0.61      | 11.06 | 0.35 | 6.52  |
|        | WMR | 0.71      | 5.43  | 0.25 | 1.83  |
| NZD    | D2  | 0.52      | 4.11  | 0.54 | 2.53  |
|        | WMR | 0.25      | 3.14  | 0.71 | 3.16  |
| SGD    | D2  | 0.27      | 2.95  | 0.55 | 1.17  |
|        | WMR | 0.18      | 2.65  | 0.72 | 2.47  |

Note: Table shows descriptive statistics for standardized relative bid-ask spreads (divided by mean) for spreads calculated for data from Reuters Matching (D2) and from the WMR 4pm fix. Source: Thomson Reuters TickHistory and Datastream.

that spreads reflect the trading conditions in different venues. However, the existence of these differences complicates the task of accurately accounting for trading costs when computing returns. While Lyons (2001) pointed out that WMR bid-ask spreads were generally larger than those in the interbank market, they are smaller than the spreads faced by end-users with typically-sized trades. Lustig et al. (2011) show that the return on a carry trade investment strategy with a high monthly turnover is reduced by half when applying WMR bid-ask spreads. For typical trade sizes, the effect would be even larger. Of course, for trading strategies with low turnover, the problem is smaller. At the very least, it is unclear whether spreads sourced from WM/Reuters, which are often used to indicate the size of transaction costs, adequately represent the costs investors face.

Figure 7: EBS Limit Prices, April 2007
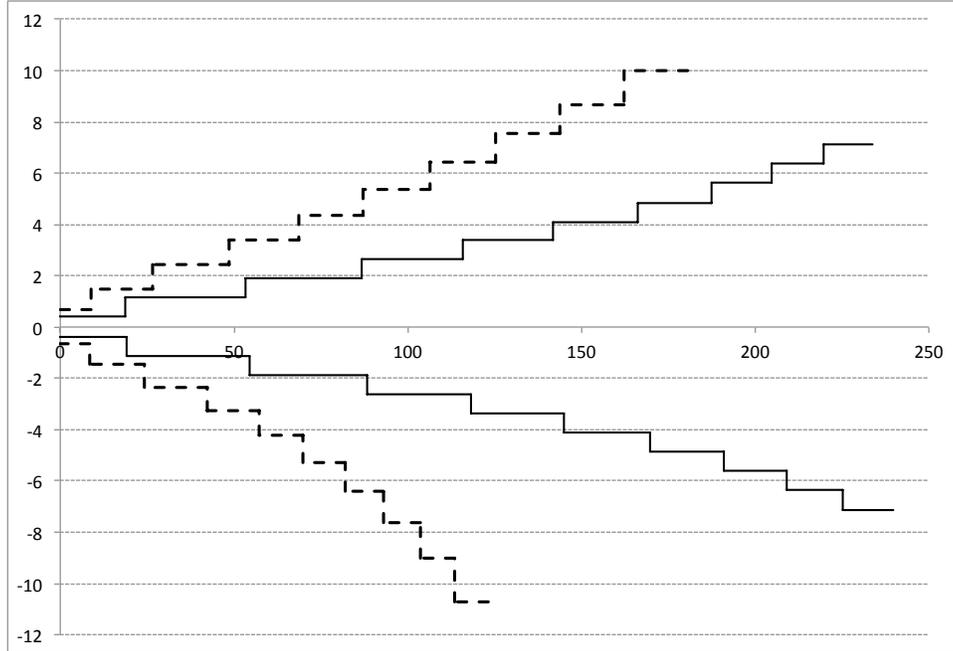


Notes: Average ask and bid limit prices in basis points (relative to mid-point be-
tween the best bid and ask prices) against the average volume of limit orders, in
millions of USD for EURUSD (solid) and USDJPY (dashed).

Spreads computed from the best bid and ask prices on the EBS and Reuters platforms provide
a limited picture of the true trading costs investors face because they give no indication of the size
of a trade that can be executed at those prices. Information on how trading costs vary according to
the size of a trade can be found from the structure of limit orders on the EBS platform. Figure 7
plots the average ask and bid limit orders against the average cumulative volume for the orders in
April 2007 for the EURUSD and USDJPY.[7] The limit prices are plotted in basis points relative to
the mid-point between the best bid and ask prices, and volumes are in millions of USD. The plots
show that spreads between the best bid and ask prices (sometimes referred to as the inside spread)
only represent the trading costs for trades of less than 20 million in the case of the EURUSD,
and 10 million in the case of USDJPY. Larger trades would incur disproportionately higher costs
because they would have to be matched, in part, by limit orders with prices that are further from
the mid-point. For example, most of a market order to purchase EUR worth USD 200 million would

---

[7]Averages are computed from limit prices and volumes sampled each minute between 7:00 am and 5:00 pm London
time on each trading day.

be matched against limit orders with prices ranging up to five basis points above the mid-point. Thus the slopes of the limit order plots provide a visual guide to how transaction costs rise with the size of trades. In addition, the width of the plots indicates the maximum size of trades that could be executed.

The plots in Figure 7 clearly illustrate the limitations of the inside spread as a measure of trading costs. In particular, the figure shows that while the inside spreads are quite similar for the two currency pairs, the actual costs of trading large amounts would be very different. For example, the spreads in EURUSD and USDJPY markets at USD 100 million are approximately 5 and 15 basis points, respectively. Furthermore, the depth of limit orders in the USDJPY market is much less than in the EURUSD. The lower plots show that the bid limit orders can only absorb JPY sales worth USD 120 million, whereas the limits can absorb EUR sales worth up to USD 230 million. In sum, Figure 7 shows how small differences in the inside spread between currency pairs can mask significant differences in the costs of executing large trades.

In several respects, the plots in Figure 7 understate the degree to which costs increase with the trade size across most currency pairs and during most time periods. First, the EURUSD and USDJPY are by far the most heavily traded currency pairs (see, e.g., Table 5) and their limit prices have tighter spreads and greater depth on the EBS platform than on other trading venues. Second, not only are the inside spreads on other currency pairs larger than those plotted here, but the slopes of the limit order books are higher and their depth is smaller. Thus the plots in Figure 7 understate the degree to which costs rise with trade size for most currency pairs. Third, the EBS and Reuters platforms were still the single most important trading venues in 2007, accounting for approximately 30 percent of spot trading volume. Since then, their share of volume has declined to approximately 10 percent by 2015 while the share of trading on Single- and Multi-bank platforms has risen to almost 50 percent (see Figure 3).

This shift in the relative importance of different trading venues has been accompanied by a change in trading costs. For example, Figure 8 plots limit prices for the EURUSD and USDJPY on the EBS platform in April of 2015. These plots show that by 2015 the slopes of the limit order books had increased for both currency pairs, so trading costs increased more with the size of the trade.[8] The plots also show that there was a marked reduction in the depth of the limit order book for both currency pairs. It is important to note that other trading venues, such as Single-bank platforms, do not provide investors with alternative sources of liquidity that offset the reduction

---

[8] In the case of the EURUSD, the inside spread widened from 0.8 basis points to 1.09 basis points but the other spreads widened far more. For example, at USD 50 million the spread increased from 2.28 basis points in 2007 to 7.37 basis points in 2015.

in liquidity on the EBS and Reuters systems. Because the same dealer-banks are major liquidity suppliers across multiple platforms, when demand for liquidity is unusually high in one venue, they reduce their supply of liquidity to other venues. Such behavior creates a so-called liquidity mirage.

Figure 8: EBS Limit Prices, April 2015
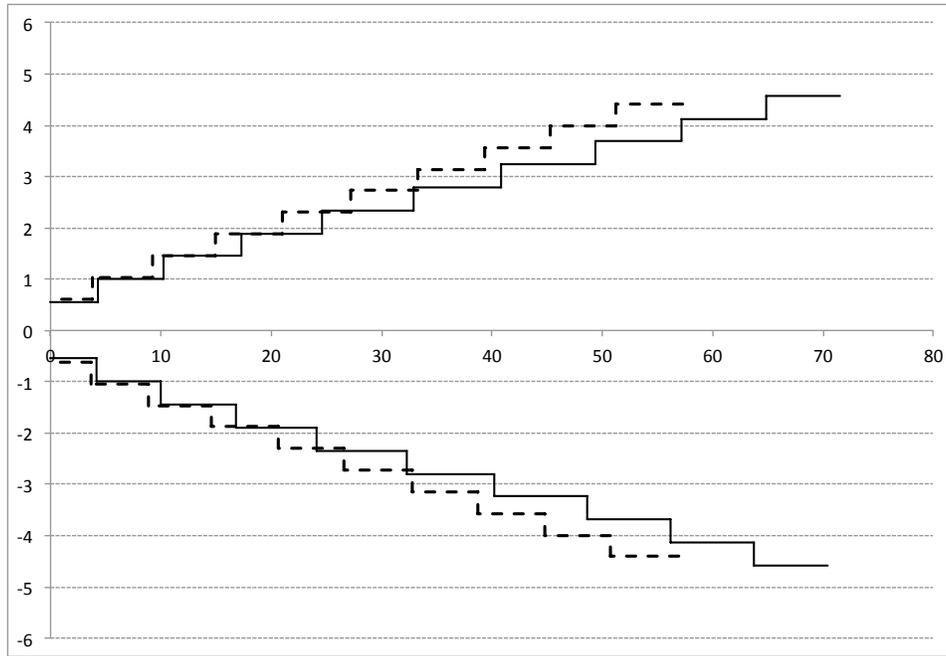


Notes: Average ask and bid limit prices in basis points (relative to mid-point between the best bid and ask prices) against the average volume of limit orders, in millions of USD for EURUSD (solid) and USDJPY (dashed).

We can gain further insight into the size of trading costs implied by the structure of limit orders depicted in Figure 8 by computing the actual costs of trading against the EURUSD and USDJPY limit orders that existed on the EBS platform during April 2015. For this purpose, we computed the Value Weighted Average Price (VWAP) of market orders to purchase and sell FX valued between 10 and 100 million USD using the set of limit orders that were present on EBS at the start of every minute between 7:00 am and 5:00 pm London time for each trading day in April 2015. In these calculations, we prioritized the matching of the market order with the best limit orders in the same way as EBS, and then compare the VWAP for the market order against the mid-point between the best bid and ask limit prices. Summary statistics from this calculations are reported in Table 4.

21

Table 4: Execution Costs

| Trade Size | A: Market Purchase | | | | B : Market Sales | | | |
| (million USD) | Execution Percentage | VWAP Percentiles 10 | 50 | 90 | Execution Percentage | VWAP Percentiles 10 | 50 | 90 |
|---|---|---|---|---|---|---|---|---|
| **EURUSD** | | | | | | | | |
| 10 | 100.00 | 0.613 | 0.912 | 1.227 | 100.00 | -1.226 | -0.917 | -0.618 |
| 20 | 99.96 | 0.905 | 1.270 | 1.658 | 99.93 | -1.657 | -1.274 | -0.917 |
| 30 | 99.71 | 1.163 | 1.603 | 2.048 | 99.63 | -2.046 | -1.607 | -1.181 |
| 40 | 98.50 | 1.401 | 1.915 | 2.396 | 98.50 | -2.399 | -1.917 | -1.429 |
| 50 | 91.31 | 1.603 | 2.182 | 2.650 | 91.03 | -2.650 | -2.184 | -1.640 |
| 60 | 66.10 | 1.689 | 2.332 | 2.766 | 65.77 | -2.763 | -2.338 | -1.750 |
| 70 | 37.20 | 1.659 | 2.368 | 2.840 | 37.31 | -2.844 | -2.389 | -1.754 |
| 80 | 21.51 | 1.613 | 2.363 | 2.935 | 20.33 | -2.921 | -2.393 | -1.703 |
| 90 | 13.66 | 1.604 | 2.359 | 3.038 | 12.45 | -3.021 | -2.406 | -1.651 |
| 100 | 9.26 | 1.587 | 2.387 | 3.129 | 8.26 | -3.101 | -2.439 | -1.640 |
| **JPYUSD** | | | | | | | | |
| 10 | 100.00 | 0.712 | 0.979 | 1.248 | 100.00 | -1.266 | -0.991 | -0.716 |
| 20 | 99.90 | 1.044 | 1.370 | 1.695 | 99.92 | -1.711 | -1.383 | -1.065 |
| 30 | 99.31 | 1.365 | 1.757 | 2.121 | 99.34 | -2.141 | -1.773 | -1.392 |
| 40 | 93.04 | 1.658 | 2.124 | 2.497 | 93.18 | -2.515 | -2.143 | -1.707 |
| 50 | 54.48 | 1.774 | 2.321 | 2.685 | 52.13 | -2.693 | -2.343 | -1.845 |
| 60 | 22.25 | 1.618 | 2.355 | 2.819 | 18.65 | -2.828 | -2.391 | -1.685 |
| 70 | 12.47 | 1.529 | 2.371 | 2.962 | 9.64 | -2.989 | -2.400 | -1.555 |
| 80 | 8.57 | 1.456 | 2.406 | 3.068 | 6.61 | -3.109 | -2.444 | -1.516 |
| 90 | 6.30 | 1.445 | 2.438 | 3.180 | 4.87 | -3.205 | -2.465 | -1.499 |
| 100 | 4.82 | 1.412 | 2.455 | 3.262 | 3.70 | -3.297 | -2.491 | -1.433 |

Notes: The table reports the execution percentage for market purchase and market sales orders with sizes ranging from 10 to USD 100 million based on existing limit orders on EBS during April 2013. The table also shows the 10th, 50th and 90th percentiles of the VWAP distribution for executed orders, where the VWAP is measured in basis points relative to the mid-point of the best limit bid and ask prices.

The table reports percentiles of the VWAP distribution for trades of different sizes, together with the probability that there was sufficient depth in the limit orders to allow the trade to fully execute. Two aspects of the statistics stand out. First, there is considerable variation in the VWAPs, particularly for larger-sized trades. Execution costs for the same trade size vary considerably with changes in market conditions. Second, the execution probabilities fall with trade size. There is sufficient depth in limit orders to allow market orders worth USD 40 million or less to almost

always execute. However, as trade size increases beyond USD 40-50 million, the probability of execution falls off sharply. Figure 4 showed that trades in major currency pairs worth more than USD 50 million are not uncommon, so this drop in execution probability is economically significant.

The VWAP statistics understate the true costs of trading large amounts when the execution probabilities are low. Trades too big to execute against existing limit orders must be split. This is likely to increase the cost of trading because splitting a large order into a sequence of smaller market orders can produce an adverse shift in limit orders before all the market orders are executed. For example, a sequence of market purchases can induce an upward shift in ask limit orders as liquidity suppliers react to the arrival of the first market orders in the sequence. As a consequence, the cost of executing a USD 100 million purchase will be greater than 10 times the cost of a USD 10 million purchase.

Several points emerge from this discussion of trading costs that are relevant to the interpretation of research studying the behavior of portfolio returns. First, the costs of executing significantly sized spot trades are vastly different across currencies and through time. Second, inside spreads provide a very incomplete picture of the true costs of trading meaningful amounts. Changes in the slope and depth of limit orders can substantially change the actual costs of trading without there being a sizable change in the inside spreads. Third, information on inside spreads that is available from commonly used databases is not representative of the spreads quoted to market participants in different trading venues. In particular, spreads computed from the 4:00 pm Fix window are poor approximations to the spreads applicable for the overwhelming majority of spot trades. Finally, there is no clear evidence that trading costs have decreased over time. Indeed, as the EBS and Reuters platforms have lost their dominance as trading venues, it appears that the costs associated with trading major currencies in large volumes have actually increased in recent years.

Taken together, these observations pose a challenge for interpreting the behavior of portfolio returns. It would be convenient if we could interpret their behavior as coming from a world in which trading costs were negligible because the returns we measure would correspond to the returns facing investors. Unfortunately, the available evidence suggests that this is a poor approximation. In reality, there is a wedge between the returns we measure and those that are available to investors who want to commit large sums to a multi-currency strategy. One of the major challenges facing researchers in FX microstructure is to characterize these wedges in a way that allows us to draw more accurate inferences from the behavior of measured portfolio returns.

## 4.2   Competition and Regulation

Until recently, researchers in FX Microstructure paid little attention to issues related to competition and regulation. Since FX trading is geographically dispersed around the world, it has not been subject to the same degree of regulatory oversight by national regulators as trading in other markets; particularly those that developed on exchanges, such as the equity and commodity markets. Despite this lack of regulatory oversight, the implicit assumption of most researchers was that FX trading took place in a highly competitive environment. Indeed, the workhorse microstructure models of FX trading assume that dealer-banks act competitively.

However, in recent years, government regulators and enforcement authorities have conducted investigations into whether many of the world's largest dealer-banks had been acting anti-competitively. These investigations have taken place across the globe – including in the United States, United Kingdom, European Union, Switzerland, Germany, Hong Kong, Singapore, Australia, and New Zealand. In addition, multiple law-suites have been brought before the courts in the United States and Canada alleging that dealer-banks engaged in anti-competitive behavior that harmed investors. By the end of 2018, these investigations and law-suites produced fines and settlements totaling over $11 billion. These developments provide a new impetus for FX microstructure researchers to consider competition and regulation issues.

Since trading in the FX market is dispersed across many trading platforms and involves many different types of participants, there is no simple way to measure the degree to which trading across the market is competitive. However, some information is available from the surveys conducted by Euromoney Magazine on the importance of the large dealer-banks. We present this information in Figure 9. In particular, the figure plots the Herfindal index (a standard measure of market concentration) for overall FX trading by top-30 dealer-banks, and for the trades of specific end-users: levered investors, real money, non-financial and non-dealer banks. We also plot the index for trading on Single-bank platforms, a trading venue that has emerged as the single largest venue for FX spot trades in recent years.

Figure 9: Market concentration (Herfindahl index) for overall market and segments
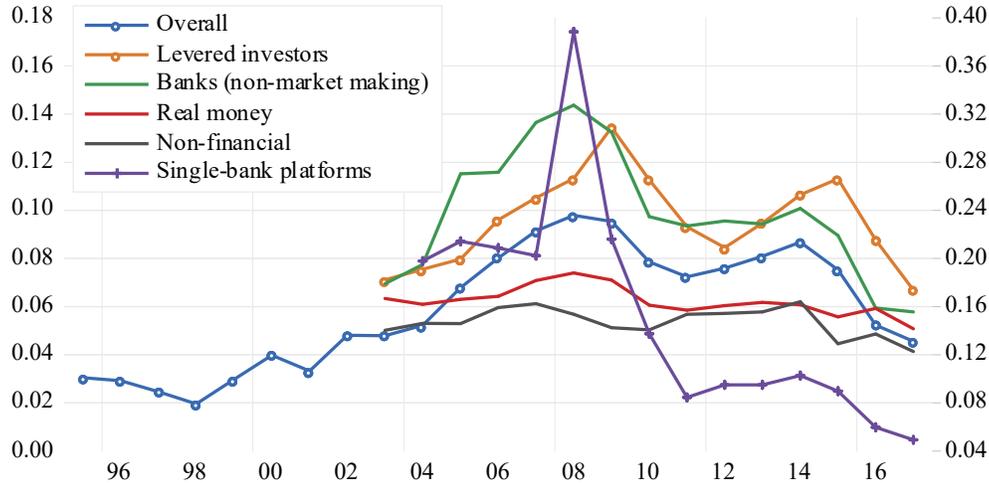


Note: Graph show market concentration, by Herfindahl index, for overall market, customer segments (lhs scale) and bank's Single-bank platforms (rhs scale). Source: Euromoney Magazine annual FX survey.

There are several noteworthy features of the plots. First, overall FX trading became a good deal more concentrated among the largest dealer-banks between 2000 and the 2008 financial crisis; but since the crisis, the concentration of overall trading has fallen. Second, this rise and then fall in the concentration of overall market trading is also reflecting in the trading of the levered-investors and non-dealer banks. These changing concentration patterns have yet to be investigated by researchers. The third noteworthy feature concerns the large changes in the market concentration of trading on Single-bank platforms. The plots indicate that trading was heavily concentrated on a few platforms before and during the 2008 financial crisis, but since then market concentration has fallen sharply. Trades on Single-bank platforms are less transparent to all but the dealer-banks running the platform than are the trades on the EBS and Reuters platforms, so it unclear whether the liquidity supplied by the dealer-banks operating these venues can be considered highly competitive during the years when trading on the platforms was heavily concentrated. Again, the implications of these changing concentration patterns on Single-bank platforms have yet to be be examined by researchers.

The investigations conducted by government regulators and enforcement authorities provide more concrete evidence against the view that trading in FX is always highly competitive. Initially, these investigations focused on the behavior of dealer-banks around the determination of FX

25

benchmarks, particularly the ECB Fix at 2:15 PM (CET) and the WMR Fix at 4:00 PM (GMT). More specifically, between 2014 and 2015 reports issued by the U.S. Department of Justice, the Commodity Futures Trading Commission, New York Department of Financial Services, the U.K. Financial Conduct Authority, and the Swiss Financial Market Supervisory Authority all concluded that dealer-banks engaged in a range of collusive conduct aimed at manipulating the benchmarks.[9] This conduct, which was facilitated by dealers' use of electronic chat rooms, included: (i) the sharing of information about positions and orders before a Fix, (ii) the coordination of trading at the time of a Fix in order to manipulate the determination of the benchmark, and (iii) the netting orders between dealers so as not to execute such orders at the time of a Fix. In response to these findings, The Financial Stability Board (an international body that monitors and makes recommendations about the global financial system) recommended reforms to the WMR Fix methodology, which led to a widening of the calculation window from one to five minutes in February 2015. In a recent study, Evans et al. (2018) find that this reform made the WMR Fix less susceptible to manipulation but also harder to replicate via actual trades within the window.

The investigations also concluded that the dealer-banks had engaged in other forms of collusive conduct that were unrelated to the determination of FX benchmarks. This conduct included: (i) the triggering ("jamming") of stop-loss orders where customer stop orders were actively triggered to the bank's advantage, (ii) the front running of customer orders by executing aligned orders in advance, (iii) the partial filling of customer orders where at least a part of the profitable transaction was credited to the bank, (iv) colluding on the prices to quote particular customers, and (v) agreeing to refrain from trading where one colluding bank had greater need to buy or sell than the others. In addition, the U.S Department of Justice and the Federal Reserve Board have indicted and placed lifetime bans on more than a dozen individual FX traders. Of particular note is the criminal conviction of the HSBC trader Mark Johnson in 2017 for committing wire fraud and wire fraud conspiracy as part of a scheme to front-run a $3.5 billion GBPUSD transaction in 2011.[10] Finally, class action law-suites filed in U.S. courts on behalf of investors alleging violations of the anti-trust

---

[9]Information on these investigations can be found at: DOJ Press Release, Five Major Banks Agree to Parent Level Guilty Pleas (May 20, 2015), https://www.justice.gov/opa/pr/five-major-banks-agree-parent-level-guilty-pleas. CFTC Press Release No. 7056-14, CFTC Orders Five Banks to Pay over $1.4 Billion in Penalties for Attempted Manipulation of Foreign Exchange Benchmark Rates (Nov. 12, 2014), http://www.cftc.gov/PressRoom/PressReleases/pr7056-14. DFS Press Release, NYDFS Announces Barclays to Pay $2.4 Billion, Terminate Employees for Conspiring to Manipulate Spot FX Trading Market (May 20, 2015), http://www.dfs.ny.gov/about/press/pr1505201.htm. FCA Press Release, FCA fines five banks 1.1 billion for FX failings and announces industry-wide remediation program (Nov. 12, 2014), http://www.fca.org.uk/news/fca-fines-five-banks-for-fx-failings. FINMA Press Release, FINMA sanctions foreign exchange manipulation at UBS (Nov. 12, 2014), http://www.finma.ch/e/aktuell/pages/mm-ubs-devisenhandel-20141112.aspx.

[10]See, U.S. v. Johnson, No. 16-cr-457-NCG-1 (E.D.N.Y.).

Sherman Act have produced settlements in excess of $2.3 billion.

Overall, these developments represent a direct challenge to the assumption, found in all microstructure models of FX trading, that dealer-banks always act unilaterally when making their trading decisions.[11] But, at the same time, they open up several new avenues of FX microstructure research. One avenue concerns detection. Can researchers devise statistical tests to detect anti-competitive behavior? What data is necessary to conduct such tests? A second avenue concerns the promotion of competition. End-users now have access to many different trading venues. How should they assess the relative benefits of using one venue rather than another in a way that supports a competitive trading environment across the entire market? The third avenue concerns regulation. National regulators do have authority over the actions of dealer-banks operating within their jurisdictions, so there is scope for multilateral regulation of the FX trading activity. Researchers in FX microstructure have unique expertise that should inform whether and how such multilateral regulation of the FX market evolves. The recent work on the reformed WMR Fix by Evans et al. (2018) represents a contribution in this direction, but in our view, there is much more that can be done.

# 5   Conclusion

In this article, we have presented a brief overview of research on the Microstructure of Foreign Exchange Markets. We described the many institutional changes in FX trading that have occurred since the 1980s but argued that the key economic features of trading have remained relatively unchanged and are still well represented by microstructure models in the literature. We have also argued that microstructure and macro exchange-rate models are complementary research tools and showed how these modeling approaches are linked. Finally, we provided a microstructure perspective on two recent areas of interest in exchange-rate economics: the behavior of returns on currency portfolios, and questions of competition and regulation.

# References

Albuquerque, Rui, Eva De Francisco, and Luis B. Marques. 2008. "Marketwide Private Information in Stocks: Forecasting Currency Returns." *Journal of Finance* 63 (5): 2297–2343.

---

[11]Recall that the daily spot rates found in popular databases are actually sourced from the 4:00 pm WMR Fix, so the implications of collusive trading around the Fix extend beyond the traditional domain of FX microstructure. If dealer-banks traded collusively around the major Fixes in a way that significantly affected the actual determination of the Fix benchmarks (before 2015), interpreting this data within a purely competitive paradigm becomes problematic.

Berger, D.W., A.P. Chaboud, S.V. Chernenko, E. Howorka, and J.H. Wright. 2008. "Order flow and exchange rate dynamics in electronic brokerage system data." *Journal of International Economics* 75 (1): 93–109.

Bjønnes, Geir H. and Neophytos Kathiziotis. 2016. "Price Discovery in Fragmented Electronic Markets — The Case of FX." Tech. rep., University of Hamburg.

Bjønnes, Geir H., Dagfinn Rime, and Haakon O. Aa. Solheim. 2005. "Liquidity Provision in the Overnight Foreign Exchange Market." *Journal of International Money & Finance* 24 (2): 177–198.

Breedon, Francis, Dagfinn Rime, and Paolo Vitale. 2016. "Carry Trades, Order Flow, and the Forward Bias Puzzle." *Journal of Money, Credit and Banking* 48 (6): 1113–1134.

Burnside, A. Craig, Martin S. Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo. 2011. "Do Peso Problems Explain the Returns to the Carry Trade?" *Rev. Financial Studies* 24 (3): 853–891.

Chinn, Menzie D. and Michael J. Moore. 2011. "Order Flow and the Monetary Model of Exchange Rates: Evidence from a Novel Data Set." *Journal of Money, Credit & Banking* 43 (8): 1599–1624.

Colacito, Riccardo. and Mariano M. Croce. 2011. "Risks for the Long Run and the Real Exchange Rate." *Journal of Political Economy* 119 (1): 153–181.

Evans, Martin D. D. 2010. "Order Flows and The Exchange Rate Disconnect Puzzle." *Journal of International Economics* 80 (1): 58–82.

———. 2011. "The Microstructure of Currency Markets." In *Encyclopedia of Financial Globalization*, edited by G. Caprio. Elsevier.

———. 2018. "FX Trading and Exchange Rate Disconnect Puzzle." Working Paper 18-18-21, Georgetown University.

Evans, Martin D. D. and Richard K. Lyons. 2002. "Order flow and exchange rate dynamics." *Journal of Political Economy* 110 (1): 170–180.

———. 2005. "Meese-Rogoff Redux: Micro-Based Exchange-Rate Forecasting." *American Economic Review. Papers & Proceedings* 95 (2): 405–414.

———. 2008. "How is macro news transmitted to exchange rates?" *Journal of Financial Economics* 88 (1): 26–50.

———. 2013. "Exchange Rate Fundamentals and Order Flow." *Quarterly Journal of Finance* 02 (04): 1250018.

Evans, Martin D. D., Peter O'Neill, Dagfinn Rime, and Jo Saakvitne. 2018. "Fixing the Fix? Assessing the Effectiveness of the 4pm Fix Benchmark." *FCA Occasional Paper* .

Evans, Martin D. D. and Dagfinn Rime. 2012. "Micro Approaches to Foreign Exchange Determination." In *Handbook of Exchange Rates*, edited by Jessica James, Ian Marsh, and Lucio Sarno. John Wiley & Sons.

———. 2016. "Order flow information and spot rate dynamics." *Journal of International Money and Finance* 69 (Supplement C): 45 – 68. SI - Tribute Jim Lothian.

Hassan, Tarek A. 2013. "Country Size, Currency Unions, and International Asset Returns." *The Journal of Finance* 68 (6): 2269–2308.

King, Michael, Carol Osler, and Dagfinn Rime. 2012. "Foreign exchange market structure, players and evolution." In *Handbook of Exchange Rates*, edited by Ian Marsh Jessica James and Lucio Sarno. John Wiley & Sons.

King, Michael, Lucio Sarno, and Elvira Sojli. 2010. "Timing exchange rates using order flow: The case of the Loonie." *Journal of Banking & Finance* 34 (12): 2917–2928.

Love, Ryan and Richard Payne. 2008. "Macroeconomic News, Order Flows and Exchange Rates." *Journal of Financial and Quantitative Analysis* 43 (02): 467–488.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan. 2011. "Common Risk Factors in Currency Markets." *Review of financial studies* 24 (11): 3731–3777.

Lustig, Hanno and Adrien Verdelhan. 2007. "The cross section of foreign currency risk premia and consumption growth risk." *American Economic Review* 97 (1): 89–117.

Lyons, Richard K. 1995. "Tests of Microstructural Hypothesis in the Foreign Exchange Market." *Journal of Financial Economics* 39 (2-3): 321–351.

———. 2001. *The Microstructure Approach to Exchange Rates*. Cambridge, MA: M I T Press.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf. 2012a. "Carry Trades and Global Foreign Exchange Volatility." *Journal of Finance* 67 (2): 681–718.

———. 2012b. "Currency Momentum Strategies." *Journal of Financial Economics* 106 (3): 660–684.

Òscar Jordàa and Alan M. Taylor. 2012. "The carry trade and fundamentals: Nothing to fear but FEER itself." *Journal of International Economics* 88 (1): 74–90.

Osler, Carol L. 2009. "Market Microstructure, Foreign Exchange." In *Encyclopedia of Complexity and System Science*, edited by Robert A. Meyers. New York: Springer, 5404–5438.

Rime, Dagfinn, Lucio Sarno, and Elvira Sojli. 2010. "Exchange Rate Forecasting, Order Flow and Macroeconomic Information." *Journal of International Economics* 80 (1): 72–88.