HFTs and Dealer Banks: Liquidity and Price Discovery in FX Trading

Abstract

By investigating dealer banks and high-frequency traders (HFTs) in foreign exchange markets, this study sheds light on the distinct yet complementary roles of "traditional" and "new" market makers in over-the-counter markets. Using message-level data, our findings reveal that these two types coexist by carrying out complementary roles. HFTs excel in processing *public* information, while dealers are skilled in managing *private* information. Specifically, HFTs provide resilient liquidity during market-wide volatility spikes, whereas dealer liquidity is robust in informational events such as scheduled macroeconomic announcements or policy regime changes. HFTs contribute to the majority of the information share through frequent *quote* updates, which incorporate public information. In contrast, dealers contribute to price discovery through *trades* that impound private information.

1 Introduction

Financial markets have undergone significant structural changes in the last two decades. As trading becomes more electronic, automated, and fast-paced, "new" market participants like high-frequency traders (HFTs)[1] have been replacing "traditional" ones like dealer banks. This trend is particularly prominent in all-to-all markets such as equity markets in which HFTs have accounted for a large share of liquidity provision (Menkveld, 2016). There is also growing evidence that HFTs have made inroads in over-the-counter (OTC) markets such as foreign exchange (FX) markets (Bank for International Settlements, 2011; Chaboud, Chiquoine, Hjalmarsson, and Vega, 2014).

A central problem of market making is how to manage adverse selection risk, that is, how to avoid losses from trading with informed traders. Glosten and Milgrom (1985) and Kyle (1985) show that market makers extract the asymmetrically-observable information embedded in the trades they engage in, which they use to update their prices. This information is referred to as *private* information in the literature, as it is revealed only *after* the trade occurs. Recent research (see e.g., Budish, Cramton, and Shim, 2015; Foucault, Hombert, and Roşu, 2016; Malceniece, Malcenieks, and Putniņš, 2019; Aquilina, Budish, and O'Neill, 2022) has emphasised the role of *public* information in modern market making. This is information that is revealed *ahead of* trading, via signals that are symmetrically observable by market participants, such as price changes in correlated assets or limit order book imbalances. Although there is fast-growing literature on how HFTs respond to public and private information in all-to-all markets (notably equity markets), related analysis on other types of markets is sporadic.

In this paper, we exploit the over-the-counter (OTC) structure of FX markets to better understand the roles of the new and traditional market makers. On the one hand, the dealer-to-dealer (D2D) segment is predominantly electronic and automated, in which HFTs can have speed and technology advantages. On the other hand, dealer banks still play a dominant role in the dealer-to-client (D2C) segment, providing agency execution services, single-dealer platforms, as well as other client relationship services (Bank for International Settlements, 2018).[2] Dealers could gain

[1]High-frequency traders are a subset of algorithmic traders who trade on their proprietary accounts (SEC, 2010). We detail the classification of the traders in Section 2.2.

[2]See Chaboud, Rime, and Sushko (2022) for an overview of the evolution of FX market structure. They note the

private information through explicit learning (e.g., via price discrimination, as argued in Collin-Dufresne, Hoffmann, and Vogel (2019)) or implicit learning (e.g., via internalised client trades or trade execution for informed clients).[3]

The contribution of this paper is twofold. First, we provide new evidence on the microstructure of the FX spot market. We find that in the D2D segment of FX spot trading, both HFTs and dealer banks are important and active market makers. This result is in contrast to equity markets, where traditional market makers play a more negligible role. We find HFTs are skilled in processing public information whereas dealers excel in managing private information. Due to the different information advantages, HFTs and dealers complement each other. HFTs contribute the majority of the information share, mostly through frequent *quote* updates which incorporate public information. In contrast, dealers contribute to price discovery primarily through *trades* that impound private information. Given the large size of the FX market (more than \$ 5 trillion daily volume) and its importance in the real economy, understanding such a coexistence structure is of first-order importance for policymakers. Second, as HFTs have made inroads into the previously dealership-only markets, there are growing concerns about HFTs' low risk-bearing capacity and the lack of flexibility of human adjustment, which could affect liquidity resilience (see e.g., Chaboud, Rime, and Sushko, 2022). Our results suggest that HFTs provide resilient liquidity during market-wide volatility spikes, but they withdraw liquidity in situations in which information asymmetry is more likely and the interpretation of fundamental information is crucial, that is ahead of scheduled macro announcements and in regime-shifting events such as the 2015 Swiss franc 'de-peg' event. We find that prior to scheduled macro announcements HFTs significantly reduce their liquidity provision while HFTs almost completely withdrew their liquidity through the regime change. In contrast, dealer banks provide robust liquidity throughout the scheduled macro announcements and the de-peg event.

The proprietary limit-order-book data[4] we use have two crucial features. First, each limit order-

unique structure of the FX market arose out of banks needing to provide the credit necessary to facilitate the T+2 settlement of large FX transactions.

[3]As dealer flows are composed of trades done on behalf of clients, as well as their own trades, we are unable to determine where this private information originates from, or whether dealers "learn" from their client flow. It is also worth noting that "private information" in an academic context merely refers to information about future prices. This can be from information generated from public information sources, such as analyst economic modeling or macroeconomic forecasts, for example.

[4]The data is from Refinitiv FX Spot Matching. It was also used in FCA Occasional Paper No. 46 (Evans, O'Neill,

book message includes a trader identifier, allowing us to classify the order submitter into a category, such as dealers or HFTs. Second, each limit-order-book message is timestamped to a high precision of a millisecond. Thus, we are able to build the complete and accurate play-by-play of the limit order book and compute the liquidity provision of each trader category on an order-level basis. In addition, it allows us to compute the information share of trades and quotes by trader category. To the best of our knowledge, we are the first to provide such participant-level liquidity and price discovery statistics in the FX market.

We report the key findings in three parts. We start by analysing the conventional metrics such as bid-ask spread (price of liquidity provision) and top-of-book depth (quantity of liquidity provision). We find that, compared to dealer banks, for GBP/USD (AUD/USD), HFTs' relative bid-ask spread is 34% (40%) lower and their top-of-book depth is 54% (131%) higher. HFTs' order-book liquidity is less sensitive to volatility spikes: in response to a large impulse of the stock-market volatility index (VIX), HFTs' bid-ask spreads increase by about 1%, which is only half of the response by dealers. However, a mixed picture of liquidity resiliency emerges. HFTs significantly reduce their liquidity provision ahead of informational events. During the minute before the scheduled macro announcement, HFTs' bid-ask spread widens by over 30%, which is much larger than the 10% widening by dealers. Nonetheless, HFTs are quick to replenish their order-book liquidity provision after announcements occur — when uncertainty is resolved. During a regime-shifting event – the 2015 Swiss franc 'de-peg' – we find that HFTs almost completely withdrew their liquidity provision immediately after the 'de-peg' announcement; in contrast, dealer banks continued providing liquidity.[5]

In the second step, to complement the conventional metrics that focus on passive 'resting liquidity', we examine liquidity provision in a general sense: traders can be considered as providing liquidity if they trade against transitory pricing errors (the so-called "leaning against the wind"), regardless of whether they are on the passive or aggressive side.[6] By estimating a state space

Rime, and Saakvtine, 2018). The sample covers two major currency pairs of GBP/USD and AUD/USD from 28 October 2012 to 5 June 2015.

[5]Breedon, Chen, Ranaldo, and Vause (2018) study the same event and focus on the trading in the Swiss franc pairs on EBS, another major D2D market. They contrast algorithmic traders (including both banks and HFTs) and human traders and find the former withdrew liquidity while the latter did the opposite.

[6]The idea is intuitive, by trading against transitory pricing errors they reduce the frequency of large, temporary price dislocations, a characteristic of an illiquid market. For example, during a flash crash, traders who buy, either aggressively via market(able) orders or passively via limit orders, can be viewed as providing liquidity as they trade

model, we show that dealers' aggressive trade flows create much larger pricing errors than HFTs'. For GBP/USD (AUD/USD), 1 million USD of dealers' aggressive trade flow, on average, leads to a pricing error of 0.087 (0.072) basis points, while HFTs' leads to only 0.008 (0.022) basis points — a tenth of dealers'. In addition, HFTs' passive trade flows absorb, on a per-dollar basis, larger pricing errors than dealers'. For GBP/USD (AUD/USD), 1 million USD of HFTs' passive trade flows, on average, reduce pricing errors by 0.087 (0.076) basis points, while dealers' reduce pricing errors by only 0.017 (0.021) basis points. The above results indicate that, beyond order-book liquidity provision, HFTs contribute to market liquidity in a broader sense by passively absorbing the pricing errors created by dealers' liquidity demands.

The third part of the results relates to price discovery. We find that dealers' *trades* are more informative than HFTs' while HFTs' *quotes* are, in general, more informative than dealers'. For GBP/USD (AUD/USD), a dealer's trade leads to an average permanent price impact of 0.31 (0.42) basis points, which is around 20% higher than HFTs'. In contrast, an HFT's limit order that improves the best bid or ask price leads to an average permanent price impact of 0.28 (0.45) basis points, which is around 25% higher than dealers'. In terms of information share, which weighs the permanent price impact of a trade or order variable by its own variance, HFTs' quote messages contribute over 50% to overall price discovery. This result is perhaps not surprising as HFTs account for the majority of top-of-book quote updates, but demonstrates that dealers and HFTs contribute to price discovery in distinct ways.

Our paper relates to several strands of literature. First, on the microstructure of exchange rate dynamics, Evans and Lyons (2002) and Evans (2002) identify order flow as an important determinant of exchange rates. Further studies link exchange rates to aggregated order flows (Payne, 2003; Bjønnes and Rime, 2005; Breedon and Vitale, 2010; Evans, 2010), or disaggregated order flows (Evans and Lyons, 2006; Breedon and Vitale, 2010; Cerrato, Sarantis, and Saunders, 2011; Osler, Mende, and Menkhoff, 2011; Breedon and Ranaldo, 2013; Menkhoff, Sarno, Schmeling, and Schrimpf, 2016). Leveraging the granularity of our dataset, we are able to construct not only disaggregated order flows, but also quote updates by dealers and HFTs. This allows us to understand the complementary roles of the traditional and new market makers and to further dis-

against a negative, transitory pricing error (Menkveld, 2013; Brogaard, Hendershott, and Riordan, 2014).

aggregate the contribution of order flows and quote updates to exchange rates for the first time.[7] Interestingly, we find that quote updates, specifically HFT quote updates, contribute a larger share to price discovery than order flows.

Second, a major focus of FX research has been on price discovery with limited research on liquidity provision. Mancini, Ranaldo, and Wrampelmeyer (2013), Karnaukh, Ranaldo, and Söderlind (2015), and Ranaldo and Santucci de Magistris (2019) document overall market liquidity of certain currency pairs. A handful of papers examine individual liquidity providers, Bjønnes and Rime (2005) characterise the inventory management of four inter-bank dealers and Bjønnes, Rime, and Solheim (2005) find that non-financial customers are the main liquidity providers during overnight hours. Several papers focus on liquidity during extreme events such as the Swiss franc "de-peg" in January 2015 (Breedon, Chen, Ranaldo, and Vause, 2018b) and the pound sterling flash crash in October 2016 (Bank for International Settlements, 2017; Noss, Pedace, Tobek, Linton, and Crowley-Reidy, 2017). We make two important contributions to the FX liquidity provision literature. First, we decompose the order-book liquidity provision of each trader category, comparing and contrasting HFTs with dealers. Second, we provide more general evidence of order-book liquidity provision by dealers and HFTs via a vector autoregressive model with exogenous variables approach (VARX) — examining their response to adverse, but less extreme, market conditions rather than focusing only on rare and extreme events.

Third, the paper relates to the high-frequency trading literature. While HFTs have been extensively studied in the equities market[8], their role in OTC markets, such as FX, is not well documented. In the equities market, past literature has shown that HFTs play an important role in both liquidity provision (eg Menkveld, 2013; Korajczyk and Murphy, 2019; Van Kervel and Menkveld,

[7] Order flows are believed to reflect either private information of the market participants (for example, their private assessment of the state of the economy) or their disagreement on the interpretation of public macroeconomic news. So in addition to only studying the bilateral relationship between exchange rate and order flow, as in the above papers, several papers set out to study the interplay between exchange rate dynamics, macroeconomic news, and order flows. For example, Evans and Lyons (2005) obtain retail flow data from Citibank and show macro news induces retail trades even days after the announcement. Froot and Ramadorai (2005) obtain institutional flow data from StateStreet, showing that it is only related to transitory returns, not permanent ones. Love and Payne (2008) show that order flows contribute a significant share (roughly one-third) of price discovery, even under scheduled macroeconomic announcements which are publicly and simultaneously released to all market participants. Evans and Lyons (2008) consider the indirect impact of macroeconomic news on prices through order flow variance and show that macroeconomic news explains 30% of aggregate daily price variance. Rime, Sarno, and Sojli (2010) add similar empirical results and show that order flow is related to current and expected macroeconomic fundamentals.

[8] See Menkveld (2016), for example, for a comprehensive review on the economics of high-frequency trading.

2019) and price discovery (eg Brogaard, Hendershott, and Riordan, 2014). However, it is not clear whether these findings apply to the FX market, given its unique market structure. The FX market is highly fragmented and dealers play a dominant role, especially in the D2C segment. With a strong OTC client network, dealers are more exposed to privately informed traders and may themselves become informed as a result. In contrast, HFTs are mainly active on primary inter-dealer platforms such as Refinitiv and EBS and thus only see a small fraction of total market flow.[9] As a comparison, in the US equities market, where most existing HFTs studies are based, HFTs can not only observe the majority of lit flows by subscribing to proprietary feeds of lit exchanges but also some dark flows through payment-for-order-flow arrangements with retail brokers.[10] Our findings reveal several key differences in liquidity provision and price discovery by HFTs and dealer banks. That being said, HFTs' roles in the FX market remain rather similar compared to the equities market.

The remainder of the paper proceeds as follows. Section 2 provides an overview of the dataset and introduces our trader classification scheme. Section 3 examines liquidity provision by dealers and HFTs. Section 4 documents dealers' and HFTs' role in price discovery. Section 5 concludes.

2 Data

2.1 Data overview

We use a proprietary dataset from Refinitiv FX Spot Matching ("Matching" hereafter), which contains all limit-order-book event messages from its matching engine (for example, new limit order submissions, cancellations, executions).[11] Messages are timestamped to the millisecond. The sample covers 2 major currency pairs predominantly traded on Matching, GBP/USD and AUD/USD, for approximately 2.5 years from 28 October 2012 to 5 June 2015. Matching runs for approximately 5.5 days each week, from 04:45 Sydney Time on Monday and continuously until 18:00 New York Time on Friday. Trades on Matching are all wholesale as the minimum trade size is 1m units of the base currency.

[9] Several HFTs have set up their own D2C platforms recently, however.

[10] Payment-for-order-flow arrangements are banned in the UK, though several HFT firms operate 'Systematic Internalisers' on equity markets which allow them to interact with some client flow directly.

[11] The same data was used in FCA Occasional Paper No. 46 (Evans, O'Neill, Rime, and Saakvtine, 2018).

Compared with other datasets used in the existing FX literature, the notable feature of our dataset is its granularity. First, it contains participant identifiers, based on which, we can classify them into different categories. In particular, we are able to cleanly identify HFTs which increasingly play an important role in FX markets. We provide a detailed account of our classification strategy in Section 2.2 below.

Other papers using disaggregated data include Chaboud, Chiquoine, Hjalmarsson, and Vega (2014), who obtain trade data from EBS where counterparties are labeled as either human or computer. However, computers can refer to a rather broad category and include both HFTs and buy-side executions algorithms (EAs). The EBS dataset used by Breedon, Chen, Ranaldo, and Vause (2018) is more granular, it includes trades whose counterparties are classified into human, bank-algorithmic and PTC-algorithmic (principal trading community), as well as anonymous quotes. However, as their study focuses on the Swiss franc de-peg event, their dataset covers only a short period of about three weeks around the event. A recent dataset from CLS used in, for example, Ranaldo and Somogyi (2021) sorts participants into four distinct categories: corporate, funds, non-bank financial firms and banks. However, in these studies it is not clear to which category HFTs belong.[12]

Second, our dataset contains not only executions, but also all other order-book event messages such as new limit order submissions and cancellations. Thus we are able to build a limit order-book *with trader identity information* and attribute order-book liquidity provision to each trader category. For example, we can compute each trader category's own bid-ask spread and top-of-book depth. Moreover, we can characterise the specific way participants contribute to price discovery: whether it is through their quote updates or through their trading.

We acknowledge that the dataset has its own limitations. The FX market is a highly fragmented OTC market with heterogeneous trading platforms and Matching only covers one particular segment of the brokered inter-dealer market. So compared with other recent studies using CLS data, our data is less representative in terms of global activity. However, Cespa, Gargano, Riddiough, and Sarno (2021) show that the correlation between Refinitiv volume and CLS is quite high across different currency pairs. Moreover, HFTs are mostly active in primary inter-dealer platforms such

[12] There are several papers which use disaggregated client trades data from a specific dealer bank (Evans and Lyons, 2006; Breedon and Vitale, 2010; Cerrato, Sarantis, and Saunders, 2011; Osler, Mende, and Menkhoff, 2011).

as EBS and Refinitiv Matching. So Matching is the most relevant place to study the interactions between dealers and HFTs.

Other datasets we use in the analysis include Bloomberg's macroeconomic news announcements data which include news items timestamped to the nearest second. In addition, we obtain minute-by-minute VIX data for the US, UK and Australian stock markets from Bloomberg.

2.2 Trader classification

As mentioned above, one crucial advantage of our dataset is that it contains participant information. Specifically, each message is associated with a four character Terminal Controller Identifier (Dealing) Code (TCID), which reconciles to the legal entity name of the trading firm. It is worth noting that participants trading through dealers as prime-brokerage clients are separately identified.[13] Based on such information, we classify participants into the following groups: (1) dealers ("Dealers"); (2) non-dealing commercial banks ("Commercial Banks"); (3) proprietary trading firms and hedge funds (PTFs & HFs); (4) other participants such as commercial firms, custodians and private banks ("Other"). Evans, O'Neill, Rime, and Saakvtine (2018) use the same dataset and we refer the reader to their paper for a detailed description of all participant groups. We further identify HFTs from PTFs & HFs TCIDs based on both a speed and a position rule specified below.

Speed rule: HFTs are fast. Inspired by Aquilina, Budish, and O'Neill (2022), we measure the speed of a TCID by computing its reaction time in the following two order-book events:

- "Add-Take" event: submitting an aggressive market(able) order after a new limit order is submitted by a different TCID which improves the best bid or ask price or adds depth to an existing best bid or ask price
- "Take-Cancel" event: submitting a cancel order after a partial or full execution of a resting limit order at the same price

Specifically, for each TCID, we first identify all occurrences of the two order-book events, if

[13] For example, a client that has a large investment bank as their prime-broker is identified as a separate trading participant, rather than being amalgamated within the prime-broker's flow.

any, in each week[14] and compute reaction times. Then for each week we take the 1% quantile of the reaction time distribution and use the average calculated across each of these weeks as our final speed measure for the TCID.

Inventory rule: HFTs carry low inventory. One defining characteristic of HFTs is that they carry a low, if not zero, overnight position (for example, Menkveld, 2013). To construct a measure of inventory that is comparable across TCIDs, we use the ratio of end-of-day position to volume.

Combining both the speed and inventory rule, a PTF or hedge fund TCID is classified as a HFT if it satisfies:

$$\begin{gathered}\text{End-of-day position to volume} < 0.2 \\ \underline{\text{and}} \\ \text{Add-Take speed} < 10\text{ms} \quad \underline{\text{or}} \quad \text{Take-Cancel speed} < 10\text{ms}\end{gathered} \tag{1}$$

The 0.2 threshold for the inventory rule is taken from Brogaard, Hendershott, and Riordan (2019). The 10-millisecond threshold for the speed rule is determined by our observation that it captures most of the legal entities that are self described or can be considered HFTs. While this reaction time is slower than in Aquilina, Budish, and O'Neill (2022), for example, market data from "Matching" is disseminated on a much slower frequency than in equities markets.[15]

2.3 Summary statistics of liquidity and trading of GBP/USD and AUD/USD

Before presenting the liquidity provision statistics of dealers and HFTs, we first show several high-level summary statistics of the liquidity and trading of GBP/USD and AUD/USD, our 2 sample currency pairs.

In Table 1a, we report summary statistics of several common liquidity metrics for the 2 currency pairs. It shows that both are very liquid with high trading volume, tight bid-ask spread and deep top-of-book depth. For example, GBP/USD has an average daily volume (ADV) of about 7.7 billion GBP, a half relative quoted spread of 0.99 basis points and a top-of-book depth of 6.2m GBP. Then in Table 1b, we report summary statistics of several common trading metrics by trader

[14] Note that when looking for the response message in both events, we restrict ourselves to 500-millisecond windows after the triggering messages as we think it is sufficiently long for fast traders.

[15] See: www.fx-markets.com/technology/market-data/2474307/thomson-reuters-is-making-its-matching-data-feed-10-times-faster.

Table 1. Summary statistics of liquidity and trading of GBP/USD and AUD/USD.

(a) Summary statistics by currency pair. *Volume* is market trading volume. *# Trade* is total number of trades. *RQS* is time-weighted (half) relative quoted spread. *DepthTop* is time-weighted top-of-book depth. Numbers reported are daily averages across all sample days.

	# Days	Volume (mil)	# Trade	RQS (bp)	DepthTop (mil)
GBP/USD	638	7721	5329	0.99	6.20
AUD/USD	630	13189	7962	1.36	10.55

(b) Summary statistics by trader category. *Total Trading Volume* is a trader category's trading volume (double-counted, that is, the sum of both aggressive and passive volume). *Total Trading Volume Share* is a trader category's trading volume (double-counted) as a percentage of total trading volume (double-counted). *Passive Trading Volume* is a trader category's passive trading volume (that is, when on the passive side of a trade). *Passive Trading Volume Share* is a trader category's passive trading volume as a percentage of total trading volume. *Passiveness* is a trader category's passive trading volume as a percentage of its own total trading volume. Numbers reported are daily averages across all sample days.

	Category	# Days	Total Trading Volume (mil)	Total Trading Volume Share (%)	Passive Trading Volume (mil)	Passive Trading Volume Share (%)	Passiveness (%)
GBP/USD	Dealer	638	7030	45.3	4323	56.3	62.2
	HFT	638	5173	33.5	1351	17.0	25.2
	Commercial Bank	638	1737	11.2	1195	15.6	69.7
	Non-HFT PTFs & HF	638	521	3.4	358	4.7	69.8
	Other	638	980	6.5	495	6.4	51.3
AUD/USD	Dealer	630	10972	41.4	5906	44.9	54.3
	HFT	630	8685	33.1	3161	23.5	35.6
	Commercial Bank	630	4962	18.8	3039	23.2	62.0
	Non-HFT PTFs & HF	630	741	2.8	498	3.8	68.3
	Other	630	1016	4.0	585	4.5	58.0

category. There are two notable observations. First, for both currency pairs, dealers and HFTs combined contribute to over 70% of trading volume. Second, HFTs are more aggressive (being on the aggressive side of a trade) than dealers and commercial banks. For example, for GBP/USD (AUD/USD), 74.8% (64.4%) of HFTs' trading volume is aggressive. In contrast, dealers' aggressive volume is 37.8% (45.7%) of their total trading volume in GBP/USD (AUD/USD).[16]

In addition, Figure 1 plots the intraday trading volume on Matching, showing a striking intraday seasonality. Although Matching runs continuously throughout the day, the market volume

[16] As a comparison, studies in the equities market which report similar statistics show that HFTs are less aggressive: 40% for Nordic stocks (Hagströmer and Nordén, 2013), 50% for E-mini (Kirilenko, Kyle, Samadi, and Tuzun, 2017) and 47% for Canadian stocks (Brogaard, Hendershott, and Riordan, 2019).

Figure 1. Intraday passive volume. This figure plots the intraday passive volume decomposed to different trader categories at 10-minute intervals.

(a) GBP-USD

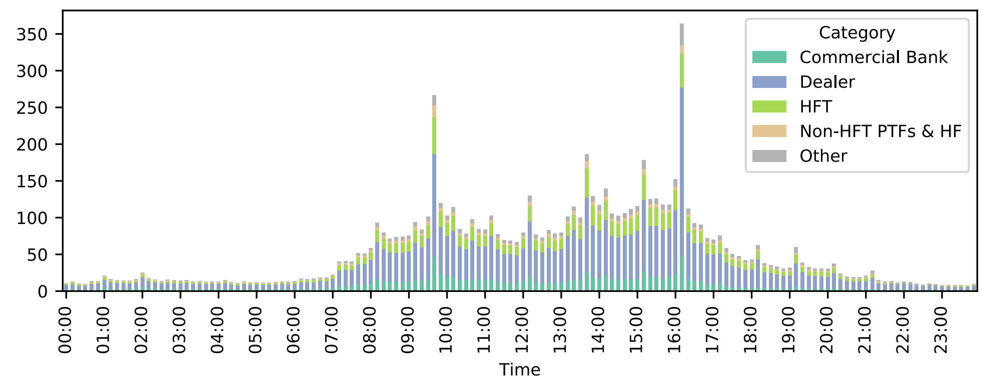

(b) AUD-USD

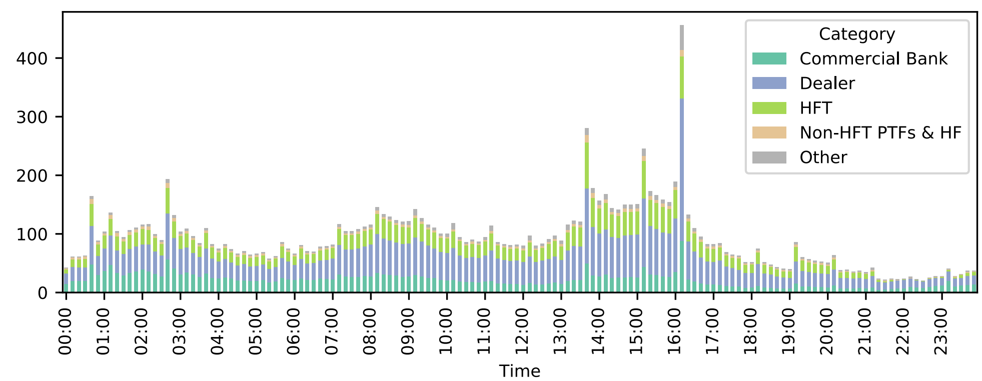

of the 2 currency pairs is largely concentrated during their respective primary trading hours. For GBP/USD, most trading volume happens during UK and US market hours, which are roughly between 07:00 and 20:00 London Time.[17] In contrast, intraday trading volume in AUD/USD spreads more smoothly from early Asian/Australian market hours towards the end of US market hours, roughly between 0:00 and 20:00 London Time. In all of our following analyses, we restrict our sample to the respective primary trading hours of the 2 currency pairs, that is, 07:00 - 20:00 for GBP/USD and 00:00 - 20:00 for AUD/USD.

[17]Without further mentioning, all times refer to London Local Time.

Table 2. Overview order-book liquidity provision by trader category. This table presents several stylised facts of the liquidity provision of different trader categories. *RQS* is the (half) relative bid-ask spread. *DepthTop* is the top-of-book depth. *DepthTopShr* is a trader category's top-of-book depth as a fraction of the total top-of-book depth. Means and standard deviations are computed across daily averages.

	Measure	RQS (bp)		DepthTop (mil)		DepthTopShr	
		Mean	SD	Mean	SD	Mean	SD
GBP/USD	Dealer	1.34	0.29	2.30	0.56	0.33	0.06
	HFT	0.88	0.15	3.54	0.84	0.50	0.07
AUD/USD	Dealer	1.97	0.51	2.71	0.62	0.25	0.05
	HFT	1.18	0.23	6.25	1.67	0.55	0.08

3 Liquidity provision roles of dealers and HFTs

3.1 Order-book liquidity provision measures

The textbook definition of a liquidity supplier is any trader who adds limit orders to the limit order-book. And a trader who quotes a narrower bid-ask spread and contributes more top-of-book depth (or depth accumulated across several levels of the order-book) is viewed as providing more liquidity. As all order-book event messages in our dataset have legal entity identification, we are able to build a limit-order-book with trader identity information and thus compute two common liquidity metrics, bid-ask spread and top-of-book depth, *for each trader category separately*. Specifically,

- The relative (half) quoted spread of trader category i at time t is computed as:

$$RQS_{it} = \frac{BO_{it} - BB_{it}}{2 \times Mid_t} \tag{2}$$

  where BB_{it} (BO_{it}) is the best bid (ask) price of trader category i at time t and Mid_t is the market midquote.

- The top-of-book depth contribution by trader category i, $DepthTop_{it}$, is computed as:

$$DepthTop_{it} = Q_{it}^{BO_t} + Q_{it}^{BB_t} \tag{3}$$

  where $Q_{it}^{BB_t}$ ($Q_{it}^{BO_t}$) is the total quantity of all resting orders of trader category i at the *market* best bid (ask) at t.

Table 2 reports the daily average and standard deviation of the above two order-book liquidity metrics by dealers, HFTs and other trader categories. It shows that in terms of order-book liquidity metrics, HFTs quote the tightest (half) bid-ask spread (0.88 basis points for GBP/USD and 1.18 basis points for AUD/USD) and contribute the largest top-of-book depth (3.54m GBP/USD and 6.25m for AUD/USD). In relative terms, for GBP/USD (AUD/USD), HFTs' bid-ask spread is about 34% (40%) lower and their top-of-book depth is about 54% (131%) higher than dealers'. So on average or during normal market times, we find HFTs provide better order-book liquidity than dealers. However, it is worth noting that in this study we focus only on a subset of the total FX market, ignoring platforms operated by dealers (single-dealer platforms) in the 'dealer-to-client' segment. As these venues comprise a much larger share of the total FX market than the inter-dealer venue used in this study, dealers remain the dominant provider of liquidity in the FX market as a whole.

3.2 Order-book liquidity provision and adverse market conditions: a VARX approach

Although HFTs provide better order-book liquidity during normal market times, a key unanswered question is whether they continue to do so during adverse market conditions. For example, periods of volatility (as evidenced from VIX spikes) or ahead of scheduled macroeconomic news announcements. Practitioners, academics and regulators have voiced their concern about the fleeting nature of market liquidity, especially in fast-paced electronic markets (Bank for International Settlements, 2018): market liquidity may be abundant when markets are tranquil but may quickly disappear when volatility increases. It is a particular concern in a market where both dealers and HFTs are endogenous liquidity providers with no formal liquidity provision obligations.

Here we focus on two types of adverse market conditions. We note that our sample period does not include recent periods of high volatility, such as the 2022 UK's Gilt market volatility, the 2020 COVID pandemic and the UK's 2016 EU membership referendum. Nonetheless, various other extreme events are included, such as the 2015 Swiss 'de-peg' event. The first group of events we examine are periods with a significant increase in market-wide volatility: an increase in the VIX — an index of the implied volatility of options on the S&P500 index. If a positive VIX shock implies a long-lasting increase in market volatility, then market makers are more exposed to inventory risks, and reduce their liquidity provision in response. Second, we turn to periods ahead of scheduled macroeconomic news announcements. These include significant announcements like "non-farm payrolls" which have a significant impact on FX prices. The difference between these periods is that the first period represents a more long-term increase in volatility, which affects a wide range of assets. The second period represents a discrete period in time where uncertainty is high, but is immediately resolved and affects a much smaller range of assets. We expect HFTs and dealers to have differing abilities to handle these volatility types. HFTs are more successful at responding to 'public information', which will be more abundant in high-VIX periods with sustained movements in a larger set of correlated securities. Dealers will be more successful at managing 'private information' with their broader client networks, which will be more abundant ahead of macroeconomic news releases.[18]

VARX model specification To examine the response of dealers' and HFTs' order-book liquidity provision to adverse market conditions, we follow Menkveld, Yueshen, and Zhu (2017) and use a vector-autoregressive model with exogenous variables (VARX) specified as follows:

$$y_t = \alpha + \Phi_1 y_{t-1} + \cdots + \Phi_p y_{t-p} + \Psi z_t + \varepsilon_t \tag{4}$$

where t indexes time intervals, y_t is a vector of endogenous variables and z_t is a vector of exogenous variables believed to be determined outside the system of endogenous variables. While $\Phi_1, \ldots, \Phi_p$ captures the dynamics within the system of endogenous variables, Ψ captures the contemporaneous impact of the exogenous variables. ε_t is a vector of white noise, that is, $E(\varepsilon_t) = 0$, $E(\varepsilon_t \varepsilon_t') = \Sigma_\varepsilon$ and $E(\varepsilon_t \varepsilon_s') = 0$ for $s \neq t$. The common way of summarizing the dynamics within the VARX system is to compute the impulse response functions (IRFs). Impulse responses to endogenous variables can be obtained by converting the vector autoregressive (VAR) part of the VARX model into its corresponding vector moving-average (VMA) form and the VMA coefficients are by definition impulse responses. The impulse responses to exogenous variables can be easily obtained by scaling

[18] See Aquilina, Budish, and O'Neill (2022) for a theoretical and empirical method of separating public and private information.

the shocks by Ψ[19].

Endogenous variables The vector of endogenous variables we include in our VARX model is

$$y = \begin{pmatrix} LqtPro_{Dealer} & LqtPro_{HFT} & Vlm & TrdImb & Vol \end{pmatrix}'$$

where $LqtPro_{Dealer}$ and $LqtPro_{HFT}$ are the order-book liquidity provision measures (that is, RQS_{it} or $DepthTop_{it}$) for dealers and HFTs respectively. *Vlm* is the market trading volume. *TrdImb* is the market-wide absolute trade imbalance (that is, buyer-initiated trades minus seller-initiated trades). *Vol* is the (endogenous) currency rate volatility measured by the difference between the highest and lowest midquote during the interval, normalised by the average of the two.[20]

Construct composite VIX innovations The most commonly used proxy for market-wide volatility is the stock-market VIX index. As we have 2 currency pairs, GBP/USD and AUD/USD, there are 3 relevant volatility indices: VIX for the US market, VFTSE for the UK market and AVIX for the Australian market.[21] As Figure 2 shows, the three volatility indices have different trading hours and none of them can fully cover the primary trading hours for either currency pair (08:00 - 20:00 for GBP/USD; 00:00 - 20:00 for AUD/USD). So we construct a composite time series of VIX shocks in the following manner. First, we fit an autoregressive (AR) model[22] to the first difference of each volatility index and obtain its innovations: $dVIX^*$, $dVFTSE^*$ and $dAVIX^*$. Second, to make it comparable between volatility indices, we standardise each innovation time series so that it has zero mean and unit variance. Then we are ready to define the composite VIX innovation for GBP/USD and AUD/USD respectively. For GBP/USD, $dCompVIX^* = dVFTSE^*$ between 08:00 and 13:30; $dCompVIX^* = dVIX^*$ between 13:30 and 20:00. For AUD/USD, $dCompVIX^* = dAVIX^*$ between 00:00 and 08:00; $dCompVIX^* = dFTSE^*$ between 08:00 and 13:00; $dCompVIX^* = dVIX^*$

[19]To compute the confidence intervals of the impulse response to the exogenous variables, we follow Menkveld, Yueshen, and Zhu (2017) and use a Monte Carlo simulation approach. Specifically, confidence intervals are obtained based on ten thousand Monte Carlo draws from the distribution of the system parameters $\Phi_1, \ldots, \Phi_p$. Significance level thresholds are chosen at 5%, two-sided.

[20]We use the high-minus-low volatility as opposed to realised volatility because the short time interval of one minute makes the estimation of realised volatility difficult.

[21]VFTSE (AVIX) is constructed based on the implied volatilities of the FTSE 100 (S&P/ASX 200) constituent stocks.

[22]The lag order is chosen by Bayesian Information Criterion.

Figure 2. VIX timelines. This figure illustrates the trading hours of different volatility indices and our constructed composite VIX innovations. *CompVIX* stands for the composite VIX we construct. *VIX*, *VFTSE* and *AVIX* are the volatility indices constructed based on the implied volatilities of the constituent stocks of the U.S. S&P 500 index, U.K. FTSE 100 index and Australian S&P/ASX 200 respectively.

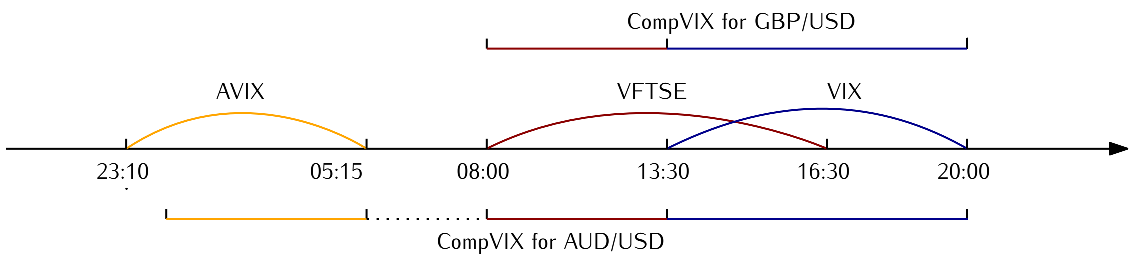

between 13:30 and 20:00. In Figure 2 we provide a schema of the VIX construction. Note that for the period between 13:30 and 16:30 when *both* VFTSE and VIX are available, we take innovations in VIX given the importance of the US market. Last, we take the non-negative part of the composite VIX innovations: $dCompVIX^*_+(t) = \max\{0, dCompVIX^*(t)\}$.

Exogenous variables The vector of exogenous variables we include in the VARX model are

$$z = \begin{pmatrix} \mathbb{1}_{dCompVIX^*_+} & News_{\text{-}5min} & \cdots & News_{5min} \end{pmatrix}'$$

where $\mathbb{1}_{dCompVIX^*_+}$ is a dummy variable which equals one when $dCompVIX^*_+$ is above its 90% quantile. We use the VIX innovation dummy in our baseline VARX model in order to focus on periods with large and positive volatility spikes.[23] $News_{\text{-}5min}, \ldots, News_{5min}$ is a set of per-minute time dummies constructed from the Bloomberg news dataset. For example, $News_{1min}$ equals one for the 1-minute interval immediately after the news announcement and zero otherwise. $News_{\text{-}1min}$ indicate the minute right before the announcement. Other news time dummies are defined similarly.

VARX implementation details First, we estimate the VARX model at 1-minute frequency, which is the highest frequency we can obtain for the volatility indices. As a result, we aggregate all endogenous variables at the 1-minute frequency. For example, for stock variables such as RQS_{it} and $DepthTop_{it}$ which are first computed tick by tick, we aggregate them at the 1-minute

[23]We have tried the 75% quantile for the construction of the VIX dummy variable and the results are qualitatively similar.

Table 3. Summary statistics of the endogenous and exogenous variables in VARX. This table shows the summary statistics of the endogenous and exogenous variables used in the VARX model. RQS_{Dealer} and RQS_{HFT} are the (half) relative bid-ask spread quoted by dealers and HFTs respectively. $DepthTop_{Dealer}$ and $DepthTop_{HFT}$ are the top-of-book depth provided by dealers and HFTs respectively. *Vlm* is market volume. *TrdImb* is trade imbalance, that is, buyer-initiated trade volume minus seller-initiated trade volume. *Vol* is the high-minus low volatility. All variables are sampled at 1-minute frequency.

		N	Mean	SD	Min	Q25	Q50	Q75	Max
GBP/USD	RQS_{Dealer}	454801	2.66	1.25	0.00	1.76	2.41	3.41	54.98
	RQS_{HFT}	454801	1.76	0.78	0.58	1.46	1.72	1.95	81.95
	$DepthTop_{Dealer}$	454801	2.32	2.98	0.00	1.04	1.79	2.79	202.60
	$DepthTop_{HFT}$	454801	3.56	1.65	0.00	2.43	3.34	4.46	15.47
	Vlm	454801	9.16	18.70	0.00	1.00	4.00	11.00	1192.00
	TrdImb	454801	3.63	6.52	0.00	1.00	2.00	4.00	648.00
	Vol	454801	1.73	1.58	0.00	0.89	1.33	2.24	89.56
	$\mathbb{1}_{dCompVIX^*_+}$	454801	0.10	0.30	0.00	0.00	0.00	0.00	1.00
AUD/USD	RQS_{Dealer}	573431	4.13	2.17	0.00	2.41	3.60	5.44	48.00
	RQS_{HFT}	573431	2.45	0.84	0.95	2.05	2.33	2.85	75.26
	$DepthTop_{Dealer}$	573431	2.70	3.22	0.00	1.16	2.00	3.24	182.16
	$DepthTop_{HFT}$	573431	6.13	3.06	0.00	3.89	5.76	8.03	23.96
	Vlm	573431	10.77	23.16	0.00	1.00	5.00	12.00	1560.00
	TrdImb	573431	4.31	8.56	0.00	1.00	2.00	5.00	808.00
	Vol	573431	2.11	2.07	0.00	1.06	1.66	2.75	163.56
	$\mathbb{1}_{dCompVIX^*_+}$	573431	0.10	0.30	0.00	0.00	0.00	0.00	1.00

frequency by computing time-weighted averages. Second, we remove any intraday seasonality by regressing each of the endogenous variables on 10-minute interval dummies and only use their residuals in the VARX estimation. Third, we follow Hasbrouck (1991) and insert missing values during the overnight periods. Essentially, we only use intraday variation in our estimation and consider each trading day independently. Fourth, to choose the optimal number of lags for the VARX model, we apply the Bayesian Information Criterion (BIC) on each day and obtain the mode lag order across all days. This is to avoid over-fitting the model, given we have a large number of 1-minute frequency intraday observations for a three-year sample. We find the optimal lag period to be one.

Summary statistics of the endogenous and exogenous variables in VARX Table 3 reports the summary statistics of all endogenous and exogenous variables used in the VARX model. Consistent with the daily measures reported in Table 2, HFTs quote a narrower bid-ask spread and supply

Figure 3. Impulses responses of order-book liquidity to a positive and large VIX impulse. This figure plots the impulse responses of dealers' and HFTs' *RQS* and *DepthTop* to $\mathbb{1}_{dCompVIX^*_+}$, the VIX dummy variable which equals one when the VIX increase is beyond its 90% quantile. *RQS* and *DepthTop* are in natural logs so that the unit of measure is in percentages.

(a) GBP/USD

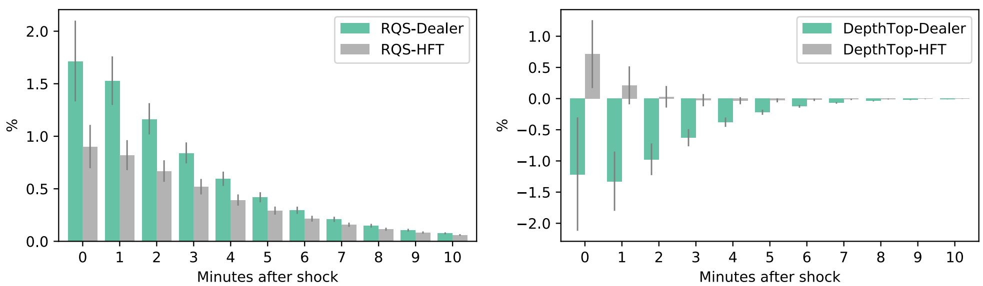

(b) AUD/USD

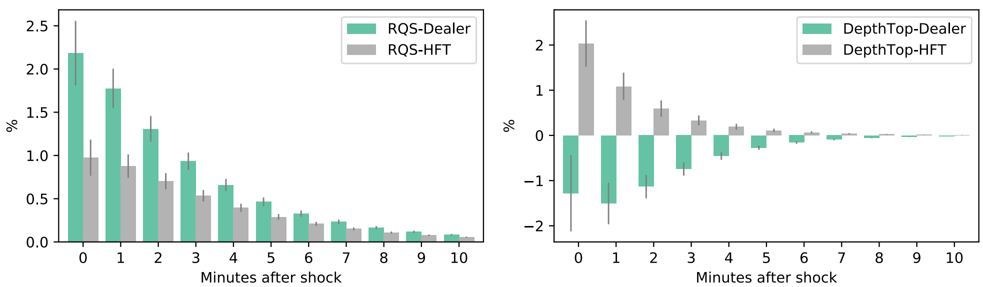

more top-of-book depth at the 1-minute frequency. One interesting observation is that, relative to the mean, the standard deviation of HFTs' top-of-book depth is much smaller than dealers', indicating that their top-of-book depth is, in general, more stable.

Order-book liquidity response to a positive and large VIX impulse We start by examining the response of dealers' and HFTs' order-book liquidity provision to a positive and large VIX impulse. Figure 3 plots the impulse responses of dealers' and HFTs' in relative quoted spread (*RQS*) and top-of-book depth (*DepthTop*) to an impulse of $\mathbb{1}_{dCompVIX^*_+}$, our volatility dummy variable which equals one when the VIX increase is beyond its 90% quantile.[24] It shows that, dealers' liquidity provision is more sensitive to a contemporaneous positive and large VIX impulse compared with HFTs'. Dealer relative quoted spread widens by about 1.75% for GBP/USD and 2% for AUD/USD, which are roughly double the magnitudes of HFTs'. The same is true for depth, dealers' decreases by

[24]To make it comparable between dealers and HFTs, we use the natural logs of both liquidity measures so that we measure their sensitivity to the VIX dummy.

Figure 4. Order-book liquidity response to scheduled macroeconomic news. This figures plots the coefficients on the news dummies Ψ which measure changes in dealers' and HFTs' order-book liquidity around macroeconomic news announcements compared with non-news periods.

(a) GBP/USD

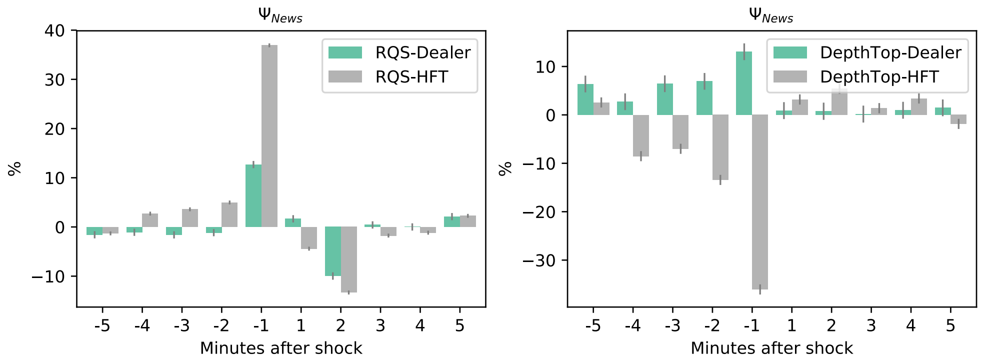

(b) AUD/USD

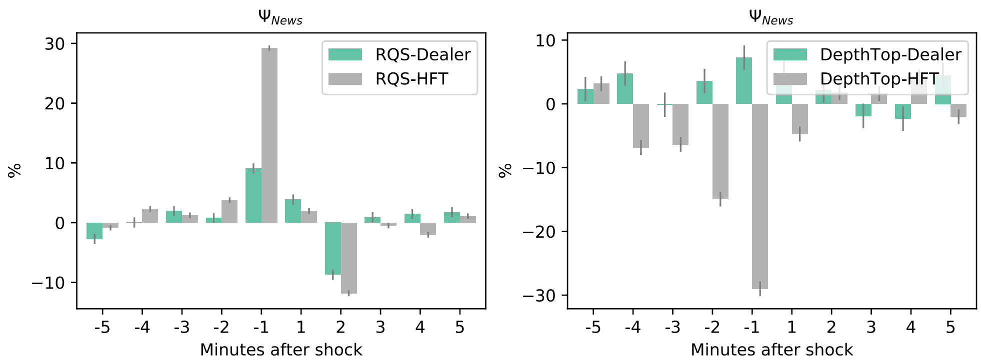

about 1% for both currency pairs while HFTs' increases. The increase in HFTs' top-of-book depth is explained by a displacement effect. When dealers widen their bid-ask spread more than HFTs, their quotes, although at worse prices, become the new BBO.

Order-book liquidity around scheduled macroeconomic news announcements We turn to periods around the macroeconomic news announcements. Figure 4 plots the coefficients on the news dummies, which represent the percentage changes in the dealers' and HFTs' order-book liquidity provision measures in the ten-minute window around news announcements, *relative to* non-news periods. The results show that, ahead of announcements, HFTs significantly withdraw their liquidity: for GBP/USD, their relative quoted spread increases by over 30% during the minute before the announcement and their top-of-book depth decreases by over 30%. In comparison, dealers' relative quoted spread increases by a much smaller magnitude of about 10%. As a result, their top-of-book depth increases by about 10%. However, it should not necessarily be interpreted as dealers providing more liquidity ahead of scheduled macroeconomic announcements. A displacement effect, wherein dealer bid-ask spreads widen less than HFTs' such that their BBO becomes the market BBO occurs. It should be noted that HFTs are quick to replenish their order-book liquidity right after the announcements. In the first minute after an announcement, HFTs' relative quoted spread even becomes tighter than non-news periods for GBP/USD and only slightly higher than non-news-periods for AUD/USD.

Order-book liquidity response to (endogenous) market conditions Last, Figure 5 plots the impulse responses of dealers' and HFTs' relative quoted spread and top-of-book depth to shocks in endogenous market condition variables. Although we can not draw causal conclusions, it is nonetheless meaningful to examine the associations between the two. We note several interesting observations. First, dealers' order-book liquidity is positively correlated with market volume. One standard deviation increase in market volume is associated with a 3% decrease in dealers' bid-ask spread and a 6% increase in their top-of-book depth in the following minute. Second, dealers' order-book liquidity provision is negatively associated with trade imbalance and (endogenous) currency rate volatility. One standard deviation increase in (endogenous) volatility is associated with a 4% increase in dealers' bid-ask spreads and 6% decreases in their top-of-book depth. Results for AUD/USD are qualitatively similar.

In summary, we interpret the above results as dealers and HFTs possessing different comparative advantages in market-making in the FX market. Dealers with higher expertise and more information from their OTC client network are better at coping with discrete periods of high information asymmetry that affect a single currency. In contrast, HFTs with faster and more advanced market-making algorithms can better navigate through more prolonged periods of uncertainty that affect a broader range of securities.

3.3 Liquidity provision during an extreme-volatility event

In the previous section, we find that compared with dealer banks, HFTs' order-book liquidity provision is less sensitive to stock-market VIX spikes. However, we caution that this finding does

Figure 5. Impulse responses of order-book liquidity to market conditions shocks. This figure plots the impulse response of dealers' and HFTs' order-book liquidity to market condition shocks. *Vlm* is market volume. *TrdImb* is market-wide absolute trade imbalance, that is, buyer-initiated trade volume minus seller-initiated trade volume. *Vol* is the exchange rate volatility measured as the difference of the highest and lowest midquote during the 1-minute interval, normalised by the average of the two. All market condition variables are standardised so that they have zero mean and unit variance. Both *RQS* and *DepthTop* in natural logs.

(a) GBP/USD

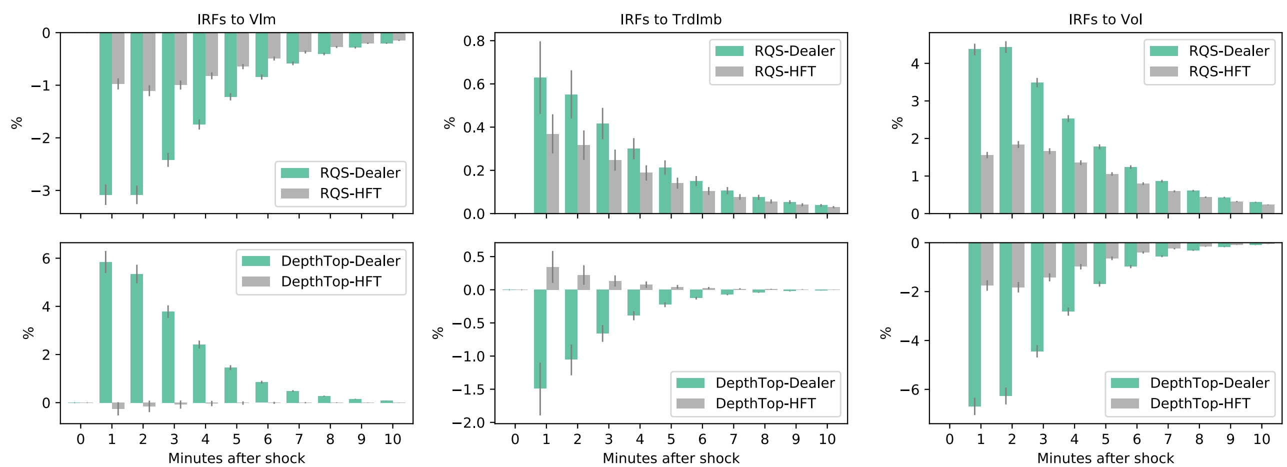

(b) AUD/USD

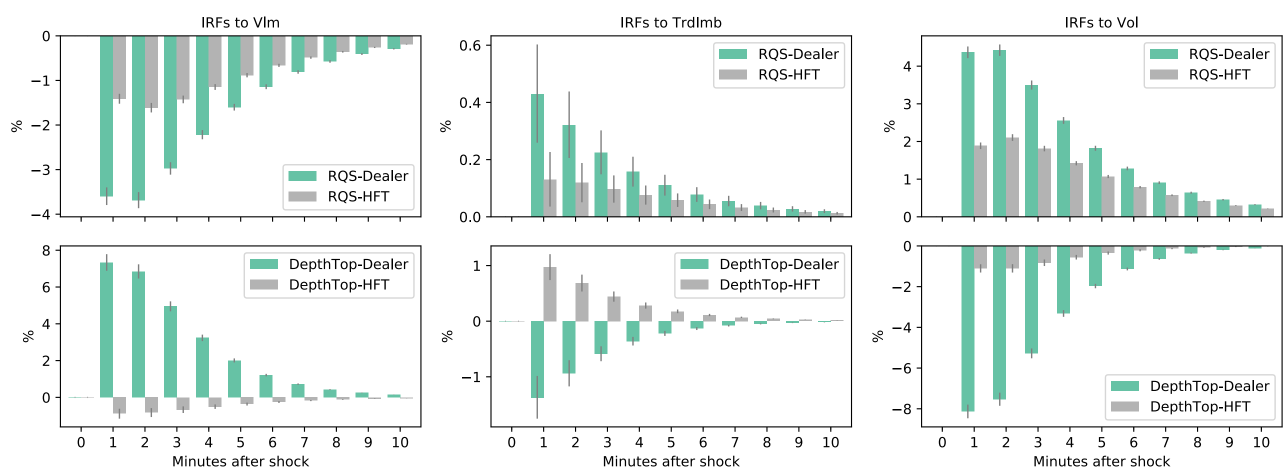

not necessarily generalise to volatility spikes in the FX market. While we select the VIX as an exogenous source of volatility, shocks in the stock market might not propagate to the FX market. In addition, although we focus on large VIX spikes defined as 1-minute increases beyond the 90% quantile, this cutoff may not be high enough to elicit only severe reactions. Thus, market events captured by our VARX analysis are likely to be a mix of both mild-volatility and extreme-volatility events.

To provide anecdotal evidence of the response of dealer banks' and HFTs' liquidity provision during extreme-volatility events, we focus on an extreme event in our sample period: the Swiss franc 'de-peg' on January 15, 2015. Around 9:30 UK time, the Swiss National Bank (SNB) announced unexpectedly that it will discontinue its three-year-old peg of 1.2 Swiss francs per euro. As a result, the Swiss franc soared by around 30 percent against the euro in a few minutes after the announcement. Although the announcement concerns the value of the Swiss franc and thus does not directly affect our two sample currencies, GBP and AUD, the de-peg event was highly unexpected and caused large losses to some market participants, affecting the FX market as a whole.

Figure 6 plots the order-book liquidity provision of dealer banks and HFTs, respectively, during the two-hour window around the announcement. The most striking result is that HFTs almost completely withdrew their liquidity provision immediately after the announcement, and only re-entered the market around half an hour later when they perceived market risk dropped back to its normal bands. In contrast, dealer banks remained present in the order book, contributing virtually all passive volume. This is likely due to their client network obligations: in order to continue serving their clients, they had to remain present in the inter-dealer platform to help them unwind positions.

3.4 General liquidity provision: a state space approach

In the previous section we show that, on average, HFTs provide better order-book liquidity, measured by bid-ask spread and top-of-book depth, than dealers. In addition, HFTs' order-book liquidity provision is more robust to sustained high volatility periods. However, being merely on the passive side of a trade is not consistent with liquidity provision in a more general sense. For example, traders who aggressively buy (that is, with marketable buy orders) during a flash crash can be considered as supplying liquidity as they trade against transitory pricing errors (that is, trade "against the wind").[25] The challenge to such a definition of liquidity supply is that one needs to first identify pricing errors. While pricing errors are evident during flash crashes or rallies by definition, it is not trivial to identify them during market normal times as price changes caused by trade flows almost always reflect both new information and mispricing.

State space model specification To identify pricing errors, we follow Brogaard, Hendershott, and Riordan (2014) and estimate a state space model specified as below:

$$
\begin{aligned}
\text{Midquote:} \quad & p_t = m_t + s_t \\
\text{Efficient price:} \quad & m_t = m_{t-1} + w_t \\
\text{Efficient price innovation:} \quad & w_t = \Sigma_i \lambda_i^j \widetilde{x}_{i,t}^j + \mu_t, \quad \mu_t \sim N(0, \sigma_\mu^2) \\
\text{Pricing error:} \quad & s_t = \phi s_{t-1} + \Sigma_i \psi_i^j x_{i,t}^j + v_t, \quad v_t \sim N(0, \sigma_v^2)
\end{aligned}
\tag{5}
$$

In the specification above, the log midquote p_t consists of 2 components: efficient price m_t and pricing error s_t. The efficient price m_t is a martingale[26] with innovation w_t, which represents new information and itself consists of 2 parts: the first part $\Sigma_i \lambda_i^j \widetilde{x}_{i,t}^j$ reflects private information revealed through trades; the second part μ_t reflects public information through, for example, quote revisions by market makers without trades. $\widetilde{x}_{i,t}^j$ represents the trade innovation[27] for trader group i and trade type j (trade by a trader group can either be aggressive or passive). Trade innovations, instead of trades themselves, are used in the efficient price equation because only unexpected trades contain new information. λ_i captures the permanent price impact of trader group i's trade flow, that is, impact on efficient price innovation, and a positive value suggests trader group i's trade flow are informed, contributing to price discovery.

The pricing error s_t is a stationary auto-regressive process of lag order one and consists of a trade-related component, $\Sigma_i \psi_i^j x_{i,t}^j$ where $x_{i,t}^j$ represents the signed trade flow for trader group i and

[25] Another subtle way of supplying liquidity can be that if HFTs or dealers trade aggressively to exploit arbitrage opportunities (such as triangular arbitrage) resulting from transient demand or supply shocks, they implicitly act as liquidity providers. Foucault, Kozhan, and Tham (2017) model and document such channel in detail.

[26] A martingale is a stochastic process for which the conditional expected value of the next observation, given all the past observations, is equal to the most recent observation, that is, $E(X_{n+1} \mid X_1, \ldots, X_n) = X_n$.

[27] The trade innovations are estimated through a VAR model with trade variables from all trader groups included in the state space model.

trade type j, and a non-trade-related component v_t. The trade-related component can be caused by, for example, price over- or under-reaction to trade flows or price pressures from market makers' inventory management. Our parameter of interest is ψ_i, which captures the transitory price impact of trader group i's trade flow, that is, impact on the pricing error. A negative value indicates that trader group i trades against the pricing error, or in a general sense, provides liquidity.

When estimating the state space model, we include trade flows from dealers, HFTs and all other trader categories combined[28]. In addition, we estimate the model with aggressive and passive trade flows separately so that we can distinguish the impact of each separately on the pricing error for each trader category. To be consistent with the structural VAR analysis above, we estimate the state space models at 1-minute frequency. The estimation is done for each day and statistical inferences are based on estimates across days.

Pricing errors and trade flows of dealers and HFTs Panel A of Table 4 reports the estimation results of the state space model with traders' *aggressive* trade flows. Focusing on the pricing error equation we find that, on a per dollar basis, dealers' aggressive trade flows create much larger pricing errors than HFTs'. For GBP/USD (AUD/USD), 1m USD of a dealers' aggressive trade flow, on average, leads to a pricing error of 0.087 (0.072) basis points while HFTs' leads to only 0.008 (0.022) basis points. Turning to the estimation results with *passive* order flows in Panel B of Table 4, we find that, on a per dollar basis, HFTs' passive trade flow reduces more pricing errors than dealers'. For GBP/USD (AUD/USD), 1m USD of HFTs' passive trade flow on average reduces pricing errors by 0.087 (0.076) basis points while dealers' reduces pricing errors by a much smaller magnitude of only 0.017 (0.021) basis points. In summary, the above results suggest that HFTs, in general, appear to beneficially impact market liquidity by passively absorbing the pricing errors created by dealers' liquidity-demanding trade flow.

As a robustness check, we use a simple and non-parametric way of examining whether dealers and HFTs trade in or against the direction of pricing errors in the Appendix. Specifically, we follow Brogaard et al. (2018) and focus on periods of extreme price movements (EPMs), de-

[28] As we focus on the comparison between dealers and HFTs, we merge all the other groups into a "Other" group. It should be noted that the "Other" group include very heterogeneous groups of traders such as principal trading firms and hedge funds which are not classified as HFTs (see our classification details in Section 2.2). It also includes some other small subgroups such as custodian, pension funds and etc.

Table 4. State space results. This table reports the estimation results of the efficient price innovation equation and the pricing error equation of the state space model specified in Equation 5. Panel A and B report the estimation results with aggressive trade flows and passive trade flows respectively. The estimation is done day by day and parameters are averages across sample days. * indicates that the difference between dealers and HFTs is significant at 1% significance level.

Panel A: Aggressive flow			Panel B: Passive flow		
Currency	GBP/USD	AUD/USD	Currency	GBP/USD	AUD/USD
Panel A1: Efficient price equation			Panel B1: Efficient price equation		
λ_{Dealer}^{Aggr}	0.087*	0.072*	λ_{Dealer}^{Pass}	−0.027*	−0.033*
	(0.003)	(0.002)		(0.002)	(0.002)
λ_{HFT}^{Aggr}	0.008	0.022	λ_{HFT}^{Pass}	−0.149	−0.126
	(0.003)	(0.003)		(0.009)	(0.005)
λ_{Other}^{Aggr}	0.067	0.065	λ_{Other}^{Pass}	−0.08	−0.068
	(0.005)	(0.004)		(0.004)	(0.003)
$\sigma(\tilde{x}_{Dealer}^{Aggr})$	4.418	6.119	$\sigma(\tilde{x}_{Dealer}^{Pass})$	5.523	5.942
$\sigma(\tilde{x}_{HFT}^{Aggr})$	4.221	4.719	$\sigma(\tilde{x}_{HFT}^{Pass})$	2.196	3.382
$\sigma(\tilde{x}_{Other}^{Aggr})$	2.372	3.36	$\sigma(\tilde{x}_{Other}^{Pass})$	3.387	4.815
σ_μ	0.996	1.114	σ_μ	0.975	1.099
Panel A2: Pricing error equation			Panel B2: Pricing error equation		
ψ_{Dealer}^{Aggr}	0.048*	0.049*	ψ_{Dealer}^{Pass}	−0.017*	−0.021*
	(0.003)	(0.002)		(0.002)	(0.002)
ψ_{HFT}^{Aggr}	0.016	0.024	ψ_{HFT}^{Pass}	−0.087	−0.076
	(0.003)	(0.003)		(0.009)	(0.005)
ψ_{Other}^{Aggr}	0.033	0.046	ψ_{Other}^{Pass}	−0.045	−0.036
	(0.005)	(0.004)		(0.003)	(0.003)
$\sigma(x_{Dealer}^{Aggr})$	4.66	6.551	$\sigma(x_{Dealer}^{Pass})$	5.835	6.328
$\sigma(x_{HFT}^{Aggr})$	4.353	4.903	$\sigma(x_{HFT}^{Pass})$	2.277	3.538
$\sigma(x_{Other}^{Aggr})$	2.502	3.555	$\sigma(x_{Other}^{Pass})$	3.572	5.196
ϕ	0.621	0.672	ϕ	0.638	0.651
σ_v	0.613	0.873	σ_v	0.597	0.852

fined as 1-minute intervals where the absolute midquote return is larger than its 99.5-th sample quantile. Moreover, as we are interested in periods with transitory pricing errors (that is, flash crashes/rallies), we only focus on transitory EPMs where the midquote reverts more than two thirds within the following 10 minutes. Results from the EPM analysis are largely consistent with what we find in the state space analysis: during EPMs, dealers on average trade aggressively in the direction of EPMs while HFTs trade passively against them.

4 Price discovery roles of dealers and HFTs

In the previous two sections, we examine the liquidity provision roles of dealers and HFTs in the FX spot market, both in the traditional sense of order-book liquidity provision and in a general sense of trading against transitory pricing errors. Now we turn to their roles in price discovery, another important dimension of market functioning. We try to answer the following two questions: first, is it dealers or HFTs that contribute a larger share to price discovery? Second, do they contribute to price discovery in the same or a different manner? We start by assessing the informativeness of their aggressive trades using a simple price impact measure and then estimating a more robust structural VAR model to study the informativeness of *both* their aggressive trades and quote updates in the limit order-book.

4.1 Trade informativeness: Simple price impact measure

To measure the information content of a trade, we compute its relative price impact (RPI) defined as below:

$$\mathrm{RPI}_t = \frac{d_i\,(\mathrm{Mid}_{t+\Delta_t} - \mathrm{Mid}_t)}{\mathrm{Mid}_t} \tag{6}$$

where t indexes the transaction time of the trade. Mid_t represents the prevailing midquote *just before* the trade. $\mathrm{Mid}_{t+\Delta_t}$ represents the prevailing midquote a time period of Δ_t after the trade. Thus, in order to compute RPI, one has to choose a proper Δ_t. While a Δ_t that is too short leaves insufficient time for the market to fully learn the information contained in the trade and to reach the equilibrium price, a too-long Δ_t might capture other confounding information events unrelated to the trade. For robustness, we compute RPI over three different time horizons: 10 seconds, 30 seconds and 1-minute after the trade. d_i is the trade direction indicator, that is, $d_i = 1$ for buy trades and $d_i = -1$ for sell trades.

Table 5 reports, by trader category pair, the average and standard deviation of daily RPI across our sample days. Note that we first compute RPI trade by trade and then aggregate it to the daily frequency by computing its volume-weighted average. The results show that, given the same passive side, the price impact is the largest when dealers are on the aggressive side. In addition, out of all reported trader category pairs, the price impact is the largest when dealers aggressively

Table 5. Simple price impact measures. This table reports, by trader category pair, the summary statistics of the relative price impact (*RPI*) computed over three different time horizons of 10 seconds, 30 seconds and 1-minute. The trader category appearing first (second) in a pair is on the passive (aggressive) side of the trade. For example, HFT → Dealer represent trades where dealers aggressively take HFTs' quotes. We first compute RPIs trade by trade and then aggregate them at the daily frequency by computing their volume-weighted averages. Means and standard deviations are computed across daily measures.

		RPI-10 Sec (bp)		RPI-30 Sec (bp)		RPI-60 Sec (bp)	
		Mean	SD	Mean	SD	Mean	SD
Currency Pair	Category Pair						
GBP/USD	Dealer→Dealer	0.49	0.15	0.52	0.20	0.53	0.27
	Dealer→HFT	0.41	0.08	0.41	0.10	0.41	0.12
	Dealer→CB	0.40	0.20	0.43	0.31	0.44	0.41
	HFT→Dealer	0.69	0.22	0.69	0.23	0.66	0.26
	HFT→HFT	0.51	0.17	0.51	0.21	0.52	0.26
	HFT→CB	0.50	0.26	0.48	0.34	0.47	0.42
	CB→Dealer	0.56	0.24	0.59	0.35	0.59	0.40
	CB→HFT	0.49	0.14	0.51	0.19	0.52	0.24
	CB→CB	0.47	0.36	0.48	0.58	0.47	0.71
AUD/USD	Dealer→Dealer	0.66	0.22	0.69	0.24	0.70	0.29
	Dealer→HFT	0.59	0.11	0.61	0.13	0.61	0.16
	Dealer→CB	0.59	0.23	0.61	0.29	0.61	0.36
	HFT→Dealer	0.84	0.26	0.84	0.31	0.83	0.33
	HFT→HFT	0.66	0.19	0.66	0.21	0.67	0.25
	HFT→CB	0.68	0.26	0.69	0.34	0.67	0.40
	CB→Dealer	0.76	0.32	0.79	0.41	0.80	0.58
	CB→HFT	0.66	0.15	0.68	0.22	0.68	0.26
	CB→CB	0.66	0.31	0.69	0.40	0.70	0.54

take HFTs' quotes (that is, the row of HFT→Dealer). The price impact results above indicate that dealers' trades are most informed. As trades are considered as reflecting private information, it suggests that dealers are most privately informed.

While the RPI measure is simple and intuitive, it does not completely rule out the possibility that we might capture some transitory price impacts, given the trade-off on the choice of Δ_t as mentioned above. In addition, it can not speak to the information content of traders' other order-book activities such as submitting new limit orders or cancellations. Thus we turn to a structural VAR model in the next section for a more rigorous treatment.

4.2 Trade and quote informativeness: A structural VAR approach

To study the price discovery roles of dealers and HFTs, we estimate a structural vector autoregressive (structural VAR) model with midquote return and a series of trade and order variables (Hasbrouck, 1991; Brogaard, Hendershott, and Riordan, 2019; Fleming, Mizrach, and Nguyen, 2018). Each trade/order variable indicates a specific limit-order-book activity performed by one trader category, for example, HFTs improving the current best bid or ask (that is, increasing the bid or decreasing the ask). Benefiting from the highly precise timestamps, our SVAR model can accurately characterise the interactions between trade and order variables. Thus the informativeness of each trade and order variable can be estimated by the cumulative impulse responses of the midquote return to a unit shock in the variable. In addition, the SVAR model allows for a further decomposition of the efficient price innovation variance and thus can assign an information share to each trade and order variable, representing its contribution to price discovery. Thus, we are able to examine whether it is dealers or HFTs that contribute more to price discovery and whether this originates through their trades or quote updates.

Structural VAR specification A general structural VAR model can be specified as follows:

$$Ay_t = \alpha + \Phi_1 y_{t-1} + \cdots + \Phi_p y_{t-p} + \varepsilon_t \tag{7}$$

where $\Phi_1 \ldots \Phi_p$ are standard VAR system matrices capturing the lead-lag relations between the endogenous variables. ε_t is the vector of structural innovations and satisfies the following conditions: $\mathrm{E}(\varepsilon_t) = 0$; $\mathrm{E}(\varepsilon_t \varepsilon_t') = \Sigma_\varepsilon$; $\mathrm{E}(\varepsilon_t \varepsilon_s') = 0$ for $s \neq t$.

Following Hasbrouck (1991) and its extensions by Fleming, Mizrach, and Nguyen (2018) and

Brogaard, Hendershott, and Riordan (2019), we specify our structural VAR model as below:

$$y_t = \begin{pmatrix} r_t & x_t' \end{pmatrix}', \qquad A_0 = \begin{pmatrix} 1 & -a_{0,1} & -a_{0,2} & \cdots & -a_{0,k} \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \cdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix}. \tag{8}$$

The endogenous vector y_t consists of r_t, the log midquote return, and x_t, a vector of trade and order variables, for example, a new limit order adding depth at the BBO by HFT. In addition, as is commonly assumed in the microstructure literature, trade and order variables are assumed to affect returns contemporaneously but not vice versa.

After we have estimated the structural VAR model, we can easily obtain the vector moving average (VMA) representation to compute the impulse responses of return and trade variables to shocks in the structural innovations:

$$y_t = \Theta(L)\varepsilon_t = \Theta_0\varepsilon_t + \Theta_1\varepsilon_{t-1} + \Theta_2\varepsilon_{t-2} + \cdots \tag{9}$$

where $\Theta(L)$ is the polynomial of the lag operator $\Theta(L) = \Theta_0 + \Theta_1 L + \Theta_2 L^2 + \cdots$. The permanent price impact (PPI) of a trade/order variable k is defined as the cumulative impulse responses of the midquote return to a unit shock in the trade/order variable, that is,

$$\text{PPI}_k = \frac{\sum_{j=0}^{\infty} \partial r_{t+j}}{\partial \varepsilon_{k,t}} = [\Theta(1)]_{1,k} \tag{10}$$

where $[\Theta(1)]_{1,k}$ denotes the $(1,k)$-*th* element of $\Theta(1)$, the impulse response of the midquote return to trade/order variable k.

In addition to permanent price impact, we can compute the so-called 'information shares' of the trade/order variables via the approach of random walk decomposition (See Hasbrouck, 1991, for detailed proofs). The information share measure weighs the permanent price impact of a trade/order variable $[\Theta(1)]_{1,k}$ by its own structural innovation variance, $\sigma^2_{\varepsilon_k}$. So given two trade/order variables having the same permanent price impact, the information share of the one which arrives at the market more frequently will be larger. Mathematically, the information share (IS) of trade/order variable k to price discovery is computed as:

$$\text{IS}_k = \frac{[\Theta(1)]_{1,k}^2 \sigma_{\varepsilon_k}^2}{\sum_k [\Theta(1)]_{1,k}^2 \sigma_{\varepsilon_k}^2} \tag{11}$$

Structural VAR implementation details Following Brogaard, Hendershott, and Riordan (2019), we implement the structural VAR estimate as below. First, we utilise event time, which measures time as a sequence of order-book events; second, estimation is conducted separately for each date with statistical inferences calculated from the set of daily estimates; third, the number of lags in the VAR representation P is chosen to be five and impulse responses are truncated at 50 lags; fourth, following Brogaard, Hendershott, and Riordan (2019), all trader and order variables are binary. In other words, we only use the direction of trades and orders, not their sizes.

Construct trade and order variables Each trade and order variable is characterised by both the trader category responsible for initiating it: *HFT*, *Dealer* and *Other*[29], and the trade/order type (to be defined below): *Trade*, *BBO Improve Limit*, *BBO Worsen Cancel*, *BBO-Depth Add Limit*, *BBO-Depth Remove Cancel*, *Non-BBO-Depth Add Limit*, *Non-BBO-Depth Remove Cancel*. So we have in total 21 (3×7) trade and order variables and each corresponds to a specific order-book action performed by a trader category. Note that each trade and order variable is directional and its sign depends on whether it is a buy or sell (for trades) and whether it changes the bid or the ask side of the order book (for limit and cancel orders). Below we define each trade/order type in detail, with their signs in brackets.

- *Trade*: Market(able) buy orders (+1) or sell orders (-1) resulting in trades
- *BBO Improve Limit*: Limit orders increasing the best bid (+1) or decreasing the best ask (-1)
- *BBO Worsen Cancel*: Cancel orders decreasing the best bid (-1) or increasing the best ask (+1)
- *BBO-Depth Add Limit*: Limit orders adding depth at the current best bid (+1) or at the best ask (-1)

[29] As with the state space analysis, we merge all other traders other than dealers and HFTs into the Other category.

Table 6. Relative frequency of order book messages by trader category and order type. This table reports the relative frequency of all trade and order variables. *Trade* refers to trade executions. *BBO Improve Limit* refers to new limit orders which improve BBO (that is, lower ask or high bid), *BBO Worsening Cancel* refers to cancel orders which worsen BBO (that is, higher ask or lower bid). *BBO-Depth Add Limit* refers to new limit orders that add depth at the BBO and *BBO-Depth Cancel* refers to cancel orders that remove depth at the BBO. *Non-BBO-Depth Add Limit* refers to new limit orders that add depth at other price levels worse than BBO and *Non-BBO-Depth Remove Cancel* refers to cancel orders that remove depth at other price levels worse than BBO. All relative frequencies are averages across all sample days and in percentages.

	Trader Category	Dealer	HFT	Other	Sum
Curency	Order Type				
GBP/USD	Trade	0.62	1.18	0.35	2.15
	BBO Improve Limit	1.14	2.51	0.61	4.26
	BBO Worsen Cancel	0.63	1.95	0.35	2.94
	BBO-Depth Add Limit	3.21	9.73	1.16	14.10
	BBO-Depth Remove Cancel	2.00	9.64	0.74	12.38
	Non-BBO-Depth Add Limit	2.39	13.35	15.65	31.38
	Non-BBO-Depth Remove Cancel	3.03	13.54	16.21	32.79
	Sum	13.02	51.90	35.08	100.00
AUD/USD	Trade	0.89	1.47	0.54	2.91
	BBO Improve Limit	0.87	1.81	0.48	3.16
	BBO Worsen Cancel	0.29	1.15	0.16	1.60
	BBO-Depth Add Limit	3.10	15.54	1.61	20.25
	BBO-Depth Remove Cancel	1.74	15.20	1.00	17.95
	Non-BBO-Depth Add Limit	2.50	12.17	11.80	26.46
	Non-BBO-Depth Remove Cancel	3.19	12.15	12.33	27.67
	Sum	12.59	59.49	27.92	100.00

- *BBO-Depth Remove Cancel*: Cancel orders removing depth at the current best bid (-1) or at the best ask (+1)
- *Non-BBO-Depth Add Limit*: Limit orders adding depth at price levels lower than the current best bid (+1) or at price levels higher than the best ask (-1)
- *Non-BBO-Depth Remove Cancel*: Cancel orders removing depth at price levels lower than the current best bid (-1) or at price levels higher than the best ask (+1)

In Table 6 we present the relative frequencies of all trade and order variables defined above. There are several notable observations. First, HFTs account for the largest share of all order-book messages, about 52% for GBP/USD and 60% for AUD/USD. Second, HFTs are more active than dealers at the top of the order book, either by improving/worsening their BBO or by adding/removing the depth at the BBO. It is perhaps not surprising as the summary statistics reported in Table 2 show that on average HFTs quote a narrower bid-ask spread and supply more depth at the BBO.

Trade and order interactions between dealers and HFTs Before turning to permanent price impacts of trade and order variables, we first look at their mutual cumulative impulse responses reported in Table 7. They reveal in detail how dealers and HFTs respond to each other's order-book activities. The bottom-line is that HFTs are highly responsive to dealers' order-book activities but dealers are not responsive to HFTs'.[30]

We now provide economic interpretations of the key results in Table 7 for GBP/USD. Results for AUD/USD are very similar. First, a 1 unit innovation in dealers' *Trade* leads to 0.36 and 0.55 units of expected response in HFTs' *BBO-Depth Add Limit* and *BBO-Depth Cancel Limit* order respectively. It shows that HFTs, on average, respond to dealers' trades by both adding and removing depth at the BBO, in the same direction of the trade. For example, after a buy trade from dealers (of direction +1), HFTs are expected to cancel their depth on the ask (of direction + 1) and add depth on the bid (of direction +1). Second, a 1 unit innovation in dealers' *BBO Improve Limit* order leads to -0.42 units of expected response in HFTs' *Trade*. This means that HFTs are expected to respond to dealers' BBO improve limit orders by trading in the opposite direction. For example, when dealers set a new best bid, HFTs are typically expected to send a market sell order to take it. The above trade and order impulse responses provide insights into typical HFT strategies: first, responding to an incoming trade, HFT market makers typically adjust their top-of-book depth to mitigate adverse selection; second, HFT arbitrageurs or HFT market makers who use aggressive orders for inventory management are expected to closely monitor the order book dynamics and take liquidity when the bid-ask spread is improved (tightened) by dealers.

Permanent price impact of dealers' and HFTs' trades and quotes Table 8 shows the cumulative impulse responses of the midquote return to all trade and order variables, which is a proxy for

[30]Note that we mark the cumulative impulse responses larger than 0.2 in bold face and all of them appear in the diagonal blocks (within-trader-category responses, that is, dealers' responses to dealers and HFTs' responses to HFTs) and in the lower-left block (HFTs' responses to dealers).

Table 7. Trade and order impulse responses. This table reports the trade impulse responses (that is, permanent price impact of trade and order variables), cumulative over 50 periods and in basis points. *Trade* refers to trade executions. *BBO Improve Limit* refers to new limit orders which improve BBO (that is, lower ask or high bid), *BBO Worsening Cancel* refers to cancel orders which worsen BBO (that is, higher ask or lower bid). *BBO-Depth Add Limit* refers to new limit orders that add depth at the BBO and *BBO-Depth Cancel* refers to cancel orders that remove depth at the BBO. *Non-BBO-Depth Add Limit* refers to new limit orders that add depth at other price levels worse than BBO and *Non-BBO-Depth Remove Cancel* refers to cancel orders that remove depth at other price levels worse than BBO. Impulse responses are cumulated over 50 event messages. Cumulative impulse responses larger than 0.2 in absolute terms are marked in bold face.

			Dealer					HFT				
			Trade	BBO Improve Limit	BBO Worsen Cancel	BBO Depth Add Limit	BBO Depth Remove Cancel	Trade	BBO Improve Limit	BBO Worsen Cancel	BBO Depth Add Limit	BBO Depth Remove Cancel
GBP-USD	Dealer	Trade	1.09	-0.07	-0.03	0.01	-0.02	0.01	-0.01	-0.01	0.01	0.00
		BBO Improve Limit	0.00	1.29	**-0.37**	-0.01	-0.08	-0.18	-0.01	-0.02	-0.01	0.00
		BBO Worsen Cancel	0.00	**-0.22**	1.19	0.01	0.05	-0.01	0.03	-0.01	0.03	0.03
		BBO-Depth Add Limit	0.03	0.03	**-0.47**	1.24	**-0.51**	0.02	0.03	-0.01	0.05	0.02
		BBO-Depth Remove Cancel	0.01	0.05	0.11	**-0.38**	1.29	0.00	0.07	0.03	0.00	0.02
	HFT	Trade	-0.05	**-0.42**	-0.04	0.04	0.02	1.33	0.06	-0.04	0.08	0.00
		BBO Improve Limit	0.10	-0.06	-0.13	0.06	0.02	0.07	1.04	**-0.27**	0.11	0.08
		BBO Worsen Cancel	0.05	0.10	0.00	0.08	0.01	0.04	**-0.21**	1.05	0.06	0.08
		BBO-Depth Add Limit	**0.36**	0.05	-0.09	**0.20**	0.08	0.10	**0.46**	0.03	1.36	0.07
		BBO-Depth Remove Cancel	**0.55**	**0.60**	**0.25**	0.11	0.09	**0.20**	**0.34**	**0.32**	-0.08	1.37
AUD-USD	Dealer	Trade	1.12	-0.11	-0.05	0.02	-0.02	0.00	-0.04	-0.02	0.02	0.00
		BBO Improve Limit	-0.01	1.33	**-0.31**	-0.01	-0.05	-0.18	-0.01	-0.01	-0.01	0.00
		BBO Worsen Cancel	0.01	-0.11	1.15	0.01	0.02	0.00	0.01	-0.01	0.01	0.01
		BBO-Depth Add Limit	0.02	0.04	**-0.46**	1.12	**-0.29**	0.00	0.02	0.00	0.03	0.02
		BBO-Depth Remove Cancel	-0.01	0.05	0.05	**-0.21**	1.12	0.01	0.06	-0.01	0.01	0.02
	HFT	Trade	-0.02	**-0.55**	-0.06	0.05	0.04	1.37	0.01	-0.05	0.09	0.01
		BBO Improve Limit	0.04	-0.04	-0.12	0.04	0.03	0.07	1.01	**-0.26**	0.05	0.08
		BBO Worsen Cancel	0.03	0.04	-0.01	0.05	0.01	0.00	-0.15	1.03	0.04	0.05
		BBO-Depth Add Limit	**0.59**	0.08	0.08	**0.25**	0.16	0.06	**0.66**	**0.22**	1.63	**0.30**
		BBO-Depth Remove Cancel	**0.66**	**1.04**	**0.38**	**0.26**	0.11	**0.43**	**0.70**	**0.21**	-0.05	1.76

Table 8. Return impulse responses. This table reports the cumulative impulse responses of midquote returns to all trade and order variables as specified in Equation 10. All cumulative impulse responses are computed over 50 events and are in basis points. *Trade* refers to trade executions. *BBO Improve Limit* refers to new limit orders which improve BBO (that is, lower ask or high bid), *BBO Worsening Cancel* refers to cancel orders which worsen BBO (that is, higher ask or lower bid). *BBO-Depth Add Limit* refers to new limit orders that add depth at the BBO and *BBO-Depth Cancel* refers to cancel orders that remove depth at the BBO. *Non-BBO-Depth Add Limit* refers to new limit orders that add depth at other price levels worse than BBO and *Non-BBO-Depth Remove Cancel* refers to cancel orders that remove depth at other price levels worse than BBO. Column "Dealer - HFT" reports the difference between dealers and HFTs. * indicates the difference is statistically significant at 1% significance level.

		Dealer	HFT	Other	Dealer - HFT
GBP-USD	Trade	0.31	0.25	0.27	0.06*
	BBO Improve Limit	0.23	0.28	0.24	−0.05*
	BBO Worsen Cancel	0.19	0.22	0.21	−0.03*
	BBO-Depth Add Limit	0.07	0.09	0.06	−0.02*
	BBO-Depth Remove Cancel	0.01	0.06	0.02	−0.05*
	Non-BBO-Depth Add Limit	0.02	0.01	0.00	0.01
	Non-BBO-Depth Remove Cancel	0.00	0.05	0.00	−0.05*
AUD-USD	Trade	0.42	0.34	0.35	0.08*
	BBO Improve Limit	0.35	0.45	0.35	−0.10*
	BBO Worsen Cancel	0.34	0.38	0.38	−0.04*
	BBO-Depth Add Limit	0.09	0.09	0.07	0.00
	BBO-Depth Remove Cancel	0.02	0.09	0.03	−0.07*
	Non-BBO-Depth Add Limit	0.01	0.00	0.01	0.01
	Non-BBO-Depth Remove Cancel	0.00	0.06	0.00	−0.06*

their permanent price impacts, or informativeness. First, consistent with the findings in Brogaard, Hendershott, and Riordan (2019), we find *Trades*, *BBO Improve Limit* and *BBO Worsen Cancel* have a much larger permanent price impact than other order variables. Second, in order to compare the informativeness of dealers and HFTs, we compute the difference between them for each trade and order variable, which we report in Column "Dealer - HFT". While dealers' trades have a larger permanent price impact than HFTs', most HFTs' order variables have a larger price impact than dealers'. In addition, while dealers' trades have a larger price impact than their quotes, HFTs' quotes are more informative than their trades. This contrasts with results found in the equities market where Brogaard, Hendershott, and Riordan (2019) show that both HFTs' trades and BBO improving limit orders have a larger permanent price impact than non-HFTs'. While the difference may arise from their trader classification, it may also point to differences in FX markets. In their paper, traders are classified into two groups: HFTs and non-HFTs and the latter can include

Table 9. Information shares. This table reports the information shares of all trade and order variables as specified in Equation 11. *Trade* refers to trade executions. *BBO Improve Limit* refers to new limit orders which improve BBO (that is, lower ask or high bid), *BBO Worsening Cancel* refers to cancel orders which worsen BBO (that is, higher ask or lower bid). *BBO-Depth Add Limit* refers to new limit orders that add depth at the BBO and *BBO-Depth Cancel* refers to cancel orders that remove depth at the BBO. *Non-BBO-Depth Add Limit* refers to new limit orders that add depth at other price levels worse than BBO and *Non-BBO-Depth Remove Cancel* refers to cancel orders that remove depth at other price levels worse than BBO.

		Dealer	HFT	Other	Sum
GBP/USD	Trade	9.05	10.98	3.48	23.51
	BBO Improve Limit	8.73	25.85	3.75	38.33
	BBO Worsen Cancel	3.02	11.89	1.26	16.17
	BBO-Depth Add Limit	1.93	9.22	0.55	11.70
	BBO-Depth Remove Cancel	0.05	4.99	0.05	5.09
	Non-BBO-Depth Add Limit	0.09	0.47	0.03	0.59
	Non-BBO-Depth Remove Cancel	0.03	4.54	0.04	4.61
	Sum	22.90	67.94	9.16	100.00
AUD/USD	Trade	11.80	12.55	4.79	29.14
	BBO Improve Limit	7.66	24.01	4.42	36.09
	BBO Worsen Cancel	2.22	10.30	1.54	14.06
	BBO-Depth Add Limit	1.51	8.04	0.53	10.08
	BBO-Depth Remove Cancel	0.07	6.65	0.09	6.81
	Non-BBO-Depth Add Limit	0.03	0.15	0.04	0.22
	Non-BBO-Depth Remove Cancel	0.03	3.54	0.02	3.59
	Sum	23.32	65.24	11.43	100.00

uninformed liquidity traders.

Information shares of dealers' and HFTs' trades and quotes For a trade or order variable, we can compute its information share (Equation 11) instead of its permanent price impact (Equation 10). As explained above, the information share of a trade/order variable takes into account not only its permanent price impact, but also its contribution to the return structural variance. For example, with two trade/order variables having the same permanent price impact, the one that occurs more frequently is more likely to contribute a larger share to the return structural variance and thus a larger information share. Table 9 reports the information shares of all trade and order variables. It turns out that HFTs' price-changing limit and cancel orders contribute nearly half of overall price discovery. Perhaps it is not surprising as HFTs are much more active in their quote updates compared with dealers.[31]

To summarise, the following results reveal that dealers and HFTs contribute to price discovery in distinct patterns. While HFTs contribute more to price discovery through their active quote updates, incorporating public information, dealers contribute more through their trades, incorporating private information. Such a difference might arise due to the unique two-tiered structure of the FX market. While HFTs are mostly active on primary inter-dealer platforms such as Refinitiv, they are largely absent in the dealer-to-client segment. As a result, HFTs rely more on public information sources such as public data feeds from trading platforms of the security in question, or closely related assets such as FX futures contracts.

5 Conclusion

In this paper we examine the liquidity provision and price discovery roles of dealers and HFTs in the FX spot market. We find that HFTs, on average, provide more order-book liquidity than dealers: they quote a narrower bid-ask spread and supply more depth at the top of the order book. In addition, we find HFTs' order-book liquidity provision is less sensitive to market-wide volatility spikes. In contrast, dealers' order-book liquidity provision is more resilient ahead of discrete single-security volatility, in the form of scheduled macroeconomic news announcements. The above results suggest that dealers and HFTs have different strengths at market making in the FX market: dealers are likely able to derive more information from their large OTC client network or draw on their higher expertise, making them more able to manage discrete private information-asymmetry periods. HFTs, with faster and more advanced market-making algorithms, can better navigate market-wide high-volatility periods with comparatively larger amounts of public information revelation. We further find that dealers' aggressive trade flows create pricing errors which are absorbed by HFTs' passive trade flows. In other words, HFTs further contribute to FX inter-dealer market liquidity by meeting dealers' liquidity demands. We caution that we only focus on dealers and HFTs' liquidity provision during adverse, but less extreme market conditions. Whether

[31] Is is worth noting that it is not contradictory to previous results that show trades have a larger price impact. The information share of a trade or order variable is essentially its price impact, weighted by its innovation variance. Thus given the same price impact, a more active order type is expected to have a higher weight and thus higher information share.

the above findings hold during more extreme stress periods that are not in our sample, such as the 2020 COVID pandemic, is a question for future researchers. For example, anecdotal evidence from the 2015 Swiss franc 'de-peg' event shows that facing extreme volatility, HFTs pulled out of the market altogether. On price discovery, we find HFTs contribute the majority share, as a result of their top-of-book quote updates. Consistent with this, HFTs' quotes are more informative than their trades while we find the opposite for dealers. The above results suggest that dealers and HFTs contribute to price discovery in distinct ways.

In this paper we characterise how HFTs and dealers perform different functions in the FX market. We show that HFTs, while being relatively new to the FX inter-dealer market, are important providers of liquidity that is, mostly, resilient and through their quote updates, provide a majority share of price discovery. It is worth noting that in this study we focus only on a subset of the total FX market, ignoring platforms operated by dealers (single-dealer platforms) in the 'dealer-to-client' segment. As these venues comprise a much larger share of the total FX market than the inter-dealer venue used in this study, dealers are still the dominant provider of liquidity in the FX market as a whole.

Figure 6. Liquidity provision during the Swiss franc de-peg event. For each sub-figure, the top panel plots the midquote of the currency pair (left-axis) and passive volume by trader category (right-axis). The middle panel plots the market-wide relative quoted spread (left-axis) and top-of-book depth by trader category (right-axis). The bottom panel plots the top-five-level depth by trader category (right-axis). The dotted vertical line indicates the announcement time of the de-peg decision.

(a) GBP-USD

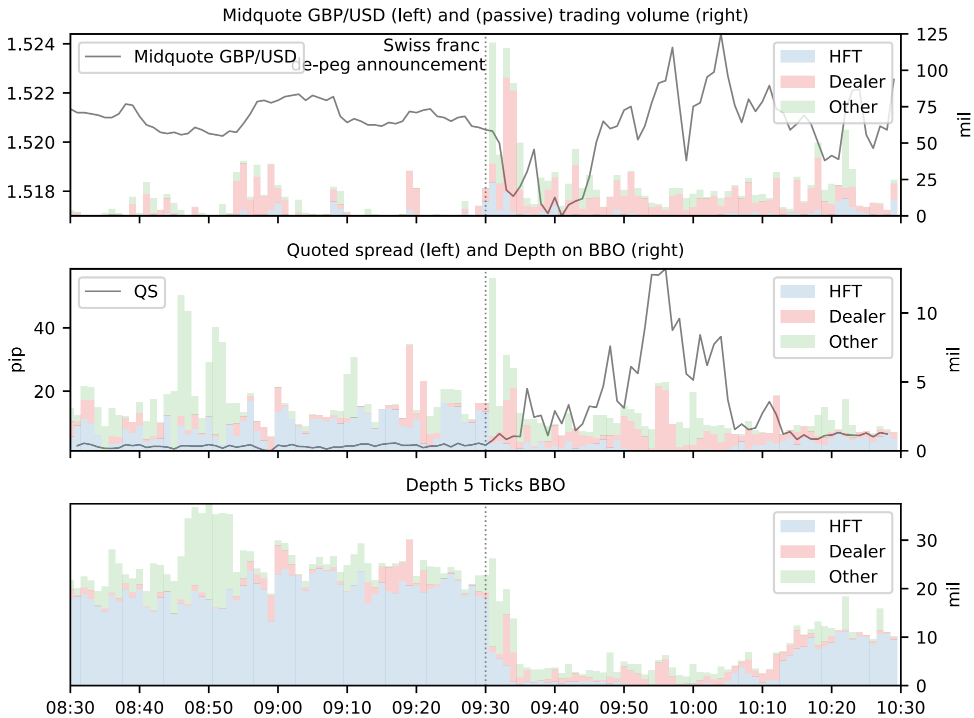

(b) AUD-USD

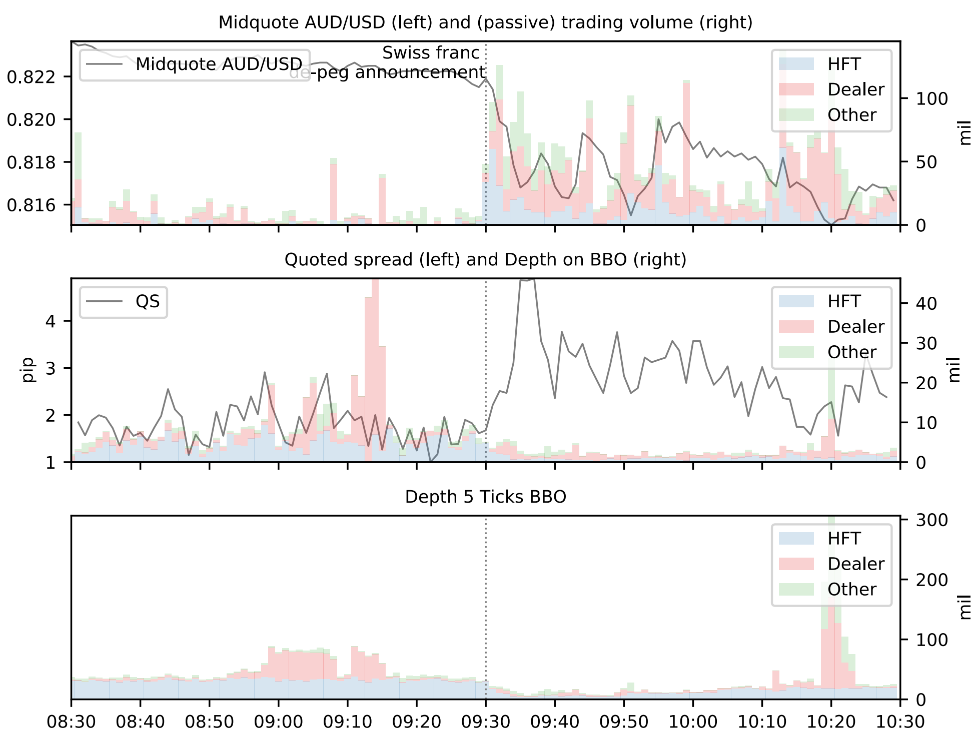

A Liquidity provision during extreme pricing errors

In Section 3.4, we estimate a state space model and show that HFTs on average trade against the pricing errors created by dealers' aggressive trade flows. However, it is not clear whether the same pattern remains during periods of extreme pricing errors (that is, flash crashes/rallies). To identify such periods, we follow a simple strategy as in Brogaard et al. (2018) and zoom in on periods of transitory and extreme price movements (EPMs). A transitory EPM defined over 1-minute frequency has to meet two conditions: first, the absolute midquote return over the 1-minute interval has to be above its 99.5-th quantile of all 1-minute intervals of our sample; second, the midquote has to revert more than two thirds within the next ten minutes. Based on the above algorithm, we find 244 (405) transitory EPMs for GBP/USD (AUD/USD).

Table A1. Summary statistics of EPMs. This table reports the summary statistics of several market condition variables during EPMs versus average 1-minute interval of the full sample. *AbsRet* is absolute return. *Vlm* is market trading volume. *TrdImb* is absolute trade imbalance (that is, buyer-initiated volume minus seller-initiated volume). *Vol* is the high-low-difference volatility (that is, difference between the highest and lowest midquote normalised by the average of the two.)

		EPMs			Full sample		
		Mean	SD	Q50	Mean	SD	Q50
GBP/USD	AbsRet (bp)	7.85	1.79	7.32	1.03	1.25	0.65
	Vlm (mil)	83.43	93.94	57.00	8.51	17.75	4.00
	TrdImb (mil)	21.01	29.54	11.00	3.48	6.38	2.00
	Vol (bp)	9.00	3.04	8.44	1.81	1.62	1.53
AUD/USD	AbsRet (bp)	9.44	2.16	8.77	1.19	1.45	0.67
	Vlm (mil)	91.15	93.31	64.00	10.33	21.00	4.00
	TrdImb (mil)	29.53	38.42	18.00	4.23	8.12	2.00
	Vol (bp)	10.44	3.76	9.74	2.08	1.85	1.67

Table A1 reports the summary statistics of several market condition variables during a transitory EPM versus an average 1-minute interval. It shows that we are indeed looking at extreme, tail events. Take GBP/USD for an example, the average absolute midquote return over a transitory EPM is 7.85 basis points, more than seven times larger than that over an average 1-minute interval. In addition, average trading volume during a transitory EPM is 83.43m, which is almost ten times higher than that over an average 1-minute interval.

Table A2. Trade flows around transitory EPMs (that is, mini flash crashes/rallies). This table reports, by trader category, their average (directional) aggressive ("Aggr"), passive ("Pass") and net ("Net") trade flows across all transitory EPMs (that is, mini flash crashes/rallies) during the 1-minute interval when an EPM happens and three minutes preceding and after the EPM. *, ** and *** indicate significance level at 10%, 5% and 1% respectively.

		Window	-3 Min	-2 Min	-1 Min	EPM	1 Min	2 Min	3 Min
GBP/USD	Dealer	Aggr	1.35***	0.99***	2.9***	15.91***	0.48	−2.3***	−1.46***
		Pass	−0.1	−0.41	−0.92	−2.1	2.28	−0.62	−0.5
		Net	1.25***	0.58	1.98***	13.8***	2.75***	−2.92***	−1.97***
	HFT	Aggr	−0.22	−0.11	0.04	−0.77	−2.01***	0.95**	0.98**
		Pass	−0.46***	−0.54***	−1.15***	−6.02***	0.08	0.86***	0.75***
		Net	−0.68	−0.65	−1.11*	−6.79***	−1.93***	1.81***	1.73***
	Commercial Banks	Aggr	−0.1	−0.16	0.3*	0.74	−0.54	0.1	−0.2
		Pass	−0.72***	−0.58***	−1.44**	−5.97***	0.41	0.67**	0.36*
		Net	−0.82***	−0.74***	−1.14*	−5.23***	−0.12	0.78	0.16
AUD/USD	Dealer	Aggr	1.35***	1.03***	3.9***	18.21***	0.37	−2.27***	−1.38***
		Pass	−0.04	−0.38	−2.45***	−5.33***	3.45***	1.4***	0.75*
		Net	1.32***	0.65	1.45***	12.87***	3.81***	−0.86	−0.63
	HFT	Aggr	0.11	−0.21	1.21***	−0.15	−2.49***	0.43	0.34
		Pass	−0.76***	−0.87***	−1.35***	−10.26***	−1.0***	0.45	0.8***
		Net	−0.65	−1.08***	−0.14	−10.42***	−3.49***	0.88*	1.14**
	Commercial Banks	Aggr	0.4**	0.35*	0.78***	5.83***	−0.71**	−0.43*	−0.45**
		Pass	−0.84***	−0.58**	−2.39***	−7.03***	0.53	0.49	0.0
		Net	−0.45	−0.23	−1.61***	−1.19	−0.18	0.05	−0.44

References

Aquilina, Matteo, Eric Budish, and Peter O'Neill (Feb. 2022). "Quantifying the High-Frequency Trading "Arms Race"". In: *The Quarterly Journal of Economics* 137.1, pp. 493–564.

Bank for International Settlements (2011). *High-frequency trading in the foreign exchange market*. Tech. rep. September, pp. 1–37.

— (2017). "The sterling 'flash event' of 7 October 2016". In: October 2016.

— (2018). *Monitoring of fast-paced electronic markets*. Tech. rep. September.

Bjønnes, Geir H. and Dagfinn Rime (2005). "Dealer behavior and trading systems in foreign exchange markets". In: *Journal of Financial Economics* 75.3, pp. 571–605.

Bjønnes, Geir H., Dagfinn Rime, and Haakon O.Aa Solheim (2005). "Liquidity provision in the overnight foreign exchange market". In: *Journal of International Money and Finance* 24.2, pp. 175–196.

Breedon, Francis, Louisa Chen, Angelo Ranaldo, and Nicholas Vause (2018a). "Judgement Day: Algorithmic Trading Around the Swiss Franc Cap Removal". In: *SSRN Journal*.

— (2018b). "Judgement Day: Algorithmic Trading Around the Swiss Franc Cap Removal". In: *SSRN Electronic Journal* 711.

Breedon, Francis and Angelo Ranaldo (2013). "Intraday patterns in FX returns and order flow". In: *Journal of Money, Credit and Banking* 45.5, pp. 953–965.

Breedon, Francis and Paolo Vitale (2010). "An empirical study of portfolio-balance and information effects of order flow on exchange rates". In: *Journal of International Money and Finance* 29.3, pp. 504–524.

Brogaard, Jonathan, Terrence J. Hendershott, and Ryan Riordan (2014). "High-frequency trading and price discovery". In: *Review of Financial Studies* 27.8, pp. 2267–2306.

— (2019). "Price Discovery without Trading: Evidence from Limit Orders". In: *Journal of Finance* 74.4, pp. 1621–1658.

Brogaard, Jonathan et al. (2018). "High frequency trading and extreme price movements". In: *Journal of Financial Economics* 128.2, pp. 253–265.

Budish, Eric B., Peter Cramton, and John J. Shim (2015). “The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response”. In: *The Quarterly Journal of Economics* 130.4, pp. 1547–1621.

Cerrato, Mario, Nicholas Sarantis, and Alex Saunders (2011). “An investigation of customer order flow in the foreign exchange market”. In: *Journal of Banking and Finance* 35.8, pp. 1892–1906.

Cespa, Giovanni, Antonio Gargano, Steven J Riddiough, and Lucio Sarno (2021). “Foreign Exchange Volume”. In: *The Review of Financial Studies* I, pp. 1–7.

Chaboud, Alain, Dagfinn Rime, and Vladyslav Sushko (2022). “The Foreign Exchange Market”. In: *SSRN Electronic Journal*, pp. 159–189.

Chaboud, Alain P., Benjamin Chiquoine, Erik Hjalmarsson, and Clara Vega (2014). “Rise of the machines: Algorithmic trading in the foreign exchange market”. In: *Journal of Finance* 69.5, pp. 2045–2084.

Collin-Dufresne, Pierre, Peter Hoffmann, and Sebastian Vogel (2019). “Informed traders and dealers in the FX forward market”. In:

Evans, Martin D.D. (2002). “FX trading and exchange rate dynamics”. In: *Journal of Finance* 57.6, pp. 2405–2447.

— (2010). “Order flows and the exchange rate disconnect puzzle”. In: *Journal of International Economics* 80.1, pp. 58–71.

Evans, Martin D.D. and Richard K. Lyons (2002). “Order flow and exchange rate dynamics”. In: *Journal of Political Economy* 110.1, pp. 170–180.

— (2005). “Do currency markets absorb news quickly?” In: *Journal of International Money and Finance* 24.2, pp. 197–217.

— (2006). “Understanding order flow”. In: *International Journal of Finance and Economics* 11.1, pp. 3–23.

— (2008). “How is macro news transmitted to exchange rates?” In: *Journal of Financial Economics* 88.1, pp. 26–50.

Evans, Martin D.D., Peter O’Neill, Dagfinn Rime, and Jo Saakvtine (2018). “Fixing the Fix? Assessing the Effectiveness of the 4pm Fix”. In: *FCA Occasional Paper No. 46*.

Fleming, Michael J., Bruce Mizrach, and Giang Nguyen (2018). "The microstructure of a U.S. Treasury ECN: The BrokerTec platform". In: *Journal of Financial Markets* 40, pp. 2–22.

Foucault, Thierry, Johan Hombert, and Ioanid Roşu (2016). "News Trading and Speed". In: *Journal of Finance* 71.1, pp. 335–382.

Foucault, Thierry, Roman Kozhan, and Wing Wah Tham (2017). "Toxic arbitrage". In: *Review of Financial Studies* 30.4, pp. 1053–1094.

Froot, Kenneth A. and Tarun Ramadorai (2005). "Currency returns, intrinsic value, and institutional-investor flows". In: *Journal of Finance* 60.3, pp. 1535–1566.

Glosten, Lawrence R. and Paul R. Milgrom (1985). "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders". In: *Journal of Financial Economics* 14.1, pp. 71–100.

Hagströmer, Björn and Lars Nordén (2013). "The diversity of high-frequency traders". In: *Journal of Financial Markets* 16.4, pp. 741–770.

Hasbrouck, Joel (1991). "The Summary Informativeness of Stock Trades: An Econometric Analysis". In: *Review of Financial Studies* 4.3, pp. 571–595.

Karnaukh, Nina, Angelo Ranaldo, and Paul Söderlind (2015). "Understanding FX Liquidity". In: *Review of Financial Studies* 28.11, pp. 3073–3108.

Kirilenko, Andrei, Albert S. Kyle, Mehrdad Samadi, and Tugkan Tuzun (2017). "The Flash Crash: High-Frequency Trading in an Electronic Market". In: *Journal of Finance* 72.3, pp. 967–998.

Korajczyk, Robert A. and Dermot Murphy (2019). "High-frequency market making to large institutional trades". In: *Review of Financial Studies* 32.3, pp. 1034–1067.

Kyle, Albert S. (1985). "Continuous Auctions and Insider Trading". In: *Econometrica* 53.6, p. 1315.

Love, Ryan and Richard Payne (2008). "Macroeconomic news, order flows, and exchange rates". In: *Journal of Financial and Quantitative Analysis* 43.2, pp. 467–488.

Malceniece, Laura, Kārlis Malcenieks, and Tālis J. Putniņš (2019). "High frequency trading and comovement in financial markets". In: *Journal of Financial Economics* 134.2, pp. 381–399.

Mancini, Loriano, Angelo Ranaldo, and Jan Wrampelmeyer (2013). "Liquidity in the foreign exchange market: Measurement, commonality, and risk premiums". In: *Journal of Finance* 68.5, pp. 1805–1841.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf (2016). “Information Flows in Foreign Exchange Markets: Dissecting Customer Currency Trades”. In: *Journal of Finance* 71.2, pp. 601–634.

Menkveld, Albert J. (2013). “High frequency trading and the new market makers”. In: *Journal of Financial Markets* 16.4, pp. 712–740.

— (2016). “The Economics of High-Frequency Trading: Taking Stock”. In: *Annual Review of Financial Economics* 8.1, pp. 1–24.

Menkveld, Albert J., Bart Z. Yueshen, and Haoxiang Zhu (2017). “Shades of darkness: A pecking order of trading venues”. In: *Journal of Financial Economics* 124.3, pp. 503–534.

Noss, Joseph, Lucas Pedace, Ondrej Tobek, Oliver B. Linton, and Liam Crowley-Reidy (2017). “The October 2016 Sterling Flash Episode: When Liquidity Disappeared from One of the World’s Most Liquid Markets”. In: *SSRN Electronic Journal* 687.

Osler, Carol L., Alexander Mende, and Lukas Menkhoff (2011). “Price discovery in currency markets”. In: *Journal of International Money and Finance* 30.8, pp. 1696–1718.

Payne, Richard (2003). “Informed trade in spot foreign exchange markets: An empirical investigation”. In: *Journal of International Economics* 61.2, pp. 307–329.

Ranaldo, Angelo and Paolo Santucci de Magistris (2019). “Trading Volume, Illiquidity and Commonalities in FX Markets”. In: *SSRN Electronic Journal*.

Ranaldo, Angelo and Fabricius Somogyi (2021). “Asymmetric information risk in FX markets”. In: *Journal of Financial Economics* 140.2, pp. 391–411.

Rime, Dagfinn, Lucio Sarno, and Elvira Sojli (2010). “Exchange rate forecasting, order flow and macroeconomic information”. In: *Journal of International Economics* 80.1, pp. 72–88.

SEC (2010). *Concept Release on Equity Market Structure*. 34.

Van Kervel, Vincent and Albert J. Menkveld (2019). “High-Frequency Trading around Large Institutional Orders”. In: *Journal of Finance* 74.3, pp. 1091–1137.