

**Finance and Economics Discussion Series  
Divisions of Research & Statistics and Monetary Affairs  
Federal Reserve Board, Washington, D.C.**

**The evolution of price discovery in an electronic market**

**Alain Chaboud, Erik Hjalmarsson, Filip Zikes**

**2020-051**

Please cite this paper as:

Chaboud, Alain, Erik Hjalmarsson, and Filip Zikes (2020). “The evolution of price discovery in an electronic market,” Finance and Economics Discussion Series 2020-051. Washington: Board of Governors of the Federal Reserve System, <https://doi.org/10.17016/FEDS.2020.051>.

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

# The evolution of price discovery in an electronic market

Alain Chaboud      Erik Hjalmarsson      Filip Zikes\*

June 5, 2020

## Abstract

We study the evolution of the price discovery process in the euro-dollar and dollar-yen currency pairs over a ten-year period on the EBS platform, a global trading venue used by both manual and automated traders. We find that the importance of market orders decreases sharply over that period, owing mainly to a decline in the information share from manual trading, while the information share of market orders from algorithmic and high-frequency traders remains fairly constant. At the same time, there is a substantial, but gradual, increase in the information share of limit orders. Price discovery also becomes faster, suggesting improvements in market efficiency. The results are consistent with theoretical predictions that in more efficient markets, informed traders tend to use more limit orders.

*JEL Classification:* F31; G14; G15.

*Keywords:* High-frequency trading; Limit orders; Price discovery.

---

\*Chaboud, [alain.p.chaboud@frb.gov](mailto:alain.p.chaboud@frb.gov), Division of International Finance, Federal Reserve Board; Hjalmarsson, [erik.hjalmarsson@economics.gu.se](mailto:erik.hjalmarsson@economics.gu.se), Department of Economics, University of Gothenburg; Zikes, [filip.zikes@frb.gov](mailto:filip.zikes@frb.gov), Division of Financial Stability, Federal Reserve Board. We are grateful Bohui Zhang (discussant) and the participants at the 14th Central Bank Workshop on the Microstructure of Financial markets for useful comments and suggestions. Avery Dao provided excellent research assistance. The views expressed in this paper are those of the authors and should not be interpreted as representing the views of the Federal Reserve Board or any other person associated with the Federal Reserve System. Hjalmarsson gratefully acknowledges financial support from the Swedish Research Council (Vetenskapsrådet) under grant 2014-01429.

# 1 Introduction

The past few decades have seen a major transformation in many important financial markets. Floor trading with dedicated market makers has given way to electronic limit order markets, where any trader can act as both market maker and market taker. The traditional classification of liquidity providers as “uninformed” and liquidity takers as “informed” has become less relevant in these markets. Price discovery may be driven by both liquidity makers and takers revealing information through their trading actions as in the theoretical models of Parlour (1998) and Roşu (2016). In addition, the advent of algorithmic and high-frequency trading has led to a setting where the vast majority of trades are now conducted with a computer trader on at least one side of the transaction. These changes have led to a need for research that helps us understand how the changing nature of the trading process has affected the way information is impounded into asset prices (see O’Hara, 2015, for a discussion).

In this study, we use a long time series of foreign exchange (FX) trading in the euro-dollar and dollar-yen currency pairs on the Electronic Broking Services (EBS) platform, with data spanning from 2008 to 2017. We use the standard empirical tool for evaluating price discovery—the Vector Autoregression (VAR) framework introduced by Hasbrouck (1991a,b)—but we extend it by studying the impact of several types of market orders and by also estimating the information content of limit orders. First, as the data allow us to observe trader types, we separate the contribution to price discovery from market orders of three different types of traders: manual traders (labeled as Manual), algorithmic traders in the employ of a bank (labeled as Bank-AT), and algorithmic traders in the employ of “non-banks” (labeled as HFT). Second, using both trade and quote data, we construct three types of limit order flow: price-improving orders, price-matching orders, and price-worsening orders, and study the contribution of these limit order flow variables to price discovery<sup>1</sup>. We then estimate the price discovery models separately for each month of the sample and report the evolution of the estimates over time. The long sample and our emphasis on the time-varying nature of price discovery affords a novel perspective in this literature, as previous studies rely on short samples that cannot reveal any long-term trends.

Our main results for the euro-dollar currency pair can be summarized as follows

---

<sup>1</sup>Recent empirical work, including Hautsch and Huang (2012), Cont, Kukanov, and Stoikov (2014), Fleming, Mizrach, and Nguyen (2017), and Brogaard, Hendershott, and Riordan (2018), propose similar ways of extending the traditional empirical model to allow for limit orders.

(the results for the dollar-yen are qualitatively the same but the quantitative estimates differ). First, the importance of market orders in price discovery decreases sharply over the sample period. In the beginning of the sample, in 2008, market orders explain upwards of 50 percent of the variation in the efficient price.<sup>2</sup> By the end of the sample period, in 2017, this number has fallen to around 20 percent. Second, we show that this steep drop in the importance of market orders can almost exclusively be attributed to orders submitted by Manual traders, as opposed to algorithmic and high-frequency traders (i.e., Bank-AT and HFT). In particular, the share of price discovery attributable to Manual market orders drop from around 30 percent to almost zero percent during the sample. Third, the price variation that can be attributed to limit orders increases from about 25 percent to around 50 percent from the beginning to the end of the sample. Finally, the overall speed of price discovery increases during the sample period, such that the transition to the new equilibrium price following an order event happens more quickly, consistent with the increased presence of computer-driven trading.<sup>3</sup> We evaluate the speed of price discovery by introducing a new measure in this context, labeled the  $\pi$ -life of the price impact function (Fanelli and Paruollo, 2010).<sup>4</sup> The  $\pi$ -life measures how quickly the price impact function converges to its long-run (permanent) value, and can thus be seen as a measure of how quickly the price reaches its new equilibrium after an order shock. More generally, the  $\pi$ -life can be viewed in this context as a proxy for market efficiency, as it essentially reflects how quickly new information is impounded into prices.

The empirical results strongly support the notion that in modern electronic markets both market orders and limit orders are used by informed traders, in the sense

---

<sup>2</sup>Formally, there are no pure “market” orders on EBS, as all orders contain a limit price. However, some orders are immediately marketable (they result in an immediate trade), and we refer to these as market orders.

<sup>3</sup>These results are generally in line with the literature that explicitly studies whether HFTs and other algorithmic traders contribute to price efficiency, and which finds that HFTs seem to push prices towards efficiency. See, for instance, Hendershott and Riordan (2013), Brogaard, Hendershott, and Riordan (2014), Chaboud, Chiquoine, Hjalmarsson, and Vega (2014), and Benos, Brugler, Hjalmarsson, and Zikes (2017).

<sup>4</sup>The  $\pi$ -life measure introduced by Fanelli and Paruollo (2010) can be viewed as a generalization of the commonly used concept of half-life. The half-life of a shock to a stationary variable captures how long it takes for half of the effect of that shock to die out. For non-stationary variables, such as price processes, the effect of a shock need not die out over time and there may be a long-run permanent effect; the  $\pi$ -life subsequently measures how long it takes before the effect of a shock gets within some pre-specified range of this permanent response.

that both types of orders contribute substantially to permanent price changes. This conclusion is in line with the studies by Fleming, Mizrach, and Nguyen (2017) and Brogaard, Hendershott, and Riordan (2018), who analyze price discovery in bond and equity markets, respectively. However, the results presented here allow us to study how the relative importance of market and limit orders in price discovery has evolved over time, with the evidence pointing to an increasing use of limit orders for informed trading over the period 2008 to 2017, while the importance of market orders decreases. Importantly, this decrease seems to be associated almost exclusively with market orders generated by only one type of trader, manual traders. These orders become both relatively less common—as a proportion of all market orders, they drop from about 65 percent to around 15 percent in the euro-dollar during the sample—and less informative, as the permanent price impact of a given-sized Manual trade decreases over the sample period. This evidence is consistent with the increasing adoption of computer-based algorithms by banks to execute the large orders of institutional investors.

Our study contributes to the growing empirical literature on price discovery in limit order markets, highlighting the apparent ongoing shift from informed liquidity taking to informed liquidity provision. The theoretical microstructure literature has long recognized that the view of liquidity providers as uninformed is likely overly simplistic, but the challenges to building models with strategic liquidity provision have proven significant.<sup>5</sup> In a recent theoretical study by Roşu (2016), it is shown that informed traders prefer limit orders when their information advantage is small. The increased speed of price discovery is consistent with improvements in market efficiency during our sample period, possibly as a result of the increase in algorithmic trading participation (see Footnote 3). This would suggest that private information, broadly interpreted, has become more difficult to obtain, and one might therefore expect an increased use of strategic liquidity provision.

From a practical perspective, our results have important implications for measuring the degree of information asymmetry in the FX market and for applications that rely on such measures. The standard approach in the literature has been to use Kyle- and Amihud-type measures (Kyle, 1985, Amihud, 2002), which are designed

---

<sup>5</sup>Contributions to the modelling of limit order markets include, among others, Glosten (1994), Chakravarty and Holden (1995), Parlour (1998), Foucault (1999), Foucault, Kadan, and Kandel (2005), Goettler, Parlour, and Rajan (2005, 2009), Kaniel and Liu (2006), Roşu, (2009, 2016), and Ricc , Rindi, and Seppi (2018).

to capture the price impact of informed market order flow. Our results imply that such measures likely have, over time, increasingly underestimated the degree of information asymmetry in the FX market, as informed traders have, over time, also increasingly relied on limit rather than market orders to impound information into exchange rates. Thus, the information content of both market and limit orders should be considered to obtain a comprehensive picture of information asymmetry in the FX market.

Moreover, our findings suggest that strategic liquidity provision might also be challenging in practice, as it appears that market participants have only gradually developed the skills necessary to implement the strategy. Previous studies have typically examined the response of market participants, and the resulting impact on market quality, of various structural changes, such as technological upgrades and trading protocols. While we also find that such changes have an immediate impact on the price discovery process in the FX market, the nature of price discovery has been gradually changing over the course of many years, even in the absence of such interventions. This “non-stationarity” of the price discovery process needs to be kept in mind when running various regression analyses, because it may easily lead to spurious inference.

During our sample period, the fraction of computer-driven algorithmic and high-frequency trading increased substantially, and manual traders became a minority on the EBS platform. In addition, several changes to the market rules occurred on EBS, including a minimum quote life (MQL) rule, a decimalization rule, a partial reversal of this decimalization rule, and a latency floor.<sup>6</sup> Although these changes to the market rules do *not* explain the long-run patterns described above, some still affected price discovery in important ways. In particular, changes in tick size had clear and significant effects. The “decimalization” policy implemented in March 2011 decreased the tick size by a factor of 10 and was subsequently partially reversed about a year and a half later, when the tick size was increased by a factor of 5. Chaboud, Dao, and Vega (2018) explicitly analyze the effects of these changes on various aspects of HFT participation and trading activities. As they point out, the most directly observable effect of decimalization was the increased relative *taking* by HFTs, whereas

---

<sup>6</sup>In accordance with EBS’s own nomenclature, as well as previous academic literature on this topic, the term “decimalization” here refers to the reform on EBS whereby the tick size was decreased by a factor of 10, not a switch from a “fractional” to a “decimal” system.

the making activities of HFTs remained fairly constant around decimalization. Our results provide an interesting complement to the findings in Chaboud, Dao, and Vega (2018), as we show that decimalization also had a clear and significant impact on how different types of limit orders contributed to price discovery.

The rest of the paper is organized as follows. In Section 2, we provide a brief overview of the institutional features of the EBS platform. In Section 3, we describe the data and introduce our methodology. In Section 4, we report our main analysis of time-varying price discovery in the EBS market and in Section 5 we discuss the effects of structural changes on the EBS platform on the price discovery process. Section 6 concludes.

## 2 The EBS platform

### 2.1 Brief overview

Trading activity in the foreign exchange market is spread across a large number of trading venues using a variety of technologies. But two electronic trading platforms, EBS Market and Thomson Reuters Matching, both central limit order books (CLOBs), are at the core of the global interdealer spot market, with prices from these two platforms widely viewed as the reference exchange rates at any moment of the day. The euro-dollar and dollar-yen currency pairs, the two currency pairs with the highest trading volume, trade primarily on EBS, and the price discovery process for these exchange rates is therefore concentrated on EBS.<sup>7</sup> In the analysis below, we present and discuss results for the euro-dollar, by far the most actively traded of all currency pairs. The corresponding results for the dollar-yen, the second-most traded currency pair, which are all qualitatively similar, are presented in the Online Appendix.

The EBS Market, as an “interdealer” system, is widely used by foreign exchange dealing banks around the globe. Since 2004, however, it has also been accessible to non-banks under prime-brokerage arrangements with some of the dealing banks. Therefore, a number of HFTs as well as a few large hedge funds and commodity trad-

---

<sup>7</sup>BIS (2018) confirm the view that while the main interdealer platforms (EBS and Thomson Reuters) are subject to strong competition, they remain central to price discovery in their respective primary currencies.

ing advisors (CTAs) also trade on the platform. Trading instructions can be entered on EBS on a specialized keyboard (the “Manual” traders), or through a computer interface. Importantly, the banks trading on EBS access the system both manually and through a computer interface, while, with very few exceptions, non-bank trading occurs almost exclusively via computer interface. The manual/computer and the bank/non-bank breakdown form the basis for the classification of EBS counterparties into the three types seen in our data: The manual traders, the bank algorithmic traders, and the non-banks, with the vast majority of that group’s activity coming from HFTs. Broadly speaking, these categories also represent the slow traders, the faster traders, and the fastest traders. We refer to these three trader groups as Manual, Bank-AT, and HFT.

## 2.2 Structural changes on EBS

Over our sample period, EBS implemented a number of important structural changes on its main trading platform. Many of these changes were designed to address the interaction between the different types of traders coexisting on EBS, in particular the balance between manual traders and HFTs. We briefly highlight four of these changes, as we will later discuss whether they affected the patterns of price discovery that we observe.

On June 15, 2009, EBS implemented a “minimum quote life” (MQL), also known as a minimum resting time. This prevented all traders from canceling limit orders they had placed in the CLOB before 250 milliseconds had passed since the initial submission of the order. The measure was reportedly introduced to address the concerns of some manual traders who complained they were having difficulties hitting some quotes before they disappeared from the order book (sometimes referred to as the “flickering quote” problem).

On March 7, 2011, EBS added an extra decimal to the precision of the quotes in its system, reducing the tick size by a factor of 10. For example, while the closest bid and ask quotes for euro-dollar could previously be, say, 1.2345 and 1.2346, respectively, after decimalization they could now be 1.23450 and 1.23451.<sup>8</sup> At the time, several other (smaller) foreign exchange trading platforms (mainly “bank-to-

---

<sup>8</sup>The last digit of the exchange rate before decimalization was known as a “pip.” After decimalization, the last digit was called a “decimal pip.” Finally, the increase in tick size in 2012 introduced the “half pip.”

customer” electronic platforms) had added an additional decimal to the quoted price, allowing for the possibility of smaller transaction costs. Before the decimalization in 2011, the minimum tick size on EBS was often binding, with the bid-ask spread at its minimum more than half of the time. After decimalization, the tick size was essentially never binding (Chaboud, Dao, and Vega, 2018).

On September 24, 2012, 18 months after the reduction in tick size, EBS partially reversed course and increased the tick size by a factor of 5. The last decimal remained, but it could now only be 0 or 5.

Finally, on March 3, 2014 and February 17, 2014, for the euro-dollar and dollar-yen, respectively, EBS introduced a “latency floor,” another measure likely designed to address the impact of fast traders on the platform. The latency floor imposes a small delay (randomly set at a few milliseconds) on incoming messages before they are released and incorporated in the CLOB. Importantly, during each of these short delays, the incoming messages are batched and, within each batch, their order is randomized before they are released to the CLOB. EBS explained at the time that the latency floor was designed to reduce the pure advantage of speed in the trading process, lowering the risk of a wasteful technological arms race.<sup>9</sup>

### 3 Data and methodology

#### 3.1 Data

We use quote and trade data, for the euro-dollar and dollar-yen exchange rates, from EBS for the period spanning January, 2008, to December, 2017. The quote data specify the best bid and ask prices, as well as the amounts or “depths” (in millions of the base currency)<sup>10</sup> available to trade at these prices. These are binding orders, and the quotes thus represent the true bid and ask prices in the market at a given

---

<sup>9</sup>The EBS platform operates continuously from early Monday morning in Australia and New Zealand through the end of the trading day on Friday afternoon in the United States. The platform is closed over the weekend. The market-rule changes were all implemented on Monday mornings, when the EBS system opens after its regular week-end closing. EBS customers were informed well in advance and given full details of the new measures.

<sup>10</sup>The euro-dollar exchange rate is quoted in dollars per euro, and the euro is thus the “base” currency. In contrast, the dollar-yen rate is expressed in yen per dollar, with the dollar as the base currency. The minimum trade size on EBS is 1 million of the base currency, and any market or limit order must be in multiples of millions of the base currency.

time. From January 2008 through August 2009, these data are recorded at a 250 millisecond (ms) frequency, which corresponds to the fastest quote update frequency available to market participants at the time. From August 2009 through the end of our sample, the sampling and broadcast frequency increases to every 100ms, again the highest price and order book update frequency available to market participants over that period. Although EBS is a 24-hour market, we focus our analysis on the most active hours between 3:00 am ET to 3:00 pm ET, Monday through Friday (as in Chaboud, Dao, and Vega, 2018), which roughly correspond to the busiest trading hours in London and New York and accounts for a large share of the daily trading volume.

In addition, we also have data on all completed trades during the sample period. For all executed trades, these data specify the transaction price and the traded amount. Importantly, the trade record also specifies the type of maker and taker in each trade (i.e., Manual, Bank-AT, or HFT), as well as whether the trade was a buy or a sell of the base currency from the perspective of the taker. The trades are time-stamped with millisecond precision throughout the sample period, where the time stamp indicates the exact time at which the order hits the order book. In terms of calculating the market and limit order flows, this time stamp convention is a significant improvement over using the time of the trade confirmation. As the taker and maker to a trade on EBS can be in different regions of the world, there could be a substantial lag in confirming the trade, potentially up to 150 milliseconds.<sup>11</sup> Such delays would particularly affect the correct calculation of limit order flows as it depends on an exact matching of incoming markets orders and changes in the order book, as seen below.

### **3.2 Participation of different trader groups**

The identification of the maker and taker type in each trade leads to 9 possible maker-taker combinations (i.e., Manual maker/Manual taker, Manual maker/Bank-AT taker, and so forth). Figure 1 shows the relative participation rates of these 9 different combinations in euro-dollar trading, measured as the fraction of transacted volume in a given month that can be attributed to a specific maker-taker combi-

---

<sup>11</sup>EBS operates interconnected matching engines in London, New York and Tokyo.

nation.<sup>12</sup> For instance, the line labeled Manual vs. Bank-AT (Panel B) shows the relative volume attributable to trades with a Manual maker and a Bank-AT taker, with the naming convention following the logic that the maker posts the quote before the taker hits it. As is seen, the most dramatic change over time occurs for the pure manual trades (Manual vs. Manual, Panel A), which drop in fraction from over 40 percent in 2008, to less than 5 percent in 2017. Thus, by the end of the sample, an algorithmic trader (i.e., Bank-AT or HFT) is involved on at least one side of a transaction in over 95 percent of the traded euro-dollar volume.

Panel A in Figure 2 shows the fraction of volume attributable to a given type of taker. The Manual line in Figure 2, Panel A, thus represents the sum of the three lines in Panel A in Figure 1 (and the Bank-AT and HFT lines represent the sums of the three lines in Panels B and C of Figure 1, respectively). Manual taking decreases dramatically, dropping from 65 to 15 percent during the sample. At the same time, HFT taking increases from 15 to 50 percent, while Bank-AT taking increases somewhat less, from about 20 to 35 percent. Panel B in Figure 2 shows the analogue results when grouping on the type of maker. For instance, the Manual line in Figure 2, Panel B, represents the sum of the Manual vs. Manual, Manual vs. Bank-AT, and Manual vs. HFT lines in Panels A, B, and C, respectively, in Figure 1. Manual making also drops substantially, from around 70 percent at the beginning of the sample to around 25 percent at the end. The biggest increase in making is seen for Bank-AT, which increase their making share from about 15 percent to almost 40 percent during the sample period. HFT making increases in the first part of the sample, but then drops somewhat until it starts to pick up again during the last two years of the sample.

### 3.3 Variable definitions

We use the above data to construct measures of returns and order flows. Let  $p_t^b$  and  $p_t^a$  be the top-of-the-book bid and ask prices at time  $t$ , respectively, and let  $q_t^b$  and  $q_t^a$  denote the associated depths available at these prices. We define the mid-quote return as  $r_t = \log p_t^m - \log p_{t-1}^m$ , where  $p_t^m$  is the mid-quote, i.e., the arithmetic average of the best bid and ask prices. Let  $v_t^b$  and  $v_t^s$  denote the buyer- and seller-initiated volume during the time interval  $[t-1, t]$ , respectively, and let  $v_t = v_t^b + v_t^s$  denote the total volume. We break down the volumes by trader type, and write  $v_t^{b,i}$  and

---

<sup>12</sup>Trading in dollar-yen exhibit similar participation patterns, as seen in the Online Appendix.

$v_t^{s,i}$ ,  $i \in \{\text{Manual, Bank-AT, HFT}\}$  for buyer- and seller-initiated volume by Manual traders, Bank-ATs, and HFTs. The market order flow generated by trader type  $i$  is defined by  $m_t^i = v_t^{b,i} - v_t^{s,i}$  and the total market order flow is given by  $m_t = v_t^b - v_t^s$ . The market order flow, as is standard, thus measures the net buy volume, as seen from the perspective of the active (taking) side of the market.

As discussed previously, our data show the state of the limit order book every 100ms (250ms in 2008 and 2009). By combining the observed changes in the limit order book, between each 100ms (250ms) interval, with the trade records for the market orders, we infer measures of the actual limit order activity that occurred at the top of the book during a given time period. We label these active changes in the order book as the “limit order flow”. As shown by Brogaard, Hendershott, and Riordan (2018), limit orders further out in the order book seem to have virtually no impact on price discovery, and we restrict attention to orders at the best bid and ask prices.

We follow a definition similar to that used in Cont, Kukanov, and Stoikov (2014), and specify the top-of-the-book bid limit order flow,  $l_t^b$  as

$$l_t^b = \begin{cases} q_t^b & \text{if } p_t^b > p_{t-1}^b, \\ q_t^b - q_{t-1}^b + v_t^s & \text{if } p_t^b = p_{t-1}^b, \\ -(q_{t-1}^b - v_t^s) & \text{if } p_t^b < p_{t-1}^b. \end{cases} \quad (1)$$

Here  $q_t^b$  is the quantity available to trade (i.e., depth) at the best bid price at the end of period  $t$ . If  $p_t^b > p_{t-1}^b$ , such that the best bid price increases from period  $t - 1$  to  $t$ , the time  $t$  bid limit order flow is simply equal to the amount posted at that new price. In this case the bid limit order is price improving, and it is natural to define the top-of-book bid limit order flow as equal to the depth at the new best price.

If the best bid price remains unchanged from period  $t - 1$  to  $t$ , the bid limit order flow is defined as the additional amount posted at that price, controlling for the changes in depth that occurred because of sell market orders.<sup>13</sup>

Finally, if the best bid price decreases from  $t - 1$  to  $t$ , the bid limit order flow

---

<sup>13</sup>For example, suppose the best bid price remains the same from period  $t - 1$  to  $t$ , the depth at the best bid price is 5 million euro at time  $t - 1$  and 4 million euro at time  $t$ , and that a market sell order of 2 million euro is executed between  $t - 1$  and  $t$ . In this case, a bid limit order of 1 million euro must also have been posted at the best price in this period, since otherwise the depth would have been only 3 million euro at time  $t$ . The 1 million euro limit order thus constitute the active change in the order book, after controlling for changes induced by market orders.

is defined as the negative of the depth available at time  $t - 1$  (at price  $p_{t-1}^b > p_t^b$ ), after controlling for the amount of any market sell orders between  $t - 1$  and  $t$ . This definition stems from the fact that we try to measure the active limit orders at the best available bid price during the period, which in this case is equal to  $p_{t-1}^b$ . Since the best bid price decreases to  $p_t^b < p_{t-1}^b$  by the end of period  $t$ , there must have been either an active withdrawal of the limit orders available at the old best price, and/or these limit orders must have been matched to incoming market sell orders during the period.<sup>14,15</sup>

The bid limit order flow can thus be seen as having three parts, the price-improving part,  $l_t^{b,1} := q_t^b 1\{p_t^b > p_{t-1}^b\}$ , the price-matching (quantity-changing only) part,  $l_t^{b,2} := (q_t^b - q_{t-1}^b + v_t^s) 1\{p_t^b = p_{t-1}^b\}$ , and the price-worsening part,  $l_t^{b,3} := -(q_{t-1}^b - v_t^s) 1\{p_t^b < p_{t-1}^b\}$ . We define analogous quantities for the ask side of the book and then define the net price-improving, price-matching, and price-worsening limit order flows by  $l_t^j = l_t^{b,j} - l_t^{a,j}$ ,  $j = 1, 2, 3$ , respectively. An overall bid limit order flow can be defined as  $l_t^b = l_t^{b,1} + l_t^{b,2} + l_t^{b,3}$ , and analogously, an overall ask limit order flow can be defined as  $l_t^a = l_t^{a,1} + l_t^{a,2} + l_t^{a,3}$ . The total limit order flow is then defined by  $l_t = l_t^b - l_t^a$ . However, given the rather different nature of the three components of the limit order flow, along with previous empirical results (e.g., Brogaard, Hendershott, and Riordan, 2018), we keep the three parts of the limit order flow separate in the empirical analysis.

### 3.4 Empirical methodology

The main econometric tool for empirically examining price discovery has long been the VAR framework introduced by Hasbrouck (1991a,b). In the standard price discovery VAR, prices are driven by trades (market orders) and the only two variables entering the VAR are returns and market order flow. The standard VAR model can be understood in the context of a designated market maker, who adjusts his quotes to reflect the expected information content of a given trade, along the lines of Kyle (1985). Such a setting is quite removed from modern limit order markets, where all

---

<sup>14</sup>To illustrate, suppose the depth at the best bid price  $p_{t-1}^b$  at time  $t - 1$  is equal to 3 million euro, and there is a market sell order for 1 million euro between time  $t - 1$  and  $t$ . At time  $t$ , the best bid price has decreased by 1 pip. Between time  $t - 1$  and  $t$ , there must therefore have been an active withdrawal of 2 million euro from the order book at the price  $p_{t-1}^b$ , and the price-worsening limit order flow is equal to  $-(q_{t-1}^b - v_t^s) = -(3 - 1) = -2$  million euro.

<sup>15</sup>The calculation of the price-worsening limit order flow does not explicitly ensure against positive outcomes and in the cases when it does end up positive, we cap the value at zero.

types of traders are free to use any kind of order and limit orders might very well be informed. Recent work by Fleming, Mizrach, and Nguyen (2017) and Brogaard, Hendershott, and Riordan (2018) extend this framework to also allow for a price impact of limit orders (earlier works by, for instance, Hautsch and Huang, 2012, and Cont, Kukanov, and Stoikov, 2014, reflect similar ideas). We follow this line of thought and extend the standard price discovery VAR introduced by Hasbrouck (1991a,b) to incorporate limit orders as well as market orders.

In particular, throughout the paper, we will be working in a standard structural VAR framework,

$$\mathbf{A}\mathbf{y}_t = \sum_{i=1}^p \mathbf{B}_i \mathbf{y}_{t-i} + \mathbf{D}^{1/2} \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim iid(\mathbf{0}, \mathbf{I}), \quad (2)$$

where  $\mathbf{y}_t = (r_t, \mathbf{x}_t)'$  is the vector of endogenous variables,  $\mathbf{A}$  is a matrix of structural parameters,  $\mathbf{B}_i$ ,  $i = 1, \dots, p$  are unrestricted lag coefficient matrices, and  $\mathbf{D}$  is a diagonal variance-covariance matrix. In the standard traditional price discovery VAR,  $\mathbf{x}_t$  is simply a scalar representing some measure of market order flow. In the current framework,  $\mathbf{x}_t$  is in general a vector that also includes measures of limit order activity. Our specification of the structural  $\mathbf{A}$  matrix in the general case follows the same scheme as in Brogaard, Hendershott, and Riordan (2018), and is detailed in the context of each model in Section 4.

The structural VAR in (2) is used to capture several different aspects of price discovery. Specifically, the information content of the market and limit order flow variables is measured by the so-called permanent price impact and information shares. The former captures the ultimate price impact of a given order book event, whereas the latter reflects the contributions of the various order flows to the variance of the permanent component of the price process. In addition, we also measure how quickly the information in market and limit orders is impounded into the price, using the long-run  $\pi$ -life introduced by Fanelli and Paruollo (2010). These three price discovery measures are defined in detail below.

Provided the process in (2) is stationary, it admits an infinite-order vector moving-average (MA) representation,

$$\mathbf{y}_t = \mathbf{C}_0 \boldsymbol{\varepsilon}_t + \mathbf{C}_1 \boldsymbol{\varepsilon}_{t-1} + \mathbf{C}_2 \boldsymbol{\varepsilon}_{t-2} + \dots \quad (3)$$

The MA terms  $\mathbf{C}_0, \mathbf{C}_1, \dots$ , are  $k \times k$  matrices, where  $k$  is the dimension of  $\mathbf{y}_t$ . Below, we denote the individual elements of matrices in non-bold lowercase letters, such that  $c_{l,ij}$  denotes the elements of  $\mathbf{C}_l$ . By the BN decomposition (Beveridge and Nelson, 1981), equation (3) can be re-stated as

$$\mathbf{y}_t = \tilde{\mathbf{C}}(1) \varepsilon_t + \eta_t, \quad (4)$$

where  $\tilde{\mathbf{C}}(1) \equiv \sum_{l=0}^{\infty} \mathbf{C}_l$  is the so-called long-run MA matrix. For the return equation,  $\tilde{\mathbf{C}}(1) \varepsilon_t$  represents the part of returns coming from the permanent random walk component of the price process (the “efficient” price), and  $\eta_t$  represents the part coming from transient “noise” in the price. Since  $\varepsilon_t \sim iid(\mathbf{0}, \mathbf{I})$ , the variance of the permanent part of  $\mathbf{y}_t$ , also referred to as the long-run variance, is given by

$$\mathbf{\Omega} = \tilde{\mathbf{C}}(1) \tilde{\mathbf{C}}(1)'. \quad (5)$$

The efficient return variance—that is, the variance in the returns coming from the permanent random walk component, excluding any variance coming from the noise component  $\eta_t$ —is given by the first diagonal element of  $\mathbf{\Omega}$ :

$$\omega_{11} = \tilde{c}_{11}(1)^2 + \dots + \tilde{c}_{1k}(1)^2. \quad (6)$$

The information share of variable  $i$  is now defined as the relative contribution of variable  $i$  to the long-run variance of the returns,

$$IS_i = \frac{\tilde{c}_{1i}(1)^2}{\sum_{j=1}^k \tilde{c}_{1j}(1)^2} = \frac{\tilde{c}_{1i}(1)^2}{\omega_{11}}. \quad (7)$$

That is, the information share for variable  $i$  captures the fraction of the long-run variance of the return process that can be explained by variations in variable  $i$ . The long-run variance represents the variations due to fluctuations that are not of a transient nature, and thus captures the variations coming from the unobserved random-walk efficient price process.<sup>16</sup>

---

<sup>16</sup>The information share is related but not identical to the variance decomposition of the returns. The latter decomposes the *total* variation of returns into parts that can be attributed to the other variables in the VAR. The information share decomposes the variations in returns that comes from the random-walk part of the price process.

Similarly, the permanent price impact captures the total change in the efficient price, following a trade event. The permanent price impact of shock  $i$  is simply defined as the long-run cumulative impulse response function for returns, given a shock to equation  $i$ . The long-run cumulative impulse response to the return equation, following a standard deviation shock to equation  $i$ , is equal to

$$IRF_{1,i}(\infty) = \sum_{l=0}^{\infty} c_{l,1i} = \tilde{c}_{1i}(1). \quad (8)$$

If the shock is instead fixed to unity, the permanent price impact is equal to

$$IRF_{1,i}^{\text{unit}}(\infty) = \tilde{c}_{1i}(1) D_{jj}^{-1/2} \equiv \tilde{c}_{1i}^{\text{unit}}(1). \quad (9)$$

Thus,  $IRF_{1,i}(\infty)$  reflects the impact of a typical-sized trade, whereas  $IRF_{1,i}^{\text{unit}}(\infty)$  reflects the price impact of a 1-million base currency trade, which is the minimum but also the most common trade size on the EBS system. In the empirical section, we focus on the impact of a 1-million base currency trade shock, since keeping the size of the shock identical across different sample periods and different types of orders makes comparison of price impacts more straightforward.

The ex-ante variance of the long-run impulse response for a shock  $\varepsilon_{t,i}$  to equation  $i$ , is given by

$$Var(IRF_{1,i}(\infty)) = Var\left(\left(\sum_{l=0}^{\infty} c_{l,1i}\right)\varepsilon_{t,i}\right) = \left(\sum_{l=0}^{\infty} c_{l,1i}\right)^2 = \tilde{c}_{1i}(1)^2, \quad (10)$$

which equals the long-run variance contribution of variable  $i$  to returns. The information share can thus be seen as the relative contribution to the variance of the efficient (random walk) price, as measured by the ultimate (long-run) price impacts of trades.

---

<sup>17</sup>The notion of information shares is also used in the related literature on price discovery in many markets (Hasbrouck, 1995), where information shares are derived in a cointegrated model of different prices for the same asset or cash flow. Similar to here, these information shares also reflect relative contributions to the efficient price. However, in studies of cointegrated prices across markets the information shares are usually reported as a min-max range, rather than as a spot estimate. This convention reflects the problem of fully identifying the cointegrated VAR system underlying the information shares. Such indeterminacy is less of a problem in the current setting. The price impact and information share are both functions of the structural matrix  $\mathbf{A}$ . The assumptions imposed on  $\mathbf{A}$  specify that the direction of causality is from orders to prices, rather than vice versa, and the issue of identification is therefore much less problematic than in the cointegrated many-prices/one-asset setting. The specification of  $\mathbf{A}$  is defined explicitly for each model in Section 4.

That is, the price impact measures the effect of a given trade, conditional on that trade actually occurring. The information share, on the other hand, measures the total effect of all trades that actually do occur. Alternatively put, the information shares reflect the relative ex-ante variances of the impulse response functions, a point originally made by Hasbrouck (1991b).<sup>17</sup>

Finally, we also define the long-run  $\pi$ -life of Fanelli and Paruollo (2010), which measures the number of periods it takes for the cumulative impulse-response function to enter a band

$$[\tilde{c}_{1i}^{\text{unit}}(1) - \pi|\tilde{c}_{1i}^{\text{unit}}(1)|, \tilde{c}_{1i}^{\text{unit}}(1) + \pi|\tilde{c}_{1i}^{\text{unit}}(1)|]. \quad (11)$$

The (long-run)  $\pi$ -life measures the speed of price discovery, and can be viewed as a generalization of the common concept of half-life. In a stationary system, where the impact of a shock must die out after some time, the half-life measures the number of periods it takes before the impact of a shock is reduced by one half. The price discovery VAR in (2) is assumed stationary but we are interested in the *cumulative* price impact of orders, which measures the total impact of an order on the (log) price, a non-stationary variable. The  $\pi$ -life thus measures the number of periods it takes for the cumulative impulse response to get within a  $\pm(100 \times \pi)\%$  interval of the ultimate long-run response, as illustrated in Figure 3. While it might be tempting to set  $\pi$  very small, such that one measures speed of convergence to the almost final price, specifying too small a  $\pi$  has the negative effect of introducing a lot of noise in the results. In particular, for small  $\pi$ , the  $\pi$ -life measure starts picking up on small deviations rather than the broad trend in the cumulative impulse response function. The usual half-life definition corresponds to setting  $\pi = 0.5$ . In the empirical analysis here, we set the value of  $\pi$  equal to 0.25, which appears to be a good trade-off between the noise effect and the aim to measure how long it takes for convergence to the new equilibrium price.

Although our data are regularly spaced (100ms or 250ms), we estimate the model in “quasi-event” time, where an event occurs if  $\Delta p_t^b \neq 0$ ,  $\Delta p_t^a \neq 0$ ,  $\Delta q_t^b \neq 0$ ,  $\Delta q_t^a \neq 0$ , or  $v_t \neq 0$ . Thus, we discard all 100ms (250ms) periods where no trades or no changes at the top of the book occur. We estimate the model month-by-month and use  $p = 50$  lags.<sup>18</sup>

---

<sup>18</sup>Diagnostic analysis (on a subset of the data) suggests that the results are not sensitive to this

Figure 4 shows the month-by-month average duration time between events in the euro-dollar. As is seen, it is mostly in the range between 500ms and 700ms, with the exception of a temporary rise in 2013 and 2014, during which the duration time roughly doubles for a short period. Thus, apart from this short episode, the duration time between events is fairly similar across the sample, and comparisons in event time across different sample months are broadly similar to comparisons in calendar time. In particular, the estimates of the  $\pi$ -life will be presented in event time units.

## 4 Price discovery results

In the following sub-sections, we present the empirical price discovery results based on the structural VAR specification formulated in equation (2), estimated using data between 2008 and 2017 for euro-dollar trading on EBS. As mentioned previously, the dollar-yen results are qualitatively similar to those for the euro-dollar, and for ease of exposition we focus exclusively on the euro-dollar results in the main text of the paper. The dollar-yen results are found in the Online Appendix.

Results for three different models are presented. The first two models use only market order flow and returns, with the second model decomposing the market order flow into components coming from Manual traders, Bank-ATs and HFTs. The final model allows for not only market order flows to impact prices, but also limit order flows.<sup>19</sup>

The empirical analysis thus starts with the simplest possible price discovery VAR, and then adds either limit orders or a finer classification of market orders. The results are presented in graphical form and shown in Figures 5-7. The graphs show the month-by-month estimates of the three price discovery measures introduced above, and thus illustrate the changes in the price discovery process in the euro-dollar that occurs over the 10-year long sample.

---

particular choice of lags, and similar results appear to be obtained if one increases the number of lags to 100 or even 200. However, significantly decreasing the number of lags (e.g., to 10) does appear to substantially affect the results. Since inclusion of too *few* lags in a VAR can result in biased estimates, whereas inclusion of too many lags simply results in lack of efficiency, the current lag choice of  $p = 50$  seems reasonable.

<sup>19</sup>The data do not allow us to break the limit order flow into components stemming from the three different trader groups.

## 4.1 The standard market order model

We begin with the simplest specification, given by the standard Hasbrouck SVAR. In this model, the variable vector  $\mathbf{y}_t$  and the structural coefficient matrix  $\mathbf{A}$  in equation (2) are given by,

$$\mathbf{y}_t = \begin{pmatrix} r_t \\ m_t \end{pmatrix}, \quad \mathbf{A} = \begin{pmatrix} 1 & -a_m \\ 0 & 1 \end{pmatrix}. \quad (\text{M1})$$

Here  $r_t$  denotes returns (log price changes) and  $m_t$  denotes the market order flow. Figure 5 reports the monthly estimates of the permanent price impact, information share, and  $\pi$ -life, as defined in Section 3.4. As is seen in Panel A of figure 5, the average permanent price impact of a 1-million base-currency market order is around 0.1 basis points. The price impact exhibits a substantial peak towards the end of 2008 and beginning of 2009, coinciding with the global financial crisis. Such an increased price reaction in times of market stress is consistent with previous evidence that documents time-variation in the price impact of order flow and argues that it relates to both investor behavior and sentiment (Berger, Chaboud, and Hjalmarsson, 2009). There is also some suggestion that the price impact is on a downward trend over the sample period, although this pattern is not very strong.

In contrast, the information share for market orders fell markedly over the sample period (Panel B). The information share can be viewed as a measure of how much of permanent price moves (or moves in the “efficient” price) that can be explained by market orders. In the beginning of the sample in 2008, between 40 and 50 percent of the (permanent) movements in the euro-dollar exchange rate could be explained by market orders. At the end of the sample in 2017, only about 20 percent could be explained in the same way. The overall contribution of market orders to price movements thus drops substantially during the sample period. In a pure limit order market, both limit orders and market orders likely contribute to price discovery and the results in Panel B thus suggest that limit orders might have become increasingly important for price discovery during recent years. We verify this hypothesis when considering the model with limit orders in the next sub-section.

The remaining Panel C in Figure 5 shows the  $\pi$ -life of the cumulative impulse response function, which measures how quickly the price converges to the new equilibrium following a market order. The  $\pi$ -life is measured in event time, and a  $\pi$ -life of, say, 5 periods, thus indicate that it takes about 5 “events” (trades or quote updates)

before the new equilibrium price is reached. There is a clear downward trend in the  $\pi$ -life during the first three years or so of the sample. The drop is quite substantial in relative terms, going from a  $\pi$ -life of around 5 event periods to 1 period. Given the natural limit at zero for the  $\pi$ -life, it is perhaps not surprising that the  $\pi$ -life remains reasonably constant after the initial large drop. The market thus seems to have become somewhat quicker in incorporating new information from trades; Figure 4 shows that the calendar length of events is reasonably stable over the sample period, such that the trend in the  $\pi$ -life measured in event time is representative of the trend in calendar time.

Before moving on to the model with limit order flows, we first separate the market order flow into the parts coming from the different trader types: Manual (M), Bank-AT (B), and HFT (H). Equation (2) is subsequently estimated with

$$\mathbf{y}_t = \begin{pmatrix} r_t \\ m_t^M \\ m_t^B \\ m_t^H \end{pmatrix}, \quad \mathbf{A} = \begin{pmatrix} 1 & -a_m^M & -a_m^B & -a_m^H \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (\text{M2})$$

The structural matrix  $\mathbf{A}$  allows for all three order flows to impact returns contemporaneously, but we rule out any contemporaneous interactions between the three order flows. Such a specification of the contemporaneous interaction of the variables is analogous to the specification scheme in Brogaard, Hendershott, and Riordan (2018), who also rule out all contemporaneous effects except those on returns.

The results are reported in Figure 6, which is organized in the same way as Figure 5 but with each panel containing separate results for the three order flows. The permanent price impacts (Panel A) of all three order flows evolve somewhat similarly over time; recall that in all cases, the price impacts reflect the permanent price changes following a 1-million base currency trade. However, whereas the price impact coefficients for trades from Bank-ATs and HFTs track each other very closely, the price impact for the Manual trades does exhibit a somewhat distinct pattern. The price impact for the Manual trades starts considerably higher than for the other two types of trades at the beginning of the sample, around the time of the Global Financial Crisis, but then decreases across almost the whole sample period and ends up lower than for the other trade types at the end of the sample.

The information shares (Panel B) show an even clearer pattern for the decreasing importance of Manual trades, with the information share for Manual order flow dropping drastically over the sample period, from over 30 percent in 2008 to near zero percent in 2017. In contrast, the Bank-AT and HFT shares do not change much on average from beginning to end of sample. The  $\pi$ -lives (Panel C) for all three order types exhibit similar trends to those seen for the total order flow model.

The results from the models M1 and M2 both show that the overall contribution of market orders to price variation has decreased substantially in recent years, as seen from the drop in the overall information share displayed in Panel B in Figure 5. The corresponding panel in Figure 6 further shows that this drop is in fact mostly due to a decreasing importance of Manual trades in the price discovery process. The information share is a summary measure of how much of the (permanent) variation in prices that can be attributed to or explained by movements in another variable. Roughly speaking, a drop in the information share due to a given variable can therefore be related to either a decreasing impact of each movement in that variable, and/or a decreasing overall activity in that variable. The price impact coefficients indicate the importance of a given trade when it actually occurs, and Panel A in Figure 2 indicate the relative frequency of the three different types of market orders. As is seen, both the price impact of a given-sized Manual market order as well as the relative participation of Manual market takers have decreased over the sample, and both of these aspects likely contribute to the large drop in the information share of Manual market orders. This is consistent with the fact that, over the years of our study, large dealers increasingly came to rely on automated execution algorithms when bringing to the foreign exchange market the large trades of institutional investors. This moved a large share of that order flow, widely viewed as informed, to the Bank-AT category, leaving mainly less-informed flow, such as trades arising from routine corporate transactions, to manual execution (Mehkhoff, Sarno, Schmeling, and Schrimpf, 2016).

## 4.2 Limit orders

We now proceed to allow for limit orders to affect price discovery. Specifically, the three components of the limit order flow enter the VAR model separately, and the

variable vector  $y_t$  and the structural parameter matrix are specified as,

$$\mathbf{y}_t = \begin{pmatrix} r_t \\ l_t^1 \\ l_t^2 \\ l_t^3 \\ m_t \end{pmatrix}, \quad \mathbf{A} = \begin{pmatrix} 1 & -a_l^1 & 0 & -a_l^3 & -a_m \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}. \quad (\text{M3})$$

Recall that  $l_t^1$  are price-improving limit orders,  $l_t^2$  are price-matching (quantity-changing only) limit orders at the current best bid or ask, and  $l_t^3$  are price-worsening limit orders, as defined in equation (1). By construction, the price-matching limit orders ( $l_t^2$ ) are associated with zero contemporaneous price changes, and the corresponding structural parameter is set equal to zero. Otherwise, we follow the same identifying assumption as before, and restrict the structural matrix  $\mathbf{A}$  such that the only contemporaneous effects in the model are from the different types of orders to returns.

A price-changing limit order will by definition impact the mid-quote price (i.e., a change in  $r_t$ ). However, unless the limit order is in some sense informed, such a price impact should not be permanent and therefore subsequently be reversed. By considering the permanent price impact and the information share of the efficient price process the VAR framework allows for the identification of these long-run effects, which should equal zero if the effect of limit orders is merely mechanical rather than information driven.

In addition to the three models described above (M1, M2, and M3), we also estimated a model where the separate limit orders ( $l_t^1, l_t^2, l_t^3$ ) as well as the separate market orders ( $m_t^M, m_t^B, m_t^H$ ) were included. However, the results from this model essentially told the same story as that obtained from models M2 and M3 together. For brevity, we therefore omit the results from that model.

The results for model M3 are reported in Figure 7. Somewhat surprisingly perhaps, there is no clear ranking between the price impacts of the differing types of orders (Panel A). The impact of a market order is often somewhat larger than those for the different limit orders, but this difference is often small and seemingly de-

---

<sup>20</sup>Kozhan, Moore, and Payne (2014) extends the Evans and Lyons (2002) model of order-flow driven exchange-rate determination to also incorporate limit orders. In an empirical illustration, using daily or hourly samples from Reuter's FX interdealer platform, they find that limit orders have a smaller price impact than market orders, but the quantitative differences are quite small and similar to those found here.

creasing over the sample period. This result is in contrast to Hautsch and Huang (2012), who find that market orders have substantially greater price impacts than limit orders, but it is closer to the findings of Kozhan, Moore, and Payne (2014).<sup>20</sup> The most interesting results are perhaps those for the information shares (Panel B). In line with the results in Figure 5, the information share for market orders drop from about 50 to 20 percent from the beginning to the end of the sample. At the same time, the importance of price-improving limit orders ( $l_t^1$ ) increases, and their information share increases from around 15 to 25 percent. The information shares due to price-worsening orders ( $l_t^3$ ) and price-matching limit orders ( $l_t^2$ ) increase a bit during the last few years in the sample, but these trends appear less strong than that for price-improving limit orders. There is also a sharp drop in the information share for price-matching limit orders coinciding with the decimalization reform on the EBS platform, and we discuss this result in more detail in Section 5 below.

Panel C indicates that prices typically reach a new equilibrium faster after a price-changing limit order than after a market order or a price-matching limit order. In fact, the  $\pi$ -life is in many periods equal to zero for these types of orders, suggesting that the new equilibrium price is reached almost instantaneously. While such an immediate effect of a price-changing limit order is of a fairly mechanical nature—the best bid or ask price is changed when the limit order is posted, which automatically induces a price impact—it is important to note that unless that price change was information-driven, it should subsequently be reversed. The  $\pi$ -life indicates speed of convergence to the permanent price impact, and therefore strongly indicates that such a reversal does not occur.

### 4.3 Information shares over time

As a comparison of the three different VAR specifications analyzed above, it is useful to further consider the information shares from these models. Of the three different price discovery measures that we consider, the information share is arguably the most central, as it reveals how much of the permanent variation in prices can be explained by different aspects of the trading process. Figure 8 sums up the information share results for the three different models. In Panel A, the information share for the market order flow in the standard price discovery VAR, M1, is plotted together with the total information share for the three different types of market orders in model M2. Panel

B shows the total information share due to the three different limit order flows in model M3, along with the information share for the market order flow in this model. Panel C plots the total information share for all orders in model M3. In both Panels B and C, the information share for the market orders in the simplest model, M1, is plotted as a comparison to the more elaborate models.

Panel A in Figure 8 shows that a disaggregation of the market order flow into trades due to different types of traders does not alter the overall fraction of exchange rate variance that can be explained by market orders. That is, while separating trades into different categories reveal interesting patterns, such as the sharp drop in the importance of Manual trades, it does not appear to help much in terms of the overall explanatory power of the price discovery model. Panel B clearly illustrates the declining importance of market orders and the increasing importance of limit orders in the price discovery process. Panel C highlights that these effects mostly cancel out over time, and the overall information share of market and limit orders have remained fairly constant across time.<sup>21</sup>

## 5 Price discovery and market-rule changes

In the previous section, we showed how various facets of the price discovery process evolved over the period 2008 to 2017. It is possible that these changes were, at least in part, related to market-rule changes on the EBS platform. Recall that during our sample period, four key alterations were made to the rules governing the trading process on EBS: (i) Minimum quote life (MQL, June 15, 2009); (ii) Decimalization (March 7, 2011); (iii) “Half-pip” (September 24, 2012); (iv) Latency floor (March 3, 2014).<sup>22</sup> Each of these are described in Section 2. In this section, we formally test for structural breaks in our measures of price discovery around these events. We again present results only for the euro-dollar currency pair here and relegate the corresponding results for the dollar-yen, which are qualitatively similar, to the Online

---

<sup>21</sup>The overall fraction of price variation that can be attributed to the trading process in this empirical framework is therefore around 75 percent. If the estimated VAR provided a complete characterization of the quote-revision process, the total information share for all orders and trades should equal 100 percent; in a limit order market, all price changes must be the result of posted orders. Like all econometric models, however, the VAR provides only a simplified approximation, and the unexplained part of return movements thus reflect aspects of the price discovery process, such as non-linearities, that are not accounted for in the VAR model.

<sup>22</sup>The latency floor was introduced on February 17, 2014, for the dollar-yen.

Appendix.

Specifically, our break analysis proceeds as follows. For every event, we estimate the SVAR separately for the 10 trading days before and after the event and test the null hypothesis that the price discovery measure implied by the SVAR estimates are identical in the before and after periods. Statistical inference is based on a bootstrap procedure, where we first estimate the SVAR using the full 20 days in the event window, thereby imposing the null hypothesis. The residuals from this SVAR are saved. We then draw bootstrap samples using the full-sample parameter estimates and residuals, fit the SVAR separately to the 10 days before and after the event, and calculate the associated price discovery measures. This procedure is repeated 500 times and p-values are calculated for the null hypothesis of no change, using the bootstrapped distribution of the difference in the before and after measures; using the bootstrapped standard error together with the assumption of asymptotic normality of the difference produces very similar results. The rule changes typically do not occur at the exact beginning of a month, and the month-by-month estimates presented in Figures 5-7 therefore tend to mix observations from both before and after a given rule change. In contrast, the estimates and tests presented in Tables 1 and 2 offer a clean before and after view.

The decimalization policy on EBS reduced the tick size by a factor of 10, enabling changes in the quoted bid and ask prices that were of an order of magnitude smaller than before and thus giving market makers considerably greater freedom to “fine-tune” their posted prices. In addition, prior to decimalization, the inside spread in the euro-dollar was at its minimum about 70 percent of the time (Chaboud, Dao, and Vega, 2018), and the minimum tick size was thus likely binding a great deal of the time. The considerably finer pricing grid, along with a previously binding tick size, makes it quite plausible that decimalization might have changed the way limit orders are strategically posted, and as a consequence changed their contribution to price discovery.

Chaboud, Dao, and Vega (2018) document that the most clearly noticeable effect of decimalization is a sharp increase in relative HFT taking, while there is seemingly little effect on relative making of the three trader groups, as also seen here in Figure 2. Interestingly, this shift in relative taking behavior does not seem to have any clear effect on the information share of market orders due to HFTs (Panel B in Figure 6). The effects of decimalization are instead most clearly evident in Figure 7, which

shows the results for the model with limit orders (M3). Following decimalization, the information shares for price-matching limit orders ( $l_t^2$ ) decrease dramatically in importance, with the information share dropping from 11 to 1 percent. There is also some indication that the information share for price-improving limit orders ( $l_t^1$ ) increase, whereas that for price-worsening limit orders ( $l_t^3$ ) remains mostly constant, suggesting that the finer quote grid is primarily used to slightly under-cut other traders on price. The  $\pi$ -life for price-matching orders increase, such that it takes a little longer for prices to adjust to the new equilibrium.

Following decimalization, there is thus a clear substitution away from simply altering the amount posted at the current best bid or ask, to posting orders that improve upon the best price. Consequently, the information share for price-matching orders drops, whereas the information share for price-improving orders seems to increase. The price impact of price-matching orders also decreases quite substantially, indicating that not only did this type of order become less frequent, but also less informative when actually used.<sup>23</sup>

The “half-pip” rule, implemented on September 24, 2012, represents a partial reversal of the decimalization event, increasing the tick size from 0.1 to 0.5 pips. One might therefore expect the effects seen from the decimalization to be partially reversed when the half-pip policy is put in place. This conjecture is somewhat supported by the empirical results. The strongest results are again seen for the information share of the price-matching limit orders ( $l_t^2$ ), which is indeed partially reversed with an increase from 2 to 8 percent. There is also some suggestion that the information share for the price-worsening limit orders ( $l_t^3$ ) might increase (from about 5 to 7 percent), but this is less visibly obvious. However, the information share for the price-improving ( $l_t^1$ ) orders does not exhibit a clear shift. The not-so-strong reversal effects might be due to the fact that the half-pip tick size is in most cases not binding. According to Chaboud, Dao, and Vega (2018), during the decimalization period, the inside spread in the euro-dollar is less than 1 pip about 70 percent of the time, but less than 0.5 pip only about 20 percent of the time. Thus, the 0.5 pip size might in most times not be a binding constraint, and the effects of going from 0.1 to 0.5 pips might therefore be limited.

The last two structural changes on the EBS platform—the MQL rule and the

---

<sup>23</sup>Riccó, Rindi, and Seppi (2018) show in a theoretical model how the tick size can have an effect on strategic limit order placement.

latency floor—appear to have had little or no discernible effect, and we do not discuss these further.

## 6 Conclusion

We document the changing nature of price discovery in an important electronic limit order market, using a long sample of data in the euro-dollar and dollar-yen currency pairs on the EBS platform that allows us to observe the trading activity of several groups of traders: Manual traders, bank algorithmic traders, and non-bank algorithmic traders (HFTs). The data span ten years, from 2008 to 2017, a period which saw a large increase in participation by automated computer trading. Our study thus provides a history of the recent changes in the price discovery process and discusses some of the mechanisms behind these changes.

We show that manual liquidity-taking trading has steadily decreased in importance across the entire sample period, from being a primary source of price discovery in 2008 to near irrelevance in 2017. Our results also strongly illustrate that, in modern limit order markets, price discovery occurs not only through market orders, but also through limit orders. These results are in line with recent works analyzing bond and stock markets (Fleming, Mizrach, and Nguyen, 2017, and Brogaard, Hendershott, and Riordan, 2018). Unlike previous studies, we are able to illustrate how the relative importance of market and limit orders has changed over an extended period of time, with market orders playing an ever smaller role in price discovery. At the beginning of the sample, in 2008, market orders contributed about twice as much to price discovery as limit orders, whereas at the end of the sample in 2017, this relationship was almost reversed. Our analysis therefore highlights that while it is reasonable to assume that “informed” traders will use limit orders strategically, the relative importance of market and limit orders in terms of moving the efficient price is far from fixed and given.

Roşu (2016) shows in a theoretical model how informed traders use limit orders rather than market orders when their information advantage is small. As markets become more efficient and “private” information, in a broad sense, becomes more difficult to obtain, theory would then suggest that informed traders should gradually switch from market orders to limit orders. As the increases in the speed of price discovery, which we document over our sample period, are consistent with the market

becoming increasingly more efficient, our main results therefore provide support for the theory associating market efficiency with an increased use of limit orders for informed trading.

## References

- Amihud, Y., 2002, Illiquidity and stock returns: Cross-section and time series effects, *Journal of Financial Markets* 5(1): 31–56.
- Benos, E., Brugler, J., Hjalmarsson, E., and F. Zikes, 2017, Interactions among High-Frequency Traders, *Journal of Financial and Quantitative Analysis*, 52(4): 1375–1402.
- Berger, D., Chaboud, A., and E. Hjalmarsson, 2009, What Drives Volatility Persistence in the Foreign Exchange Market?, *Journal of Financial Economics*, 94(2): 192-213.
- Beveridge, S., and C.R. Nelson, 1981, A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the ‘business cycle’. *Journal of Monetary Economics*, 7(2): 151-174.
- BIS, 2018, Monitoring of fast paced electronic markets, Markets Committee Report, Bank for International Settlements.
- Brogaard, J., Hendershott, T. and R. Riordan, 2014, High frequency trading and price discovery, *Review of Financial Studies*, 27(8): 2267–2306.
- Brogaard, J., Hendershott, T., and R. Riordan, 2018, Price Discovery without Trading: Evidence from Limit Orders, *Journal of Finance*, Forthcoming.
- Chaboud, A., Chiquoine, B., Hjalmarsson, E. and C. Vega, 2014, Rise of the machines: Algorithmic trading in the foreign exchange market, *Journal of Finance*, 69(5): 2045–2084.
- Chaboud, A., Dao, A., and C. Vega, 2018, Tick Size Change, Information Advantage, and High Frequency Trading, Working Paper, Federal Reserve Board.
- Chakravarty, S., and C. Holden, 1995, An Integrated Model of Market and Limit Orders, *Journal of Financial Intermediation*, 4(3): 213-241.
- Cont, R., Kukanov, A., and S. Stoikov, 2014, The Price Impact of Order Book Events, *Journal of Financial Econometrics*, 12(1): 47–88.
- Evans, M.D.D., and R.K. Lyons, 2002, Order Flow and Exchange Rate Dynamics, *Journal of Political Economy*, 110(1): 170–180.
- Fanelli, L. and P. Paruolo, 2010, Speed of Adjustment in Cointegrated Systems, *Journal of Econometrics*, 158: 130–141.

- Fleming, M. J., Mizrach, B., and G. Nguyen, 2017, The Microstructure of a U.S. Treasury ECN: The BrokerTec Platform, *Journal of Financial Markets*, forthcoming.
- Foucault, T., 1999, Order Flow Composition and Trading Costs in a Dynamic Limit Order Market, *Journal of Financial Markets*, 2(2): 99-134.
- Foucault, T., O. Kadan, and E. Kandel, 2005, Limit Order Book as a Market for Liquidity, *Review of Financial Studies*, 18(4): 1171-1217.
- Glosten, L., 1994, Is the Electronic Open Limit Order Book Inevitable?, *Journal of Finance*, 49(4): 1127-1161.
- Goettler, R., C. Parlour, and U. Rajan, 2005, Equilibrium in a Dynamic Limit Order Market, *Journal of Finance*, 60(5): 2149-2192.
- Goettler, R., C. Parlour, and U. Rajan, 2009, Informed Traders and Limit Order Markets, *Journal of Financial Economics*, 93(1): 67-87.
- Hasbrouck, J., 1991a, Measuring the information content of stock trades, *Journal of Finance*, 46(1): 179-207.
- Hasbrouck, J., 1991b, The summary informativeness of stock trades: An econometric analysis, *Review of Financial Studies*, 4(3): 571-595.
- Hasbrouck, J., 1995, One security, many markets: Determining the contributions to price discovery, *Journal of Finance*, 50(4): 1175-1199.
- Hasbrouck, J., and R.M. Levich, 2019, FX Liquidity and Market Metrics: New Results Using CLS Bank Settlement Data, Working Paper, NYU Stern School of Business.
- Hautsch, N., and R. Huang, 2012, The market impact of a limit order, *Journal of Economic Dynamics and Control*, 36(4): 501-522.
- Hendershott, T. and R. Riordan, 2013, Algorithmic Trading and the Market for Liquidity, *Journal of Financial and Quantitative Analysis*, 48(4): 1001-1024.
- Kaniel, R., and H. Liu, 2006, So What Orders Do Informed Traders Use?, *Journal of Business*, 79(4): 1867-1913.
- Kozhan, R., Moore, M.J., and R. Payne, 2014, Market Order Flows, Limit Order Flows, and Exchange Rate Dynamics, Working paper, University of Warwick.
- Kyle, A., 1985, Continuous Auctions and Insider Trading, *Econometrica*, 53(6): 1315-1335.
- Menkhoff, L., L. Sarno, M. Schmeling, A. Schrimpf, 2016, Information Flows in For-

- eign Exchange Markets: Dissecting Customer Currency Trades, *Journal of Finance*, 71, 601–634.
- O’Hara, M., 2015, High Frequency Market Microstructure, *Journal of Financial Economics*, 116(2): 257-270.
- Parlour, C., 1998, Price Dynamics in Limit Order Markets, *Review of Financial Studies*, 11(4): 789-816.
- Riccó, R., Rindi, B., and D.J. Seppi, 2018. Information, Liquidity, and Dynamic Limit Order Markets, Working Paper, Carnegie Mellon University.
- Roşu, I., 2009, A Dynamic Model of the Limit Order Book, *Review of Financial Studies*, 22(11): 4601-4641.
- Roşu, I., 2016, Liquidity and Information in Order Driven Markets, Working paper, HEC Paris.

Table 1: Structural break test results for model M2 with euro-dollar trading. The table shows the results for tests of structural breaks around the MQL, Decimalization, “Half-pip”, and Latency floor events. Panel A shows the results for the price impact coefficients, Panel B shows the results for the information shares, and Panel C shows the results for the  $\pi$ -life. The “Before” and “After” values represent the point estimates for each of the price discovery measures using the 10-day sample periods immediately-before and immediately-after the events, respectively. The p-values are obtained from a bootstrap simulation, as described in the main text. For the  $\pi$ -life in Panel C, some p-values are missing. These represent cases where the  $\pi$ -lives before and after the event were identical. Since the  $\pi$ -life is a discrete variable, no meaningful p-value exists in these cases.

	MQL		Decimalization		“Half-pip”		Latency floor	
	Before	After	Before	After	Before	After	Before	After
A. Price impact								
$m^M$	0.136	0.113	0.066	0.058	0.052	0.053	0.031	0.039
$m^B$	0.082	0.078	0.060	0.058	0.070	0.065	0.055	0.066
$m^H$	0.080	0.088	0.082	0.073	0.079	0.071	0.066	0.073
B. Information shares								
$m^M$	0.260	0.217	0.150	0.126	0.074	0.101	0.040	0.046
$m^B$	0.097	0.106	0.125	0.130	0.139	0.154	0.131	0.145
$m^H$	0.100	0.145	0.227	0.200	0.168	0.172	0.196	0.172
C. $\pi$ -life								
$m^M$	4.000	3.000	0.000	1.000	1.000	1.000	0.000	0.000
$m^B$	4.000	4.000	2.000	4.000	2.000	2.000	2.000	2.000
$m^H$	3.000	3.000	1.000	2.000	1.000	1.000	1.000	1.000

Table 2: Structural break test results for model M3 with euro-dollar trading. The table shows the results for tests of structural breaks around the MQL, Decimalization, “Half-pip”, and Latency floor events. Panel A shows the results for the price impact coefficients, Panel B shows the results for the information shares, and Panel C shows the results for the  $\pi$ -life. The “Before” and “After” values represent the point estimates for each of the price discovery measures using the 10-day sample periods immediately-before and immediately-after the events, respectively. The p-values are obtained from a bootstrap simulation, as described in the main text. For the  $\pi$ -life in Panel C, some p-values are missing. These represent cases where the  $\pi$ -lives before and after the event were identical. Since the  $\pi$ -life is a discrete variable, no meaningful p-value exists in these cases.

	MQL			Decimalization			“Half-pip”			Latency floor		
	Before	After	p-val	Before	After	p-val	Before	After	p-val	Before	After	p-val
A. Price impact												
$l^1$	0.138	0.069	0.000	0.075	0.051	0.008	0.101	0.050	0.004	0.092	0.054	0.084
$l^2$	0.100	0.053	0.040	0.070	0.013	0.000	0.036	0.033	0.000	0.046	0.028	0.116
$l^3$	0.121	0.063	0.052	0.053	0.032	0.320	0.059	0.034	0.980	0.055	0.032	0.392
$m$	0.272	0.139	0.000	0.141	0.081	0.160	0.141	0.073	0.000	0.114	0.072	0.000
B. Information shares												
$l^1$	0.117	0.107	0.256	0.112	0.162	0.000	0.170	0.164	0.500	0.195	0.172	0.236
$l^2$	0.070	0.073	0.576	0.111	0.011	0.000	0.024	0.080	0.000	0.055	0.052	0.760
$l^3$	0.077	0.079	0.848	0.046	0.055	0.100	0.048	0.067	0.000	0.062	0.052	0.408
$m$	0.507	0.495	0.324	0.426	0.418	0.536	0.336	0.364	0.024	0.313	0.317	0.840
C. $\pi$ -life												
$l^1$	0.000	1.000	0.050	1.000	0.000	0.010	0.000	0.000	-	0.000	0.000	-
$l^2$	5.000	4.000	0.150	2.000	6.000	0.000	5.000	3.000	0.000	2.000	3.000	0.242
$l^3$	2.000	3.000	0.116	2.000	0.000	0.000	0.000	1.000	0.036	1.000	0.000	0.258
$m$	4.000	4.000	-	1.000	2.000	0.014	2.000	2.000	-	1.000	1.000	-

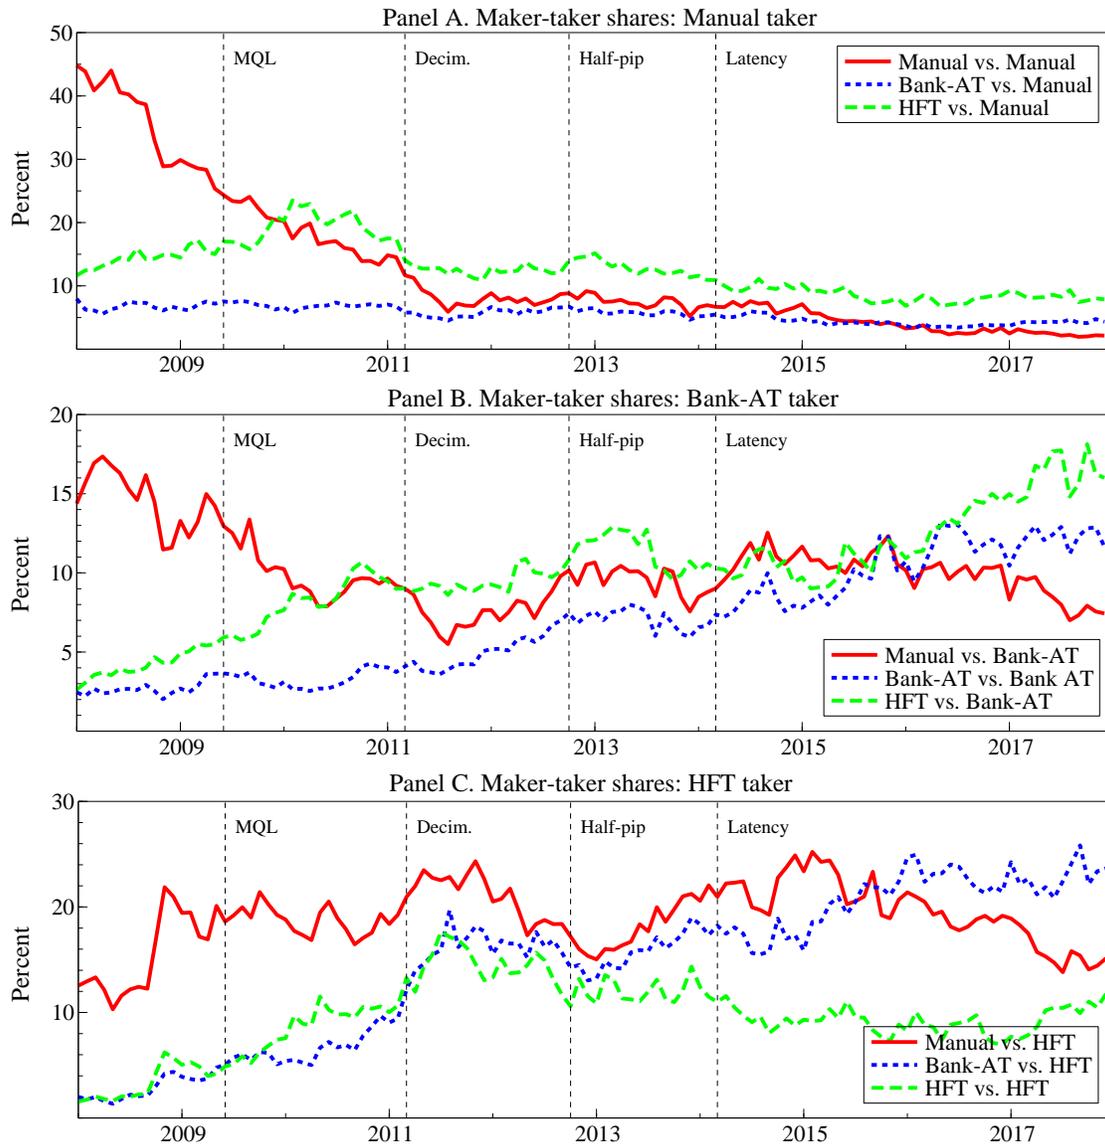


Figure 1: Detailed maker-taker shares for euro-dollar trading. The graphs show the relative contribution, to the overall trading volume, of the nine possible maker-taker pairs. The figure plots monthly averages.

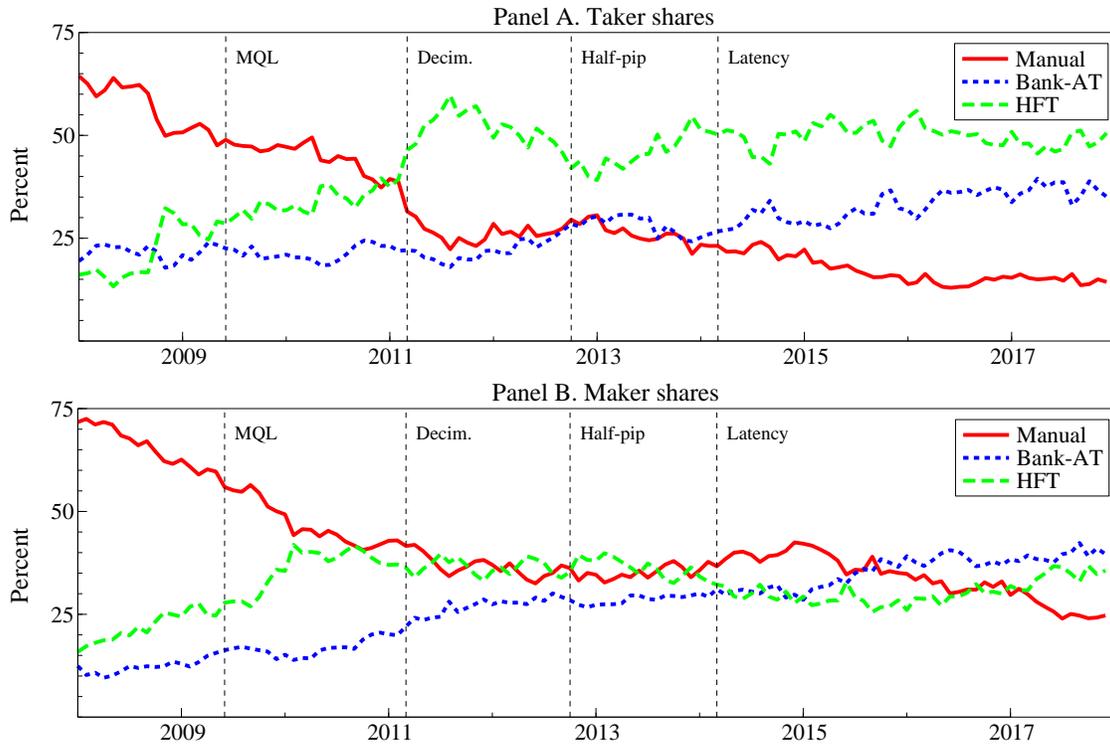


Figure 2: Aggregated maker-taker shares for euro-dollar trading. The graphs show the relative contribution, to the overall trading volume, of the three different types of makers and takers: Manual, Bank-AT, and HFT. In Panel A, the relative contributions for each type of maker are shown, and in Panel B, the relative contributions for each type of taker are shown. The figure plots monthly averages.

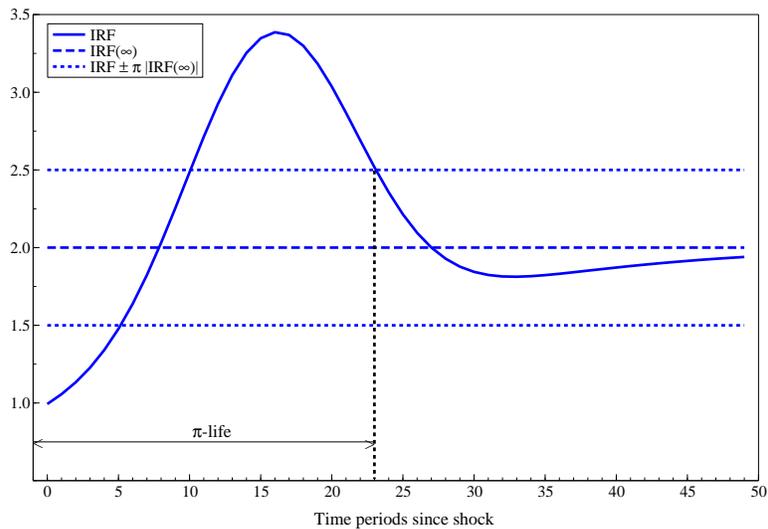


Figure 3: Illustration of the  $\pi$ -life measure. The graph illustrates the speed-of-adjustment measure,  $\pi$ -life, as a function of the shape of the cumulative impulse response function. The solid line represents an illustrative cumulative impulse-response function, plotted as a function of time-since-the-shock. The dashed line represents the permanent (long-run) impact, and the dotted lines indicate an interval of width  $2\pi$  ( $\pi = 0.25$ ), centered around the long-run impact. The  $\pi$ -life is defined as the number of periods it takes before the cumulative impulse-response function permanently enters this interval.

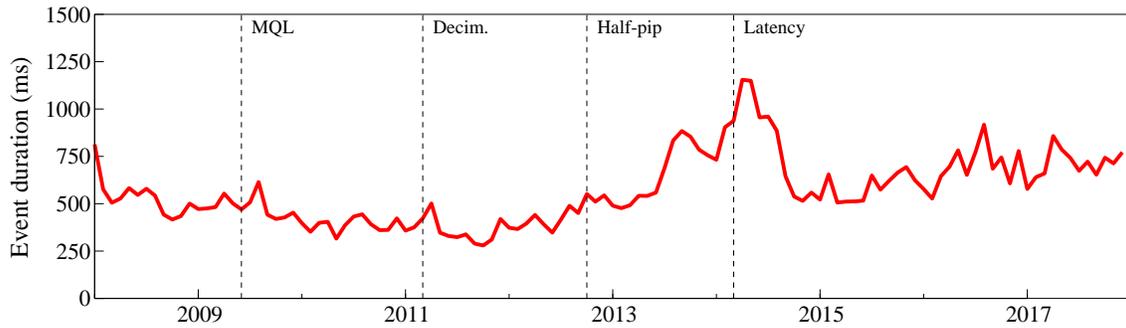


Figure 4: Event duration for euro-dollar trading. The graph shows the average time, in milliseconds, between the occurrence of two trade events (trades or changes at the top of the order book). The figure plots averages calculated for each month in the sample.

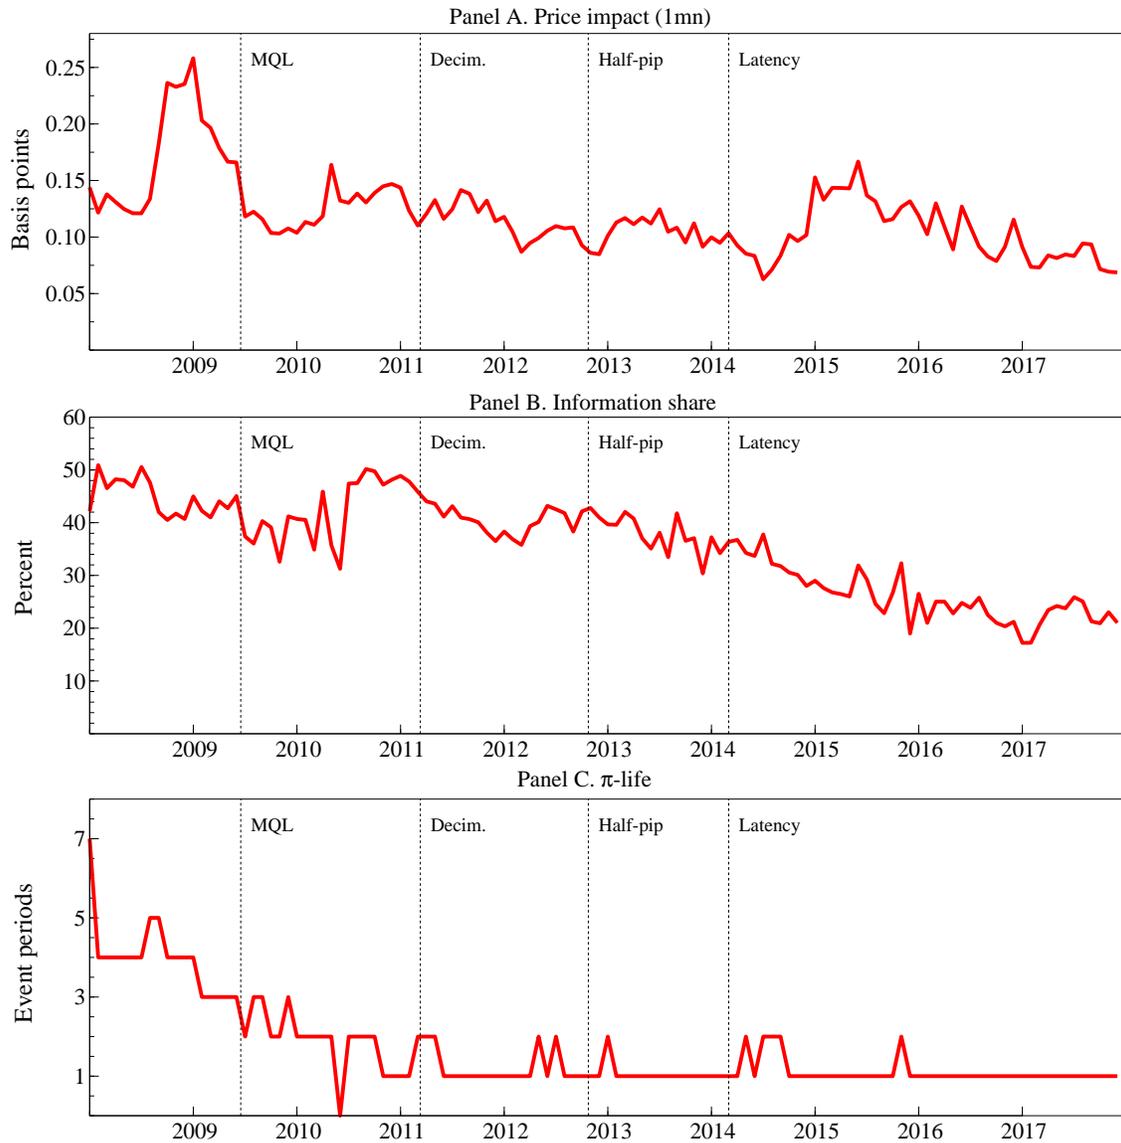


Figure 5: Results for model M1 with euro-dollar trading. The graphs show the month-by-month estimates from model M1. Panel A plots the estimates of the price impact, measured in basis points, following a 1-million Euro order. Panel B plots the estimates of the information share for the market orders, and Panel C plots the estimates of the  $\pi$ -life for the market orders (in event time).

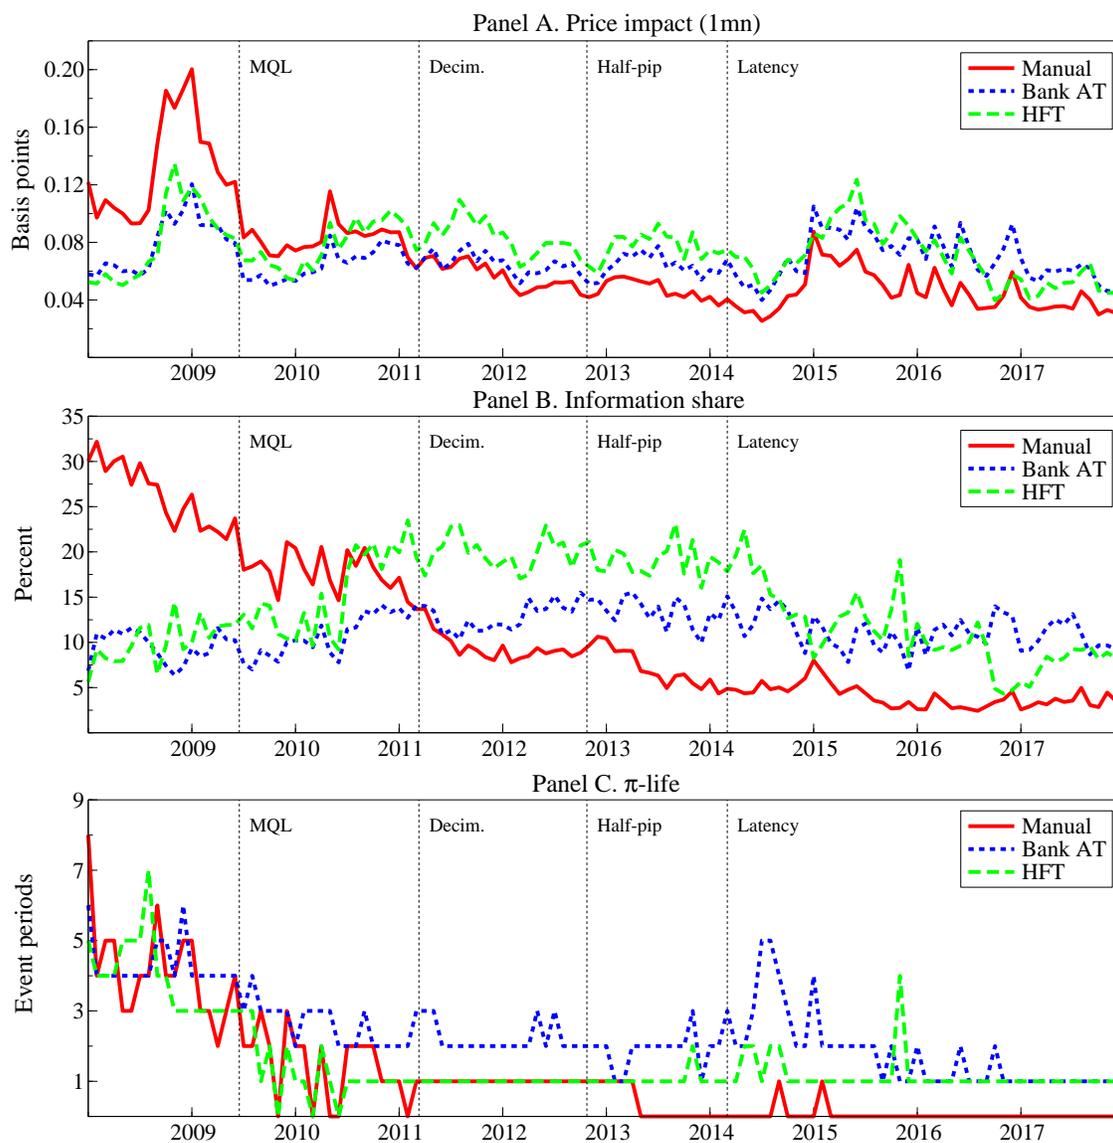


Figure 6: Results for model M2 with euro-dollar trading. The graphs show the month-by-month estimates from model M2. Panel A plots the estimates of the price impact, measured in basis points, following a 1-million Euro order. Panel B plots the estimates of the information share for the market orders, and Panel C plots the estimates of the  $\pi$ -life for the market orders (in event time). The solid lines correspond to Manual market orders, the dotted lines correspond to Bank-AT market orders, and the dashed lines correspond to HFT market orders.

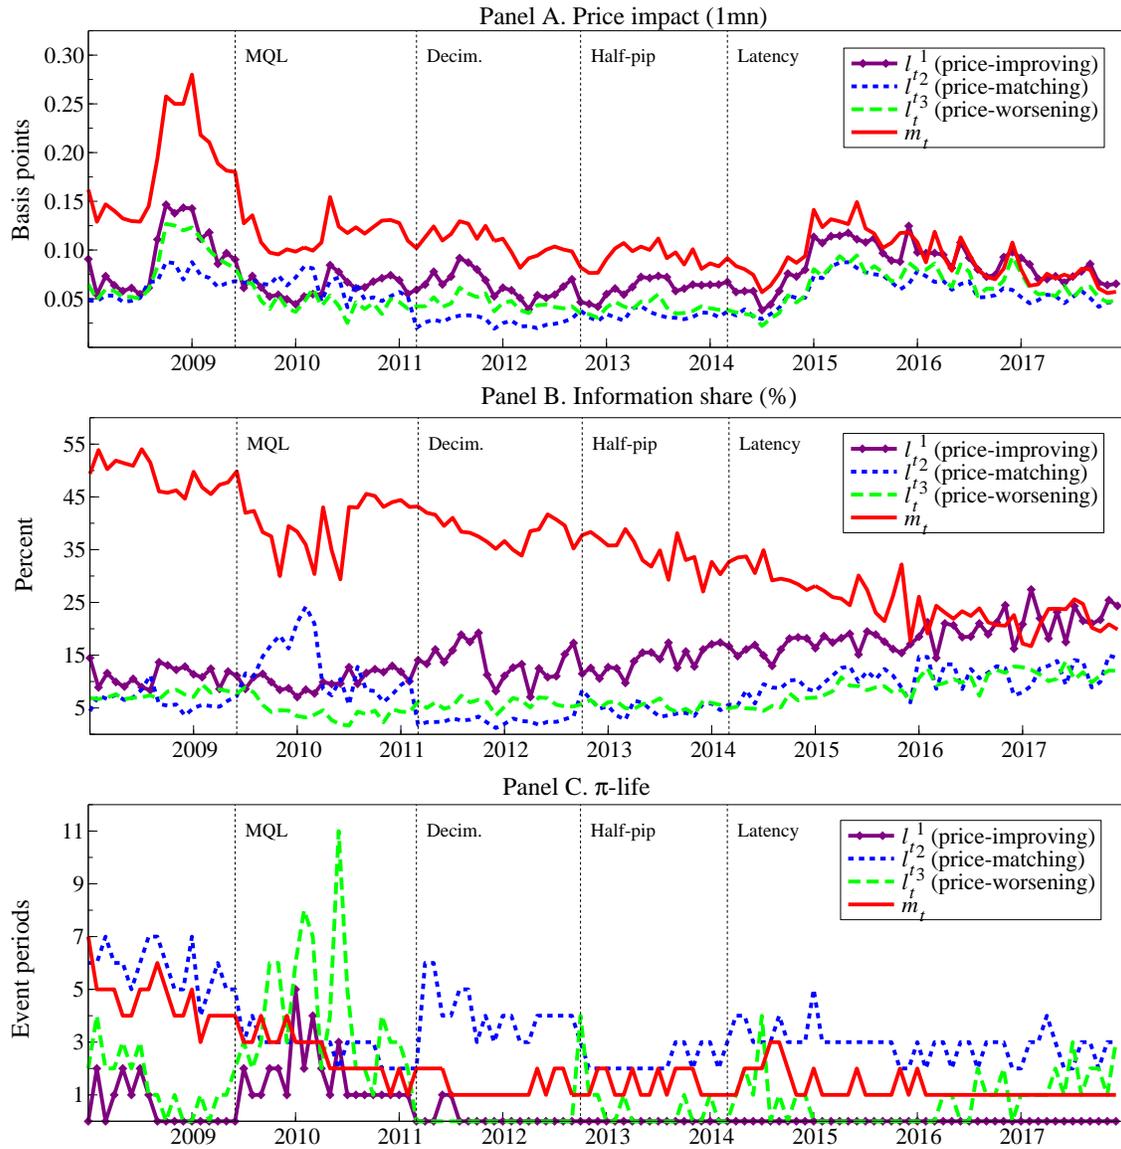


Figure 7: Results for model M3 with euro-dollar trading. The graphs show the month-by-month estimates from model M3. Panel A plots estimates of the price impact, measured in basis points, following a 1-million Euro order. Panel B plots the estimates of the information shares for the market and limit orders, and Panel C plots the estimates of the  $\pi$ -life for the market and limit orders (in event time). The solid-diamond lines correspond to price-improving limit orders ( $l_t^1$ ), the dotted lines correspond to price-matching limit orders ( $l_t^2$ ), the dashed lines correspond to price-worsening limit orders ( $l_t^3$ ), and the solid lines correspond to market orders ( $m_t$ ).

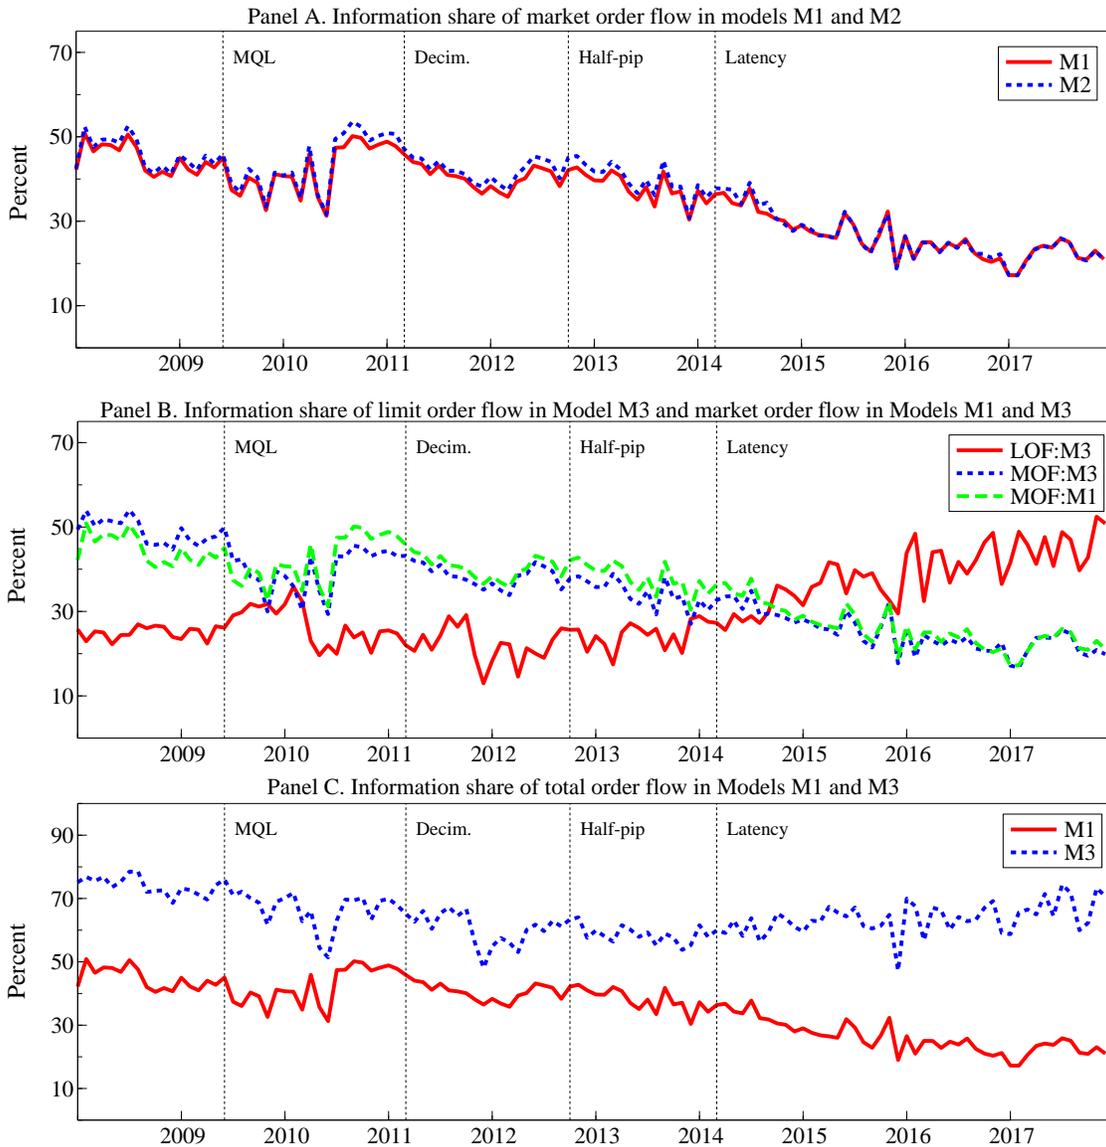


Figure 8: Information shares for euro-dollar trading. The graphs show the month-by-month estimates of different information shares implied by models M1, M2, and M3. Panel A shows the information share for market orders in model M1 (solid line), along with the total information share for all market orders in model M2 (dotted line). Thus, the dotted line represents the sum of the information shares for the Manual, Bank-AT, and HFT market orders in model M2. Panel B shows the total information share for the limit orders in model M3 (dashed line), the information share for the market orders in model M3 (dotted line), and the information share for market orders in model M1 (solid line). Thus, the dashed line represents the sum of the information shares for the price-improving, price-matching, and price-worsening limit orders in M3. Panel C shows the total information share for all orders in model M3 (dotted line), along with the information share for market orders in model M1 (solid line). Thus, the dotted line represents the sum of the information shares for the price-improving, price-matching and price-worsening limit orders, as well as the market orders, in model M3.