Human vs. High Frequency Traders in the Interbank FX Market, Role of Tick Size☆

Soheil Mahmoodzadeh

Faculty of Economics, University of Cambridge, Sidgwick Avenue, Cambridge CB3 9DD, United Kingdom

Ramazan Gençay

Department of Economics, Simon Fraser University, Burnaby, BC, V5A 1S6, Canada

Abstract

This paper studies the changes in market quality due to the smaller tick size in the interbank foreign exchange market. Coupled with the lower tick size, the special composition of traders and their order placement strategies have provided a suitable environment for high frequency traders (HFT's) to implement the sub-penny jumping strategy to front-run human traders. We show that spread was reduced following the introduction of decimal pip pricing. However, the benefit of the reduction in spread was mostly absorbed by the HFT's. Market depth was also reduced significantly due to the occupation of the top of the order book by HFT's. The new environment changed the market maker-market taker composition between different types of traders and also altered the price impact of the order flow.

Keywords: Interbank Foreign Exchange Market, Tick Size, Market Quality
JEL: F31, G14, G15

1. Introduction

In this paper, we study changes in market quality in the Electronic Broking Services (EBS) following the adoption of decimal pip tick size. EBS is the leading interbank foreign exchange (FX) market, and it is used mainly to trade the major currency pairs EUR/USD, USD/JPY, EUR/JPY, USD/CHF, and EUR/CHF in units of millions. EBS and Reuters are the predominant sources of interbank liquidity in the FX market. In March 2011, EBS

☆The authors would like to thank three anonymous referees for their valuable inputs. We are deeply grateful to Alain Chaboud and Angelo Ranaldo, for their generosity and extensive discussions. We are also grateful for discussions, comments and suggestions made by Alec Schmidt, Michael Tseng, Mark Bruce (EBS head of strategic innovation), Oliver Linton, Jakub Rojcek, and seminar participants at the 11^{th} World Congress of the Econometric Society, Simon Fraser University, and the University of Victoria. Ramazan Gençay gratefully acknowledges financial support from the Natural Sciences and Engineering Research Council of Canada and the Social Sciences and Humanities Research Council of Canada.

Email addresses: `sm2179@cam.ac.uk` (Soheil Mahmoodzadeh), `rgencay@sfu.ca` (Ramazan Gençay)

decided to reduce the tick size, minimum price movement, on major currency pairs from a pip (four decimals) to a decimal pip (five decimals). For example, if the tick size is a pip, equal to 0.0001, and the EUR/USD best bid is 1.39940, then a buyer can improve this price by placing an order with a price of 1.39950. However, if the tick size is a decimal pip, equal to 0.00001, the buyer can also place an order with a price of 1.39941. The decision to shift to decimal pip pricing was mainly driven by a competitive effort to match some smaller trading platforms to attract more HFT's. Reuters, the main competitor of EBS, still uses pip pricing.

We argue how the specific structure of the EBS market has helped HFT's implement the sub-penny jumping strategy to take advantage of the lower tick size. Sub-penny jumping is a type of front-running strategy in which the sub-penny jumper trades in front of and on the same side of a large, patient trader by improving the price by the smallest possible amount (the tick size).[1] In the EBS market, there are two main types of traders, human and algorithmic traders (see Chaboud et al. (2014) and Schmidt (2012)). Human traders are individuals who trade at the trading desks of major banks. Unlike in the equity market, human traders play a vital role in the interbank FX market for providing liquidity. They were either market taker or market maker in about 65% of EUR/USD transactions in 2011.[2]

Using difference-in-difference estimators, we find that spread, as a measure of liquidity cost, decreased following the introduction of decimal pip pricing. However, we argue that, due to the implementation of the sub-penny jumping strategy, the benefit of the reduction in the spread was mostly absorbed by HFT's. The absolute value of realized spread also decreased toward zero meaning that the market maker would suffer less adverse price move. HFT's, using the sub-penny jumping strategy, occupied the top of the order book with smaller orders, which led to the significant drop in market depth. The tick size change from pip to decimal pip also changed the market maker-market taker composition of human and algorithmic traders in the EBS market. While the share of the computer maker-computer taker from total trades increased, the share of human maker-human taker decreased after the tick size change. The new environment also altered the informational content of trades

[1] In equity markets, exchanges are forbidden from accepting, ranking or displaying orders in price increments smaller than a penny. However, sub-penny trading is allowed if trades are executed on Alternative Trading Systems such as dark pools, or internalized by broker-dealers.

[2] The market maker provides liquidity by posting quotes and the market taker trades at the quoted price.

influencing the relationship between returns and order flows measured by the price impact regression. Finally, based on empirical considerations, we argue that the optimal tick size for EUR/USD in the EBS market is between a pip and a decimal pip. Not surprisingly, the reduction in tick size was partially reversed from a decimal pip to half pip in September 2012. Half pip is decimal pip pricing in which traders can use only '0" or '5" as the last digit. Market quality measures improved upon adoption of half pip pricing by EBS.

The remainder of the paper is organized as follows: Section 2 provides a brief literature review, Section 3 describes the different datasets that we have used, Section 4 provides the tick size changes in EBS market from 2004 to 2013, Section 5 explains the structure of the EBS market, Section 6 describes the sub-penny jumping, Section 7 provides empirical evidence for the existence of sub-penny jumping, Section 8 discusses the effects of tick size changes on interbank FX market quality, and Section 9 concludes.

2. Literature Review

High frequency trading practices are becoming the predominant feature in the context of financial markets owing to the technological developments.[3] The increasing pervasiveness of HFT means that it is a central issue of modern financial markets to better understand the economic role of HFT's and their impact on market quality. While some studies have supported the beneficial role of HFT (see among others, Terrence et al. (2011), Hasbrouck and Saar (2013), Brogaard et al. (2014), and Chaboud et al. (2014)), other studies have highlighted the harmful aspects of it (see for example Hirschey (2013), Biais and Woolley (2011), and Kirilenko et al. (2015)).

The algorithmic trading started in January 2004 in EBS market. The participation has grown steadily and reached 20% at the end of 2005 and around 60% at the of 2007 (see Chaboud et al. (2014)). In 2011, computer traders were either market taker or market maker of the EUR/USD transactions in about 85% of all trades. In this paper, we study how the tick size changes altered the HFT strategies and consequently market quality measures. Only a few papers have examined the effects of tick size on the interbank FX market. Using proprietary data from EBS, Schmidt (2012) documents that human traders did not use the last digit very often under decimal pip pricing. He also provides a taxonomy of types of

[3]For an overview of the literature on HFT, see the surveys by Jones (2013) and O'Hara (2015).

EBS customers and their order placement characteristics. Lallouache and Abergel (2014) analyze the EBS data, EUR/USD and USD/JPY distributions, and report price clustering at prices ending in "0" and "5" after March 2011. They argue that automated traders take price priority by submitting limit orders one tick ahead of clusters. However, they do not provide insights into why these traders take such priority. The observation that emerges from these two papers is that the EBS market microstructure changed significantly after the introduction of decimal pip pricing.

Empirical studies of stock exchanges generally find that a tick size reduction is associated with a decline in both spread and depth.[4] Ready (1999) presents empirical evidence that dealers on the NYSE impose adverse selection costs on standing limit orders by selectively stepping ahead of these orders to interact with incoming marketable orders. Goldstein and Kavajecz (2000) examine the NYSE change to sixteenths and find that even if the effective spread generally declines, under sixteenths, depth decreases throughout the limit order book. Bessembinder (2000) finds that spreads are reduced in NASDAQ market when tick size decreases. Jones and Lipson (2001) show that when NYSE lowered its minimum price increment on most stocks from eighths to sixteenths of a dollar in 1997, quoted and effective spreads declined, but realized execution costs for institutional trades increased. Bacidore et al. (2003) use NYSE system order data to examine changes in trader behavior, displayed liquidity supply and execution quality around the reduction of the minimum tick size in U.S. equity markets. Among other things, the authors find that while the inside bid-ask spread tightens, displayed liquidity deeper in the limit order book is reduced. Bessembinder (2003) finds that both spreads and intraday return volatility decreased after decimalization, with the reduction in quoted spreads being stronger in heavily traded large-capitalization NASDAQ stocks. Bourghelle and Declerck (2004) show that a relatively larger tick size encourages traders to submit and expose limit orders, while a smaller tick size induces frequent undercutting strategies. Ahn et al. (2007) find that the spread declined significantly after the Tokyo Stock Exchange (TSE) introduced a change in tick size for stocks traded within certain price ranges. Reductions in spreads are larger for more liquid stocks with

[4] There has been a long debate about the effects of tick size on financial markets. Recently, on June 2014, the Securities and Exchange Commission (SEC) ordered a plan to implement a targeted one year pilot program that will widen the tick size for certain small capitalization stocks to assess effects on market quality.

larger tick size reductions. Cai et al. (2008) conclude that a tick size reduction has no general effect on the TSE because trading volume, the number of shares traded, and the average trade size react differently. In the theoretical and empirical literature on tick size in equity markets, there is also an interesting debate on whether the effects of a tick size reduction may depend on the liquidity of a stock (see, for example, Bourghelle and Declerck (2004) and Goldstein and Kavajecz (2000)). Similar considerations also apply to interbank foreign exchange markets, where tick size is always smaller for more liquid pairs. The effects of tick size changes have led to theoretical studies of the choice of a suitable minimum tick size. These include, among others, Harris (1994), Anshuman and Kalay (1998), Cordella and Foucault (1999), Alexander and Zabotina (2005), Kadan (2006), and Ascioglu et al. (2010).

3. Data Description

Banks and large financial institutions trade currencies with each other on two interbank electronic trading platforms: EBS and Reuters. In practice, EBS is the leading liquidity provider for EUR/USD, USD/JPY, EUR/JPY, USD/CHF and EUR/CHF, and Reuters is the primary trading venue for commonwealth and emerging market currencies. To study the effects of tick size changes on interbank FX market, we use four EBS data sets as following:

EBS level 1, 2004-2008. This data includes one level quote (best buy and sell) and deal records at one second frequency.[5] This data has been used in Section 4 to find the minimum tick sizes in the EBS market from January 2004 to December 2008.

EBS level 10, 2009-2011. This data includes ten level quotes and deal records at 100 milliseconds frequency. This data is used in Section 4 to find the minimum tick sizes in the EBS market from January 2009 to December 2011. The most important tick size changes have been implemented in March 2011 and the data covers that period. This data enable us to analyze different levels of the order book and it has been used in analysis related to pip and decimal pip tick sizes in Section 4, Section 5, Section 7, Section 8.1, and Section 8.2.

EBS level 10, 2013. This data includes ten level quotes and deal records at 100 milliseconds frequency from January to February 2013. We obtained the tick size changes in 2012 from

[5] "Deal" is a term used by EBS. It is synonymous with trade.

EBS. The tick sizes from January to February 2013 in Section 4 are consistent with EBS 2012 information. We use this data to analyze different levels of the order book related to the half pip tick size in Section 4, Section 5, Section 7, Section 8.1, and Section 8.2.

In the three datasets above, the dealt prices are the highest buying or lowest selling deal price between two consecutive snapshots of the order book, rounded to second or 100 milliseconds. There are 22.5 and 225 million snapshots of the order books for each year per major currency for second and millisecond frequencies, respectively. The data do not include dealer identifications and there are no hidden orders. In EBS terminology, EUR/USD denotes the amount of local currency, USD, required to buy (or sell) one unit of the base currency, EUR. Orders are submitted in units of millions of the base currency.[6] We excluded thin weekend trading periods and holidays because liquidity during those periods may have been extremely limited. We also controlled for daylight savings and standard time. Similar conventions were adopted by Andersen et al. (2003) and Chaboud et al. (2004).

EBS Maker-Taker, 2011. This transaction data provides the minute-by-minute last price, aggregated volume, and direction of trades broken down into categories specifying the human (H) and computer (C) traders from January to June 2011. It provides eight categories specifying of human and computer makers and takers: HH_b, HH_s, HC_b, HC_s, CH_b, CH_s, CC_b, and CC_s, where b denotes buyer initiated trade and s denotes seller initiated trade. The market maker posts quotes before the market taker (initiator) choose to trade at the quoted price. The sum of buyer and seller initiated volumes provides volumes attributable to the human maker-human taker HH, computer maker-human taker CH, human maker-computer taker HC, and computer maker-computer taker CC. The difference between the buyer and seller initiated volumes provides the order flow. The sum of the four order flows gives the total order flow. This data has been used in Section 8.3 and Section 8.4.

4. Minimum Tick Size

The minimum tick size is the smallest possible increment between quoted prices in the limit order book. Tick size has important effects on liquidity provision and transaction costs.

[6] EBS introduced small currencies in early 2010 and allowed participants to trade in increments of 100,000 units of the base currency. The trading volume in small currencies is very minimal. Therefore, we do not analyze them.

A large tick size increases trading costs by widening the spread. The liquidity effect arises because small tick size makes it easier for traders to step ahead of an existing limit order, increasing the cost for traders who provide liquidity. It is a challenge for markets to choose the optimal tick size to balance the liquidity and transaction cost effects.

We have analyzed our datasets to find the tick size changes on EBS from 2004 to 2013.[7] Figure 1 illustrates the tick size changes in the EBS market.[8] For instance, tick size was a pip (4 decimal points) for EUR/USD until March 2011. It was then reduced to a decimal pip (5 decimal points). Finally, it was changed to a half pip (4.5) in September 2012. Half pip pricing is a special case of decimal pip pricing in which the fifth digit can only be "0" or "5". The pip (decimal pip) tick size is 2 (3) decimal points when the local currency is JPY, because currencies are traded versus 100 JPY. Figure 1 indicates a general pattern that goes from pip to decimal pip pricing, then reverts to half pip pricing over a two-year period. The original change from pip to decimal pip pricing on EBS began in September 2010 with less active currency pairs, presumably to test the reactions of market participants. In October 2010, EBS reduced tick size for the next group of most commonly traded currency pairs: EUR/GBP, GBP/USD, and USD/CAD. Finally, EBS reduced the tick size for the five major currency pairs –EUR/USD, USD/JPY, EUR/JPY, USD/CHF and EUR/CHF– to decimal pip in March 2011. The decision to lower the tick size was mainly driven to match some smaller trading platforms to attract more HFT's in a competitive market. This move generated intense debate among the two main types of traders. Although decimal pip pricing was welcomed by HFT's, human traders believed that HFT's already had an unfair advantage, which was enhanced by the smaller tick size.[9] After EBS shifted to decimal pip pricing, it eventually took the view that it risked losing more business by continuing with decimal pip pricing. As a consequence, there was a reversion to half pip pricing for most currency pairs in September 2012. Table 1 provides the dates for the tick size changes in the EBS market.

[7] As far as we know, there has not been any change regarding tick sizes on EBS in 2014 and 2015.

[8] We have excluded currency pairs with minimal activities.

[9] For example, see pages 14-15 from the report by the Bank of International Settlements (BIS): http://www.bis.org/publ/mktc05.pdf.

Figure 1: EBS Minimum Tick Size Changes, 2004-2013

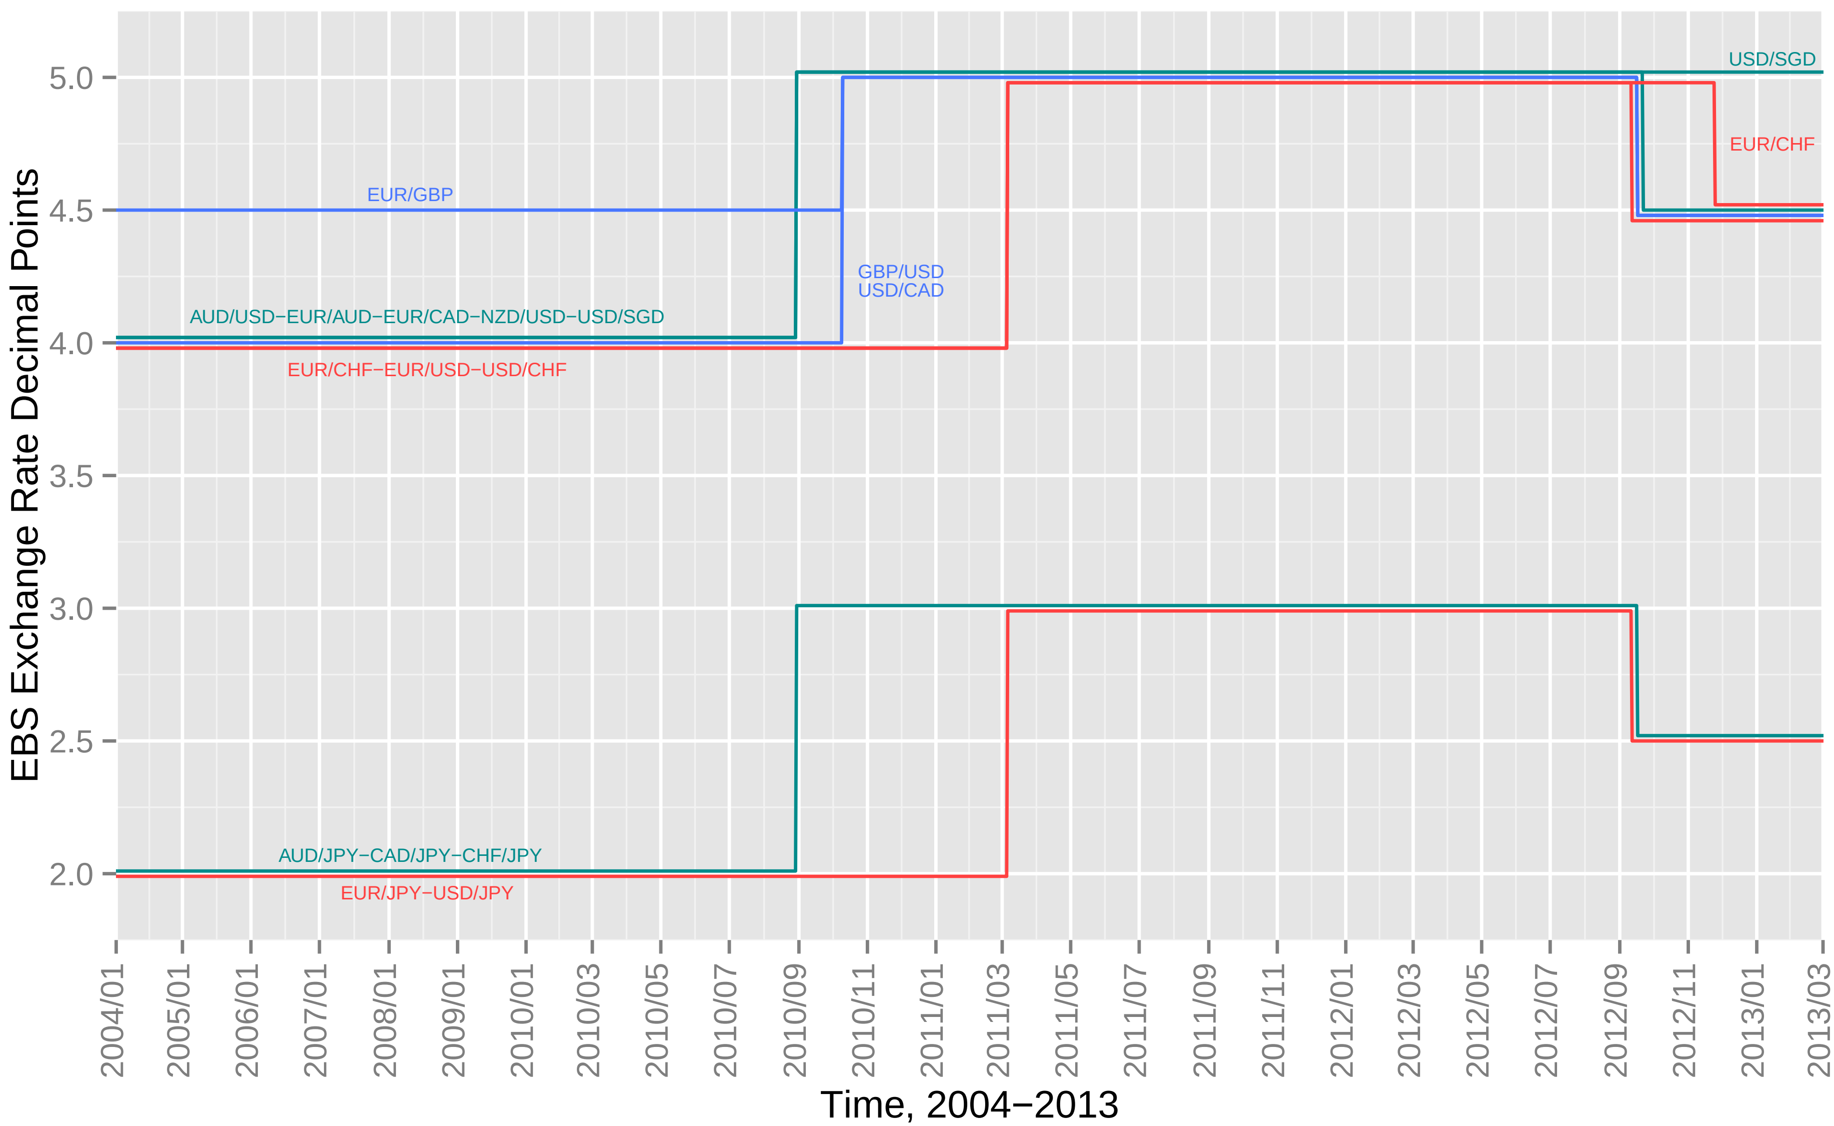

Notes: Figure 1 illustrates the tick size changes in the EBS market from 2004 to 2013 for different currency pairs. The vertical axis shows the decimal points used in the exchange rates. As an example, EUR/USD tick size was 0.0001 (one pip) until March 2011. Thus, if the best bid was 1.39940, then a trader can improve this price by adding the tick size and placing an order with the price of 1.39950. However, the trader was not allowed to place an order with a price of 1.39941 because the added value of 0.00001 (a decimal pip) was less than the tick size of 0.0001. The pip (decimal pip) tick size is 2 (3) decimal points when the local currency is JPY, as the exchange rates are based on 100 JPY. The 2.5 and 4.5 decimal points refer to half pip pricing, which is a special case of decimal pip pricing in which the last digit could only be "0" or "5". For example, the tick size of 4.5 for the EUR/USD exchange rate means that the tick size is one decimal pip (5 decimal points), and hence, the fifth digit can only be "0" or "5".

Table 1: EBS Minimum Tick Size Changes, 2004-2013

	2010/08/30	2010/10/10	2011/03/07	2012/09/17	2012/11/25
EUR/GBP		$hp \rightarrow dp$		$dp \rightarrow hp$	
AUD/USD	$p \rightarrow dp$			$dp \rightarrow hp$	
EUR/AUD	$p \rightarrow dp$			$dp \rightarrow hp$	
EUR/CAD	$p \rightarrow dp$			$dp \rightarrow hp$	
NZD/USD	$p \rightarrow dp$			$dp \rightarrow hp$	
USD/SGD	$p \rightarrow dp$				
EUR/CHF			$p \rightarrow dp$		$dp \rightarrow hp$
EUR/USD			$p \rightarrow dp$	$dp \rightarrow hp$	
USD/CHF			$p \rightarrow dp$	$dp \rightarrow hp$	
AUD/JPY	$p \rightarrow dp$			$dp \rightarrow hp$	
CAD/JPY	$p \rightarrow dp$			$dp \rightarrow hp$	
CHF/JPY	$p \rightarrow dp$			$dp \rightarrow hp$	
EUR/JPY			$p \rightarrow dp$		
USD/JPY			$p \rightarrow dp$		
GBP/USD		$p \rightarrow dp$		$dp \rightarrow hp$	
USD/CAD		$p \rightarrow dp$		$dp \rightarrow hp$	

Notes: Table 1 provides the tick size changes on EBS from 2004 to 2013, p refers to pip, dp indicates decimal pip, and hp refers to half pip.

5. EBS Market Structure

In this section, we provide information about the EBS market structure to help explain how, given the lower minimum tick size, the structure of the interbank FX market facilitated HFT's taking advantage of human traders. Our analysis is based on the EUR/USD currency pair. The other major currency pairs –USD/JPY, USD/CHF, EUR/CHF, and EUR/JPY– are mostly similar to EUR/USD.[10] There are two main types of traders in the EBS market: "automated" traders, who use an automated interface (AI) to place orders without human intervention; and "human" traders, who use GUI-based access for order management. Human traders are individuals who trade, using keyboard, at the trading desks of major banks. Unlike in the equity market, these traditional traders play a vital role in the interbank FX market for providing liquidity. As depicted in Figures 23 and 24, in about 65% of all trades, human traders were either taker or maker of the transactions before the tick size change in March 2011. The important role of human traders distinguishes the structure of the interbank FX market from other markets. EBS first allowed AI into the market in 2004 (see

[10]The full results for USD/JPY, USD/CHF, EUR/CHF, and EUR/JPY are available upon request.

Chaboud et al. (2014) for more details). The main component of an AI in the EBS market is the professional trading community (PTC), which places orders at very high frequency.[11] Using propriety data Schmidt (2012) provides details regarding the ecology of the EBS market in 2011. The PTC typically uses an order size of one million and almost never uses an order size exceeding four million. Human traders place larger orders, for example, approximately 5% of orders submitted by them reach the size of ten million. The fill ratio is defined as the ratio of dealt quotes to submitted quotes. This ratio is more than 50% for human traders and approximately 8% for the PTC. The high fill ratio for human traders implies that these traders typically do not cancel their orders. Human traders did not widely use decimal pip pricing after the tick size change. Approximately 80% of the human traders used decimal pip pricing in less than 20% of their orders. In contrast, AI traders adopted decimal pip pricing widely in their quoted prices.[12]

Stylized Fact: Human traders (slow traders), place larger orders at prices ending in "0" (under pip pricing) and on average do not cancel half of their orders.

There are some reasons why human traders were reluctant to use decimal pip pricing. Prior to March 2011, the value of a tick size was $100 for an order size of one million under pip pricing. This value fell to $10 under decimal pip pricing. Furthermore, traditional human traders had been accustomed to pip pricing over many years and found it difficult to adapt to decimal pip pricing. If human traders did not use decimal pip pricing and used "0" as the fifth digit, there should be a price clustering at zero after the tick size change. Figures 2 and 3 provide distributions of the last digit of prices under pip and decimal pip pricing. The last digit distribution is nearly uniform under pip pricing, but there is a price clustering at zero under decimal pip pricing. A weaker price clustering is also found for prices ending in "5". The reason for this may be that some human traders used "5" to partially adapt to decimal pip pricing. If human traders typically used zero for the last digit and placed larger orders, the volume distributions of prices ending in zero and non-zero digits should also differ.

[11] Automated Traders are further divided into three subgroups: Professional Trading Community (PTC), Bank AI, and Aggregators. PTC are high frequency market makers, usually in the market by prime brokerage. Bank AI are algorithmic trading operations of banks. Aggregators consolidate liquidity from multiple retail liquidity providers.

[12] The details of order sizes, fill ratio, and order's last digits could be found in Figure 1, Figure 2, and Table 4 in Schmidt (2012) .

Figure 2: EUR-USD, Last Digit Distribution, Pip Tick Size

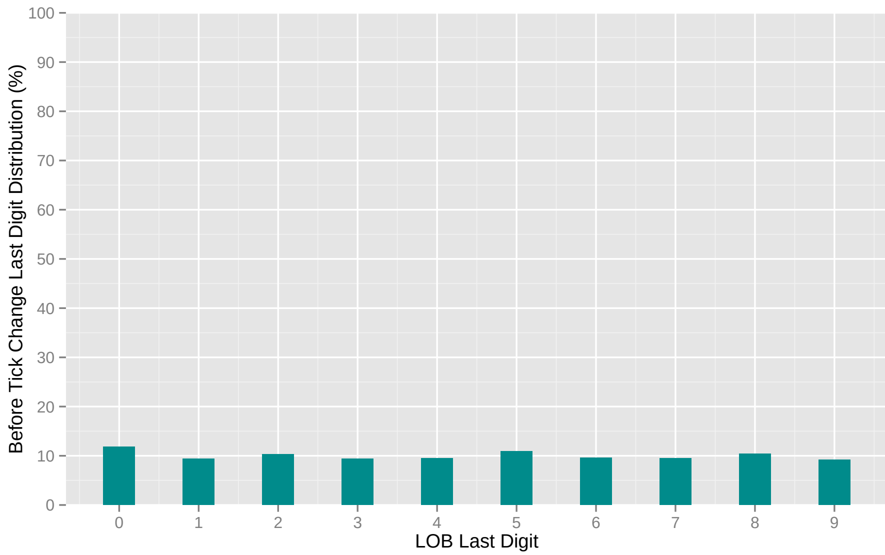

Figure 3: EUR-USD, Last Digit Distribution, Decimal Pip Tick Size

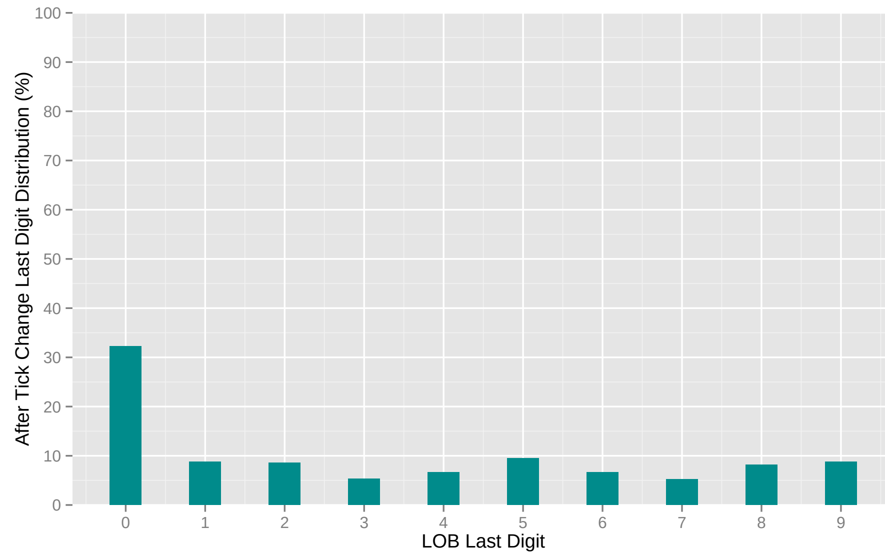

Notes: Figures 2 and 3 show the limit order book prices' last digit distributions under pip and decimal pip pricing, respectively. This distribution was nearly uniform under pip pricing, implying that traders used all digits equally for the last digit of prices. However, human traders did not use the last digit under decimal pip pricing, which means they often used "0" for the fifth digit. This behavior created price clustering at prices ending in zero.

Figure 4: EUR-USD, Volume Distribution, Decimal Pip, Last Digit=0

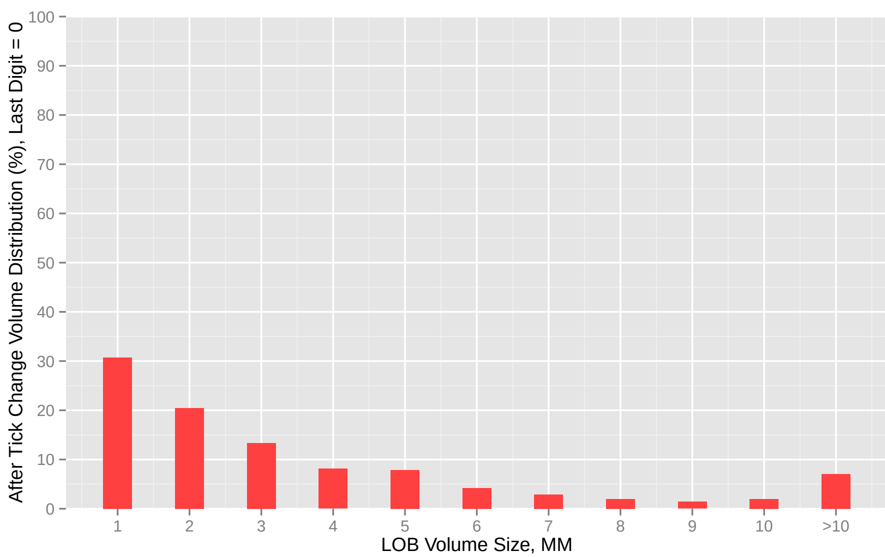

Figure 5: EUR-USD, Volume Distribution, Decimal Pip, Last Digit≠0

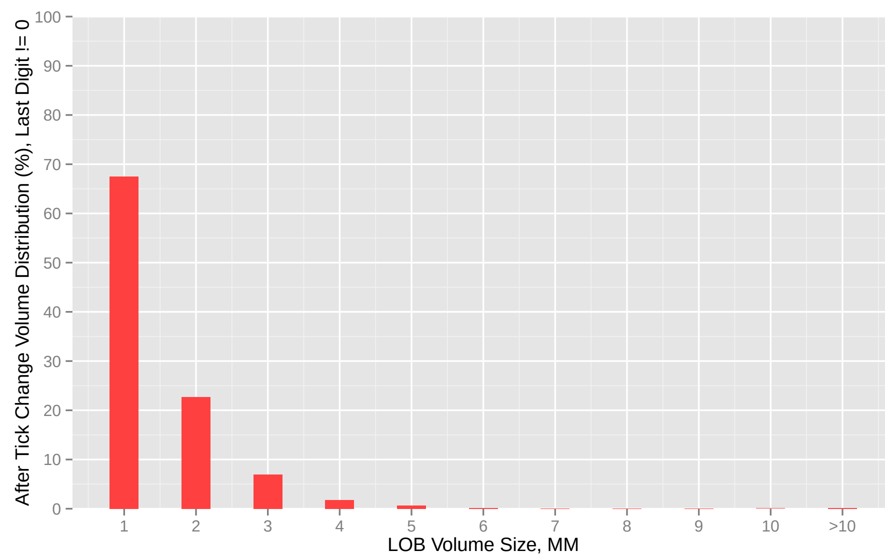

Notes: Because human traders usually use zero for the last digit and place large orders, the volume distributions for prices ending in zero and non-zero digits differ. Figures 4 and 5 show that prices ending in non-zero digits are found in smaller volumes, whereas orders with prices ending in zero are found in larger volumes. This pattern indicates that human traders usually place larger orders at prices ending in zero.

Figures 4 and 5 show that orders with prices ending in non-zero digits come in smaller volumes, whereas orders with prices ending in zero come in larger volumes. A human trader who utilizes a price ending in "0" and trades in larger volumes is a patient market maker, given his high fill ratio. The described structure of the EBS market creates an ideal setting for HFT's to front-run human traders, using the sub-penny jumping strategy. This has been discussed in details in the following section. Sub-penny jumping is not the only strategy implemented by HFT's to front-run other traders. However, the particular order management strategy employed by human traders greatly facilitates the implementation of sub-penny jumping by HFT's.

6. Sub-Penny Jumping

Human traders are slow traders, they submit larger orders at prices ending in zero under decimal pip pricing, and do not cancel their orders very often. In this section, we explain how HFT's use this information to implement the sub-penny jumping strategy to exploit the smaller tick size. Figure 6 depicts possible sub-penny jumping in the EBS market. Suppose that a human trader places a 9 million buy order for EUR/USD at a round price of 1.39940. A sub-penny jumper, who can place an order at very high frequency, would then place a small order, such as a 2 million buy limit at 1.39941. The difference between this price and the human trader's price is equal to the tick size, i.e., a decimal pip. Prior to March 2011, the minimum cost of unwinding against a standing one million limit order was $100 and fell to $10 afterwards. If the sub-penny jumper buys at his placed order, he will move to the sell side and can place a sell order at 1.39945. There are then three possibilities. If another trader places a buy market order at 1.39945, the sub-penny jumper's profit would be $80. If a 2 million buy limit order comes in at a price of 1.39943 (which is higher than 1.39941), the sub-penny jumper could cancel his order and place a sell market order at 1.39943. If the sub-penny jumper is successful in selling his share, then his profit would be $40. However, if the sub-penny jumper thinks that the market is drifting away from his position, he can sell back to the human trader at a price of 1.39940, and his loss would be $20. If there is competition between HFT's to front-run human traders, this would limit potential profits. HFT's also use other order anticipation strategies. We have found empirical evidence that HFT's used the sub-penny jumping strategy frequently under decimal pip pricing.

Figure 6: Sub-Penny Jumping, HT- HFT

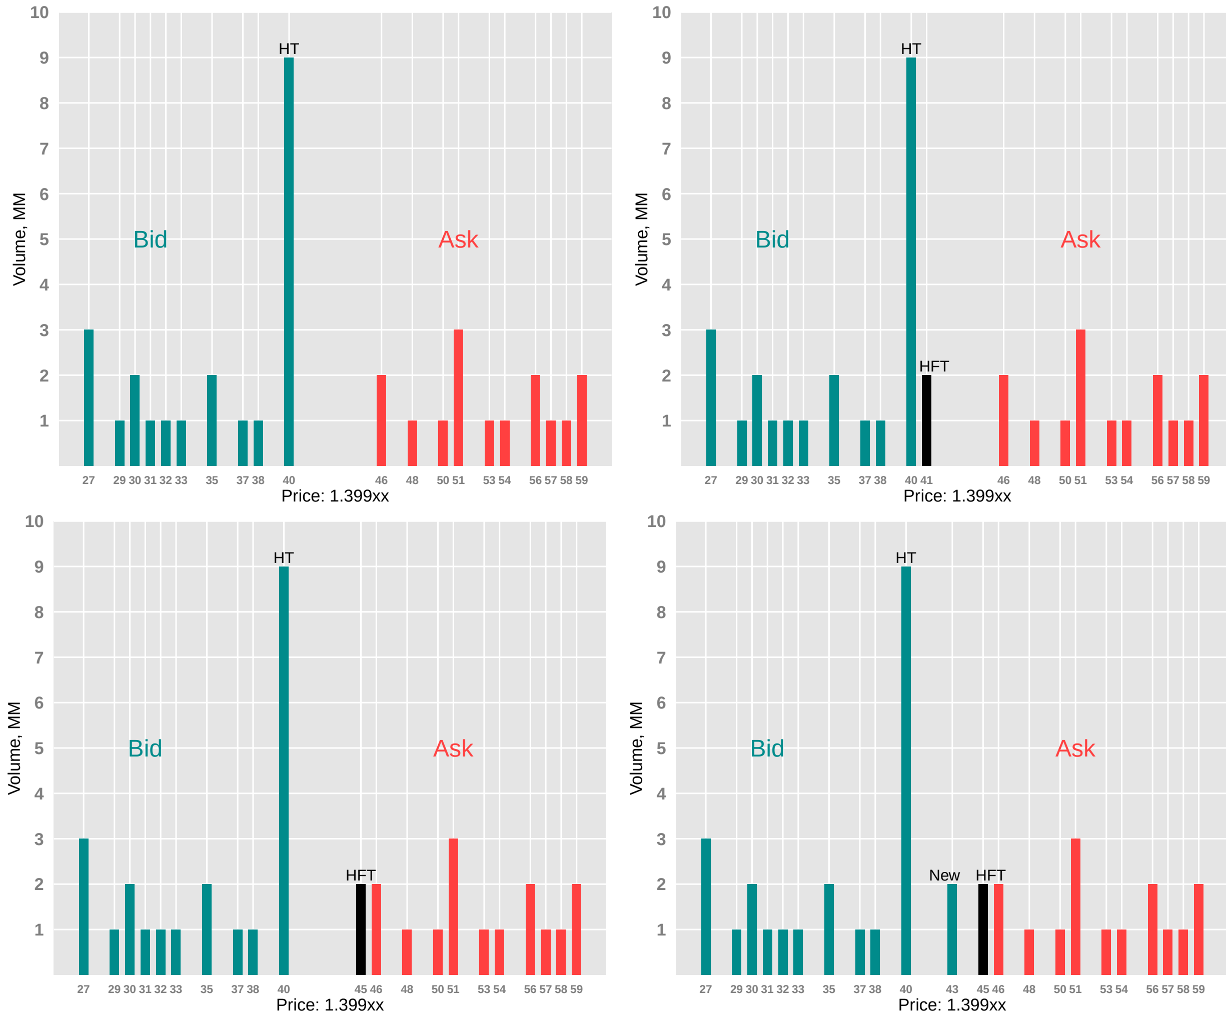

Notes: Figure 6 provides an example of sub-penny jumping strategy implemented by HFT's in EBS market after the introduction of decimal pip pricing. *HT* stands for human trader and *HFT* refers to the high frequency trader.

Let us now consider the general case of sub-penny jumping in the interbank FX market. Suppose that a human trader and a sub-penny jumper place buy orders at p and $p + \tau$ or sell orders at p and $p - \tau$, where τ is the minimum tick size. When the sub-penny jumper buys or sells his order, he will move to the other side of the order book. If the human trader does not cancel or adjust the order and other traders do not fill the human trader's order, then the loss is bounded at the rate of return, $a = \frac{p-(p+\tau)}{p+\tau} = \frac{-\tau}{p+\tau}$ for the buying sub-penny jumper and $b = -\frac{p-(p-\tau)}{p-\tau} = -\frac{\tau}{p-\tau}$ for the selling sub-penny jumper. However, if prices move in favor of the sub-penny jumper, such a favorable price change would be profitable. We will discuss the case of a buying sub-penny jumper, as the selling case is very similar. Suppose that the rate of return has a standard normal distribution: $x \sim N(\mu, \sigma^2)$. This is a simple reduced-form model and not the equilibrium model since the distribution of the return on the foreign currency is assumed to be exogenously given. This is a reasonable assumption given the liquidity of FX markets and the fact that the focus of our analysis is empirical rather than theoretical. Once the sub-penny jumper trades, the orders he front-runs protect him from serious losses. If prices move in his favor, the sub-penny jumper profits to the full extent of the price changes. The returns are unbounded above and limited below by $a = \frac{p-(p+\tau)}{p+\tau} = \frac{-\tau}{p+\tau}$.

Proposition: If the minimum tick size decreases, then the conditional expectation of returns increases and the conditional variance of returns decreases.[13]

$$E(x|x > a) > E(x) \text{ and } \frac{dE(x|x > a)}{d\tau} < 0$$

$$Var(x|x > a) < Var(x) \text{ and } \frac{dVar(x|x > a)}{d\tau} > 0$$

With a smaller tick size, the sub-penny jumper must improve prices by a smaller price increment to trade ahead of the human traders. Therefore, the tick size is the price that the sub-penny jumper must pay to front-run the human traders. However, sub-penny jumpers do not pay this price to the human traders; instead, they pay it to the traders with whom they trade to establish their positions. These traders would have traded with the human traders if the sub-penny jumper had not front-run them. A sub-penny jumper will trade profitably only if the standing orders that they front-run do not cancel their orders, adjusting them quickly, and other traders do not fill the human trader's order. If these options are no longer available, the sub-penny jumper will have difficulty when prices move against him. Based on the EBS market structure presented in Section 3, HFT's know that orders with prices ending in zero and larger order sizes indicate the presence of human traders to front-run. These orders are not canceled or adjusted very often, and due to the larger volume, they cannot be filled quickly. In general, HFT's post inside quotes by front running the best prices, and human traders post outside spread. Then, when the sub-penny jumper trades, if the market starts to move against him, he trades with human trader. This occurs systematically, due to the order management strategy employed by human traders, which differentiates this market from the equity market.

[13]We refer the reader to Appendix A for the proof of this proposition.

7. Sub-Penny Jumping, Deal and Quote Distributions

We study tick sizes in the EBS market in three periods. Tick size was a pip until March 2011, when it was reduced to a decimal pip. It was then changed to a half pip in September 2012. As shown in Figure 7, distributions of orders' last digits were almost uniform with pip tick size for both the bid and ask sides. The x axis shows the last digits ranging from "0" to "9". The y axis indicates the limit order book level, ranging from level "1" to level "10". The term "level" refers to occupied price levels, which implies that the difference between two levels is not necessarily the minimum tick size. For example if the best price is 1.3994 at the best buy (level 1), the next price (level 2) could be 1.3996 where the difference is more than the minimum tick size. Frequency distributions are given on the z axis.

Figure 7: EUR-USD, LOB Last Digit Distributions, Pip Tick Size

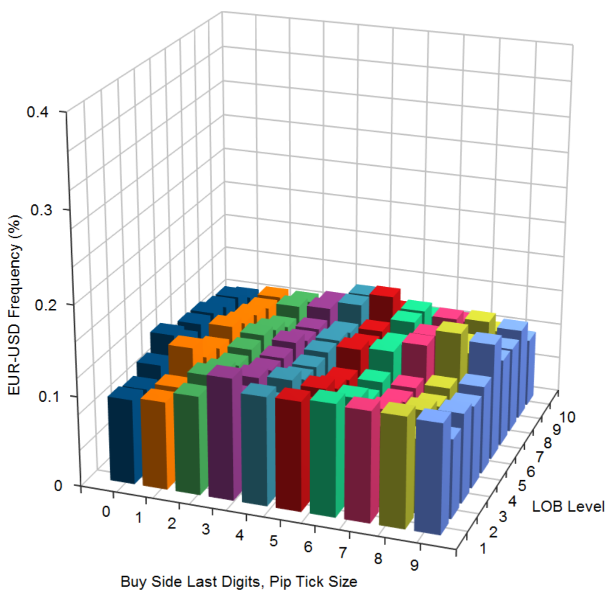

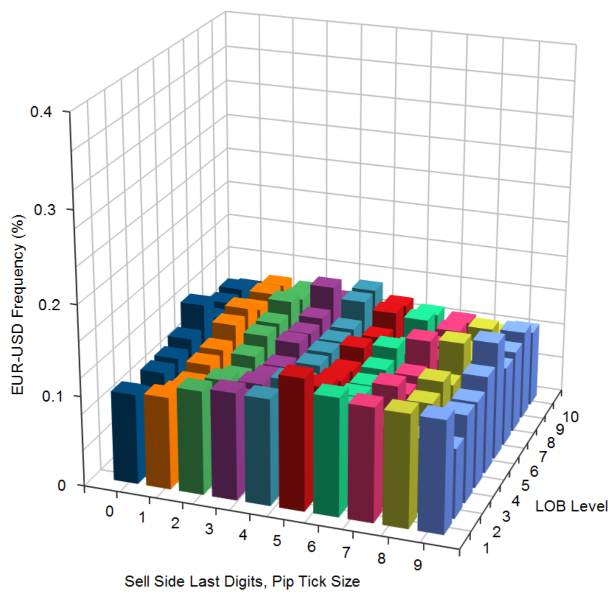

Notes: The price's last digit distributions were almost uniform for all levels of the bid and ask sides of the order book with pip tick size, as shown in Figure 7.

Because human traders were reluctant to use decimal pip pricing, there is price clustering of quotes in all limit order book levels at prices ending in zero. Furthermore, if a buying human trader places an order with a price ending in "0" (e.g., 1.39940), then the buying sub-penny jumper who wants to front-run the human trader should place an order with a price ending in "1" (1.39941). If there are other sub-penny jumpers, they would place their orders with prices ending in "2", "3", etc. However, greater distance from "0" means less expected profit; therefore, we would anticipate a decreasing distribution of the number of orders for the buy side of the limit order book at all levels. Figure 8 shows such distributions for the ten levels of the order book with different last digits. For example, at the level-one buy side of the limit order book, 28% of orders have a last digit of "0", and 18% of the orders have a last digit of "1". This ratio starts to fall and reaches 2.5% for orders with a last digit of "9". The shapes of the distributions at other levels of the bid side are similar to the shape at level-one.

Figure 8: EUR-USD, LOB Last Digit Distributions, Decimal Pip Tick Size

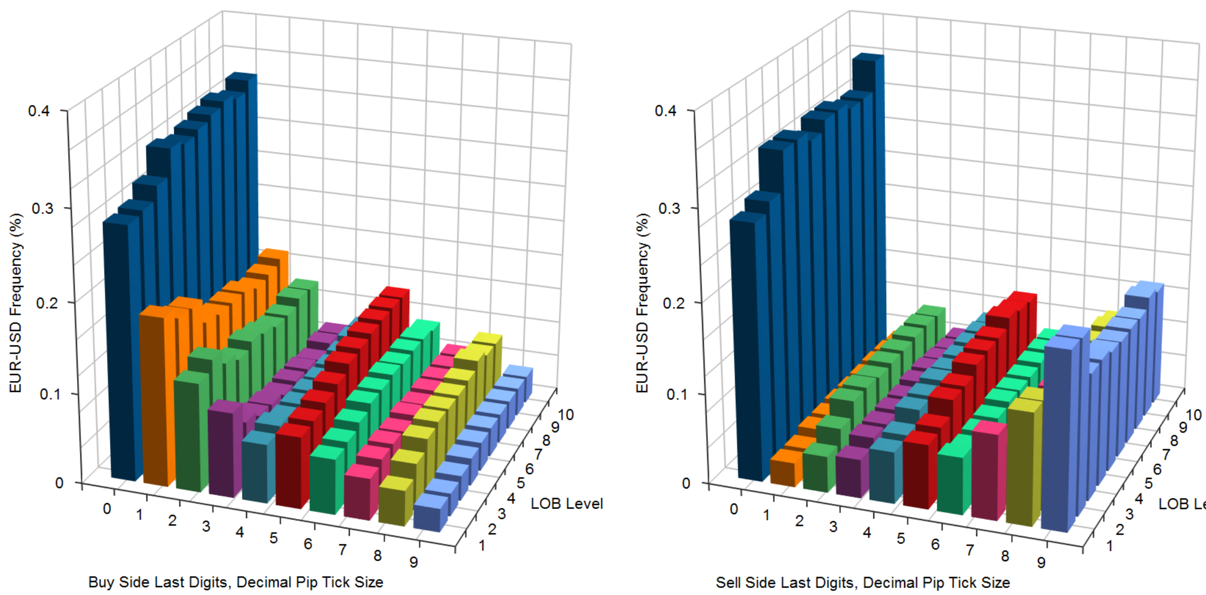

Notes: Figure 8 illustrates the bid and ask prices' last digit distributions for the ten levels of the limit order book with decimal pip tick size.

If a sub-penny jumper improves upon the order of a human trader by one tick while bidding, he would place an order with a last digit of "1". However, if an HFT improves on a human trader's price (e.g., 1.39940) on the ask side, he should place a price ending in "9" (1.39939). Therefore, we would expect an increasing distribution (excluding "0") for all levels of the sell side of the limit order book. For example, in Figure 8, we see that approximately 28% of the orders have a last digit of "0", but only around 2.7% of the orders have a last digit of "1". This ratio starts to rise and reaches 19.5% for orders with a last digit of "9". We observe the same distributions at other levels of the ask side. Distributions for decimal pip pricing also show weaker price clustering at prices ending in "5". It appears that some human traders placed such orders to adopt decimal pip pricing partially without the complexity of using all the digits. However, the volume sizes used with this digit are smaller, and there was no opportunity for HFT's to front-run the digit "5". The different patterns on the ask and buy sides provide evidence of sub-penny jumping in the EBS market under decimal pip pricing. When the tick size changed to a half pip, traders had only two options, "0" and "5", for the last digit of their orders. If the last digit of the best order is "0", then the last digit of the second best would be "5". Consequently, we observe more clustering at "0" at odd levels in Figure 9.

Figure 9: EUR-USD, LOB Last Digit Distributions, Half Pip Tick Size

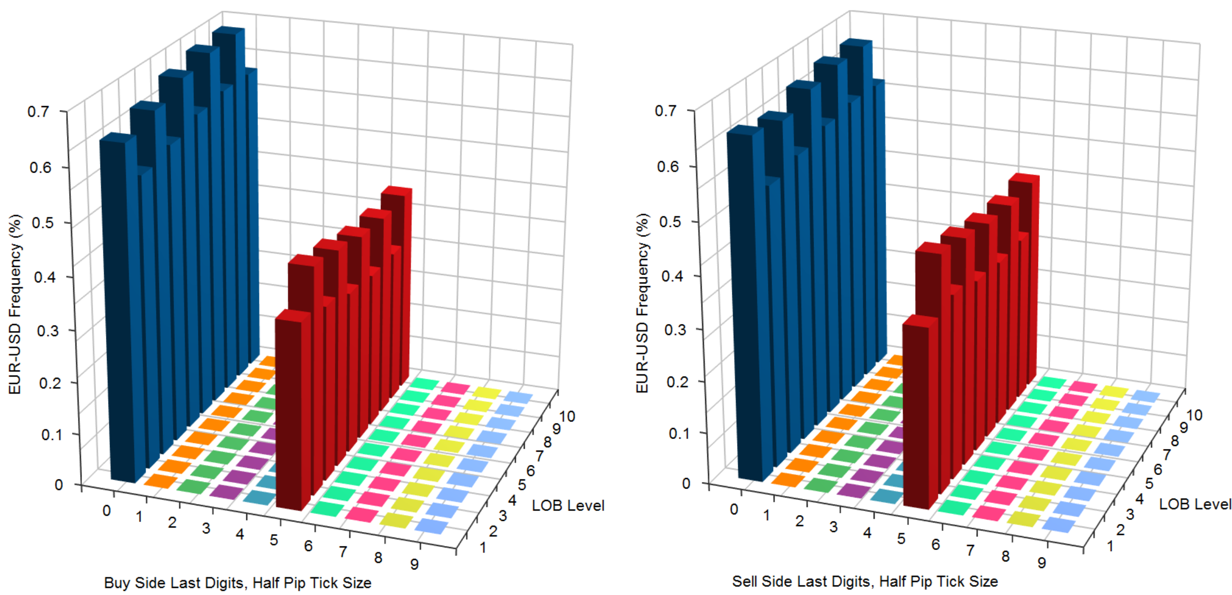

Notes: Figure 9 shows the bid and ask prices' last digit distributions for the ten levels of the limit order book with half pip tick size

Volume distributions also provide evidence of sub-penny jumping. When HFT's front-run human traders, they secure their positions by placing orders with smaller sizes. We know from Section 3 that HFT's typically use one million orders and almost never use an order size exceeding four million. As a result, we should observe smaller volumes at the top of the order book under decimal pip pricing. The order book volume distributions under pip and decimal pip pricing are given in Figures 10 and 11, respectively. At level one of the buy side under pip pricing, 16% of orders had a size of one million, and 11% had a size of more than ten million. However, under decimal pip pricing, 61% of orders had a size of one million, and only 0.7% of orders had a size of more than ten million. Changes at the other levels are more significant. At level two with pip pricing, 3% of orders had a size of one million, and 44% had a size of more than ten million. However, under decimal pip pricing, 62% of orders had a size of one million, and only 1% of orders had a size of more than ten million. The changes on the sell side of the order book are similar to those on the bid side.

Figure 10: EUR-USD, LOB Volume Distributions, Pip Tick Size

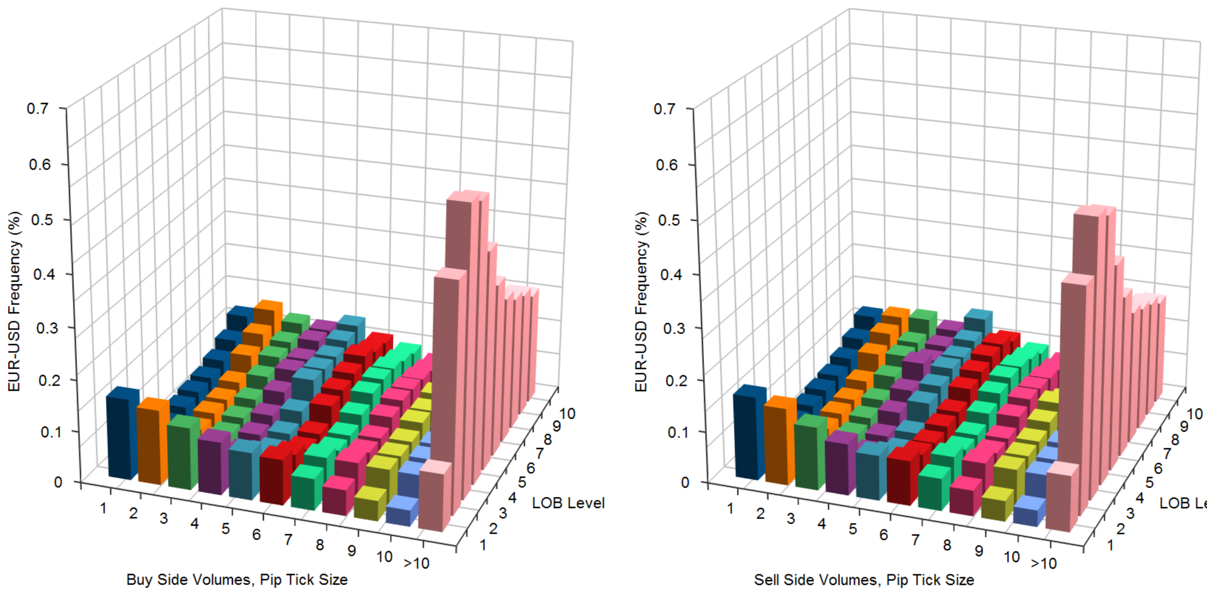

Notes: Figure 10 illustrates the order book volume distributions with pip tick size change. The existence of human traders at the top of the book provides larger volumes.

Figure 11: EUR-USD, LOB Volume Distributions, Decimal Pip Tick Size

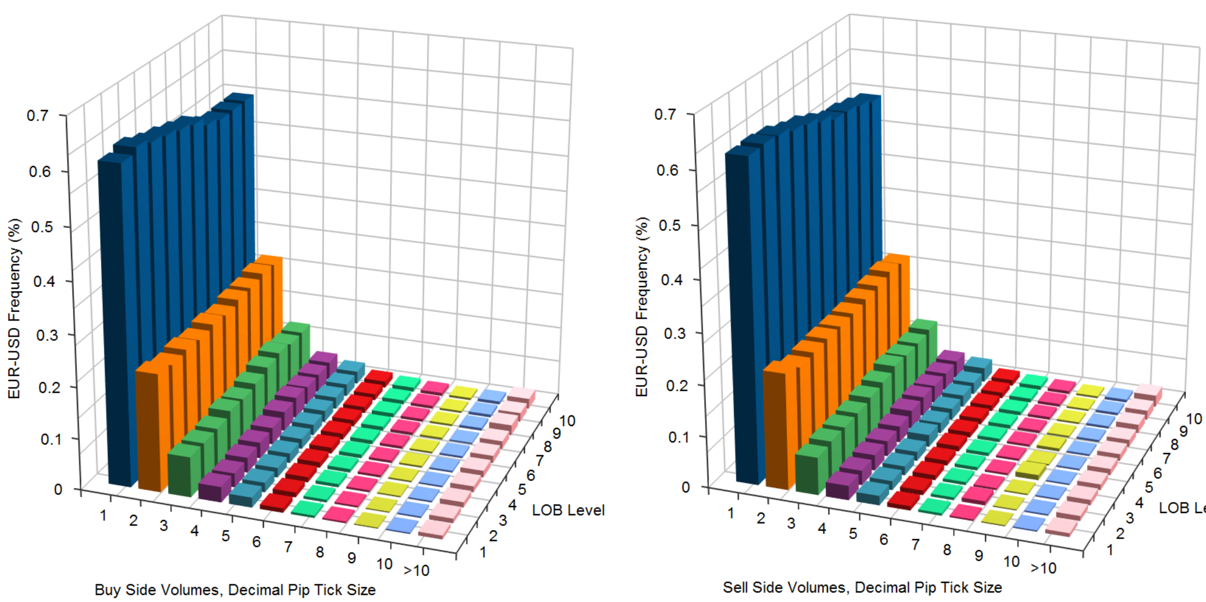

Notes: The order book volume distributions with decimal pip tick size are given in Figure 11. HFT's have occupied the top of the order book by smaller volumes.

Comparing Figures 10 and 11, we observe that large volumes nearly disappeared from the top of the order book. Along with other reasons, some human traders might also use smaller orders with decimal pip pricing which could partially explain the change in the order book. However, based on Section 5, human traders did not stop using larger orders.

[Figure 12](#) indicates that there are more large volumes at the top of the order book after reversion to half pip pricing in September 2012 comparing to decimal pip pricing. The strong symmetry between the ask and bid sides is an interesting empirical finding.

Figure 12: EUR-USD, LOB Volume Distributions, Half Pip Tick Size

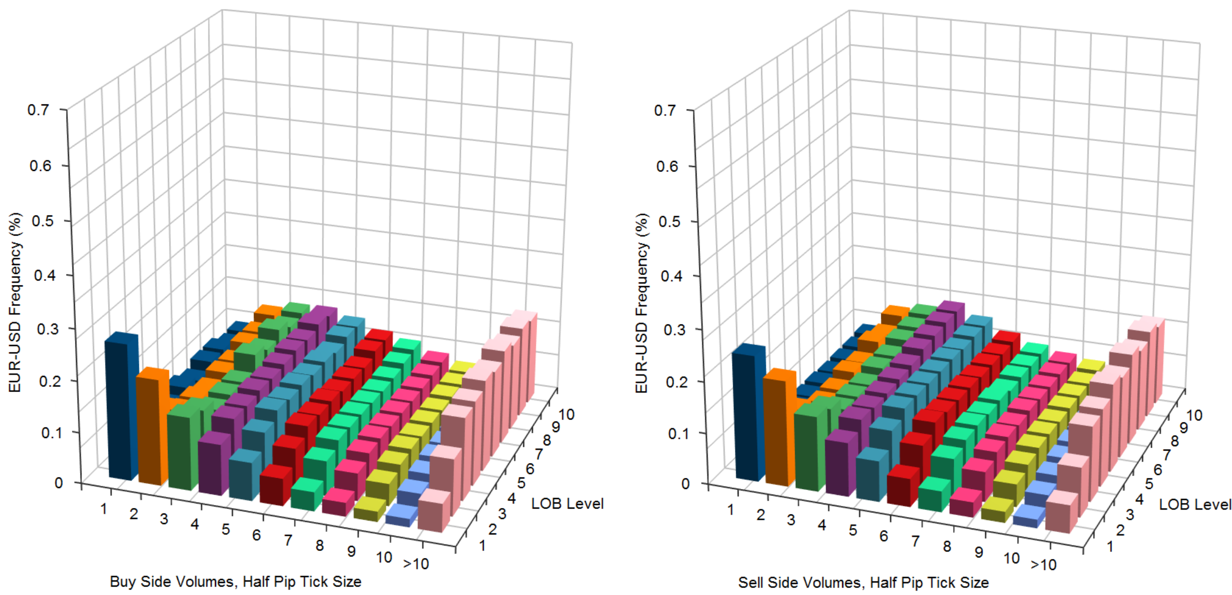

Notes: Figure 12 provides the order book volume distributions with half pip tick size.

Patterns similar to those of the order books are present in the deal prices. We examine the trades and split them into trades initiated by buyers and sellers. Figure 13 shows the distribution of the last digit of deal prices for buyer-initiated and seller-initiated transactions for different tick sizes. The x axis shows the last digits of deal prices ranging from "0" to "9". The y axis indicates tick sizes: pip, decimal pip, and half pip. Frequency distributions are given on the z axis. The distributions of the last digits of deal prices are nearly uniform under pip pricing for both buyer- and seller-initiated deals. This is seen in the uniform order book last digit distribution in Figure 7. This distribution is increasing for buyer-initiated trades and decreasing for seller-initiated trades under decimal pip pricing, a finding that is consistent with the decreasing buy side last digit distribution and the increasing sell side last digit distribution shown in Figure 8. There is also price clustering at prices ending in "5" under decimal pip pricing. More deals occur at prices ending in "0" under half pip pricing. The patterns for half pip tick size are consistent with Figure 9. The shape of the order book has also changed the distributions of deal volumes. Figure 14 shows that in buyer-initiated trades, 57% of deals occurred with a volume of one million under pip pricing, 72% under decimal pip pricing, and 65% under half pip pricing. We observe the same pattern in seller-initiated trades. Figure 14 shows that more deals occurred with larger volumes under pip pricing, as they were available at the top of the order book.

Figure 13: EUR-USD, Deal Last Digit Distributions with Different Tick Sizes

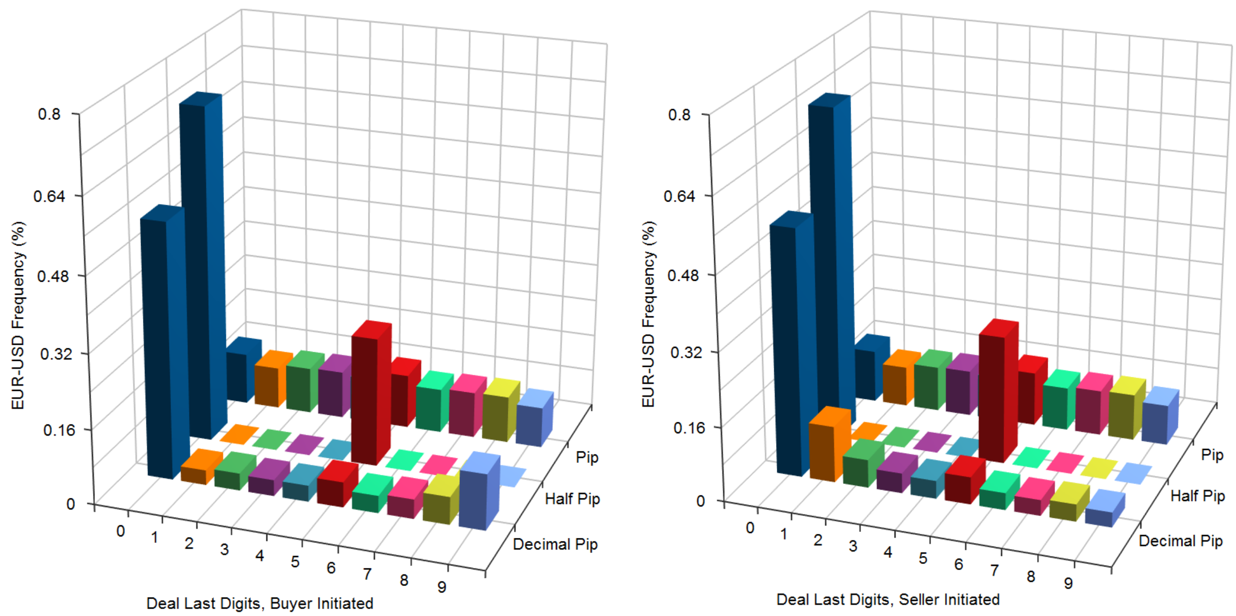

Notes: The deal last digits' distributions for pip, decimal pip, and half pip are given in Figure 13 for both buyer and seller initiated deals.

Figure 14: EUR-USD, Deal Volume Distributions with Different Tick Sizes

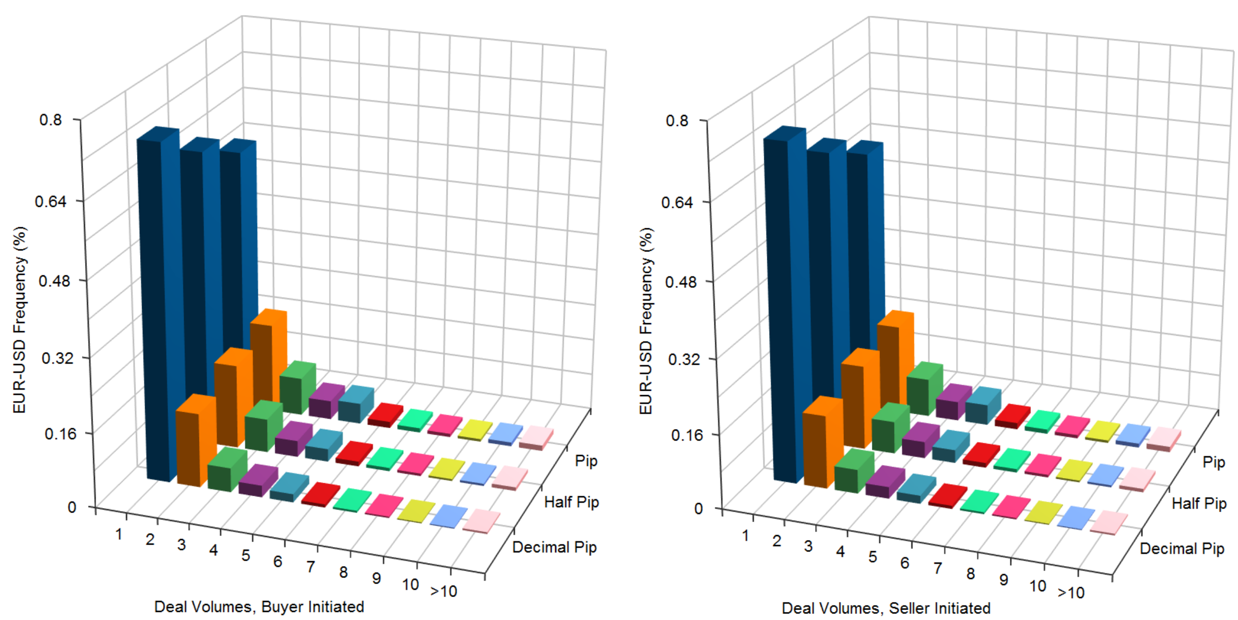

Notes: The deal volumes distributions for pip, decimal pip, and half pip are given in Figure 14 for both buyer and seller initiated deals.

8. Market Quality

In this section we study the maker quality measures, spread, realized spread and market depth. The change in the market maker-market taker compositions also changed the permanent price impact measured by VAR model of Hasbrouck (1991).

8.1. Spread & Realized Spread

Spread, defined as the difference between the best ask and the best bidding prices, is a measure of liquidity cost. Table 2 provides summary statistics for spread under different tick sizes. The average spread, which was *0.00015* under pip pricing, decreased to *0.00013* under decimal pip pricing. Interestingly, it further decreased to *0.00010* when EBS adopted a half pip tick size, which is considered a larger tick size than decimal pip tick size.

Table 2: Spread Summary Statistics

	Min	Max	Median	Mean	Std.dev
Pip (p)	0.00010	0.00250	0.00010	0.00015	0.00006
Decimal Pip (dp)	0.000010	0.00400	0.00012	0.00013	0.00007
Half Pip (hp)	0.000050	0.00995	0.00010	0.00010	0.00005

We test the following hypotheses regarding the spread mean under different scenarios. All test are strongly rejected with all *p-values* less than *0.00001*.

$$\begin{cases} H_0 & \mu_{dp} \geq \mu_p \\ H_a & \mu_{dp} < \mu_p \end{cases} \qquad \begin{cases} H_0 & \mu_{hp} \geq \mu_p \\ H_a & \mu_{hp} < \mu_p \end{cases} \qquad \begin{cases} H_0 & \mu_{hp} \geq \mu_{dp} \\ H_a & \mu_{hp} < \mu_{dp} \end{cases}$$

Figure 15 indicates whether the spread was binding under pip and decimal pip pricing.[14] The results show that on average, the spread was binding (equal to *0.0001*) more than 50% of the time under pip pricing. This means that there was pressure on the spread to decrease, which could partially justify the adoption of decimal pip tick size. Figure 15 also illustrates that approximately 1% of the time, the spread was binding and equal to *0.00001* under decimal pip pricing. This suggests that the decimal pip tick size was too small for EUR/USD in the EBS market. Empirically, this means that the optimal tick size for EUR/USD should be between a pip and a decimal pip, namely, a half pip, where the fifth digit can only be

[14] If the tick size is τ and the spread is equal to τ, spread is binding.

"0" or "5". In Figure 16, we demonstrate the binding spreads with half pip tick size. On average, between 15% and 20% of the time, the spread was binding. Compared with other tick sizes, the binding spread under half pip pricing is not as high as 50% or as low as 1%.

Figure 17 shows the one-hour average spread for 2011. In the graph, each dot represents a one-hour average spread, and we have depicted all 24 observations (hours) across the day. The gaps represent weekends which we have excluded. When EBS changed the tick size from pip to decimal pip in March 2011, indicated by the dashed line in the graph, there was a significant draw down in the spread. Technically, traders had more options inside the spread, which was under greater pressure before. Figure 18 provides the one-hour average spread for January-Feburary 2013 under half pip tick size.

To test the effects of the lower tick size on the spread, we use difference-in-difference (DID) estimation. DID is typically used to identify the effects of a specific policy intervention or treatment. The idea behind the DID approach is that if an intervention has an effect, the difference between the unaffected group (the control group) and the group directly affected by the intervention (treatment group) should change after the policy intervention. Then, one compares the difference in outcomes between the two groups before and after the intervention. We chose EUR/USD as a treatment group to test our hypothesis. The best control group would be the same currency pair in the Reuters market on the condition that the spread in Reuters was not affected by the tick size change in the EBS market. Unfortunately, we do not have access to this data set. Our next best options are the next most commonly traded currency pairs in the EBS market, namely, EUR/GBP, AUD/USD and GBP/USD. The tick size did not change for these currency pairs in the time period under consideration. The control groups are not perfect since these currency pairs are primarily traded at the Reuters, and hence may have low liquidity at EBS. We could not study the change in the spread from decimal pip to half pip because, as shown in Figure 1, EBS simultaneously changed the tick size for both our treatment and control groups. We have considered other factors that may have changed in or around March 2011 that may potentially impact our analysis. One of those factors is changes in the tick size on other venues. Furthermore, If the HFT's moved activities from the control to the treatments groups, the control groups would have been affected indirectly by the tick size change. We did not find significant changes in the deals and the state of the limit order book of the control groups.

Figure 15: EUR-USD Binding Spread, Pip & Decimal Pip

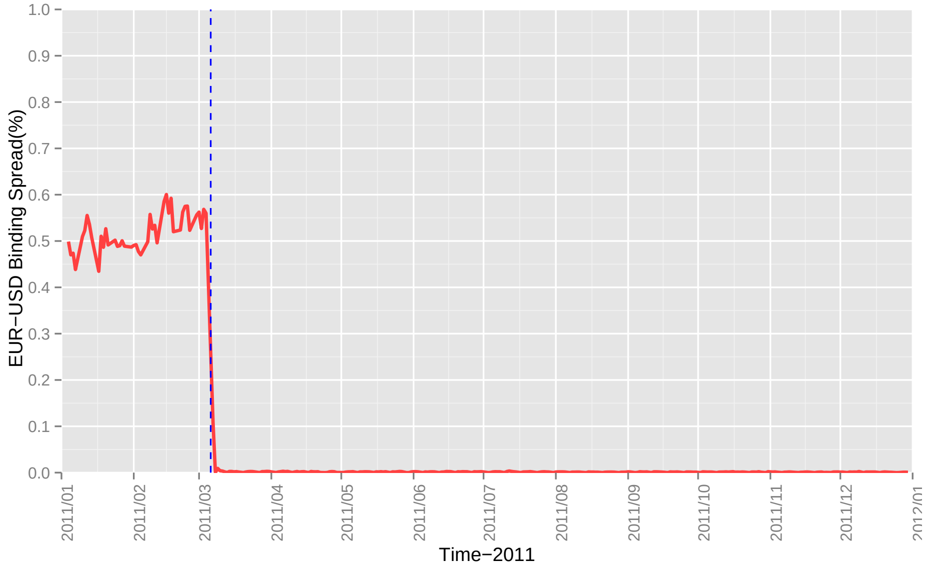

Figure 16: EUR-USD Binding Spread, Half Pip

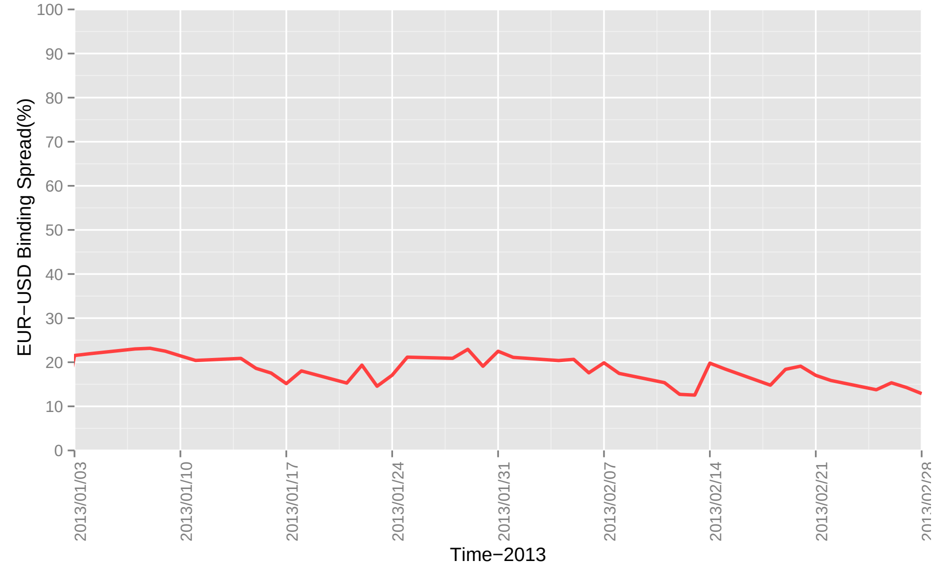

Notes: Figure 15 illustrates the binding EUR/USD spread under pip and decimal pip tick sizes. The dashed line indicates the time of the tick size change in March 2011. Figure 16 shows the binding spread for the same currency under half pip tick size.

Figure 17: EUR-USD Spread, Pip & Decimal Pip Tick Size

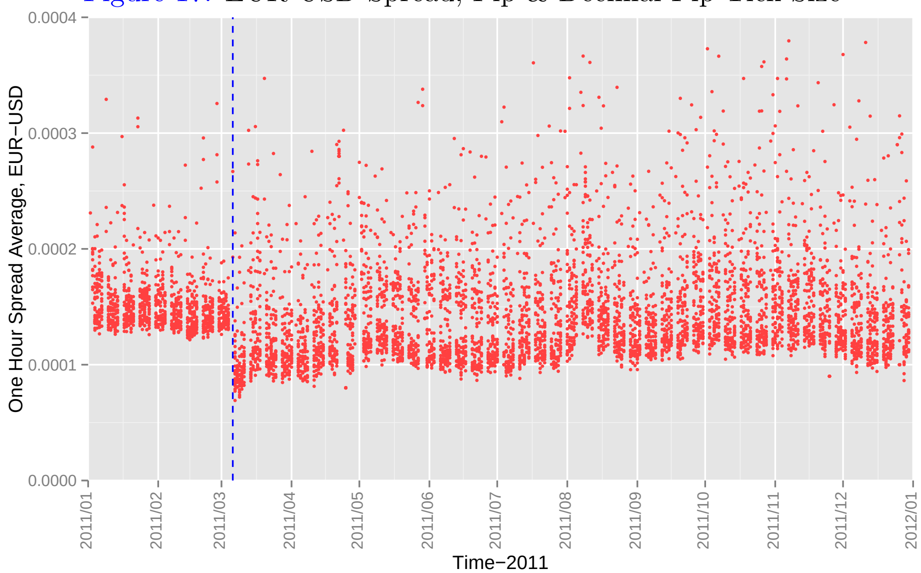

Figure 18: EUR-USD Spread, Half Pip Tick Size

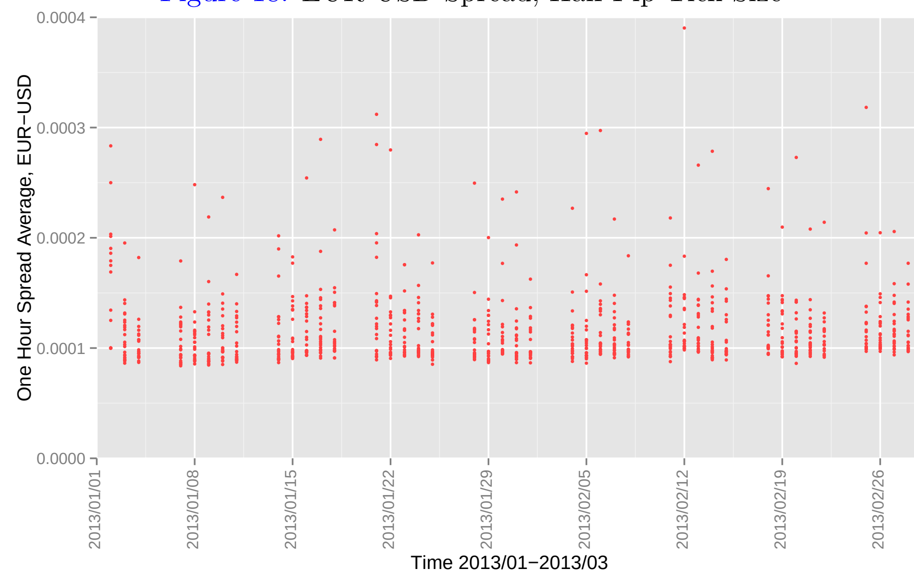

Notes: Figure 17 illustrates the one-hour frequency EUR/USD spread under pip and decimal pip tick sizes for 2011. Figure 18 illustrates the one-hour frequency EUR/USD spread under half pip tick size for 2013.

Assuming that the treatment effect is stationary over time, we define the DID model by:

$$S_t = \beta_0 + \beta_1 P + \beta_2 T + \beta_3 PT + \varepsilon_t \tag{1}$$

where T is a treatment variable equal to one for EUR/USD and zero for the control group; P is the post-treatment indicator, equal to one after the tick size change and zero before the change. T controls for permanent differences between the treatment and control groups, with β_2 capturing this variation. Similarly, P controls for trends common to both the control and treatment groups, and β_1 captures this variation. The variation that remains is captured by β_3. Conditional means corresponding to the four combinations of T and P produce Table 3.

Table 3: Conditional Mean Estimates from the DID Regression Model

	After Tick Size Change	Before Tick Size Change	Difference
Treatment	$\beta_0 + \beta_1 + \beta_2 + \beta_3$	$\beta_0 + \beta_2$	$\beta_1 + \beta_3$
Control	$\beta_0 + \beta_1$	β_0	β_1
Difference	$\beta_2 + \beta_3$	β_2	β_3

The DID regression results for the one-hour EUR/USD average spread are provided in Table 4 with different control groups. In all cases, the estimates of β_3 are negative and significant at the one percent level. We also used the one-minute average spread frequency in robustness checks. The results, presented in Table 5, show that the estimates of β_3 are significant at the one percent level. The estimates are close to the difference between the average spread, with the pip and decimal pip equal to *-0.00002*. The DID regression results indicate that the spread decreased after the tick size change due to the lower minimum tick size. Reduction in the minimum tick size from pip to decimal pip enables HFT's to post more aggressively priced limit orders in the EBS market, leading to tighter quoted spreads. However, by reducing the cost of implementing sub-penny jumping strategies, HFT's occupied the top of the order book and the benefit of the reduction in the spread was mostly absorbed by them. We then estimate the realized spread defined as:

$$RS_t = 2q_t(p_t - m_{t+s}) \tag{2}$$

Where p_t is the deal price at time t, m_{t+s} is the midpoint at time $t+s$ and q_t is the trade indicator, equal to 1 for buyer-initiated trades and -1 for seller-initiated trades. We chose the $s = 5$ seconds window for the midpoint.

Table 4: Difference-in-Difference Regressions of EUR-USD Spread, 1 Hour Frequency

Control Group	EUR-GBP (T=4,236)				AUD-USD (T=4,012)				GBP-USD (T=3,670)			
	Estimate	Std. Error	t-value	Pr(>|t|)	Estimate	Std. Error	t-value	Pr(>|t|)	Estimate	Std. Error	t-value	Pr(>|t|)
β_0	0.000262	$2 * 10^{-6}$	128.77	$2 * 10^{-16}$	0.000224	$2 * 10^{-6}$	119.82	$2 * 10^{-16}$	0.000348	$2 * 10^{-6}$	146.58	$2 * 10^{-16}$
β_1	0.000015	$3 * 10^{-6}$	5.65	$2 * 10^{-8}$	0.000013	$2 * 10^{-6}$	5.77	$8 * 10^{-9}$	-0.000012	$3 * 10^{-6}$	-4.06	$5 * 10^{-5}$
β_2	-0.000113	$2 * 10^{-6}$	-53.69	$2 * 10^{-16}$	-0.000076	$2 * 10^{-6}$	-38.78	$2 * 10^{-16}$	-0.000203	$2 * 10^{-6}$	-83.65	$2 * 10^{-16}$
β_3	-0.000030	$3 * 10^{-6}$	-10.32	$2 * 10^{-16}$	-0.000032	$2 * 10^{-6}$	-13.19	$2 * 10^{-16}$	-0.000007	$3 * 10^{-6}$	-2.17	$3 * 10^{-2}$

All coefficients are significant at the 1 percent level.

Table 5: Difference-in-Difference Regressions of EUR-USD Spread, 1 Minute Frequency

Control Group	EUR-GBP (T=250,880)				AUD-USD (T=242,544)				GBP-USD (T=219,448)			
	Estimate	Std. Error	t-value	Pr(>|t|)	Estimate	Std. Error	t-value	Pr(>|t|)	Estimate	Std. Error	t-value	Pr(>|t|)
β_0	0.000258	$3 * 10^{-7}$	918.53	$2 * 10^{-16}$	0.000219	$2 * 10^{-7}$	946.95	$2 * 10^{-16}$	0.000340	$3 * 10^{-7}$	1048.10	$2 * 10^{-16}$
β_1	0.000015	$4 * 10^{-7}$	40.00	$2 * 10^{-16}$	0.000016	$3 * 10^{-7}$	55.43	$2 * 10^{-16}$	-0.000007	$4 * 10^{-7}$	-16.58	$2 * 10^{-16}$
β_2	-0.000110	$3 * 10^{-7}$	-364.32	$2 * 10^{-16}$	-0.000072	$3 * 10^{-7}$	-279.44	$2 * 10^{-16}$	-0.000196	$3 * 10^{-7}$	-572.04	$2 * 10^{-16}$
β_3	-0.000031	$4 * 10^{-7}$	-74.52	$2 * 10^{-16}$	-0.000034	$3 * 10^{-7}$	-101.64	$2 * 10^{-16}$	-0.000011	$5 * 10^{-7}$	-24.25	$2 * 10^{-16}$

All coefficients are significant at the 1 percent level.

Notes: Table 4 provides the difference-in-difference estimation results for the one-hour EUR/USD spread average. We have used EUR/GBP, AUD/USD and GBP/USD as control groups because they are the busiest currency pairs after the major currency pairs. Activities in other currency pairs are usually sparse, which makes them inappropriate for use in our control group. The estimates of β_3 are negative and significant at the one percent level, which indicates that the spread decreased after the introduction of decimal pip pricing. Table 5 shows the DID results for the one-minute EUR/USD spread average. All β_3 coefficients are negative and significant at the one percent level.

In case of buyer-initiated trade, a positive realized spread means that the market maker can unwind his short position to make a profit. A negative realized spread means that he has suffered an adverse price move and must revise his estimate of the fundamental value accordingly. Similar reasoning applies to seller-initiated trades. The absolute value of realized spread has decreased toward zero meaning that the market maker would suffer less adverse price move. Figure 19 provides the realized spread under pip and half pip pricing in 2011. The graph shows the one-hour average realized spread of EUR/USD exchange rates. Figure 20 provides the realized spread under half pip pricing in 2013.

8.2. Market Depth

Market depth is the amount available in the limit order book. This quantity could also be interpreted as the size an order must reach to move the market's best available price by a given amount. Generally, traders prefer a deep market because then there will be less price impact. In analyzing both the bid and ask sides of the market, we have calculated EUR/USD market depth of *0.0001* and *0.0002* under pip, decimal pip and half pip pricing. This means how much volume is available between and including the best bid, or sell, and the next price with *0.0001* and *0.0002* differences. Figure 21 shows the daily averages of ask side depth. As before, the dashed line indicates the tick size change from pip to decimal pip in March 2011. Each circle indicates the market depth of *0.0001*, and each triangle shows a market depth of *0.0002*. For example, orders of *$15-$20* million were necessary to move the best bid by *0.0001* under pip pricing before March 2011. However, an overall order size of *$10-$15* million was sufficient to move the best price by the same amount after the introduction of decimal pip pricing. Figure 21 shows that the introduction of decimal pip pricing reduced market depth significantly. We found similar results for the buy side of the order book.

There are some reasons why market depth worsened after the introduction of decimal pip pricing. First, the HFT's implemented the sub-penny jumping strategy and occupied the top of the order book with smaller volumes. Second, some human traders may also have switched from larger to smaller orders to adapt to the tick size change. When EBS changed the tick size to a half pip, there was a significant improvement in market depth, as shown in Figure 22.

Figure 19: EUR-USD Realized Spread, Pip & Decimal Pip

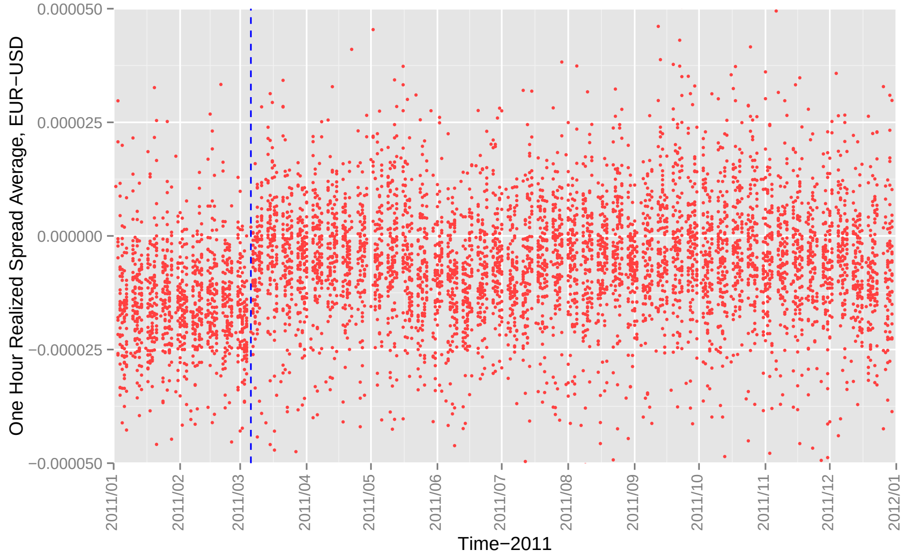

Figure 20: EUR-USD Realized Spread, Half Pip

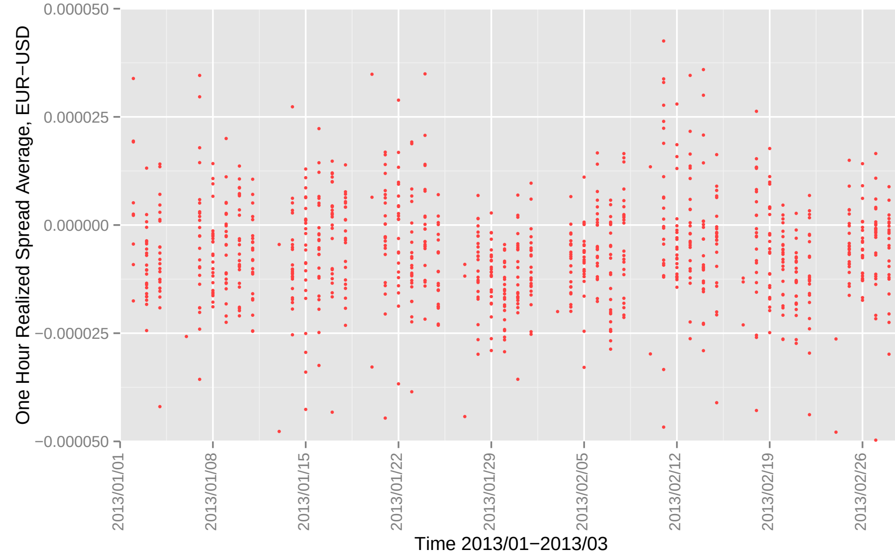

Notes: Figure 19 indicates the one hour realized spread average under pip and decimal pip pricing in 2011. Figure 20 shows the realized spread under half pip pricing in 2013.

Figure 21: EUR-USD Market Depth, Pip & Decimal Pip

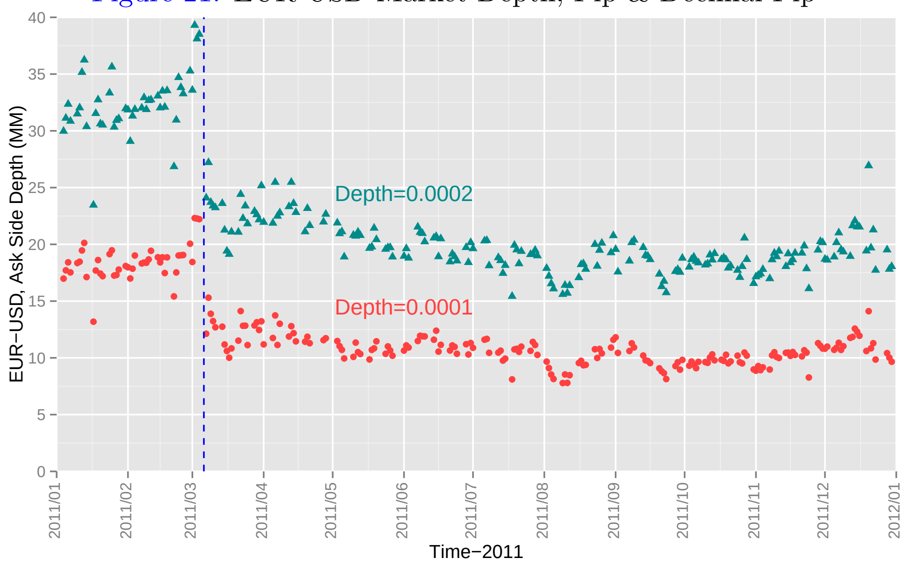

Figure 22: EUR-USD Market Depth, Half Pip

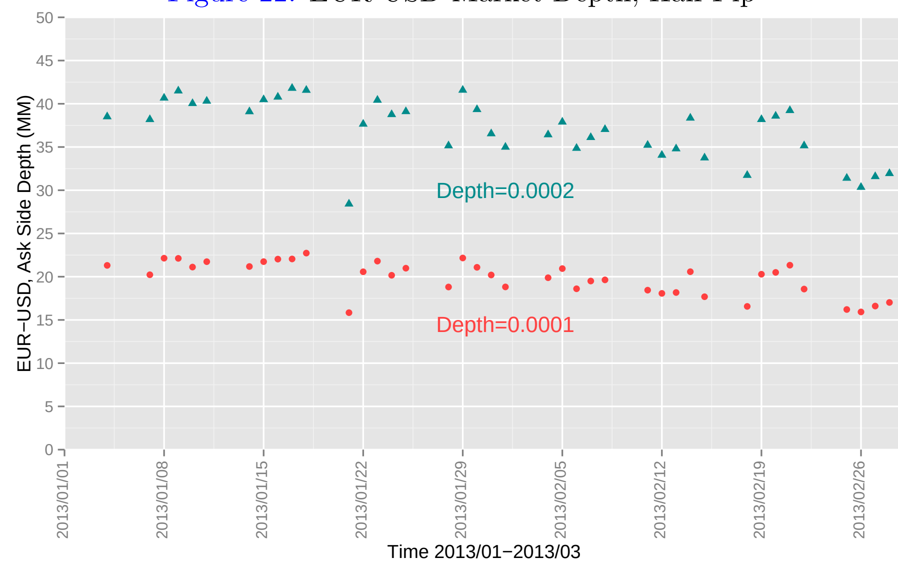

Notes: Figure 21 shows EUR/USD daily average buy side depth under pip and decimal pip pricing in 2011. Each circle shows which order size is necessary to move the best price by *0.0001*, and each triangle shows a market depth of *0.0002*. Figure 22 illustrates sell side depth under half pip pricing in 2013.

8.3. Maker-Taker Structure

In this section, we use the transaction data broken down into categories specifying the human (H) and computer (C) traders. The data provides eight categories HH_b, HH_s, HC_b, HC_s, CH_b, CH_s, CC_b, and CC_s, where b denotes buyer initiated trade and s denotes seller initiated trade. By adding both buyer and seller initiated volumes, we construct the four groups: human maker-human taker HH, computer maker-human taker CH, human maker-computer taker HC, and computer maker-computer taker CC. This enables us to investigate if there is a change in market maker-market taker composition between humans and algorithmic traders after the introduction of the decimal pip pricing.

Figure 23 shows the human maker-human taker -HH- and computer maker-computer taker -CC- ratios. The trade between only humans decreased from about 15% under the pip tick size to about 10% under the decimal pip tick size. Meanwhile, the trade between only computers has increased gradually from 30-35% to 40-45%. The decreasing trend of HH means that humans trade less with each other, or, they have less access to each other in the market. At the same time, increasing CC indicates that computers trade with each other more often. One possible explanation for the mentioned observations is as follows: lower minimum tick size enabled HFT's to implement sub-penny jumping strategies and post more aggressively limit orders. As a result, HFT's occupied the top of the order book and trade more with each other (higher CC). Human traders did not adopt the decimal pip pricing as much as HFT's and consequently were placed more at the back of the order book. This means that HFT's lowered the execution probability of human maker-human taker trades (lower HH).

Figure 24 provides the computer maker-human taker -CH- and human maker-computer taker -HC- ratios. The CH decreased from about 25% under pip pricing to 15-20% under decimal pip pricing. At the same time period, HC had an increasing trend. The volumes of CH trades did not change considerably after the tick size change. With higher total volume after the tick size change, the decline in CH share is therefore due to the fact that the shares of other groups have increased overall. The findings regarding HC needs more careful considerations. Our results in previous sections and Schmidt (2012) indicate that the manual traders were placed more at the back of the order book, indicating possibly a decreasing trend for HC.

Figure 23: EUR-USD, HH and CC

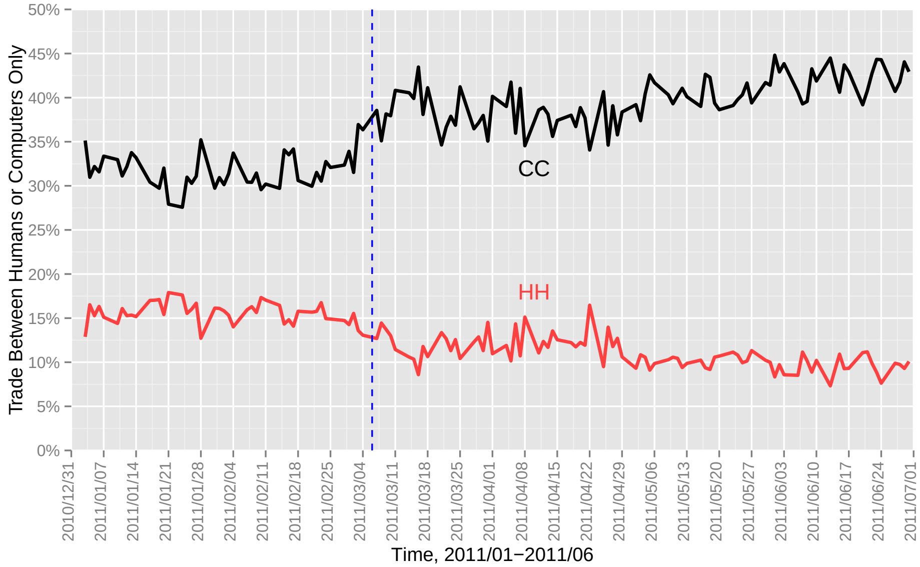

Figure 24: EUR-USD, HC and CH

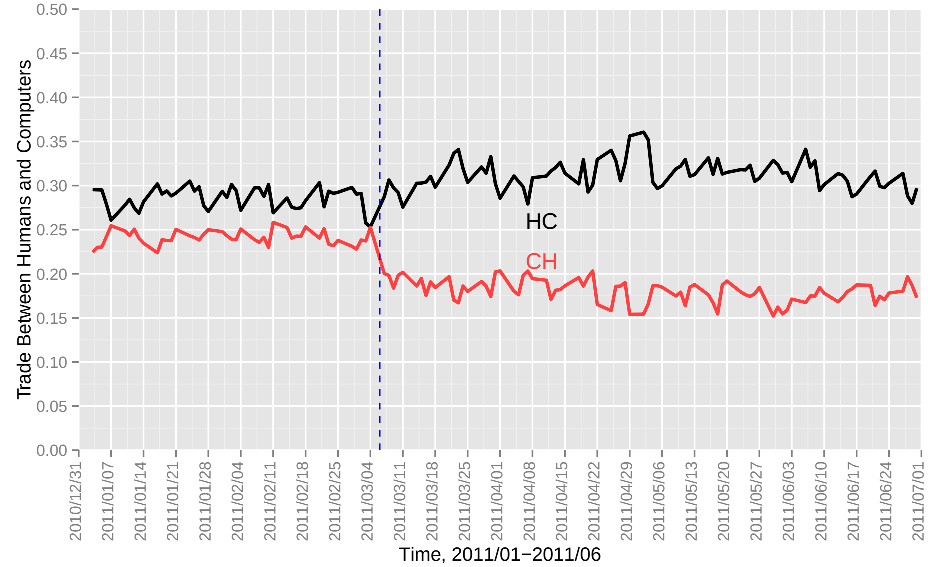

Notes: Figure 23 shows the human maker-human taker -HH- and computer maker-computer taker -CC- ratios. Figure 24 show the computer maker-human taker -CH- and human maker-computer taker -HC- ratios

Figure 25: Lag Selection for the Order Flow

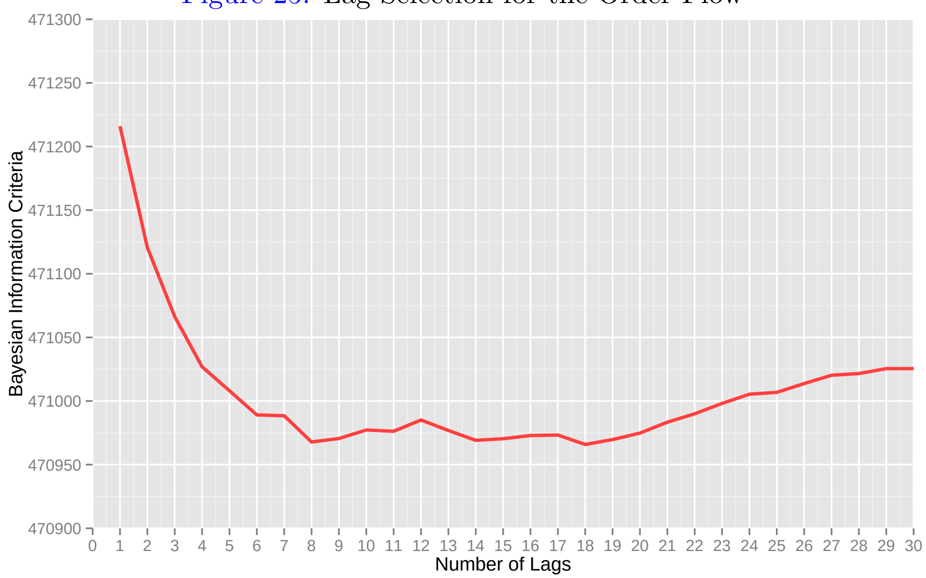

Figure 26: Lag Selection for the Returns

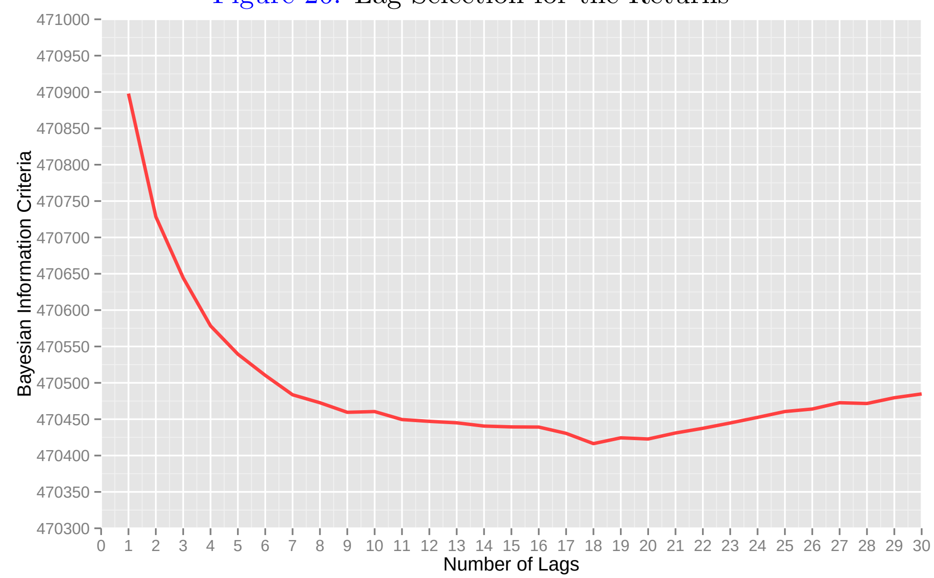

Notes: Figures 25 and 26 represent the values of the Bayesian Information Criteria (BIC) for the lag selection of the order flow and the returns in Equation (4).

One possible answer for increasing HC is that the manual traders' resting orders get hit when the price moves, and then they become the market maker because they are slow in revising their prices. However, the limit order book data with trader identification, if available, should be used to explain this empirical finding in further depth.

8.4. Permanent Price Impact

To examine the permanent price impact under which information is incorporated into prices, we use the VAR model of Hasbrouck (1991). The goal is to find if the tick size change has altered the information contents of trades. We use the data categorized to HH_b, HH_s, HC_b, HC_s, CH_b, CH_s, CC_b, and CC_s, where b denotes buyer-initiated trade and s denotes seller-initiated trade. The difference between the buyer and seller-initiated volumes provides the order flow. The sum of the four order flows gives the total order flow in a given minute. Let r_t, the log return of the price at time t, denote the change in the market maker's estimate of fundamental value. For simplicity, we multiply r_t by 10,000 times so that the returns could be interpreted as the percentage change in the exchange rate multiplied by 100. OF_t is the aggregated order flow equal to the net of buyer and seller-initiated trading volumes. We start with the basic assumption that trades have linear price impact (see Kyle (1985) for example):

$$r_t = \beta_t OF_t + \varepsilon_t \tag{3}$$

where ε_t is information available to the market maker in addition to information from order flow. Ensuring exogeneity of ε_t and its interpretation as the market maker's information uncorrelated with order flow requires controlling for microstructure frictions. For example, inventory-control effects may cause past trades to influence the market maker's current trade price. Furthermore, delayed adjustment to information on the part of the market takers means order flow may have lagged effects on trades. This suggests an infinite-order autoregressive specification:

$$r_t = \alpha_0 + \sum_{i=1}^{\infty} \alpha_i r_{t-i} + \sum_{i=1}^{\infty} \beta_i OF_{t-i} + \beta_0 OF_t + \gamma\, OF_t * D_t + \varepsilon_t \tag{4}$$

where D_t is a dummy variable equal to one if the tick size is decimal pip and zero if the tick size is pip. Including an intercept means that the regression demean the variables. The coefficient γ captures the price impact of the tick size change. The negative sign would means the lower price impact and the positive sign means the higher price impact of the tick size change. To find the number of lagged variables, we compare the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) up to 30 lags of past returns and order flows. Figure 25 reports BIC for the lags of order flow and Figure 26 provides the BIC for the lag of returns. Based on the values of AIC and BIC, we choose the 18 lags for the returns and 8 lags for the order flows. This provides us the following equation:

$$r_t = \alpha_0 + \sum_{i=1}^{18} \alpha_i r_{t-i} + \sum_{i=1}^{8} \beta_i OF_{t-i} + \beta_0 OF_t + \gamma\, OF_t * D_t + \varepsilon_t \tag{5}$$

Table 6 reports estimated coefficients and Heteroskedasticity and autocorrelation adjusted p-values for the VAR specification in Equation (5).

Table 6: Estimated Coefficients for the Order Flow Equation in the VAR

	Estimate	Sd.	t-value	Pr(>|t|)		Estimate	Sd.	t-value	Pr(>|t|)
α_0	0.0292	0.0062	4.729	0.0000	α_{15}	-0.0063	0.0031	-2.024	0.0430
α_1	-0.0783	0.0056	-13.966	0.0000	α_{16}	-0.0060	0.0031	-1.960	0.0500
α_2	-0.0194	0.0046	-4.261	0.0000	α_{17}	-0.0082	0.0030	-2.786	0.0053
α_3	-0.0144	0.0044	-3.282	0.0010	α_{18}	-0.0052	0.0030	-1.720	0.0854
α_4	-0.0128	0.0043	-2.987	0.0028	β_0	0.0507	0.0010	51.820	0.0000
α_5	-0.0102	0.0042	-2.444	0.0145	β_1	-0.0030	0.0005	-6.480	0.0000
α_6	-0.0078	0.0040	-1.957	0.0504	β_2	-0.0012	0.0004	-3.158	0.0016
α_7	-0.0105	0.0040	-2.658	0.0079	β_3	-0.0008	0.0004	-2.273	0.0230
α_8	-0.0020	0.0039	-0.515	0.6067	β_4	-0.0007	0.0004	-2.064	0.0390
α_9	-0.0088	0.0040	-2.229	0.0258	β_5	-0.0006	0.0004	-1.568	0.1169
α_{10}	-0.0059	0.0036	-1.661	0.0967	β_6	-0.0007	0.0003	-2.183	0.0291
α_{11}	-0.0089	0.0033	-2.730	0.0063	β_7	-0.0002	0.0003	-0.585	0.5589
α_{12}	-0.0067	0.0031	-2.168	0.0302	β_8	-0.0011	0.0003	-3.175	0.0015
α_{13}	-0.0068	0.0031	-2.235	0.0254	γ	-0.0039	0.0017	-2.324	0.0201
α_{14}	-0.0073	0.0032	-2.269	0.0233					

Notes: $n = 127,531$. Table 6 reports estimated coefficients and heteroskedasticity and autocorrelation adjusted p-values for the VAR specification in Equation (5).

As expected, the estimate of β_0 (*0.0507*) is positive and significant at one percent level. This means that the positive order flow (excess demand) increases the rate of the return in the market. The coefficient of the tick size effect, γ (*-0.0039*), is negative and significant at one percent level. This indicates that the price impact of the trades decreased after the introduction of the decimal pip pricing. The lower price impact is always interpreted as an improvement in the market quality. However, one may argue that the lower price impact comes with a cost of human trades being front-run by HFT's.

9. Conclusions

EBS is the main interdealer market for the currency pairs EUR/USD, USD/JPY, EUR/JPY, USD/CHF, and EUR/CHF. EBS decided to change the tick size, the minimum price improvement, from pip pricing (four decimal points) to decimal pip pricing (five decimal points) in March 2011. This decision changed the EBS market structure significantly. Our analysis shows that the new environment enabled HFT's to front-run human traders using the sub-penny jumping strategy. Unlike the equity market, human traders play a vital role in the interbank FX market for providing liquidity. They were either market taker or market maker in about 65% of EUR/USD transactions in 2011. Using a difference-in-difference regression, we find that the spread decreased after the introduction of decimal pip tick size. However, the benefit of the reduction in the spread was mostly absorbed by HFT's. The absolute value of the realized spread also decreased meaning that the market maker would suffer less adverse price move. There was also a significant drop in market depth. The tick size change altered the market maker-market taker composition of human and algorithmic traders. While the share of the computer maker-computer taker from total trades increased, the share of human maker-human taker decreased after the tick size change. The new market structure also changed the informational content of trades leading to less price impact.

Appendix A.

We consider the general case of sub-penny jumping in the interbank FX market. Suppose that x is the rate of return and $x \sim N(\mu, \sigma^2)$. We will discuss the case of a buying sub-penny jumper, as the selling case is very similar. Once the penny jumper trades, the orders he front-runs protect him from serious losses on his position. The returns are unbounded on one side and limited on the other side at $a = \frac{p-(p+\tau)}{p+\tau} = \frac{-\tau}{p+\tau}$.

$$f(x|x > a) = \frac{f(x)}{Prob(x > a)} = \frac{f(x)}{1 - F(a)} = \frac{\phi(\alpha)}{1 - \Phi(\alpha)} \;\; where \;\; \alpha = \frac{a - \mu}{\sigma}$$

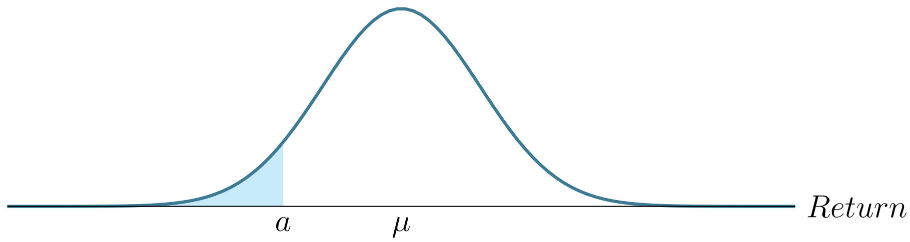

The sub-penny jumping leads to higher conditional expectation of returns with smaller variance.

$$E(x|x > a) = \mu + \sigma\lambda(\alpha), \text{ where } \lambda(\alpha) = \frac{\phi(\alpha)}{1 - \Phi(\alpha)} > 0$$

$$\lambda(\alpha) > 0 \rightarrow \mu + \sigma\lambda(\alpha) > \mu \rightarrow E(x|x > a) > E(x)$$

$$Var(x|x > a) = \sigma^2(1 - \delta(\alpha)) \text{ where } \delta(\alpha) = \lambda(\alpha)[\lambda(\alpha) - \alpha]$$

$$0 < \delta(\alpha) < 1 \rightarrow Var(x|x > a) < Var(x).$$

If the minimum tick size decreases from pip pricing τ_1 to decimal pip pricing τ_2, the loss limit will shift from $a(\tau_1)$ to $a(\tau_2)$ since

$$a = \frac{p - (p + \tau)}{p + \tau} = \frac{-\tau}{p + \tau} \rightarrow \frac{da(\tau)}{d\tau} = \frac{-p}{(p + \tau)^2} < 0.$$

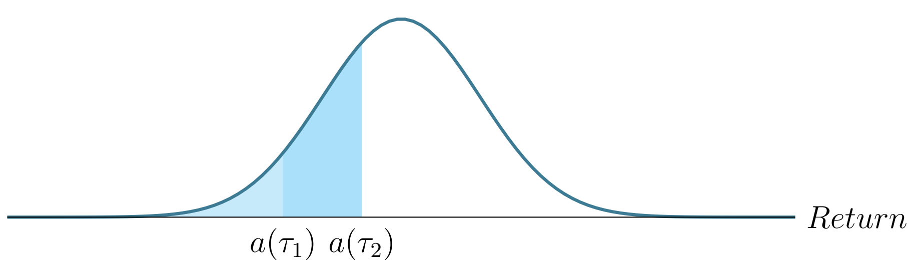

$$\frac{dE(x|x>a)}{d\tau} = \frac{dE(x|x>a)}{d\alpha}\frac{d\alpha}{da}\frac{da}{d\tau} = \frac{d\lambda(\alpha)}{d\alpha}\frac{-p}{(p+\tau)^2}$$

$$\lambda(\alpha) = \frac{\phi(\alpha)}{1-\Phi(\alpha)} \to \frac{d\lambda(\alpha)}{d\alpha} = \frac{-\alpha\phi(\alpha)[1-\Phi(\alpha)] + \phi(\alpha)\phi(\alpha)}{[1-\Phi(\alpha)]^2}$$

$$\to \frac{d\lambda(\alpha)}{d\alpha} = \lambda(\alpha)[\lambda(\alpha)-\alpha] = 0 < \delta(\alpha) < 1 \to \frac{dE(x|x>a)}{d\tau} < 0$$

$$\frac{dVar(x|x>a)}{d\tau} = \frac{dVar(x|x>a)}{d\alpha}\frac{d\alpha}{da}\frac{da}{d\tau} = -\sigma\frac{d\delta(\alpha)}{d\alpha}\frac{-p}{(p+\tau)^2}$$

$$\frac{d\delta(\alpha)}{d\alpha} = \frac{d\lambda(\alpha)}{d\alpha}[\lambda(\alpha)-\alpha] + [\frac{d\lambda(\alpha)}{d\alpha} - 1]\lambda(\alpha) = \lambda(\alpha)[(\lambda(\alpha)-\alpha)^2 + \lambda(\alpha)(\lambda(\alpha)-\alpha) - 1]$$

$$\to \frac{d\delta(\alpha)}{d\alpha} = \lambda(\alpha)[(\lambda(\alpha)-\alpha)(\lambda(\alpha)-\alpha+\lambda(\alpha)) - 1]$$

Case 1: $\lambda(\alpha) \geq 1$

$$\to \lambda(\alpha) - \alpha > 1 \text{ and } \lambda(\alpha) - \alpha + \lambda(\alpha) > 1 \text{ since } \alpha < 0 \quad \to \frac{d\delta(\alpha)}{d\alpha} > 0$$

Case 2: $0 < \lambda(\alpha) < 1$

$$\delta(\alpha) = \lambda(\alpha)(\lambda(\alpha)-\alpha) < 1 \to (\lambda(\alpha)-\alpha) > \frac{1}{\lambda(\alpha)} > \frac{1}{(\lambda(\alpha)-\alpha+\lambda)} \quad \to \frac{d\delta(\alpha)}{d\alpha} > 0$$

$$\frac{d\delta(\alpha)}{d\alpha} > 0 \to \frac{dVar(x|x>a)}{d\tau} > 0 \ \square$$

References

Ahn, H. J., J. Cai, K. Chan, and Y. Hamao (2007). Tick size change and liquidity provision on the Tokyo Stock Exchange. *Journal of the Japanese and International Economics 21*, 173–194.

Alexander, K. and T. Zabotina (2005). Is it time to reduce the minimum tick sizes of the E-Mini futures? *Journal of Futures Markets 25*, 79–104.

Andersen, T. G., T. Bollerslev, D. Francis, and V. Clara (2003). Micro effects of macro announcements: Real-time price discovery in foreign exchange. *American Economic Review 93*, 38–62.

Anshuman, V. R. and A. Kalay (1998). Market making with discrete prices. *Review of Financial Studies 11*, 81–109.

Ascioglu, A., C. Comerton-Forde, and T. H. McInish (2010). An examination of minimum tick sizes on the Tokyo stock exchange. *Japan and the World Economy 22*, 40–48.

Bacidore, J., R. H. Battalio, and R. H. Jennings (2003). Order submission strategies, liquiditysupply, and trading in pennies on the New York Stock Exchange. *Journal of Financial Markets 6*, 337–362.

Bessembinder, H. (2000). Tick size, spreads, and liquidity: An analysis of NASDAQ securities trading near ten dollars. *Journal of Financial Intermediation 9*, 213–239.

Bessembinder, H. (2003). Trade execution costs and market quality after decimalization. *The Journal of Financial and Quantitative Analysis 38*, 747–777.

Biais, B. and P. Woolley (2011). High frequency trading. *Manuscript, Toulouse University, IDEI.*.

Bourghelle, D. and F. Declerck (2004). Why markets should not necessarily reduce the tick size. *Journal of Banking and Finance 28*, 373–398.

Brogaard, J., T. Hendershott, and R. Riordan (2014). High frequency trading and price discovery. *Review of Financial Studies 27*, 2267–2306.

Cai, J., Y. Hamaob, and R. Y. Ho (2008). Tick size change and liquidity provision for Japanese stock trading near ¥1000. *Japan and the World Economy 20*, 19–39.

Chaboud, A., B. Chiquoine, E. Hjalmarsson, and C. Vega (2014). Rise of the machines: Algorithmic trading in the foreign exchange market. *Journal of Finance 69*, 2045–2084.

Chaboud, A. P., S. V. Chernenko, E. Howorka, R. S. Krishnasami Iyer, D. Liu, and J. H. Wright (2004). The high-frequency effects of U.S. macroeconomic data releases on prices and trading activity in the global interdealer foreign exchange market. *Board of Governors of the Federal Reserve System, International Finance Discussion Papers Number 823*.

Cordella, T. and T. Foucault (1999). Minimum price variations, time priority and quote dynamics. *Journal of Financial Intermediation 8*, 141–173.

Goldstein, M. A. and K. A. Kavajecz (2000). Eighths, sixteenths, and market depth: Changes in tick size and liquidity provision on the NYSE. *Journal of Financial Economics 56*, 125–149.

Harris, L. E. (1994). Minimum price variations, discrete bid-ask spreads, and quotation sizes. *The Review of Financial Studies 7*, 149–178.

Hasbrouck, J. (1991). Measuring the information content of stock trades. *Journal of Finance 46*, 179–207.

Hasbrouck, J. and G. Saar (2013). Low-latency trading. *Journal of Financial Markets 16*, 646–679.

Hirschey, N. (2013). Do high-frequency traders anticipate buying and selling pressure? *Technical report, London Business School*.

Jones, C. M. (2013). What do we know about high-frequency trading? *Columbia Business School Research Paper*.

Jones, C. M. and M. L. Lipson (2001). Sixteenths: Direct evidence on institutional execution costs. *Journal of Financial Economics 59*, 253–278.

Kadan, O. (2006). So who gains from a small tick size? *Journal of Financial Intermediation 15*, 32–66.

Kirilenko, A., M. Samadi, A. S. Kyle, and T. Tuzun (2015). The flash crash: The impact of high frequency trading on an electronic market. *University of Maryland Working Paper*.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica 53*, 1315–1335.

Lallouache, M. and F. Abergel (2014). Tick size reduction and price clustering in a FX order book. *Physica A: Statistical Mechanics and its Applications 416*, 488–498.

O'Hara, M. (2015). High frequency market microstructure. *Journal of Financial Economics 116*, 257–270.

Ready, M. (1999). The specialist's discretion: stopped orders and price improvement. *Review of Financial Studies 12*, 1075–1112.

Schmidt, A. B. (2012). Ecology of the modern institutional spot FX: The EBS market in 2011. *SSRN: http://ssrn.com/abstract=1984070*.

Terrence, H., C. M. Jones, and A. J. Menkveld (2011). Does algorithmic trading improve liquidity? *The Journal of Finance 66*, 1–33.