

**Essays on Market Microstructure
and
Foreign Exchange Market**

by

Soheil Mahmoodzadeh

M.Sc., Allameh Tabatabai University, 2009

B.Sc., Shahed University, 2006

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in the

Department of Economics
Faculty of Arts and Social Sciences

© Soheil Mahmoodzadeh 2015

SIMON FRASER UNIVERSITY

Summer 2015

All rights reserved.

However, in accordance with the *Copyright Act of Canada*, this work may be reproduced without authorization under the conditions for “Fair Dealing.” Therefore, limited reproduction of this work for the purposes of private study, research, criticism, review and news reporting is likely to be in accordance with the law, particularly if cited appropriately.

APPROVAL

Name: Soheil Mahmoodzadeh
Degree: Doctor of Philosophy
Title of Thesis: Essays on Market Microstructure and Foreign Exchange Market

Examining Committee: Dr. Simon Woodcock, Associate Professor
Chair

Dr. Ramazan Gençay
Senior Supervisor
Professor

Dr. Alain Chaboud
Supervisor
Economist, Board of Governors of the Federal Reserve
System

Dr. Robert Jones
Supervisor
Professor

Dr. Jamal Nazari
Internal Examiner
Assistant Professor

Dr. Ilias Tsiakas
External Examiner
Associate Professor, Department of Economics and
Finance, University of Guelph

Date Defended: August 4, 2015

Abstract

High frequency trading (HFT) has become a predominant feature of financial markets. This thesis studies different aspects of the HFT in the in the Electronic Broking Services (EBS) interbank foreign exchange (FX) market.

The first paper of this Thesis (Chapter 1) studies changes in the spread, market depth and degree of adverse selection due to the lower minimum tick size. The main conclusion is that the reduction in spread was mostly absorbed by the HFTs, whereas the manual traders were pushed back from the top of the order book and experienced longer execution times. Manual market makers were willing to cross the spread and act as market takers changing the informational content of the order flow. Market depth was reduced significantly following the introduction of decimal pip pricing.

The second paper of this Thesis (Chapter 2) presents the effect of tick size change on the adverse selection problem in the EBS market. Econometric analysis of serial correlation properties of jumps in exchange rates, and of the spread leads to the conclusion that adverse selection is reduced by tick size change. Similar cleavage occurs before and after tick size change in an empirical adverse selection proxy. This chapter sheds light on trading behavior of market participants.

The third paper of this thesis (Chapter 3) discusses the properties of triangular arbitrage opportunity in the EBS market. The results cast into question current understanding of triangular arbitrage in the literature, specifically in relation to algorithmic trading. The increasing presence of algorithmic traders does not offer significant improvement in speed of price discovery by quickly consuming the triangular arbitrage opportunities. Rather, algorithmic trading influences the creation of triangular arbitrage by two countervailing effects.

To Mitra

“What you seek is seeking you.”

— RUMI

Acknowledgments

Throughout my PhD, I have been blessed to receive insights, support and encouragement from my main supervisor Ramazan Gencay. My heartfelt thanks for his patience and generosity. I am very grateful to my supervisor, Alain Chaboud, who I learned from him the interbank FX market. I learned many topics from Robert Jones throughout our discussions about the financial markets. The conversations that I had with my supervisors about my research significantly improved my work. The second and third chapter of the thesis is in collaboration with Michael Tseng who I enjoyed working with him. I would like to acknowledge the insightful feedbacks I have received from: Chris Muris, Simon Woodcock, and seminar participants at the Simon Fraser University and University of Victoria. All omissions and mistakes are my own.

Contents

Approval	ii
Abstract	iii
Dedication	iv
Quotation	v
Acknowledgments	vi
Contents	vii
List of Tables	ix
List of Figures	x
1 Tick Size Change in the Interbank FX Market	1
1.1 Abstract	1
1.2 Introduction	2
1.3 Literature Review	3
1.4 Data Description and Market Structure	5
1.5 Sub-Penny Jumping	11
1.6 Sub-Penny Jumping, Deal and Quote Distributions	15
1.7 Tick Size and Interbank FX Market Quality	21
1.7.1 Spread	21
1.7.2 Adverse Selection	24
1.7.3 Market Depth	28
1.8 Conclusions	28

2	White Noise Jumps and Adverse Selection in FX Market	30
2.1	Abstract	30
2.2	Introduction	31
2.3	Literature Review	34
2.4	Description of Data	36
2.5	Jumps	37
2.5.1	Semi-martingale model of price	37
2.5.2	Estimation of volatility and jumps	38
2.5.3	Microstructure noise	38
2.5.4	Confirmation of Stylized Facts	39
2.5.5	Jumps Before and After Tick Size Change	42
2.5.6	Serial Correlation Properties of Jumps	42
2.6	Components of quoted spread	46
2.7	Realized Spread	47
2.8	Market participant behavior	49
2.9	Conclusions	53
3	Unprofitable Arbitrage, Speed Is Not of the Essence	55
3.1	Abstract	55
3.2	Introduction	56
3.3	Data Description	59
3.4	Occurrence of Triangular Arbitrage	60
3.5	Profit and Duration of Triangular Arbitrage	63
3.6	Non-profitability	65
3.6.1	Left over problem	65
3.6.2	Non-exploitation	66
3.7	Tick Size Effect	74
3.7.1	Perturbation Toward Triangular Arbitrage	74
3.7.2	Pressure on spread	76
3.8	Conclusions	76
	References	78
	Appendix A Truncated Rate of Return	82

List of Tables

1.1	Spread Summary Statistics	21
1.2	Conditional Mean Estimates from the DID Regression Model	24
1.3	Difference-in-Difference Regressions of EUR-USD Spread	25
1.4	Difference-in-Difference Regressions of EUR-USD Spread	25
2.1	EUR/USD Extreme Volatility Periods	40
2.2	Jump Component of EUR/USD Process Before and After Tick Size Change	45
2.3	Estimated EUR/USD components of spread, 5-seconds, Sample size = 163, 069.	46
2.4	Size of Average Order (Million) at the Best Bid	51
2.5	Size of Average Order (Million) at the Best Ask	51
2.6	Queue Jumping in One Trading Hour After Tick Size Change	52
3.1	Size of Left Overs, Rout 1	66
3.2	Size of Left Overs, Rout 2	67

List of Figures

1.1	EBS Minimum Tick Size Changes, 2009-2013	6
1.2	EUR-USD, Last Digit Distribution, Pip Tick Size	10
1.3	EUR-USD, Last Digit Distribution, Decimal Pip Tick Size	10
1.4	EUR-USD, Volume Distribution, Decimal Pip, Last Digit=0	10
1.5	EUR-USD, Volume Distribution, Decimal Pip, Last Digit \neq 0	10
1.6	Sub-Penny Jumping, MT- HFT	12
1.7	EUR-USD, LOB Last Digit Distributions, Pip Tick Size	15
1.8	EUR-USD, LOB Last Digit Distributions, Decimal Pip Tick Size	16
1.9	EUR-USD, LOB Last Digit Distributions, Half Pip Tick Size	17
1.10	EUR-USD, LOB Volume Distributions, Pip Tick Size	17
1.11	EUR-USD, LOB Volume Distributions, Decimal Pip Tick Size	18
1.12	EUR-USD, LOB Volume Distributions, Half Pip Tick Size	19
1.13	EUR-USD, Deal Last Digit Distributions with Different Tick Sizes	19
1.14	EUR-USD, Deal Volume Distributions with Different Tick Sizes	20
1.15	EUR-USD Binding Spread, Pip & Decimal Pip	22
1.16	EUR-USD Binding Spread, Half Pip	22
1.17	EUR-USD Spread, Pip & Decimal Pip Tick Size	22
1.18	EUR-USD Spread, Half Pip Tick Size	22
1.19	EUR-USD Realized Spread, Pip & Decimal Pip	27
1.20	EUR-USD Realized Spread, Half Pip	27
1.21	EUR-USD Market Depth, Pip & Decimal Pip	27
1.22	EUR-USD Market Depth, Half Pip	27
2.1	RV - BP, EUR/USD	39
2.2	RV - BP, USD/JPY	39
2.3	Intraday Seasonality of EUR/USD, 2010	41

2.4	Intraday Seasonality of EUR/USD, 2011	41
2.5	Integrated Volatility of EUR/USD, 2010	41
2.6	Integrated Volatility of EUR/USD 2011	41
2.7	Jump Sizes of EUR/USD, Before Tick Size Change	43
2.8	Jump Sizes of EUR/USD, After Tick Size Change	43
2.9	Jump Inter-Arrival Times Before Tick Size Change, EUR/USD	43
2.10	Jump Inter-Arrival Times After Tick Size Change, EUR/USD	43
2.11	Daily Number of Jumps of EUR/USD, Before Tick Size Change	44
2.12	Daily Number of Jumps of EUR/USD, After Tick Size Change	44
2.13	Jump Inter-Arrival Times, Before Tick Size Change	44
2.14	Jump Inter-Arrival Times, After Tick Size Change	44
2.15	EUR/USD Realized Spread, Five Second Frequency	48
2.16	EUR/USD Adverse Selection Proxy, Five Second Frequency	48
2.17	EUR/USD Realized Spread, Ten Second Frequency	48
2.18	EUR/USD Adverse Selection Proxy, Ten Second Frequency	48
2.19	Best Bid Last Digits, Before Tick Change	51
2.20	Best Bid Last Digits, After Tick Change	51
2.21	Best Ask Last Digits, Before Tick Change	51
2.22	Best Ask Last Digits, After Tick Change	51
3.1	Daily Arbitrage by Duration, EUR-USD-JPY Rout 1	62
3.2	Daily Arbitrage by Duration, EUR-USD-JPY Rout 2	62
3.3	Daily Arbitrage by Seconds, EUR-USD-JPY Rout 1	62
3.4	Daily Arbitrage by Seconds, EUR-USD-JPY Rout 2	62
3.5	Triangular Arbitrage Duration EUR-USD-JPY	63
3.6	Triangular Arbitrage Duration Distribution EUR-USD-JPY	64
3.7	Triangular Arbitrage Profit EUR-USD-JPY	65
3.8	Triangular Arbitrage Profit Distribution EUR-USD-JPY	65
3.9	After Arbitrage Trade Codes 2004, Rout 1	69
3.10	After Arbitrage Trade Codes 2004, Rout 2	69
3.11	Before Arbitrage Trade Codes 2004, Rout 1	69
3.12	Before Arbitrage Trade Codes 2004, Rout 2	69
3.13	After Arbitrage Trade Codes 2005, Rout 1	70

3.14	After Arbitrage Trade Codes 2005, Rout 2	70
3.15	Before Arbitrage Trade Codes 2005, Rout 1	70
3.16	Before Arbitrage Trade Codes 2005, Rout 2	70
3.17	After Arbitrage Trade Codes 2006, Rout 1	71
3.18	After Arbitrage Trade Codes 2006, Rout 2	71
3.19	Before Arbitrage Trade Codes 2006, Rout 1	71
3.20	Before Arbitrage Trade Codes 2006, Rout 2	71
3.21	After Arbitrage Trade Codes 2007, Rout 1	72
3.22	After Arbitrage Trade Codes 2007, Rout 2	72
3.23	Before Arbitrage Trade Codes 2007, Rout 1	72
3.24	Before Arbitrage Trade Codes 2007, Rout 2	72
3.25	After Arbitrage Trade Codes 2006, Rout 1	73
3.26	After Arbitrage Trade Codes 2008, Rout 2	73
3.27	Before Arbitrage Trade Codes 2006, Rout 1	73
3.28	Before Arbitrage Trade Codes 2008, Rout 2	73
3.29	EUR-USD Increment to Create Arbitrage	74
3.30	Change in EUR-USD to Create Arbitrage, Rout 1	75
3.31	Change in EUR-USD to Create Arbitrage, Rout 2	75
3.32	Binding Spread, 2004-2008	75

Chapter 1

Tick Size Change in the Interbank FX Market

1.1 Abstract

This paper studies changes in the spread, market depth and degree of adverse selection due to the lower minimum tick size (which was changed from a pip to a decimal pip) in the Electronic Broking Services (EBS) interbank foreign exchange (FX) market. Coupled with the lower tick size, the special composition of traders and their order placement strategies have provided an exceptional opportunity for high frequency traders (HFTs) to implement the sub-penny jumping strategy to front-run manual traders. Our analysis shows that the lower tick size in the interbank FX market has enabled HFTs to be more aggressive in employing the sub-penny jumping strategy due to limited losses. Using difference-in-difference (DID) regressions, we show that spread as a liquidity cost was reduced after the tick size change. However, the benefit of the reduction in spread was mostly absorbed by the HFTs, whereas the manual traders were pushed back from the top of the order book and experienced longer execution times. As a result, manual market makers were willing to cross the spread and act as market takers under decimal pip pricing changing the informational content of the order flow. Market depth was reduced significantly following the introduction of decimal pip pricing due to the occupation of the top of the order book by HFTs and the change in the behavior of the manual market makers, who were forced to become market takers.

1.2 Introduction

In this paper, we study changes in market quality due to the lower minimum tick size in the EBS market and the role of HFTs. EBS is the leading interbank 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 decided to reduce the tick size on major currency pairs from a pip to a decimal pip, where tick size is the minimum price movement. 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 HFTs. Reuters, the main competitor of EBS, still uses pip pricing.

We show how the specific structure of the EBS market has helped HFTs 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, manual and high frequency traders. Unlike in the equity market, manual traders constitute approximately 75% of all EBS customers in EUR/USD currency. The composition of the FX interbank market entails that the tick size problem is different from the equity market. Manual traders use keyboards for order management, place large orders and do not cancel their orders. HFTs use this information by trading in front of and on the same side as manual traders. Once a sub-penny jumper front-runs a manual trader, he is protected from serious losses on his position because in the event of an adverse price movement, the sub-penny jumper limits his losses by trading with the manual trader. Given the presence of a manual trader with a large order, and a high fill ratio, small tick size limits the potential loss of a penny jumper. However, if there is a favorable price move, the sub-penny jumper profits to the full extent of the price change. Therefore, the gain of the sub-penny jumper is unbounded on the upside and bounded on the downside. The sub-penny

¹In equity markets, traders are forbidden from accepting, ranking or displaying orders in price increments smaller than a penny. The concern is that sub-penny increments are used to gain priority over other essentially same-priced quotes. However, sub-penny trading is allowed if it does not result from executions of quotations in sub-penny increments (trades executed on Alternative Trading Systems such as dark pools, or internalized by broker-dealers).

jumper profits at the expense of the manual trader by taking liquidity that otherwise would have gone to the manual trader. We do not assume that front running by HFTs is the only consequence of the tick size change. However, we argue that the order management strategy employed by manual traders greatly facilitates implementation of sub-penny jumping by HFTs.

We also discuss the effects of the tick size change on spread, market makers' adverse selection and market depth. Using difference-in-difference estimators, we find that, as a measure of liquidity cost, spread 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 HFTs. As a result, manual traders were pushed back in the limit order book, facing longer execution times. In this environment, manual market makers were willing to become market takers under decimal pip pricing to overcome the disadvantage they faced with respect to execution time. This changed the information content of the order flow, which has been empirically observed in the significant change in realized spread as a proxy for adverse selection. These changes led to the notable drop in market depth following adoption of decimal pip pricing. HFTs, using the sub-penny jumping strategy, occupied the top of the order book with smaller orders, while manual traders with larger orders shifted from market making to market taking.

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 a pip in September 2012. Half a pip is decimal pip pricing in which traders can use only '0' or '5' as the final digit. Market quality measures improved upon adoption of half pip pricing by EBS.

The remainder of the paper is organized as follows: [Section 1.3](#) provides a brief literature review. [Section 1.4](#) describes the data and the EBS market structure, [Section 1.5](#) describes sub-penny jumping, [Section 1.6](#) provides empirical evidence for the existence of sub-penny jumping. [Section 1.7](#) discusses the effects of tick size changes on interbank FX market quality. [Section 1.8](#) concludes.

1.3 Literature Review

Only a few papers have examined the effects of tick size on the interbank foreign exchange market. Using proprietary data from EBS, [Schmidt \(2012\)](#) documents that manual traders did not use the last digit very often under decimal pip pricing. He also provides a taxonomy

of types of EBS customers and their order placement characteristics. [Lallouache and Abergel \(2014\)](#) analyze the data distributions of EUR/USD and USD/JPY 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. However, reasons for this change and its consequences remain unclear.

Empirical studies of stock exchanges generally find that a tick size reduction is associated with a decline in both spread and depth.² [Bessembinder \(2000\)](#) finds that spreads are reduced when tick size decreases. [Jones and Lipson \(2001\)](#) show that when the New York Stock Exchange (NYSE) lowered its minimum price increment on most stocks from eighths of a dollar to sixteenths of a dollar in 1997, quoted and effective spreads declined, but realized execution costs for institutional trades increased. [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. [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. [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 stocks with larger tick size reductions and higher trading activity. [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. [Bourghelle and Declerck \(2004\)](#) show that a change in tick size generated neither a lower liquidity provision for large trades nor a change in the spread on the Paris Bourse. 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,

²There 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.

where tick size is always smaller for more liquid pairs. [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. This suggests that a time and place advantage can be used to implement a penny jumping strategy. 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 Zobotina \(2005\)](#), [Kadan \(2006\)](#), and [Ascioglu et al. \(2010\)](#).

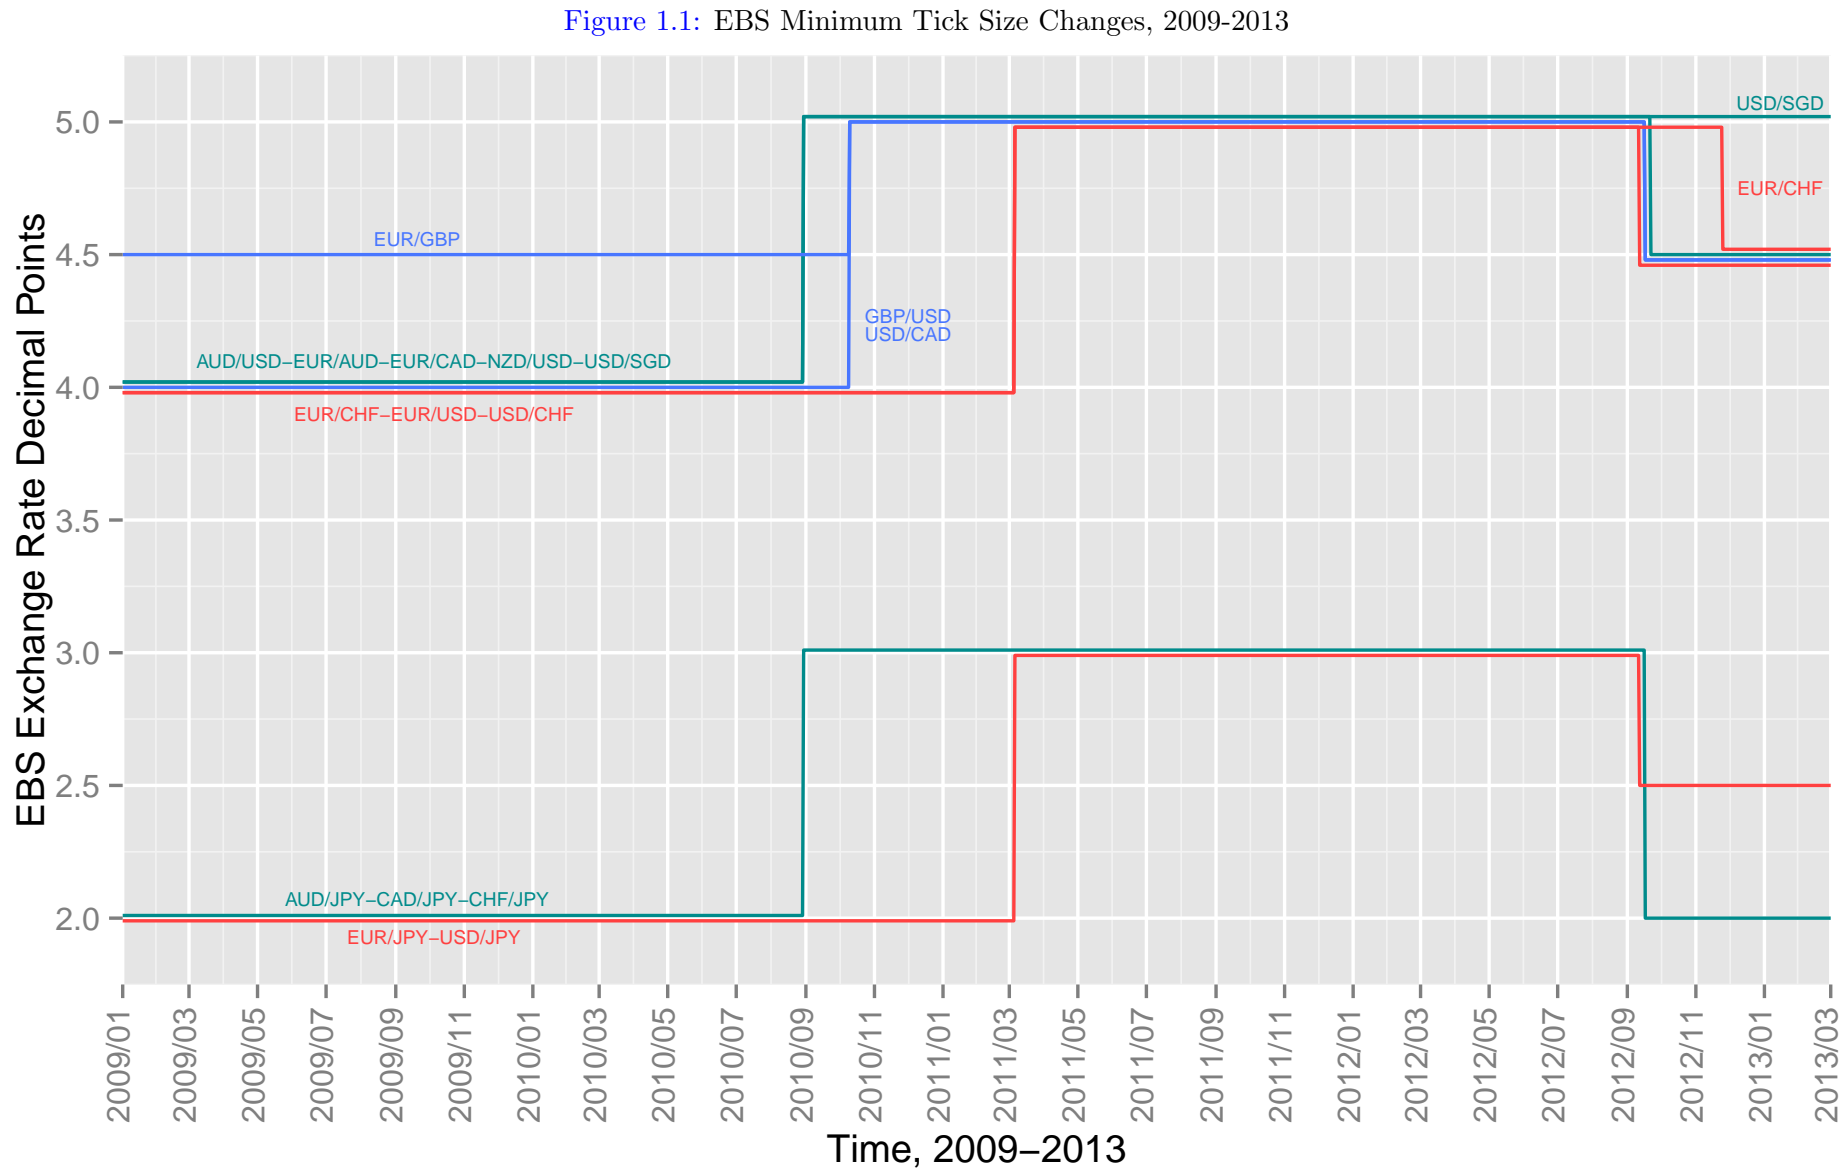
1.4 Data Description and Market Structure

Banks trade currencies with each other on two wholesale 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. We use EBS data, which includes 10 levels of “Quotes” (buy, sell) and “Deal” records at 100 milliseconds each.³ There are around 225 million snapshots of the order books for each year per major currency. The dealt prices are the highest buying or lowest selling deal price between two consecutive snapshots of the order book, rounded to 100 milliseconds. 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.⁴ 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. \(2004\)](#). To find the minimum tick size changes from 2009 to 2013 in the EBS market, we analyzed approximately 800 gigabytes of millisecond limit order book data for 60 currency pairs. [Figure 1.1](#) illustrates the tick size changes in the EBS market.⁵ 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.

³“Deal” is a term used by EBS. It is synonymous with “trade”.

⁴EBS 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.

⁵We have excluded currency pairs with minimal activities.



Notes: Figure 1.1 illustrates the tick size changes in the EBS market from 2009 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 minimum 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 Japanese Yen. 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”.

Half-pip pricing is a special case of decimal pip pricing in which the fifth digit can only be “0” or “5”.⁶ Figure 1.1 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 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 to decimal pip pricing in March 2011. This move aroused intense debate among the two main types of traders. Although decimal pip pricing was welcomed by HFTs, manual traders believed that HFTs already had an unfair advantage, which was enhanced by the smaller tick size. After EBS shifted to decimal pip pricing to attract more HFTs, 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.

In this section, we provide information about the EBS market structure to help explain how, given a lower minimum tick size, the structure of the interbank FX market greatly facilitates the implementation of sub-penny jumping by HFTs. 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.⁷

There are two main types of traders in the EBS market, Manual Traders and Automated Traders.⁸ “Manual” traders use GUI-based access for order management.⁹ They 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 in the provision of liquidity. “Automated” traders use an automated interface (AI) to place orders without human intervention.¹⁰ The main component of an AI in the EBS market is the professional trading community (PTC), which places orders at very high frequency. EBS first allowed AI into the market in 2005. The same classification is adopted by Schmidt (2012) and Chaboud et al. (2014).

The typical order management of manual and automated traders has been discussed in Schmidt (2012) with EBS proprietary data. Using EBS client identity data, Schmidt (2012)

⁶The pip (decimal pip) tick size is 2 (3) decimal points when the local currency is JPY.

⁷The full results are available upon request.

⁸See BIS 2011 report: <http://www.bis.org/publ/mktc05.pdf> for more details.

⁹See <http://www.ebs.com/access-methods/ebs-workstation.aspx>.

¹⁰See <http://www.ebs.com/access-methods/ebs-ai.aspx>.

makes the following empirical observations regarding the ecology of the EBS market in 2011:

Empirical Observation 1. *Manual traders are the main EBS customers. They place orders less frequently and are slower to react than are HFTs.*

Manual traders account for approximately 75% of all EBS customers in EUR/USD, whereas for example, the PTC accounts for 16.6% of customers in the same currency pair. Even though the manual traders dominate the number of EBS customers, they place only 3.7% of daily orders, whereas the PTC places 61.6% of daily orders. This is due to the fact that automated traders are capable of placing orders at a frequency far exceeding that for manual traders.

Empirical Observation 2. *Manual traders place large orders, whereas HFTs place smaller orders.*

The PTC typically uses an order size of one million and almost never uses an order size exceeding four million. In contrast, manual traders place large orders. For example, approximately 5% of orders submitted by manual traders reach a size of ten million. We should emphasize that this does not necessarily mean that all manual traders use large orders in the market. This is the typical behavior of the manual traders, however, some manual traders may find it better strategy to use smaller orders.

Empirical Observation 3. *Manual traders do not cancel their orders, in contrast to HFTs; therefore, manual traders have a high fill ratio.*

The fill ratio is defined as the ratio of dealt quotes to submitted quotes. This ratio is more than 50% for manual traders and approximately 7% for the PTC. The high fill ratio for manual traders implies that these traders typically do not cancel their orders. It is a strategy by HFTs to send and cancel their orders. In contrast, manual traders are very slow and they are not capable to implement such a strategy.

Empirical Observation 4. *Manual traders place orders at prices ending in “0” under decimal pip pricing.*

Manual traders did not widely use decimal pip pricing after the tick size change. Approximately 80% of the manual traders used decimal pip pricing in less than 20% of their orders. In contrast, AI traders have adopted decimal pip pricing very well, only 12% of them used decimal pips in less than 20% of their orders.

The following statement summarizes the manual traders’ order placement strategy.

Hypothesis 1. *Manual traders place large orders at prices ending in “0” (under pip pricing) and on average do not cancel half of their orders.*

There are specific reasons why manual 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 manual traders became accustomed to pip pricing over many years and found it difficult to adapt to decimal pip pricing. If manual traders did not use decimal pip pricing and used “0” as the fifth digit, there should be price clustering at zero after the tick size change. [Figures 1.2](#) and [1.3](#) show distributions of the final digit of prices under pip and decimal pip pricing. The final digit distribution is nearly uniform under pip pricing, but there is price clustering at zero under decimal pip pricing.

Figure 1.2: EUR-USD, Last Digit Distribution, Pip Tick Size


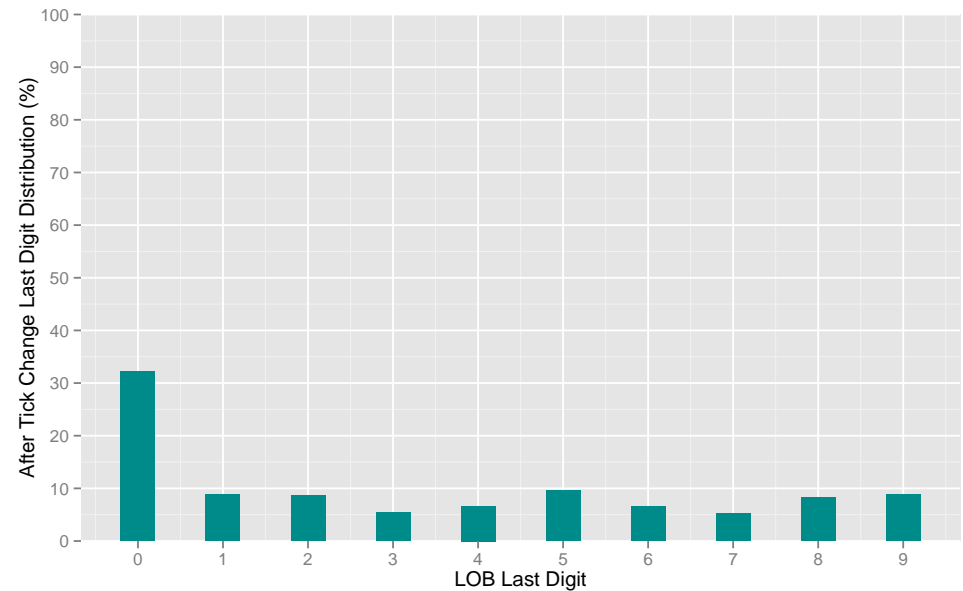


Figure 1.3: EUR-USD, Last Digit Distribution, Decimal Pip Tick Size



Notes: Figures 1.2 and 1.3 show the limit order book price's 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 final digit of prices. However, manual traders did not use the final 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 1.4: EUR-USD, Volume Distribution, Decimal Pip, Last Digit=0

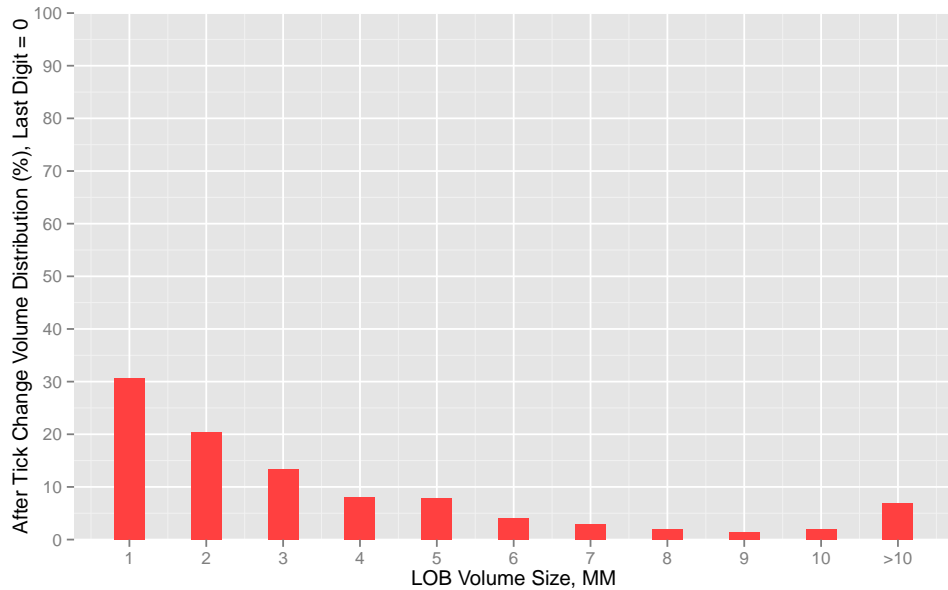



Figure 1.5: EUR-USD, Volume Distribution, Decimal Pip, Last Digit≠0



Notes: Because manual traders usually use zero for the final digit and place large orders, the volume distributions for prices ending in zero and non-zero digits differ. Figures 1.4 and 1.5 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 manual traders usually place large orders at prices ending in zero.

A weaker price clustering is also found for prices ending in “5”. The reason for this may be that some manual traders used “5” to partially adapt to decimal pip pricing. If manual traders typically used zero for the last digit and placed large orders, the volume distributions of prices ending in zero and non-zero digits should differ. [Figures 1.4](#) and [1.5](#) 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.

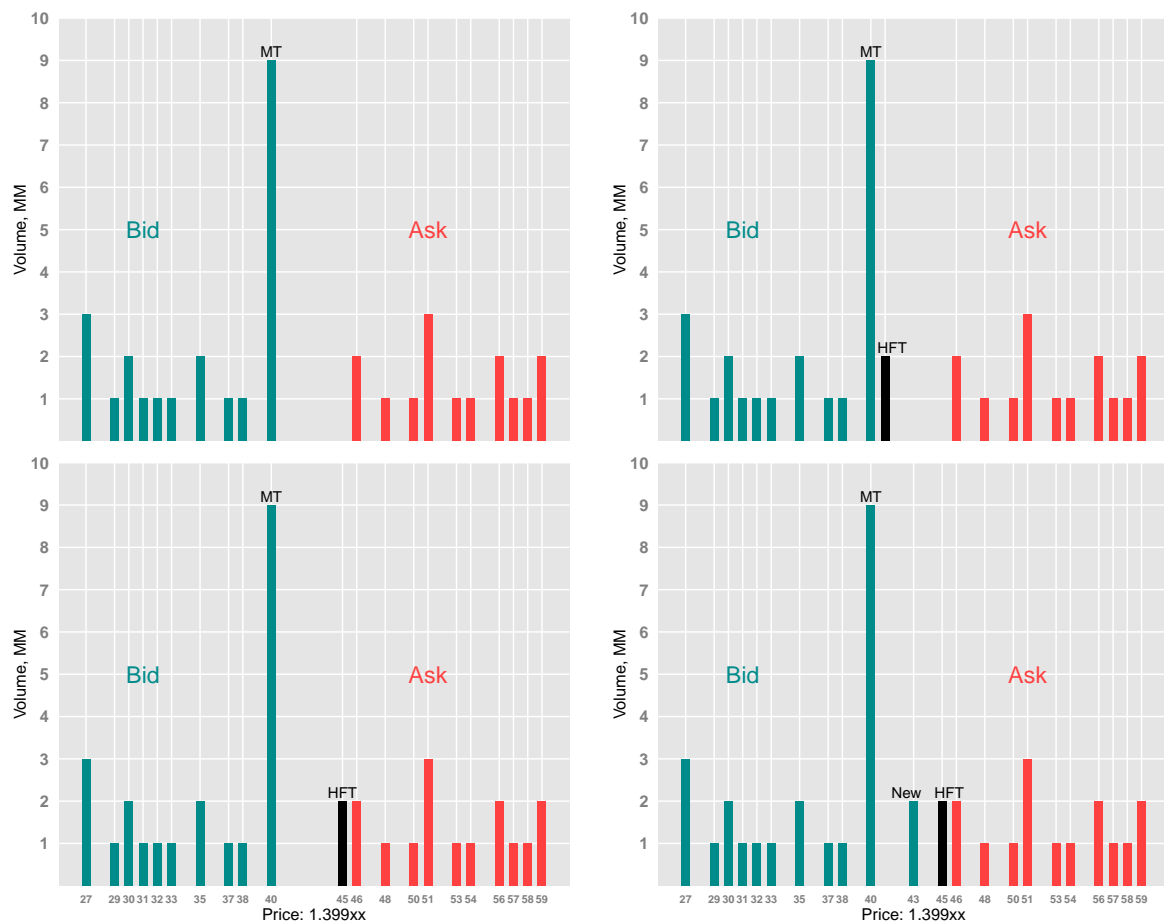
The described structure of the EBS market creates an ideal setting for HFTs to front-run manual traders, using the sub-penny jumping strategy. A manual trader who utilizes a price ending in “0” and trades in large volumes is a patient market maker, given his high fill ratio. Using this information, HFTs can front-run manual traders, as manual traders nearly guarantee that HFTs’ losses will be limited by the tick size. Sub-penny jumping is not the only strategy implemented by HFTs to front-run other traders. However, the particular order management strategy employed by manual traders greatly facilitates the implementation of sub-penny jumping by HFTs.

1.5 Sub-Penny Jumping

Manual traders are the main EBS costumers. They are slow traders; they submit large orders at prices ending in zero under decimal pip pricing; and they do not cancel their orders very often. In this section, we explain how HFTs use this unique market structure (very different from the equity market) to implement the sub-penny jumping strategy to exploit the smaller tick size. [Figure 1.6](#) depicts possible sub-penny jumping in the EBS market. Suppose that a manual 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 manual 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 manual trader at a price of 1.39940, and his loss would be \$20. As explained above, manual traders were reluctant to engage in sub-penny jumping, thus avoiding use of decimal pip pricing. This makes the sub-penny strategy even more profitable for the HFTs. However, there are also costs associated with a front running strategy. The HFT will not front-run with the same size order, as the manual trader's order could be partially filled. There is also competition between HFTs to front-run manual traders, which would limit potential profits. HFTs also use other order anticipation strategies to trade, so they must determine the best time to implement a sub-penny jumping strategy. We find compelling evidence that HFTs used the front running strategy frequently under decimal pip pricing.


Figure 1.6: Sub-Penny Jumping, MT- HFT



Let us now consider the general case of sub-penny jumping in the interbank FX market. Suppose that a manual 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 manual trader does not cancel or adjust the order and other traders do not fill the manual 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 selling case is very similar. Suppose that the rate of return has a standard normal distribution: $x \sim N(\mu, \sigma^2)$. In what follows, f, F will denote the pdf and cdf, respectively. Similarly, ϕ, Φ will denote the pdf and cdf of the standard normal distribution. Once the 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 $a = \frac{p-(p+\tau)}{p+\tau} = \frac{-\tau}{p+\tau}$.

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



Proposition 1. $E(x|x > a) > E(x)$ and $\text{Var}(x|x > a) < \text{Var}(x)$. The sub-penny jumping leads to higher conditional expectation of returns with smaller conditional variance.¹¹

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

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

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


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

Proposition 2. $\frac{da(\tau)}{d\tau} < 0$. The smaller the tick size is, the higher the loss limit will be.

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

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

¹¹See the Appendix A. ?? for details.



Proposition 3. $\frac{dE(x|x > a)}{d\tau} < 0$, $\frac{dVar(x|x > a)}{d\tau} > 0$. *The lower the tick size is, the higher the conditional expectation of returns will be and the lower the conditional variance of returns will be.*

$$\frac{dE(x|x > a)}{d\tau} = \frac{dE(x|x > a)}{d\alpha} \frac{d\alpha}{da} \frac{da}{d\tau} = \frac{d\lambda(\alpha)}{d\alpha} \frac{-p}{(p + \tau)^2} < 0 \text{ since } 0 < \frac{d\lambda(\alpha)}{d\alpha} = \delta(\alpha) < 1 \quad (1.1)$$

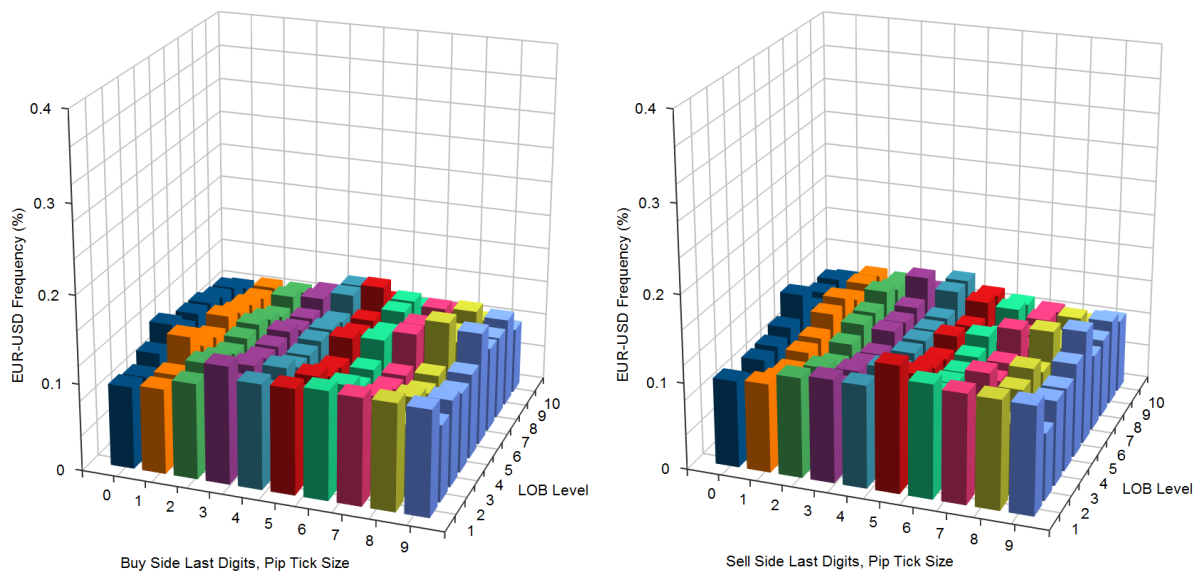
$$\frac{dVar(x|x > a)}{d\tau} = \frac{dVar(x|x > a)}{d\alpha} \frac{d\alpha}{da} \frac{da}{d\tau} = -\sigma \frac{d\delta(\alpha)}{d\alpha} \frac{-p}{(p + \tau)^2} > 0 \text{ since } \frac{d\delta(\alpha)}{d\alpha} > 0 \quad (1.2)$$

Equations (1.1) and (1.2) show that a smaller tick size makes sub-penny jumping more profitable, as the conditional expected rate of return increases, and the conditional variance decreases. With a lower tick size, the sub-penny jumper must improve prices by a smaller price increment to trade ahead of the manual traders. Therefore, the tick size is the price that the sub-penny jumper must pay to front-run the manual traders. However, sub-penny jumpers do not pay this price to the manual traders; instead, they pay it to the traders with whom they trade to establish their positions. These traders would have traded with the manual 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 manual 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's structure presented in Section 1.4, HFTs know that orders with prices ending in zero and large order sizes indicate the presence of manual traders to front-run. These orders are not canceled or adjusted very often, and they cannot be filled quickly. In general, HFTs post inside quotes, and manual traders post the outside spread. Then, if the market starts to move, HFTs trade against the outside spread. This occurs systematically, due to the order management strategy employed by manual traders, which differentiates this market from the equity market.

1.6 Sub-Penny Jumping, Deal and Quote Distributions

We study tick size in the EBS market in three periods. Tick size was a pip (4 decimal points) until March 2011, when it was reduced to a decimal pip (5 decimal points). It was then changed to a half-pip in September 2012. Half-pip pricing is decimal pip pricing in which the fifth digit can only be “0” or “5”. We first discuss the limit order book last digit and volume distributions then turn to the same distributions in deals. As shown in [Figure 1.7](#), distributions of orders final digits were almost uniform with pip tick size for both the bid and ask sides. The x axis shows the final 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. Frequency distributions are given on the z axis. The uniform distribution implies that sub-penny jumping was not a common strategy of HFTs under pip pricing. If an HFT engages in sub-penny jumping, he would face a large tick size when prices move against him.

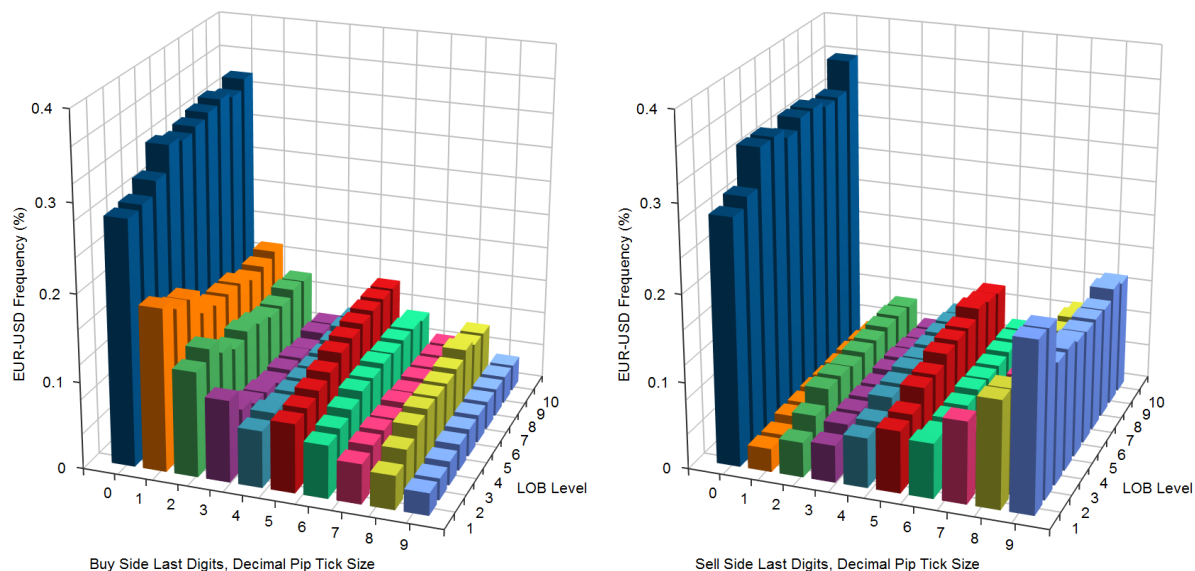
[Figure 1.7](#): EUR-USD, LOB Last Digit Distributions, Pip Tick Size



Because manual 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 manual 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 manual 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 expect a decreasing distribution of the number of orders for the buy side of the limit order book at all levels. [Figure 1.8](#) shows such distributions for the ten levels of the

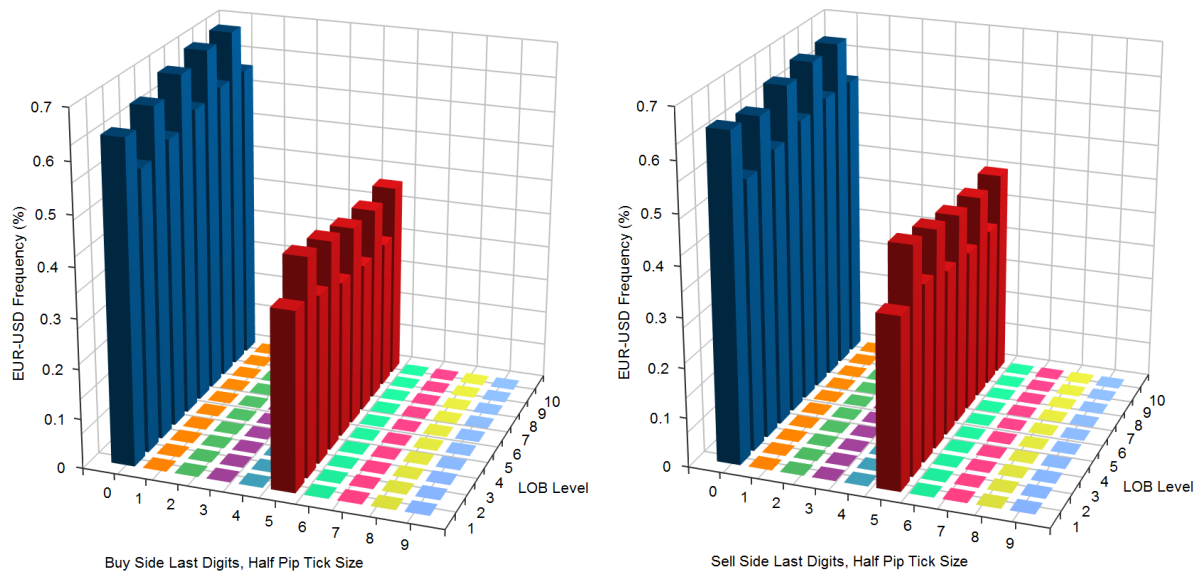
order book with different final digits. For example, at the level-one buy side of the limit order book, 28% of orders have a last digit of “0”, and 18% of the orders have a last digit of “1”. This ratio starts to fall and reaches 2.5% for orders with a last digit of “9”. The shapes of the distributions at other levels of the bid side are similar to the shape at level-one.

Figure 1.8: EUR-USD, LOB Last Digit Distributions, Decimal Pip Tick Size




If a sub-penny jumper improves upon the order of a manual trader by one tick while bidding, he would place an order with a last digit of “1”. However, if an HFT improves on a manual 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 1.8, we see that approximately 28% of the orders have a last digit of “0”, but only approximately 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 manual 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 small, and there was no opportunity for HFTs 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 1.9.

Figure 1.9: EUR-USD, LOB Last Digit Distributions, Half Pip Tick Size



Volume distributions also provide evidence of sub-penny jumping. When HFTs front-run manual traders, they secure their positions by placing orders with smaller volumes. We know from [Section 1.4](#) that HFTs 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 1.10](#) and [1.11](#), respectively.

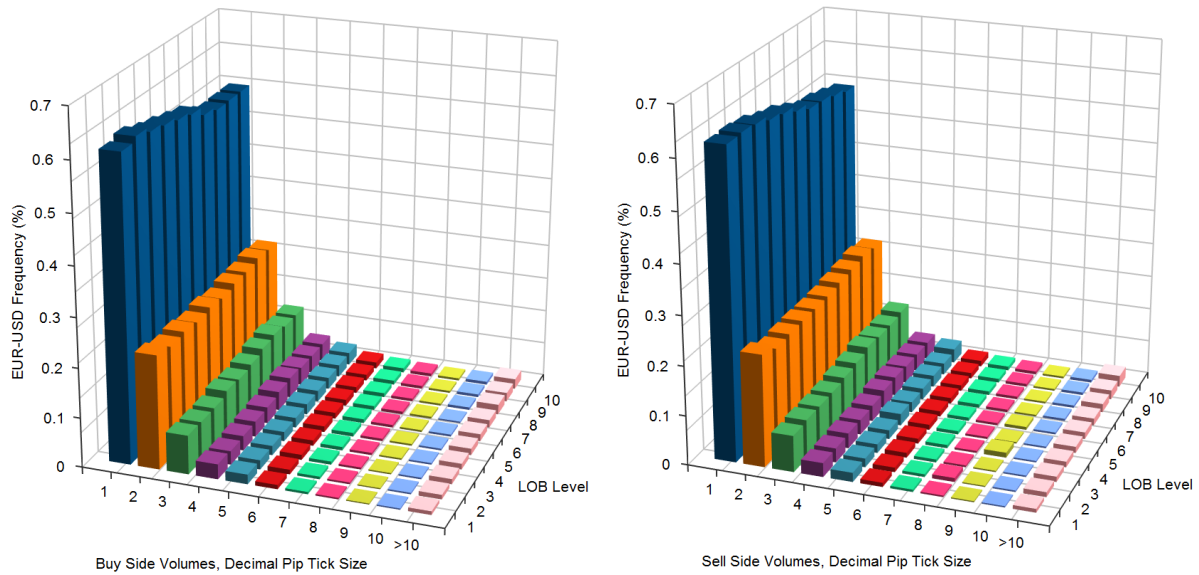
Figure 1.10: EUR-USD, LOB Volume Distributions, Pip Tick Size



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. Comparing [Figures 1.10](#) and [1.11](#), we observe that large volumes nearly disappeared from the top of the order book after the introduction of decimal pip pricing. There are various explanations for the significant change at the top of the order book.

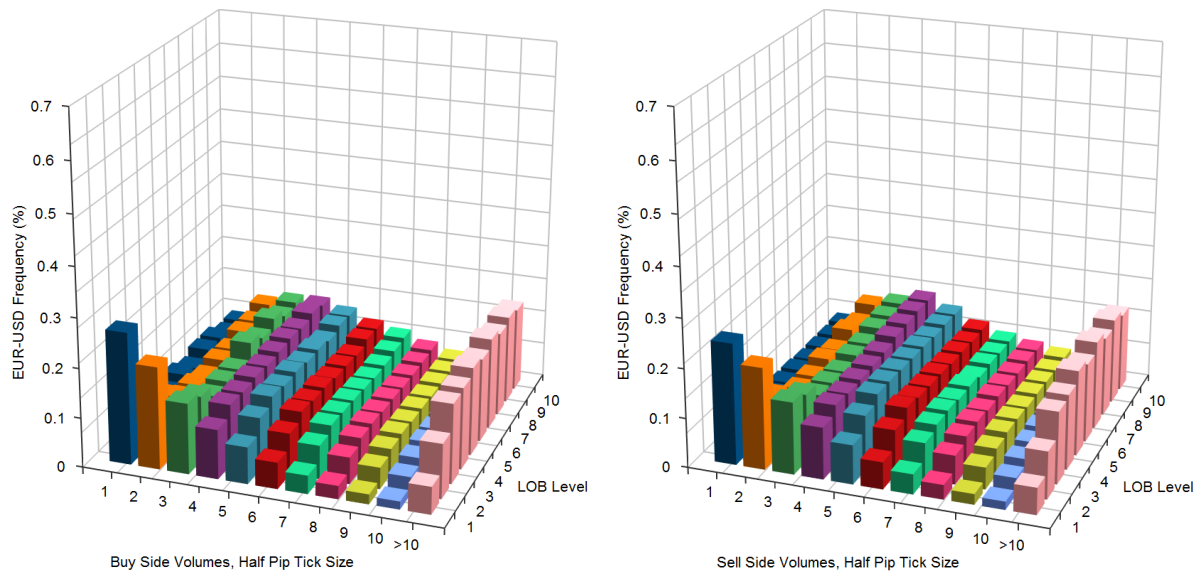
[Figure 1.11](#): EUR-USD, LOB Volume Distributions, Decimal Pip Tick Size



The implementation of sub-penny jumping by HFTs pushed back manual traders in the order book, which led to smaller trading volumes. Furthermore, some manual traders started to switch from providing liquidity to demanding liquidity. Given that we observe mainly order sizes of 1 million under decimal pip pricing, HFTs take more executions away from manual traders under decimal pip pricing than under pip pricing. A manual trader with a large order to fill should wait longer in the back of the order book under decimal pip pricing. Instead, he might prefer to use market orders to solve the execution time problem. Thus, the large orders would disappear from the order book.


Along with other reasons, the relative percentage of small orders submitted by manual traders might also increase after the adoption of decimal pip pricing by EBS. Although this could also partially explain the change in the order book, the volume distribution of manual traders indicates that they did not stop using large orders. [Figure 1.12](#) indicates that there are more large volumes at the top of the order book after reversion to half-pip pricing in September 2012. The symmetry in the order book for the ask and bid sides is an interesting empirical finding.

Figure 1.12: EUR-USD, LOB Volume Distributions, Half Pip Tick Size



Patterns similar to those of the order books are present in the deal price. We examine the trades and split them into trades initiated by buyers and by sellers. Figure 1.13 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 final 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.

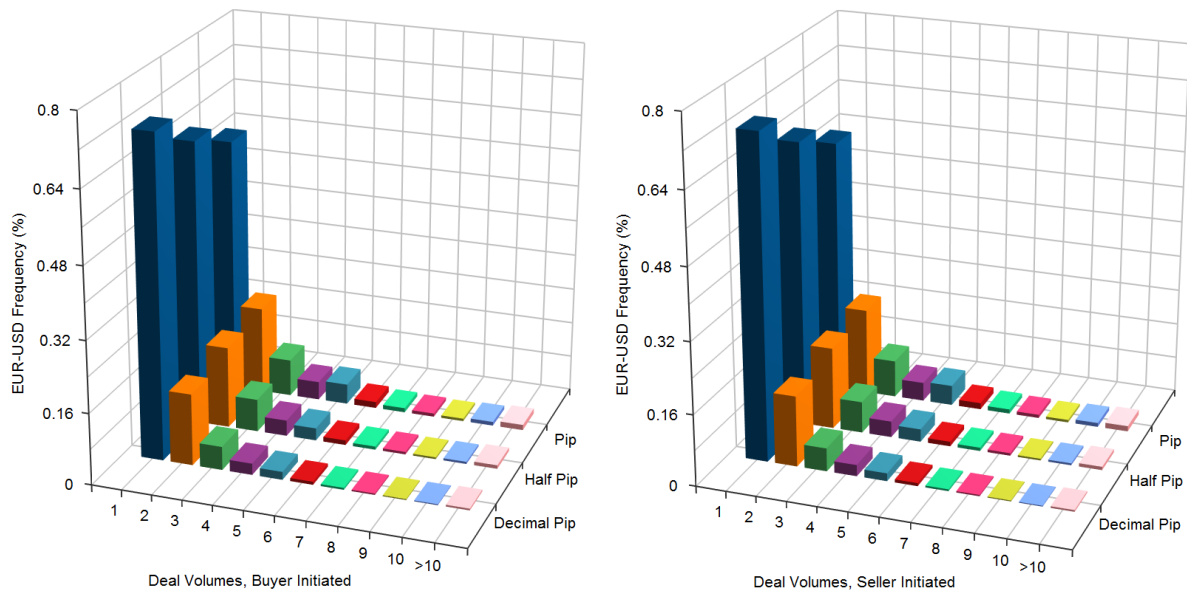
Figure 1.13: EUR-USD, Deal Last Digit Distributions with Different Tick Sizes



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 1.7. 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 1.7. 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.

Figure 1.14: EUR-USD, Deal Volume Distributions with Different Tick Sizes



The shape of the order book has also changed the distribution of deal volumes. Figure 1.14 shows that in buyer-initiated trades, 57% of deals occurred with a volume of one million under pip pricing, 72% of deals occurred with a volume of one million under decimal pip pricing, and 65% of deals occurred with a volume of one million under half pip pricing. We observe the same pattern in seller-initiated trades. Figure 1.14 shows that more deals occurred with larger volumes under pip pricing, as larger volumes were available at the top of the order book.

1.7 Tick Size and Interbank FX Market Quality

1.7.1 Spread

Spread, defined as the difference between the best bid and the best asking price, is one measure of liquidity cost. [Table 1.1](#) 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 1.1](#): Spread Summary Statistics


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

[Figure 1.15](#) indicates whether the spread was binding under pip pricing. The results show that on average, the spread was binding more than 50% of the time and equal to the tick size (i.e., 0.0001). This means there was pressure on the spread to decrease, which could partially justify the adoption of decimal pip tick size. [Figure 1.15](#) 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 “0” or “5”. In [Figure 1.16](#), we demonstrate the binding spreads with half pip tick size. On average, between 15% and 20% of the time, the spread was binding under half pip pricing. Compared with other tick sizes, the binding spread under half pip pricing is not as high as 50% or as low as 1%.

[Figure 1.17](#) shows the one-hour average spread for 2011. This frequency better represents the change in spread from pip to decimal pip tick size. 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. We have also excluded holidays and negative or zero spreads.¹² 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.

¹²If there is no bilateral credit between a buyer and a seller at the top of the limit order book, the trade does not occur, and the spread could be negative or zero.

Figure 1.15: EUR-USD Binding Spread, Pip & Decimal Pip



Notes: Figure 1.15 illustrates the binding EUR/USD spread under pip and decimal pip tick size. The dashed line indicates the time of the tick size change in March 2011. Figure 1.16 shows the binding spread for the same currency under half pip tick size.

Figure 1.16: EUR-USD Binding Spread, Half Pip




Figure 1.17: EUR-USD Spread, Pip & Decimal Pip Tick Size





Figure 1.18: EUR-USD Spread, Half Pip Tick Size



Notes: Figure 1.17 illustrates the one-hour frequency EUR/USD spread under pip and decimal pip tick size for 2011. Figure 1.18 illustrates the one-hour frequency EUR/USD spread under half pip tick size for 2013.

Technically, traders have more options inside the spread, which was under greater pressure before. In [Figure 1.18](#), the variance of the spread is lower with half pip pricing. [Figure 1.18](#) represents the spread with half pip pricing, using the same frequency. We will test the following hypothesis regarding the spread:

Hypothesis 2. *Decimal pip pricing narrowed the bid-ask spread in the EBS market.*

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. Tick size did not change for these currency pairs in the time period under consideration. We could not study the change in the spread from decimal pip to half pip because, as shown in [Figure 1.1](#), 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 minimum tick size on other venues. The relationship between EUR/USD and the rest of the market with regard to trading the control currency pairs remained constant across the tick size reduction.¹³

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

$$S_t = \beta_0 + \beta_1 P + \beta_2 T + \beta_3 PT + \epsilon_t \quad (1.3)$$

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

¹³If the HFTs 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.

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 1.2](#).

[Table 1.2](#): 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 1.3](#) for different control groups. In all cases, using different control groups, the estimates of β_3 are negative and significant at the one percent level. Consequently, the spread decreased after the introduction of decimal pip pricing. We also used the one-minute average spread frequency in robustness checks. The results, presented in [Table 1.4](#), show that the estimates of β_3 are all -0.00003 and 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. At the same time, our findings, presented in [Figures 1.10](#) and [1.11](#), indicate that manual traders were pushed back in the order book. Inasmuch as HFTs occupied the top of the limit order book with the sub-penny jumping strategy, the number of large orders (which is a proxy for manual traders) decreased significantly at the top of the order book after the tick size change. As a result, the benefit of the reduction in the spread was mostly absorbed by the HFTs.

1.7.2 Adverse Selection

Reduction in the tick size from pip to decimal pip enables some liquidity suppliers to post more aggressively priced limit orders, leading to tighter quoted spreads. However, by reducing the cost of implementing penny jumping strategies, the HFTs occupied the top of the order book for the provision of liquidity. One might expect manual traders to adjust their order submission strategies in response to the lower tick size. However, manual traders continued to use pip pricing even with the decimal pip tick size. Therefore, they were pushed back in the order book, meaning that HFTs lowered the execution probability of limit orders by manual traders. Consequently, the longer execution time reduced the incentive for manual traders to provide liquidity, forcing manual traders to act more as market takers than market makers. This subpopulation changed the information content of the order flow and adverse selection.

Table 1.3: Difference-in-Difference Regressions of EUR-USD Spread

Control Group	EUR-GBP ($T=4,236$), $R^2 = 0.65$				AUD-USD ($T=4,012$), $R^2 = 0.59$				GBP-USD ($T=3,670$), $R^2 = 0.84$			
	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	0.000001	175.70	$2 * 10^{-16}$	0.000224	0.000001	182.57	$2 * 10^{-16}$	0.000347	0.000001	221.86	$2 * 10^{-16}$
β_1	0.000015	0.000002	7.30	$3 * 10^{-13}$	0.000013	0.000002	7.30	$3 * 10^{-13}$	-0.000012	0.000002	-5.489	$4 * 10^{-8}$
β_2	-0.000113	0.000002	-53.72	$2 * 10^{-16}$	-0.000113	-0.000076	-43.60	$2 * 10^{-16}$	-0.000202	0.000002	-91.388	$2 * 10^{-16}$
β_3	-0.000030	0.000003	-10.32	$2 * 10^{-16}$	-0.000030	-0.000032	-13.21	$2 * 10^{-16}$	-0.000006	0.000003	-2.193	$3 * 10^{-2}$

All coefficients are significant at the 1 percent level.

Table 1.4: Difference-in-Difference Regressions of EUR-USD Spread

Control Group	EUR-GBP ($T=250,880$), $R^2 = 0.59$				AUD-USD ($T=242,544$), $R^2 = 0.54$				GBP-USD ($T=219,448$), $R^2 = 0.79$			
	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	$2 * 10^{-7}$	1202.23	$2 * 10^{-16}$	0.000219	0.000001	1284.30	$1 * 10^{-16}$	0.000339	$2 * 10^{-7}$	1052.84	$2 * 10^{-16}$
β_1	0.000015	$3 * 10^{-7}$	50.44	$3 * 10^{-13}$	0.000016	0.000002	65.89	$2 * 10^{-16}$	-0.000006	$3 * 10^{-7}$	46.02	$2 * 10^{-16}$
β_2	-0.000110	$3 * 10^{-7}$	-363.46	$2 * 10^{-16}$	-0.000072	-0.000076	-297.45	$2 * 10^{-16}$	-0.000113	$3 * 10^{-7}$	-744.21	$2 * 10^{-16}$
β_3	-0.000031	$4 * 10^{-7}$	-74.51	$2 * 10^{-16}$	-0.000034	-0.000032	-101.85	$3 * 10^{-16}$	-0.000030	$4 * 10^{-7}$	-32.56	$2 * 10^{-2}$

All coefficients are significant at the 1 percent level.

Notes: Table 1.3 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 1.4 shows the DID results for the one-minute EUR/USD spread average. All β_3 coefficients are negative, equal to -0.00003 and significant at the one percent level.

This is empirically observed in the significant change in the realized spread as a proxy for adverse selection.

$$\text{Realized Spread}_t = 2q_t(p_t - m_{t+s}) \quad (1.4)$$

$$q_t = \begin{cases} 1 & : \text{Buyer Initiated Trade} \\ -1 & : \text{Seller Initiated Trade} \end{cases}$$

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. The idea behind the realized spread is that if an informed trader buys, the midpoint later should increase, resulting in a negative realized spread. If an informed trader sells, the midpoint should decrease later, and given the trade sign, the realized spread would also be negative in this situation. [Figure 1.19](#) 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. The dashed line indicates the change in tick size from pip to decimal pip in March 2011. The realized spread was usually negative under pip pricing, indicating a higher degree of adverse selection. However, the realized spread increased significantly when the tick size changed from pip to decimal pip. All things equal, lower adverse selection is preferred to higher adverse selection. However, there is a cost associated with lower adverse selection in the EBS market. Market depth decreased significantly after the introduction of decimal pip pricing due to the occupation of the top of the order book by HFTs and the change in the behavior of manual market makers, who were forced to become market takers. Market depth will be discussed in the next section. [Figure 1.20](#) provides the realized spread under half pip pricing in 2013. Comparing the decimal pip and half pip tick sizes, the realized spread is lower under half pip pricing. This is consistent with the relationship between sub-penny jumping and the degree of adverse selection in the EBS market. Because the penny-jumping strategy is less profitable for HFTs under half pip pricing, manual traders would face shorter execution times and be willing to provide more liquidity in the market.

Figure 1.19: EUR-USD Realized Spread, Pip & Decimal Pip





Figure 1.20: EUR-USD Realized Spread, Half Pip



Notes: Figure 1.19 indicates the one hour realized spread average under pip and decimal pip pricing in 2011. Figure 1.20 shows the realized spread under half pip pricing in 2013.

Figure 1.21: EUR-USD Market Depth, Pip & Decimal Pip





Figure 1.22: EUR-USD Market Depth, Half Pip



Notes: Figure 1.21 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 1.22 illustrates buy side depth under half pip pricing in 2013.

1.7.3 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 found a EUR/USD market depth of 0.0001 and 0.0002 under pip, decimal pip and half pip pricing. [Figure 1.21](#) shows the daily averages of buy side depth. As before, the dashed line indicates the tick size change from pip to decimal pip in March 2011. Each circle shows the order size that is necessary to move the best price by 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 1.21](#) shows that the introduction of decimal pip pricing reduced market depth significantly. We found similar results for the ask side of the order book.

There are various reasons why market depth worsened after the introduction of decimal pip pricing. First, the HFTs implemented the sub-penny jumping strategy and occupied the top of the order book with smaller volumes. Furthermore, there was less incentive for manual traders to supply liquidity due to the longer execution times, which worsened market depth. Some manual traders may also have switched from large to smaller orders to adapt to tick size changes. When EBS changed the tick size to a half pip, there was a significant improvement in market depth, as shown in [Figure 1.22](#).

1.8 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, i.e., the minimum price improvement, from pip pricing (four decimal points) to decimal pip pricing (five decimal points) for quoted prices in March 2011. This decision changed the EBS market's microstructure significantly. Our analysis shows that the EBS market's structure enabled high frequency traders to front-run manual traders using the sub-penny jumping strategy. Manual traders typically place large orders in the order book at prices ending in "0", and they do not cancel their orders. Using this information, HTFs place orders in front of manual traders by improving prices by the amount of the minimum tick size. Once the sub-penny jumper trades, the orders he front-runs

protect him from serious losses on his position. If the price moves in his favor, the sub-penny jumper profits to the full extent of the price change. Therefore, the returns are unbounded for HFTs on one side and limited on the other side. We also show that the lower tick size helped HFTs become more aggressive in sub-penny jumping. Using a difference-in-difference regression, we find that the spread as a liquidity cost decreased after the introduction of decimal pip pricing. However, due to the implementation of the sub-penny jumping strategy by HFTs, manual traders were pushed back in the order book, and the benefit of the smaller tick size was absorbed by HFTs. This increased the execution time for manual traders, and they were willing to act more as liquidity makers than liquidity takers. This subpopulation changed the information content of the order flow and the degree of adverse selection in the market. This is empirically observed in the significant change in the realized spread as a proxy for adverse selection. Market depth decreased under decimal pip pricing on both the bid and ask sides. HFTs used sub-penny jumping with smaller orders, and manual traders were willing to become market takers under decimal pip pricing. We found consistent results for the period when EBS changed the tick size from a decimal pip to a half pip.

Chapter 2

White Noise Jumps and Adverse Selection in FX Market

2.1 Abstract

We study the degree of adverse selection in the main inter-dealer foreign exchange platform, the Electronic Broking Services (EBS). Tick-by-tick limit order book data from EBS from 2010 to 2011 is examined and particular attention paid to a tick size change instituted in the EBS market on March 7, 2011. Based on a multi-faceted analysis, we show that this microstructure event drastically changed the informational relationship between market makers and market takers. The reduced tick size has a clear whitening effect on the jump time series of quoted exchange rates. We argue that this indicates abatement of adverse selection problem for the market maker. Other evidence supports this conclusion. There is a reduction in the adverse selection component of the spread. A downward cleavage also occurs before and after tick size change in an empirical adverse selection proxy. We uncover the reason behind this decline of adverse selection by examining the shape of the limit order book and behavior of market participants. Our results form a case study on how market microstructure affects information asymmetry in the presence of algorithmic trading activity.

2.2 Introduction

The spot foreign exchange (FX) market daily turn over reached to 1.5 trillion USD in 2013, approximately 7 times that of equity markets.¹ A significant portion of this volume, most recently estimated at 35%, are interdealer trades. Given the immensity of the FX market and potential impact of exchange rates on world economy, how information is propagated in the FX market is an important question. In contrast to the highly fragmented equity market, the interbank foreign exchange spot market consists mainly of two venues, where large banks and institutions trade with each other in units of million. Reuters Matching is the primary trading venue for commonwealth and emerging market currencies and Electronic Broking Services (EBS) is the leading liquidity provider for currency pairs EUR/USD, USD/JPY, EUR/JPY, USD/CHF and EUR/CHF.² On March 7, 2011, EBS decided to change tick size, i.e. minimum price increment, from four to five decimal places –from *pip* to *decimal pip* pricing– in FX market vernacular.³ In this paper, we investigate how this microstructure event altered the informational relationship between market participants –in particular the algorithmic traders and their manual counterparts– in the EBS market. We present our results on the EUR/USD currency pair, which accounts for 28% of global FX turnover.⁴

The impact of algorithmic trading on market quality has garnered increasing scrutiny, with particular impetus stemming from events such as the May 6, 2010 flash crash, and the Oct 15, 2014 bond market flash crash. For example, [Hendershott et al. \(2011\)](#) shows that algorithmic trading enhances the informativeness of quotes and improves liquidity in NYSE limit order books. Using publicly available NASDAQ data, [Hasbrouck and Saar \(2013\)](#) concludes that increased algorithmic trading activity lowers short-term volatility, decrease spreads, and increase depth. Also on NASDAQ, [Brogaard et al. \(2012\)](#) find that algorithmic traders enhances price efficiency. On the other hand, [Biais and Woolley \(2011\)](#) discusses some tactics used by algorithmic traders for market manipulation, e.g. “stuffing”, “smoking”, and “spoofing”.⁵ In this paper we examine the impact of algorithmic trading via both traditional microstructure

¹See the report of Bank for International Settlements (BIS): <http://www.bis.org/publ/rpfx13fx.pdf>. The spot market makes up 37% of the global FX market. Other currency related instruments include FX swaps, forwards, currency swaps, and options.

²FX market convention is to list base currency first. For example, EUR/USD should be read as US dollar per Euro.

³“Pip” is abbreviation for Price Increment Point.

⁴See the BIS report: <http://www.bis.org/publ/rpfx10.pdf>. Individually the US dollar and Euro are involved in approximately 75% and 46% of all spot transactions respectively. Results for other four major currency pairs do not differ qualitatively from EUR/USD and are available upon request.

⁵For example, “spoofing” is a bait-and-switch tactic which involves submitting and canceling orders with no intention of execution, with the goal of swaying prices in a favorable direction.

measures of information efficiency and a detailed empirical analysis of trader behavior. This is the first paper to reconcile the apparent improvement in market quality due to the presence of algorithmic traders and their predatory market making activity. The tick size change serves as a mechanism that reveals the different preferences of algorithmic and manual traders.

While algorithmic trading has become a pre-requisite for market making in equity markets, manual traders still play an important role in the interbank foreign exchange market as liquidity suppliers.⁶ Given the monopoly of the EBS platform in the interbank market in major currencies, manual traders –in particular the manual market makers– did not to relocate their trading after tick size change. Akin to a coordination game, an individual manual trader has no incentive to transfer his liquidity elsewhere and there is positive network externality only if all manual traders coordinate to move their trading activity *en masse*. It is difficult to envision such a coordinated move given the current structure of the inter-dealer market. On the contrary, limit order book evidence points to manual traders remaining in the market after tick size change. This is due to the FX market practice of trading via “vehicle currencies” thereby concentrating liquidity in the major currency pairs. Furthermore, examining the order book shows that the decimal pip pricing scheme further delineates the behavior of manual and algorithmic traders. The absence of liquidity transfer means that our results is a case study on how a microstructure event affect the information structure between market participants while controlling for variation in trader population.

From the microstructure point of view, order flow represents noisy information flow, from liquidity consumers/market takers to liquidity providers/market makers (see. e.g. the seminal papers Kyle (1985), Glosten and Milgrom (1985) and their descendants). In a limit order book market, the jump component of price directly captures the most conspicuous part of market taker’s liquidity consumption. Jumps in prices are caused by large market orders.⁷ Serial correlation properties of jump component of price therefore indicate the information content of large order flows. For example, staggered jumps in the same direction arriving in a clustered manner show that the market orders responsible for the jumps are driven by market taker’s private information. Conversely, a jump time series which is white noise means large market orders convey no information. Such jumps are caused by market orders placed for exogenous liquidity reasons. Analyzing the price process spanning the entire two years of 2010 to 2011, we

⁶For example, the Swiss National Bank is an active manual trader on the EBS platform.

⁷To the best of our knowledge, price jumps caused by market makers adjusting quotes by large discrete increments –known as price “gapping”– are rare occurrences for electronic order book markets in general, apart from singular catastrophic events like the Flash Crash. We do not observe any price gapping in our data.

find that the jump time series reject strongly the null hypothesis of no serial correlation before tick size change and does not reject after tick size change. This suggests that the severity of the market maker's adverse selection problem is lessened significantly after decimal pip tick size is introduced.

In addition to the price process, we also analyze the limit order book during the same period. Decomposing the quoted spread into components due to transaction cost, inventory risk, and adverse selection shows that the adverse selection component decreases after tick size change. We also consider an empirical proxy of adverse selection defined using the effective spread and realized spread. We find clear negative level shift in the proxy time series after tick size change. Therefore, adverse selection decrease unconditionally for market makers, not just conditional on large orders.

This evident reduction in adverse selection can be explained by market participant behavior before and after tick size change. We show that, after tick size change, the shape of the limit order book exhibits pronounced price clustering at previous tick size. Our data suggest strongly that this peculiar shape is due to manual market makers' failure to adapt to new decimal tick size. The refusal of manual market makers to adapt to new decimal tick size accentuates the easiness with which algorithmic market makers can employ the queue jumping strategy already made available by tick size change. Conceding the top of the book forces manual traders who are formerly market makers to cross the spread and take the market, thereby diluting –significantly, our results show– the informational content of incoming market orders.

The rest of this paper is organized as follows. [Section 2.3](#) reviews relevant literature covering the three strands of technical part of our analysis: high frequency continuous-time econometrics, limit order book econometrics, and tick size change. [Section 2.4](#) gives a description of EBS order book data. [Section 2.5](#) considers the structural change in the jump component of the price processes before and after tick size change. En route to our analysis of serial correlation of jumps, we confirm stylized facts regarding the foreign exchange market. [Section 2.6](#) discusses components of the quoted spread before and after tick size change. In addition to adverse selection, we confirm the market maker's inventory risk decrease after tick size change as expected. [Section 2.7](#) discusses the realized spread before and after tick size change and its implication on information asymmetry. [Section 2.8](#) explains the observed reduction in adverse selection by presenting an analysis of the shape of the order book and its implications regarding market participants behavior. [Section 2.9](#) concludes.

2.3 Literature Review

As appropriate for our high-frequency data, we model the price process by a continuous-time semi-martingale.⁸ The econometric analysis of such processes is first considered by [Barndorff-Nielsen and Shephard \(2004\)](#) which separates quadratic variation into its continuous and jump components. The quadratic variation is estimated by realized variance and integrated volatility estimated by realized bipower variation, with the difference between the two provides a consistent estimate of the jump contribution to price variation. [Barndorff-Nielsen and Shephard \(2006a\)](#) provide an asymptotic distribution theory to construct non-parametric tests for the presence of jumps; this is the test we use in testing for jumps. Finite sample refinements have been offered by [Huang and Tauchen \(2005\)](#) and [Barndorff-Nielsen and Shephard \(2006b\)](#). Studies that apply this methodology to data (see e.g. [Andersen et al. \(2007\)](#), [Andersen et al. \(2007\)](#), and [Andersen et al. \(2010\)](#) in the context of equity markets) have been confined to jump and volatility estimation. Rather than merely detecting jumps, we focus on the informational implications of jump behavior. We do, however, confirm stylized facts regarding seasonality of high frequency markets *en route*.

In dealing with the trade-off between approaching the high-frequency limit and facing contamination by microstructure noise, we adopt the same approach of [Andersen et al. \(2007\)](#) in comparing the difference between realized volatility and bipower variation, which consistently estimates the jump contribution to price variation in the absence of microstructure noise, across different sampling frequencies. Our sampling frequency is every 30 seconds; we are not aware of any other studies that samples the price process at a frequency higher than every 5 minutes. An alternative approach introduced by [Jacod et al. \(2009\)](#), [Podolskija and Vetter \(2009\)](#) and [Podolskij and Vetter \(2009\)](#) is to exploit the data at the highest frequency but uses local pre-averaging to produce noise reduced observations. [Christensen et al. \(2014\)](#) construct noise –and outliers– robust estimator using pre-averaging method. Rather than merely seeking to escape the frequency zone occupied by microstructure noise, local pre-averaging requires statistical assumptions on the nature of microstructure noise that might not be empirically justified for the EBS market we consider.

Microstructure theory divides the quoted spread into three components: expected loss due to adverse selection ([Glosten and Milgrom \(1985\)](#), [Copeland and Galai \(1983\)](#)), inventory risk ([Stoll \(1978\)](#), [Amihud and Mendelson \(1980\)](#), [Ho and Stoll \(1981\)](#)) and transaction cost ([Roll](#)

⁸Asset pricing theory postulates that in an arbitrage-free market, prices necessarily follow a semi-martingale.

(1984)). The trade indicator model of [Huang and Stoll \(1997\)](#) has been used extensively in analyzing components of the spread, in the context of equity markets. For example, [Van Ness et al. \(2002\)](#) examine the NYSE and NASDAQ spreads for the same stocks. An alternative way to estimate adverse selection in the market is [Easley and O'Hara \(1987\)](#) (see also [Easley et al. \(2015\)](#)), probability of informed trading model. Their volume-based approach bypasses the spread and is suitable for markets where algorithmic brokerage is prevalent. Considering the make-up of market participants in the inter-dealer FX market, we use the [Huang and Stoll \(1997\)](#) model in decomposing the spread.

The microstructure perspective on foreign exchange markets originated with [Lyons \(1995\)](#), which tests validity of structural microstructure models and confirms both the informational and inventory-control aspects of market maker in reaction to incoming order flow. In this paper, we focus on the informational aspect. [Payne \(2003\)](#), using the variance decomposition methodology of [Hasbrouck \(1991\)](#), showed that adverse selection contributed to 60% of quoted spread in one week USD/DEM data from Reuters D2002-2 dealing system. [Bjønnes and Rime \(2005\)](#) find also that adverse selection component makes up for a large portion of spread in [Huang and Stoll \(1997\)](#) trade indicator model, although dealer's own prices do not reflect inventory control effort. In this paper, we investigate the change of adverse selection component due to a specific microstructure event, namely tick size change.

Tick size drew the attention of microstructure literature as the minimum price increment in the US equity markets moved from one-eighth to one-sixteenth, and finally decimal pricing (see, for example, [Angel \(1997\)](#), [Goldstein and Kavajecz \(2000\)](#), and [Schultz \(2000\)](#)). Extensive studies have been done on the relationship between tick size and informational structure for equity markets, with a diverse array of conclusions reached. [Harris \(1994\)](#) advances empirically the position that small tick size benefits professional traders at the expense of large order traders and public traders who use limit orders. [Gibson et al. \(2003\)](#) examines NYSE data from approximately two months before and after decimalized pricing of January 29 2001. Also using methodology of [Huang and Stoll \(1997\)](#), they found that adverse selection component of the spread did not change significantly across decimalization, with the possible explanation that pre-decimalization spread was artificially too large, allowing dealers to charge excessive order processing cost. To the best of our knowledge, there is no literature discussion on effect of tick size on adverse selection in the FX market. We provide a multi-faceted analysis on effect of tick size change on adverse selection and an empirical explanation, based on a detailed examination of the limit order book, for the observed results.

Our analysis of the limit order book shows that price clustering becomes a distinguishing characteristic of manual trader behavior only after tick size change. The price clustering phenomenon has previously been observed in equity markets (see e.g. [Huang and Stoll \(2001\)](#) and [Ohta \(2006\)](#)) and derivative markets ([Gwilym et al. \(1998\)](#)). Microstructure literature offers four conjectures as possible reasons for price clustering. According to the price resolution hypothesis proposed by [Ball et al. \(1985\)](#), traders resist quoting at higher resolution when facing greater uncertainty about an asset’s fundamental value. The negotiation hypothesis of [Lawrence \(1991\)](#) posits that a small set of prices eases the negotiation process by precluding frivolous offers and counter-offers. The attraction hypothesis of [Goodhart and Curcio \(1990\)](#) states that traders quote-cluster at certain prices due to their specific preferences. Finally, the collusion hypothesis proposes that dealers may collude to quote at larger price increments in order to get larger profits ([Christie and Schults \(1994\)](#)). Price resolution, negotiation and collusion hypotheses can be immediately rejected in our context. The attraction hypothesis is similar in spirit to what we observe. On the price clustering aspect of our analysis, this paper is different from previous studies in that we show the price clustering preference of a sub-population of market makers—the manual traders—is revealed by tick size change.

2.4 Description of Data

The data used in this study is the EBS limit order book at highest resolution available, which includes 10 levels of quotes on both the bid and ask sides at 100 milliseconds frequency for EUR/USD currency pair. This is the same as tick-by-tick snapshots of limit order book seen by traders. The deal time is rounded to the nearest 100 millisecond and only best buyer or seller initiated transactions are reported. Orders in EBS must be submitted in units of millions of the base currency. The period we analyze is the entire two year period from January 2010 to December 2011. There are approximately 500 million snapshots of the EUR/USD order book and approximately 18 million recorded deals. FX markets trade continuously and each trading day in EBS is 24 hours beginning and ending at 17:00 US Eastern Standard Time (21:00 Greenwich Mean Time). We exclude thin weekend trading periods and holidays as the liquidity tend to be extremely limited during these periods. The time stamps in the data are in GMT which varies due to daylight savings. We control for daylight savings time and standard time.⁹

⁹Similar conventions were adopted by [Andersen et al. \(2003\)](#) and [Chaboud et al. \(2004\)](#).

2.5 Jumps

2.5.1 Semi-martingale model of price

According to asset pricing theory, an arbitrage-free price process must be a semi-martingale. Following Andersen et al. (2007) and Ait-Sahalia and Jacod (2009), for econometric tractability we assume the prices, currency exchange rates in our case, follows an Itô-semimartingale. In this section, we summarize the model of price and econometric methodology.

An Itô semi-martingale is a stochastic process of the form¹⁰

$$y(t) = \alpha + \int_0^t \sigma(s)dw(s) + \sum_{i=1}^{N(t)} j_i$$

where the summand processes are independent and

1. α is a process for which almost all paths are continuous and of finite variation.
2. w is standard Brownian motion.
3. The Itô integrand σ , the *spot volatility process*, is pathwise strictly positive, càdlàg, and locally bounded away from zero.
4. $N(t)$ is a finite activity, simple counting process with, for all $t > 0$, $N(t) < \infty$ almost surely and $\{j_i\}$ is a countable family of non-zero random variables.¹¹
5. (α, σ) is independent of w .

Assumptions 1 and 2 describe general Itô-semimartingales. The strict positivity of the spot volatility process can be assumed per Assumption 3 because thin weekends and holidays are excluded from our data. Assumption 4 specifies that, with probability 1, sample paths of the price process have finitely many jumps on $[0, t]$ for all $0 < t < \infty$.¹² This assumption is empirically justified by that average time-between-trades in our data is approximately 2.5 seconds. Assumption 5 precludes leverage effect, i.e. the negative correlation between volatility and returns. The FX market is subject to factors such as central bank interventions that make the existence of leverage effect not apparent.

¹⁰Adaptedness with respect to an underlying filtration is assumed throughout this section. We suppress notation for readability. Similarly, stopping times are measurable with respect to the underlying filtration.

¹¹ $\{j_i\}$, and the inter-arrival times of $N(t)$ are therefore time series and analyzable using discrete-time techniques, as is done in Section 2.5.6.

¹²A semimartingale can have infinitely many jumps on a compact interval, e.g. an infinite activity Lévy process.

2.5.2 Estimation of volatility and jumps

Denote the jump component of $y(t)$, $\sum_{i=1}^{N(t)} j_i = y(t) - y(t_-)$, by $\Delta y(t)$. The *integrated variance* of $y(t)$ is $c(t) = \int_0^t \sigma^2(s) ds < \infty$ for all $t < \infty$.¹³ The *quadratic variation*, or *square bracket process*, of $y(t)$ is defined as

$$[y](t) = c(t) + \sum_{s \in [0, t]} (\Delta y(s))^2$$

and can be consistently estimated by *realized volatility*

$$[y](t) = \text{plim}_{M \rightarrow \infty} \sum_{j=1}^M (y(t_j) - y(t_{j-1}))^2$$

where $t_0 = 0 < t_1 < \dots < t_M = t$ are stopping times with $\lim_{M \rightarrow \infty} \sup_{1 \leq j \leq M} t_j - t_{j-1} \rightarrow 0$ almost surely.¹⁴ We use the bipower variation technology of [Barndorff-Nielsen and Shephard \(2006a\)](#) for separating integrated volatility and jump contribution.¹⁵ Define the notation

$$\mu_r = \mathbb{E}[|u|^r] = 2^{\frac{r}{2}} \frac{\Gamma(\frac{1}{2}(r+1))}{\Gamma(\frac{1}{2})}, \text{ where } u \sim \mathcal{N}(0, 1)$$

where $u \sim \mathcal{N}(0, 1)$ and Γ denotes the Gamma function. For $r \in (0, 2)$,

$$\frac{1}{\mu_r \mu_{2-r}} \{y_M\}_i^{[r, 2-r]} \xrightarrow{p} \int_{h(i-1)}^{hi} \sigma^2(u) du$$

The difference between realized volatility and bipower variation can therefore be used to detect jumps. According to [Barndorff-Nielsen and Shephard \(2006a\)](#), under the null that the sample paths have no jumps,

$$\frac{\log(\sum_{j=1}^{\lfloor \frac{t}{\delta} \rfloor - 1} y_j^2) - \log(\frac{1}{\mu_1^2} \sum_{j=1}^{\lfloor \frac{t}{\delta} \rfloor - 1} |y_j| |y_{j+1}|)}{(0.6091 \cdot \max\{\frac{\delta}{t}, \frac{\sum_{j=1}^{\lfloor \frac{t}{\delta} \rfloor - 3} |y_j| |y_{j+1}| |y_{j+2}| |y_{j+3}|}{(\sum_{j=1}^{\lfloor \frac{t}{\delta} \rfloor - 1} |y_j| |y_{j+1}|)^2}\)}\right)^{\frac{1}{2}}$$

converges in law to $\mathcal{N}(0, 1)$.

2.5.3 Microstructure noise

To choose a sampling frequency, we compare the difference between realized volatility $[y]$ and bipower variation $\frac{1}{\mu_1^2} \{y_M\}_i^{[1, 1]}$ across different frequencies. In the absence of microstructure noise, the difference consistently estimates the quadratic variation of the jump component:



¹³By Itô isometry, $c(t) = \text{Var}(\int_0^t \sigma(s) dw(s))$.

¹⁴In our specific case, we choose to sample at regular time intervals. So the stopping times are in fact deterministic.

¹⁵Informally, the bipower variation does not see the contribution of jump component because large jumps do not occur between two adjacent intervals as the intervals become sufficiently small.

$$\{y_M\}_i^{[2]} - \frac{1}{\mu_r \mu_{2-r}} \{y_M\}_i^{[r, 2-r]} \xrightarrow{p} \sum_{k=N(h(i-1))+1}^{N(h(i))} j_k^2$$

The stabilization of the difference with respect to frequency therefore indicate absence of microstructure noise. For all five major currency pairs in the EBS market, hourly averages over the entire year of 2011 are computed for both realized volatility and bipower variation. It is interesting to observe that the difference $\{y_M\}_i^{[2]} - \frac{1}{\mu_r \mu_{2-r}} \{y_M\}_i^{[r, 2-r]}$ stabilizes around the same frequency, every 30 seconds, for all five pairs. [Figure 2.1](#) shows the results for EUR/USD and [Figure 2.2](#) for USD/JPY. The magnitude of the stabilized gap, however, decreases with respect to the degree of liquidity empirically observed in the market. The most liquid currency pair, EUR/USD has the smallest difference between realized volatility and bipower variation at the stabilized frequency of 30 seconds. In other words, empirically, microstructure noise disappears at the same frequency for all five major currency pairs. The more liquid a currency pair, the more of its variation is due to the continuous part of the exchange rate process, instead of the jump component.



2.5.4 Confirmation of Stylized Facts

En route to our results on serial correlation, we confirm two stylized facts regarding the foreign exchange market. Unlike the U-shaped volatility smile of (US) equity market, the FX market intraday volatility typically is hump-shaped with two peaks (see also [Gençay et al. \(2001\)](#), Figures 6.1 and 6.2). While trading can be done at any time, peak volatility occur at the opening times of US and European markets. There may be a smaller third peak that coincides with the opening of Asian markets. Two sample days, one before and one after tick size change,

are shown in Figures 2.3 and 2.4. Furthermore, extreme high volatility –orders of magnitude larger than that of the typical intraday seasonality– tend to coincide with macroeconomic news releases or otherwise significant economic events. We show some representative instances of such news-driven high volatility periods for EUR/USD currency pair. Figures 2.5 and 2.6 show the estimated hourly integrated volatility of EUR/USD of 2010 and 2011 respectively. The integrated volatility time series show clearly visible peaks, the eight highest of which are circled in red. All eight periods of extreme high volatility coincide with significant economic news or events. Table 2.1 lists them below in chronological order, along with coincident economic events. The hourly periods of estimated extreme high volatility are given in Greenwich Mean Time (GMT), as is customary in FX market. All macroeconomic news announcements are captured on the hour. For example, the Federal Reserve announcement regarding the Fed Funds rate are consistently made within a few minutes of 2:15pm US Eastern Standard Time (EST) of FOMC meeting day, which is 18:15 GMT. The joint announcement on 2011/11/30 by the Federal Reserve, Bank of Canada, Bank of England, Bank of Japan, European Central Bank, and Swiss National Bank occurred on 8:00am EST, which is 13:00 GMT as shown in the table.¹⁶ Similarly, the Flash Crash occurred on 2:45pm EST, which is 18:45 GMT.

Table 2.1: EUR/USD Extreme Volatility Periods

Date	GMT	Corresponding economic Event
2010/05/06	18:00-19:00	The Flash Crash.
2010/05/19	14:00-15:00	Germany surprises the market by unilaterally banning short-selling against stocks and bonds, including sovereign bonds.
2010/05/20	17:00-18:00	European stocks plunge amid split over response to sovereign debt crisis.
2010/11/03	18:00-19:00	The Federal Reserve announces major quantitative easing plan to buy \$600 billion in long-term treasuries over the next eight months.
2011/06/29	12:00-13:00	The Greek parliament passes new austerity package measures amid unrest.
2011/08/09	18:00-19:00	The Federal Reserve announces it intends to keep interest rate at exceptionally low levels –between 0 to 0.25 percent– through mid-2013.
2011/09/06	08:00-09:00	The Swiss National Bank announces decision to no longer tolerate EUR/CHF below CHF 1.2.
2011/11/30	13:00-14:00	The Federal Reserve and central banks around the world announce joint policy to alleviate the Eurozone crisis.

¹⁶See, for example, <http://www.federalreserve.gov/newsevents/press/monetary/20111130a.htm> and <https://www.ecb.europa.eu/press/pr/date/2011/html/pr111130.en.html> for corresponding press releases of the Federal Reserve and European Central Bank respectively.

Figure 2.3: Intraday Seasonality of EUR/USD, 2010





Figure 2.4: Intraday Seasonality of EUR/USD, 2011



Notes: Figures 2.3 and 2.4 Intraday seasonality of EUR/USD volatility before and after tick size change.

Figure 2.5: Integrated Volatility of EUR/USD, 2010





Figure 2.6: Integrated Volatility of EUR/USD 2011



Notes: 2010-2011 estimates of hourly integrated volatility of EUR/USD before and after tick size change.

2.5.5 Jumps Before and After Tick Size Change

We first show the empirical findings from applying the jump estimation methodology outlined in [Section 2.5.1](#). For each currency pair, we estimate the size and timing of the jumps for the two years 2010 to 2011 at the frequency of every 30 seconds. In other words, given the observed sample path of $y(t)$, we estimate the realization of the jump time series $\{j_i\}$ (indexed by random arrival times of the counting process $N(t)$). [Figures 2.7](#) and [2.8](#) show the histogram of estimated realization of $\{j_i\}$ for EUR/USD, i.e. the empirical distribution of jump sizes before and after tick size change, respectively. Not surprisingly, introduction of decimal pip cuts out a neighborhood around zero from the distribution of jump sizes. [Figures 2.10](#) and [2.11](#) show the histograms of estimated inter-arrival times before and after tick size change. Both can be reasonably fitted by exponential distributions. [Figures 2.11](#) and [2.12](#) shows the daily time series of number of jumps, spanning the years 2010-2011.

2.5.6 Serial Correlation Properties of Jumps

The serial correlation properties of the jump component Δy of the currency pair EUR/USD are reflected in the serial correlation properties of the jump time series $\{j_i\}$ and inter-arrival times of jumps. We test $\{j_i\}$ for serial correlation and fit appropriate time series models to the inter-arrival times. The data generating process of jumps undergoes unambiguous structural change in serial correlation properties before and after tick size change. The reduction (complete dissipation, in the case of $\{j_i\}$) of serial correlation in both jump time series and inter-arrival times means the jump component $\sum_{s \in [0,t]} \Delta y(t)$ of exchange rate becomes significantly more martingale-like after tick size change.¹⁷

Inter-arrival times Inter-arrival times exhibit autocorrelation both before and after tick size change. Fitting auto-regressive models to the before and after inter-arrive times series yields AR(27) and AR(8) models, respectively. Therefore the autoregressive lag decreases sharply across tick size change. Diagnostic tests of residuals, shown in [Figures 2.13](#) and [2.14](#), confirm model specification of the autoregressive model. The arrivals of jumps are significantly less clustered after tick size change.

¹⁷For example, jump sizes and inter-arrival times of the compound Poisson process –a basic building block of Lévy processes– are both white noise time series.

Figure 2.7: Jump Sizes of EUR/USD, Before Tick Size Change





Figure 2.8: Jump Sizes of EUR/USD, After Tick Size Change



Notes: Figures 2.7 and 2.8 show the histograms of jump sizes before and after tick size change.

Figure 2.9: Jump Inter-Arrival Times Before Tick Size Change, EUR/USD





Figure 2.10: Jump Inter-Arrival Times After Tick Size Change, EUR/USD



Notes: Figures 2.9 and 2.10 provide the histograms of jump inter-arrival times before and after tick size change. Exponential distributions are fitted, with results shown in Table 2.2.

Figure 2.11: Daily Number of Jumps of EUR/USD, Before Tick Size Change





Figure 2.12: Daily Number of Jumps of EUR/USD, After Tick Size Change



Notes: Figures 2.11 and 2.12 illustrate the number of daily jumps before the tick size change respectively.

Figure 2.13: Jump Inter-Arrival Times, Before Tick Size Change

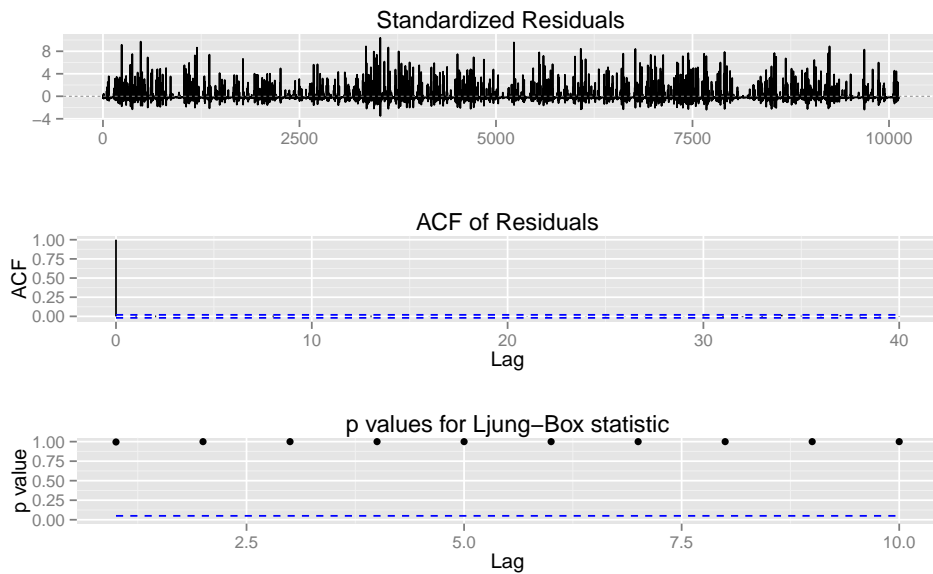



Figure 2.14: Jump Inter-Arrival Times, After Tick Size Change



Notes: Autoregressive time series fit and diagnostics for jump inter-arrival times. The Akaike Information Criterion achieves the minimum value of zero for the AR orders chosen.

Distribution of jump sizes Jump size distributions are centered near zero both before and after tick size change. The kurtosis of the distribution shows heavy tail both before and after tick size change. Nonparametric Mira test for symmetry rejects the unconditional jump size distributions before tick size change at $p < 2.2 \times 10^{-16}$ and does not reject the distributions after tick size change at 1% level of significance. Jumps sizes having a symmetric distribution centered at zero imply that, in the long run, large market orders have zero permanent price impact, implying no information content.

Jump time series Taking the time series dimension into account, the Box-Ljung serial correlation test rejects the jump sizes time series $\{j_i\}$ before tick size change at 1% significance level with $p = 0.005544$. After tick size change, the jump sizes time series has $p = 0.8185$. The time series sample sizes are 10150 and 4214 before and after respectively. The lag chosen for the Box-Ljung test is 40.¹⁸ Results are summarized in Table 2.2. A jump process driven by market taker's private information necessarily has serial correlation both in direction and arrival times. Less clustering of jump arrivals, increased symmetry and disappearance of serial correlation in jump sizes all point to a reduction in information content in jumps, in both time and direction dimensions. We can conclude that for the EUR/USD currency pair, large liquidity consumption after tick size change is much less likely to be driven by market taker's private information than before.

Table 2.2: Jump Component of EUR/USD Process Before and After Tick Size Change

Inter-arrival times	Before	After
Daily average number of jumps	33.28105	19.60094
Arrival intensity (exponential fit)	0.0003982155	0.0002353829
Autoregressive order	27	8
Jump sizes	Before	After
Mean	6.286406×10^{-6}	1.059269×10^{-5}
Standard deviation	0.000368071	0.0004490323
Skewness	1.191105	-0.1275948
Kurtosis	67.60528	4.543458
Mira symmetry test	$p < 2.2 \times 10^{-16}$	$p = 0.01928$
Box-Ljung Test	$p = 0.005544$	$p = 0.8185$
Box-Pierce Test	$p = 0.3558$	$p = 0.468$
Durbin-Watson Test	$p = 0.1767$	$p = 0.2297$

¹⁸The less powerful non-portmanteau Box-Pierce and Durbin-Watson tests are also performed. Although they are unable to reject, both p -values increase after tick size change.

2.6 Components of quoted spread

Jumps in currency prices only directly captures large liquidity consumption activity. To provide further evidence of curtailment of information asymmetry across tick size change, we examine the limit order book using the model of [Huang and Stoll \(1997\)](#). The model posits that

$$\Delta P_t = \delta + \frac{S}{2}Q_t + (\alpha + \beta - 1)\frac{S}{2}Q_{t-1} - \alpha(1 - 2\pi)\frac{S}{2}Q_{t-2} + \epsilon_t$$

where P_t is the transaction price, S is the traded spread, Q_t is the trade indicator process, 1 if buyer initiated and -1 if seller initiated. The parameter α is the portion of S due to adverse selection, β is the portion of S due to inventory risk, and π is the probability of a trade flow reversal. The trade indicator process is assumed to follow a Markov process

$$E[Q_{t-1}|Q_{t-2}] = (1 - 2\pi)Q_{t-2}$$

The error term ϵ_t contains both public information and the difference between traded spread S and the quoted spread. The latter may include, for instance, rounding error. The model is estimated using generalized method of moments (GMM). To enter a frequency where microstructure effects are explicitly present, we choose to aggregate the data every 5 seconds. [Table 2.3](#) shows the results for EUR/USD. The traded spread before the tick size change is estimated to be approximately 6.4 decimal pips, while the minimum possible quoted spread is 1 pip. This agrees with the fact that there was binding pressure on the spread prior to tick size change. Before tick size change, the quoted spread was binding approximately 60% of the time.¹⁹ The estimates of α and β sum to greater than 1 is also likely due to pressure on quoted spread. The average quoted spread is approximately twice that of estimated spread. Scaling the estimates of α , and β , before tick size change by 0.5 still shows a significant reduction in adverse selection component of spread across tick size change. After tick size change the estimated traded spread is around 8.5 decimal pips, which is very close to the average quoted spread.

[Table 2.3](#): Estimated EUR/USD components of spread, 5-seconds, Sample size = 163,069.

Components of the Spread	Before	After
α	0.85	0.17
β	1.05	0.23
π	0.21	0.31
S	0.00006	0.00008
δ	0.00014	-0.00059

All estimates are significant at 5% level of significance.

¹⁹In contrast, the binding probability of quote spread after tick size change is 1%.

The estimated adverse selection component α decreases across tick size change. Estimates of the other parameters are also of interest. The inventory risk component β also decreases across tick size change. This is as we expect, since decimalization means market maker suffers smaller loss in disposing potentially toxic order flow. The probability of trade flow reversal π increases after tick size change. Therefore trade direction becomes less persistent not just for large liquidity consumers, as shown in [Section 2.5.6](#), but for all liquidity consumers. The change in α and π are consistent with our hypothesis that tick size change attenuates the market maker's adverse selection problem.

2.7 Realized Spread

In addition to econometric evidence, we consider the empirical measure of realized spread of the order book before and after tick size change. Realized spread at time t , denoted by RS_t , is defined:

$$RS_t = 2Q_t(P_t - M_{t+s})$$

where P_t is the transaction price at time t , Q_t is the trade direction indicator, and M_{t+s} is the midpoint at time $t+s$ for some chosen time interval s . RS_t is the difference between current deal price and the quoted midpoint at a future time. After a transaction at time t , price movements favorable to the market maker from t to $t+s$ results in positive RS_t , and vice versa. According to our analysis in [Section 2.5.3](#), microstructure effect ceases to be present at frequency lower than every 30 seconds. Our computation shows that the realized spread exhibits the same behavior across tick size change at all frequencies higher than every 30 seconds, that is, under different degrees of microstructure effect.

Figure 2.15: EUR/USD Realized Spread, Five Second Frequency





Figure 2.16: EUR/USD Adverse Selection Proxy, Five Second Frequency



Notes: Realized spread is measured at 5 sec lag in Figure 2.15 and computed using 1.3 million deals before the tick size change and 7.3 million deals after. Time series of daily averages are plotted. Tick size change is demarcated by blue line. Adverse selection proxy before and after tick size change in Figure 2.16 shows clear downward level shift.

Figure 2.17: EUR/USD Realized Spread, Ten Second Frequency





Figure 2.18: EUR/USD Adverse Selection Proxy, Ten Second Frequency



Notes: Realized spread is measured at 10 sec lag in Figure 2.17 and computed using 1.3 million deals before the tick change and 7.3 million deals after. Time series of daily averages are plotted. Tick size change is demarcated by blue line. Adverse selection proxy before and after tick size change in Figure 2.18 shows clear downward level shift.

Figures 2.15 and 2.17 show the daily average realized spread of EUR/USD currency pair for 2011, before and after tick size change at frequencies of every 5 and 10 seconds. There is a clear positive shift in realized spread across tick size change. For EUR/USD, price movement for the market maker tend to be unfavorable with negative realized spread before tick size change, while realized spread is nearly zero after tick size change. The effective spread at time t is defined by $ES_t = |P_t - M_t|$. ES_t measures the revenue of the market maker from supplying immediacy. As a proxy for adverse selection of market making, we use the limit order book statistic

$$ES_t - \frac{RS_t}{2}$$

i.e. revenue from supplying immediacy minus loss due to adverse price move. Similar to that of realized spread, the behavior of the adverse selection proxy is consistent with respect to different levels of microstructure effect. Figures 2.16 and 2.18 show a clear downward shift for EUR/USD across tick size change at frequencies of every 5 and 10 seconds.

2.8 Market participant behavior

First we fix terminology by using the classification of traders in FX market given in the BIS 2011 report.²⁰ Market participants are divided into two major categories, Manual Traders and Automated Traders. Manual Traders use proprietary EBS workstations, for manual order management.²¹ Automated Traders place orders algorithmically with little or no human intervention.²² Automated Traders are capable of placing orders at a frequency far exceeding that for Manual Traders. Manual and automated market makers make up two distinct species of liquidity providers on the EBS platform. Manual market makers place limit orders for inventory or liquidity reasons, whereas their automated counterparts engage in opportunistic market making. Perhaps surprisingly, in sharp contrast with equity markets, manual trader presence dominates the interbank FX market. Using EBS client identity data, Schmidt (2012) makes the following observations:

- 75% of all traders in the EUR/USD pair are manual traders in 2011.

²⁰See <http://www.bis.org/publ/mktc05.pdf> for more details. The same classification is adopted by EBS (see e.g. Schmidt (2012) and Chaboud et al. (2014)).

²¹See <http://www.ebs.com/access-methods/ebs-workstation.aspx> for details on EBS workstations provided to Manual Traders.

²²See <http://www.ebs.com/access-methods/ebs-ai.aspx> for details on EBS interface technology for automated trading.

- The orders of manual market makers are filled in about 50% of the time before cancellation. In contrast, algorithmic market makers cancel 93% of their quotes.
- Manual market makers place large limit orders while automated market makers tend to submit orders of the minimum size one million. In fact, all orders larger than 4 million are from the manual market makers.

The clear reduction in adverse selection we observed in [Section 2.5.6](#), [Section 2.6](#), and [Section 2.7](#) can be explained by market participant behavior before and after tick size change. To infer market participant behavior, we now undertake an analysis on the evolution of shape of the limit order book. While we do not possess trader identities, a snapshot-to-snapshot inspection of the book shows clearly the behavior of automated market makers. Distributions of order sizes and quote price placements obtained from our anonymous data set are both consistent with snapshot-to-snapshot activity of automated traders and known characteristics of traders cited above. The tick size change revealed distinct preferences of the two species of market makers. Our results indicate that, while automated market makers engage in queue jumping after tick size change, manual market makers did not make use of newly available decimal prices in placing quotes.

Tick size change made one additional decimal place, the fifth, available to the market maker. The best bid and ask prices are predominantly concentrated at the old pip pricing levels after tick size change. Before tick size change, the last digits of the best bid prices are distributed uniformly as shown in [Figures 2.19](#) and [2.20](#). The clear uniform distribution shows that all market makers make equal use of available prices in placing quotes. The distribution of last digits undergoes a clear, and somewhat surprising change after tick size change. After tick size change, the last digits of best limit prices are concentrated at 0, around 30% for both the best ask and the best bid as shown in [Figures 2.21](#) and [2.22](#). While we have summarized the distribution of all snapshots of the limit order book in our data before and after tick size change respectively, the two distinct distributions are stable at the daily level. This points to a significant portion of market makers who did not adapt to decimal pip pricing.

Our analysis suggests strongly that it is the manual traders who did not adapt. [Table 2.4](#) shows the average order size at the best bid. Before tick size change, the average order size is uniformly distributed with respect to last digits, suggesting again that all traders make equal use of all available price levels when submitting quotes. After tick size change, average order size at prices with last digit zero, that is, at old pip pricing levels, is twice as large as those

at the newly available decimal levels. In fact, orders placed at the newly available decimal pip levels have average size very close to the minimum order size of one million. Table 2.5 shows very similar results for the best ask. As automated market makers tend to submit orders of minimum size, this supports our claim that manual traders price-clustered at pip pricing –and, occasionally, half-pip pricing– levels.

Figure 2.19: Best Bid Last Digits, Before Tick Change




Figure 2.20: Best Bid Last Digits, After Tick Change




Figure 2.21: Best Ask Last Digits, Before Tick Change




Figure 2.22: Best Ask Last Digits, After Tick Change




Table 2.4: Size of Average Order (Million) at the Best Bid

Last digit of quoted price	0	1	2	3	4	5	6	7	8	9
Before tick size change	1.59	1.43	1.57	1.52	1.54	1.61	1.54	1.60	1.50	1.41
After tick size change	2.12	1.06	1.05	1.08	1.08	1.48	1.15	1.11	1.07	1.10

Table 2.5: Size of Average Order (Million) at the Best Ask

Last digit of quoted price	0	1	2	3	4	5	6	7	8	9
Before tick size change	1.53	1.32	1.62	1.59	1.45	1.54	1.53	1.54	1.43	1.48
After tick size change	2.19	1.23	1.22	1.10	1.09	1.32	1.06	1.04	1.07	1.09

Automated market makers, on the other hand, have no reason to not take advantage of decimal pip pricing. Indeed, we show that automated market makers engages in queue jumping. Queue jumping exploits the likely favorable price movement precipitated by a large static order, which also limits the loss for the queue jumper under unfavorable price movement. Static limit orders placed by manual traders at old pip pricing, with last digit zero, are vulnerable to queue jumping strategy by automated traders. On the ask side, such queue jumping is done by submitting an order with last digit 9 and the buy side by an order with last digit 1. Looking at ask side through the course of one trading hour, 36,000 snapshots of limit order book, at tick frequency, after tick size change shows that 43% of limit orders are at pip pricing levels (see [Table 2.6](#)). Conditional on the best ask, the probability that the next snapshot shows one tick size improvement is 55%, with average order size around minimum order size. One or two tick size improvement makes up 70% of next-snapshot possibilities at the best ask, all with around average minimum order size. Similar patterns can be seen on the buy side of the book. As time between snapshots is 100 milliseconds, this directly exposes queue jumping activity by automated traders.

[Table 2.6](#): Queue Jumping in One Trading Hour After Tick Size Change

Number of limit order book snapshots	36,000
Best ask at old pip pricing	15,490
One tick price improvement next snapshot	8,485 (avg order size 1.06)
Two tick price improvement next snapshot	2,289 (avg order size 1.04)
Best bid at old pip pricing	14,668
One tick price improvement next snapshot	7,767 (avg order size 1.03)
Two tick price improvement next snapshot	2,055 (avg order size 1.02)

The decreasing pattern shown in [Figure 2.20](#) of last digit distribution of best ask price can be explained by the interaction we describe between traders. An initial static quote might be placed at a pip price by a manual traders. Automated traders queue-jump first the manual trader quote then leap frog each other as they compete for top of the book. [Figure 2.22](#) tells the same story on the other side of the book. The already dramatic inequality in speed between manual and automated traders is accentuated by the tick size change. Clinging to old pip pricing makes limit orders submitted by manual market makers prey to the queue jumping strategy. The resulting delay in order execution compels manual market makers to cross the spread and pay their algorithmic counterparts a premium for timely execution. Decimalization of tick size therefore changes the information balance between market makers and takers in the market by effectively forcing a sub-population of market makers into market taking.

2.9 Conclusions

Analyzing both the price process and the limit order book, we showed that the degree of information asymmetry exhibited a discrete change across tick size change and provided an explanation on the causality relationship between these two events. A smaller tick size has two textbook counteracting effects on volume of informed order flow. As a larger minimum tick means the spread is more likely to bracket the fundamental value, decimalization incentivizes the market taker to become more informed. On the other hand, as smaller tick size encourages predatory market making and makes informed trading less profitable, the market taker has less incentive to acquire information. However, the ecology of the interdealer FX markets, with no evidence of trader emigration or immigration across tick size change, gives rise to a unique response driven by interaction between two distinct subspecies of market makers. Automated market makers crowd out their manual counterparts from the top of the book and force them to cross the spread.

Current literature concerning the issue of the effect of algorithmic trading on market efficiency for the FX market have generally been favorable. For example, using EBS data from 2004 to 2008, [Chaboud et al. \(2014\)](#) showed empirically that high frequency returns are serially uncorrelated and conclude that algorithmic trading activity improves market efficiency. There has also been some evidence that algorithmic trading contributes to the speeding up of the price discovery of exchange rates with respect to macroeconomic news (see e.g. [Andersen et al. \(2003\)](#) and [Faust et al. \(2007\)](#)), which is partially corroborated by our analysis of periods of extreme volatility shown in [Table 2.1](#).²³ To the best of our knowledge, this is the first paper that considers market efficiency and algorithmic trading from a microstructure perspective by analyzing comprehensively the impact of a specific microstructure event. In contrast to previous studies, our findings do not support the position that algorithmic trading improves market efficiency. In our high frequency continuous-time setting, the price semi-martingale is the independent sum of a finite-variation process, a continuous martingale, and a jump process. In other words, the exchange rate has a locally riskless component, an informationally efficient component, and a jump component. While the unequivocal whitening of the jump component moves exchange rate closer to the martingale property that characterizes an informational efficient market ([Fama and French \(1988\)](#)), this is not a verification of market efficiency—despite observing uniform

²³[Chaboud et al. \(2014\)](#), [Andersen et al. \(2003\)](#) and [Faust et al. \(2007\)](#) all use discrete time models and sample at lower frequencies. According to EBS itself: automated trading “...is a key component of the professional spot FX market place, offering efficient price discovery and 24-hour access to tight liquidity...”—<http://www.ebs.com/access-methods/ebs-ai.aspx>.

reduction in traditional measures of adverse selection in addition to whitening of jumps. This abatement of adverse selection is an artifact of market microstructure, rather than genuine dissemination of information across market participants. Our study highlights that traditional microstructure metrics, though often indicative of the general market conditions, need to be complemented by analysis of market participant behavior to obtain a meaningful picture of the state of market, especially given the increasing pervasiveness of high-speed electronic trading.

Chapter 3

Unprofitable Arbitrage, Speed Is Not of the Essence

3.1 Abstract

This paper studies the properties of triangular arbitrage opportunity in the interbank foreign exchange (FX) market. We examine trading activity from Electronic Broking Services (EBS), the main interdealer spot FX platform for major currencies. EBS limit order book data of the entire five year period from 2004 to 2008 at tick frequency (second-by-second) is analyzed. Our results cast into question current understanding of triangular arbitrage in the literature, specifically in relation to algorithmic trading. For microstructure reasons, The increasing presence of algorithmic traders does not offer significant improvement in speed of price discovery by quickly consuming the triangular arbitrage opportunities. A single trader cannot exploit triangular arbitrage opportunity in the interdealer market. Rather, algorithmic trading influences the creation of triangular arbitrage by two countervailing effects. First, the narrowing of the spread reduces the price improvement necessary to create triangular arbitrage. Second, the minimum tick size, if it exceeds the necessary price improvement, prevents triangular arbitrage from being created by an algorithmic trader or otherwise. Our results show that the second effect dominates the first.

3.2 Introduction

The prevalence of algorithmic trading have spurred a growing body of literature studying its impact on various aspects of market behavior. [Hendershott et al. \(2011\)](#) show that algorithmic trading enhances the informativeness of quotes and improves liquidity in NYSE limit order books. They also find that for large stocks in particular, algorithmic trading narrows spreads, reduces adverse selection and trade-related price discovery. Using publicly available NASDAQ data, [Hasbrouck and Saar \(2013\)](#) concludes that increased algorithmic trading activity lowers short-term volatility, decrease spreads, and increase depth on. Also on NASDAQ, [Brogaard et al. \(2012\)](#) find that algorithmic traders enhances price efficiency. Not all commentaries on the effect of algorithmic trading have been positive. [Hirschey \(2013\)](#) finds that the algorithmic traders on NASDAQ use their speed to step ahead of future order flow at the expense of slower traders. [Biais and Woolley \(2011\)](#) discusses some tactics used algorithmic traders for market manipulation, e.g. “stuffing”, “smoking”, and “spoofing”.

Because no arbitrage conditions in the FX market are relatively easy to define, the literature has found them useful as measures of market efficiency in empirical investigations. Triangular arbitrage means taking advantage of mis-pricing of three currencies by completing a three-leg round trip trade. The decreasing frequency of triangular arbitrage opportunities in interbank FX market has been consistently cited as evidence that algorithmic trading improves market efficiency. However, we find that the impact of algorithmic trading is effectively blunted by market microstructure. The market reacts the same way to the appearance of triangular arbitrage opportunities as algorithmic trader presence increased from 2004 to 2008. There is no discernible decreasing trend in conditional distribution of neither durations nor profits of triangular arbitrage opportunities from 2004 to 2008. An individual trader, algorithmic or otherwise, cannot exploit triangular arbitrage opportunity in the interdealer market due to the severe left over problem. In fact, for the market collectively, our analysis shows that very little exploitation of triangular arbitrage occurs by either manual or algorithmic traders. Triangular arbitrage opportunities are most likely removed by market maker’s quote adjustment, although market maker’s reaction time is not necessarily faster than before. To the extent that trading occurs between the appearance and disappearance of triangular arbitrage opportunities, it occurs in a single currency pair. The same pattern describes the market reaction to triangular arbitrage in 2004, when algorithmic trading is extremely minimal, and persists after 2005, when algorithmic

trading is introduced into EBS. Therefore, contrary to claims made by existing literature, algorithmic trading does not expedite the elimination of triangular arbitrage opportunities after they appear. Neither does the increasing presence of algorithmic traders lead to improvement in the speed of price discovery.

The data does show, however, a gradual downward trend in the number of triangular arbitrage opportunities.¹ Rather than the elimination of triangular arbitrage *ex post*, our results shows one should focus on the effects of algorithmic trading on the creation of triangular arbitrage. Two countervailing effects exist. First, algorithmic trading activity indirectly reduces the price improvement necessary to create triangular arbitrage. After the introduction of algorithmic trading, a triangular arbitrage opportunity is more likely to be caused by a market maker placing quotes inside the spread, instead of a market order on the other side of the limit order book. On the other hand, as the prices becomes more easily perturbed into triangular arbitrage, the minimum tick size –or pip, in FX vernacular– becomes a barrier to the creation of triangular arbitrage.² If the spread equal one pip and the price improvement necessary to create triangular arbitrage is less than one pip, then no limit order, placed by either manual or algorithmic trader, can lead to a triangular arbitrage opportunity. Our results show that the second effect slightly dominates the first, leading to a gradual decrease in triangular arbitrage opportunities. As the increasing presence of algorithmic traders move exchange rates closer to the knife-edge condition of no arbitrage, they also move rates to microstructure barriers to the creation of arbitrage.

Most commonly used data sets in studying triangular arbitrage come from the two major FX platforms, EBS or Reuters.³ The microstructure obstructions pointed out above in examining this issue are entirely ignored in the literature.⁴ The mere fact that triangular arbitrage opportunities in the FX market are not profitable and no attempt is made to profit from them, as we will show, makes such conclusions questionable. For example, while we agree with [Chaboud et al. \(2014\)](#) that algorithmic trading leads to a reduction in the number of triangular arbitrage

¹When counting the number of triangular arbitrage opportunities, we take their durations into account. For example, if triangular arbitrage of the same route appears for 10 consecutive seconds, it is counted as a single opportunity with a duration of 10 seconds. Alternatively one can count 10 triangular arbitrage opportunities, each of one second. The latter approach is common in the literature and would show a more dramatic reduction in triangular arbitrage opportunities (see e.g. [Chaboud et al. \(2014\)](#)). We view this as misleading, in light of our results.

²“Pip” is abbreviation for Price Increment Point.

³In terms of comprehensiveness, our data set is comparable to [Chaboud et al. \(2014\)](#) and Ito et al (NBER). [Chaboud et al. \(2014\)](#) uses EBS data from 2003 to 2007, while the sample period of Ito et al is 1999 to 2010.

⁴Trades made on the Reuters platform must also be in units of millions of base currency and face similar tick size constraints. The tick size on Reuters was simialr to EBS equal to 1 pip.

opportunities, our analysis contradicts the unsupported conclusion that algorithmic traders in the FX market detect triangular arbitrage opportunities to profit from them.⁵ Marshall et al (2008 AFA meeting paper) concludes from 2005 EBS data that profitable triangular arbitrage opportunities exists. In casting their analysis in the context of Grossman-Stiglitz paradox, they assumed that traders do in fact actively attempt to exploit triangular arbitrage, without examining actual trader behavior. Kozhan and Wah Tham (2012) offer a model of execution risk faced by high frequency traders in exploiting triangular arbitrage, with empirical results based on EBS 2004 data. The underpinning assumption is again that the triangular arbitrage opportunities are profitable, but subject to risky execution.⁶

The same unique microstructure frictions in making linkages to conclusions regarding arbitrage, in different empirical contexts or theoretical models. In a model of arbitrage between different asset classes, Oehmke (2012) shows that a more liquid market has higher competition among arbitrageurs, which should in turn leads to less arbitrage opportunities. Also in a setting of multiple asset, Kondor (2009) reaches a similar conclusion; more arbitrageur activity turns riskless arbitrage into risky bets. This points to needs for accounting for microstructure empirically, theoretically, and in linking the two strands.

The remainder of the paper is organized as follows: Section 3.3 describes the data and the EBS market. Section 3.4 gives the definition of triangular arbitrage opportunity and shows the number arbitrage opportunities during the five year period examined. Section 3.5 discusses the durations and profitability in the absence of microstructure friction –with particular attention on the effect, or lack thereof, of increasing presence of algorithmic trading from Jan 2004 onwards. Section 3.6 computes the severity of the left-over problem and shows triangular arbitrage round trip trades are never carried out by examining the limit order books between the appearance and disappearance of triangular arbitrage opportunities. Section 3.7 shows that, as algorithmic traders move prices closer to triangular arbitrage, the pressure on the spread due to tick size becomes a barrier to triangular arbitrage. Section 3.8 concludes.

⁵The left over problem is alluded to in Chaboud et al. (2014) in a footnote but ignored in their analysis.

⁶Kozhan and Wah Tham (2012) assumes that an (high frequency) arbitrageur exploit triangular arbitrage by continuously acquire growing inventories, which he then unwinds at some arbitrary time later. This contradicts both microstructure theory and empirical evidence. Risk averse market makers are inventory neutral due to inventory risk (see e.g. inventory risk (Stoll (1978), Amihud and Mendelson (1980), Ho and Stoll (1981)). Moreover, the spot rate, say, 20 minutes, from the time left over inventory is acquired is unknown.

3.3 Data Description

The spot foreign exchange (FX) market daily turn over reached to 1.5 trillion USD in 2013, approximately 7 times that of equity markets. The spot market makes up 37% of the global FX market. Other currency related instruments include FX swaps, forwards, currency swaps, and options. A significant portion of this volume, most recently estimated at 35%, are interdealer trades. The US dollar is the dominant vehicle currency and it was on one side of 87% of all trades in April 2013. The euro remains the second most important currency worldwide. The Japanese yen significantly expanded its share in global FX trading.⁷ Banks trade currencies with each other on two interdealer electronic trading platforms, namely, EBS and Thomson Reuters. The decision whether to use EBS or Reuters is usually driven by the currency pair. 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.

The data set used in this study is the EBS limit order book from 2004 to 2008, which includes best bid, best offer, and the deal prices at one second frequency. The deal prices are the highest buying deal price and the lowest selling deal price between two snapshots rounded to the second. It is important to mention that the refreshing rate of the limit order book was at one second frequency. This means that traders had no information about the order book information between two consecutive snapshots. Orders are submitted in units of millions of the base currency.⁸ EBS defines the currencies based on the base and local currencies, and EUR/USD denotes the amount of local currency USD required to buy (or sell) one unit of the base currency EUR. The market is continuous, however, we have excluded thin weekend trading periods and holidays because the liquidity may have been extremely limited. The time stamps in the data are in GMT which varies due to daylight savings. We controlled for daylight savings time (summer) and standard time (winter). Similar conventions were adopted by Andersen et al. (2003) and Chaboud et al. (2004).

The tick size, i.e. minimum price movement, was a pip equal to 0.0001 on all major currency pairs from 2004 to 2008. If the tick size is 0.0001 and the EUR/USD best bid is 1.39940 , then a buyer could improve this price by placing an order with the price of 1.39950 .⁹

There are two main types of traders in the EBS market: “automated” traders, who use an

⁷See the report of Bank for International Settlements (BIS): <http://www.bis.org/publ/rpfx13fx.pdf>.

⁸EBS has also introduced small currencies in 2011 where the minimum trade is the hundred thousands of units of base currency.

⁹The pip tick size is 2 decimal points when the local currency is JPY, as the exchange rates are based on 100 JPY.

automated interface (AI) to place orders without human intervention; and “manual” traders, who use GUI-based access for order management. Manual 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 in the provision of liquidity. The main component of an AI in the EBS market is the professional trading community (PTC), which places orders at very high frequency. EBS first allowed AI into the market in 2004. The activities of these algorithmic traders has increased since then.

3.4 Occurrence of Triangular Arbitrage

Classical arbitrage exploits difference between exchange rates and the rates implied by no-arbitrage conditions.¹⁰ Consider the situation where one initially holds x_i EUR. If he sells the EUR and buys USD, converts these dollars into JPY and finally converts the JPY into x_f EUR then if $x_f > x_i$ a triangular arbitrage profit is realized.

Suppose there are three currencies, C_1 , C_2 and C_3 in the market. They can create the following six situations for arbitrage opportunities. For example, in the first equation, a trader could sell C_1 to buy C_2 , sell C_2 to buy C_3 and finally sell C_3 to buy C_1 . If the amount of currency C_1 at the end is higher than the starting amount, there would be a triangular arbitrage opportunity.

$$\begin{aligned} C_1 \rightarrow C_2 \rightarrow C_3 \rightarrow C_1, & \quad C_1 \rightarrow C_3 \rightarrow C_2 \rightarrow C_1 \\ C_2 \rightarrow C_1 \rightarrow C_3 \rightarrow C_2, & \quad C_2 \rightarrow C_3 \rightarrow C_1 \rightarrow C_2 \\ C_3 \rightarrow C_1 \rightarrow C_2 \rightarrow C_3, & \quad C_3 \rightarrow C_2 \rightarrow C_1 \rightarrow C_3 \end{aligned}$$

Considering the above equations, the last four situations are just different forms of the first two equations. For example, in the third equation, if we start from C_1 , we are repeating the second equation. As we mentioned in the previous section, C_1/C_2 denotes the amount of local currency C_2 required to buy (or sell) one unit of the base currency C_1 in EBS market. As a result, C_1/C_2 would be different from C_2/C_1 when finding the arbitrage opportunities. If a trader has C_1 and want to exchange for C_2 , he can either sell C_1 or buy C_2 . These are not the same in LOB. Selling C_1 means he goes to bid side, buying C_2 means he goes to ask side. The base currency of the book determines which side he goes to. For example if the currency pairs C_1/C_2 , C_2/C_3 and C_3/C_1 are available for the trade, then we have to check two arbitrage

¹⁰Classical arbitrage in FX market consists of two types: triangular arbitrage and covered interest rate (CIR) arbitrage. In this paper we consider the former. The latter is similar in spirit and involve trading two currencies to exploit disparity between domestic and foreign interest rates while hedging against exchange rate risk by a forward contract.

conditions, buy–buy–buy ($C_3/C_1 * C_2/C_3 * C_1/C_2$) and sell–sell–sell ($C_1/C_2 * C_2/C_3 * C_3/C_1$). As an instance, in the first condition a trader can start with C_1 to buy C_3 , use C_3 to buy C_2 and finally use C_2 to buy C_1 back. However, if C_1/C_2 , C_2/C_3 and C_1/C_3 are available for trades, then we have to check sell–sell–buy ($C_1/C_2 * C_2/C_3 * C_1/C_3$) and buy–buy–sell ($C_1/C_2 * C_2/C_3 * C_1/C_3$).

The major currencies EUR/USD, USD/JPY, EUR/JPY, USD/CHF and EUR/CHF provides the following arbitrage opportunities:

$$EUR/USD - USD/JPY - EUR/JPY \quad (3.1)$$

$$EUR/CHF - EUR/JPY - CHF/JPY \quad (3.2)$$

$$CHF/JPY - USD/CHF - USD/JPY \quad (3.3)$$

We will provide the results based on the EUR–USD–JPY.¹¹ The combinations of these three currency pairs could create two routes (we would call them route 1 and route 2 afterwards) for triangular arbitrage opportunities. These routes can be identified through [Equations \(3.4\)](#) and [\(3.5\)](#). In [Equation \(3.4\)](#), we calculate whether a trader could make a profit initially by selling USD and buying EUR, selling EUR buying JPY, and finally buying USD with selling JPY. In [Equation \(3.5\)](#), there could be a profit by selling JPY and buying EUR, selling EUR buying USD, and finally buying JPY with selling USD.

$$\left(r_b^{EUR/USD} * r_b^{USD/JPY} \right) / r_s^{EUR/JPY} \quad (3.4)$$

$$r_b^{EUR/JPY} / \left(r_s^{EUR/USD} * r_s^{USD/JPY} \right) \quad (3.5)$$

When counting the number of triangular arbitrage opportunities, we can take their durations or number of seconds into account. If triangular arbitrage of the same route appears for 10 consecutive seconds, it is counted as a single opportunity with a duration of 10 seconds. Alternatively one can count 10 triangular arbitrage opportunities, each of one second. The latter approach is commonly in the literature and would show a more dramatic reduction in triangular arbitrage opportunities. [Figures 3.1](#) and [3.2](#) illustrate the daily frequency of arbitrage opportunities for route 1 and route 2 respectively considering the duration. The frequency of arbitrage opportunities decreases, albeit slowly, from 2004 to 2008. [Figures 3.3](#) and [3.4](#) indicates the number of daily arbitrage counting the seconds that there is arbitrage. The second type of the arbitrage counting indicates suggest the stronger relation between the increasing activities by algorithmic trading and reduction in the number of arbitrages.

¹¹The full results are available upon request.

Figure 3.1: Daily Arbitrage by Duration, EUR-USD-JPY Rout 1




Figure 3.2: Daily Arbitrage by Duration, EUR-USD-JPY Rout 2




Figure 3.3: Daily Arbitrage by Seconds, EUR-USD-JPY Rout 1





Figure 3.4: Daily Arbitrage by Seconds, EUR-USD-JPY Rout 2



3.5 Profit and Duration of Triangular Arbitrage

This section provides and examines the magnitude of the triangular arbitrage opportunities meaning if there is an arbitrage, how long it takes. Figure 3.5 shows the histograms of arbitrage durations for route 1 and route 2 of EUR-USD-JPY from 2004 to 2008. The x axis labels the arbitrage duration in seconds. The y axis indicates the year from 2004 to 2008. The z axis gives the number of arbitrage opportunities in a given year with given duration.

Figure 3.5: Triangular Arbitrage Duration EUR-USD-JPY

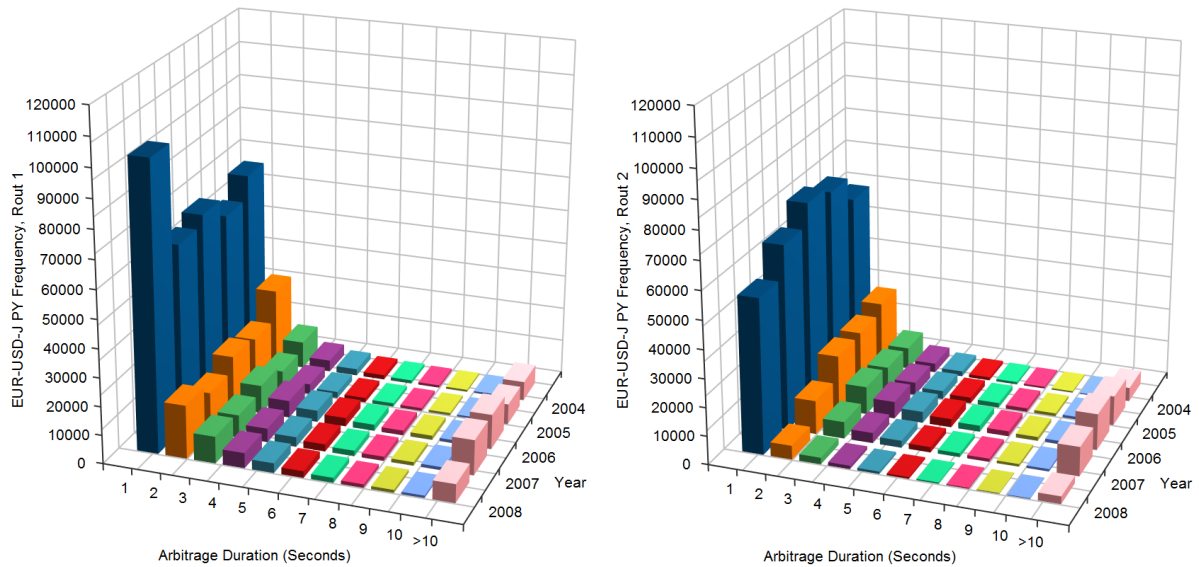
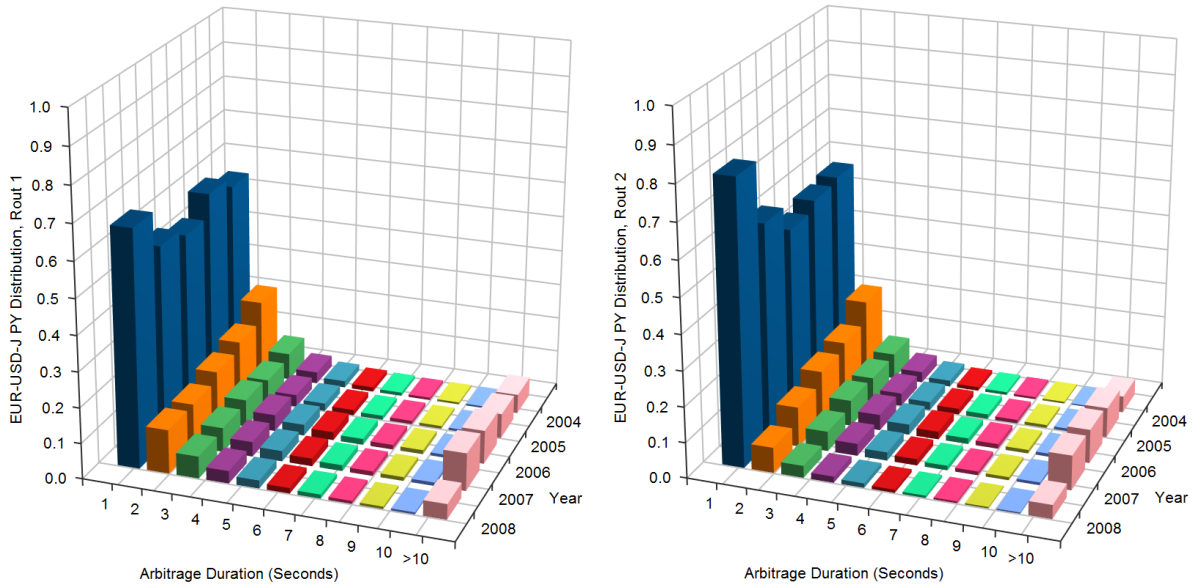


Figure 3.6 shows frequency of arbitrage durations from 2004 to 2008. No discernible pattern emerges from the distribution of Figure 3.6 to justify the claim that increased algorithmic trading activity shorten the duration of arbitrage opportunities, conditional on their appearance. For route 1, the proportion of 1-second durations and ≥ 10 second durations actually experienced decrease and increase respectively from 2004 to 2007.¹² That is, the proportion of shortest durations actually decreased and that of longest duration increased. For route 2, the proportion of 1 second durations remain approximately constant and that of longest duration increased during the same period. Thus limit book data provides no evidence that algorithmic trading contributes to price efficiency by quickly eliminating triangular arbitrage opportunities when they appear.¹³

¹²Due to the financial crisis in 2008, all aspects of the market behave accordingly including the existence of arbitrage opportunities.

¹³The fraction of trading volume involving algorithmic trading grows from near %3 in 2004 to near %60 by the end of 2007 for EUR/USD. See Chaboud et al. (2014) for details.

Figure 3.6: Triangular Arbitrage Duration Distribution EUR-USD-JPY



Next we examine the profitability of triangular arbitrage opportunities in the absence of microstructure impediment to their exploitation. The histogram of arbitrage profits are given in Figure 3.7 for route 1 and route 2 respectively. Figure 3.8 shows the distribution for the two routes in the same order. In order to have a useful distribution, we round them to the nearest basis point. 1 basis point profit is 0.0001 dollar per dollar of currency being arbitrated. Generally speaking, the profits are clustered at 1 or 2 basis point. For route 1, there is a significant drop in the proportions of 1 basis point profits, the smallest profit possible, between 2004-2005 and 2006-2008. The timing of the drop does not align with the initial influx of algorithmic traders into the market in 2005. The direction of the drop also runs counter to the claim that arbitrage profit is reduced by competition among algorithmic traders for that profit. For route 2, the distribution of profits remains approximately constant during the five years.

Overall decrease in the number of arbitrage opportunities, which is arguably contributable to algorithmic trading activity, does not coincide with the decrease in their duration or magnitude, conditional on their appearance. This leads to the question on whether arbitrage opportunities are in fact profited by traders. We consider this issue next.

Figure 3.7: Triangular Arbitrage Profit EUR-USD-JPY

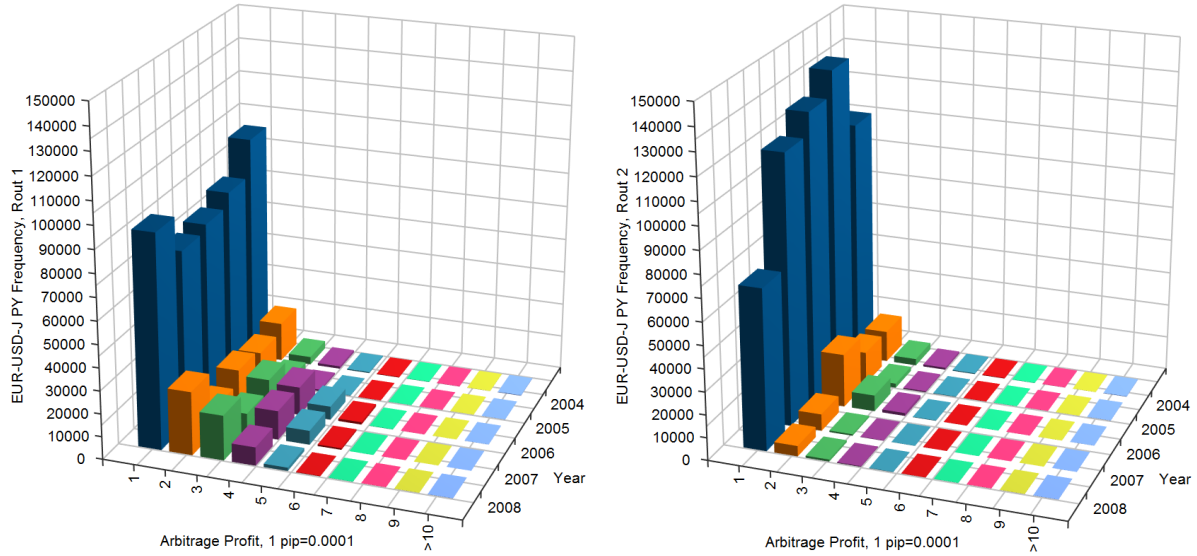
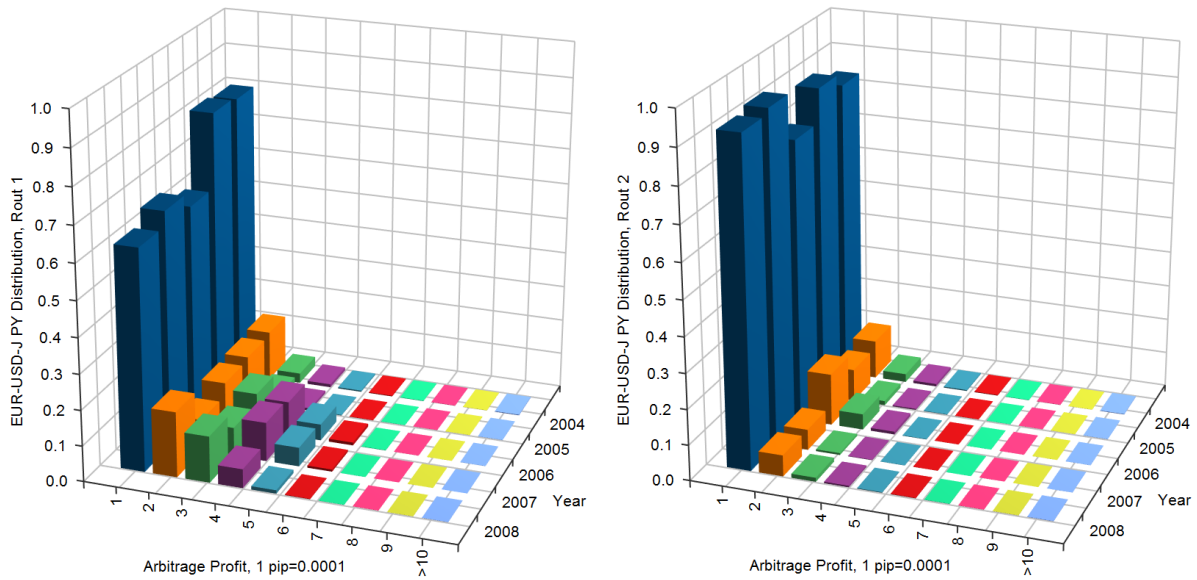


Figure 3.8: Triangular Arbitrage Profit Distribution EUR-USD-JPY



3.6 Non-profitability

3.6.1 Left over problem

In the EBS market, orders are submitted in units of millions of the base currency. Orders must have a size of $K \cdot 10^6$, where K is a positive integer. Since most arbitrage opportunities occur with highest possible degree of decimalization, one or two tick sizes, the non-decimalized order block size requirement means that any trader attempting to arbitrage must confront the left-over problem. For example if we revisit the scenario of route 1, the arbitrageur should buy KV units of USD in the second transaction, which are not necessarily equal to the exact amount of his EUR holdings.

Table 3.1: Size of Left Overs, Rout 1

	Initial	EUR	USD	JPY	Ave.Profit
2004	1 million	—	—	—	—
	2 million	-1000000	477553.70	82414687	104.14785
	3 million	-1000000	680440.47	60384674	156.22178
	4 million	-1000000	647922.85	63392694	208.29570
2005	1 million	—	—	—	—
	2 million	-1000000	530664.00	79991956	115.44232
	3 million	-1000000	769248.00	54178125	173.16348
	4 million	-1000000	447144.75	87656828	230.88464
2006	1 million	—	—	—	—
	2 million	-1000000	477553.70	82414687	104.14785
	3 million	-1000000	680440.47	60384674	156.22178
	4 million	-1000000	647922.85	63392694	208.29570
2007	1 million	—	—	—	—
	2 million	-1000000	724809.69	753476910	322.13921
	3 million	-1000000	204405.74	136730024	483.20881
	4 million	-1000000	449619.38	108045219	644.27841
2008	1 million	—	—	—	—
	2 million	-1000000	326567.13	121184668	225.79432
	3 million	-1000000	601888.96	930253070	338.69148
	4 million	-1000000	362351.05	117367533	451.58864

The average profits and left overs are given in [Tables 3.1](#) and [3.2](#) with different order sizes for rout 1 and rout 2 respectively. For example, considering rout 1 and 2 million initial EUR, the average profit is around \$100. However, the arbitrageur should increase his inventory significantly to make such profit. Comparing the results for both routs, there is more average profit with higher left over in rout 1 and lower average profit with lower left over in rout 2. [Kozhan and Wah Tham \(2012\)](#) assumes that an arbitrageur exploit triangular arbitrage by continuous acquire growing inventories which he then unwinds at some arbitrary time later. This assumption is erroneous since the arbitrageur does not have information about the future dynamic of the market.

3.6.2 Non-exploitation

If traders, algorithmic or otherwise, actively exploit triangular arbitrage opportunities, the required round-trip three-leg trades can be observed between the appearance and disappearance of triangular arbitrage opportunity in the three limit order books. During the 1 second between

Table 3.2: Size of Left Overs, Rout 2

	Initial	EUR	USD	JPY	Ave.Profit
2004	1 million	0	242032.40	26114030	53.95654
	2 million	0	484064.80	52228060	107.9131
	3 million	0	676120.29	73174846	161.8696
	4 million	0	647227.84	70627534	215.8262
2005	1 million	0	237309.88	26141579	43.736642
	2 million	0	474619.76	52283158	87.473283
	3 million	0	691618.48	76312010	131.20992
	4 million	0	576933.65	65334339	174.94657
2006	1 million	0	264274.93	30475331	69.232966
	2 million	0	528549.86	60950661	138.46593
	3 million	0	792071.30	91339219	207.69890
	4 million	0	22532946	22532946	276.93186
2007	1 million	0	379727.43	44392470	38.190957
	2 million	0	759454.85	88784939	76.381913
	3 million	0	337333.34	39487697	114.57287
	4 million	0	518909.7	60225846	152.76383
2008	1 million	0	464182.06	48845082	43.017989
	2 million	0	603551.82	63893994	86.035977
	3 million	0	502394.69	52330931	129.05397
	4 million	0	539525.50	57518599	172.07066

snapshots, there are 64 possibilities jointly for the three limit order books. Each book has 4 possibilities –buy, sell, both, or no trade. We construct the 6 digit code using EUR/USD_{buy} , EUR/JPY_{sell} , USD/JPY_{buy} , EUR/USD_{sell} , EUR/JPY_{buy} , USD/JPY_{sell} . The first three books are related to the rout 1 and second three books indicate the rout 2. If we observe a trade in a specified direction, we put “1” otherwise “0”. In this way, we would have $2^6 = 64$ possibilities for the three limit order books. Such detail separation between trade possibilities is crucial to observe the arbitrage cleaning process by trade. If there is an arbitrage opportunity in rout 1 and a trade clear it, we should observe $111xxx$ most likely, where x could be “0” or “1”.¹⁴ Figure 3.9 provides the distribution of 2004 deals for rout 1 which does not show such a pattern. If there is an arbitrage in rout 2, we should observe $xxx111$ when it disappears by triple trades.¹⁵ However, there is no such a pattern in Figure 3.10. Distribution for other years reveal the same fact.

¹⁴ $111000, 111001, 111010, 111011, 111100, 111101, 111110, 111111$

¹⁵ $000111, 001111, 010111, 011111, 100111, 101111, 110111, 111111$

This means that the algorithmic trading did not increase the occurring of the triple deals when arbitrage disappears. The arbitrage opportunities are not profited by traders which is shown by the lack of triple deals. This also confirms why arbitrage duration did not change conditional on their appearance. Therefore, the increasing presence of algorithmic traders does not offer significant improvement in speed of price discovery by quickly consuming the triangular arbitrage opportunities. There is a need for different reasoning out of existing literature to explain the overall decrease in the number of arbitrage opportunities, when algorithmic trading activity has increased.

Figure 3.9: After Arbitrage Trade Codes 2004, Rout 1




Figure 3.10: After Arbitrage Trade Codes 2004, Rout 2




Figure 3.11: Before Arbitrage Trade Codes 2004, Rout 1




Figure 3.12: Before Arbitrage Trade Codes 2004, Rout 2




Figure 3.13: After Arbitrage Trade Codes 2005, Rout 1




Figure 3.14: After Arbitrage Trade Codes 2005, Rout 2




Figure 3.15: Before Arbitrage Trade Codes 2005, Rout 1

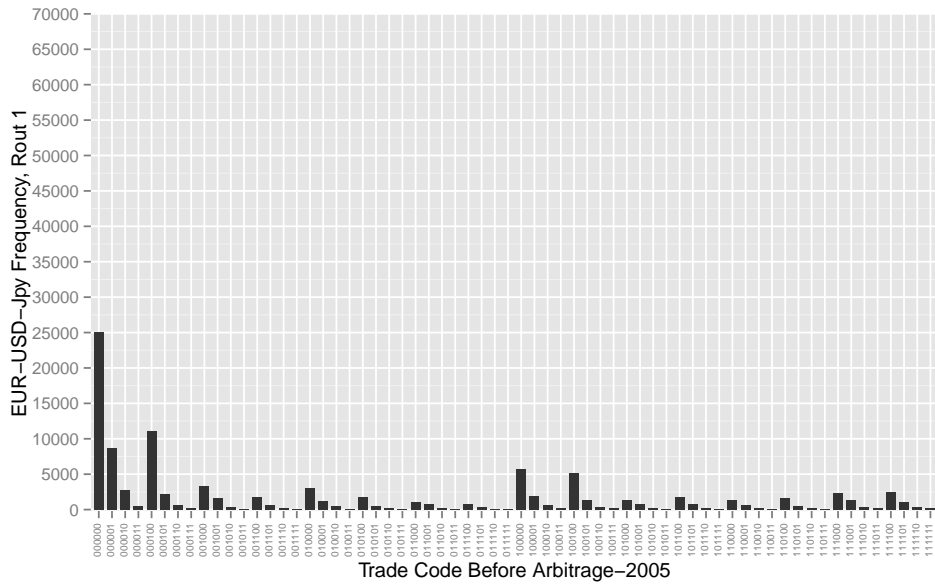


Figure 3.16: Before Arbitrage Trade Codes 2005, Rout 2




Figure 3.17: After Arbitrage Trade Codes 2006, Rout 1




Figure 3.18: After Arbitrage Trade Codes 2006, Rout 2




Figure 3.19: Before Arbitrage Trade Codes 2006, Rout 1




Figure 3.20: Before Arbitrage Trade Codes 2006, Rout 2




Figure 3.21: After Arbitrage Trade Codes 2007, Rout 1




Figure 3.22: After Arbitrage Trade Codes 2007, Rout 2




Figure 3.23: Before Arbitrage Trade Codes 2007, Rout 1




Figure 3.24: Before Arbitrage Trade Codes 2007, Rout 2




Figure 3.25: After Arbitrage Trade Codes 2006, Rout 1




Figure 3.26: After Arbitrage Trade Codes 2008, Rout 2

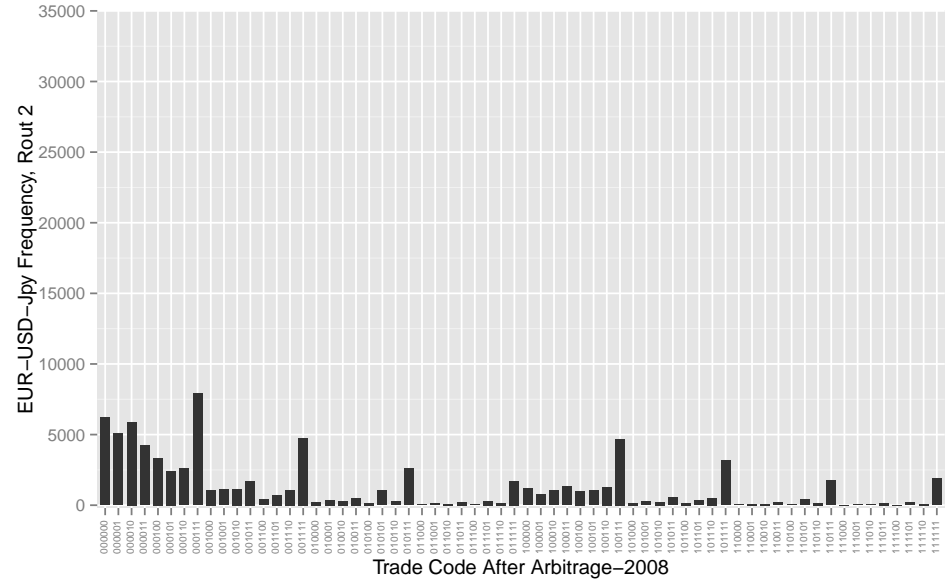


Figure 3.27: Before Arbitrage Trade Codes 2006, Rout 1





Figure 3.28: Before Arbitrage Trade Codes 2008, Rout 2

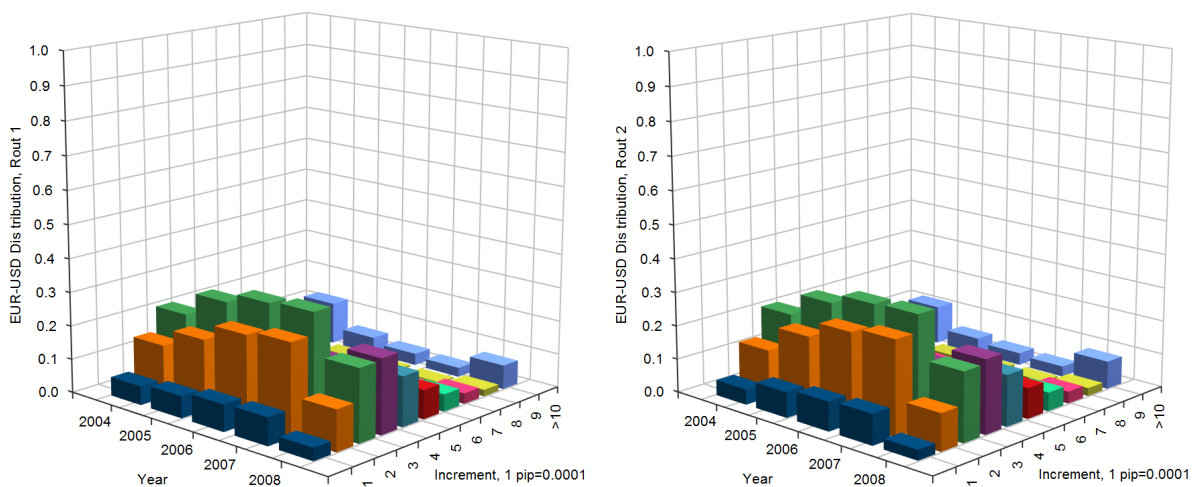


3.7 Tick Size Effect

3.7.1 Perturbation Toward Triangular Arbitrage

The results presented in the last section shows that the algorithmic trading does not affect the elimination of triangular arbitrage significantly. Therefore, we focus on the creation of triangular arbitrage to explain the overall decrease in the number of arbitrage opportunities. Algorithmic trading influences the creation of triangular arbitrage by two countervailing effects. If the tick size exceeds the necessary price improvement for creating arbitrage opportunity, it would prevent triangular arbitrage from being. Furthermore, the narrowing of the spread reduces the price improvement necessary to create triangular arbitrage. Figure 3.29 indicates the frequency of the EUR/USD increments necessary to create triangular arbitrage opportunity. The x axis indicates the year from 2004 to 2008. The y axis labels the increment in pips, equal to 0.0001 , to create arbitrage. z axis gives the frequency in a given year with given increment. The graph indicates that the smaller increment in one currency could create arbitrage which is the evidence for the first effect of algorithmic trading.¹⁶

Figure 3.29: EUR-USD Increment to Create Arbitrage



Figures 3.30 and 3.31 provide the daily average of EUR/USD improvement necessary to create triangular arbitrage. The trend is decreasing until the end of the 2007 when the financial crisis hit the markets.

¹⁶The different behavior in 2008 is due to the financial crisis. The spread increased significantly as a response to the uncertainty in this year.

Figure 3.30: Change in EUR-USD to Create Arbitrage, Rout 1




Figure 3.31: Change in EUR-USD to Create Arbitrage, Rout 2





Figure 3.32: Binding Spread, 2004-2008



Notes: Figures 3.30 and 3.31 demonstrates the daily average of the EUR/USD increments needed to create triangular arbitrage opportunity. This value has decreased after the introduction of algorithmic trading in EBS market and it has increased starting the financial crisis. Figure 3.32 indicates the increasing binding spread, however, it decreases during the financial crisis.

3.7.2 Pressure on spread

We argued that the price improvement for creating arbitrage opportunity has decreased after introducing the algorithmic trading to the market. This would increase the possibility of more arbitrages. However, there is a second effect. If the spread become smaller and the price increment cross the spread, this would prevent triangular arbitrage from being created. [Figure 3.32](#) indicates the increasing binding spread at 0.0001 . This proportion starts at %60 in 2004 and reaches the %90 in 2007. The financial crisis widens the spread and decreases the binding spread in 2008.

3.8 Conclusions

We have shown that, taking into account microstructure friction, a causal relationship between algorithmic trading and lack of triangular arbitrage comes from a different channel than that commonly believed. The claim that the reduction of triangular arbitrage opportunities in the spot FX market serves as evidence that algorithmic trading activity aids in the price discovery process is rendered questionable by our results. The left-over problem makes it unlikely that any trader, manual or algorithmic, would exploit a triangular arbitrage opportunity –and indeed no individual trader does. Neither are the durations of triangular arbitrage opportunities reduced algorithmic trading activity. While it may be the case that algorithmic traders are better at guarding against placing quotes that would create triangular arbitrage, the relatively large tick size makes this an unprovable proposition for the data examined.¹⁷ The resulting decrease in triangular arbitrage opportunities in the presence of algorithmic trading is therefore primarily an artifact of the tick size.

Our results highlight the fact that examining the impact of high frequency trading requires not only high frequency data, but also higher resolution in considering market microstructure. When market microstructure effectively forbids the exploitation of apparently profitable arbitrage, the textbook conclusion that lack of arbitrage is due to the increasing number of arbitrageurs is no longer valid.¹⁸ Non-execution also makes discussion of execution irrelevant

¹⁷It could well be that algorithmic traders make triangular arbitrage opportunities more likely, as they move prices closer to triangular arbitrages on average. For example, during the last four months of 2008, the spread is less likely to bind. This removes the tick size consideration and indeed we see more triangular arbitrage opportunities than any other consecutive fourth months period in the five years examined, including all of 2004 when there is no algorithmic trading. However, this could also be related to the higher volatility, due to the financial crisis, of the market during that period.

¹⁸From a similar perspective, [Lamont and Thaler \(2003\)](#) posits that market rules such as short sale constraints can also impede arbitrage.

(see e.g. [Kollias and Metaxas \(2001\)](#) and [Kozhan and Wah Tham \(2012\)](#)).

Discussion on other aspects of price discovery for which these microstructure issues are not relevant are not impacted by our results. For example, it remains true that an increase in algorithmic trading is associated with a decrease in the serial correlation of high-frequency returns as shown in [Chaboud et al. \(2014\)](#). We can also confirm that public information, such as macroeconomic news release, are more quickly incorporated into prices by algorithmic traders.

Bibliography

- 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.
- Ait-Sahalia, Y. and J. Jacod (2009). Estimating the degree of activity of jumps in high frequency data. *Annals of Statistics* 37, 2202–2244.
- Alexander, K. and T. Zobotina (2005). Is it time to reduce the minimum tick sizes of the E-Mini futures? *Journal of Futures Markets* 25, 79–104.
- Amihud, Y. and H. Mendelson (1980). Dealership market: Market making with inventory. *Journal of Financial Economics* 8, 31–53.
- Andersen, T. G., T. Bollerslev, and F. X. Diebold (2007). Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility. *Review of Economics and Statistics* 89, 701–720.
- Andersen, T. G., T. Bollerslev, and D. Dobrev (2007). No-arbitrage semi-martingale restrictions for continuous-time volatility models subject to leverage effects, jumps and i.i.d. noise: Theory and testable distributional implications. *Journal of Econometrics* 138, 125–180.
- 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.
- Andersen, T. G., T. Bollerslev, P. Frederiksen, and M. Ørregaard Nielsen (2010). Continuous-time models, realized volatilities, and testable distributional implications for daily stock returns. *Journal of Applied Econometrics* 25, 233–261.
- Angel, J. (1997). Tick size, share prices, and stock splits. *Journal of Finance* 52, 655–681.
- 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. McNish (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, liquidity supply, and trading in pennies on the New York Stock exchange. *Journal of Financial Markets* 6, 337–362.
- Ball, C. A., W. N. Torous, and A. E. Tschoegl (1985). The degree of price resolution: The case of the gold market. *Journal of Futures Markets* 5, 29–43.
- Barndorff-Nielsen, O. E. and N. Shephard (2004). Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics* 2, 1–37.
- Barndorff-Nielsen, O. E. and N. Shephard (2006a). Econometrics of testing for jumps in financial economics using bipower variation. *Journal of Financial Econometrics* 4, 1–30.

- Barndorff-Nielsen, O. E. and N. Shephard (2006b). Impact of jumps on returns and realised variances: Econometric analysis of time-deformed Levy processes. *Journal of Econometrics* 131, 217–252.
- 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*.
- Bjønnes, G. H. and D. Rime (2005). Dealer behavior and trading systems in foreign exchange markets. *Journal of Financial Economics* 75, 571–605.
- 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 (2012). High frequency trading and price discovery. *Technical Report, University of California Berkeley*.
- 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*.
- Christensen, K., R. C. A. Oomen, and M. Podolskij (2014). Fact or friction: Jumps at ultra high frequency. *Journal of Financial Economics* 114, 576–599.
- Christie, W. G. and P. H. Schults (1994). Why do NASDAQ market makers avoid odd-eighth quotes? *The Journal of Finance* 49(5), 1813–1840.
- Copeland, T. E. and D. Galai (1983). Information effects on the bid-ask spread. *Journal of Finance* 38, 1457–1469.
- Cordella, T. and T. Foucault (1999). Minimum price variations, time priority and quote dynamics. *Journal of Financial Intermediation* 8, 141–173.
- Easley, D., M. Lopez de Prado, and M. O’Hara (2015). Discerning information from trade data. *NBER Working Paper, <http://ssrn.com/abstract=1989555>*.
- Easley, D. and M. O’Hara (1987). Price, trade size, and information in securities markets. *Journal of Financial Economics* 19, 69–90.
- Fama, E. F. and K. R. French (1988). Permanent and temporary components of stock prices. *Journal of Political Economy* 96, 246–273.
- Faust, J., J. H. Rogers, S.-Y. B. Wang, and J. H. Wright (2007). The high-frequency response of exchange rates and interest rates to macroeconomic announcements. *Journal of Monetary Economics* 54, 1051–1068.
- Gençay, R., M. Dacorogna, U. A. Müller, O. Pictet, and R. Olsen (2001). *An Introduction to High-Frequency Finance*. Academic Press.

- Gibson, S., R. Singh, and V. Yerramilli (2003). The effect of decimalization on the components of the bid-ask spread. *Journal of Financial Intermediation* 12, 121–148.
- Glosten, L. R. and P. R. Milgrom (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics* 14, 71–100.
- 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.
- Goodhart, C. and R. Curcio (1990). Asset price discovery and price clustering in the foreign exchange market. *Unpublished paper, London School of Economics*.
- Gwilym, O., A. Clare, and S. Thomas (1998). Extreme price clustering in the london equity index futures and options markets. *Journal of Banking & Finance* 22, 1193–1206.
- 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). The summary informativeness of stock trades: An econometric analysis. *Review of Financial Studies* 4, 571–595.
- Hasbrouck, J. and G. Saar (2013). Low-latency trading. *Journal of Financial Markets* 16, 646–679.
- Hendershott, T., C. M. Jones, and A. J. Menkveld (2011). Does algorithmic trading improve liquidity? *Journal of Finance* 66, 1–33.
- Hirschey, N. (2013). Do high-frequency traders anticipate buying and selling pressure? *Tech. rep., London Business School*.
- Ho, T. and H. Stoll (1981). Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics* 9, 47–73.
- Huang, R. D. and H. R. Stoll (1997). The components of the bid-ask spread: A general approach. *Review of Financial Studies* 10, 995–1034.
- Huang, R. D. and H. R. Stoll (2001). Tick size, bid-ask spreads, and market structure. *The Journal of Financial and Quantitative Analysis* 36, 503–522.
- Huang, X. and G. Tauchen (2005). The relative contribution of jumps to total price variance. *Journal of Financial Econometrics* 3, 456–499.
- Jacod, J., Y. Li, P. A. Mykland, M. Podolskij, and M. Vetter (2009). Microstructure noise in the continuous case: The pre-averaging approach. *Stochastic Processes and their Applications* 119, 2249–2276.
- 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.
- Kollias, C. and K. Metaxas (2001). How efficient are fx markets? Empirical evidence of arbitrage opportunities using high-frequency data. *Applied Financial Economics* 11, 435–444.
- Kondor, P. (2009). Risk in dynamic arbitrage: The price effects of convergence trading. *The Journal of Finance* 64, 631–655.
- Kozhan, R. and W. Wah Tham (2012). Execution risk in high-frequency arbitrage. <http://ssrn.com/abstract=2030767>.

- 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.
- Lamont, O. A. and R. H. Thaler (2003). The law of one price in financial markets. *Journal of Economic Perspectives* 17, 191–202.
- Lawrence, H. (1991). Stock price clustering and discreteness. *Review of Financial Studies* 4, 389–415.
- Lyons, R. (1995). Tests of microstructural hypotheses in the foreign exchange market. *Journal of Financial Economics* 39, 321–351.
- Oehmke, M. (2012). Gradual arbitrage. <http://ssrn.com/abstract=1364126>.
- Ohta, W. (2006). An analysis of intraday patterns in price clustering on the Tokyo stock exchange. *Journal of Banking & Finance* 30, 1023–1039.
- Payne, R. (2003). Informed trade in spot foreign exchange markets: An empirical investigation. *Journal of International Economics* 61, 307–329.
- Podolskij, M. and M. Vetter (2009). Estimation of volatility functionals in the simultaneous presence of microstructure noise and jumps. *Bernoulli* 15, 599–924.
- Podolskija, M. and M. Vetter (2009). Bipower-type estimation in a noisy diffusion setting. *Stochastic Processes and their Applications* 119, 2803–2831.
- Ready, M. (1999). The specialist's discretion: stopped orders and price improvement. *Review of Financial Studies* 12, 1075–1112.
- Roll, R. (1984). A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance* 39, 1127–1139.
- Schmidt, A. B. (2012). Ecology of the modern institutional spot FX: The EBS market in 2011. SSRN: <http://ssrn.com/abstract=1984070>.
- Schultz, P. (2000). Stock splits, tick size, and sponsorship. *Journal of Finance* 55, 429–450.
- Stoll, H. (1978). The supply of dealer services in security markets. *Journal of Finance* 33, 1133–1151.
- Van Ness, B. F., R. A. Van Ness, and R. S. Warr (2002). Is the adverse selection component really higher on the NYSE/Amex than on the Nasdaq? *Journal of Business Finance & Accounting* 29, 807–824.

Appendix A

Truncated Rate of Return

Suppose that the rate of return that the sub-penny jumper faces has a standard normal distribution: $x \sim N(\mu, \sigma^2)$. In what follows, f, F will denote the pdf and cdf, respectively. Similarly, ϕ, Φ will denote the pdf and cdf of the standard normal distribution.


$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(x-\mu)^2}, \quad \phi(z) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2}$$

$\phi'(z) = -z\phi(z)$ and $\phi(-z) = \phi(z)$. The associated cumulative distribution function is

$$\Phi(z) = Pr(Z \leq z) = \int_{-\infty}^z \phi(t) dt$$

Note that $\Phi'(z) = \phi(z)$ and $\Phi(-z) = 1 - \Phi(z)$. Once the penny jumper trades, the orders he front-runs protect him from serious losses on his position. If prices move in his favor, the sub-penny jumper profits to the full extent of the price changes. The returns are unbounded on one side and limited on the other side at $a = \frac{p-(p+\tau)}{p+\tau} = \frac{-\tau}{p+\tau}$.

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



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

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

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

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

The Moment Generating Function (MGF) of this distribution is

$$\begin{aligned} M(t) &= E[e^{tx}|x > a] = \frac{\int_a^{+\infty} e^{tx} f(x) dx}{1 - \Phi(\alpha)} \\ \int_a^{+\infty} e^{tx} f(x) dx &= \frac{1}{\sqrt{2\pi\sigma^2}} \int_a^{+\infty} e^{tx} e^{-\frac{1}{2\sigma^2}(x-\mu)^2} dx \\ &= \frac{1}{\sqrt{2\pi\sigma^2}} \int_a^{+\infty} e^{-\frac{1}{2\sigma^2}[(x-(\sigma^2 t + \mu))^2 - (\sigma^2 t + \mu)^2 + \mu^2]} dx \\ &= e^{-\frac{1}{2\sigma^2}[-(\sigma^2 t + \mu)^2 + \mu^2]} \frac{1}{\sqrt{2\pi\sigma^2}} \int_a^{+\infty} e^{-\frac{1}{2\sigma^2}[(x-(\sigma^2 t + \mu))^2]} dx \\ &= e^{\mu t + \frac{\sigma^2 t^2}{2}} \int_a^{+\infty} \frac{1}{\sigma} \phi\left(\frac{x - (\sigma^2 t + \mu)}{\sigma}\right) dx \\ &= e^{\mu t + \frac{\sigma^2 t^2}{2}} [-\Phi\left(\frac{a - (\sigma^2 t + \mu)}{\sigma}\right)] \\ \rightarrow M(t) &= E[e^{tx}|x > a] = e^{\mu t + \sigma^2 t^2 / 2} \frac{1 - \Phi(\alpha - \sigma t)}{1 - \Phi(\alpha)} \end{aligned}$$

Using MGF, the expected value of the left truncated return distribution would be


$$\begin{aligned} E(x|x > a) &= M'(t)|_{t=0} = \mu + \sigma \lambda(\alpha), \quad \text{where } \lambda(\alpha) = \frac{\phi(\alpha)}{1 - \Phi(\alpha)} > 0 \\ &\text{since } \lambda(\alpha) > 0 \rightarrow \mu + \sigma \lambda(\alpha) > \mu \rightarrow E(x|x > a) > E(x). \end{aligned}$$

To find the variance we need to calculate the following expression first:

$$\begin{aligned} E(x^2|x > a) &= M''(t)|_{t=0} = \mu^2 + \sigma^2 + \sigma^2 \frac{-\phi'(\alpha)}{1 - \Phi(\alpha)} - 2\mu\sigma \frac{-\phi(\alpha)}{1 - \Phi(\alpha)} \\ \rightarrow E(x^2|x > a) &= \mu^2 + \sigma^2 + \sigma^2 \alpha \lambda(\alpha) + 2\mu\sigma \lambda(\alpha) \quad \text{Since } -\phi'(\alpha) = \alpha \phi(\alpha) \\ \text{Var}(x|x > a) &= E(x^2|x > a) - E^2(x|x > a) = \sigma^2 [1 + \alpha \lambda(\alpha) - \lambda^2(\alpha)] \\ \text{Var}(x|x > a) &= \sigma^2 (1 - \delta(\alpha)) \quad \text{where } \delta(\alpha) = \lambda(\alpha) [\lambda(\alpha) - \alpha] \\ &\text{since } 0 < \delta(\alpha) < 1 \rightarrow \text{Var}(x|x > a) < \text{Var}(x). \end{aligned}$$

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

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



$$\begin{aligned}\frac{dE(x|x > a)}{d\tau} &= \frac{dE(x|x > a)}{d\alpha} \frac{d\alpha}{da} \frac{da}{d\tau} = \frac{d\lambda(\alpha)}{d\alpha} \frac{-p}{(p+\tau)^2} \\ \lambda(\alpha) &= \frac{\phi(\alpha)}{1-\Phi(\alpha)} \rightarrow \frac{d\lambda(\alpha)}{d\alpha} = \frac{-\alpha\phi(\alpha)[1-\Phi(\alpha)] + \phi(\alpha)\phi(\alpha)}{[1-\Phi(\alpha)]^2} = \lambda(\alpha)[\lambda(\alpha) - \alpha] = \delta(\alpha) \\ \frac{d\lambda(\alpha)}{d\alpha} &= 0 < \delta(\alpha) < 1 \rightarrow \frac{dE(x|x > a)}{d\tau} < 0\end{aligned}$$

$$\begin{aligned}\frac{dVar(x|x > a)}{d\tau} &= \frac{dVar(x|x > a)}{d\alpha} \frac{d\alpha}{da} \frac{da}{d\tau} = -\sigma \frac{d\delta(\alpha)}{d\alpha} \frac{-p}{(p+\tau)^2} \\ \frac{d\delta(\alpha)}{d\alpha} &= \frac{d\lambda(\alpha)}{d\alpha} [\lambda(\alpha) - \alpha] + \left[\frac{d\lambda(\alpha)}{d\alpha} - 1 \right] \lambda(\alpha) = \lambda(\alpha)[(\lambda(\alpha) - \alpha)^2 + \lambda(\alpha)(\lambda(\alpha) - \alpha) - 1] \\ &\rightarrow \frac{d\delta(\alpha)}{d\alpha} = \lambda(\alpha)[(\lambda(\alpha) - \alpha)(\lambda(\alpha) - \alpha + \lambda(\alpha)) - 1]\end{aligned}$$

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

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

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

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

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