

# Stock Exchange Competition in a High Frequency World

**Author:**

Chen, Hao Ming

**Publication Date:**

2019

**DOI:**

<https://doi.org/10.26190/unsworks/21622>

**License:**

<https://creativecommons.org/licenses/by-nc-nd/3.0/au/>

Link to license to see what you are allowed to do with this resource.

Downloaded from <http://hdl.handle.net/1959.4/64958> in <https://unsworks.unsw.edu.au> on 2024-05-01

# Stock Exchange Competition in a High Frequency World

Hao Ming Chen

A thesis in fulfillment of the requirements for the degree of  
Doctor of Philosophy



School of Banking & Finance

UNSW Business School

November 2019



**FOR OFFICE USE ONLY**    Date of completion of requirements for Award:

# Certificate of Originality

I, Hao Ming Chen, hereby declare that this submission is my own work and to the best of my knowledge it contains no materials previously published or written by another person, or substantial proportions of material which have been accepted for the award of any other degree or diploma at UNSW or any other educational institution, except where due acknowledgement is made in the thesis. Any contribution made to the research by others, with whom I have worked at UNSW or elsewhere, is explicitly acknowledged in the thesis.

I also declare that the intellectual content of this thesis is the product of my own work, except to the extent that assistance from others in the project's design and conception in style, presentation and linguistic expression is acknowledged.

Signed .....

Date .....

## INCLUSION OF PUBLICATIONS STATEMENT

UNSW is supportive of candidates publishing their research results during their candidature as detailed in the UNSW Thesis Examination Procedure.

**Publications can be used in their thesis in lieu of a Chapter if:**

- The student contributed greater than 50% of the content in the publication and is the “primary author”, i.e. the student was responsible primarily for the planning, execution and preparation of the work for publication
- The student has approval to include the publication in their thesis in lieu of a Chapter from their supervisor and Postgraduate Coordinator.
- The publication is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in the thesis

Please indicate whether this thesis contains published material or not.

☐

*This thesis contains no publications, either published or submitted for publication*

☒

*Some of the work described in this thesis has been published and it has been documented in the relevant Chapters with acknowledgement*

☐

*This thesis has publications (either published or submitted for publication) incorporated into it in lieu of a chapter and the details are presented below*

### CANDIDATE'S DECLARATION

I declare that:

- I have complied with the Thesis Examination Procedure
- where I have used a publication in lieu of a Chapter, the listed publication(s) below meet(s) the requirements to be included in the thesis.

Name	Signature	Date (dd/mm/yy)

## **COPYRIGHT STATEMENT**

'I hereby grant the University of New South Wales or its agents a non-exclusive licence to archive and to make available (including to members of the public) my thesis or dissertation in whole or part in the University libraries in all forms of media, now or here after known. I acknowledge that I retain all intellectual property rights which subsist in my thesis or dissertation, such as copyright and patent rights, subject to applicable law. I also retain the right to use all or part of my thesis or dissertation in future works (such as articles or books).'

'For any substantial portions of copyright material used in this thesis, written permission for use has been obtained, or the copyright material is removed from the final public version of the thesis.'

Signed .....

Date .....

## **AUTHENTICITY STATEMENT**

'I certify that the Library deposit digital copy is a direct equivalent of the final officially approved version of my thesis.'

Signed .....

Date .....

# **Stock Exchange Competition in a High Frequency World**

**Hao Ming Chen**

Submitted for the degree of Doctor of Philosophy

November 2019

## **Abstract**

This dissertation examines the impacts of competition among stock exchanges in a high frequency equities trading environment on liquidity provision, gains from trade for different trader types and externalities on other exchange venues. New securities regulation has facilitated the entry of stock exchanges that compete to provide alternative secondary market trading venues. Concurrently, major advancements in technology, including algorithmic trading, co-located market access and low latency communication networks, have dramatically increased the speed of trading. As a result, modern equities trading is both fast and fragmented.

This dissertation presents three empirical studies. The first study assesses the entry of a new trading venue into a monopolistic trading environment. Competition improves liquidity as the entrant reduces the incumbent's market power, which lowers trading fees that liquidity suppliers pass on via narrower bid-ask spreads. Market efficiency improves following the entrant exchange's launch, especially for participants with smart order routing technology to access both exchanges.

The second study investigates the introduction of a speed bump on an ex-

isting stock exchange, which provides some fast liquidity suppliers with guaranteed millisecond-level latency advantages to avoid order flow driven adverse selection. Profits from liquidity provision increase immensely for these beneficiaries. Slower liquidity suppliers, and those on other exchanges, face higher adverse selection, which is partially absorbed into lower profits and partially passed on via wider bid-ask spreads, increasing implicit transaction costs.

The third study examines a trading platform speed upgrade to the microsecond level on the main stock exchange in a competitive trading environment. The upgrade allows faster liquidity replenishment by high frequency traders and increases the venue's market share. However, quicker trade completion also enables fast liquidity suppliers on other venues to fade their quotes after observing trades on the upgraded venue, reducing the accessibility of liquidity across venues.

The dissertation's findings demonstrate that competition among stock exchanges in a high frequency world can have both intended consequences of attracting order flow and reducing market frictions, as well as unintended consequences that affect liquidity on other trading venues and the gains from trade for different trader types.



# Acknowledgments

Firstly, I would like to express my gratitude to my supervisor Sean Foley, who has guided me at each stage of this journey with thoughtful advice, inspirational ideas and endless encouragement. Additionally, I thank my supervisors Thomas Ruf and Jerry T. Parwada, who have provided exceptional guidance and support throughout my candidature.

Secondly, I would like to thank Michael Aitken for the opportunity to undertake an industry-sponsored PhD, and outstanding mentorship throughout the program. I am grateful for the financial support I received from Capital Markets CRC and the Ontario Securities Commission (OSC), where I was a visiting researcher during the inception stages of the third and fourth chapters of this dissertation. I thank Paul Redman and Tracey Stern at the OSC for welcoming me during my visit to Toronto, as well as sharing their knowledge on Canadian equity market structure.

Thirdly, I would also like to thank my friends and fellow students Peter O'Neill, Yi Ping Lin, Drew Harris, Daniel Rigney, Alex Sacco, Ming Ying Lim and Ivy Zhou, who made my candidature enjoyable and engaging.

Finally, I am indebted to my parents for their persistent support over the years and my brother, Henry, for meticulously proofreading this dissertation.

# Contents

<b>Abstract</b>	<b>ii</b>
<b>Certificate of Originality</b>	<b>iv</b>
<b>Acknowledgments</b>	<b>v</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Entrant Trading Venues . . . . .	7
1.2 Asymmetric Speed Bumps . . . . .	8
1.3 Matching Engine Latencies . . . . .	10
<b>2 The Impact of Fragmentation, Exchange Fees and Liquidity</b>	
<b>Provision on Market Quality</b>	<b>12</b>
2.1 Introduction . . . . .	12
2.2 Institutional Details . . . . .	15
2.2.1 Chi-X Introduction . . . . .	15
2.2.2 Australian Equity Market Structure Regulation . . . . .	16
2.2.3 Trading Fee Comparison . . . . .	17
2.3 Literature Review . . . . .	18
2.4 Data and Research Design . . . . .	22
2.4.1 Continuous Analysis of Transaction Costs . . . . .	24

2.4.2	Difference-in-Differences Natural Experiment . . . . .	25
2.4.3	Impact of Competition on Smaller Stocks . . . . .	27
2.4.4	Impact of Trading Fee Reduction . . . . .	27
2.4.5	Transaction Cost Measures . . . . .	28
2.4.6	Informational Efficiency Measures . . . . .	31
2.4.7	Measures of Pre-Trade Fragmentation . . . . .	32
2.5	Summary Statistics and Results . . . . .	33
2.5.1	Summary Statistics . . . . .	33
2.5.2	Impact of Market Fragmentation . . . . .	36
2.5.3	Intraday Analysis of Cross-Market Dynamics . . . . .	46
2.5.4	Impact of Trading Fee Reduction . . . . .	49
2.5.5	Impact on Price Efficiency and Volatility . . . . .	52
2.5.6	Robustness Tests . . . . .	53
2.6	Conclusion . . . . .	57

### **3 The Value of a Millisecond: Harnessing Information in Fast, Fragmented Markets 60**

3.1	Introduction . . . . .	61
3.2	Institutional Details . . . . .	67
3.2.1	TSX Alpha Relaunch and Speed Bump . . . . .	68
3.2.2	Trading Fee Change and Minimum Size for Delay-Exempt Orders . . . . .	70
3.2.3	Removal from Order Protection Rule . . . . .	72
3.2.4	Decommissioning of TMX Select . . . . .	73
3.3	Data and Liquidity Metrics . . . . .	74
3.3.1	Data . . . . .	74
3.3.2	Standard Liquidity Metrics . . . . .	76

3.3.3	Novel Liquidity Metrics . . . . .	78
3.3.4	Summary Statistics . . . . .	81
3.4	Methodology, Results and Discussion . . . . .	85
3.4.1	Speed Bump Mechanics and Fleeting Liquidity . . . . .	88
3.4.2	Multi-Venue Trades, NBBO Liquidity Depletion and Or- der Flow Segmentation . . . . .	95
3.4.3	Traded Liquidity Metrics per Venue . . . . .	103
3.4.4	Liquidity Provider Metrics on Inverted Fee Venues . . . .	110
3.4.5	Impacts on Consolidated Liquidity . . . . .	112
3.5	Conclusion . . . . .	121
3.6	Appendix I: Speed Bump Designs . . . . .	122
3.6.1	Asymmetric Fixed Duration Speed Bumps . . . . .	124
3.6.2	Asymmetric Random Duration Speed Bumps . . . . .	126
3.6.3	Symmetric Fixed Duration Speed Bumps . . . . .	129
3.6.4	Symmetric Random Duration Speed Bumps . . . . .	130
3.7	Appendix II: Novel Metric Development . . . . .	131
3.7.1	Trade Direction Classification . . . . .	131
3.7.2	Timestamp Synchronization Benchmarking . . . . .	134
3.7.3	Trade String Construction . . . . .	136
<b>4</b>	<b>Trading Engine Speed, Fleeting Liquidity and Competition for Order Flow</b>	<b>143</b>
4.1	Introduction . . . . .	143
4.2	Trading Infrastructure Speed . . . . .	147
4.2.1	Faster Trading Platforms in a Consolidated Market . . .	147
4.2.2	Faster Trading Platforms in a Fragmented Market . . . .	151
4.3	Institutional Details . . . . .	156

4.4	Data and Market Efficiency Metrics . . . . .	159
4.4.1	Data . . . . .	159
4.4.2	Standard Liquidity Metrics . . . . .	161
4.4.3	Novel Liquidity Metrics . . . . .	165
4.4.4	Summary Statistics . . . . .	170
4.5	Methodology, Results and Discussion . . . . .	171
4.5.1	Fleeting Liquidity Across Venues . . . . .	171
4.5.2	Stale Quote Sniping . . . . .	177
4.5.3	Liquidity Provision on TSX . . . . .	180
4.5.4	TSX Market Share . . . . .	185
4.5.5	Consolidated Quoted and Traded Liquidity . . . . .	190
4.5.6	Liquidity Metrics Across Venues . . . . .	194
4.5.7	Liquidity Metrics Across Broker Types . . . . .	197
4.5.8	Immediate Impacts on Liquidity Metrics . . . . .	199
4.6	Conclusion . . . . .	203
4.7	Appendix . . . . .	205
4.7.1	Timestamp Synchronization Benchmarking . . . . .	205
<b>5</b>	<b>Conclusions</b>	<b>208</b>
	<b>References</b>	<b>213</b>

# List of Figures

2.1	Chi-X Market Share . . . . .	18
2.2	Quoted Spreads, Effective Spreads and Price Impact . . . . .	36
2.3	Quoted Spreads by Tick Constraint . . . . .	37
3.1	SOR Interaction with Trading Venues . . . . .	70
3.2	Quote Fade Across Venues . . . . .	89
3.3	Trade, Fade and Stay for Liquidity on Alpha . . . . .	94
3.4	Realized Spreads for Single-Venue and Multi-Venue Trades . . .	96
3.5	Trading Volume Composition . . . . .	98
3.6	Broker Market Shares on Alpha . . . . .	102
3.7	Adverse selection and Realized Spreads . . . . .	104
3.8	Duration of Locked and Crossed Markets . . . . .	137
3.9	Trade String Construction . . . . .	142
4.1	Trading Platform Latency and Multi-Market Run Games . . . .	153
4.2	Quote Fade on Non-TSX Venues . . . . .	174
4.3	TSX Time at NBBO . . . . .	182
4.4	HFT Broker Queue Priority on TSX . . . . .	184
4.5	TSX Market Share . . . . .	188
4.6	Timestamp Synchronization Benchmarking . . . . .	207

# List of Tables

2.1	Summary Statistics . . . . .	35
2.2	ASX200 Liquidity Metrics as Chi-X Market Share Increases . .	40
2.3	ASX200 Liquidity Metrics by Tick Constraint as Chi-X Market Share Increases . . . . .	42
2.4	Liquidity Metrics for Incremental Additions . . . . .	43
2.5	Liquidity Metrics for Full Universe Introduction . . . . .	45
2.6	ASX200 Intraday Liquidity Duplication . . . . .	48
2.7	ASX200 Liquidity Metrics After Trading Fee Reduction . . . . .	51
2.8	Price Efficiency Metrics . . . . .	54
2.9	Robustness Tests . . . . .	56
2.10	ASX200 Liquidity Metrics as Chi-X Quote Quality Increases . .	58
3.1	Summary Statistics . . . . .	83
3.2	Trade Strings Summary Statistics . . . . .	86
3.3	Adverse Selection Metric Correlation Table . . . . .	87
3.4	Quote Fade on Each Venue . . . . .	91
3.5	NBBO Depletion and SOR on Each Venue . . . . .	100
3.6	Traded Liquidity Metrics on Each Venue . . . . .	107
3.7	Traded Liquidity Metrics on Each Venue and Fast Orders . . . .	109
3.8	Traded Liquidity Metrics on Inverted Venues . . . . .	113

3.9	Consolidated Liquidity Metrics . . . . .	116
3.10	Consolidated Liquidity Metrics by Quartile . . . . .	118
3.11	Consolidated Liquidity Metrics and Fast Orders . . . . .	120
3.12	Trade Initiation Direction Example 1 . . . . .	133
3.13	Trade Initiation Direction Example 2 . . . . .	134
3.14	Trade String Example . . . . .	140
4.1	Summary Statistics . . . . .	172
4.2	Quote Fade Across All Stocks . . . . .	176
4.3	Quote Fade by Quartile . . . . .	178
4.4	High Frequency Competition Among Liquidity Takers . . . . .	181
4.5	TSX Liquidity Provision . . . . .	186
4.6	TSX Market Share . . . . .	191
4.7	Consolidated Liquidity Metrics . . . . .	193
4.8	Comparison Across Venues . . . . .	196
4.9	Comparison Across Brokers . . . . .	200
4.10	Immediate Impacts Within One Week . . . . .	202



# Preface

Some of the work described in this thesis has been presented at academic conferences and published as joint work in refereed journals.

- The work that forms the basis of Chapter 2 has been published in the *Journal of Empirical Finance* as:

Aitken, Michael, Haoming Chen, and Sean Foley. 2017. “The Impact of Fragmentation, Exchange Fees and Liquidity Provision on Market Quality.” *Journal of Empirical Finance* 41 (March): 140–60.

- The work that forms the basis of Chapter 3 is an unpublished working paper with the same title, co-authored with Sean Foley, Michael Goldstein and Thomas Ruf, and has been presented at the 2018 American Finance Association Meeting, 2016 NBER Market Microstructure Meeting and 2016 Northern Finance Association Conference, among other conferences.
- The work that forms the basis of Chapter 4 is a sole-authored unpublished working paper with the same title and has not been presented at any conferences.

# Chapter 1

## Introduction

This dissertation examines the impacts of competition among stock exchanges in a high frequency equities trading environment on liquidity provision, gains from trade for different trader types and externalities on other exchange venues. It examines the entry of a competing stock exchange into a previously consolidated trading environment, and then turns to focus on how stock exchanges use trading speed as a market design feature to compete for different types of order flow.

A key development in equities market structure over the past decade has been the emergence of new stock exchanges as alternative public trading venues for securities that are listed on another stock exchange. Since the mid-2000s, securities regulations such as Regulation National Market System (Reg NMS) in the United States and the Markets in Financial Instruments Directive (MiFID) in Europe have facilitated the entry of new stock exchanges with unlisted trading privileges. Previously, securities could only be traded on their listing market. In a floor trading environment, Pagano (1989) argues that a centralized market structure is optimal due to the network externality of liquidity begetting liquidity. In an electronic trading environment, entry of new trading

venues creates multiple order book queues, which increases aggregate market depth (Foucault and Menkveld, 2008). O’Hara and Ye (2011) postulate that although equities trading in the United States is heavily fragmented across multiple trading venues, smart order routing and trade-through prohibition virtually recreate the network benefits of consolidated trading. Bernales et al. (2018) contend that improved market quality arises from competition in market design, rather than merely fragmentation as a result of new entrants.

To compete for fragmented order flow, stock exchanges have adopted a variety of market design features that attract different types of traders. Recent stock exchange innovations include dark trading without pre-trade transparency to reduce information leakage (Comerton-Forde and Putniņš, 2015), different trading fee structures for liquidity suppliers and demanders (Battalio et al., 2016), co-location of market participant computers with the exchange’s matching engine to reduce messaging delays (Brogaard et al., 2015), premium data feeds to receive market information more quickly (Goldstein et al., 2018), complex order types that are designed to achieve advantageous trading outcomes (Macey and Swensen, 2017) and discriminatory artificial delays on order processing and market data dissemination (Chakrabarty et al., 2019). In turn, these market design innovations facilitate and impact competition among traders.

A contemporaneous and interlinked development has been the rapid advancement in trading technology, which has facilitated algorithmic and high frequency trading (HFT) and increased the speed of the trading process. Menkveld (2016) surveys the HFT literature and explains that HFTs and entrant trading venues rely on each other to flourish, as entrant venues provide lower trading fees that are important for fee-sensitive proprietary traders and HFTs provide attractive quotes on these venues. Menkveld (2013) examines

the trading activity of a large HFT that predominantly provides liquidity and trades similar volumes on an incumbent venue and an entrant venue, indicative of a cross-market strategy. Malceniece et al. (2019) utilizes the incremental entry of an entrant venue across Europe as an instrumental variable for variation in HFT activity, finding increases in return co-movement and liquidity.

In its 2010 Concept Release on Equity Market Structure, the United States Securities and Exchange Commission (SEC) documented several features of HFTs, who are characterized as proprietary trading firms that trade large volumes, use high-speed and sophisticated software for generating and submitting orders to trading venues, use co-location services and fast market data feeds, trade over very short time horizons, frequently cancel orders soon after submission and minimize inter-day inventory.<sup>1</sup>

Numerous empirical studies find that algorithmic and high frequency traders exert positive impacts on market efficiency. Algorithmic trading improves liquidity (Hendershott et al., 2011). HFTs supply liquidity when it is scarce and consume it when it is plentiful (Carrion, 2013), facilitate price discovery when they demand liquidity (Brogaard et al., 2014b), mitigate transient order book imbalances (Jarnecic and Snape, 2014), reduce intraday price volatility (Hagströmer and Nordén, 2013), and improve linkages between futures and stock markets (Zhang, 2018). Other empirical studies find that HFTs can also have detrimental impacts on market efficiency, including exacerbating volatility during periods of market stress (Kirilenko et al., 2017), increasing trading costs for institutional investors by inferring their presence and trading in the same direction as their large orders (Kervel and Menkveld, 2019, Korajczyk and Murphy, 2019) and exploiting systematic arbitrage opportunities, which imposes adverse selection on liquidity suppliers (Foucault et al., 2017).

---

<sup>1</sup>See <https://www.sec.gov/rules/concept/2010/34-61358.pdf>

More recently, the theoretical literature has shifted to focus on competition for speed among HFT liquidity suppliers. Jovanovic and Menkveld (2016) argue that HFT market makers react faster to quantitative information, reducing the adverse selection risk of limit orders and improving welfare. Ait-Sahalia and Sağlam (2017) model the market making strategy of an HFT that receives an imprecise signal on incoming order flow direction and seeks to capture the spread without accumulating excessive inventory, finding that welfare increases with HFT speed. Hoffmann (2014) suggests that fast liquidity suppliers post more competitive quotes because their speed reduces adverse selection risks, but slower liquidity suppliers post less competitive quotes, which reduces their participation. Empirically, Baron et al. (2019) examine speed competition among HFTs, finding that faster HFTs earn larger revenues.

With speed being important for determining trading outcomes, traders may invest excessively in fast trading technologies, perpetuating a winner-takes-all arms race to be the fastest trader. Budish et al. (2015) identify mechanical arbitrage opportunities that arise in continuous limit order books with serial processing, which is currently the predominant exchange venue design, and results in a socially wasteful arms race for speed among traders. Several alternative market designs that could reduce speed competition are proposed. Biais et al. (2015) suggest that while fast traders invest in speed to reduce their search costs in fragmented markets, they do not internalize the negative externalities of the resulting increased adverse selection faced by slow traders. There is a long history of traders seeking faster methods of information transmission between geographically dispersed markets, starting with carrier pigeons and the telegraph in the nineteenth century (Foucault and Moinas, 2018). More recently, HFTs have invested generously in transmission technology between exchanges to capture fleeting trading opportunities, with innovations rang-

ing from more efficient fiber optic cables to microwave connectivity (Laughlin et al., 2014). Shkilko and Sokolov (2016) examine episodes of precipitation that disrupt microwave transmission, equalizing all traders to the same speed level on fiber, which reduces adverse selection and improves liquidity.

The theoretical literature has also examined the role of stock exchanges in increasing trading speeds. Pagnotta and Philippon (2018) develop a model in which faster trading venues reduce search costs, allowing for higher trading fees and attracting latency sensitive traders. Menkveld and Zoican (2017) find that faster trading venues increase adverse selection risk. Wang (2018) posits that limit orders on faster trading venues are less likely to be traded-through, as they are more likely to be observed by other traders. Conspicuously, these models examine exchange trading speed as the total time between a trader’s decision to trade and the trade’s completion, which includes trader latency and order transmission latency, in addition to exchange order processing latency.

The purpose of this dissertation is to examine the characteristics of competition among stock exchanges in a high-speed trading environment and investigate research questions that have been suggested in the recent literature. O’Hara (2015) examines contemporary issues in high frequency market microstructure and suggests that market linkages across fragmented trading venues and fairness proxied by participation are two especially important areas for future market microstructure research. I develop a method to quantify market linkages, to extend the microstructure toolkit, and then apply it to examine the impacts of stock exchange speed innovations. Foucault and Moinas (2018) review the theoretical and empirical literature on faster traders, exchanges and information transmission in financial markets, which has increased the speed of trade matching and information dissemination. Faster trading increases adverse selection, competition among liquidity providers and trade completion

speed, resulting in mixed overall impacts on market quality. They suggest that further evidence is required on the magnitude of benefits from faster trading, including increased competition between liquidity providers and lower search costs to find a trading counterparty. I conduct empirical analysis to examine these issues and quantify their impacts, utilizing stock exchange innovations in a competitive trading environment as exogenous shocks.

Much of this dissertation examines the impact of exchange competition on market liquidity, measured primarily via bid-ask spreads, as a proxy for implicit transaction costs. The canonical literature decomposes bid-ask spreads into their core components of order processing costs, inventory holding costs and adverse selection costs (Huang and Stoll, 1997). I examine how these individual components incurred by liquidity suppliers can be altered by stock exchange competition. I contend that lower exchange trading fees reduce order processing costs, while HFT liquidity suppliers face lower adverse selection costs if they can receive and act on order flow information more quickly, as well as lower inventory holding costs if they are more successful at achieving order book queue priority.

Closely related to this dissertation is Budish et al. (2019), which develops a model of stock exchange competition in which exchanges compete on trading fees but have market power in selling speed technology. Although an exchange can improve welfare by adopting batch auctions rather than the extant continuous order processing market design, to eliminate market inefficiencies and ‘fix the market’, the minimal or even negative private incentives prevent social benefits from being realized. This dissertation supports their hypothesis, by demonstrating that several recent stock exchange innovations continue to facilitate speed competition among traders, which contributes to the high frequency trading arms race.

This dissertation is structured as three empirical studies. The first study examines the introduction of stock exchange competition into a consolidated trading environment. Competition initially manifests as reduced trading fees on the incumbent venue, followed by the entry of endogenous liquidity providers, and finally the competitor’s entry. The second study examines a ‘speed bump’ innovation on an existing stock exchange, which increases the relative speed differentials between various types of traders to segment single-venue trades that typically have lower adverse selection costs. The third study investigates an immense reduction in exchange order processing and information dissemination latency in a fragmented trading environment, which can be utilized by fast traders to improve their gains from trade.

## **1.1 Entrant Trading Venues**

Securities regulators have facilitated competition among trading venues with new regulations, such as Reg NMS in the United States and MiFID in Europe. Similar regulations have been introduced to enable the entry of competing stock exchanges in Australia and Canada. Subsequently, entrant trading venues such as Chi-X have emerged to compete against incumbent stock exchanges in numerous countries (He et al., 2015).

The second chapter of this dissertation examines the incremental entry of Chi-X as an alternative trading venue into the previously consolidated Australian equities trading environment, where securities were only traded on the Australian Securities Exchange (ASX). Following announcement of Chi-X’s entry, ASX lowered its trading fees, decreasing order processing costs, which liquidity suppliers passed on via narrower bid-ask spreads. Liquidity and price discovery improve following Chi-X’s launch, especially for participants with



smart order routing technology to access both exchanges. Liquidity improvements increase with higher Chi-X market share over time. Quoted spreads narrow for stocks that were not previously constrained by the minimum tick size, while quoted depths increase for those that were. Most of the liquidity improvements are due to intermarket queue jumping (Foucault and Menkveld, 2008) and the quoting activities of two endogenous liquidity providers (Menkveld, 2013), who entered the Australian equities market soon before Chi-X's launch.

Chi-X launched in Australia with two market design innovations. Firstly, its liquidity pool incorporates hidden order types that do not have pre-trade transparency but offer price improvement over displayed orders, which contrasts with ASX's structure of offering hidden orders in a separate liquidity pool to the displayed orders. Secondly, orders on its market are good-for-day and purged at the end of each trading day, which provides endogenous liquidity providers more opportunity to gain queue priority at the start of each trading day.

## 1.2 Asymmetric Speed Bumps

It is a trader's speed relative to his competitors, rather than his absolute speed, that determines his profits from trading (Baron et al., 2019). In the quest for market participants seeking to trade faster than their competitors, stock exchanges have devised speed bumps, or systematic order processing delays, to provide selected traders with guaranteed speed advantages by slowing down other traders. IEX and NYSE American in the United States, and TSX Alpha and Aequitas NEO in Canada, have introduced speed bumps to delay the processing of some orders. The theoretical literature has also proposed speed bumps as mechanisms to de-emphasize the speed race in the trading

process (Baldauf and Mollner, 2019).

Stock exchange speed bumps can differ across two attributes, firstly whether the speed bump duration is fixed or randomized and secondly whether the delay is applied to all orders or only some orders. Each speed bump design seeks to address specific issues that arise from the nuances of frictions in continuous limit order book market design. Consequently, there are vastly different impacts on the types of order flow that the trading venue intends to attract, the impacts on gains from trade for traders who are delayed by the speed bump or able to bypass it, and spillovers onto other trading venues.

The third chapter of this dissertation examines the introduction of an asymmetric randomized 1 to 3 millisecond speed bump on TSX Alpha that only applies to liquidity demanders. In the new high frequency competitive trading environment, order size (O’Hara, 2015) and order dispersion across venues (Malinova and Park, 2017) are proxies for informed trading. I find that in this environment, the speed bump enables fast liquidity suppliers on Alpha to avoid order flow driven adverse selection by canceling their orders immediately in response to large trades on other venues, which immensely increases the profitability of their liquidity provision. Bid-ask spreads on Alpha do not narrow, indicating lack of competition among extremely fast liquidity suppliers. Consistent with theoretical predictions, adverse selection increases for slower liquidity providers on Alpha (Han et al., 2014) and liquidity providers on other venues (Biais et al., 2015), which is partially absorbed into lower profits and partially incorporated into wider bid-ask spreads.

A key contribution of this chapter is the development of new empirical methods to expand the empirical market microstructure toolkit. O’Hara (2015) identifies several data quality issues that affect the robustness of empirical research, including incorrectly sequenced trade and quote data, which obfuscates

trade initiation direction classification, and potential lack of timestamp synchronization across trading venues. Budish et al. (2015) also identify lack of clock synchronization as a market data limitation in continuous time trading. To the best of my knowledge, this dissertation proposes the first method developed in the academic literature to benchmark timestamp synchronization in high frequency datasets.

### 1.3 Matching Engine Latencies

Facilitated by technological innovation, stock exchanges around the world periodically upgrade their trading platforms to increase the speed at which incoming orders are processed and market data is disseminated. Over the last decade, the time taken has decreased tremendously from seconds, to milliseconds and now microseconds (Pagnotta and Philippon, 2018). Stock exchanges promote these upgrades as significant improvements to the operation of their trading venue. In a consolidated trading environment, the empirical literature has generally found market efficiency improvements as trading speed increases (Conrad et al., 2015, Riordan and Storckenmaier, 2012), while more recent theoretical literature has outlined the potential for higher adverse selection risk (Menkveld and Zoican, 2017). There has been limited analysis on the impacts of faster stock exchanges in a fragmented trading environment.

The fourth chapter of this dissertation examines the impact of a trading engine upgrade on the dominant Toronto Stock Exchange (TSX) in the fragmented Canadian equities trading environment. The upgrade reduces the latency of round-trip order processing and message dissemination between traders and the exchange from 2.3 milliseconds to 26 microseconds. I exploit the upgrade's immense magnitude and staggered rollout by the first letter of

stock codes to observe its immediate impacts. TSX increased its market share and fast liquidity providers were better able to achieve queue priority. The impact on liquidity was immaterial. However, the upgrade had the unintended consequence of enabling fast liquidity suppliers on other venues to fade their quotes sooner after observing trades on TSX, reducing the accessibility of liquidity across venues. In line with higher fleeting liquidity elsewhere, relative adverse selection costs on those venues decreases. The necessary conditions for predatory HFT sniping activity also increase.

For tractability, most of the theoretical literature has modeled trading speed from the perspective of sequential ‘batches’, which represent the frequency with which traders ‘visit’ the market and can submit or cancel orders, after which the exchange matches buyers and sellers. Shorter gaps between batches represent faster exchange order processing. Where necessary, this chapter departs from this convention to discuss the more representative institutional detail of stock exchanges that process orders serially in continuous time upon receipt (Budish et al., 2015).

## Chapter 2

# The Impact of Fragmentation, Exchange Fees and Liquidity Provision on Market Quality

### 2.1 Introduction

Explicit and implicit transaction costs are two of the major financial frictions facing traders today. The proliferation of alternative trading venues has led to a global fragmentation of order flow, bringing competition to bear on both financial frictions. While trading venue competition can drive explicit trading fees down, the positive network externalities of liquidity suggest that a consolidated venue may be the optimal structure for financial markets. We examine the impact of the reduction in these frictions due to the introduction of competition for equities trading in the Australian market.

We examine the separate impacts of two mechanisms that Foucault and Menkveld (2008) argue may result in lower bid-ask spreads when additional trading venues are present. Specifically, these two mechanisms are: compe-

tition between market makers duplicating their limit order schedules across marketplaces; and reduced explicit costs resulting from competition between exchanges. We show that while reductions in explicit transaction costs are sufficient to reduce bid-ask spreads, the primary benefits of fragmentation arise from increased competition between market makers across venues, specifically resulting from increased queue jumping and order duplication. We also document the occurrence of these two mechanisms using intra-day cross-market dynamics.

To distinguish between these two competing effects, we analyze two types of events: the reduction in explicit transaction fees charged by the Australian Securities Exchange (ASX) in the absence of competition in 2010; and the incremental introduction of different stocks on Chi-X, an alternative equities trading venue, into the Australian equities market between 2011 and 2013. The separate introduction of these changes allows us to examine the impact of each independently of the other. We investigate the impact of competition in equities trading on a variety of market quality measures, including quoted, effective and realized spreads, price impact, market depth, tick size constraint, Amihud (2002)’s illiquidity measure and several measures of price efficiency.

We have four main findings, confirming and extending the results of studies in Europe (Foucault and Menkveld, 2008) and the United States (O’Hara and Ye, 2011), which find that competition reduces both explicit and implicit transaction costs. First, we document that liquidity suppliers pass on reductions in exchange trading fees in 2010, reducing financial frictions and providing empirical evidence for one of Foucault and Menkveld (2008)’s theoretical channels. Second, we observe a further improvement in market quality after Chi-X’s introduction as a competing venue, with benefits increasing in entrant market share. Third, the vast majority of quoted spread reductions occur in

the stocks that were least tick-constrained, with tick-constrained stocks instead experiencing greater improvements in market depth. Fourth, analyzing the cross-market dynamics in high-frequency, we document queue jumping and order duplication when spreads are tick-constrained, or order book queues are long on the incumbent market. Our results are consistent with the queue-jumping hypothesis of Foucault and Menkveld (2008). These results indicate that the introduction of competition in the secondary trading of equities is welfare improving for a range of market participants, such as fund managers and retail traders.

However, there is a caveat to our results, because in some cases the increase in consolidated liquidity comes at the expense of the incumbent market. Participants who do not have a connection to the new entrant, likely smaller entities, may in this situation only observe a market efficiency detriment from competition. Compared to the mixed literature on the impact of competition between venues, it appears that both competition for explicit trading fees and competition among market makers on two separate exchanges are necessary to generate positive outcomes.

Our contribution has two unique and novel features. The first is the high-frequency nature of the environment we examine. Most studies assessing fragmentation of equity market order flow examine time periods when high-frequency trading (HFT) was still in its infancy. As O'Hara and Ye (2011) argue, the widespread use of smart order routers (SORs) leads fragmented markets to become “virtually consolidated” because participants are able to access orders across all venues almost simultaneously. We analyze a market where HFT is an established and pervasive component of the market.

The second unique feature of our analysis is the examination of the intra-day cross-market dynamics of traders across the entrant and incumbent mar-

kets. Dividing the day into 5-minute buckets, we are able to observe how prevailing market conditions impact the trading and quoting behavior of participants in an intra-day setting, demonstrating the mechanisms by which fragmentation improves liquidity – namely queue jumping and order duplication at times of high liquidity demand, when there are long queues on the incumbent exchange and quoted spreads are constrained by the minimum tick size.

The remainder of this chapter is organized as follows. Section 2.2 documents the market structure and introduction of competition to Australia. Section 2.3 reviews the related literature. Section 2.4 discusses the data and research design. Section 2.5 presents the results of our empirical analysis and examines the robustness of our main findings. Section 2.6 concludes.

## **2.2 Institutional Details**

### **2.2.1 Chi-X Introduction**

The ASX was the monopoly provider of an exchange trading platform for the equities and other financial securities that it lists and did not face competition until 2011, when Chi-X entered the Australian market. Chi-X conducted a staggered rollout, with an initial “soft launch” of six highly liquid stocks and two exchange traded funds (ETFs) on 31 October 2011.<sup>1</sup> The remaining ASX200 constituents and ASX-listed ETFs began trading on Chi-X on 9 November 2011. Between Chi-X’s ASX200 launch in November 2011 and the expansion to trading the full universe of ASX-listed securities on 3 May 2013, trading eligibility in 57 securities was incrementally introduced from 19 December 2011 to 15 March 2013. Of these 57 securities, 29 were not associ-

---

<sup>1</sup>Stock codes are BHP, CSL, LEI, ORG, QBE and WOW. ETF codes are STW and ISO.



ated with concurrent inclusion or deletion from an S&P/ASX index series, or an initial public offering.<sup>2</sup> On 3 May 2013, eligibility for Chi-X trading was introduced for the full universe of ASX securities. However, no trades were recorded on Chi-X among these stocks until 12 June 2013. We analyze changes in market quality arising from these introductions in three groups: the first analyzes the introduction of the ASX200 constituents; the second the incremental introduction of 29 securities using a difference-in-differences analysis with controls for stocks not trading on Chi-X at that time, while the third analyses the commencement of trading among the smaller securities in 2013.

### **2.2.2 Australian Equity Market Structure Regulation**

The introduction of a second marketplace required the establishment of regulations relating to how brokers are required to submit their orders to the market to obtain the best outcome for their clients. The Australian Securities and Investments Commission (ASIC) introduced a “best-execution” regime under ASIC Market Integrity Rules (Competition in Exchange Markets) 2011, similar to that in operation in Europe. This regime allows brokers flexibility to consider price, probability of fill, cost, speed and several qualitative metrics when deciding how to route client orders to different trading venues, and does not explicitly require brokers to connect to Chi-X. This framework differs significantly from that in the United States, where trade-through prohibitions require trades to occur on the venue displaying the best price, even if the quantity available is very small, so brokers must connect to all exchanges. While both ASX and Chi-X order books operate a “price-time” priority matching algorithm, inter-market time priority is not enforced across venues. This means

---

<sup>2</sup>Stock codes are ANP, FML, MOY, PEN, SBL, WHN, GRR, HZN, MNC, NXS, ROC, SAR, ABU, AJA, CFU, GOR, HIL, NST, PEM, PRR, SSN, TGR, UNS, WTP, EXG, BRG, ASB, SRX and CMW.

that if long queues exist on the ASX order book, but the Chi-X order book is empty, a trader could choose to start a new and shorter queue on Chi-X at the same price. Such a feature is important for explaining an incentive to queue-jump between markets.

Figure 2.1 presents the evolution of Chi-X’s market share in Australia over the first year since its introduction. Similar to the global evolution of Chi-X as an alternative trading venue, its market share in Australian equities began at a very low level, remaining below 2% of total daily on-exchange turnover over the first 6 months of operation. Within one year, total on-exchange market share had exceeded 5% and by late 2013 it exceeded 10%. This growth in market share was rapid compared to other countries, as documented by He et al. (2015). The gradual introduction and increase in market share provide market participants the opportunity to become familiar with the Chi-X trading system functionality over time. Significant variation in both the cross-section and time-series of fragmentation facilitates the causal analysis of its impact on market quality.

### 2.2.3 Trading Fee Comparison

On 1 July 2010, after the announcement of Chi-X’s intention to establish a competing stock exchange in Australia, the ASX reduced fees for the trading services that would be subject to competition.<sup>3</sup> The trading fee for on-market trades during the continuous trading session was reduced from 0.28 basis points to 0.15 basis points, while the trading fee for off-market trade reporting was decreased from 0.075 basis points to 0.05 basis points.<sup>4</sup> Chi-X introduced a

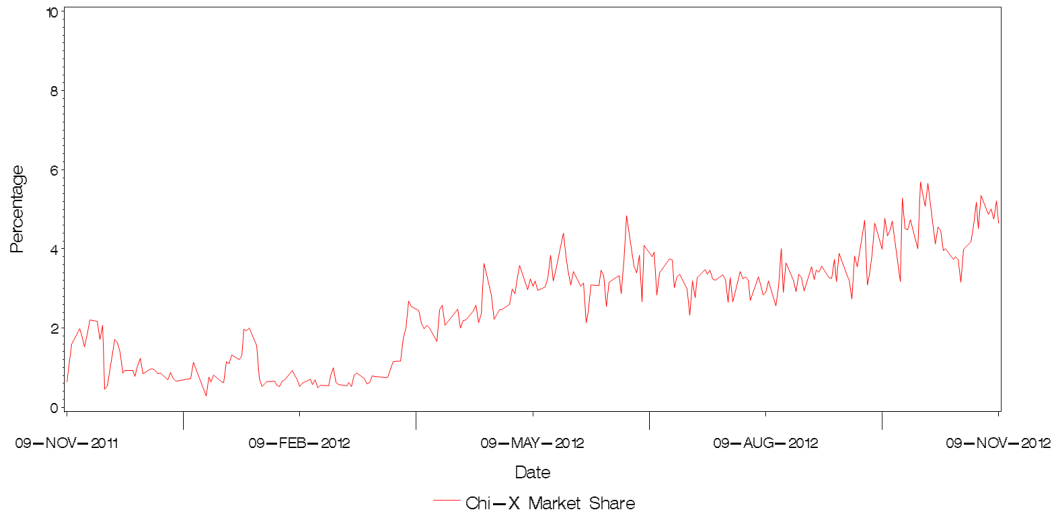
---

<sup>3</sup>ASX Market announcement, see [https://www.asx.com.au/documents/investor-relations/20100603\\_asx\\_fees\\_and\\_rebates.pdf](https://www.asx.com.au/documents/investor-relations/20100603_asx_fees_and_rebates.pdf).

<sup>4</sup>We examine only on-market trading activity and exclude trades that do not interact with the limit order book. Rose (2014) provides additional institutional details for off-market trade reporting in the Australian equities market.

Figure 2.1: **Chi-X Australia Market Share of Trading Turnover**

This figure presents the percentage of daily total on-market trading turnover in S&P/ASX 200 index constituent securities that was traded on Chi-X Australia over its first year of its operation, commencing from 9 November 2011. On-market trades are those that interact with the limit order book.



distinction between maker and taker fees, where a “maker” is a liquidity supplier, who submits a non-marketable limit order and a “taker” is a liquidity demander, who submits a marketable order, similar to the trading fee structures prevalent in North America.<sup>5</sup> Chi-X’s trading fees at launch were 0.06 basis points for the maker and 0.12 basis points for the taker, as well as 0.04 basis points for reporting off-market trades. Notably, trading during the ASX opening and closing auctions, which did not face competition, experienced no reduction in trading fees and remained at 0.28 basis points.

## 2.3 Literature Review

Two main strands of literature investigate the impact of competition on market quality in equities trading. The first assesses the impact of competition

---

<sup>5</sup>For further information on maker-taker trading fees, see Malinova and Park (2015).

between brokers and separate listing venues. The second examines the impact of order flow fragmentation in the U.S. and European equities markets on market quality.

Equity market structure has two potentially competing impacts on transaction costs - economies of scale reduce costs in consolidated markets, while competition between fragmented markets drives costs down (Hamilton, 1979). Securities exchanges may be natural monopolies, with significant economies of scale in clearing, settlement and infrastructure provision due to the high fixed and low marginal costs associated with matching orders. Pagano (1989) shows that these network externalities result in order flow gravitating towards one single, dominant exchange. In contrast, Economides (1996) argues that competitive forces are necessary to promote operating efficiencies and to ensure that exchanges do not earn excessive monopolistic profits.

Numerous studies of new market entrants provide support for the notion that competition reduces implicit transactions costs by narrowing bid-ask spreads. These include studies of off-board markets (Cohen and Conroy, 1990), the entry of additional broker dealers (Battalio et al., 1997), international cross listings (Domowitz et al., 1998), and direct competition among options trading exchanges (Fontnouvelle et al., 2003). A number of studies, however, identify significant costs of competition. Bessembinder and Kaufman (1997) provide evidence that fragmentation allows cream skimming of uninformed traders, increasing bid-ask spreads. Bennett and Wei (2006) demonstrate that stocks transferring from the NASDAQ dealer market to the more consolidated NYSE experience significantly decreased spreads, with liquidity improving most for stocks with the greatest increases in the level of consolidation. Gajewski and Gresse (2007) find that spreads in European stocks are lower in centralized, consolidated electronic order-driven markets than in hybrid markets where

orders are fragmented between an order book and competing dealers. The varying market structures and the different types of competition examined in each of these studies are likely to have contributed to the mixed findings.

Implementation of the Regulation National Market System (Reg NMS) in the United States in 2005 and the Markets in Financial Instruments Directive (MiFID) in Europe in 2007 facilitated the rapid fragmentation of global securities markets. O’Hara and Ye (2011) examine the fragmentation facilitated by Reg NMS, finding reduced effective spreads and enhanced price discovery. They argue that fragmented trading across U.S. venues benefits from virtual consolidation, due to smart order routing and trade-through prohibitions.

Foucault and Menkveld (2008) propose two mechanisms by which competition may reduce trading costs. Firstly, price competition from the entrant exchange may encourage the incumbent venue to reduce explicit trading fees and hence order processing costs for market makers. Secondly, “queue jumping” on the entrant exchange can occur in the absence of inter-market time priority across venues, which intensifies competition between market makers, increasing total quoted depth and potentially reducing bid-ask spreads. Foucault and Menkveld (2008) conduct a joint empirical test of these hypotheses in the Dutch equities market, finding increased depth as well as narrower or unchanged quoted spreads. An identified limitation is the opposing effects of queue jumping induced by the entrant exchange, which is expected to reduce quoted depth on the incumbent exchange, and the simultaneous reduction of fees on the incumbent exchange, which should increase quoted depth. In the Australian equities environment, a fee reduction is implemented prior to the entrant market’s launch, allowing us to disentangle these competing effects.

Further evidence on the beneficial impacts of market fragmentation is provided by Chlistalla and Lutat (2011) in France and Gresse (2017) for LSE

and Euronext listed securities, with stocks that are actively traded on entrant venues found to have narrower quoted spreads, increased quoted depth and reduced round trip execution costs for a specified transaction size. In a global study, He et al. (2015) also document that the introduction of an alternative equities trading venue both improves market depth and reduces bid-ask spreads in the majority of countries analyzed.

Our study of equities trading fragmentation is also related to the emerging literature on fragmentation across lit and dark order books, especially because both the entrant and incumbent markets have dark trading functionality without pre-trade transparency. Kwan et al. (2015) show that US dark pools provide participants the ability to obtain a finer pricing grid, particularly when prices are constrained by the minimum tick size, leading to lower transactions costs. Similarly, Buti et al. (2015) show that dark pools can facilitate “queue jumping” in front of lit orders, which is similar to the case made by Foucault and Menkveld (2008) for lit market fragmentation. Degryse et al. (2015) find that while lit market fragmentation is beneficial, dark fragmentation is detrimental to market quality. Critically, the benefits of lit market fragmentation only accrue to those who connect to the entrant markets, with the incumbent market often experiencing a deterioration of liquidity. Foley and Putniņš (2016) provide further evidence on the role of dark trading, finding that two-sided dark trading improves market quality by encouraging competition between liquidity suppliers, consistent with the theoretical model of Boulatov and George (2013).

The literature has also examined the trading characteristics of market makers and liquidity providers, as well as their impact on market efficiency. Goldstein and Nelling (1999) find that increased competition between market makers on NASDAQ reduces overall trading costs. More recently, Anand

and Venkataraman (2016) document the existence of endogenous liquidity providers (ELPs) who are present in the market with two-sided quotes, although they do not have explicit obligations to supply liquidity. They find that reducing inventory holding costs leads to higher ELP participation. Similarly, Menkveld (2013) finds that the liquidity providing activity of a large high frequency trader (HFT) utilizes a cross-market strategy, with similar turnover in both the incumbent and entrant market. With 78% of its trades originating from limit orders, trading fees significantly impact the HFT's profitability. We extend Menkveld (2013)'s analysis of HFT liquidity provision to its impact on aggregate market liquidity.

## 2.4 Data and Research Design

Data on trades and quotes is obtained from Thomson Reuters Tick History (TRTH), which is provided by the Securities Industry Research Centre of Asia-Pacific (SIRCA) and time-stamped to the millisecond. Trade records include price and volume, as well as flags for off-market trades, dark trades, and trades executed in the opening and closing auctions. Quote data with millisecond time-stamps is available for all 10 levels of the visible order-book. This data provides information on the price and quantity at each price level, and is updated each time a trade or order cancellation, amendment or entry updates the order book. Broker identifiers for ASX trades are sourced from SIRCA's Australian Equities Tick History database. For each trade, the buying and selling broker are identified, as well as the trade initiator. Similar trade data with broker identifiers is provided by Chi-X.

Trades and quotes prior to 10:10am and after 4:00pm are removed to exclude the opening and closing auction process. We also exclude off-market

trade reports as they do not interact with limit order books. The sample of stocks utilized to examine the reduction in fees includes all 180 stocks that remain in the S&P/ASX 200 index over the period covering the six months before and after the fee changes, specifically from 1 January to 31 December 2010. We exclude 14 stocks that were not in the ASX200 Index for the whole of 2010, 1 stock that trades below 10c (due to the large reduction in minimum tick size at this price threshold) and 5 stocks that were newly listed or delist due to bankruptcy or takeover.

We analyze the impact of fragmentation using four empirical specifications. The first method utilizes Chi-X's percentage of total on-market turnover for each stock-day across ASX200 component securities, as a proxy for the level of fragmentation. This approach allows us to identify changes in market efficiency as order flow fragmentation increases over the first year of Chi-X's operation, as well as differences in fragmentation levels in the cross-section of stocks. We include all 170 stocks that remain ASX200 constituents for the whole year. We exclude 11 stocks that were removed from the ASX200, 4 stocks that were listed after 9 November 2010, 3 stocks that trade below 10c (due to the large reduction in minimum tick size at this price threshold) and 12 stocks that delist due to bankruptcy or takeover.

The second method utilizes the staggered introduction of 29 stocks along with a matched sample of non-Chi-X traded securities as a natural experiment to undertake a difference-in-differences approach to examine the impact of the introduction of competition.<sup>6</sup> For these 29 stocks, we examine the period from three months prior to the first introduction to three months after the last introduction, being 10 October 2011 to 6 May 2013.

In contrast with stock-day observations utilized in the first specification, the

---

<sup>6</sup>We follow Goldstein et al. (2007) who face a similar situation in US Bonds.



third method utilizes stock-day-interval observations constructed across five-minute buckets, to analyze the impact of high frequency cross-market dynamics on liquidity provision and order book duplication. Eligibility for trading on Chi-X was expanded to the full universe of ASX-listed stocks on 3 May 2013. However, competitive quoting did not commence immediately.

Our final method identifies and examines 80 All Ordinaries Index securities that commenced competitive quoting on Chi-X on 12 June 2013, using a period from 12 December 2012 to 11 December 2013.

#### 2.4.1 Continuous Analysis of Transaction Costs

Our first specification adopts a measurement of variation in Chi-X market share to analyze the impact of fragmentation on liquidity. We calculate market share as the percentage of on-market turnover per security on each venue per day, capturing variation in fragmentation in both the cross-section and time-series. We avoid potential issues of endogeneity between fragmentation and liquidity by using the 5-day moving average of Chi-X trading, 1-day lagged.

We analyze data for all trading days from 9 November 2011 to 8 November 2012, being the first year of Chi-X's operation, using equations of the form:

$$\begin{aligned}
EfficiencyMetric_{i,d} = & \beta_1 Chi - XMAVG_{i,d} + \beta_2 Price_{i,d} \\
& + \beta_3 Turnover_{i,d} + \beta_4 Volatility_{i,d} + \beta_5 MedTick_{i,d} + FE_i + \epsilon_{i,d}
\end{aligned} \tag{2.1}$$

where  $EfficiencyMetric_{i,d}$  is the market efficiency metric of interest for stock  $i$  on day  $d$  and  $Chi - XMAVG_{i,d}$  is the lagged 5-day moving average Chi-X market share. For regressions where the liquidity metric is presented in basis points,  $Price_{i,d}$  is the inverse of the time-weighted NBBO midpoint

price, similar to Hendershott et al. (2011). In all other regressions,  $Price_{i,d}$  is the natural logarithm of the daily time-weighted midpoint price of the security.  $Turnover_{i,d}$  is the natural logarithm of on-market trading turnover for stock  $i$  on day  $d$ .<sup>7</sup>  $Volatility_{i,d}$  is calculated as the daily high price minus the low price, divided by their average.  $LowTick_{i,d}$  ( $MedTick_{i,d}$ )<sup>8</sup> takes a value of one if the price of stock  $i$  on day  $d$  is less than 10c (between 10c and \$2), which corresponds to having a tick size of 0.1c (0.5c), and zero otherwise, with our base-case being a stock with a price greater than \$2, for a tick size of 1c.<sup>9</sup>  $FE_i$  represents stock-specific fixed effects that control for time-invariant heterogeneity in liquidity at the stock-level, which are unrelated to fragmentation.  $\epsilon_{i,d}$  is an error term.

## 2.4.2 Difference-in-Differences Natural Experiment

Our second specification examines the impact of fragmentation on liquidity using a treatment group of 29 stocks for which competition was incrementally introduced between Chi-X's launch of trading in the ASX 200 components and subsequent roll-out of the full universe of stocks. As these 29 stocks were introduced in a staggered fashion, we are able to utilize a difference-in-differences approach, using a matched set of non-competition securities to control for changes in market efficiency characteristics that are driven by factors other than the introduction of competition. Each stock is matched to an ASX listed stock which did not have Chi-X connectivity during the period in a manner similar to Huang and Stoll (1996) that minimizes the sum of squared relative

---

<sup>7</sup>Following Benston and Hagerman (1974), we apply a natural logarithmic transformation to price and volume, to minimize the impacts of large right tail observations.

<sup>8</sup>Due to the exclusion of the one stock priced less than 10c from our sample the  $LowTick_{i,d}$  variable is not necessary for our ASX200 sample, but is used in our other specifications.

<sup>9</sup>The importance of considering the impact of tick size on market quality variables is documented by Ahn et al. (2002).

differences in turnover and price during October 2011:

$$MatchingScore_{C,N} = \sum_{j=1}^2 \left( \frac{x_j^C - x_j^N}{0.5 * (x_j^C + x_j^N)} \right)^2 \quad (2.2)$$

The superscript  $C$  indexes stocks experiencing competition, while the superscript  $N$  indexes non-competition stocks.<sup>10</sup> Equation 2.3 specifies the regression model estimated:

$$\begin{aligned} EfficiencyMetric_{i,d} = & \beta_0 + \beta_1 Competition_{i,d} + \beta_2 Treatment_i \\ & + \beta_3 Competition_{i,d} * Treatment_i + \beta_4 Price_{i,d} + \beta_5 Turnover_{i,d} \\ & + \beta_6 Volatility_{i,d} + \beta_7 LowTick_{i,d} + \beta_8 MedTick_{i,d} + \epsilon_{i,d} \end{aligned} \quad (2.3)$$

where  $EfficiencyMetric_{i,d}$  is the market efficiency metric of interest for stock  $i$  on day  $d$ , and  $Competition_{i,d}$  is an indicator variable equal to one for stock-day observations in the post-competition period (for both treatment and control stocks), and zero otherwise, with the pre-competition period constituting our base case.  $Treatment_i$  is an indicator variable equal to one for stocks that receive Chi-X introduction, and zero for the control stocks. The interaction term  $Competition_{i,d} * Treatment_i$  is our main explanatory variable of interest and captures the marginal effects of being a treatment stock in the post-competition period. The remaining variables take the same values as in Equation 2.1.

---

<sup>10</sup>The median differences between the competition and matched non-competition stocks' price and average traded dollar volume are less than 8% and 20% respectively, suggesting the matching is relatively precise. Larger differences for turnover are observed due to the relatively low average turnover among these non-index securities.

### 2.4.3 Impact of Competition on Smaller Stocks

Our third specification utilizes the expansion of trading on Chi-X to all ASX-listed securities in 2013. Although eligibility for quoting and trading on Chi-X was implemented on 3 May 2013, competitive quoting only begins on 12 June 2013. We identify 80 securities among the All Ordinaries Index components where Chi-X quoted at the NBBO price for at least 5% of each trading day during the calendar month after 12 June 2013, and were not previously eligible to trade on Chi-X. We analyze data for all trading days from 12 December 2012 to 11 December 2013, being six months on either side of the commencement of competitive quoting on Chi-X. We use an event study to examine the impact of the introduction of competition, where  $Competition_d$  is equal to zero prior to 12 June 2013 and one after. Equation 2.4 presents the regression specification, with all other variable definitions following from Equation 2.1.

$$\begin{aligned} EfficiencyMetric_{i,d} = & \beta_1 Competition_d + \beta_2 Price_{i,d} + \beta_3 Turnover_{i,d} \\ & + \beta_4 Volatility_{i,d} + \beta_5 LowTick_{i,d} + \beta_6 MedTick_{i,d} + FE_i + \epsilon_{i,d} \end{aligned} \quad (2.4)$$

### 2.4.4 Impact of Trading Fee Reduction

Our final specification examines the separate impact of a reduction in exchange fees. While there are two potential channels by which competition may reduce bid-ask spreads, the typically contemporaneous nature of exchange fee reductions and the launch of a competing venue limits the ability to determine which channel's effects dominate. Prior to the introduction of Chi-X, the ASX reduced explicit trading fees by almost 50% – from 0.28 to 0.15 basis points. The reduction of exchange fees on 1 July 2010 allows us to independently test

the impact of exchange fee reductions on liquidity. We use an event study to test the impact of the fee reduction, where  $FeeReduction_d$  is equal to zero prior to the fee reduction and one after. Our regression specification is provided in Equation 2.5 below, with all other variables defined as in Equation 2.1:

$$\begin{aligned}
EfficiencyMetric_{i,d} = & \beta_1 FeeReduction_d + \beta_2 Price_{i,d} \\
& + \beta_3 Turnover_{i,d} + \beta_4 Volatility_{i,d} + \beta_5 MedTick_{i,d} + FE_i + \epsilon_{i,d}
\end{aligned} \tag{2.5}$$

## 2.4.5 Transaction Cost Measures

In Equation 2.1, 2.3, 2.4 and 2.5,  $EfficiencyMetric_{i,d}$  represents the liquidity measure of interest. We examine quoted spreads, quoted depth, tick constraint, effective spreads, price impacts, realized spreads and Amihud's illiquidity measure, as well as several informational efficiency measures.

The majority of recent empirical literature has focused on relative spreads in basis points, rather than absolute spreads in cents. Noting the importance of tick size changes identified in Bessembinder (2000), as well as the severely tick-constrained quoted spreads observed in our sample with a median of 1.065 ticks among ASX 200 stocks in the post-entry period, we focus on spreads in tick increments.

The spread metric definitions below are presented in absolute dollar terms. These are converted into tick increments by dividing by the prevailing tick size, or basis points by dividing by the prevailing NBBO midpoint price. Quoted spreads and quoted depths are time-weighted per stock-day, while effective spreads, price impacts and realized spreads are turnover-weighted per stock-day. Absolute quoted spreads are calculated as:

$$QuotedSpread = NBOPrice - NBBPrice \quad (2.6)$$

where  $NBOPrice$  is the lowest ask price and  $NBBPrice$  is the highest bid price prevailing across both venues, which is also referred to as the National Best Bid and Offer (NBBO) prices, at the time of observation in the stock of interest. Time-weighted quoted spread per stock-day is constructed by weighting all quoted spreads across the consolidated market by the percentage of the trading day that the respective spreads were active.

One mechanism by which the introduction of competition for order flow between venues could decrease spreads is by increasing the proportion of the trading day during which quoted spreads are constrained by the minimum tick size. This percentage is impacted less by extremely large spread observations. We define the quoted spread at each point in time as being constrained if it is equal to the minimum tick size and then calculate the proportion of each stock-day for which the quoted spread was constrained.

Quoted dollar depth is the value that can be immediately traded at the NBBO across both venues, and is constructed by multiplying the price by the volume available at the NBBO, as described in Equation 2.7 below:

$$QuotedDepth = NBBPrice * NBBVolume + NBOPrice * NBOVolume \quad (2.7)$$

where  $NBOPrice$  is the lowest ask price across both venues,  $NBOVolume$  is the aggregate number of shares available to be bought at that price across both venues,  $NBBPrice$  is the highest bid price across both venues and  $NBBVolume$  is the aggregate number of shares available to be sold at that

price across both venues. We then weight each quoted depth observation for the stock-day by the percentage of the trading day for which that depth observation was active.

We also examine effective spreads, which measure the implicit transaction costs for liquidity demanders, being the difference between the transaction price and the NBBO midpoint price at the time of the transaction,  $NBBOMP_t$ . Effective spread is defined as:

$$EffectiveSpread = 2 * Direction * (TradePrice - NBBOMP_t) \quad (2.8)$$

*Direction* is +1 for buyer-initiated trades and -1 for seller-initiated trades. Buyer and seller-initiation flags are identified explicitly in our ASX and Chi-X datasets.<sup>11</sup>

Price impacts measure the directional change in NBBO midpoint price following each trade, as a gauge of the informativeness of the trade and market resiliency. It reflects the implicit transaction cost paid by liquidity demanders, less the residual portion attributed to liquidity supplier profits. Following Conrad et al. (2015), we present price impacts calculated twenty seconds after each trade. Price impact is defined as:

$$PriceImpact = 2 * Direction * (NBBOMP_{t+20sec} - NBBOMP_t) \quad (2.9)$$

where  $NBBOMP_t$  is the NBBO midpoint at the time of the trade,  $NBBOMP_{t+20sec}$

---

<sup>11</sup>Lack of time stamp synchronicity across venues may result in negative effective spreads being calculated, due to the national best bid price being higher than the national best offer price at the time of the trade. Recognizing the non-negative cost of demanding liquidity, we set effective spreads lower than zero to be equal to zero for each trade.

is the NBBO midpoint prevailing twenty seconds after the trade, and *Direction* is the trade initiation direction indicator, being +1 for buyer-initiated trades and -1 for seller-initiated trades.

Realized spreads measure the trading profits attributable to liquidity provision at a specified time horizon, being the effective spread less any price impact incurred over that time horizon. It is defined as:

$$\begin{aligned} RealizedSpread &= EffectiveSpread - PriceImpact \\ &= 2 * Direction * (TradePrice - NBBOMP_{t+20sec}) \end{aligned} \quad (2.10)$$

Finally, we calculate Amihud (2002)'s measure of illiquidity to examine changes in the market impact of trades. A lower illiquidity ratio indicates that the market was more able to absorb volume shocks without incurring large price movements over each hour of the trading day. This metric is logarithmically normalized and defined as:

$$ILLIQ_{i,d} = \ln(1 + 10,000 * \frac{1}{H_{i,d}} \sum_{h=1}^H \frac{|r_{i,h,d}|}{\$Vol_{i,h,d}}) \quad (2.11)$$

where  $H_{i,d}$  is the number of hourly buckets per stock day,  $r_{i,h,d}$  is the absolute return on stock  $i$  for day  $d$  during hour  $h$  and  $\$Vol_{i,h,d}$  is the dollar value transacted in that same period.

## 2.4.6 Informational Efficiency Measures

We examine the autocorrelations and standard deviations of mid-quote returns, as well as variance ratios, as high frequency measures of informational efficiency. Non-zero mid-quote return autocorrelations indicate short-term return predictability in stock prices, due to prices deviating from a stochastic



random walk. We calculate the absolute value of first-order mid-quote return autocorrelations at intraday frequencies of 10 seconds and 5 minutes, similar to Hendershott and Jones (2005):

$$Autocorrelation_k = |Corr(r_{k,\tau}, r_{k,\tau-1})| \quad (2.12)$$

where  $r_{k,\tau}$  is the  $\tau$ -th mid-quote return of frequency  $k$ . Taking the absolute value measures the magnitude of both momentum and reversal in stock returns, with larger values indicating greater price inefficiency.

We calculate the standard deviation of mid-quote returns at frequencies of 10 seconds and 5 minutes. A lower return standard deviation indicates less noise in the price discovery process and fewer deviations from the fundamental value due to trading frictions.

Finally, we examine deviations of stock price returns from a random walk, using the property that return variances in an efficient market increase linearly in time. Variance ratios are constructed following the process outlined in Lo and MacKinlay (1988):

$$VarianceRatio_{l,kl} = \left| \frac{\sigma_{kl}^2}{k\sigma_l^2} - 1 \right| \quad (2.13)$$

where  $\sigma_l^2$  and  $\sigma_{kl}^2$  are the variances of l-second and kl-second mid-quote returns for a given stock-day. We utilize the following (l,kl) combinations: (10 seconds, 5 minutes), (1minute, 30 minutes).

#### 2.4.7 Measures of Pre-Trade Fragmentation

Whilst traded market share captures the extent of competition across venues from the perspective of liquidity demanders, competition may also occur among

liquidity suppliers across venues. We therefore develop two metrics of pre-trade market fragmentation. The first measures the percentage of time during which the entrant market is quoting at the national best bid or offer price, as a measure of its quote competitiveness.

$$Chi-XAtNBBO = \frac{1_{Chi-XBestBidPrice=NBBPrice} + 1_{Chi-XBestOfferPrice=NBOPrice}}{2} \quad (2.14)$$

The second measures the fragmentation of pre-trade displayed order volumes at the national best bid and offer prices:

$$Chi-XNBBODepth = \frac{Chi-XDepthAtNBB + Chi-XDepthAtNBO}{TotalDepthAtNBB + TotalDepthAtNBO} \quad (2.15)$$

These measures are used in our robustness tests as alternate (pre-trade) measures of Chi-X market share. They avoid depending on trades actually occurring on Chi-X, alleviating any potential biases that solely relying on post-trade measures may generate.

## 2.5 Summary Statistics and Results

### 2.5.1 Summary Statistics

Table 2.1 provides summary statistics for variables used in the three regression models. Panel A includes the liquidity metrics and Panel B includes the control variables. The first sample contains ASX200 securities during the year following Chi-X's introduction. Quoted, effective and realized spreads averaged 1.166, 0.971 and 0.359 tick increments, respectively. An effective

spread below one tick indicates that a portion of trading activity was executed with price improvement against hidden orders. Share prices average \$7.95 and bid-ask spreads were constrained by the minimum tick size 87% of the day. The second sample contains 29 securities where eligibility for trading on Chi-X was incrementally introduced. Mean quoted, effective and realized spreads are higher for these smaller stocks, at 1.475, 1.309 and 0.839 tick increments, respectively. Indicative of their smaller size, the average share price was \$0.88, with 57% of stock-day observations having a share price between \$0.10 and \$2, and 32% of observations below \$0.10. The final sample contains 80 All Ordinaries securities that commenced competitive quoting on Chi-X on 12 June 2013. Quoted, effective and realized spreads averaged 1.34, 1.18 and 0.66 tick increments, respectively. These securities tended to have narrower spreads than the incremental additions, although both groups averaged \$0.48 million of daily turnover per security.

Figure 2.2 displays equal-weighted quoted spread, effective spread and price impact measures for ASX200 securities during the period beginning one year prior to the introduction of competition in equities trading and ending one year after. Quoted and effective spreads, as well as price impacts, all decline following Chi-X's introduction and over the period.

Figure 2.3 displays quoted spreads over the one year prior to and following Chi-X's launch for ASX200 securities, grouped by tick constraint tercile. Large reductions in quoted spreads are observed among the least constrained tercile, where quoted spreads were equal to the minimum tick size less than 84% of the time, proving preliminary univariate evidence that implicit trading costs decline after Chi-X's entry.

Table 2.1: **Summary Statistics**

This table reports summary statistics for stock-day observations of dependent and independent variables across the three main regression specifications. The first specification consists of ASX200 securities over the one-year period following Chi-X's introduction. The second specification consists of 29 securities that became incrementally eligible for trading on Chi-X (excluding new listings and index rebalances) with an observation period starting three months prior to the first addition and ending three months after the final addition. The third specification consists of 80 All Ordinaries securities that began quoting competitively on Chi-X on 12 June 2013, with an observation period of six months on either side of this date. Quoted spreads are time-weighted across both ASX and Chi-X. Effective spreads are computed based on the prevailing NBBO for transactions across both markets. Price impact is constructed by comparing the NBBO midpoint at the time of each trade with that after twenty seconds. Realized spreads are equal to the effective spreads minus the price impacts. All spreads are calculated in tick increments and basis points. Depth is constructed as the time-weighted dollar value of orders available at the NBBO aggregated across both markets. The percentage of the trading day during which quoted spreads are constrained at the minimum tick size is also reported. Amihud illiquidity is the magnitude of return per hour divided by the turnover transacted in that hour, averaged for each hour to construct a daily metric and then logarithmically transformed. Chi-X market share is the proportion of total daily on-market trading turnover that was traded on Chi-X. Price is the daily time-weighted NBBO midpoint. Ln turnover is the natural logarithm of the daily trading turnover. Medium (low) tick size is equal to one if the daily time-weighted midpoint is between ten cents and two dollars (below ten cents) and zero otherwise. Volatility is calculated daily as the high price minus the low price, as a percentage of the time-weighted midpoint price.

	ASX200 Securities			29 Incremental Stocks			Final 80 Stocks		
	9 Nov 2011 – 8 Nov 2012			10 Oct 2011 – 6 May 2013			12 Dec 2012 – 11 Dec 2013		
	Mean	Median	Std Dev	Mean	Median	Std Dev	Mean	Median	Std Dev
Panel A: Liquidity Metrics									
Quoted Spread (ticks)	1.166	1.065	0.325	1.475	1.031	1.470	1.339	1.106	0.723
Effective Spread (ticks)	0.971	0.935	0.254	1.309	1.000	1.149	1.181	1.000	0.651
Realized Spread (ticks)	0.359	0.371	0.333	0.839	0.756	1.272	0.660	0.643	0.760
Price Impact (ticks)	0.613	0.548	0.456	0.470	0.318	0.899	0.521	0.396	0.707
Quoted Spread (bps)	32.00	27.77	27.26	263.9	179.2	374.0	200.1	164.0	172.3
Effective Spread (bps)	27.21	22.83	24.29	227.3	166.8	267.6	182.7	145.9	167.8
Realized Spread (bps)	12.64	8.356	17.11	158.6	94.28	257.8	113.8	73.25	160.0
Price Impact (bps)	14.57	10.67	15.58	68.63	31.37	150.8	68.93	29.26	124.2
Depth (\$ '0,000s)	41.11	14.23	149.6	14.32	7.232	19.00	11.00	3.949	50.51
Ln Depth	11.95	11.87	1.231	11.16	11.19	1.244	10.69	10.58	1.061
Constrained %	87.31	93.62	16.34	81.02	96.91	29.49	79.36	90.04	26.00
Amihud Illiquidity	1.394	1.228	0.961	2.711	2.474	2.076	3.032	2.932	2.137
Panel B: Control Variables									
Chi-X Market Share (%)	2.605	1.920	2.604	2.272	0.000	5.500	5.902	0.000	11.17
Price (\$)	7.954	3.662	10.93	0.882	0.273	1.744	1.282	0.348	3.855
Turnover (\$ millions)	15.77	5.105	34.21	0.485	0.170	0.932	0.484	0.146	1.205
Ln turnover	15.48	15.45	1.439	11.99	12.04	1.639	11.86	11.89	1.691
Medium Tick Size	0.286	0.000	0.452	0.572	1.000	0.495	0.713	1.000	0.452
Low Tick Size	0.000	0.000	0.000	0.316	0.000	0.465	0.136	0.000	0.343
Volatility (%)	2.226	1.852	1.523	3.404	2.759	3.198	3.804	3.046	3.736

Figure 2.2: **Quoted Spreads, Effective Spreads and Price Impact Across S&P/ASX 200 Securities**

This figure presents daily average quoted spreads, effective spreads and price impacts across S&P/ASX 200 index constituent securities. The observation period begins on 9 November 2010, one year prior to Chi-X Australia’s entry, and spans two years. Metrics are presented in tick increments since there is a substantial level of tick constraint across most securities in the sample. Quoted spreads have a lower bound of one tick increment. Effective spreads may be smaller than one tick if some trades interact with dark orders that offer price improvement.



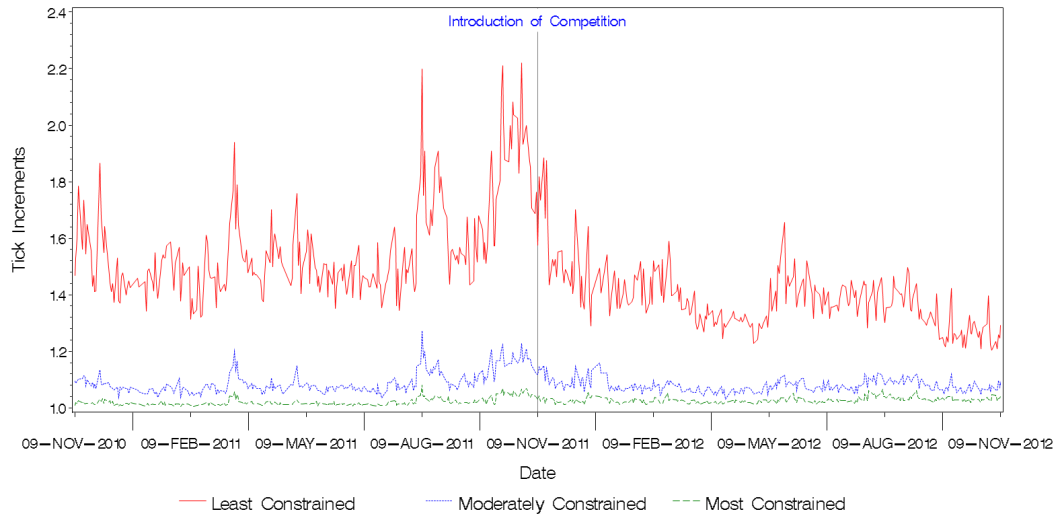
## 2.5.2 Impact of Market Fragmentation

Our first set of results examines the impact of competition by utilizing Chi-X’s market share as a continuous variable with cross-sectional and time series variation. This is constructed over the calendar year starting from the introduction of Chi-X on 9 November 2011. Consequently, the level of fragmentation experienced in each stock-day varies significantly. We find that increasing fragmentation of trading activity reduces quoted spreads. An increase in fragmentation of 10% leads to a reduction of 5% of one tick increment in the quoted spread amongst the ASX 200 securities, as well as quoted spreads being constrained by the minimum tick size 2.8% more frequently.

Table 2.2 presents the results of our analysis. Panel A considers the con-

**Figure 2.3: Quoted Spreads across S&P/ASX 200 Securities by Tick Constraint Tercile**

This figure presents the average quoted spreads for SP/ASX 200 index constituent securities by tick constraint tercile. The observation period begins on 9 November 2010, one year prior to Chi-X Australia’s entry, and spans two years. The equally-sized most constrained, moderately constrained and least constrained groups had a quoted spread of one tick increment more than 95.8% of the time, 84.0% to 95.8% of the time and less than 84.0% of the time, respectively, during the pre-competition period.



solidated orders and trades across both the ASX and Chi-X, whilst Panel B considers orders and trades on the ASX market only. Degryse et al. (2015) suggest that fragmentation may cause a detriment to participants who are only able to access liquidity on the incumbent exchange. Separate presentation of metrics constructed from only the incumbent market allows us to determine how the changes may differentially impact the incumbent versus the entrant market, as well as being an additional robustness for potential issues arising from asynchronous time stamps on each market.

Furthermore, in circumstances where consolidated liquidity across multiple venues is potentially difficult to access, the impact on trading costs is ascertained by changes in liquidity on the incumbent venue only. van Kervel (2015) constructs a theoretical model in which liquidity suppliers duplicate their or-

ders across multiple venues and subsequently cancel their residual orders on other venues after trading on any venue, and empirically confirms this prediction. He also finds that the incidence of fleeting liquidity is higher when there are fewer fast liquidity demanders, less algorithmic trading activity, larger average order sizes, lower turnover and higher realized volatility. Similarly, the third chapter of this dissertation finds that fleeting liquidity is higher when turnover is lower, realized volatility is higher and incoming marketable orders consume all the displayed orders at a price level. In these circumstances, consolidated liquidity in a fragmented trading environment may be more elusive.

We find larger reductions in effective spreads than quoted spreads as Chi-X market share increases, at 0.08 ticks and 0.05 ticks respectively on the consolidated market for a 10% increase in Chi-X market share, representing reduced costs for demanding liquidity. Realized spreads also decline, indicating lower profits for liquidity providers. This implies that traders incurring spread costs are benefited regardless of trading venue, whilst liquidity suppliers on Chi-X are able to attract potentially less-informed order flow. Our results are qualitatively similar when examining changes in spread metrics measured in basis points, which are not reported for brevity. However, no statistically significant change in quoted depth on ASX is observed as fragmentation increases, consistent with order flow migration to Chi-X increasing global liquidity but not local ASX liquidity. An increase in Chi-X market share of 10% is also associated with a 16% increase in quoted depth on the consolidated market, consistent with the duplication of limit orders by market makers documented by Foucault and Menkveld (2008). If volume is duplicated in stocks that are tick constrained on the ASX in order to “jump” time priority, this will result in greater increases in consolidated depth than in ASX-only depth. This increase in depth appears to increase the market’s ability to absorb large trades, lower-

ing the Amihud illiquidity metric on the consolidated market. Overall, these results are consistent with increased competition among liquidity suppliers.

The establishment of an exchange with lower trading fees may also encourage entry of fee-sensitive liquidity providers. Consistent with the lower trading fees on Chi-X, two endogenous liquidity providers (ELPs) entered the Australian equities market in 2011 shortly prior to Chi-X's launch. These ELPs were on the liquidity supplying side of 81% of turnover traded on the entrant venue but only 1% of turnover on the incumbent venue. Their dominance on the entrant venue mirrors the findings of Menkveld (2013) for a large HFT liquidity provider in the Dutch equities market.

The descriptive statistics from Table 2.1 indicate that the median quoted spread of ASX200 stocks over the year following Chi-X's entry was 1.065 tick increments, being constrained by the minimum tick size 94% of the time. Since most of the stocks are tick constrained for much of the trading day, in Table 2.3 we separate our analysis by the degree to which quoting activity was tick constrained in the one year prior to the introduction of competition. We construct terciles of stocks by tick constraint, with the first tercile constrained by the minimum tick size less than 84.0% of the time, the third tercile constrained more than 95.8% of the time, and the second tercile falling between the two. For a Chi-X market share of 10%, the least constrained stocks saw a 0.17 tick reduction in NBBO quoted spreads, whilst no significant change is observed among the two more constrained terciles.

However, these constrained stocks may experience an increase in quoted depth. Indeed, an increase of 10% in entrant market share leads to increases in consolidated quoted depths of between 22% and 23% for the two more constrained stock terciles. No significant change is observed in market depth on the ASX for these terciles. Most of the increases in quoted depth for the



Table 2.2: **Liquidity Metrics for ASX 200 Stocks as Chi-X Market Share Increases**

This table reports changes in quoted spreads, effective spreads, realized spreads and price impacts in tick increments, logarithmic quoted depths, percentage of time quoted spreads are tick constrained and normalized Amihud illiquidity for ASX 200 constituent stocks in the calendar year after Chi-X's introduction. The econometric specification in Equation 2.1 expresses the liquidity metric for stock  $i$  on day  $d$  as the sum of a stock-specific mean, 5-day moving average Chi-X market share, control variables for price, volume, volatility and tick size, and an error term. Panel A presents the consolidated market, including both ASX and Chi-X data, whilst Panel B reports ASX data only. The observation period runs from 9 November 2011 to 8 November 2012. We calculate the change in liquidity metrics for each percentage of market share captured by Chi-X, multiplied by 100, as well as for changes in each control variable, and add a “\*/\*\*/\*\*” to the t-statistic if it is significantly different at the 90%/95%/99% levels. We double cluster standard errors by stock and date.

	Quoted Spread	Effective Spread	Realized Spread	Price Impact	Ln Depth	Constrained %	Amihud Illiquidity
Panel A: Consolidated Market							
Chi-X	-0.46	-0.78	-0.94	0.16	1.63	28.23	-0.86
Market Share	(-3.05)***	(-7.01)***	(-6.39)***	(0.82)	(3.25)***	(3.43)***	(-2.22)**
Price	0.20	0.14	-0.14	0.28	-0.70	-8.72	0.12
	(1.84)*	(2.11)**	(-2.24)**	(2.22)**	(-7.72)***	(-3.40)***	(1.60)
Turnover	-0.03	-0.04	0.02	-0.06	0.48	1.60	-0.41
	(-2.86)***	(-4.94)***	(2.14)**	(-3.90)***	(14.75)***	(4.30)***	(-20.54)***
Volatility	0.03	0.03	-0.03	0.06	-0.22	-1.60	0.18
	(4.05)***	(5.66)***	(-5.35)***	(5.79)***	(-12.25)***	(-7.47)***	(15.57)***
Medium	0.12	0.07	-0.10	0.17	-0.58	-7.67	0.18
Tick	(2.62)***	(2.35)**	(-3.32)***	(3.04)***	(-6.89)***	(-6.13)***	(3.33)***
Adjusted R2	5.0%	3.9%	1.2%	2.7%	20.9%	5.1%	8.7%
# Obs	42,587	42,587	42,587	42,587	42,587	42,587	42,587
Panel B: ASX Only							
Chi-X	-0.32	-0.52	-1.06	0.54	0.12	20.09	-0.35
Market Share	(-2.23)**	(-4.70)***	(-7.14)***	(2.65)***	(0.25)	(2.49)**	(-0.90)
Price	0.21	0.16	-0.14	0.29	-0.70	-9.04	0.13
	(1.89)*	(2.11)**	(-2.21)**	(2.20)**	(-7.59)***	(-3.55)***	(1.62)
Turnover	-0.03	-0.04	0.02	-0.06	0.48	1.58	-0.41
	(-2.80)***	(-4.80)***	(2.19)**	(-3.88)***	(14.44)***	(4.24)***	(-20.55)***
Volatility	0.03	0.03	-0.03	0.06	-0.22	-1.62	0.18
	(4.00)***	(5.46)***	(-5.37)***	(5.68)***	(-12.01)***	(-7.55)***	(15.43)***
Medium	0.13	0.08	-0.10	0.18	-0.59	-7.93	0.18
Tick	(2.68)***	(2.38)**	(-3.32)***	(3.02)***	(-6.81)***	(-6.35)***	(3.36)***
Adjusted R2	5.1%	3.6%	1.2%	2.8%	20.0%	4.9%	8.5%
# Obs	42,587	42,587	42,587	42,587	42,587	42,587	42,587

most tick constrained stocks occur on Chi-X, supporting the queue-jumping hypothesis of Foucault and Menkveld (2008).

Our second set of results investigate the impact of competition by utilizing Chi-X's staggered entry in 29 securities as a natural experiment, via a difference-in-differences regression framework. Table 2.4 documents the impact of Chi-X's introduction for these stocks on measures of liquidity.

Significant declines of 0.35 ticks are observed for quoted spreads at the NBBO quotes. Effective spreads decline by 0.25 ticks on the consolidated market. Similar reductions are observed for realized spreads, while very modest reductions that are not statistically significant are observed for price impacts. These findings indicate a reduction in transaction costs and lower profits to liquidity suppliers, but no change in the informativeness of trading. Quoted depths at the best bid and ask prices, the level of tick constraint, and Amihud illiquidity measure all show no significant changes following the introduction of competition in these stocks. The lack of increased depth or tick constraint could be related to the small size of these securities, with these stocks less frequently constrained by the minimum tick size and usually quoting thinner depths at the NBBO.

**Table 2.3: Liquidity Metrics for ASX 200 Stocks by Tick Constraint as Chi-X Market Share Increases**

This table reports changes in quoted spreads and logarithmic quoted depths for ASX 200 constituent stocks in the calendar year after Chi-X's introduction, grouped into terciles by the proportion of time each stock's quoted spread was constrained by the minimum tick size in the year prior to Chi-X's entry. Tercile thresholds for the proportion of time at minimum tick are 84.0% and 95.8%. The econometric specification in Equation 2.1 expresses the liquidity metric for stock  $i$  on day  $d$  as the sum of a stock-specific mean, 5-day moving average Chi-X market share, control variables for price, volume, volatility and tick size, and an error term. Panel A presents the consolidated market, including both ASX and Chi-X data, whilst Panel B reports ASX data only. The observation period runs from 9 November 2011 to 8 November 2012. We calculate the change in liquidity metrics for each percentage of market share captured by Chi-X, multiplied by 100, as well as for changes in each control variable, and add a “\*/\*\*/\*\*” to the t-statistic if it is significantly different at the 90%/95%/99% levels. We double cluster standard errors by stock and date.

	Least Constrained		Moderately Constrained		Most Constrained	
	Quoted Spread	Ln Depth	Quoted Spread	Ln Depth	Quoted Spread	Ln Depth
Panel A: Consolidated Market						
Chi-X	-1.66	-0.46	-0.10	2.23	-0.02	2.23
Market Share	(-2.88)***	(-0.65)	(-1.23)	(2.63)***	(-0.63)	(4.12)***
Price	0.50	-0.45	0.05	-0.63	0.03	-0.93
	(2.75)***	(-5.62)***	(5.82)***	(-4.92)***	(5.27)***	(-8.15)***
Turnover	-0.09	0.42	-0.01	0.51	0.00	0.48
	(-4.00)***	(16.77)***	(-2.97)***	(7.51)***	(-2.83)***	(17.65)***
Volatility	0.08	-0.17	0.01	-0.23	0.01	-0.24
	(5.39)***	(-12.01)***	(6.97)***	(-6.39)***	(7.18)***	(-12.95)***
Medium	0.42	-0.74	0.07	-0.67	0.05	-0.62
Tick	(4.97)***	(-18.08)***	(12.39)***	(-5.06)***	(5.46)***	(-7.44)***
Adjusted R2	12.0%	14.8%	4.6%	22.2%	6.1%	26.9%
# Obs	14,289	14,289	14,017	14,017	14,281	14,281
Panel B: ASX Only						
Chi-X	-1.18	-1.94	-0.03	0.65	-0.01	0.79
Market Share	(-2.18)**	(-2.74)***	(-0.30)	(0.76)	(-0.38)	(1.49)
Price	0.54	-0.43	0.06	-0.64	0.03	-0.94
	(2.90)***	(-5.43)***	(6.21)***	(-5.04)***	(5.32)***	(-8.14)***
Turnover	-0.09	0.42	-0.01	0.51	0.00	0.47
	(-3.94)***	(16.97)***	(-3.12)***	(7.37)***	(-2.89)***	(17.24)***
Volatility	0.08	-0.17	0.01	-0.23	0.01	-0.24
	(5.42)***	(-12.14)***	(6.96)***	(-6.28)***	(7.21)***	(-12.68)***
Medium	0.43	-0.72	0.07	-0.67	0.05	-0.62
Tick	(5.11)***	(-17.34)***	(12.20)***	(-5.03)***	(5.50)***	(-7.49)***
Adjusted R2	12.1%	15.0%	4.7%	21.1%	6.3%	25.6%
# Obs	14,289	14,289	14,017	14,017	14,281	14,281

Table 2.4: **Liquidity Metrics for Stocks where Chi-X Trading was Incrementally Introduced**

This table reports changes in quoted spreads, effective spreads and price impacts, logarithmic quoted depths, percentage of time quoted spreads are tick constrained and normalized Amihud illiquidity for 29 ASX stocks where eligibility for trading on Chi-X was incrementally introduced, against a control sample of stocks that were not eligible to trade on Chi-X. The econometric specification in Equation 2.3 expresses the liquidity metric for stock  $i$  on day  $d$  as the sum of an intercept, an indicator variable for securities that became eligible to trade on Chi-X, an indicator variable for dates after Chi-X introduction in the treatment stocks and matched control stocks, interaction terms of these two indicator variables, control variables for price, volume, volatility and tick size, and an error term. The observation period runs from 10 October 2011, three months before the first stock commenced, to 6 May 2013, three months after the last stock commenced. We calculate the change in liquidity metrics for the stocks that became eligible to trade on Chi-X, the days after eligibility to trade on Chi-X, and the interaction term between these variables, as well as for changes in each control variable, and add a “\*/\*\*/\*\*” to the t-statistic if it is significantly different at the 90%/95%/99% levels. We double cluster standard errors by stock and date.

	Quoted Spread		Effective Spread		Realized Spread		Price Impact		Ln Depth		Constrained	Amihud
	Ticks	Bps	Ticks	Bps	Ticks	Bps	Ticks	Bps	Bps	%	%	Illiquidity
Treatment	0.33 (1.60)	2.29 (0.13)	0.25 (1.57)	0.37 (0.02)	0.25 (1.72)*	1.67 (0.12)	0.02 (0.36)	-1.72 (-0.30)	0.12 (0.86)	0.69 (0.31)	0.69 (-1.50)	-0.18 (-1.50)
Competition	-0.08 (-0.93)	25.56 (1.84)*	-0.05 (-0.87)	20.78 (1.80)*	-0.08 (-2.23)**	7.48 (0.88)	0.03 (0.58)	13.01 (2.75)***	-0.03 (-0.26)	3.14 (1.65)	-0.12 (-1.26)	-0.12 (-1.26)
Treatment *	-0.35 (-1.98)**	-43.26 (-2.35)**	-0.25 (-1.89)*	-35.36 (-2.16)**	-0.24 (-2.06)**	-25.03 (-1.88)*	-0.02 (-0.42)	-10.66 (-1.76)*	-0.10 (-0.85)	1.41 (0.61)	1.41 (0.61)	0.13 (1.06)
Competition	0.66 (7.53)***	10.87 (33.85)***	0.44 (7.28)***	10.23 (29.73)***	0.22 (4.12)***	8.52 (23.95)***	0.23 (11.55)***	1.95 (19.12)***	-1.06 (-16.12)***	-23.20 (-17.55)***	-23.20 (-17.55)***	1.23 (20.50)***
Price	-0.38 (-6.78)***	-45.13 (-7.28)***	-0.28 (-7.00)***	-33.56 (-7.28)***	-0.23 (-5.64)***	-26.39 (-7.43)***	-0.06 (-6.29)***	-7.34 (-4.12)***	0.60 (16.85)***	11.94 (19.63)***	11.94 (19.63)***	-0.56 (-16.42)***
Turnover	0.11 (5.45)***	17.34 (5.66)***	0.08 (5.81)***	12.76 (6.81)***	0.04 (3.30)***	4.45 (2.70)***	0.05 (8.34)***	8.17 (13.33)***	-0.23 (-8.94)***	-3.34 (-10.77)***	-3.34 (-10.77)***	0.34 (15.32)***
Volatility	1.12 (2.77)***	-123.06 (-7.50)***	0.56 (2.05)**	-91.34 (-7.66)***	0.22 (0.87)	-87.82 (-8.12)***	0.35 (3.92)***	-9.79 (-1.97)**	-2.59 (-9.97)***	-50.64 (-8.14)***	-50.64 (-8.14)***	3.49 (12.13)***
Low	0.20 (0.65)	62.55 (4.66)***	0.03 (0.14)	73.88 (6.23)***	-0.02 (-0.11)	51.47 (5.38)***	0.05 (0.77)	21.41 (6.27)***	-0.77 (-5.04)***	-16.98 (-4.44)***	-16.98 (-4.44)***	1.33 (6.34)***
Tick	30.8% (0.65)	76.3% (4.66)***	26.5% (0.14)	74.3% (6.23)***	9.6% (-0.11)	47.4% (5.38)***	7.8% (0.77)	9.9% (6.27)***	65.6% (-5.04)***	49.7% (-4.44)***	49.7% (-4.44)***	25.5% (6.34)***
Adjusted R2	22,288	22,288	21,705	21,705	21,705	21,705	21,705	21,705	22,288	22,288	22,288	21,539

Our third set of results analyze the introduction of competition in 2013 for the full universe of ASX-listed securities. We utilize the 80 stocks that all start quoting competitively on Chi-X on 12 June 2013, and whose quotes on Chi-X were at the NBBO at least 5% of the time in the month commencing on this date. These stocks are members of the All Ordinaries index, which includes approximately 500 of the largest stocks in Australia, but are significantly smaller than the ASX200 stocks, as evidenced in Table 2.1. Given their size, these stocks are less likely to attract ELPs, which may result in different outcomes from fragmentation. In the worst case, these stocks may experience negative externalities from the division of already-thin liquidity across multiple order books.

Table 2.5 shows that following the initiation of quoting on Chi-X, these small firms experience significant reductions in quoted, effective and realized spreads of between 0.11 and 0.20 ticks, similar to the magnitudes documented for our other samples. We also observe a significant 20% increase in consolidated depth and an increase in time quoted at minimum tick of around 4%. Further, the consolidated market is more resilient, with a reduction in Amihud illiquidity. As in our other samples, most of the liquidity improvements are on the consolidated market. The reduction in ASX spreads, whilst still significant, is not as pronounced. We see an increase in price impact, implying some migration of uninformed traders to Chi-X. Similarly, the increase in depth is less significant, and there is no significance in the Amihud illiquidity reduction, implying that resiliency increases on the consolidated market, consistent with the benefits of competition accruing to participants who are able to access both entrant and incumbent markets.

**Table 2.5: Liquidity Metrics for Full Universe Introduction**

This table reports changes in quoted spreads, effective spreads, realized spreads and price impacts in tick increments, logarithmic quoted depths, percentage of time quoted spreads are tick constrained and normalized Amihud illiquidity for All Ordinaries constituent stocks that commence competitive quoting on Chi-X on 12 June 2013, with an event window of six months on either side of this date. The econometric specification in Equation 2.4 expresses the liquidity metric for stock  $i$  on day  $d$  as the sum of a stock-specific mean, an indicator variable equal to one after Chi-X competitive quoting commencement and zero prior, control variables for price, volume, volatility and tick size, and an error term. Panel A presents the consolidated market, including both ASX and Chi-X data, whilst Panel B reports ASX data only. The observation period runs from 12 December 2012 to 11 December 2013. We calculate the difference in liquidity metrics between the pre- and post- competition periods, as well as for changes in each control variable, and add a “\*/\*\*/\*\*” to the t-statistic if it is significantly different at the 90%/95%/99% levels. We double cluster standard errors by stock and date.

	Quoted Spread	Effective Spread	Realized Spread	Price Impact	Ln Depth	Constrained %	Amihud Illiquidity
Panel A: Consolidated Market							
Competition	-0.11 (-4.37)***	-0.19 (-6.57)***	-0.20 (-8.22)***	0.01 (0.54)	0.20 (3.28)***	4.16 (4.43)***	-0.13 (-2.49)**
Price	0.37 (6.14)***	0.30 (5.46)***	0.07 (2.00)**	0.24 (6.18)***	-0.71 (-7.58)***	-20.39 (-10.37)***	1.11 (11.13)***
Turnover	-0.17 (-5.30)***	-0.16 (-6.16)***	-0.12 (-5.38)***	-0.04 (-3.60)***	0.43 (11.78)***	6.91 (8.69)***	-0.43 (-10.92)***
Volatility	0.06 (4.28)***	0.05 (4.79)***	0.02 (2.36)**	0.03 (6.19)***	-0.15 (-7.39)***	-2.36 (-7.23)***	0.24 (13.37)***
Low Tick	1.02 (7.54)***	0.17 (1.34)	-0.02 (-0.21)	0.19 (2.97)***	-2.50 (-16.24)***	-62.05 (-11.79)***	3.62 (20.05)***
Medium	0.30 (5.63)***	-0.06 (-0.70)	-0.04 (-0.68)	-0.02 (-0.38)	-0.63 (-9.03)***	-19.86 (-7.62)***	0.76 (6.06)***
Adjusted R2	9.4%	8.5%	3.3%	1.6%	28.2%	15.1%	9.2%
# Obs	19,018	19,013	19,013	19,013	19,018	19,018	18,890
Panel B: ASX Only							
Competition	-0.08 (-3.83)***	-0.13 (-5.20)***	-0.18 (-7.86)***	0.05 (3.10)***	0.11 (1.74)*	3.14 (3.29)***	-0.04 (-0.67)
Price	0.39 (6.53)***	0.28 (5.19)***	0.04 (1.35)	0.25 (5.96)***	-0.69 (-7.52)***	-20.90 (-10.57)***	1.08 (10.62)***
Turnover	-0.17 (-5.28)***	-0.16 (-6.63)***	-0.12 (-6.10)***	-0.04 (-3.34)***	0.42 (11.47)***	7.01 (8.75)***	-0.42 (-10.57)***
Volatility	0.06 (4.30)***	0.05 (4.88)***	0.02 (2.28)**	0.03 (6.03)***	-0.15 (-7.27)***	-2.38 (-7.23)***	0.24 (12.62)***
Low Tick	1.04 (7.67)***	0.14 (1.08)	-0.05 (-0.61)	0.19 (2.60)***	-2.43 (-15.63)***	-62.00 (-11.81)***	3.52 (17.06)***
Medium	0.32 (5.71)***	-0.06 (-0.68)	-0.06 (-1.09)	0.00 (0.04)	-0.61 (-9.63)***	-19.96 (-7.29)***	0.76 (5.50)***
Adjusted R2	9.0%	6.6%	2.8%	1.5%	26.6%	14.7%	8.5%
# Obs	19,018	18,998	18,998	18,998	19,018	19,018	18,859

### 2.5.3 Intraday Analysis of Cross-Market Dynamics

Examining liquidity at the daily level, our results thus far are consistent with the duplication of limit orders across the two trading venues. However, within each trading day there is likely to be substantial variation in the magnitude of queue jumping and limit order duplication, as well as marketable order routing choices between the two venues, driven by changes in the relative arrival rate of market and limit orders on each venue.

To investigate intraday liquidity dynamics across the trading venues, we undertake a high frequency analysis of changes in liquidity, fragmentation and order routing in response to market conditions across the venues for ASX200 securities during the first year of Chi-X trading. Trading characteristics are observed by participants in the market, and include quoted depths on the incumbent market, the level of tick constraint and the difference between price impacts across trading venues. Variations in these metrics signal to participants the likelihood of obtaining favorable trading outcomes on the entrant market compared with the incumbent market, and therefore affect the routing of limit orders and market orders to the venues. These metrics also enable a more granular analysis of the mechanisms by which fragmentation reduces quoted spreads and increases quoted depths.

Each trading day from 10:10am to 4:00pm is divided into 70 equally-sized 5-minute intervals over which the liquidity metrics and control variables are constructed. The trading characteristics are lagged by one interval to address endogeneity concerns. Equation 2.16 outlines the regression specification, with control variables taking the same meaning as in Equation 2.1.  $FE_t$  is a time-interval-specific fixed effect that controls for heterogeneity in liquidity at specific times throughout each trading day, which are unrelated to cross-market

dynamics.

$$\begin{aligned}
EfficiencyMetric_{i,t,d} = & \beta_1 TradingCharacteristic_{i,t-1,d} + \beta_2 Price_{i,t,d} \\
& + \beta_3 Turnover_{i,t,d} + \beta_4 Volatility_{i,t,d} + \beta_5 MedTick_{i,t,d} + FE_i + FE_t + \epsilon_{i,t,d}
\end{aligned} \tag{2.16}$$

Several new metrics are constructed to enable our intraday analysis, including a marketable order routing ratio and Chi-X NBBO depth share change. By comparing the fragmentation of marketable orders against the fragmentation of non-marketable orders, we are able to approximate the propensity for marketable orders to be routed to the entrant market, relative to the availability of limit order volume at competitive prices:

$$RoutingRatio = \frac{Chi - X Turnover Market Share}{Chi - X NBBO Depth Share} \tag{2.17}$$

A routing ratio metric higher than one indicates that the entrant market's share of trading volume exceeds its share of displayed liquidity, due to higher routing preference for marketable orders. This metric is normalized to values between zero and one, using the transformation  $Routing = \frac{RoutingRatio}{RoutingRatio+1}$ .

Finally, we examine the change in fragmentation of displayed liquidity across adjacent time intervals, by defining the Chi-X NBBO depth share change as changes in the metric defined in Equation 2.15, from the previous time interval to the current time interval. An increase in Chi-X's NBBO depth share indicates that a larger proportion of total consolidated depth is being quoted on Chi-X in the current period, compared with the previous period. Since the metric measures changes in displayed depth fragmentation across the two venues, it is independent of overall quoted depth.



Table 2.6: **Intraday Liquidity Duplication in ASX 200 Stocks**

This table reports changes in intraday cross-market liquidity dynamics for ASX 200 constituent stocks in the calendar year after Chi-X's introduction, using 5-minute intervals per stock. The econometric specification in Equation 2.16 expresses the intraday liquidity duplication metric for stock  $i$  in interval  $t$  on day  $d$  as the sum of a stock-specific mean, an interval-specific mean, lagged intraday quoting characteristic, control variables for price, volume, volatility and tick size, and an error term. The intraday quoting characteristic is either ASX dollar depth or the percentage of time that quoting is tick constrained. The observation period runs from 9 November 2011 to 8 November 2012. We calculate the change in intraday liquidity dynamics for changes in the quoting characteristic, as well as for changes in each control variable, and add a “\*/\*\*/\*\*” to the t-statistic if it is significantly different at the 90%/95%/99% levels. We double cluster standard errors by stock and date.

	Chi-X Depth Share Change	Chi-X Market Share	Routing	Chi-X Depth Share Change	Chi-X Market Share	Routing
ASX Depth	0.578 (12.06)***	0.611 (6.70)***	2.543 (11.65)***			
Tick				0.007 (11.98)***	0.009 (8.34)***	0.061 (13.65)***
Constraint						
Price	0.179 (2.33)**	-1.225 (-1.75)*	-1.159 (-0.73)	0.016 (0.74)	-1.379 (-2.02)**	-1.699 (-1.21)
Turnover	0.012 (4.40)***	-0.247 (-7.52)***	-0.188 (-3.40)***	0.065 (11.20)***	-0.193 (-5.86)***	0.014 (0.24)
Volatility	0.424 (9.35)***	-0.617 (-4.13)***	-3.046 (-8.62)***	0.051 (2.78)***	-0.965 (-5.64)***	-4.061 (-9.30)***
Medium	0.28 (3.97)***	-0.924 (-1.59)	-0.331 (-0.25)	0.033 (2.43)**	-1.169 (-2.02)**	-1.228 (-0.99)
Adjusted R2	1.9%	0.3%	0.5%	0.4%	0.2%	0.4%
# Obs	2,834,363	2,834,363	2,555,753	2,834,363	2,834,363	2,555,753

Table 2.6 presents the results of this analysis, allowing us to observe changes in the dependent variable one period in the future in response to changes in the current level of the independent variables. It provides evidence on how displayed liquidity duplicates or migrates across the two venues depending on the quoting characteristics of the securities throughout the previous five-minute interval. It also examines the change in trade executions contemporaneously with the changes in liquidity. Consistent with the queue-jumping hypothesis of Foucault and Menkveld (2008), we show that longer queues on ASX, which are represented by higher depth, lead to an increase in displayed depth share on Chi-X at the NBBO, more executions on Chi-X and a larger proportion of marketable orders being routed to Chi-X relative to its displayed depth share, in the next period. Similarly, when spreads are constrained by the minimum tick, traders are more likely to look for alternate methods of execution. This could be by “tick-splitting” in the dark as in Kwan et al. (2015), or by creating a new queue on an entrant venue. When traders are constrained by the tick size in the present period, they are more likely to post both limit and market orders on Chi-X rather than ASX in the subsequent period.

#### **2.5.4 Impact of Trading Fee Reduction**

We utilize the ASX’s reduction in trading fees on 1 July 2010 in anticipation of Chi-X’s entry to examine the extent to which reductions in explicit fees, which reduce order processing costs for liquidity providers, are passed on to liquidity demanders. Table 2.7 documents decreases in quoted spreads of 0.02 ticks and 0.34 basis points after the reduction in explicit trading fees in 2010. This compares to a reduction in explicit trading fees of 0.26 basis points (0.13 per side) by the ASX over the same period. This “passing on” of explicit fee

reductions to the market is indicative of efficient competition among liquidity suppliers and is consistent with the findings of Malinova and Park (2015) in Canada following changes in maker-taker trading fees.

Realized spreads, a measure of the revenues attributable to liquidity provision and defined as the effective spread earned less the price impact incurred, declined 0.95 basis points. The significantly higher reductions observed in effective and realized spreads, combined with increases in depth and tick constraint, along with reductions in Amihud illiquidity, imply that the reduction in fees may have facilitated increased competition between liquidity providers.

**Table 2.7: Liquidity Changes in ASX200 Stocks after ASX's 2010 Trading Fee Reduction**

This table reports changes in quoted spreads, effective spreads, realized spreads and price impacts, logarithmic quoted depths, percentage of time quoted spreads are tick constrained and normalized Amihud illiquidity for ASX200 constituent stocks from the period six months prior to until six months after ASX's 2010 trading fee reduction. The econometric specification in Equation 2.5 expresses the liquidity metric for stock  $i$  on day  $d$  as the sum of a stock-specific mean, fee reduction indicator variable equal to zero prior to ASX's trading fee reduction on 1 July 2010 and one after, controls for price, volume, volatility and tick size, and an error term. The observation period runs from 1 January 2010 to 31 December 2010. We calculate the difference in liquidity metrics between the pre- and post- fee change periods, as well as for changes in each control variable, and add a “\*/\*\*/\*\*” to the t-statistic if it is significantly different at the 90%/95%/99% levels. We double cluster standard errors by stock and date.

	Quoted Spread		Effective Spread		Price Impact		Realized Spread		Ln Depth		Constrained	Amihud
	Ticks	Bps	Ticks	Bps	Ticks	Bps	Ticks	Bps			%	Illiquidity
Fee	-0.02	-0.34	-0.05	-1.26	-0.04	-0.27	0.00	-0.95	0.11	1.27		-0.09
Reduction	(-2.68)***	(-2.50)**	(-6.83)***	(-6.33)***	(-3.30)***	(-0.79)	(-0.30)	(-2.97)***	(4.55)***	(3.29)***		(-4.97)***
Price	0.11	50.09	0.08	47.92	0.23	13.06	-0.14	35.19	-1.08	-6.52		0.28
	(5.43)***	(181.25)***	(4.82)***	(44.64)***	(9.32)***	(6.75)***	(-8.90)***	(12.11)***	(-10.53)***	(-8.26)***		(4.70)***
Turnover	-0.03	-1.05	-0.04	-1.61	-0.07	-2.12	0.03	0.48	0.52	1.94		-0.38
	(-4.99)***	(-8.16)***	(-7.84)***	(-5.58)***	(-7.16)***	(-6.43)***	(3.42)***	(1.45)	(13.32)***	(7.01)***		(-18.26)***
Volatility	0.03	0.74	0.02	0.59	0.08	2.56	-0.05	-1.95	-0.21	-1.70		0.18
	(7.13)***	(10.84)***	(7.12)***	(4.38)***	(11.24)***	(15.01)***	(-11.53)***	(-9.87)***	(-9.71)***	(-8.93)***		(14.77)***
Medium	0.11	-19.43	0.10	-18.74	0.25	-2.52	-0.15	-16.33	-0.54	-7.74		0.14
Tick	(6.12)***	(-50.75)***	(7.16)***	(-20.06)***	(9.21)***	(-1.54)	(-6.64)***	(-6.76)***	(-8.16)***	(-7.38)***		(2.18)**
Adjusted R2	1.8%	95.9%	3.0%	73.9%	3.8%	5.2%	2.0%	23.4%	29.5%	6.0%		9.4%
# Obs	44,678	44,678	44,678	44,678	44,678	44,678	44,678	44,678	44,678	44,678		44,678

### 2.5.5 Impact on Price Efficiency and Volatility

Market quality consists not only of liquidity metrics, but also the informational efficiency of prices. Table 2.8 reports on the impact of exchange fee reduction and competition in trading on the efficiency and volatility of the price formation process. For brevity, only the coefficient estimates and t-statistics corresponding to the competition explanatory variable are reported. Each regression model is run with a price discovery metric as the dependent variable, the competition variable as the main independent variable, as well as the same control variables for stock price, trading turnover, volatility, tick size and stock fixed effects as the main specifications previously described. The observation windows are also the same as the main specifications. Ideally, increases in competition between market makers should improve the efficiency of prices and reduce volatility by making prices more resilient. We calculate a short-term and long-term measure of the autocorrelation of absolute midpoint returns, the standard deviation of midpoint returns, and the variance ratio, which assesses how closely prices resemble a random walk.

The reduction in explicit fees by the ASX in 2010 causes a noticeable improvement in the short-term efficiency of prices, reducing short-term autocorrelation, reducing standard deviations of returns and reducing the short-run variance ratio. Turning to the commencement of Chi-X for trading ASX200 stocks, we observe significant improvements across almost all the short-term efficiency and volatility measures, regardless of whether this is measured using the incumbent or consolidated market. Given these are the most liquid, competitively quoted stocks in Australia, it is intuitive that the improvements would be observed primarily in short-run metrics.

The 29 incremental additions and the remainder of the full universe of the

ASX consist of much smaller securities. As such, it is likely they are quoted less competitively, and that they will benefit less from the fragmentation of trading. Consistent with this notion, the benefits for efficiency and volatility are concentrated in the consolidated market and are stronger for the longer horizon measures than for the shorter horizon measures. While changes in many of the metrics are insignificant, none display a significant deterioration.

These results are consistent with an increase in competition between market makers reducing volatility and improving the efficiency of price discovery, both for small and large stocks. Taken with the improvements in liquidity documented in previous sections, these results indicate that market quality overall has been improved.

### **2.5.6 Robustness Tests**

We examine the robustness of our findings for Chi-X's entry in ASX200 stocks to five separate alternative specifications. Table 2.9 presents the results from these robustness tests. For brevity, we have reported only the coefficient estimate and t-statistic on the competition variable for each regression. Each specification is run with a liquidity metric as the dependent variable, the competition variable as the main independent variable, as well as the same control variables for stock price, trading turnover, volatility, tick size and stock fixed effects as the main specification.

Specification (1) omits the first six months after Chi-X's entry to allow time for broker connectivity and order routing systems to be established, examining the six months from 9 May 2012. Specification (2) utilizes the actual level of fragmentation per stock-day, rather than the lagged moving average used to avoid the potential issue of endogeneity. Specification (3) and (4) utilize a

**Table 2.8: Changes in Price Efficiency Metrics after Competition Introduction and Fee Reduction**

This table reports changes in various return autocorrelations, return standard deviations and variance ratios for different model specifications. The econometric specification in Equation 2.4 expresses the price efficiency metric for stock  $i$  on day  $d$  as the sum of a stock-specific mean, an indicator variable equal to one after Chi-X introduction or ASX fee reduction and zero prior, control variables for price, volume, volatility and tick size, and an error term. We calculate means and event-effects based on the model estimates across each sample. The samples consist of the ASX fee reduction on 30 June 2010, introduction of Chi-X trading for ASX200 securities on 9 November 2011 and commencement of competitive quoting of All Ordinaries securities on Chi-X on 12 June 2013. These observation windows span six months prior to and six months following each event date. The final sample consist of incremental additions of securities eligible to trade on Chi-X from 10 January 2012 to 6 February 2013, with an observation period from three months prior to the first addition to three months after the last addition, using the specification in Equation 2.3. We calculate the difference between price efficiency metrics for changes in each explanatory variable, but for brevity report only the coefficient on the fee change or Chi-X introduction indicator variable, and add a “\*/\*\*/\*\*” to the t-statistic if it is significantly different at the 90%/95%/99% level. We double cluster standard errors by stock and date.

	Autocorrelations		Standard Deviations		Variance Ratios	
	10 sec	5 min	10 sec	5 min	10 sec / 5 min	1 min / 30 min
Panel A: ASX Only						
2010 Fee	-0.004	0.001	-0.839	-1.758	-0.278	0.036
Change	(-4.16)***	(0.76)	(-8.27)***	(-5.74)***	(-5.30)***	(1.27)
2011 ASX200	-0.003	0.001	-0.729	-2.192	-0.11	0.001
	(-2.21)**	(0.46)	(-7.12)***	(-6.89)***	(-2.40)**	(0.05)
2013 Full	0.000	-0.002	1.104	1.273	-0.378	0.07
Universe	(0.05)	(-0.77)	(0.67)	(0.79)	(-2.39)**	(0.67)
Incremental	-0.009	-0.011	-7.426	-7.759	0.096	-0.42
Additions	(-0.91)	(-1.68)*	(-1.33)	(-1.55)	(0.31)	(-1.90)*
Panel B: Consolidated Market						
2011 ASX200	-0.003	0.001	-0.768	-2.181	-0.153	-0.002
	(-2.25)**	(0.43)	(-7.52)***	(-6.85)***	(-3.36)***	(-0.07)
2013 Full	-0.022	-0.017	-7.376	-5.236	-0.451	-0.053
Universe	(-5.66)***	(-5.43)***	(-4.13)***	(-3.19)***	(-2.98)***	(-0.55)
Incremental	-0.012	-0.013	-8.491	-8.397	0.033	-0.469
Additions	(-1.30)	(-1.95)*	(-1.53)	(-1.67)*	(0.11)	(-2.14)**

threshold dummy-variable for fragmentation, turning from 0 to 1 when stock-day fragmentation exceeds 2% and 5% respectively. This recognizes that the effects of fragmentation might not increase linearly. Specification (5) lengthens the unit of analysis from stock-days to stock-months, motivated by the methodology of O'Hara and Ye (2011). This longer aggregation interval ensures that changes in liquidity are persistent and not driven by time-series spurious correlation. For all specifications, we find significant reductions in effective spreads and increases in tick constraint, as well as improvements in quoted spreads and depths on the consolidated market, consistent with our primary regression specification.



Table 2.9: **Robustness Tests**

This table reports changes in liquidity metrics for ASX200 constituent stocks using various model specifications. The econometric specification in Equation 2.1 expresses the liquidity metric for stock  $i$  on day  $d$  as the sum of a stock-specific mean, alternative measures of Chi-X market presence, control variables for price, volume, volatility and tick size, and an error term. We calculate means and event-effects based on the model estimates under different model specifications to quantify the level of competition. The continuous observation window in (1) runs from 6 months to 1 year after Chi-X's introduction. The model specification in (2) utilizes actual Chi-X market share, rather than a lagged moving average. The event-variables in (3) and (4) are indicators equal to one for stock-days with Chi-X market share exceeding 2% and 5% of total trading turnover, respectively, and zero otherwise. Regression analysis in (5) is conducted by aggregating stock-day observations by month, similar to the methodology in O'Hara and Ye (2011). The baseline observation period runs from 9 November 2011 to 8 November 2012. We calculate the change in liquidity metrics for changes in each explanatory variable, but for brevity report only the coefficient on the variable for Chi-X market presence, and add a “\*/\*\*/\*\*” to the t-statistic if it is significantly different at the 90%/95%/99% level. We double cluster standard errors by stock and date for (1) to (4) and by stock for (5).

	Consolidated Market					ASX Only				
	(1)	(2)	(3)	(4)	(5)	(1)	(2)	(3)	(4)	(5)
Quoted	-0.79	-0.40	-0.02	-0.01	-0.62	-0.64	-0.29	-0.02	-0.01	-0.47
Spread	(-3.28)***	(-4.48)***	(-3.38)***	(-2.29)**	(-3.14)***	(-2.92)***	(-3.68)***	(-2.65)***	(-1.60)	(-2.56)**
Effective	-0.94	-0.53	-0.03	-0.03	-1.09	-0.85	-0.34	-0.02	-0.02	-0.78
Spread	(-5.59)***	(-7.92)***	(-5.93)***	(-7.46)***	(-7.77)***	(-4.96)***	(-5.46)***	(-3.65)***	(-5.89)***	(-5.82)***
Realized	-0.50	-0.69	-0.03	-0.03	-1.28	-0.59	-0.91	-0.03	-0.04	-1.46
Spread	(-2.81)***	(-7.95)***	(-5.36)***	(-5.33)***	(-7.76)***	(-3.33)***	(-10.27)***	(-5.98)***	(-6.01)***	(-8.87)***
Price	-0.45	0.17	0.00	0.00	0.18	-0.25	0.57	0.02	0.01	0.68
Impact	(-1.56)	(1.65)*	(0.29)	(0.44)	(0.82)	(-0.87)	(5.33)***	(2.27)**	(1.56)	(3.11)***
Ln Depth	3.46	2.27	0.04	0.10	1.30	2.66	1.21	-0.01	0.05	-0.41
	(6.45)***	(7.53)***	(2.19)**	(4.51)***	(3.46)***	(5.01)***	(3.93)***	(-0.78)	(2.08)**	(-1.10)
Constraint	40.68	26.02	1.14	0.87	36.68	31.01	19.58	0.84	0.66	27.15
%	(4.94)***	(5.50)***	(3.46)***	(3.33)***	(4.34)***	(4.10)***	(4.39)***	(2.56)**	(2.65)***	(3.40)***
Amihud	-1.25	-0.96	-0.02	-0.05	-0.67	-0.80	-0.43	-0.01	-0.03	-0.08
Illiquidity	(-2.54)**	(-4.42)***	(-1.58)	(-3.09)***	(-2.13)**	(-1.64)	(-1.87)*	(-0.51)	(-1.86)*	(-0.24)

We further assess the robustness of our findings to the typically examined post-trade measures of turnover market share, by instead using pre-trade measures, being Chi-X's percentage of time quoting at the NBBO and its proportion of total NBBO depth. Apart from price impact when Chi-X participation is measured by the percentage of time quoting at the NBBO, Table 2.10 reports results that are broadly consistent with our main analysis. Spreads are uniformly decreasing in pre-trade Chi-X market share and are significant in all but one case. We also see consistent coefficient estimates for depth, tick constraint and Amihud illiquidity, though these are at times insignificant. These weaker results are consistent with the arguments of Foucault and Menkveld (2008) that in order to realize the benefits of competition among trading venues, it is important not only that alternative venues exists, but also that traders are connected to the entrant market.

## 2.6 Conclusion

The introduction of competition for equities trading in Australia has led to an increasing level of order flow fragmentation, which ultimately reduces both implicit and explicit transactions costs, which are two financial frictions on trading. Using both a natural experiment and difference-in-differences approach, we find that quoted and effective spreads for stocks exposed to competition decline as an increasing proportion of order flow migrates from the incumbent to the entrant exchange. In the year since Chi-X's launch, quoted spreads among the ASX200 index constituents averaged 1.16 ticks and declined 0.05 ticks for each 10% increase in Chi-X market share. The reduction in incumbent trading fees prior to Chi-X's entry enables us to show that while liquidity providers pass through these reductions in explicit fees, which reduce their or-



der processing costs, most of the improvements in liquidity are attributable to the entry of two new fee-sensitive electronic liquidity providers who arrive only with the new entrant market. Importantly, while we observe that fragmentation holds market-wide benefits, it can have negative consequences for liquidity on the incumbent market, which could disadvantage traders who connect only to the incumbent market.

Consistent with the intermarket queue-jumping strategies predicted by Foucault and Menkveld (2008) for stocks that were previously most tick constrained, we observe the entrant market contributing to quoted depth increases, with no significant change in quoted spreads. For stocks that were previously least tick constrained, quoted spreads declined. An intraday analysis of cross-market dynamics identifies that it is necessary for traders to both quote and trade on the entrant market for these benefits to be realized. Our intraday analysis further identifies that queue-jumping and order duplication on the entrant are most pronounced when stocks are tick constrained or when queues on the incumbent market are long.

Our findings have implications for the recent debates around tick size, fragmentation and maker-taker trading fees. Specifically, regulators and market participants should carefully consider the impact of any market design changes on the incentives of liquidity suppliers and order processing costs they incur, in addition to the level of fragmentation in trading activity.

## Chapter 3

# The Value of a Millisecond: Harnessing Information in Fast, Fragmented Markets

*“I’d say, ‘Watch closely. I am about to buy one hundred thousand shares of Amgen. I am willing to pay forty-eight dollars a share. There are currently one hundred thousand shares of Amgen being offered at forty-eight dollars a share — ten thousand on BATS, thirty-five thousand on the New York Stock Exchange, thirty thousand on Nasdaq, and twenty-five thousand on Direct Edge.’ You could see it all on the screens. We’d all sit there and stare at the screen and I’d have my finger over the Enter button. I’d count out loud to five.*

*One. Two. See, nothing’s happened. Three. Offers are still there at forty-eight. Four. Still no movement. Five. Then I’d hit the Enter button and — boom! — all hell would break loose. The offerings would all disappear, and the stock would pop higher. At which point he turned to the guys standing behind him and said, ‘You see, I’m the event. I am the news.’”*

‘Flash Boys: A Wall Street Revolt’, Lewis (2015), page 34

### 3.1 Introduction

The desire for speed to reduce search costs and realize gains from fleeting trading opportunities has led traders to seek faster access to financial markets. While marginal increases in speed have diminished over time, what persists is that speed provides an advantage to those who possess it over those who do not, leading to a perpetual high frequency trading (HFT) arms race (Budish et al., 2015). Faster traders capture most of the profits from liquidity provision (Baron et al., 2019, Roşu, 2019) and impose adverse selection costs on relatively slower counterparts (Baldauf and Mollner, 2018, Li, 2018). Stock exchanges and technology vendors have facilitated speed competition among traders with recent innovations such as co-location (Brogaard et al., 2015), microwave towers (Shkilko and Sokolov, 2016) and faster proprietary data feeds (Goldstein et al., 2018).

As communications technology approaches the speed of light (Angel, 2014), it is increasingly costly and complex for individual traders to further increase their trading speed. Several equities trading venues have implemented or proposed “speed bumps”, which apply systematic pauses before orders are processed, to change the dynamics of speed competition among traders. There are two main nuances in speed bump design. Firstly, the delay can be applied to only some types of traders or orders, rather than all orders, to adjust speed differentials among traders and alter equilibrium order matching outcomes. Secondly, the delay duration, which is typically measured in milliseconds or microseconds, can either be fixed or randomly drawn from a specified distribution. Academics, practitioners and securities regulators continue to debate the benefits and harms of various speed bump designs and are particularly concerned about the first order effects on order matching outcomes and the

second order effects on overall market quality. These uncertainties are highlighted by U.S. Securities and Exchange Commission (SEC) staff approval of the Chicago Stock Exchange’s speed bump proposal on October 19, 2017, only to be put on hold by the Commission five days later for further review,<sup>1</sup> and subsequent application withdrawal.<sup>2</sup>

In this chapter, we investigate the introduction of an asymmetric, randomized duration speed bump on the Canadian stock exchange TSX Alpha on September 21, 2015. This mechanism delays incoming order messages by a duration that is continuously distributed between 1 and 3 milliseconds but provides the option to pay a higher trading fee to enter and cancel a special type of non-marketable limit order without being delayed. Intuitively, it ensures that large liquidity demanders are unable to time their marketable orders to arrive on Alpha and other venues simultaneously. Fast liquidity suppliers on Alpha are provided with a short window during which they can cancel standing orders after observing trades elsewhere, in the spirit of the excerpt from Michael Lewis’ popular culture novel ‘Flash Boys’ quoted at the beginning of this chapter. Foucault and Moinas (2018) assert that trades sufficiently large to consume all liquidity in a stock at a price level on one venue constitute material information for pricing the stock on other venues. Malinova and Park (2017) report that multi-venue trades have double the price impact of single-venue trades, even after controlling for trade size and trader type. Therefore, Alpha’s speed bump implementation could allow fast liquidity suppliers to selectively avoid multi-venue trades with relatively higher adverse selection, to increase their profits from liquidity provision.

We start by examining the institutional details of TSX Alpha’s asymmetric

---

<sup>1</sup>See <https://www.sec.gov/rules/sro/chx/2017/34-81913-letter-from-secretary.pdf>.

<sup>2</sup>See <https://www.sec.gov/comments/sr-chx-2017-04/chx201704-4118079-171622.pdf>.

randomized duration speed bump. We analyze usage of the delay-exempt non-marketable orders and the accessibility of displayed liquidity, finding that fast liquidity providers on Alpha cancel their delay-exempt orders in response to trades on other venues. Alpha captures a larger proportion of single-venue trades as well as trades without immediate price impact, resulting in lower adverse selection. Liquidity provider profits increase immensely because they do not impound these savings into narrower bid-ask spreads, potentially due to lack of competition among fast liquidity providers on Alpha, especially those willing to post the minimum size requirement for delay exempt orders.

Adverse selection is redistributed to other venues, consistent with the theoretical model in Biais et al. (2015), where a slow venue for single-venue trades operates in parallel with fast venues for multi-venue trades. Liquidity providers on CX2, which operates a similar trading fee structure to the relaunched TSX Alpha, narrow their bid-ask spreads to compete more effectively, resulting in sharply lower profits from supplying liquidity. Bid-ask spreads on other venues naturally widen to protect against higher adverse selection, and aggregate liquidity declines. Overall, these findings are consistent with a much earlier literature on the impact of order flow segmentation through payment for order flow schemes (Chakravarty and Sarkar, 2002, Easley et al., 1996) and also with more recent studies on the segmentation of less informed order flow by dark trading venues (Comerton-Forde and Putniņš, 2015, Zhu, 2014) and venues with inverted maker-taker trading fees (Maglaras et al., 2015).

An asymmetric speed bump is an impediment on liquidity demanders, which reduces their monitoring intensity in the context of Foucault et al. (2013)'s model of maker-taker liquidity cycles. To increase its appeal to liquidity demanders, TSX Alpha adopts an inverted maker-taker trading fee structure, as well as minimum size requirements for speed bump exempt orders.



Battalio et al. (2016) suggest that in a fragmented trading environment, adverse selection and realized spreads vary across venues depending on their relative trading fee structures. To disaggregate the impact of new Alpha’s speed bump and inverted trading fee structure, we compare post-relaunch adverse selection and liquidity provision on TSX Alpha with CX2, which has an almost identical trading fee structure and similar market share. While fast liquidity providers tend to be more profitable in general (Roşu, 2019), on Alpha they achieve realized spreads that are half a cent larger than on CX2. Conversely, slower traders do worse on Alpha than on CX2, suggesting a redistribution of adverse selection not only across venues but also across speed tiers within venues (Han et al., 2014). A nuance of Alpha’s speed bump design is that traders seeking to harness the speed advantage need to send order messages with sub-millisecond reaction times.

Applying the notion that short-term price movements represent ‘news’ in modern financial markets (Foucault et al., 2016), we utilize variation in the frequency of price movements across stocks as a proxy for the information arrival rate in Du and Zhu (2017), to examine cross-sectional differences in the speed bump’s impact. Redistribution of adverse selection away from Alpha is heavily concentrated in stocks where trades frequently consume all available depth at the best price level, which mechanically moves the price. While inverted venues tend to have higher market shares in lower priced, tick constrained stocks where prices move less frequently (Yao and Ye, 2018), these differences suggest that Alpha’s speed bump complements its inverted fee structure by helping liquidity suppliers avoid adverse selection in higher priced stocks as well.

Theory has suggested that the impact of differential trading speeds on liquidity is dependent on whether it is used by liquidity providers to manage

adverse selection risk (Jovanovic and Menkveld, 2016), or by arbitrageurs to pick off stale quotes (Budish et al., 2015). An extensive empirical literature supports these predictions, with studies finding that faster liquidity providers can improve displayed liquidity (Brogaard et al., 2015, Menkveld, 2013), while faster liquidity demanders can harm liquidity (Foucault et al., 2017, Shkilko and Sokolov, 2016). Consistent with this notion, some models suggest that asymmetric speed bumps can assist liquidity suppliers in avoiding adverse selection, which enables them to increase liquidity provision, although redistribution of adverse selection to other venues creates externalities (Biais et al., 2015, Brolley and Cimon, 2018). We provide empirical evidence to assess these predictions, finding that differential speed is valuable to liquidity suppliers.

This chapter makes several methodological contributions by developing new measures that can be constructed with data that is commonly available to empirical researchers. First, we develop a methodology to aggregate related trades on different venues into trade strings, which we use to examine low latency, cross-market liquidity dynamics. The metric is constructed in a way that is robust to trade direction assignment and timestamp synchronization issues identified in O’Hara (2015). Second, we propose two innovative classification schemes to analyze a venue’s order flow composition. Similar to Malinova and Park (2017) and van Kervel (2015), we distinguish between single-venue trade strings and multi-venue trade strings, with the latter likely to have originated from a smart order router (SOR). Additionally, we define depleting trade strings as those that are immediately followed by adverse selection, which is a necessary condition for profits to be realized. These metrics are used to quantify order flow segmentation across trading venues after Alpha’s relaunch. Finally, we construct a quote fade metric that measures the accessibility of liquidity in fast, fragmented markets.

Closely related to this chapter is Anderson et al. (2018), which also investigates changes in consolidated liquidity following TSX Alpha’s relaunch with a speed bump. They find a mild reduction in effective spreads and increase in execution size on Canadian trading venues, relative to the United States, using a difference-in-differences regression approach. There are two main differences between the methodology and findings of their study and this chapter.

Firstly, this chapter focuses on segmentation of single-venue and multi-venue trades, which is the order flow characteristic that an asymmetric randomized speed bump is designed to filter. In contrast, Anderson et al. (2018) focus on segmentation of retail trades. Notably, Alpha is a public exchange venue that is accessible to all brokers and their clients, which contrasts with dark pool venues, payment for order flow arrangements or broker internalization systems that are able to segment order flow based on the identity of the trader. Examining the Canadian equities market, Malinova and Park (2017) find that retail brokers frequently utilize multi-market trades, and that retail liquidity demanders exhibit higher average price impact than institutional liquidity demanders, up to a post-trade horizon of one minute. Additionally, they find that multi-market trades have double the price impact of single-market trades, even after controlling for trade size and trader type.

Secondly, in addition to examining changes in market-wide liquidity, this chapter also investigates the competitive dynamics in liquidity provision on each trading venue, to disaggregate the effects on gains from trade for different types of traders. Anderson et al. (2018) report a large increase in average intraday volatility in the United States over their observation period but no corresponding increase in Canada. We examine changes in only Canadian market liquidity, rather than a difference-in-differences comparison with the United States, as this observation indicates that it might not be a suitable

control sample.

The remainder of this chapter is organized as follows. Section 3.2 outlines the institutional details of the Canadian equities trading landscape, the relaunched TSX Alpha’s speed bump and other concurrent market design changes. Section 3.3 describes the data and metrics. Section 3.4 discusses the empirical methodology and findings for order flow segmentation and impact on liquidity on each venue, as well as the consolidated market. Section 3.5 concludes and discusses implications for practitioners, securities regulators and future researchers. Section 3.6 provides a holistic overview of the different permutations of speed bump design. Section 3.7 provides supplementary materials on the development of novel metrics.

## **3.2 Institutional Details**

Canadian equities trading is fragmented across multiple marketplaces, with six lit trading venues and three dark trading venues. Securities are listed on the Toronto Stock Exchange, operated by the TMX Group, which retains approximately 60% of trading activity. The TMX Group also operates Alpha and TMX Select, which was decommissioned once the changes on Alpha occurred. Chi-X and CX-2 were operated by Chi-X Canada and have since been acquired by Nasdaq. Notably, CX-2 has an inverted maker/taker trading fee structure that is almost identical to that of the relaunched Alpha. Other venues include Omega, Pure Trading, Aequis Neo, Aequis Lit and three dedicated continuous dark pools, Match Now, Instinet and Liquidnet. Intermarket price priority is enforced via an order protection rule, similar to trade-through prohibitions in the United States, although it applies to all price levels and not only the top of the order book.

Unlike in the United States, internalization of retail order flow in Canada has been significantly constrained by regulation. Brokers wishing to internalize trades of less than 5,000 shares are required to provide one full tick of price improvement, or a half tick when the bid-ask spread is one tick wide (Larrymore and Murphy, 2009). This requirement has prevented the growth of retail internalization venues such as those common in the United States, which account for around 22 percent of trading (Kwan et al., 2015). As a result of this regulation and the prohibition of payment for order flow, retail orders, which are generally considered to exhibit lower adverse selection and therefore greater profitability for the liquidity supplying counterparty, remain predominantly on-exchange in Canada.

### **3.2.1 TSX Alpha Relaunch and Speed Bump**

On 21 September 2015, the Canadian stock exchange Alpha was relaunched as TSX Alpha, with the following market design changes:<sup>3</sup>

1. a randomized speed bump of 1-3 milliseconds for all orders except non-marketable Post Only limit orders,<sup>4</sup> which have a stock-specific minimum size requirement;
2. an inverted maker-taker trading fee model;
3. removal of marketplace protection under trade-through prohibitions; and
4. migration from the TMX Quantum to the TMX Quantum XA trading platform.<sup>5</sup>

---

<sup>3</sup>For brevity, we often refer to Alpha prior to the relaunch with speed bump as “old” Alpha, and as “new” Alpha after its relaunch on 21 September 2015.

<sup>4</sup>Post Only orders are rejected if they are received at a price that is marketable.

<sup>5</sup>The fourth chapter of this dissertation examines Toronto Stock Exchange’s upgrade to the Quantum XA trading platform in 2014, finding increases in market share and relative adverse selection on TSX, alongside higher fleeting liquidity on other trading venues.

Prior to Alpha’s speed bump implementation, several brokers submitted that the proposal could result in undesirable market quality consequences. For example, TD Securities suggested that *“the introduction of speed bumps on both Alpha and Aequitas will slow down the operation of smart order routers . . . aggravating quote fade across all marketplaces”*<sup>6</sup> while ITG Canada asserted that *“the new Alpha design will allow passive post only resting orders the ability to fade should they see trading on another venue.”*<sup>7</sup>

Figure 3.1 depicts the choice of large Canadian equities trading venues facing liquidity demanders after Alpha’s relaunch. Large investors who require more liquidity than what is displayed at the best price level on any single trading venue typically utilize a smart order router (SOR) to spray marketable orders across multiple trading venues simultaneously, efficiently accessing consolidated liquidity across all venues. In this spirit, O’Hara and Ye (2011) surmise that SOR and trade-through prohibition virtually re-consolidate fragmented trading venues, while practitioners report that SORs are typically able to access almost 100% of consolidated liquidity displayed at the best price.<sup>8</sup> Alpha’s asymmetric, randomized speed bump for incoming marketable orders but not delay-exempt non-marketable limit order entries or cancellations potentially enables its fast liquidity suppliers to observe the first portions of large SOR sprays being executed on other venues and subsequently cancel their limit orders on Alpha within the 1-3 millisecond window. In this way, they are able to avoid interacting with larger SOR sprays that likely impose higher adverse selection costs.


---

<sup>6</sup>See [https://www.osc.gov.on.ca/documents/en/Marketplaces/com\\_20141208\\_td-securities.pdf](https://www.osc.gov.on.ca/documents/en/Marketplaces/com_20141208_td-securities.pdf)

<sup>7</sup>See [https://www.osc.gov.on.ca/documents/en/Marketplaces/com\\_20141208\\_itg-canada-corp.pdf](https://www.osc.gov.on.ca/documents/en/Marketplaces/com_20141208_itg-canada-corp.pdf)

<sup>8</sup>For example, IEX reports a “first wave fill rates” for its SOR of over 98%. See <https://iextrading.com/stats/>.

Figure 3.1: **Schematic Diagram of SOR Interaction with Canadian Equity Markets**



As a consequence, the optimal trading strategy, from a large liquidity demander's perspective, may be to send all orders to Alpha when the desired quantity can be filled there alone, but otherwise submit to all venues simultaneously and expect fleeting liquidity on Alpha. Importantly, such concerns are much less relevant for smaller orders that do not need to access multiple venues. Malinova and Park (2017) report that multi-venue trades have double the price impact of single venue trades, even after controlling for trader type and trade size. More broadly, O'Hara (2015) suggests that large traders are informed traders, at least to the extent that they know the likely market impact of their own trading intentions. Therefore, the speed bump implementation on new Alpha could be an effective mechanism to segment single-venue order flow that has relatively lower average adverse selection.

### 3.2.2 Trading Fee Change and Minimum Size for Delay-Exempt Orders

Alpha also adopted an inverted maker-taker trading fee model when it was relaunched on 21 September 2015, which charges a fee to the liquidity supplying trader and pays a rebate to the liquidity demanding trader. The exchange earns the net difference between the maker fee and the taker rebate.

Between 1 June 2015 and its relaunch, Alpha had a conventional maker-taker trading fee structure, with a maker rebate of 0.14 cents per share and a taker fee of 0.18 cents per share. After its relaunch, Alpha adopted an introductory pricing scheme until 30 November 2015, with a maker fee of 0.10 cents per share and a taker rebate of 0.10 cents per share, for no net fee margin. After 1 December 2015, the maker fee is 0.16 cents per share for speed bump exempt Post Only orders and 0.14 cents per share for non-Post Only orders, while the taker rebate remains at 0.10 cents per share. To address concerns that the concurrent trading fee changes may confound attribution of the observed effects to the speed bump implementation, we also compare liquidity metrics on new Alpha with CX2. This venue has almost identical inverted maker-taker pricing, set at a maker fee of 0.14 cents per share and a taker rebate of 0.10 cents per share.

Additionally, imposing additional fees or obligations on speed bump exempt limit orders is not necessarily novel or unique. An asymmetric speed bump is an impediment on traders who submit marketable orders because it increases search costs, which lowers the intensity of marketable order submission in the maker-taker liquidity cycles of Foucault et al. (2013). In this model, inverted maker-taker fees can increase the monitoring intensity of traders who submit marketable orders on Alpha, to balance the arrival rate with liquidity suppliers, who have higher monitoring intensity due to systematic speed advantages. Since fragmented trading venues do not have intermarket time priority, liquidity demanders have an incentive to first route marketable orders to venues with the lowest fee or highest rebate (Battalio et al., 2016). This can offset the speed bump's hindrance, particularly for brokers that do not pass on trading fees or rebates to their clients. As Brolley and Malinova (2013) note, such a flat fee structure is common for retail brokers in Canada.



Asymmetric speed bumps that currently operate on IEX and NYSE American in the United States do not delay exchange repricing of non-displayed orders, which incur higher trading fees than lit trades arising from orders that are delayed.<sup>9</sup> Additionally, in the presence of binding tick constraints, price improving dark trades and inverted trading fee structures both offer partial-tick price differentiation to enable a finer pricing grid.

Post Only orders on Alpha have a stock-specific minimum size requirement that is updated monthly and varies between 500 and 5,000 shares for stocks in our sample, most of which require 500 or 700 shares.<sup>10</sup> Heightened requirements on liquidity providers was also included in the Chicago Stock Exchange's asymmetric speed bump proposal, which would exempt liquidity providers who commit to elevated quoting obligations from the delay. These additional requirements to encourage increased liquidity are designed to attract liquidity demanders and alter equilibrium liquidity provision outcomes.

### 3.2.3 Removal from Order Protection Rule

Alpha's removal from being a protected marketplace under the order protection rule was a condition for regulatory approval of its relaunch, because the randomized delay could substantially reduce the certainty of marketable orders accessing displayed liquidity.<sup>11</sup> Removing a venue's protected status means that brokers can choose to avoid sending orders to that venue. However, brokers in Canada are obliged to strive for best execution: if brokers were

---

<sup>9</sup>IEX's trading fee schedule is available at <https://iextrading.com/trading/fees/>. NYSE American's trading fee schedule is available at [https://www.nyse.com/publicdocs/nyse/markets/nyse-american/NYSE\\_America\\_Equities\\_Price\\_List.pdf](https://www.nyse.com/publicdocs/nyse/markets/nyse-american/NYSE_America_Equities_Price_List.pdf).

<sup>10</sup>The full list of minimum Post Only order sizes on Alpha by stock is updated monthly and available at <http://api.tmxmoney.com/en/research/minpo.csv>.

<sup>11</sup>See [https://www.osc.gov.on.ca/en/Marketplaces\\_alpha-exchange\\_20150421\\_noa-proposed-changes.htm](https://www.osc.gov.on.ca/en/Marketplaces_alpha-exchange_20150421_noa-proposed-changes.htm)

to ignore Alpha's quotes when it has visibly better prices, they would need to explain why this was beneficial for overall execution quality of client orders. Thus, regardless of whether the venue is protected, it is in both the clients' and brokers' interest to not categorically ignore the venue, particularly when the venue is offering competitive quotes. Since other venues remain protected under the order protection rule, new Alpha must still quote at the best price to execute marketable orders.

To further explore this issue, our empirical analysis shows that the quotes on new Alpha rarely stay in the presence of trades that deplete a price level of liquidity on other trading venues. Therefore, regardless of whether a SOR spray sought to access liquidity on Alpha, it fades regardless, making its inclusion or exclusion from the order protection rule something of a moot point, from the perspective of multi-venue liquidity demanders. Furthermore, it has been common practice to remove order protection for venues or order types that adopt speed bumps. Aequitas NEO, the other Canadian equities trading venue with a speed bump, is also an unprotected marketplace. In the United States, dark orders on IEX and NYSE American that derive a first order benefit from their speed bumps are not protected under trade-through prohibitions, which only protect displayed orders.

### **3.2.4 Decommissioning of TMX Select**

Concurrent with Alpha's relaunch, TMX Group decommissioned TMX Select, which was its existing trading venue with an inverted maker-taker fee structure. Prior to Alpha's relaunch, CX2, TMX Select and Omega operated inverted trading fee structures. Following Alpha's relaunch, CX2, Omega and new Alpha operate inverted trading fee structures. We do not expect any sig-

nificant market-wide impact on order flow segmentation from the change in Alpha’s trading fee model alone, because the number of inverted trading fee venues remained the same. Since trading venues with inverted fee structures tend to have lower adverse selection costs (Battalio et al., 2016), decommissioning TMX Select results in its order flow being redistributed amongst other trading venues, potentially reducing aggregate adverse selection. Anderson et al. (2018) report that TMX Select previously had a high proportion of retail trades, which was redistributed among other venues after Alpha’s relaunch. Therefore, any adverse order flow segmentation impacts arising from Alpha’s relaunch would need to overcome this effect on the consolidated Canadian equities market.

### 3.3 Data and Liquidity Metrics

#### 3.3.1 Data

The data for this study is sourced from Thomson Reuters Tick History (TRTH), which is supplied by the Securities Industry Research Centre of Asia Pacific (SIRCA). TRTH data for each trading venue includes all order book updates and trades. Bid and ask prices and sizes are reported each time there is a change in the state of the order book. The trade record includes fields for price, volume, qualifiers, and buyer and seller broker IDs. Data is stamped to millisecond granularity. We exclude off-market trade reports,<sup>12</sup> odd lot trades and dark trades, as they do not interact with the displayed limit order book. We analyze data for the four major Canadian trading venues, TSX,

---

<sup>12</sup>In addition to trades flagged with a qualifier to indicate that they were pre-arranged intentional crosses, we also exclude trades with a value above \$1 million. Trade qualifiers in TRTH data may be incomplete, and we are aware of trades exceeding \$100 million in the TRTH data without a trade reporting qualifier.

Alpha, Chi-X and CX2, which together account for more than 90% of all lit on-market trading.

TRTH also includes data for the TMX Select, Omega and Pure Trading markets, which we do not include in our sample. TMX Select and Omega both use a legacy data feed during the sample period, with timestamp inaccuracies that frequently exceed 200-300 milliseconds, making it impossible to precisely compare their trades and quotes with concurrent activity on other venues. Each venue accounts for less than three percent of trading activity. Weighing data accuracy against sample completeness, these two venues are excluded. Pure Trading is excluded as it has a market share of less than one percent.

Our observation period spans 32 weeks centered on TSX Alpha's relaunch, commencing on 1 June 2015, when trading fee changes were implemented on TSX and Alpha, and ending on 8 January 2016. We exclude the NYSE trading holidays and partial holidays of 3 July 2015, 26 November 2015 and 27 November 2015, on which equities trading activity in Canada is low, a U.S. equities market "flash crash" on 24 August 2015, Christmas Eve on 24 December 2015, as well as the quarterly S&P index rebalancing days. Although several venues operate extended trading hours, our analysis is restricted to the TSX listing market's continuous trading hours, being 9.30am to 4.00pm. Our sample consists of stocks that remain members of the S&P/TSX Composite Index across our entire sample period. We also exclude stocks FFH, CSU, VRX, CP and CCLb because they trade at above \$200, compared with less than \$130 across all other stocks in the sample. The final sample includes 219 stocks across 145 days.

### 3.3.2 Standard Liquidity Metrics

Following the empirical literature, we construct a variety of standard liquidity metrics to analyze the impact of the relaunched TSX Alpha exchange with speed bump. These metrics are aggregated at the stock-day level. We create one dataset that records the national best bid and offer (NBBO) prices and volumes and another dataset that contains the trades on each venue. On each quote update during continuous trading, the NBBO is constructed as the set of highest bid price and lowest ask price, and the aggregate volumes available at those prices, across all trading venues. Since the trades and quote updates are provided in the same file per venue, trade initiation direction can be assigned with near certainty by comparing trades with the prevailing order book. A detailed example is provided in Section 3.7.

Standard liquidity metrics can measure either quoted or traded liquidity. Quoted liquidity metrics measure the displayed liquidity that is available to market participants at each point in time across all venues and are time-weighted throughout each trading day by the duration that each quote update is prevailing, until the next quote update arrives. The quoted liquidity metrics we calculate are quoted spreads and quoted depths.

The NBBO quoted spread is calculated in absolute dollar terms as the difference between the prevailing national best bid (NBB) and national best offer (NBO) prices. We also calculate quoted spreads in relative basis point terms, by dividing the absolute quoted spread by the NBBO midpoint price. Narrower quoted spreads indicate lower implicit transaction costs on average throughout the day.

$$QuotedSpread = NBOPrice - NBBPrice \quad (3.1)$$

The NBBO quoted depth is calculated as the total displayed value quoted at the NBBO prices, across all venues. Higher quoted depths indicate that larger trades can occur at the available NBB and NBO prices.

$$QuotedDepth = NBBPrice * NBBVolume + NBOPrice * NBOVolume \quad (3.2)$$

Traded liquidity metrics compare trade prices with the reference NBBO midpoint price and are aggregated per stock-day by turnover-weighting, to proxy for the implicit cost of immediacy that was actually incurred by liquidity demanders. The traded liquidity metrics we calculate are effective spreads, realized spreads and adverse selection. Trade initiation direction is assigned by comparing trades with the prevailing order book on the same venue, with trades equal to or higher than the prevailing offer price being buyer-initiated and trades equal to or lower than the prevailing bid price being seller-initiated. The inequality is required to account for large marketable orders that can simultaneously consume liquidity at multiple price levels.

Effective spreads are calculated as twice the difference between the trade price and the prevailing NBBO midpoint price at the time of the trade,  $NBBOMP_t$ , to reflect the implicit transaction cost for small round-trip trades at the best quotes. Trade direction is +1 for buyer-initiated trades and -1 for seller-initiated trades.

$$EffectiveSpread = 2 * Direction * (TradePrice - NBBOMP_t) \quad (3.3)$$

Adverse selection is computed as twice the movement of the NBBO mid-

point price in the direction of the trade, from immediately before the trade to a specified reference time after the trade, which reflects the permanent price impact of each trade. Following Conrad et al. (2015), adverse selection is calculated at horizons of one, five, ten and twenty seconds after each trade. For brevity, we report this metric after one second as our primary result.

$$AdverseSelection = 2 * Direction * (NBBOMP_{t+1sec} - NBBOMP_t) \quad (3.4)$$

Realized spreads are calculated as twice the signed difference between the trade price and the NBBO midpoint price prevailing at a specified reference time after the trade, as a proxy for the profits earned by liquidity providers. Similar to the adverse selection metric, realized spreads are calculated at horizons of one, five, ten and twenty seconds after each trade but we report the one second metric as our primary result.

$$\begin{aligned} RealizedSpread &= EffectiveSpread - AdverseSelection \\ &= 2 * Direction * (TradePrice - NBBOMP_{t+1sec}) \end{aligned} \quad (3.5)$$

The above metric specifications for effective spreads, realized spreads and adverse selection are in absolute dollar terms. These metrics are also transformed into relative basis point terms, by dividing the metric in dollars by the NBBO midpoint price.

### 3.3.3 Novel Liquidity Metrics

Motivated by the importance of market linkages highlighted in O'Hara (2015), we construct several novel metrics using trade and quote data from

TRTH, which is commonly available to empirical researchers, to estimate the ability of liquidity demanders to access consolidated liquidity displayed across multiple venues. This section provides a brief overview of metric construction, while a more detailed explanation with examples is provided in Section 3.7. By examining the distribution of locked and crossed market durations, which securities regulations in Canada prohibit (IIROC, 2011), we estimate that the timestamps across the venues in our TRTH data sample to be precise to within 30 milliseconds at a 95% confidence level. This benchmarked interval is similar to regulatory requirements that specify a 50-millisecond clock synchronization threshold for trading venues (IIROC, 2016).

For each stock and separately for buyer- and seller-initiated trades, high frequency trade strings are constructed by grouping together all trades across venues that are recorded within 30 milliseconds of each other. Whilst timestamps for any individual trade may exhibit inaccuracies of up to 30 milliseconds in our data, trades that occur over very short time intervals are likely to be related. For example, they may originate from an individual broker's SOR, consist of different traders that are responding to the same informational event, or include some trades that are in response to the order flow information of earlier trades in that string. For trades that occurred at the best price within each string, which is generally the prevailing NBB (NBO) price at the start of the string for seller (buyer) initiated trade strings, we record the trade price, trade volume, start time and end time. We then take snapshots of the limit order books across each venue immediately before the first trade and 30 milliseconds after the last trade in the string. Only trades occurring at the best price within each string are analyzed, so we can attribute traded volume to displayed market depth.

The construction of high frequency trade strings also allows for effective



spreads, realized spreads and adverse selection to be calculated at the trade string level rather than at the individual trade level, which is the standard approach in the existing literature. This aggregated analysis may reveal further insights into transaction costs for different liquidity consuming strategies or under different market conditions throughout a trading day. We leave this potential extension to future empirical research.

Buyer (seller) initiated trades are classified as “depleting” if they originate from a trade string in which the national best offer (bid) price at the end of the string is higher (lower) than the best price traded during the string. This classification is not strictly a proxy for trade size, since an order smaller than pre-trade NBB or NBO depth can also displace an entire price level if residual orders are canceled during or immediately following the trades. Foucault and Moinas (2018) suggest that trading that removes all liquidity in a stock at a price level on one venue is relevant information for pricing that stock on other venues. Extending this notion across all venues, depleting trade strings have high information content. Our definition of depleting trade strings as being informed is analogous to calculating the traditional adverse selection metric at a virtually instantaneous horizon of 30 milliseconds, rather than a few minutes (Hendershott et al., 2011) or seconds (Conrad et al., 2015) after each trade.

Similar to Malinova and Park (2017) and van Kervel (2015), we also classify trade strings as either multi-venue or single-venue, depending on the number of venues that they include trades from. For brevity, we refer to multi-venue trade strings as “SOR”, because they may have originated from a single broker’s SOR spray that sought to access the consolidated liquidity across multiple venues.

For depleting buyer (seller) initiated trade strings, the NBBO quote fade metric is calculated as the proportion of starting liquidity displayed at the national best offer (bid) price that did not result in trades. As there is no

residual liquidity at the NBBO after a depleting trade string, the starting liquidity can either be consumed by liquidity demanders or withdrawn by liquidity suppliers. A lower bound of zero is placed on the quote fade metric per trade string, which caps trading volume at the starting depth. Trade volume could exceed starting depth if trades execute against partially hidden iceberg orders, trades occur within a millisecond after new orders are submitted, or liquidity is rapidly replenished after trades.

$$QuoteFade = 1 - \frac{TradeValue}{Max(StartingLiquidity, TradeValue)} \quad (3.6)$$

A key difference between our quote fade metric and the similar cross-market order cancellation metrics developed by Malinova and Park (2017) and van Kervel (2015) is that we only examine trades that occur at the best price and those that deplete the liquidity available at that price level. These two conditions specify the necessary parameters for latency sensitive competition between liquidity suppliers and demanders. If liquidity depletion at a price level is not imminent, liquidity demanders do not need to race to consume remaining orders at that price. Similarly, liquidity suppliers do not need to race to cancel their standing orders before prices move. Additionally, we only examine order activity at the same price, to quantify the economic cost of fleeting liquidity.

### 3.3.4 Summary Statistics

Table 3.1 presents summary statistics for a number of liquidity metrics and control variables at the stock-day level. We report the mean, median and standard deviation for each variable separately for the sixteen weeks before and after Alpha’s relaunch on 21 September 2015. In both this table and the

formal regression analysis, all observations are winsorized at the 1% level per stock and per day. Following its relaunch, 80% of order entries and 75% of order cancellations on Alpha are sufficiently large to meet the minimum Post Only order size requirements to avoid the speed bump delay, which leads to its median quote fade increasing immensely to 70% of displayed liquidity. The proportion of trades on Alpha that are part of depleting trade strings decreases from 53% to 31%. Conversely on CX2, which has a virtually identical trading fee structure to new Alpha, quote fade decreases substantially from 21% to 16% and the proportion of its trades participating in depleting trade strings increases from 31% to 37%. Relatively low levels of quote fade across all venues in the pre-event period is consistent with the supposition of O'Hara and Ye (2011) that trade-through prohibition combined with SOR in fragmented markets virtually replicates the network benefits of consolidated trading. The proportion of trades that deplete a price level increases in the consolidated market after Alpha's relaunch.

On Alpha, average effective spreads increase from 2.30 to 2.84 cents, while adverse selection decreases from 2.05 cents to 1.44 cents, making similar contributions to its increase in realized spreads from 0.25 cents to 1.38 cents. Consolidated traded liquidity metrics exhibit smaller changes. Average quoted spreads increase from 2.82 to 2.86 cents. Applying an exponential transformation to logarithmic variables, average share prices declined from \$18.93 to \$17.12, while average daily turnover per stock increased from \$8.4 million to \$9.0 million and quoted depths increased from \$60,962 to \$68,105. Market-wide volatility increased over the period, observed from the increase in both realized 1-minute intraday return volatility and the S&P/TSX 60 VIX index.

Table 3.2 presents additional summary statistics on the frequency, size and duration of trade strings constructed over the pre-event period, to provide fur-

Table 3.1: **Summary Statistics**

This table reports stock-day descriptive statistics across TSX Composite Index component securities 6 months either side of Alpha's relaunch on 21 September 2015. Alpha Fast Order Enter and Alpha Fast Order Cancel are the proportion of order submission and cancellations, respectively, that exceeded the specified post only order size requirement for that month. Quote Fade is the proportion of total displayed starting liquidity on a venue at the NBB or NBO that did not result in trades. Deplete Best is the proportion of trades that were followed by instantaneous adverse NBBO price movements. Quoted spread and Quoted Depth are time-weighted and consolidated at the NBBO prices across Alpha, Chi-X, CX-2 and TSX. Effective Spread is calculated using the prevailing NBBO midpoint. Realized spread is calculated against the NBBO midpoint one second after the trade. Adverse selection is calculated as the movement in the NBBO midpoint between the time of the trade and 1 second after the trade, in the direction of the trade. These three traded liquidity metrics are calculated for both the consolidated market and only trades on Alpha. Price is the natural logarithm of the time-weighted NBBO midpoint. Turnover is the natural logarithm of the total value of on-market trades. Volatility is the realized standard deviation of one-minute NBBO midpoint returns. VIXC is the daily opening level of the S&P/TSX 60 VIX Index, which measures the market's 30-day options implied volatility.

	Pre-Event Period			Post-Event Period		
	Mean	Median	Std. Dev.	Mean	Median	Std. Dev.
<b>A: Quote Fade</b>						
Alpha Fast Order Enter (%)	0.00	0.00	0.00	0.80	0.86	0.19
Alpha Fast Order Cancel (%)	0.00	0.00	0.00	0.75	0.81	0.19
Alpha Quote Fade (%)	14.10	11.09	11.43	63.52	70.04	24.09
Chi-X Quote Fade (%)	21.38	18.81	12.73	21.77	20.16	11.47
CX2 Quote Fade (%)	21.37	19.99	13.93	15.61	14.30	9.79
TSX Quote Fade (%)	9.05	6.54	7.80	8.96	6.55	7.75
<b>B: NBBO Depth Depletion</b>						
Total Deplete Best (%)	57.37	58.47	12.37	59.00	59.94	12.11
Alpha Deplete Best (%)	52.99	53.58	14.89	31.22	29.06	17.41
Chi-X Deplete Best (%)	57.71	58.26	13.16	60.85	61.37	13.50
CX2 Deplete Best (%)	30.53	28.89	15.76	37.32	37.05	16.68
TSX Deplete Best (%)	60.94	62.35	12.70	63.56	65.06	12.71
<b>C: Liquidity Metrics</b>						
Quoted Spread (cents)	2.82	1.67	2.80	2.86	1.73	2.86
Quoted Depth (log \$)	11.02	10.87	0.77	11.13	10.97	0.74
Total Effective Spread (cents)	2.34	1.54	2.02	2.36	1.56	2.02
Alpha Effective Spread (cents)	2.30	1.44	2.08	2.84	1.81	2.73
Total Realized Spread (cents)	0.00	-0.12	0.74	-0.04	-0.15	0.68
Alpha Realized Spread (cents)	0.25	0.06	1.00	1.38	0.84	1.89
Total Adverse selection (cents)	2.34	1.82	1.70	2.39	1.83	1.75
Alpha Adverse selection (cents)	2.05	1.56	1.60	1.44	0.99	1.51
<b>D: Control Variables</b>						
Price (log \$)	2.94	3.11	0.97	2.84	2.99	1.01
Turnover (log \$)	15.95	15.80	1.19	16.02	15.85	1.22
Volatility (basis points)	10.22	8.29	6.10	11.37	9.72	6.08
VIXC (%)	16.55	14.46	5.35	20.33	20.00	2.42

ther information on our novel trade aggregation methodology. Trade strings are categorized into four groups according to the number of venues on which trades occurred within the string. Three quarters of trade strings only interact with a single venue, although they account for only 44% of total turnover, with a median volume of only 100 shares. Trade strings with trades across all four venues account for only 1.6% of total trade strings by count, but 8.9% of total turnover due to their larger average size, with a median volume of 2,100 shares. Consistent with multi-venue trade strings imposing higher instantaneous adverse selection costs than single-venue trade strings, the proportion of trade strings that are depleting increases monotonically in the number of venues accessed, from 30.4% for single-venue to 71.6% for four-venue trade strings. In addition to being larger, there also tends to be higher displayed liquidity when multi-venue trade strings start.

While our trade string construction methodology theoretically enables spanning very long periods of time if subsequent trades are followed by another within 30 milliseconds, in practice the durations are very brief, with median string lengths of 0, 9, 13 and 17 milliseconds for those with trades on one, two, three and four venues respectively. These lengths are similar to the trade string construction methodology in Malinova and Park (2017), which groups trades originating from a single trader across venues within 5 milliseconds. Our intervals are significantly shorter than the 100 millisecond snapshots taken in van Kervel (2015) and are closer to the latencies of the fastest traders documented in Hasbrouck and Saar (2013). There is usually a time gap of many seconds between the end of one trade string and the start of the next, which ensures methodological robustness because the signal of closely related trades in rapid succession far exceeds the noise of overall intraday market activity. Interestingly, single-venue trade strings occur sooner after the end of the previous

trade string but have a longer gap until the start of the next trade string.

Following Baldauf and Mollner (2018), Table 3.3 reports the stock-day Pearson correlation coefficients on Alpha between our novel order flow segmentation metrics and established measures of informed trading, to confirm that these variables are highly correlated in the cross-section and time-series. The correlation between quote fade on Alpha and presence of orders sufficiently large to meet the speed bump exempt Post Only order size requirement exceeds 0.80. The correlation between meeting the Post Only order size requirement and participating in multi-venue trades is -0.61. The correlation between quote fade and participation in depleting trade strings is -0.63. Participation in depleting trade strings has a correlation of 0.58 with nominal adverse selection costs at a one second horizon.

### **3.4 Methodology, Results and Discussion**

We adopt the following structure to investigate the effects of TSX Alpha’s implementation of an asymmetric randomized systematic order processing delay and other concurrent changes on liquidity in the Canadian equities market. We start by measuring the first order impacts of the speed bump via changes in the ability of liquidity providers on Alpha to fade their quotes upon observing trades on other venues. We then examine whether the ability to segment order flow on Alpha impacts the relative composition and information content of trading volumes across each venue. Next, we examine the second order impacts of Alpha’s market design changes on adverse selection, liquidity provider profits and transaction costs on each venue. We also directly compare traded liquidity metrics for different types of liquidity providers on Alpha with CX2, which has a virtually identical inverted maker-taker trading fee structure. Fi-

Table 3.2: **Summary Statistics for Trade Strings**

This table presents descriptive statistics for trade strings across TSX Composite Index component securities over the 16 weeks prior to Alpha's relaunch on 21 September 2015. Statistics are reported by the number of venues on which trades were executed in each trade string. Trade strings are constructed by grouping together all trades in the same stock and same direction that are separated by less than 30 milliseconds. Frequency is the number of trade strings observed. Means and medians are presented for traded value and traded volume within each trade string. Starting depth is the total depth at the NBB or NBO at the time of the first trade in the trade string. Percentage of Count and Percentage of Traded Value are the proportion of trade strings within each venue count group, by frequency and turnover value. Depleted trade strings are those that displace all available liquidity at the NBB or NBO. String length is the time between the first and last trade within a trade string. Space before is the time between the last trade in the previous trade string and the first trade in the current trade string. Space after is the time between the last trade in the current trade string and the first trade in the next trade string. The 10th, 25th, 50th, 75th and 90th percentiles of the distributions for these time variables is presented in seconds.

	Number of Venues Accessed by Trade String			
	1	2	3	4
Frequency (#)	19,578,745	4,333,637	1,794,495	418,139
Mean Traded Value (\$)	6,527	18,068	32,551	61,936
Median Traded Value (\$)	3,396	10,922	20,880	39,930
Mean Traded Volume (#)	276	816	2,034	5,121
Median Traded Volume (#)	100	400	1,000	2,100
Mean Starting Depth (\$)	46,851	39,734	51,389	82,110
Median Starting Depth (\$)	17,888	20,796	31,407	53,405
Pct of Count (%)	74.94%	16.59%	6.87%	1.60%
Pct of Traded Value (%)	44.01%	26.96%	20.11%	8.92%
Pct Depleted by Count (%)	30.41%	55.00%	64.02%	71.62%
Pct Depleted by Traded Value (%)	46.97%	62.95%	69.56%	75.86%
P10 String Length (sec)	0.000	0.000	0.004	0.008
P25 String Length (sec)	0.000	0.003	0.009	0.011
P50 String Length (sec)	0.000	0.009	0.013	0.017
P75 String Length (sec)	0.000	0.014	0.020	0.024
P90 String Length (sec)	0.000	0.021	0.030	0.036
P10 Duration Before (sec)	0.179	0.271	0.455	0.970
P25 Duration Before (sec)	1.379	1.972	2.873	4.764
P50 Duration Before (sec)	6.909	9.170	12.000	16.098
P75 Duration Before (sec)	26.112	32.708	40.098	49.560
P90 Duration Before (sec)	75.313	90.392	106.718	125.553
P10 Duration After (sec)	0.250	0.137	0.109	0.094
P25 Duration After (sec)	1.750	1.189	0.969	0.651
P50 Duration After (sec)	8.032	6.814	6.342	5.702
P75 Duration After (sec)	29.183	26.848	25.927	25.015
P90 Duration After (sec)	82.415	77.211	75.848	75.045

**Table 3.3: Correlation Table for Novel and Established Measures of Adverse Selection**

This table reports Pearson correlation coefficients between the novel metrics of adverse selection and established metrics from the literature, for stock-day observations on Alpha over the period spanning six months either side of Alpha’s relaunch on 21 September 2015. Fast Order Enter and Fast Order Cancel are the proportion of order submission and cancellations, respectively, that exceeded the specified post only order size requirement for that month. Quote Fade is the proportion of total displayed starting liquidity at the NBB or NBO that did not result in trades. Deplete Best is the proportion of trades that were immediately followed by adverse NBBO price movements. SOR is the proportion of trades that accessed multiple venues within a 30-millisecond time horizon. Deplete Best and SOR is the proportion of trades that accessed multiple venues within a 30-millisecond time horizon and were immediately followed by adverse NBBO price movements. Adverse selection is calculated as the movement in the NBBO midpoint between the time of the trade and 1 second after the trade, in the direction of the trade. Realized spread is calculated by comparing the trade price against the NBBO midpoint one second later.

	1	2	3	4	5	6	7	8	9
1. Post-Relaunch (binary)	1.00								
2. Fast Order Enter (%)	0.95	1.00							
3. Fast Order Cancel (%)	0.94	0.98	1.00						
4. Quote Fade (%)	0.80	0.84	0.81	1.00					
5. Deplete Best (%)	-0.56	-0.54	-0.51	-0.63	1.00				
6. SOR (%)	-0.58	-0.61	-0.61	-0.51	0.39	1.00			
7. Deplete Best and SOR (%)	-0.58	-0.57	-0.55	-0.58	0.87	0.66	1.00		
8. Adverse selection (cents)	-0.19	-0.18	-0.15	-0.34	0.58	-0.07	0.32	1.00	
9. Realized Spread (cents)	0.35	0.36	0.39	0.19	-0.16	-0.45	-0.33	0.13	1.00



nally, we examine the impacts on consolidated liquidity across the Canadian equities market.

### 3.4.1 Speed Bump Mechanics and Fleeting Liquidity


Alpha’s asymmetric randomized speed bump of 1 to 3 milliseconds provides an opportunity for fast liquidity suppliers to cancel their standing limit orders ahead of incoming marketable orders, after observing an external signal, such as large trades on other venues. Figure 3.2 presents daily average quote fade across all stocks in the sample by trading venue. A significant increase in quote fade is observed on Alpha immediately after its relaunch, while quote fade decreases slightly on CX2. Higher quote fade suggests that liquidity providers are better able to avoid depleting trade strings that interact with many venues, whilst interacting with order flow sent to only one venue that is expected to have lower average adverse selection. While summary statistics and formal regression analysis utilize a symmetric interval 16 weeks before and after Alpha’s relaunch, from 1 June 2015 to 8 January 2016, all time-series figures extend to the end of January 2016 to display full monthly intervals on the horizontal axis.

We more formally test the significance of changes in quote fade on each venue using Equation 3.7:

$$\begin{aligned}
QuoteFade_{i,d,v} = & \beta_1 Post_d + \beta_2 FastOrderEntryOrCancel_{i,d} \\
& + \beta_3 NBBODEpletion_{i,d} + \beta_4 VenueCount_{i,d} + \beta_5 DepthShare_{i,d,v} \\
& + \beta_6 Price_{i,d} + \beta_7 Turnover_{i,d,v} + \beta_8 Volatility_{i,d} \\
& + \beta_9 VIXC_d + FE_i + \epsilon_{i,d,v}
\end{aligned} \tag{3.7}$$

Figure 3.2: **Quote Fade Across Venues**

This figure presents the daily average level of quote fade across venues. We construct trade strings by joining all trades for the same stock in the same direction separated by less than 30 milliseconds. A trade string is called depleting when the entire NBB or NBO depth is displaced following the trade. Among all depleting trade strings we calculate quote fade as the proportion of starting liquidity that did not result in trades.



where  $QuoteFade_{i,d,v}$  is one minus the volume traded at the best price divided by the total starting liquidity at that price, as defined in Equation 3.6, aggregated across all trade strings for stock  $i$  on day  $d$  at venue  $v$ . The key explanatory variables of interest are  $Post_d$ , an indicator variable equal to one for observations after Alpha's relaunch and zero prior, and  $FastOrderEntryOrCancel_{i,d}$  which is, alternatively, the proportion of orders entered or canceled, on Alpha for stock  $i$  on day  $d$  that met the size requirement for speed bump exempt Post Only non-marketable orders. The set of control variables includes  $NBBODepletion_{i,d}$ , the proportion of turnover that depleted a price level of liquidity,  $VenueCount_{i,d}$ , the average number of venues that each trade string accessed,  $DepthShare_{i,d,v}$ , the NBBO depth share of the particular trading venue,  $Price_{i,d}$ , the natural logarithm of the time-weighted NBBO midpoint price,  $Turnover_{i,d,v}$ , the natural logarithm of on-market turnover on the particular venue,  $Volatility_{i,d}$ , the realized intraday volatility of one minute NBBO midpoint returns and  $VIXC_d$ , the daily opening level of the S&P/TSX 60 VIX volatility index. We explicitly control for market-wide volatility in the regression specification due to its substantial increases in the post-event period.  $FE_i$  indicates stock fixed effects that control for time-invariant differences in the dependent variable for each stock, and  $\epsilon_{i,d,v}$  is an error term. Standard errors are double-clustered by stock and date.

Consistent with Figure 3.2, Table 3.4 shows that Alpha's Quote Fade jumps by 41.78% on average after the introduction of the speed bump, while the magnitude of changes on other venues is minor in comparison. High quote fade indicates that displayed orders were canceled before they could be traded against, representing inaccessible liquidity. A large increase in quote fade on new Alpha is expected because its asymmetric randomized duration speed bump makes it impossible for an SOR to time orders to arrive simultaneously

**Table 3.4: Changes in Quote Fade on Each Venue Relative to the Pre-Relaunch Period**

This table reports coefficient estimates and t-statistics for the determinants of NBBO quote fade for each of Alpha, Chi-X, CX2 and TSX across TSX Composite Index securities, after Alpha's relaunch relative to previous levels using the specification in Equation 3.7, where the quote fade, defined in Equation 3.6, for stock  $i$  on day  $d$  at venue  $v$  is expressed as the sum of either an indicator variable for the post-relaunch period or the proportion of order entries or cancellations on Alpha that met the Post Only size requirement after relaunch, and control variables for the proportion of trades that depleted NBBO liquidity, the average number of venues per trade string, each venue's NBBO depth share, natural logarithms of price and each venue's turnover, realized intraday volatility, implied inter-day volatility, a stock specific mean and an error term. The observation period runs from 1 June 2015 to 8 January 2016. We add a \*/\*\*/\*\* to the t-statistic to indicate statistical significance at the 90%/95%/99% levels, respectively. Standard errors are double clustered by stock and date.

	Alpha	Alpha	Alpha	Chi-X	CX2	TSX
Post	41.78 (45.01)***			1.18 (3.56)***	-3.64 (-7.99)***	0.08 (0.50)
Fast Order Entry		53.68 (42.99)***				
Fast Order Cancel			56.19 (42.30)***			
Deplete	0.05 (1.95)*	0.05 (2.25)**	0.05 (2.16)**	0.06 (5.25)***	0.01 (1.09)	0.03 (3.67)***
Best Venue	-8.75 (-5.37)***	-5.83 (-3.88)***	-9.03 (-5.50)***	-6.21 (-8.63)***	-2.76 (-3.22)***	-2.74 (-6.52)***
Count						
Depth	89.65 (19.02)***	52.38 (11.82)***	54.65 (11.46)***	46.54 (14.86)***	59.59 (12.39)***	-3.49 (-4.77)***
Share						
Price	-2.98 (-1.53)	-4.72 (-2.30)**	-5.72 (-2.67)***	2.82 (2.00)**	6.04 (2.85)***	2.65 (3.43)***
Turnover	-2.48 (-4.33)***	-3.15 (-6.44)***	-3.03 (-6.15)***	-4.82 (-16.14)***	-4.03 (-10.50)***	-1.94 (-9.66)***
Volatility	0.08 (1.00)	0.12 (1.87)*	0.09 (1.36)	0.28 (5.93)***	0.16 (2.79)***	0.16 (5.67)***
VIXC	-0.07 (-1.70)*	-0.01 (-0.36)	-0.01 (-0.31)	-0.06 (-2.03)**	-0.12 (-3.06)***	-0.09 (-6.75)***
Adjusted R2	76.47%	79.06%	78.59%	12.49%	11.04%	8.34%
# Obs	31,654	31,654	31,654	31,749	31,472	31,754

on multiple venues. Quote fade decreases on CX2 by 3.64% after Alpha’s relaunch, indicating that liquidity demanders become more targeted and successful at accessing its displayed liquidity at competitive prices, following the reduction in number of accessible venues. Across all venues, quote fade tends to increase with the proportion of depleting trade strings and stock-specific volatility, while it decreases with the number of venues accessed per trade string, turnover and market-wide volatility.

To validate that the documented upsurge in quote fade on Alpha is specifically due to liquidity providers canceling limit order that are exempt from the speed bump, we conduct additional regression specifications that replace the post-event indicator variable with the percentage of order entries and cancellations for each stock-day on Alpha in the post-event period that met the minimum order size requirement for undelayed Post Only orders. This order size requirement is stock-specific and time-varying, and is determined by TMX Group at the start of each month based on each stock’s price and liquidity. It varies between 500 and 5,000 shares for stocks in our sample. The coefficient estimates for the impact on Alpha’s quote fade of order entries and cancellations that meet the size requirement for speed bump delay exemption are 53.68% and 56.19% respectively. Multiplying these coefficients by the average of 79.60% and 74.61% of orders per stock-day that meet the undelayed order size requirements on Alpha after relaunch from the summary statistics in Table 3.1 yields quote fade changes of 42.73% and 41.92%, respectively. These are remarkably similar to the coefficient estimate for the post-event indicator variable in the base specification and we therefore attribute the entire increase in quote fade on Alpha to usage of speed bump exempt Post Only orders.

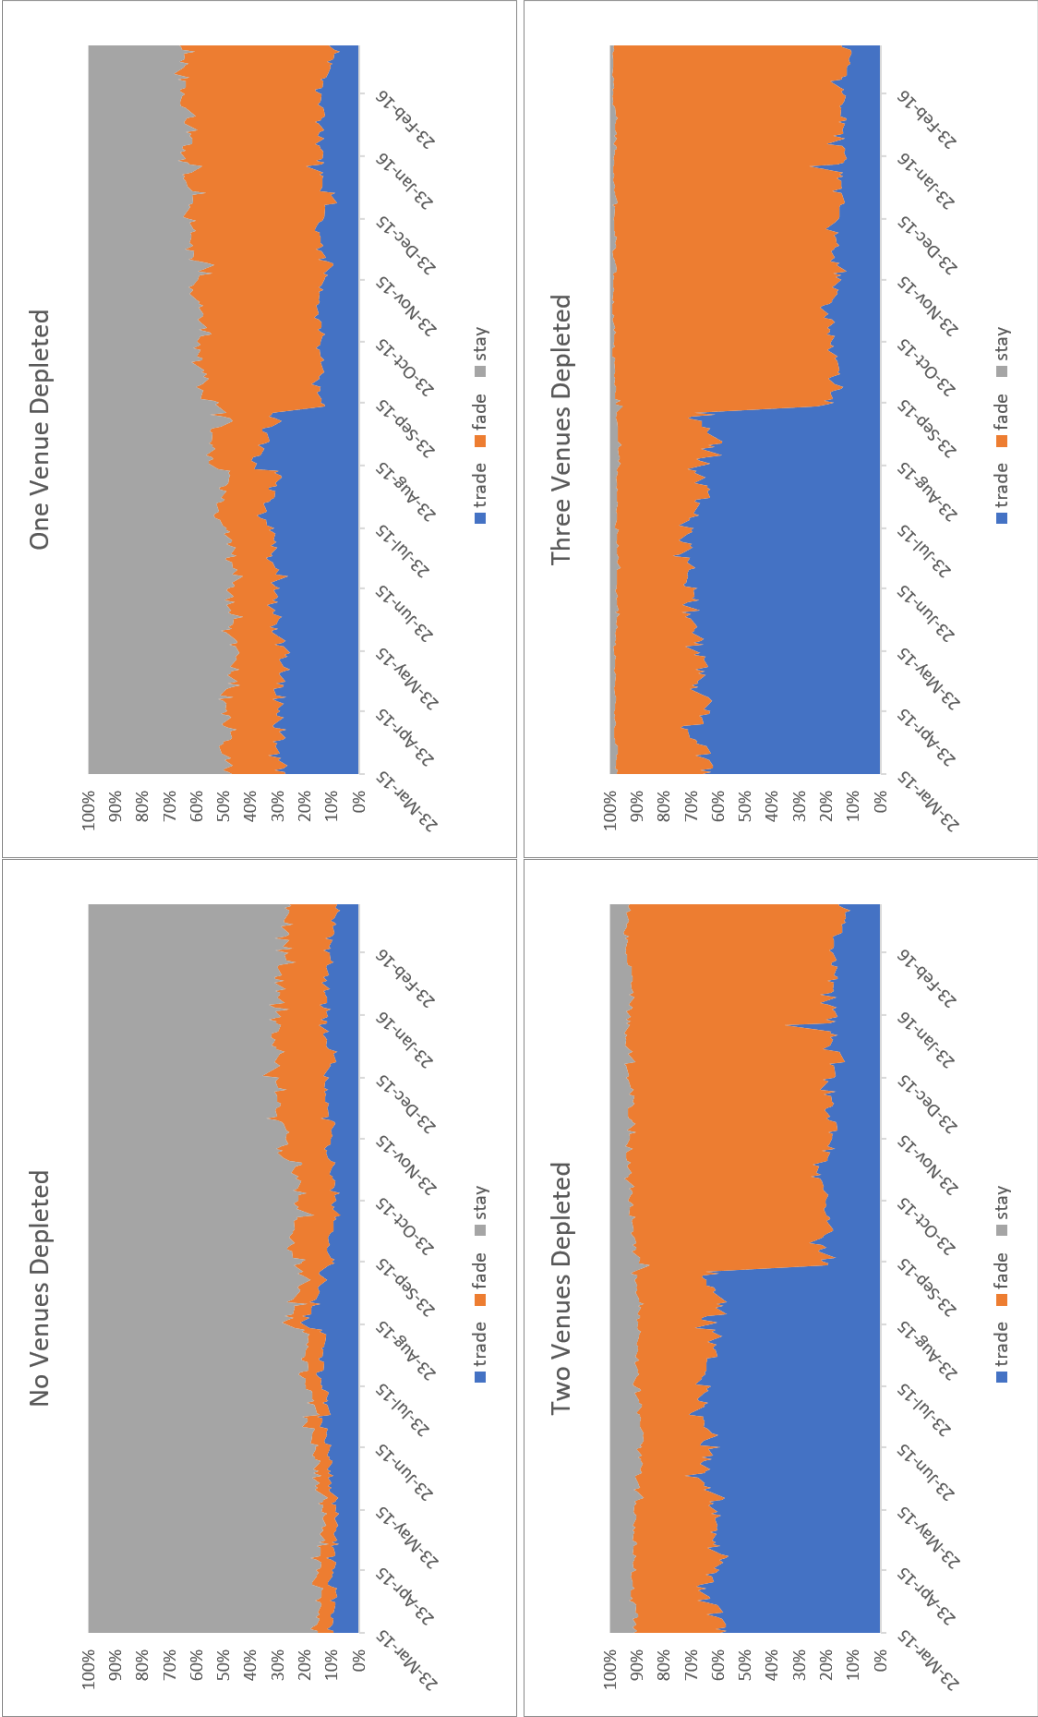
A potential limitation of the quote fade metric is that it is only defined for trade strings where all liquidity at the NBB or NBO was depleted across

all venues by the end of the trade string. Following Alpha's removal from the Order Protection Rule concurrently with the speed bump implementation, a plausible outcome is that liquidity demanders seeking to trade large volumes with a multi-venue SOR spray might exclude Alpha from their routing tables completely. This would reduce the proportion of depleting trade strings, for which quote fade is defined, and give liquidity suppliers on Alpha plenty of time to cancel their stale orders. While the summary statistics in Table 3.1 report that the proportion of turnover in depleting trade strings actually increased after Alpha's relaunch, from an average of 57% to 59%, we additionally investigate the eventual outcome of displayed liquidity on Alpha around depleting trades on other venues, to examine whether it also traded, faded before it could be traded, or stayed in the order book.

Figure 3.3 presents four area charts that display the eventual outcome of liquidity resting on Alpha at the NBB (NBO) conditional on depleting seller (buyer) initiated trades on other venues, immediately after the end of the trade string. In the post-relaunch period, there is a large increase in the proportion of starting liquidity on Alpha that fades if trade strings deplete the NBBO liquidity on two or more other venues. A very limited proportion of liquidity stays on Alpha's order book. A similar, though smaller, increase in order cancellation is observed when one of the other three trading venues is depleted of liquidity at the NBBO. Thus, the fleeting nature of orders on Alpha when confronted by trades that deplete liquidity across multiple trading venues is independent of whether liquidity demanders seeking to use multi-venue SOR sprays for large orders continue to route their marketable orders there.

Figure 3.3: **Percentage of Displayed Liquidity on Alpha that Stayed, Traded and Faded**

This quadrant of figures presents the proportion of displayed liquidity on Alpha at the national best offer (bid) price at the start of buyer (seller) initiated trade strings that resulted in trades, stayed in the order book, or were faded from the order book, conditional on the number of venues accessed during each trade string. Trade strings are constructed by grouping all trades separated by less than 30 milliseconds. The starting liquidity snapshot is taken immediately prior to the first trade in the string and the ending liquidity snapshot is taken 30 milliseconds after the last trade in the string.



### 3.4.2 Multi-Venue Trades, NBBO Liquidity Depletion and Order Flow Segmentation

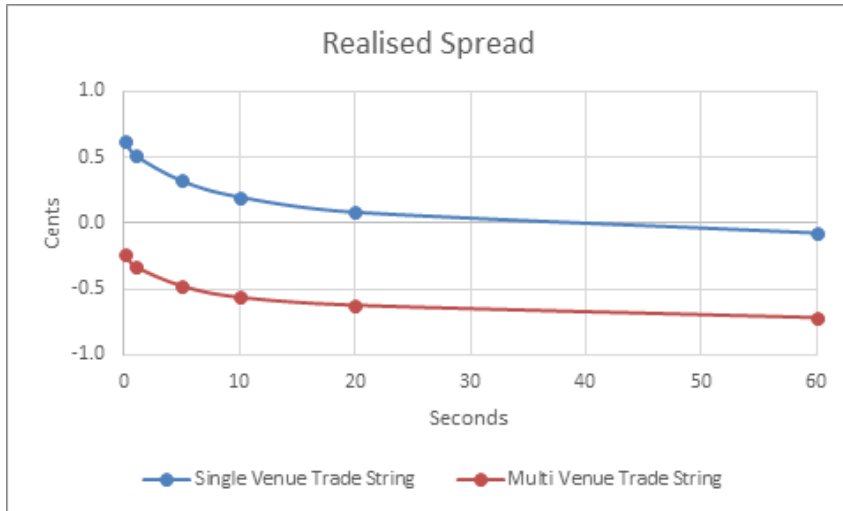
We next examine the mechanisms via which fleeting liquidity can lead to order flow segmentation. Using the classification schemes derived in Section 3.3, we first split trade strings into single venue trade strings and multi-venue trade strings, with the latter potentially originating from the SOR of a large, “fast” trader (Biais et al., 2015, van Kervel, 2015). Figure 3.4 shows that this distinction is important, by contrasting realized spreads for these two types of trade strings at several time horizons from 100 milliseconds to a minute after the end of each trade string. For both types of trade strings, the vast majority of adverse selection occurs almost instantaneously. However, multi-venue trade strings exhibit immediately negative realized spreads, while trade strings that only execute on a single venue provide liquidity suppliers with positive realized spreads for at least the first 20 seconds after the trades on average. The persistent difference between the negative realized spreads for multi-venue trade strings and the positive realized spreads for single-venue trade strings is therefore the economic value to liquidity suppliers who are able to avoid more toxic multi-venue trades. Similarly, Malinova and Park (2017) find that multi-venue trades in Canadian equities have double the price impact of single-venue trades, even after controlling for trader type and trade size. Meanwhile, Conrad and Wahal (2019) find that realized spreads are only marginally positive in the U.S. equities market and most of the long term price impact occurs within a few seconds after each trade.

In a second classification, we split trade strings into those that deplete an NBBO price level and those that do not. By construction, depleting trade strings exert instantaneous adverse selection on the liquidity supplier because



Figure 3.4: **Realized Spread within One Minute by Number of Venues Accessed**

This figure presents the average realized spreads after trade strings over 100 milliseconds, 1, 5, 10, 20 and 60 seconds. Trade strings are split into those that access only a single venue, and those that access multiple venues. Trade strings are defined as a series of trades in the same stock and direction that execute within a 30 millisecond window.



the NBBO quotes move in the direction of the trade, leading to lower realized spreads. Thus, while the two classifications are related, some single-venue trade strings are depleting while some multi-venue trade strings are not. The combination of these two classifications results in a total of four trade string categories. For brevity, we refer to multi-venue trades as SOR and single-venue trades as non-SOR.

Panel A of Figure 3.5 presents Alpha's 2-by-2 daily trade composition proportions for both depleting/non-depleting trades and SOR/non-SOR trades. Small orders are likely to be filled on one venue without depleting the NBBO. The proportion of these non-depleting, non-SOR trades increases dramatically, from 21% on old Alpha to 44% on new Alpha. Conversely, large trades across multiple venues, such as a SOR spray from an institutional investor, are more likely to exhaust all liquidity available at the NBBO. The incidence of depleting, SOR trades exhibits an almost matched decline, from 44% to 25%. Little

change is observed for both depleting non-SOR trades and non-depleting SOR trades. These observations show that liquidity suppliers on new Alpha are selective in fading their quotes to avoid multi-venue depleting trades, rather than simply all multi-venue trades. While the fraction of multi-venue trades does decline on Alpha, it still accounts for 50% of turnover on average after the relaunch, providing further evidence on the limited impact of removal from the Order Protection Rule on liquidity demander strategies. For comparison, Panel B of Figure 3.5 presents CX2's 2-by-2 trade composition, which shows that the proportion of non-depleting, non-SOR trades decreases from 44% to 38% after Alpha's relaunch, indicating greater aggressiveness by liquidity demanders on this venue.

We formally test for changes in trade composition on each venue with equations of the form:

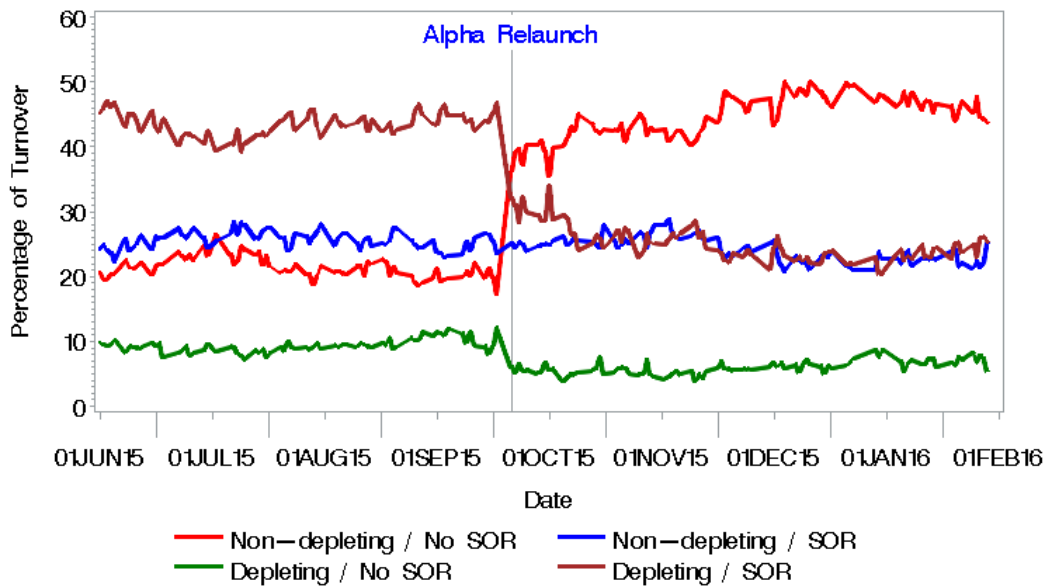
$$\begin{aligned} TradeComposition_{i,d,v} = & \beta_1 Post_d + \beta_2 Price_{i,d} + \beta_3 Turnover_{i,d,v} \\ & + \beta_4 Volatility_{i,d} + \beta_5 VIXC_d + FE_i + \epsilon_{i,d,v} \end{aligned} \quad (3.8)$$

where  $TradeComposition_{i,d,v}$  is the proportion of turnover for stock  $i$  on day  $d$  at venue  $v$  that participated in either a depleting trade string, a multi-venue trade string, or a jointly depleting and multi-venue trade string. All other variables are as described in Equation 3.7. Table 3.5 reports the results. The proportion of trading on Alpha that participated in depleting trade strings declines by 23.84%, while increases are observed across the other venues. The proportion of trading on Alpha that participated in a multi-venue trade string decreases by 20.49%. Consistent with its observed reduction in quote fade, multi-venue trade string participation increases by 3.31% on CX2. After


Figure 3.5: **Trading Volume Composition by Trade String Type on Alpha and CX2**

This figure presents a decomposition of Alpha's on-market turnover by trade string type. We construct trade strings by joining all trades in the same direction for a stock separated by less than 30 milliseconds. We distinguish between trade strings that leave the top level of quoted depth at the NBBO depleted in contrast with those that do not (undepleted). Smart order router (SOR) strings are those that execute on multiple venues.

(a) Panel A: Alpha



(b) Panel B: CX2



Alpha’s relaunch, its participation in multi-venue depleting trade strings decreases by 19.73%, while Chi-X and CX2 participation increases. We attribute the difference between the coefficient estimates for the post-relaunch changes in Alpha’s proportion of unconditioned depleting trades and multi-venue depleting trades, 4.11%, to the speed bump’s segmentation of symmetric public information released to all traders at the same time, arising from sources other than the order flow on other venues (Baldauf and Mollner, 2019, Brolley and Cimon, 2018, Budish et al., 2019). The similar magnitudes of these two coefficient estimates supports our earlier observation that Alpha’s speed bump is used by liquidity suppliers to selectively fade against depleting multi-venue trades, rather than all multi-venue trades.

Table 3.5: **Changes in NBBO Depletion and SOR on Each Venue Relative to the Pre-Relaunch Period**

This table reports coefficient estimates and t-statistics for the determinants of depleting trade proportion and multi-venue smart order routed (SOR) trade proportion for each of Alpha, Chi-X, CX2 and TSX across TSX Composite Index securities, after Alpha's relaunch relative to previous levels using the specification in Equation 3.8, where deplete best, SOR proportion and the combination of these characteristics for stock  $i$  on day  $d$  at venue  $v$  is expressed as the sum of an indicator variable for the post-relaunch period and control variables for the natural logarithms of price and each venue's turnover, realized intraday volatility, implied inter-day volatility, a stock specific mean and an error term. Trade strings are constructed by joining all trades in the same direction separated by less than 30 milliseconds. Trades within a string are called depleting if the entire NBB or NBO depth is consumed. Trades within a string are called multi-venue if the string contained trades on more than one venue. Deplete best is calculated as the proportion of turnover on a venue that consists of depleting trades, while SOR is calculated as the proportion of turnover on a venue that consists of multi-venue trades, and the interaction between deplete best and SOR is the proportion of turnover on a venue where both these conditions are met. The observation period runs from 1 June 2015 to 8 January 2016. We add a \*/\*\*/\*\*\*/ to the t-statistic to indicate statistical significance at the 90%/95%/99% levels, respectively. Standard errors are double clustered by stock and date.

	Deplete Best				SOR				Deplete Best and SOR			
	Alpha	Chi-X	CX2	TSX	Alpha	Chi-X	CX2	TSX	Alpha	Chi-X	CX2	TSX
Post	-23.84 (-39.73)***	3.47 (8.97)***	5.89 (12.56)***	3.97 (12.70)***	-20.49 (-28.83)***	-1.22 (-3.02)***	3.31 (8.09)***	-2.67 (-6.58)***	-19.73 (-35.27)***	2.06 (5.36)***	3.85 (11.54)***	-0.45 (-1.21)
Price	15.16 (10.90)***	13.25 (8.56)***	11.31 (8.67)***	19.74 (10.62)***	2.36 (0.98)	-3.26 (-2.74)***	0.78 (0.79)	-13.51 (-10.82)***	10.67 (5.38)***	5.68 (3.48)***	7.89 (7.04)***	-1.08 (-0.75)
Turnover	-2.81 (-9.36)***	-0.79 (-2.89)***	1.97 (5.21)***	-7.02 (-22.72)***	-1.22 (-3.32)***	2.67 (9.80)***	1.00 (3.69)***	3.41 (11.37)***	-2.08 (-7.39)***	0.31 (1.18)	1.51 (5.81)***	-0.97 (-3.54)***
Volatility	0.80 (17.65)***	0.64 (13.61)***	0.52 (11.11)***	1.00 (17.98)***	0.22 (4.33)***	0.08 (2.08)**	0.32 (7.39)***	-0.43 (-8.94)***	0.58 (13.08)***	0.42 (9.37)***	0.38 (9.43)***	0.13 (2.78)***
VIXC	0.18 (4.45)***	0.12 (3.38)***	0.21 (4.84)***	0.09 (3.23)***	0.06 (1.28)	0.12 (3.90)***	-0.03 (-0.84)	-0.25 (-6.24)***	0.10 (2.75)***	0.07 (2.05)**	0.13 (4.07)***	-0.18 (-5.44)***
Adjusted R2	47.31%	7.85%	14.41%	16.33%	39.50%	3.76%	5.02%	7.34%	41.96%	4.09%	9.17%	1.50%
# Obs	31,718	31,755	31,684	31,755	31,718	31,755	31,684	31,755	31,718	31,755	31,684	31,755

We provide an additional, although somewhat noisy, proxy for order flow segmentation on Alpha after its relaunch by examining the changes in composition of liquidity providing and liquidity demanding brokers. Conversations with practitioners and securities regulators indicate that two brokers handle the majority of retail order flow in Canada and a large proportion of HFT activity passes through two other brokers via their direct market access (DMA) facilities. The remaining order flow contains a mixture of various trader types. Brokers can also elect to anonymize their identity and forgo the opportunity to participate in broker preferencing for enhanced non-marketable order queue priority.

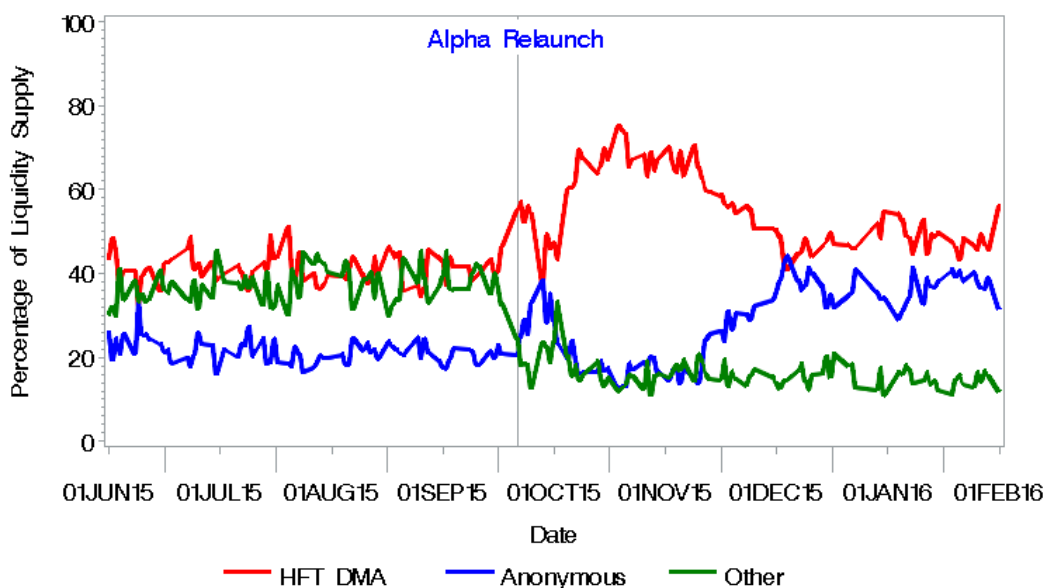
Panel A of Figure 3.6 presents the market share of liquidity supply on Alpha by broker type. The HFT brokers' combined participation in liquidity provision on Alpha increases from 41% to 57%. This increase is consistent with the requirement to continuously monitor the market with latency of less than a millisecond to benefit from the speed bump's delay exemption, which only very fast liquidity suppliers are able to do. Anonymous broker attribution also increases, from 21% to 27%. Interestingly, the daily trends of HFT broker participation are the reverse of the anonymous broker, indicating that some HFTs alternate between broadcasting and anonymizing their identity. Panel B of Figure 3.6 presents the market share of liquidity demand on Alpha by broker type. The retail brokers' combined participation in consuming liquidity increases from 17% to 27%. Retail orders tend to be smaller and less likely to exceed the size requirements of individual Post Only orders. Further, consistent with Battalio et al. (2016), the trading fee rebates on Alpha for marketable orders are attractive to retail brokers who do not pass this rebate on to their clients, which Brolley and Malinova (2013) document to be the case in Canada. These trends suggest that new Alpha enables sophisticated HFT

liquidity suppliers to partially segment marketable orders from retail investors.


**Figure 3.6: Market Share by Broker Type on Alpha**

This figure presents Alpha's market share by broker type separately for the liquidity demanding and liquidity supplying side of each trade. Retail consists of two domestic Canadian banks that are known to constitute a large proportion of retail broking activity. HFT DMA consists of two banks that offer direct market access services to low latency proprietary trading firms that act as electronic liquidity providers.

(a) Panel A: Liquidity Supplier



(b) Panel B: Liquidity Demander



### 3.4.3 Traded Liquidity Metrics per Venue

Since we have established that Alpha’s relaunch leads to segmentation of order flow that is single-venue, non-depleting and often from retail traders, we now examine the second order impacts on each venue’s traded liquidity metrics. Panels A and B of Figure 3.7 present the daily average adverse selection and net-of-fees realized spreads, respectively, at a one second horizon on each venue over the sample period. Consistent with avoiding informed order flow such as depleting trade strings, average adverse selection costs on Alpha decrease immensely from 2.05 cents to 1.44 cents. Adverse selection increases slightly on other trading venues. Most of the savings on Alpha are retained by liquidity suppliers, demonstrated by the large increase in average realized spreads on Alpha from 0.53 cents to 1.15 cents, after adjusting for liquidity provider trading fee changes. Meanwhile, net-of-fees realized spreads on CX2 decline from 1.10 cents to 0.71 cents, indicating a lower proportion of uninformed order flow arriving at this competing venue with inverted trading fees.

We formally test for changes in traded liquidity metrics across the four major venues in our sample in the wake of Alpha’s market design changes using the following regression specification:

$$\begin{aligned} LiquidityMetric_{i,d,v} = & \beta_1 Post_d + \beta_2 Price_{i,d} + \beta_3 Turnover_{i,d,v} \\ & + \beta_4 Volatility_{i,d} + \beta_5 VIXC_d + FE_i + \epsilon_{i,d,v} \end{aligned} \quad (3.9)$$


where  $LiquidityMetric_{i,d,v}$  is either the effective spread, realized spread or adverse selection for stock  $i$  on day  $d$  at venue  $v$ , and all other variables are as described in Equation 3.7. Table 3.6 presents the results. Average



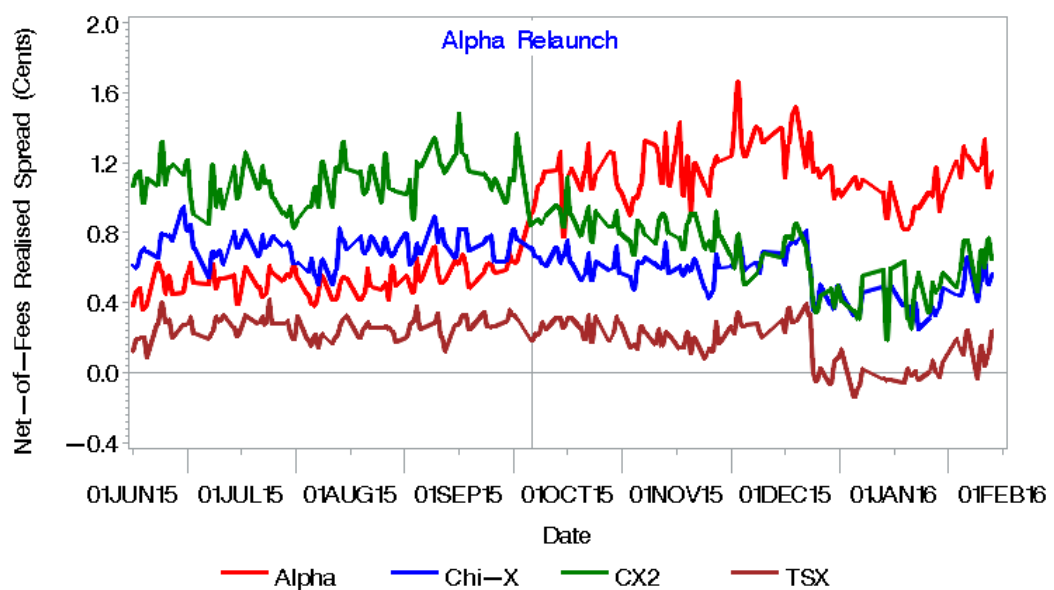
Figure 3.7: **Adverse selection and Net-of-Fees Realized Spreads per Venue**

This figure presents the average adverse selection and realized spreads per venue. Adverse selection is calculated as the directional change in NBBO midpoint price from the time of the trade to one second after the trade, as a measure of the trade's market impact. Realized spread is calculated as the directional difference between the trade price and the NBBO midpoint one second after the trade, as a proxy for the liquidity supplier's trading profits.

(a) Panel A: Adverse selection



(b) Panel B: Net-of-Fees Realized Spreads



adverse selection costs, which measure the permanent price impact of each trade at a one second horizon, decrease by 0.60 cents on Alpha after the market structure changes, consistent with segmentation of less informed order flow and in-line with theoretical predictions (Brolley and Cimon, 2018, Han et al., 2014). Effective spreads on Alpha increase by 0.55 cents, which is almost identical to the 0.56 cent decrease in liquidity taking fee, from a fee of 0.18 cents to a rebate of 0.10 cents, for a round-trip transaction. Consistent with Malinova and Park (2015), liquidity suppliers pass on changes in explicit fees, even under changes to inverted maker-taker pricing schemes.

Lower adverse selection and wider effective spreads both contribute to substantially higher realized spreads for liquidity suppliers on Alpha, which increase by 1.14 cents. Increases in liquidity supplier profits vastly exceed the change in their trading fees, from a rebate of 0.14 cents to a fee of 0.10 cents initially, and 0.16 cents after the end of the introductory pricing period at the end of November 2015. This finding contrasts with the theoretical literature (Brolley and Cimon, 2018, Han et al., 2014), which suggests that lower adverse selection should be incorporated into narrower effective spreads rather than higher realized spreads. We offer two potential explanations for this discrepancy. Firstly, there could be lack of competition among fast liquidity suppliers to quote narrower spreads, especially among those who are willing to display orders large enough to meet the size requirements for undelayed Post Only orders. Secondly, the speed bump on Alpha is only useful if it quotes orders at the same price as other venues. It does not provide the opportunity to fade quotes after observing trades on other venues at the same price if it quotes at better prices than the other venues.

Consistent with a higher concentration of residual informed order flow, adverse selection increases by between 0.22 cents and 0.24 cents on other venues

after Alpha's relaunch. Interestingly, liquidity providers on Chi-X and TSX absorb approximately half of the higher adverse selection costs into lower realized spreads and pass on the remainder as wider effective spreads. These findings are consistent with the model predictions in Brolley and Cimon (2018), where the venue without a speed bump faces higher adverse selection and increases its bid-ask spreads, as well as Biais et al. (2015), where the establishment of a slow venue that excludes SOR trades results in higher adverse selection on the fast venues that accept SOR trades. More broadly, the results are in line with existing empirical evidence suggesting that the segmentation of uninformed marketable orders to private or dark venues increases the toxicity of the remaining order flow on public lit markets (Comerton-Forde and Putniņš, 2015, Easley et al., 1996). On the other hand, realized spreads on CX2 decrease sharply by 0.33 cents due to tighter effective spreads, although the latter change is not statistically significant. We conjecture that liquidity suppliers on CX2, which has virtually the same inverted maker-taker trading fees as new Alpha, quote narrower bid-ask spreads and accept lower profits, similar to the predictions of Hoffmann (2014), to more effectively compete with new Alpha for order flow. Effective spreads tend to increase with the control variables for share price and volatility and decline with turnover, consistent with the prior literature.

Table 3.6: **Changes in Traded Liquidity Metrics on Each Venue Relative to the Pre-Relaunch Period**

This table reports coefficient estimates and t-statistics for the determinants of transaction cost measures across TSX Composite Index securities traded on Alpha, Chi-X, CX2 and TSX around Alpha's relaunch using the specification in Equation 3.9, where each transaction cost metric for stock  $i$  on day  $d$  at venue  $v$  is expressed as the sum of an indicator variable for the post-relaunch period, and control variables for the natural logarithms of price and each venue's turnover, realized intraday volatility, implied inter-day volatility, a stock specific mean and an error term. Effective spreads and realized spreads, in cents, compare the traded price with the prevailing NBBO midpoint and the NBBO midpoint after 1 second, respectively. Adverse selection is calculated as the effective spread minus the realized spread. The observation period runs from 1 June 2015 to 8 January 2016. We add a  $*/**/**$  to the t-statistic to indicate statistical significance at the 90%/95%/99% levels, respectively. Standard errors are double clustered by stock and date.

	Effective Spread				Realized Spread				Adverse Selection			
	Alpha	Chi-X	CX2	TSX	Alpha	Chi-X	CX2	TSX	Alpha	Chi-X	CX2	TSX
Post	0.55 (7.87)***	0.12 (3.04)***	-0.08 (-1.59)	0.10 (2.86)***	1.14 (13.30)***	-0.10 (-3.63)***	-0.33 (-6.05)***	-0.13 (-5.54)***	-0.60 (-15.48)***	0.22 (6.59)***	0.24 (7.54)***	0.23 (7.12)***
Price	2.17 (5.51)***	2.25 (5.42)***	2.29 (5.97)***	2.07 (5.69)***	0.65 (3.02)***	0.20 (1.10)	0.70 (2.69)***	-0.45 (-3.13)***	1.52 (5.84)***	2.05 (8.05)***	1.56 (9.69)***	2.56 (9.76)***
Turnover	-0.14 (-4.14)***	-0.41 (-9.42)***	-0.37 (-8.14)***	-0.41 (-7.94)***	-0.06 (-1.49)	-0.17 (-6.68)***	-0.24 (-5.81)***	0.15 (6.81)***	-0.08 (-3.11)***	-0.25 (-7.58)***	-0.11 (-4.16)***	-0.55 (-12.12)***
Volatility	0.07 (7.88)***	0.08 (8.51)***	0.08 (8.54)***	0.08 (7.94)***	0.00 (-0.20)	-0.01 (-2.42)**	0.02 (3.04)***	-0.04 (-11.21)***	0.07 (10.44)***	0.09 (11.08)***	0.07 (11.86)***	0.12 (11.58)***
VIXC	0.01 (3.49)***	0.01 (2.51)**	0.02 (5.36)***	0.01 (4.87)***	0.01 (2.63)***	0.01 (2.29)**	0.02 (4.29)***	0.01 (3.75)***	0.00 (1.28)	0.00 (0.62)	0.01 (2.64)***	0.01 (2.23)**
Adjusted R2	13.72%	15.48%	14.05%	16.71%	20.69%	2.50%	5.94%	4.00%	14.33%	12.82%	8.96%	21.82%
# Obs	31,723	31,755	31,688	31,755	31,723	31,755	31,688	31,755	31,723	31,755	31,688	31,755

We conduct a robustness check to ascertain whether changes in traded liquidity on each venue after Alpha’s relaunch are attributable to the segmentation effects of speed bump exempt Post Only order usage. To do this, we repeat the regression analysis described in Equation 3.9, but substitute the post-event indicator variable with the proportion of order cancellations on Alpha in the post-event period that were equal to or greater than the minimum Post Only order size requirement on each stock-day. This alternative model specification allows for cross-sectional and time-series variation in the impact of Alpha’s speed bump utilization on each venue’s traded liquidity outcomes.

Results are reports in Table 3.7. After accounting for 75% of quote cancellations on new Alpha being sufficiently large to be Post Only orders, the coefficient estimates for post-event changes in adverse selection and realized spreads across all venues are broadly consistent with the event study results reported in Table 3.6. However, several peculiarities emerge for changes in effective spreads with this model specification, compared with the event study. Effective spreads widen more strongly on Alpha, with higher statistical significance. Meanwhile, effective spread reduction on CX2 in response to Post Only order usage on Alpha is statistically significant and much larger at 0.14 cents, compared with 0.08 cents estimated in the event study. In contrast, effective spread widening on Chi-X and TSX have lower economic and statistical significance, indicating that a portion of the changes observed after Alpha’s relaunch on these venues may be attributable to factors other than usage of speed bump exempt orders.

Table 3.7: **Changes in Traded Liquidity Metrics on Each Venue in Response to Alpha Fast Order Cancellations**

This table reports coefficient estimates and t-statistics for the determinants of transaction cost measures across TSX Composite Index securities traded on Alpha, Chi-X, CX2 and TSX around Alpha's relaunch using the specification in Equation 3.9, where each transaction cost metric for stock  $i$  on day  $d$  at venue  $v$  is expressed as the sum of the proportion of order cancellations on Alpha that met the Post Only size requirement for non-delayed cancellation after relaunch and control variables for the natural logarithms of price and each venue's turnover, realized intraday volatility, implied inter-day volatility, a stock specific mean and an error term. Effective spreads and realized spreads, in cents, compare the traded price with the prevailing NBBO midpoint and the NBBO midpoint after 1 second, respectively. Adverse selection is calculated as the effective spread minus the realized spread. The observation period runs from 1 June 2015 to 8 January 2016. We add a  $*/**/**$  to the t-statistic to indicate statistical significance at the 90%/95%/99% levels, respectively. Standard errors are double clustered by stock and date.

	Effective Spread				Realized Spread				Adverse selection			
	Alpha	Chi-X	CX2	TSX	Alpha	Chi-X	CX2	TSX	Alpha	Chi-X	CX2	TSX
Fast Order	0.78	0.11	-0.14	0.09	1.49	-0.16	-0.47	-0.19	-0.73	0.27	0.31	0.29
Cancel	(8.43)***	(2.04)**	(-2.15)**	(2.00)**	(13.28)***	(-4.65)***	(-6.50)***	(-6.37)***	(-14.39)***	(6.14)***	(7.43)***	(6.66)***
Price	2.19	2.21	2.27	2.04	0.66	0.19	0.69	-0.45	1.53	2.02	1.55	2.53
	(5.59)***	(5.34)***	(5.94)***	(5.62)***	(3.19)***	(1.04)	(2.71)***	(-3.19)***	(5.97)***	(7.93)***	(9.60)***	(9.67)***
Turnover	-0.16	-0.40	-0.37	-0.40	-0.12	-0.16	-0.23	0.15	-0.04	-0.24	-0.11	-0.55
	(-5.07)***	(-9.28)***	(-8.00)***	(-7.86)***	(-3.46)***	(-6.62)***	(-5.67)***	(7.10)***	(-1.45)	(-7.43)***	(-4.12)***	(-12.09)***
Volatility	0.07	0.08	0.08	0.08	0.00	-0.01	0.02	-0.04	0.07	0.09	0.07	0.12
	(8.01)***	(8.42)***	(8.47)***	(7.87)***	(0.42)	(-2.54)**	(3.04)***	(-11.36)***	(10.35)***	(11.00)***	(11.86)***	(11.53)***
VIXC	0.01	0.01	0.02	0.02	0.01	0.01	0.02	0.01	0.00	0.00	0.01	0.01
	(3.11)***	(2.99)***	(5.56)***	(5.30)***	(2.82)***	(2.60)***	(4.46)***	(4.04)***	(0.84)	(1.05)	(2.91)***	(2.54)**
Adjusted R2	15.34%	15.31%	14.19%	16.56%	23.14%	2.69%	6.56%	4.37%	14.51%	12.75%	9.12%	21.84%
# Obs	31,723	31,755	31,688	31,755	31,723	31,755	31,688	31,755	31,723	31,755	31,688	31,755

### 3.4.4 Liquidity Provider Metrics on Inverted Fee Venues

The concurrent adoption of inverted maker-taker pricing on Alpha could confound analysis of the speed bump’s impacts when examining the market design changes with a pure event study regression methodology. Inverted trading fee venues typically have lower adverse selection costs and higher realized spreads (Battalio et al., 2016). To isolate the impact of the speed bump, we directly compare traded liquidity metrics in the post-event period on Alpha and CX2, a venue that has almost the same fee structure, across different types of liquidity supplying brokers on Alpha, particularly those that are likely fast enough to utilize the speed bump. In this setting, any differences observed are attributable to the trader’s ability to utilize Alpha’s speed bump. Each stock-day has one observation for CX2 and three observations for Alpha, separately for whether the liquidity supplying broker is one of the two brokers associated with HFT activity, the broker chose to use the anonymous flag and forgo broker preferencing by declining to reveal their identity, and all other brokers. This grouping is the same as Panel A of Figure 3.6. We do not separate liquidity suppliers on CX2 into their respective broker types because inverted trading fee venues are considered less attractive to HFT liquidity providers, who are able to use their speed to attain queue priority on venues that offer rebates for supplying liquidity (Yao and Ye, 2018), rather than needing to ‘pay-to-post’ for shorter queue times. We utilize the following regression equation:

$$\begin{aligned}
LiquidityMetric_{i,d,b} = & \beta_1 AnonOnAlpha_{i,d,b} + \beta_2 HFTDMAOnAlpha_{i,d,b} \\
& + \beta_3 OtherOnAlpha_{i,d,b} + \beta_4 FastOrderCancel_d * AnonOnAlpha_{i,d,b} \\
& + \beta_5 FastOrderCancel_d * HFTDMAOnAlpha_{i,d,b} \\
& + \beta_6 FastOrderCancel_d * OtherOnAlpha_{i,d,b} \\
& + \beta_7 Price_{i,d} + \beta_8 Turnover_{i,d,b} \\
& + \beta_9 Volatility_{i,d} + \beta_{10} VIXC_d + FE_{i,b} + \epsilon_{i,d,b}
\end{aligned} \tag{3.10}$$

where  $LiquidityMetric_{i,d,b}$  is either the effective spread, realized spread or adverse selection for stock  $i$  on day  $d$  and venue-broker type  $b$ ,  $AnonOnAlpha_{i,d,b}$  is an indicator variable equal to one if the observation corresponds to the anonymous broker ID on Alpha and zero otherwise,  $HFTDMAOnAlpha_{i,d,b}$  is an indicator variable equal to one if the observation corresponds to either of the HFT-affiliated broker IDs on Alpha and zero otherwise and  $OtherOnAlpha_{i,d,b}$  is an indicator variable equal to one if the observation corresponds to a broker ID on Alpha other than HFT or anonymous, and zero otherwise,  $Turnover_{i,d,b}$  is the natural logarithm of on-market trading turnover for stock  $i$  on day  $d$  and venue-broker type  $b$ ,  $FE_{i,b}$  indicates stock-venue-broker type fixed effects, and  $\epsilon_{i,d,b}$  is an error term. All other variables are as defined in Equation 3.7.

These results are reported in Table 3.8. Overall, adverse selection costs for anonymous brokers and HFT-affiliated brokers on Alpha are 0.63 cents and 0.12 cents lower than CX2, respectively. Therefore, the most sophisticated HFTs on Alpha are electing to maintain their anonymity. Panel A of Figure 3.6 shows that the daily passive market shares of these two groups has a strong negative correlation, indicating that some traders may be alternating between



the two groups. Interacting these indicator variables with the proportion of order cancellations that meet the Post Only order size requirement on Alpha markedly diminishes economic and statistical significance for the difference in adverse selection between liquidity suppliers on CX2 and those on Alpha who are neither HFT nor anonymous. Therefore, adverse selection on Alpha is only lower than CX2 for fast traders, confirming that the speed bump is critical for segmentation of less informed order flow.

Anonymous brokers on Alpha have effective spreads that are 0.18 cents wider than CX2 when they supply liquidity. Notably, HFT brokers on Alpha are on the passive side of trades with much wider effective spreads, averaging a difference of 0.47 cents. Anonymous brokers on Alpha earn the highest realized spreads at a one second post-trade horizon, which arises mainly from lower adverse selection, while HFT brokers are not far behind, although most of their additional profits originate from wider effective spreads. Traders on Alpha other than HFT-affiliated and anonymous brokers supply liquidity at substantially narrower effective spreads than both these groups, as well as CX2, consistent with Hoffmann (2014), and earn the lowest realized spreads. A noteworthy implication of these results is that Alpha's speed bump not only segments order flow between venues and redistributes liquidity provision profits towards Alpha, it also redistributes liquidity provision profits on Alpha. Fast and sophisticated HFTs can cherry-pick the most favorable marketable orders on Alpha, leaving slower liquidity suppliers worse off.

### **3.4.5 Impacts on Consolidated Liquidity**

Having established that Alpha's speed bump enables it to segment single-venue, non-depleting trades with lower adverse selection, increasing the prof-

**Table 3.8: Traded Liquidity Metrics for Broker Types on Alpha Compared with CX2**

This table reports coefficient estimates and t-statistics for the determinants of transaction cost across TSX Composite Index securities on CX2 compared with various types of broker accounts supplying liquidity on Alpha after the latter's relaunch, using the specification in Equation 3.10, where each liquidity metric for stock  $i$  on day  $d$  and venue-broker type  $b$  is expressed as the sum of indicator variables for anonymous, HFT DMA and other brokers on Alpha (with CX2 being the base case), interaction terms between these variables and the proportion of order cancellations on Alpha that met the Post Only size requirement, and control variables for the natural logarithms of price and each venue-broker type's turnover, realized intraday volatility, implied inter-day volatility, a stock specific mean and an error term. HFT DMA consists of two brokers that offer direct market access services to proprietary traders. Anonymous consists of all participants that chose not to broadcast their identity. All other brokers are grouped as other. Effective spreads and realized spreads compare the traded price with the prevailing NBBO midpoint and the NBBO midpoint after 1 second, respectively. Adverse selection is calculated as the effective spread minus the realized spread. The observation period runs from 21 September 2015 to 8 January 2016. We add a \*/\*\*/\*\* to the t-statistic to indicate statistical significance at the 90%/95%/99% levels, respectively. Standard errors are double clustered by stock and date.

	Effective Spread		Realized Spread		Adverse Selection	
Anon on Alpha	0.18 (2.49)**		0.83 (7.31)***		-0.63 (-10.90)***	
HFT DMA on Alpha	0.47 (9.19)***		0.74 (10.17)***		-0.27 (-8.41)***	
Other on Alpha	-0.38 (-9.41)***		-0.24 (-4.45)***		-0.12 (-3.12)***	
Anon on Alpha		0.25		1.05		-0.77
*Fast Order Cancel		(2.88)***		(7.83)***		(-11.26)***
HFT DMA on Alpha		0.61		0.94		-0.33
*Fast Order Cancel		(9.97)***		(10.38)***		(-7.75)***
Other on Alpha		-0.55		-0.47		-0.07
*Fast Order Cancel		(-9.63)***		(-6.43)***		(-1.42)
Price	1.40 (5.34)***	1.41 (5.39)***	0.03 (0.19)	0.11 (0.73)	1.39 (8.20)***	1.34 (7.79)***
Turnover	0.07 (4.98)***	0.06 (5.24)***	0.19 (10.75)***	0.17 (11.10)***	-0.12 (-10.26)***	-0.10 (-9.70)***
Volatility	0.05 (5.92)***	0.05 (5.93)***	-0.01 (-1.90)*	-0.01 (-1.59)	0.06 (10.35)***	0.06 (10.37)***
VIXC	0.00 (-0.50)	0.00 (-0.38)	-0.01 (-1.50)	0.00 (-0.79)	0.00 (0.84)	0.00 (0.27)
Adjusted R2	11.35%	12.27%	11.70%	13.04%	6.06%	6.38%
# Obs	45,632	45,632	45,632	45,632	45,632	45,632

its of its fast liquidity suppliers and redistributing adverse selection to other venues, we now investigate how these changes affect consolidated quoted and traded liquidity across all venues. We test for changes in these liquidity metrics at the NBBO level aggregated across Alpha, TSX, Chi-X and CX2 using the regression equation:

$$\begin{aligned} LiquidityMetric_{i,d} = & \beta_1 Post_d + \beta_2 Price_{i,d} + \beta_3 Turnover_{i,d} \\ & + \beta_4 Volatility_{i,d} + \beta_5 VIXC_d + FE_i + \epsilon_{i,d} \end{aligned} \quad (3.11)$$

where  $LiquidityMetric_{i,d}$  is a measure of consolidated liquidity for stock  $i$  on day  $d$ ,  $Turnover_{i,d}$  is the natural logarithm of on-market turnover aggregated across all venues,  $Price_{i,d}$  is the inverse of the time-weighted NBBO midpoint price where the liquidity metric is expressed in basis points, following Hendershott et al. (2011), and the natural logarithm of the time-weighted NBBO midpoint price otherwise, and  $\epsilon_{i,d}$  is an error term. All other variables are as defined in Equation 3.7.

Table 3.9 presents the results. Quoted spreads increase by a substantial 0.17 cents in absolute terms and 0.21 basis points in relative terms, although the latter is not statistically significant. Around half of the increases in consolidated quoted depth is contributed by larger orders on Alpha, as required for exemption from the speed bump. Although the order book becomes deeper on average, the proportion of trading activity that depletes all available depth at the NBB or NBO price level also increases, by 2.29%. Implicit costs of demanding liquidity, as measured by effective spreads, increase by 0.12 cents or 0.20 basis points, similar to the individual venue changes reported in Table 3.6 for TSX and Chi-X, which are the two larger venues. The smaller increase in

effective spreads than quoted spreads is consistent with higher quoted depth. Aggregate realized spreads at a post-trade horizon of one second decline by 0.05 cents or 0.59 basis points, while adverse selection costs at the same time horizon increase by 0.17 cents or 0.83 basis points. These changes are substantially smaller than the individual venue changes estimated for venues other than Alpha in Table 3.6, which demonstrates Alpha's considerable redistributive impact on adverse selection and liquidity supplier profits across venues.

Table 3.9: **Changes in Consolidated Liquidity Metrics Relative to the Pre-Relaunch Period**

This table reports coefficient estimates and t-statistics for the determinants of consolidated liquidity across TSX Composite Index securities around Alpha's relaunch using the specification in Equation 3.11, where each liquidity metric for stock  $i$  on day  $d$  is expressed as the sum of an indicator variable for the post-relaunch period, and control variables for price, natural logarithm of turnover, realized intraday volatility, implied inter-day volatility, a stock specific mean and an error term. An inverse transformation is applied to the control variable for price in model specifications where the liquidity metric is expressed in basis points and a natural logarithmic transformation is applied otherwise. Quoted spreads and quoted depths are time-weighted and measured using the NBBO prices and volumes. Effective spreads and realized spreads compare the traded price with the prevailing NBBO midpoint and the NBBO midpoint after 1 second, respectively. Adverse selection is calculated as the effective spread minus the realized spread. NBBO depletion is the proportion of trading volume that occurred as part of a trade string that displaced the entire NBB or NBO depth. The observation period runs from 1 June 2015 to 8 January 2016. We add a \*/\*\*/\*\*\*/ to the t-statistic to indicate statistical significance at the 90%/95%/99% levels, respectively. Standard errors are double clustered by stock and date.

	Quoted Spread			Quoted Depth			Effective Spread			Realized Spread			Adverse selection			NBBO	
	Cents	Bps	All Venues	Ex. Alpha	Cents	Bps	Cents	Bps	Cents	Bps	Cents	Bps	Cents	Bps	Depletion		
Post	0.17 (3.46)***	0.21 (1.48)	0.16 (13.74)***	0.09 (7.24)***	0.12 (3.48)***	0.20 (2.09)**	-0.05 (-2.18)**	-0.59 (-4.25)***	0.17 (5.88)***	-0.83 (6.08)***	0.17 (5.88)***	-0.59 (-4.25)***	0.17 (5.88)***	0.83 (6.08)***	2.29 (7.77)***		
Price	3.04 (6.16)***	77.92 (24.86)***	-0.09 (-1.17)	-0.12 (-1.69)*	2.16 (5.81)***	85.39 (46.93)***	-0.19 (-1.27)	45.88 (7.92)***	2.39 (9.31)***	29.41 (7.92)***	2.39 (9.31)***	45.88 (7.92)***	2.39 (9.31)***	29.41 (7.92)***	17.87 (11.05)***		
Turnover	-0.78 (-10.28)***	-2.56 (-12.62)***	0.23 (17.14)***	0.25 (20.23)***	-0.46 (-8.27)***	-1.22 (-7.61)***	0.01 (0.55)	0.57 (3.36)***	-0.47 (-10.33)***	-2.04 (-13.70)***	-0.47 (-10.33)***	0.57 (3.36)***	-0.47 (-10.33)***	-2.04 (-13.70)***	-4.64 (-14.77)***		
Volatility	0.10 (8.11)***	0.33 (11.00)***	-0.03 (-16.63)***	-0.03 (-20.25)***	0.08 (8.11)***	0.31 (12.58)***	-0.03 (-8.60)***	-0.37 (-11.51)***	0.11 (11.21)***	0.73 (22.24)***	0.11 (11.21)***	-0.37 (-11.51)***	0.11 (11.21)***	0.73 (22.24)***	0.90 (17.65)***		
VIXC	0.03 (7.28)***	0.09 (8.16)***	-0.01 (-9.42)***	-0.01 (-8.16)***	0.01 (4.44)***	0.02 (2.22)**	0.01 (3.59)***	0.08 (6.61)***	0.01 (1.99)**	-0.06 (-4.25)***	0.01 (1.99)**	0.08 (6.61)***	0.01 (1.99)**	-0.06 (-4.25)***	0.11 (4.23)***		
Adjusted R2	23.38%	66.93%	33.57%	32.59%	17.54%	64.58%	2.31%	15.26%	21.34%	37.64%	21.34%	15.26%	21.34%	37.64%	13.86%		
# Obs	31,755	31,755	31,755	31,755	31,755	31,755	31,755	31,755	31,755	31,755	31,755	31,755	31,755	31,755	31,755		

We also examine cross-sectional variation in consolidated liquidity changes after Alpha’s relaunch, across quartiles of stocks constructed from the average proportion of trades that deplete all liquidity at a price level in the pre-event period, using Equation 3.11 separately for each quartile. Menkveld and Zoican (2017) postulate that the impact of trading speed on liquidity depends on the relative arrival rate of information in a security. Since the recent literature has acknowledged that price movements in a security constitute information (Foucault et al., 2016), we contend that the proportion of trades that result in NBBO depletion encapsulates the rate of information arrival in each stock.

For brevity, Table 3.10 reports the coefficient estimates and t-statistics on the post-event indicator variable only. Adverse selection increases in both absolute and relative terms across all quartiles, with the largest absolute increase of 0.28 cents in the highest depletion quartile and the largest relative increase of 2.10 basis points in the lowest depletion quartile. Interestingly, higher losses to adverse selection are absorbed by liquidity suppliers into smaller realized spreads in the lower depletion quartiles, where realized spreads tend to be higher and more able to absorb market impact without declining to zero. In contrast, adverse selection increases are passed on to liquidity demanders via wider effective spreads in the higher depletion quartiles, as realized spreads among this group are already very small, with lower capacity to absorb further losses. Similarly, quoted spreads widen in the quartiles with higher depletion.

**Table 3.10: Changes in Consolidated Liquidity Metrics Relative to the Pre-Relaunch Period by NBBO Depletion Quartile**  
This table reports coefficient estimates and t-statistics for the determinants of liquidity across TSX Composite Index securities around Alpha's relaunch using the specification in Equation 3.11, where each liquidity metric for stock  $i$  on day  $d$  is expressed as the sum of an indicator variable for the post-relaunch period, and control variables for price, natural logarithm of turnover, realized intraday volatility, implied inter-day volatility, a stock specific mean and an error term. An inverse transformation is applied to the control variable for price in model specifications where the liquidity metric is expressed in basis points and a natural logarithmic transformation is applied otherwise. The regression models are run separately for quartiles of stocks grouped by the proportion of trades that were immediately followed by an adverse movement in the NBB or NBO price, during the pre-relaunch period. For brevity, coefficient estimates and t-statistics are only reported for the post-relaunch variable. Quoted spreads and quoted depths are measured using the NBBO prices and volumes. Effective spreads and realized spreads compare the traded price with the prevailing NBBO midpoint and the NBBO midpoint after 1 second, respectively. Adverse selection is calculated as the effective spread minus the realized spread. NBBO depletion is the proportion of trading volume that occurred as part of a trade string that displaced the entire NBB or NBO depth. The observation period runs from 1 June 2015 to 8 January 2016. We add a \*/\*\*/\*\*\*/ to the t-statistic to indicate statistical significance at the 90%/95%/99% levels, respectively. Standard errors are double clustered by stock and date.

	Quoted Spread			Quoted Depth			Effective Spread			Realized Spread			Adverse selection			NBBO	
	Cents	Bps	All Venues	Ex.	Alpha	Cents	Bps	Cents	Bps	Cents	Bps	Cents	Bps	Cents	Bps	Depletion	
Most NBBO	0.39	0.95	0.24	0.11		0.29	0.64	0.01	-0.01	0.28	0.64	0.59					
Depletion	(4.73)***	(3.94)***	(12.58)***	(5.90)***		(4.20)***	(3.56)***	(0.18)	(-0.10)	(4.74)***	(4.07)***	(1.57)					
Quartile 2	0.18	0.35	0.20	0.10		0.12	0.26	-0.05	-0.13	0.16	0.39	2.19					
	(1.37)	(1.71)*	(12.09)***	(5.48)***		(1.44)	(1.93)*	(-1.04)	(-0.78)	(2.69)***	(2.70)***	(4.39)***					
Quartile 3	0.06	0.25	0.12	0.04		0.03	0.15	-0.05	-0.22	0.08	0.36	2.23					
	(1.10)	(0.82)	(6.22)***	(2.36)**		(0.97)	(0.76)	(-1.89)*	(-1.42)	(2.82)***	(2.25)**	(5.17)***					
Least NBBO	-0.01	-0.38	0.09	0.09		0.00	-0.01	-0.10	-1.98	0.12	2.10	3.98					
Depletion	(-0.65)	(-2.16)**	(4.95)***	(4.31)***		(-0.08)	(-0.09)	(-5.16)***	(-5.92)***	(3.93)***	(5.52)***	(7.95)***					

We conduct an additional robustness test for whether changes in consolidated liquidity after Alpha’s relaunch are attributable to the usage of its speed bump exempt Post Only orders. To do this, we repeat the analysis in Table 3.9, but substitute the post-event indicator variable in Equation 3.11 with the proportion of order cancellations on Alpha in the post-event period that were equal to or greater than the minimum Post Only order size requirement for each stock-day, to allow for cross-sectional and time-series variation in the impact of Alpha’s speed bump utilization on consolidated liquidity metrics. This alternative explanatory variable approach is similar to that adopted for Table 3.7.

Table 3.11 reports the results. Noting that 75% of Alpha’s order cancellations for the average stock-day after relaunch meet the Post Only size requirements, the entire reduction in consolidated realized spreads and increases in NBBO depth on Alpha, as well as NBBO depletion, can be attributed to Alpha’s speed bump implementation. Meanwhile, the coefficient estimates for post-event changes in quoted depths excluding Alpha, quoted spreads and effective spreads are very similar in magnitude to those reported for the event study approach in Table 3.9, indicating that only a small proportion of the changes observed in these variables after Alpha’s relaunch may be attributable to influences other than its speed bump implementation, such as its adoption of an inverted trading fee schedule, or fluctuations in liquidity more broadly.



Table 3.11: **Changes in Consolidated Liquidity Metrics Relative to the Pre-Relaunch Period in Response to Alpha Fast Order Cancellations**

This table reports coefficient estimates and t-statistics for the determinants of liquidity across TSX Composite Index securities in response to fleeting liquidity cancellation after Alpha's relaunch using the specification in Equation 3.11, where each liquidity metric for stock  $i$  on day  $d$  is expressed as the sum of the proportion of order cancellations on Alpha that met the Post Only size requirement for non-delayed cancellation, and control variables for price, natural logarithm of turnover, realized intraday volatility, implied inter-day volatility, a stock specific mean and an error term. An inverse transformation is applied to the control variable for price in model specifications where the liquidity metric is expressed in basis points and a natural logarithmic transformation is applied otherwise. Quoted spreads and quoted depths are measured using the NBBO prices and volumes. Effective spreads and realized spreads compare the traded price with the prevailing NBBO midpoint and the NBBO midpoint after 1 second, respectively. Adverse selection is calculated as the effective spread minus the realized spread. NBBO depletion is the proportion of trading volume that occurred as part of a trade string that displaced the entire NBB or NBO depth. The observation period runs from 1 June 2015 to 8 January 2016. We add a \*/\*\*/\*\*\*/ to the t-statistic to indicate statistical significance at the 90%/95%/99% levels, respectively. Standard errors are double clustered by stock and date.

	Quoted Spread			Quoted Depth			Effective			Realized Spread			Adverse selection			NBBO	
	Cents	Bps	All Venues	Ex. Alpha	Cents	Bps	Cents	Bps	Cents	Bps	Cents	Bps	Cents	Bps	Cents	Bps	Depletion
Fast Order	0.18	0.21	0.23	0.10	0.13	0.18	-0.08	-0.74	0.21	0.94	2.86						
Cancel	(2.71)***	(1.14)	(16.42)***	(6.55)***	(2.79)***	(1.43)	(-2.58)***	(-4.35)***	(5.39)***	(5.95)***	(7.57)***						
Price	3.00	78.11	-0.09	-0.13	2.14	85.62	-0.19	45.57	2.37	30.00	17.64						
	(6.12)***	(25.24)***	(-1.12)	(-1.91)*	(5.75)***	(47.42)***	(-1.31)	(7.87)***	(9.24)***	(7.91)***	(11.02)***						
Turnover	-0.78	-2.55	0.22	0.25	-0.46	-1.21	0.01	0.57	-0.47	-2.03	-4.63						
	(-10.24)***	(-12.69)***	(17.47)***	(20.36)***	(-8.21)***	(-7.6)***	(0.63)	(3.38)***	(-10.31)***	(-13.58)***	(-14.81)***						
Volatility	0.10	0.33	-0.03	-0.03	0.08	0.31	-0.03	-0.37	0.11	0.73	0.90						
	(8.05)***	(11.02)***	(-16.99)***	(-20.56)***	(8.06)***	(12.55)***	(-8.68)***	(-11.54)***	(11.16)***	(22.18)***	(17.6)***						
VIXC	0.03	0.09	-0.01	-0.01	0.01	0.02	0.01	0.08	0.01	-0.05	0.12						
	(7.55)***	(8.58)***	(-9.64)***	(-7.52)***	(4.86)***	(2.65)***	(3.84)***	(6.57)***	(2.26)**	(-3.81)***	(4.64)***						
Adjusted R2	23.21%	66.91%	36.28%	32.11%	17.42%	64.56%	2.40%	15.26%	21.33%	37.52%	13.84%						
# Obs	31,755	31,755	31,755	31,755	31,755	31,755	31,755	31,755	31,755	31,755	31,755						

### 3.5 Conclusion

Several stock exchanges in North America have recently implemented speed bumps, to alter the otherwise deterministic outcomes of speed competition in financial markets where orders are processed as soon as they arrive. This chapter investigates the introduction of an asymmetric, randomized duration speed bump on TSX Alpha, which allows traders to pay a higher trading fee to use a specialized non-marketable order type that is exempt from the delay. Liquidity suppliers on new Alpha who can monitor the market in ultra-high frequency are able to harness the information contained within the order flow on other venues, to avoid multi-market trades and significantly reduce their adverse selection costs. Realized spreads increase immensely in an environment where they are otherwise very close to zero, indicating lack of competition among fast liquidity suppliers on Alpha.

Segmentation of predominantly uninformed order flow with low adverse selection increases the fraction of informed trading on the remaining venues. The venue with a similar trading fee structure to new Alpha quotes narrower spreads to compete for marketable orders, while other venues widen their spreads to offset higher adverse selection, leading to wider consolidated spreads overall. Cross-sectionally, the speed bump's impacts are most pronounced in stocks where trades frequently result in price movements, precisely the type of order flow information that the speed bump is most useful for.

Additionally, this chapter develops several novel empirical techniques that enable the analysis of cross-market linkages in a fragmented equities trading environment, which O'Hara (2015) suggests as a particularly important issue for both policy and research. We propose methods to correctly assign trade initiation direction using trade and quote data emanating from a single data feed

and benchmark timestamp synchronization across multiple trading venues. We then construct trade strings from individual trades that either originate from a smart order router or respond to the same informational event and develop a quote fade metric that quantifies the economic costs associated with fleeting liquidity.

Asymmetric speed bumps have been touted as potential remedies in the “arms race for speed” by some (Baldauf and Mollner, 2019, Biais et al., 2015) and decried by others for being “*discriminatory, anti-competitive and facially inconsistent with the fundamental objectives*” of financial market regulation.<sup>13</sup> While our findings regarding the first order impacts on traders provided with a speed advantage and second order impacts on overall liquidity take place in the context of an asymmetric randomized duration speed bump, they also generalize to other innovations that offer differential speed access, such as co-location, premium data feeds and microwave connectivity. More broadly, this chapter suggests that caution is warranted for market design innovation proposals that endow some traders with systematic speed advantages over others, whether via speed bump or otherwise.

## 3.6 Appendix I: Speed Bump Designs

Speed bumps have been implemented or proposed on several trading venues around the world, with various designs that are intended to achieve different outcomes. In this section, we outline the different types of speed bump designs, alongside related theoretical models and predictions, practitioner implementations and empirical evidence to date. Common across all speed bump designs

---

<sup>13</sup>Statement of Commissioner Dan M. Berkovitz Regarding the ICE Futures U.S., Inc. Passive Order Protection Functionality Rule, <https://www.cftc.gov/PressRoom/SpeechesTestimony/berkovitzstatement051519>.

is the purpose of altering equilibrium trading outcomes in a continuous limit order book market design where orders are processed sequentially in the order that they arrive. Budish et al. (2015) argue that continuous limit order books result in an HFT arms race for traders to arrive at the exchange sooner than their competitors, whether it be by a millisecond or a microsecond, to profit from fleeting arbitrage opportunities. Therefore, many have suggested that speed bumps may offer a mechanism to disincentivize speed races and the associated costly investment in low latency infrastructure (Baldauf and Mollner, 2019, Harris, 2013).

Speed bumps differ in their implementation along two key dimensions. Firstly, the order processing delay can be applied either symmetrically on all orders, or asymmetrically on some orders but not others. All theoretical models and practitioner proposals and implementations of asymmetric speed bumps to date slow down liquidity demanders but not liquidity suppliers. Delay exempt order amendment can be conducted either by the individual liquidity provider, which benefits only those fast enough to amend their orders within the delay duration, or automatic exchange order repricing against a reference price, which benefits both fast and slow liquidity providers. Secondly, the delay duration can either be a fixed length of time or drawn from a stochastic random variable with a distribution that exhibits specified properties. Applying permutations, there are four potential speed bump designs: asymmetric with fixed duration, asymmetric with random duration, symmetric with fixed duration and symmetric with random duration.

### 3.6.1 Asymmetric Fixed Duration Speed Bumps

Asymmetric fixed duration speed bumps currently operate on IEX and NYSE American in the United States. These exchanges apply a 350-microsecond delay to all incoming orders from traders, but not the exchange's repricing of hidden pegged orders against external reference price movements. Hidden pegged orders account for 75-80% of trading on IEX.<sup>14</sup> IEX is often referred to as having a symmetric speed bump, which is only true for the minority of pre-trade transparent trades on its market, where both sides of the trade are slowed down. This implementation provides the exchange with a head start to update the reference price at which dark trades are matched, to mitigate the potential risk of extremely fast traders observing reference price changes earlier than the exchange and picking off stale quotes before they are repriced. Aquilina et al. (2016) investigate the causes and impacts of dark pool reference price latency arbitrage, and report that it affects 4% of dark pool trades in the UK. Similarly, dark pool reference price latency arbitrage has been found to affect 4% of dark trades in Canada (Anderson et al., 2016), 1% of dark trades in Australia (ASIC, 2015) and 2.6% of trades in a specific dark pool in the United States.<sup>15</sup>

Exogenous public information releases are not limited to reference prices for dark pools. Asymmetric fixed duration speed bumps that slow down liquidity demanders relative to liquidity suppliers can also be designed to mitigate speed races between these traders on any public information that is broad-

---

<sup>14</sup>IEX trading statistics are available at <https://iextrading.com/stats/>.

<sup>15</sup>While we are not aware of any empirical evidence on the market-wide magnitude of dark pool reference price latency issues in the United States, Citadel Securities was sanctioned by the SEC for operating algorithms that explicitly sought to benefit from this market friction, which were triggered when proprietary data feeds indicated a better price than the slower consolidated SIP data feed, to provide retail orders with the inferior SIP price. These algorithms executed approximately 2.6% of retail orders handled by Citadel. See <https://www.sec.gov/litigation/admin/2017/33-10280.pdf>.

cast to recipients at the same time. Budish et al. (2015) document that the duration of potential latency arbitrage opportunities between futures markets in Chicago and ETF markets in New York has declined over time but their frequency of occurrence is persistent. Chicago Stock Exchange has proposed a 350-microsecond fixed duration speed bump that would apply to all orders except those from specified market makers, to provide a head start in the race to cancel stale quotes and avoid adverse selection. The proposal was approved by SEC staff, subsequently put on hold by the Commission and then withdrawn by the applicant, highlighting the uncertainty and caution that practitioners and securities regulators maintain on the potential impacts of asymmetric speed bumps. More recently, the U.S. Commodity Futures Trading Commission (CFTC) has approved Intercontinental Exchange’s proposal for a 3-millisecond fixed duration speed bump on liquidity demanders for some of its gold and silver futures contracts.

Budish et al. (2015) argue that while asymmetric speed bumps can resolve latency arbitrage, they are unable to address the race for queue priority in continuous limit order books, which contrasts with their proposed frequent batch auction market design. Baldauf and Mollner (2019) also propose asymmetric speed bumps to slow down liquidity demanders, in order to reduce stale quote sniping. Aoyagi (2019) suggests that asymmetric speed bumps could increase speed investments by liquidity demanding HFTs, because they reduce the marginal cost of becoming faster. Analogous to systematically slowing down liquidity demanders, Han et al. (2014) develop a model in which some fast HFT liquidity providers can observe public information sooner and cancel orders to avoid adverse selection, allowing them to quote narrower bid-ask spreads. Adverse selection increases for slower liquidity providers, who then widen their bid-ask spreads, and overall liquidity may deteriorate unless it

is provided exclusively by fast HFT liquidity providers. Similarly, Hoffmann (2014) models a dynamic limit order market where fast traders are able to revise their orders to avoid adverse selection risk, but their presence results in slow traders submitting less competitive orders.

While we are not aware of any analysis that estimates the proportion of lit market trades affected by latency arbitrage, empirical studies find market quality improvements when its potential is reduced. Chakrabarty et al. (2014) examine the SEC’s naked access ban, which can be considered a “regulatory speed bump” on HFT liquidity demanders, but not other traders, finding that it leads to lower adverse selection costs for liquidity suppliers, which are impounded into narrower bid-ask spreads. Meanwhile, Chakrabarty et al. (2019) examine the industrial organization among stock exchange groups in the United States, finding that those that adopt or propose speed bumps or other mechanisms to disincentivize speed races increase their market share and profitability.

### **3.6.2 Asymmetric Random Duration Speed Bumps**

Similar to asymmetric fixed duration speed bumps, asymmetric speed bumps with a random delay duration could help liquidity suppliers avoid adverse selection in response to exogenous public information released to all traders at the same time. Additionally, drawing the random delay duration from a range of time periods could allow delay-exempt liquidity providers to harness endogenous information from order book activity on other trading venues, such as a simultaneous smart order router spray from a single large trader. As narrated in the excerpt of Michael Lewis’ popular culture novel ‘Flash Boys’ extracted at the start of this chapter, institutional traders found that they were unable

to access consolidated liquidity across trading venues if “news” of their trades on one venue would lead liquidity suppliers on other venues to immediately cancel their quotes. Practitioners verify the occurrence of order cancellations on one venue in response to trades on other venues, leading to a portion of consolidated liquidity being unavailable, which is a major friction in a fragmented trading environment.<sup>16</sup> Baldauf and Mollner (2019) acknowledge that asymmetric random duration speed bumps can reduce the accessibility of liquidity across venues. Meanwhile, European securities regulators have recently taken a keen interest in the accessibility of liquidity across fragmented trading venues, as a critical market quality issue (AFM, 2016, Aquilina and Ysusi, 2016, ESMA, 2016).

The academic literature also highlights the importance of order flow as a potential source of information in fast and fragmented markets. Foucault and Moinas (2018) note that a trade that depletes the top-of-book liquidity in a security on one trading venue is relevant information for pricing that security on other venues. Empirically, Malinova and Park (2017) find that multi-venue trades have around double the price impact of single-venue trades, even after controlling for trade size and trader type. Therefore, liquidity suppliers who are able to avoid multi-venue trades can reduce their losses to adverse selection. More broadly, O’Hara (2015) argues that in the new high frequency world, large traders are informed traders, at least to the extent that they are aware of their own trading intentions, which can cause market imbalances.

Asymmetric randomized duration speed bumps currently operate on TSX Alpha and Aequis NEO in Canada. TSX Alpha applies a delay on the continuous interval between 1 and 3 milliseconds to all incoming orders except a

---

<sup>16</sup>For example, see <http://www.nanex.net/aqck2/4661.html> from Nanex, a data services vendor.



subset of liquidity providing orders. In contrast, Aequitas NEO applies a delay on the continuous interval between 3 and 9 milliseconds to marketable orders from participants that they categorize as latency sensitive traders. Given the broader application of TSX Alpha’s speed bump, we focus on its design below. The magnitude of impacts on Aequitas NEO is likely to be significantly smaller because the delay applies only to a subset of, rather than all, liquidity demanders. Alpha’s speed bump design is expected to enable its liquidity suppliers to fade their quotes against order flow information from other trading venues, especially in anticipation of large incoming orders that are likely to impose adverse selection. As we reported in Table 3.2, approximately 60% of trading in Canadian equities is followed by instantaneous adverse selection after the trade. For single-venue trades this proportion is 47%, compared with 63% to 76% for multi-venue trades. Therefore, an asymmetric randomized speed bump could assist liquidity providers in avoiding a colossal proportion of adverse selection, which could lead to higher profitability and willingness to quote in a competitive equilibrium.

Brolley and Cimon (2018) model an asymmetric random duration speed bump as providing execution uncertainty for liquidity demanders. Their market design simply modifies a fixed duration speed bump by applying the delay with only a certain probability. In their model, the speed bump allows liquidity providers to avoid adverse selection arising from exogenous public information, as we have described above for fixed duration speed bumps. Their model does not account for endogenous order flow information arising from trades on other venues. In the spirit of avoiding multi-venue trades, Biais et al. (2015) develop a model of fast liquidity demanders, who can simultaneously access multiple trading venues, and slow liquidity demanders, who can only visit one venue at a time. These traders are equivalent to the fast multi-venue SOR traders

and slow single-venue traders, respectively, in van Kervel (2015). Their model, with proposed fast and slow trading venues for the respective trader types, is analogous to the setting with an asymmetric randomized speed bump, where the speed bump venue is “slow” and all other venues are “fast”. In this model, slow venues have zero bid-ask spreads because slow liquidity demanders do not have private information. However, as more slow traders migrate to the slow venue, average adverse selection and bid-ask spreads on the fast venue increase, imposing negative externalities.

### **3.6.3 Symmetric Fixed Duration Speed Bumps**

Symmetric fixed duration speed bumps that apply to all orders do not change the execution outcome of individual orders after they have been received by the trading venue, as the sequence of order processing is not altered. However, this market design slows down the dissemination of market data after each order is processed, which is equivalent to adding a fixed delay to a trading venue’s aggregate order processing and data dissemination latency. While we have explained above that IEX operates an asymmetric fixed duration speed bump from the perspective its dark trades, it can be considered a symmetric fixed duration speed bump from the perspective of traders on other venues. This is because IEX delays both incoming orders from traders and outbound market data dissemination to proprietary data feeds by 350 microseconds, but not data dissemination to the public Securities Information Processor (SIP). Ding et al. (2014) find that the SIP is slower than proprietary data feeds by one to two milliseconds, far exceeding IEX’s 350-microsecond speed bump.

Applying this interpretation, Hu (2018) examines the second order impacts of IEX’s migration from dark pool to stock exchange and finds an improvement

in consolidated market quality. The operation of IEX’s speed bump from the perspective of individual trades did not change when it converted from a dark pool to a stock exchange. However, in contrast with a dark pool, a stock exchange’s quotes contribute to the consolidated SIP data feeds and are protected under trade-through prohibition rules. Hu (2018) argues that the channel via which market quality improves is that IEX’s change from dark pool to exchange slows down the SIP infrastructure, which is analogous to an exchange implementing a symmetric fixed duration speed bump as we describe above. Similarly, Chakrabarty et al. (2019) document slower trader reactions to order flow information after IEX attained exchange status. The fourth chapter of this dissertation explores stock exchange trading platform speed in more detail and investigates the impact of an increase in order processing and data dissemination speed on the Toronto Stock Exchange in 2014, which is analogous to removing a symmetric fixed duration speed bump. That chapter finds that in a fragmented trading environment, faster order processing and data dissemination on one venue can contribute to higher fleeting liquidity on other venues, which also enables lower adverse selection for liquidity suppliers who successfully fade their quotes.

### **3.6.4 Symmetric Random Duration Speed Bumps**

Speed bumps can also be designed to delay processing of all orders by a randomized duration. This design has been adopted by several major foreign exchange trading venues, including Thomson Reuters and EBS, which apply a randomized delay of several milliseconds before incoming orders are processed. Harris (2013) proposes brief random duration delays of between 0 and 10 milliseconds for all order instructions, to introduce variability into order matching

outcomes, which reduces the deterministic nature of the fastest trader always winning a speed race. Hoffmann (2014) also suggests that these small speed bumps could be an appropriate policy response to reduce the advantage of fast traders and discourage the arms race for speed.

In contrast, Budish et al. (2015) argue that symmetric fixed duration speed bumps are ineffective for resolving the market friction of speed races to trade on public information that is released to all traders at the same time, because for each stale quote there is a single liquidity supplier but many liquidity demanders. Therefore, randomizing the sequence of order processing does not increase the probability of the liquidity supplier being able to cancel the stale quote, in contrast with a model setup with one liquidity supplier and only one liquidity demander. Additionally, symmetric random duration speed bumps potentially incentivize traders to submit redundant orders, as each order submission essentially represents a lottery ticket for receiving a relatively short delay.

## **3.7 Appendix II: Novel Metric Development**

### **3.7.1 Trade Direction Classification**

Empirical measures of market quality fall into two categories. Order-based metrics, such as quoted spreads, quoted depths and midpoint return volatility, assess the state of the displayed limit order books observed throughout time. Trade-based metrics, such as effective spreads, realized spreads and adverse selection, measure the actual costs of demanding liquidity and the market impact when participants decided to trade. Trade-based metrics require the accurate assignment of either buyer- or seller- initiated trade direction. Other

than studies with access to proprietary datasets from securities regulators or exchange operators containing flags for buyer- or seller- initiation, algorithms based on Lee and Ready (1991) and Ellis et al. (2000) have been utilized to infer trade direction. Numerous studies have analyzed the accuracy of these approaches, including Bessembinder (2003) and Holden and Jacobsen (2014), with the best levels of accuracy estimated at between 70-80%.

We develop a methodology that assigns trade direction correctly with virtually full accuracy using the Thomson Reuters Tick History (TRTH) academic research database provided by the Securities Industry Research Centre of Asia-Pacific (SIRCA). Although the messages in this data source exhibit delays (latency) and variation in these delays (jitter), the ordinal sequencing of messages within a venue is preserved, allowing us to obtain an accurate snapshot of the state of the limit order book at the time of each trade. The preservation of ordinal message sequencing is verified by consistent observation of an order book update immediately after each trade, reflecting the consumption of liquidity at the displayed price. Our approach may be applied to all datasets where order and trade messages arrive sequentially in the same file, including Canadian, UK, Scandinavian and Japanese equity markets in TRTH.<sup>17</sup>

From the TRTH database, we concurrently request order book snapshots and trade messages. This data contains fields for security, date, timestamp to the millisecond, trade price, trade volume, trade qualifier, bid price, bid size, ask price and ask size. The table below provides an example for Royal Bank of Canada (RY) trading on the Toronto Stock Exchange (.TO) on 1 September 2015. We assign a message counter per venue-stock-day.

We compare each trade with the order book update from the most recent

---

<sup>17</sup>Unfortunately this methodology cannot be applied to TAQ data, because the trades are collated by separate services (CTS for trades and CQS for quotes). Holden and Jacobsen (2014) provide further details on TAQ.

Table 3.12: **Trade Initiation Direction Example 1**

ID	Time	Price	Volume	Bid Price	Bid Size	Ask Price	Ask Size
1	09:30:08.497			71.98	1000	72.00	3100
2	09:30:08.507	71.98	1000				
3	09:30:08.507			71.97	1100	72.00	3100
4	09:30:08.507			71.97	1100	71.98	200
5	09:30:08.507	71.98	200				
6	09:30:08.507			71.97	1100	72.00	3100
7	09:30:08.656			71.97	1100	71.99	200

message ID. We assign trades that occurred at or above the prevailing ask price as buyer initiated, and those at or below the prevailing bid price as seller initiated. For example, the trade in message ID 2 for 1000 shares at \$71.98 in the above table would be assigned as seller-initiated, since it occurred at the prevailing bid price. The trade in message ID 5 for 200 shares at \$71.98 would be assigned as buyer initiated, since it occurred at the prevailing ask price. Note that these trades occur at the same price and millisecond timestamp but are in opposite directions. We are confident in the accuracy of our methodology, since both trades visibly consume the liquidity available in the limit order book.

Occasionally, trades may happen outside the prevailing best bid or offer prices. These trades are direct action orders in Canada, which are similar to intermarket sweep orders in the United States. An example for Royal Bank of Canada (RY) trading on the Toronto Stock Exchange (.TO) on 16 November 2015 is provided in the table below, where a buyer-initiated trade consumed 300 shares of liquidity at the best offer price of \$74.00 and continued at execute 200 shares at \$74.01. Note that there are no order book update messages between each trade message since trades are recorded from the perspective of

the liquidity supplier, and the liquidity demander here entered a marketable buy order for 500 shares with a limit price equal to or higher than \$74.01.

Table 3.13: **Trade Initiation Direction Example 2**

ID	Time	Price	Volume	Bid Price	Bid Size	Ask Price	Ask Size
1	09:40:49.497			73.98	600	74.00	300
2	09:40:49.497	74.00	200				
3	09:40:49.497	74.00	100				
4	09:40:49.497	74.01	100				
5	09:40:49.497	74.01	100				
7	09:40:49.497			73.98	600	74.01	400

All trades with trade reporting, intentional broker crossing or odd-lot qualifiers are removed as they do not interact with the limit order book. We also discard trades that happen between the prevailing bid and offer prices, since these trades interacted with undisplayed limit orders and we are unable to match the trade to a corresponding quote to assign trade direction.

### 3.7.2 Timestamp Synchronization Benchmarking

A defining characteristic of modern low latency trading is that market events occur in microseconds or milliseconds, rather than in seconds or minutes. O'Hara (2015) argues that cross-market linkages are one of two critical issues in modern market microstructure research and regulation. In order to examine the dynamics of high frequency quoting and trading across multiple venues, empirical researchers need to either demonstrate that timestamps across venues are consistently synchronized with de minimis delays (latency), variation in delays (jitter) and batching of multiple message arrivals due to system capacity constraints (caching), or adopt a research design that ac-

counts for known magnitudes of clock drift between venues. To the best of our knowledge, we are the first empirical researchers to explicitly address the issue of clock synchronization in high frequency market data and propose a methodology for robust analysis. This methodology for benchmarking clock synchronization could be used in any multi-venue market with trade-through prohibitions, including TAQ data in the United States.

Several studies have examined high frequency quoting and trading across multiple markets, utilizing various timestamp synchronization assumptions. In their respective research designs, the implicit assumptions are that cross-venue clock synchronization is within one millisecond in Conrad et al. (2015) and Budish et al. (2015) using daily TAQ, within 5 milliseconds in Malinova and Park (2017) using regulatory data from the Investment Industry Regulatory Organization of Canada (IIROC) and within 100 milliseconds in van Kervel (2015) using TRTH data.

Assessing clock synchronization is necessary even where proprietary data is sourced from either the matching engines of individual trading venues or consolidated regulatory tape. Under IIROC's Universal Market Integrity Rules, marketplace participants shall not intentionally submit orders that would result in "locked" or "crossed" markets (IIROC, 2011). Locked (crossed) markets occur when the bid price on one venue is equal to (higher than) the offer price on another venue. This rule facilitates best execution outcomes, since it prevents an order that would be immediately executable at one venue from being routed to another venue where it is not immediately executable. Similarly, locked and crossed markets are prohibited under Rule 610 of Regulation NMS in the United States.<sup>18</sup> Utilizing this feature of the market structure, our clock synchronization benchmarking approach examines the frequency distribution

---

<sup>18</sup>See <https://www.sec.gov/rules/final/34-51808.pdf>.



of durations for which markets are locked and crossed. Most short duration observations of locked and crossed markets are likely to be driven by non-synchronicity in the timestamps.

Figure 3.8 presents the distribution of the duration of locked or crossed markets in our sample. As the entry of an order that would lock or cross the market is prohibited in Canada (and the US), the observation of any such period is likely driven by non-synchronicity in the timestamps. We use this feature to characterize an upper bound for observed asynchronicity in our data, where 30 milliseconds corresponds to the 95th-percentile in the distribution. Therefore, events that occurred simultaneously across our venues will be stamped within 30 milliseconds of each other at least 95% of the time. This interval is consistent with IIROC’s Guidance on Time Synchronization, which permits clocks to drift up to +/- 50 milliseconds from Coordinated Universal Time for both marketplaces and participants (IIROC, 2016), as well as clock synchronization requirements in the United States for recording events in NMS securities.<sup>19</sup>

We define a locked market of 0 milliseconds as occurring when quote updates were observed on multiple venues within the same millisecond, and a locked market would have occurred unless the quotes on the multiple venues were all updated on that millisecond. This is typically observed when multiple venues have quoted spreads equal to the minimum tick size both before and after a price change.

### 3.7.3 Trade String Construction

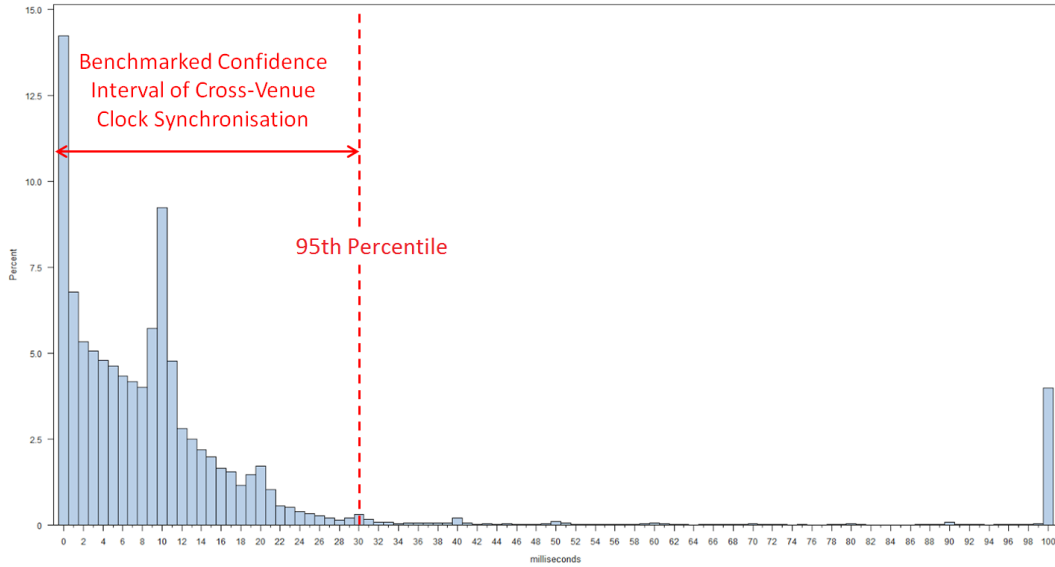
It is well documented that brokers utilize smart order routers to simultaneously access fragmented pools of liquidity across multiple venues and fill large

---

<sup>19</sup>See <https://www.sec.gov/rules/sro/finra/2016/34-77565.pdf>.

Figure 3.8: **Duration of Locked and Crossed Markets**

This figure presents a histogram of the duration of periods of locked and crossed markets using potentially asynchronous timestamps across venues. The 95th percentile of locked and crossed durations is marked with a dashed line at 30 milliseconds.



orders (for example, O’Hara and Ye (2011), van Kervel (2015) and Malinova and Park (2017)). In the previous Appendix section we established that 30 milliseconds is the maximum level of timestamp asynchronicity in our dataset, at a 95% confidence level. We construct high frequency trade strings to group together these related trades by concatenating all trades in the same direction across all venues within 30 milliseconds of another trade in that direction. Whilst trades occurring within 30 milliseconds may not be from the same smart order router spray, or even the same trader, traders acting on correlated information face the same available pool of liquidity. The signal such trading will generate, as observed by a liquidity supplier, is identical – rapid liquidity consumption across multiple venues, with a high likelihood of impending adverse selection caused by the depletion of displayed liquidity. Hence our metric will still capture the total amount of displayed liquidity that traders were able

to access in aggregate.

After constructing trade strings, we develop tools to assess the proportion of depleting trades, which measures the instantaneous information arriving in the order flow, and the accessibility of displayed liquidity across venues. Within each trade string, we are not certain of the actual sequencing of each trade, but we are confident that each trade occurred within the 30 milliseconds calculated in the previous Appendix section. We are also confident that we capture the signal of all trades in the same direction during the time interval of the trade string, although the relatively long duration with which we join trades also results in capturing a small amount of noise from unrelated trading activity.

We are aware of three studies that have attempted to identify similar high frequency trade strings:

1. van Kervel (2015) analyzes order book changes across venues around the same time as trades by aggregating trades within fixed 100 millisecond buckets (that are deemed to have occurred “simultaneously”) and taking snapshots of order books at 100 millisecond intervals. Although we use the same TRTH data source, our methodology has three advantages over van Kervel (2015):
  - (a) We avoid attributing limit order book changes before each trade to the trade, spanning half the fixed 100 millisecond interval length on average, minimizing the potential reverse causality associated with trades that occur as a result of order cancellations. Bessembinder (2003) finds that trades tend to occur immediately after order book cancellations in the opposite direction.
  - (b) We are able to group together trades that would otherwise span two

separate 100ms buckets, for example trades at 99 milliseconds and 101 milliseconds would be considered unrelated in the van Kervel (2015) method.

(c) We are able to significantly reduce the observation period span from a fixed 100 milliseconds bucket to a dynamic median that includes 1 millisecond before the trade, a typical trade string length of a few milliseconds and 30 milliseconds after the trade string. These 30 milliseconds are necessary to allow order books on all venues to update and reflect the information of the trade.

2. Malinova and Park (2017) use regulatory data from IIROC to reconstruct multi-venue trades from the same trader. They join all trades across multiple venues that are in the same direction and separated by less than five milliseconds, finding a higher incidence of liquidity fade and also greater quote sniping after these multi-venue trades, compared with single-venue trades. Although we use a data source with significantly noisier timestamps, our methodology allows the construction of trade strings with more broadly available data, and our trade string lengths are similar to theirs.
3. Whilst not explicitly specified as potential smart order router sprays, Conrad et al. (2015) use daily TAQ data to examine liquidity drawdowns containing trades with identical millisecond timestamps across multiple trading venues. They find that these liquidity drawdowns incur significant implicit transaction costs but have no lasting impact on price efficiency.

The table below presents a trade string that was constructed on 21 August 2015 for stock Royal Bank of Canada (RY), with trades across TSX, Alpha,

Chi-X and CX2. The 12 seller-initiated trades for 1,800 shares at \$73.55 were recorded within 6 milliseconds of each other and would have resembled a smart order router spray from the perspective of the liquidity suppliers. The trade string is deemed to be depleting as no bids at prices equal to or higher than \$73.55 were available on any venue immediately after the end of the trade string. The volumes traded on each venue are compared with the displayed liquidity at the \$73.55 bid price immediately prior to the start of the trade string. Quote fade is the volume of starting displayed liquidity in excess of the traded volume.

Table 3.14: **Trade String Example**

Exchange	Trade ID	Timestamp	Price	Volume	Initiator
TSX	1	09:50:36.247	73.55	500	Seller
TSX	2	09:50:36.247	73.55	100	Seller
TSX	3	09:50:36.253	73.55	100	Seller
TSX	4	09:50:36.253	73.55	200	Seller
Alpha	1	09:50:36.253	73.55	200	Seller
Chi-X	1	09:50:36.250	73.55	100	Seller
Chi-X	2	09:50:36.250	73.55	100	Seller
CX2	1	09:50:36.253	73.55	100	Seller
CX2	2	09:50:36.253	73.55	100	Seller
CX2	3	09:50:36.253	73.55	100	Seller
CX2	4	09:50:36.253	73.55	100	Seller
CX2	5	09:50:36.253	73.55	100	Seller

While our definition of trade strings theoretically allows day-long trade strings (as long as each consecutive trade follows less than 30 ms after the previous trade), in practice the duration of trade strings is very short. The median length for multi-venue trade strings is between 9 and 17 milliseconds.

Even at the 90th percentile, durations increase modestly to between 21 and 36 milliseconds, depending on the number of venues accessed.


Overall, the median length of a trade string executed across multiple venues is 11 milliseconds in our sample, which is comparable to the analysis using regulatory data from IIROC in Malinova and Park (2017). They group trades originating from a unique trader conditional on being separated by less than 5 milliseconds between each trade. These intervals are significantly smaller than the 100 millisecond snapshots taken by van Kervel (2015), and more consistent with the time horizons in which high frequency traders are known to operate, as documented in Hasbrouck and Saar (2013).

For each trade string, we take a snapshot of the limit order book across each venue 1 millisecond before the start of the first trade, since order book updates are produced to show trades consuming liquidity. We also snapshot the limit order books across all venues 30 milliseconds after the end of the last trade, to allow sufficient time for venues with potentially slower clocks to update their order books to reflect the information of the last trade. Notably, this snapshot does not overlap into the previous or the next trade string for the same stock. Buyer-initiated trade strings are compared with changes in the offer prices and sizes, while seller-initiated trade strings are compared with changes in the bid prices and sizes, on each venue. For trades that occurred at the best price within each string (generally the prevailing NBBO price at the start of the string) we record the trade price, start time and end time, as well as recording the trade volume, start price, start volume, end price and end volume on each trading venue. Only trades occurring at the best prices within each string are analyzed, to enable trade attribution to the consumption of visible liquidity at each venue's best bid or offer price. Figure 3.9 provides an example of the logic applied to constructing trade strings for the purpose of

our metrics.

**Figure 3.9: Example of a Depleting Trade String Construction**

This figure depicts an example for the construction of a trade string that depleted all available depth at the NBO, which used to examine quote fade. The depletion could be driven by both trade executions and order cancellations. Trades in a stock within 30 milliseconds of each other are grouped into the same string. At least 30 milliseconds of no trading separates each trade string. A snapshot of the order book is taken 1 millisecond prior to the first trade, with the depth at NBBO across all order books recorded.



## Chapter 4

# Trading Engine Speed, Fleeting Liquidity and Competition for Order Flow

### 4.1 Introduction

Stock exchanges around the world periodically upgrade their core trading engine infrastructure to reduce the amount of time taken for each step in the trading process to be completed, which is referred to as the exchange's latency. It includes the duration taken between a broker's order message arriving at an exchange's entry point, transmission to the gateways, conducting pre-trade risk checks, submission to the matching engine, generation of an execution report, return of the order confirmation message to the gateway and delivery to the submitting broker (Kirilenko and Lamacie, 2015). At the end of this process, order and trade information is also disseminated to other traders via market data feeds.

As new hardware and software technologies have been developed, the or-



der of magnitude of trading platform latencies has declined from seconds to milliseconds and then to microseconds (Pagnotta and Philippon, 2018). Stock exchanges promote ever lower latencies as technological advancements that reduce frictions in trading. Intuitively, faster trading speeds could benefit algorithmic traders, who can utilize the minuscule incremental speed increases that are not perceptible to human traders, to cancel stale quotes or submit new quotes at the front of the order book queue (Baldauf and Mollner, 2018).

This chapter investigates the 2014 Toronto Stock Exchange (TSX) Quantum XA trading platform upgrade, which reduces latency from 2.3 milliseconds to 26 microseconds. We examine the impacts on aggregate market-wide liquidity, as well as TSX’s competitiveness in attracting orders and trades. O’Hara (2015) asserts that market linkages and fairness are two especially important focus areas for market microstructure policy and research. In this spirit, this chapter also examines whether the upgrade impacted these dimensions of market quality, which we proxy by the accessibility of liquidity across fragmented Canadian equities trading venues and the potential presence of stale quote sniping.

The empirical literature meticulously examines the impact of electronic trading upgrades that reduce order processing latency on market efficiency across a wide range of stock exchanges. Electronic trading platform latency was initially reduced from seconds to milliseconds, with mixed findings. Hendershott and Moulton (2011) study the New York Stock Exchange (NYSE) Hybrid Market, which increases automation and reduces market order execution latency from 10 seconds to less than 1 second. Bid-ask spreads widen due to higher adverse selection costs, while price discovery efficiency increases. Conrad et al. (2015) analyze the Tokyo Stock Exchange’s introduction of colocation services and concurrent Arrowhead trading platform upgrade, which

reduces order receipt to processing latency from 1 to 2 seconds to less than 10 milliseconds. Faster order processing facilitates increased high frequency quoting activity, resulting in narrower effective spreads and improved price discovery. Jain et al. (2016) investigate the same upgrade, finding a reduction in bid-ask spreads and return volatility, as well as negative impacts on price discovery, fleeting trades and quote-to-trade ratios.

Empirical studies generally concur that trading platform latency reductions to the tens or hundreds of milliseconds improve market efficiency. McInish and Upson (2013) examine an upgrade to NYSE's computational speed from around 625 milliseconds to 75 milliseconds, finding a substantial decline in trading costs for fast liquidity demanders, in contrast with a marginal decline for slow liquidity demanders. Riordan and Storkenmaier (2012) analyze enhancements to Deutsche Boerse's Xetra trading platform, which reduces latency from 50 milliseconds to 10 milliseconds, resulting in lower adverse selection and bid-ask spreads. The contribution of quotes to price discovery increases, consistent with liquidity suppliers utilizing the faster order processing speeds to more efficiently update their quotes. Murray et al. (2016) document two separate information dissemination latency reductions on the Australian Securities Exchange, firstly from 70 to 30 milliseconds, and then to 300 microseconds. The first upgrade narrows bid-ask spreads, but there is no further reduction following the second.

Consistent with this result, other studies on the impact of exchange latency reductions to single digit milliseconds or hundreds of microseconds also yield null or mixed results. Brogaard et al. (2014a) utilize sequential upgrades of the London Stock Exchange's TradElect trading system, which reduce latency from 11 milliseconds to 3 milliseconds, as exogenous shocks that increase HFT participation, to demonstrate that HFT activity does not impact institutional

investor transaction costs. Chakrabarty et al. (2017) examine the introduction of the Spanish Stock Exchange’s SIBE-Smart trading platform, which reduces latency from around 7 milliseconds to less than 1 millisecond, attracting HFT activity, increasing effective spreads, reducing market depth and harming price discovery. However, deteriorating market quality is partially attributed to uncertainties surrounding short selling regulations proposed around the same time as the upgrade’s implementation. Ye et al. (2013) analyze a NASDAQ technology enhancement that reduces the latency of trade data dissemination from 3 milliseconds to 1 millisecond, finding no change in bid-ask spreads or price efficiency, but higher short-term volatility and order-to-trade ratios.

To the best of our knowledge, this chapter conducts the first empirical analysis on the impact of a stock exchange’s speed upgrade on order flow competition and market linkages in a fragmented trading environment. We make several contributions to the literature. Firstly, we develop a stylized schematic of multi-venue trading to show that faster order processing on one venue can increase fleeting liquidity and quote sniping on other venues. This results in additional market frictions for large traders who require access to consolidated liquidity across fragmented trading venues. Secondly, we empirically validate the schematic’s predictions and find that it alters the distribution of adverse selection across trading venues. Thirdly, we find that the upgrade enhances the venue’s attractiveness to traders, immediately increasing its market share. Finally, we extend the literature on trading platform speed upgrades to double digit microseconds, which represents the magnitude of order processing speeds among the fastest stock exchanges today. This time scale is similar to the 31-microsecond average NYSE-to-trader message transmission time that Baldauf and Mollner (2019) measure from a proprietary dataset. Overall, this chapter documents externalities from stock exchange speed competition and

contributes to the growing literature on the high frequency trading arms race.

The rest of this chapter is organized as follows. Section 4.2 outlines the purpose and impact of faster exchange trading infrastructure in consolidated and fragmented trading environments. Section 4.3 discusses the institutional details of Canadian equities trading and the TSX Quantum XA upgrade. Section 4.4 outlines the data and empirical methodology. Section 4.5 presents and discusses the results. Section 4.6 concludes. Section 4.7 provides supplementary materials.

## **4.2 Trading Infrastructure Speed**

### **4.2.1 Faster Trading Platforms in a Consolidated Market**

We start by reviewing the literature on the components of trading speed in a consolidated equities trading environment. Menkveld and Zoican (2017) develop a theoretical model of exchange trading speed and market liquidity, finding two opposing effects. Faster trading enables market makers to update quotes more frequently in response to new information, which reduces bid-ask spreads. However, market maker quotes are more likely to be picked off by predatory snipers who can also act faster on the new information, which increases adverse selection. The relative ratio of news to liquidity traders determines the overall impact. Conspicuously, Menkveld and Zoican (2017) model the trading process in discrete time, where trader arrival and order processing occur at specified rates, akin to a batch auction market. Faster trading speed is modeled as a shorter batching interval, whereby traders can visit the exchange and update their orders more frequently. In this setup, trading speed

consists of the aggregate time taken for trader decision making, order routing and exchange order processing. Foucault and Moinas (2018) survey the fast trading literature and disentangle trading speeds into informational speed and matching speed, which we now turn to discuss in detail.

It is well documented that high frequency traders (HFTs) compete on speed to observe, process and respond to new information, to capture fleeting trading opportunities. Menkveld (2013) estimates an upper bound of 1 millisecond for an identified HFT's response time. Menkveld (2018) finds that around twenty percent of trades occur in sub-millisecond clusters and these episodes impose high adverse selection costs on non-HFTs. Clustering of trades in calendar time is indicative of speed races between traders. Baron et al. (2019) examine speed competition among HFTs, finding that faster HFTs earn higher revenues. Baldauf and Mollner (2018) theoretically model and empirically verify that faster traders avoid adverse selection by canceling stale quotes, while being more likely to win the race for order book queue priority at new price levels. Kirilenko et al. (2017) find that HFTs frequently consume the last units of available liquidity before price changes, which is indicative of high speed opportunistic quote sniping activity.

After determining their trading intentions, traders need to transmit order messages to a trading venue, which is affected by the speed of their communication technology. Laughlin et al. (2014) document advances in high speed transmission networks between the ETF market in New York and the futures market in Chicago, initially with more efficient fiber optic cables, and subsequently microwave connectivity. Shkilko and Sokolov (2016) examine episodes of precipitation that disrupt microwave transmission, equalizing all traders to the same speed level on fiber, which reduces adverse selection and improves liquidity.

A major revenue stream for stock exchanges is technical connectivity, with faster speeds often available for higher subscription fees. Exchanges have developed co-location services that enable traders to place their computers in the same data center as its matching engine, to ensure deterministically low latency for traders to access the market. Brogaard et al. (2015) examine the introduction of an optional co-location upgrade on NASDAQ OMX Stockholm and find an increase in liquidity. They suggest that liquidity supplying HFTs subscribe to the fastest connection, which improves their management of adverse selection risk, and enables quoting narrower spreads.

In contrast with co-location, which stock exchanges offer as an optional subscription service to increase connectivity speed for subscribers relative to non-subscribers, upgrades to exchange trading platforms increase the speed at which orders are processed for all traders. By preserving the sequence of order processing, trading platform upgrades do not affect the outcome of individual order matching after they are received by the exchange. Consistent with this institutional detail, Budish et al. (2015) describe continuous limit order books as processing orders serially upon arrival, resulting in a speed race among traders to be the first to arrive at the exchange. Kirilenko and Lamacie (2015) collect statistics on order processing latency within the Brazilian BM&FBOVESPA Exchange and find that the round-trip time clusters around a median of one millisecond but has a large right tail. They assert that its variation has explanatory power for volatility in the price discovery process.

Importantly, Kirilenko and Lamacie (2015) affirm that trading platform latency is exogenous to and beyond the influence of individual traders. Therefore, it is not immediately clear how exchange latency can impact market efficiency in a consolidated trading environment. Foucault and Moinas (2018) suggest that unless traders apply extremely high discount rates, faster re-

alization of gains from trade does not materially increase welfare given the miniscule time scales involved. This assertion is supported by the empirical literature, which finds that exchange speed upgrades enhance market efficiency up to a threshold in the double-digit milliseconds, after which improvements are exhausted (Murray et al., 2016).

Foucault and Moinas (2018) also observe that exchange latencies for order processing and information dissemination are often intertwined. Indeed, most empirical studies we survey on exchange upgrades report latency as the aggregate of these two components. Therefore, trading platform speed may impact broader market efficiency if HFTs can utilize faster message dissemination. Foucault et al. (2013) model make and take cycles in which liquidity is supplied and then consumed, after which the process repeats. In this spirit, faster data dissemination alerts traders to the completion of each cycle sooner, increasing the intensity of liquidity supply and consumption and the rate at which gains from trade are realized. Goldstein et al. (2018) examine the introduction of the Australian Securities Exchange’s premium ITCH connectivity protocol, which provides faster access to market data for a higher monthly subscription fee. They find that HFTs are subsequently more successful at predicting price movements from order book imbalances and attaining queue priority, resulting in lower limit order execution probability for non-HFTs.

In this chapter, we examine the impact of the TSX Quantum XA trading platform upgrade on overall liquidity and order book replenishment. Following the predictions and findings in the above literature, we hypothesize no change in average liquidity throughout the trading day. However, we expect faster order book replenishment after quotes at a price level are consumed, as well as increased ability of HFTs to win the race for queue priority when new price levels are established on the TSX.

### 4.2.2 Faster Trading Platforms in a Fragmented Market

Several theoretical models examine speed competition among stock exchanges to attract order flow in a fragmented trading environment. In these models, traders decide on which one of the multiple exchanges to send their orders to. Pagnotta and Philippon (2018) develop a theoretical model to examine trading speed, market fragmentation and trading regulation. Faster trading venues reduce search costs, which attracts latency sensitive traders and enables higher trading fees. They argue that some trading venues could be too slow, so regulated minimum trading speeds could increase competition with faster trading venues and improve trader welfare.

Wang (2018) develops a model of speed competition among stock exchanges. A faster exchange attracts new price-improving limit orders, because lower exchange latency results in liquidity demanders being more likely to be aware that the new orders are available on that exchange. If traders send more orders to that venue, its market share increases. The model also predicts that faster exchange speed will result in unchanged or narrower bid-ask spreads. Wang (2018)'s empirical analysis finds that the IEX stock exchange in the United States, which has inbound and outbound speed bumps of 350 microseconds to slow down order processing and data dissemination, has much lower market share than the non-delayed trading venues. Similarly, Hu (2018) examines IEX's upgrade from dark pool to stock exchange and finds improvements in consolidated liquidity and price discovery. In contrast with dark pools, displayed quotes on stock exchanges contribute to the consolidated Securities Information Processor (SIP) data feeds, and the improvements in market efficiency are attributed to IEX's delayed order processing slowing down the SIP data feeds.



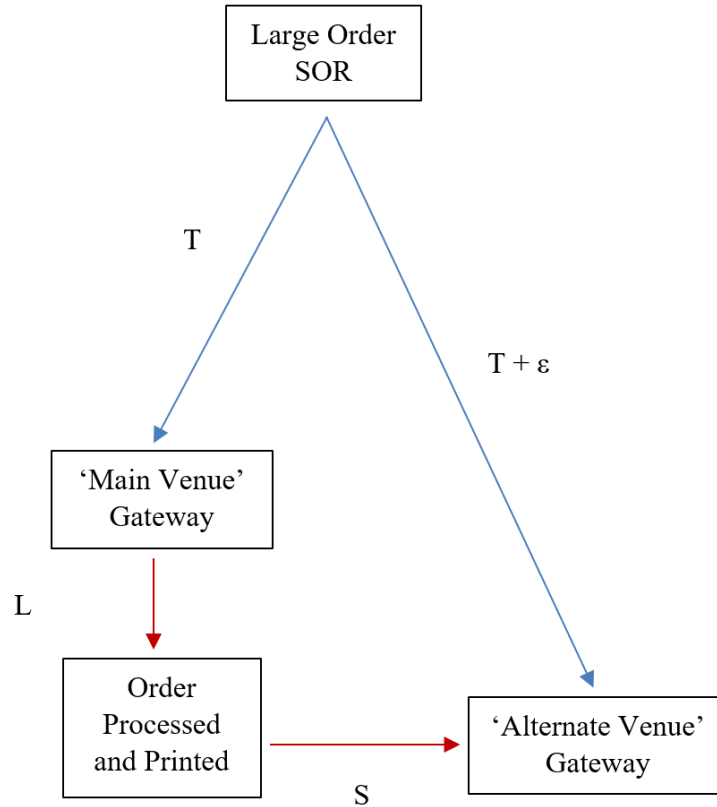
Lee (2019) develops a model of transmission latency between exchanges and competition among informed traders in a fragmented trading environment. Shorter transmission latencies result in order flow information on one venue being incorporated faster into prices on other venues, increasing informed trader competition and price discovery, but reducing informed trader profits. Additionally, the price impact of a venue's own order flow declines while price impact from other venue's order flow increases. Therefore, a latency reduction on one venue can have spillovers onto other venues in the form of higher price impact and reduced liquidity.

We extend the paradigm in Lee (2019) to consider the potential externalities of an exchange's increased order processing speed on other exchanges. In this high frequency environment, informed traders are likely to use smart order routing (SOR) technology to simultaneously spray multiple trading venues and consume all available consolidated liquidity at a given price level, to maximize the profits from their short-lived information. Figure 4.1 presents a stylized illustration of this process.

Suppose that a trader's broker wishes to buy a quantity of  $2X$  shares, and that there are  $X$  shares available at the same best offer price on each of the 'main venue' and the 'alternate venue'. To immediately buy  $2X$  shares at the best offer price, the broker will use an SOR to simultaneously spray both venues with child orders. However, in continuous time, events do not occur exactly simultaneously due to stochastic factors such as computational and transmission delays. Let  $T$  be the duration taken for a child order to reach the main venue exchange gateway and  $T + \epsilon$  be the duration taken for the other child order to reach the alternate venue exchange gateway. Let  $L$  be the duration taken by the main venue's matching engine to process the incoming order and disseminate the resulting order book update or trade via market data

Figure 4.1: **Schematic Diagram of Trading Platform Latency and Multi-Market Run Games**

This figure presents a stylized schematic diagram that outlines potential multi-market run games with outcomes that depend on order processing and data dissemination latency, communication latency between trading venues and smart order router (SOR) timing precision. A broker seeking to execute a large order immediately may use an SOR to simultaneously spray both the main venue and alternate venue. In continuous time, events do not occur “exactly” simultaneously. Assume that the order takes  $T$  units of time to reach the main venue and  $T + \epsilon$  units of time to reach the alternate venue. The main venue takes  $L$  units of time to process the order and disseminate the trade message to market data feeds. A co-located HFT observes this trade print and then takes  $S$  units of time to send either an aggressive marketable order or limit order cancellation to the alternate venue. If  $\epsilon > L + S$ , the initial investor broker SOR will be unable to access the observed liquidity on the alternate venue, due to information leakage of their trade on the main venue.



feeds. Let  $S$  be the duration taken by a co-located HFT to observe the market data feed update, process this information and send either a marketable order or limit order cancellation to the alternative venue gateway.  $S$  is affected by the time taken to transmit messages between geographically separated trading venues, as modeled in Lee (2019). In this stylized example, if the child orders do not arrive within close succession, or more formally if  $\epsilon > L + S$ , the HFT will win the speed race to have its order message arrive at the alternative venue's gateway first, resulting in opportunity costs from missed trades for the original trader.

In a fragmented trading environment, the trading platform speed of one venue can increase market frictions for large traders who need to simultaneously spray multiple venues to trade their desired volumes, if they do not invest in costly technologies to reduce  $\epsilon$  to be below this upper bound. Some practitioners have developed SOR systems that synchronize order routing across geographically dispersed trading venues to ensure that child orders arrive virtually simultaneously.<sup>1</sup> Consistent with this notion, van Kervel (2015) finds that fleeting liquidity across fragmented venues declines when more traders use SOR technology.

This chapter examines the TSX Quantum XA trading platform upgrade, which reduces exchange latency  $L$  from approximately 2,300 microseconds to 26 microseconds. Communication latency  $S$  has a lower bound of 100 microseconds if the alternate venue is Chi-X or CX2, based on speed of light constraints to traverse the geographical distance between the data centers. It is likely to be de minimus if the alternate venue is Alpha, which is located in the same data center as TSX. Therefore, the exchange upgrade provides an

---

<sup>1</sup>An example of this technology is Royal Bank of Canada's THOR SOR, see <http://www.rbc.com/newsroom/news/2016/20161027-cm-thor.html>

exogenous shock that immensely reduces the upper bound for required SOR precision  $\epsilon$ . Similar to the model in Lee (2019), the impact of a trading venue speed upgrade is asymmetric. While message transmission between the upgraded venue and other venues becomes faster, the speed of transmission in the reverse direction does not change.

Liquidity providers are incentivized to cancel their quotes when confronted by multi-venue SOR sprays because these trades are more likely to be informed and impose adverse selection. Malinova and Park (2017) find that multi-venue trades have around double the price impact of single-venue trades, even after controlling for trade size and trader type. Similarly, O’Hara (2015) remarks that in the new high frequency world, large traders are informed traders, at least because they know their own trading intentions. Additionally, Foucault and Moinas (2018) suggest that a trade that consumes all liquidity at a price level in a stock on one venue is information that is relevant to pricing that stock on other venues.

Closely related to this chapter is Malinova and Park (2017), which examines the geographical relocation of a Canadian equities trading venue to be in the same data center as the largest TSX venue. Previously the physical latency between these two venues was around 400 microseconds. After communication latency is virtually eliminated, liquidity fade and quote sniping decline, due to more accurate order synchronization by traders who utilize multi-venue trades. There are fewer opportunities for HFTs to intersperse their orders among the orders of multi-venue traders. These findings are consistent with the stylized example in Figure 4.1, as relocating the ‘alternate venue’ matching engine to the same data center as the ‘main venue’ matching engine is likely to substantially reduce or eliminate both  $\epsilon$  and  $S$ . Prior to relocation,  $\epsilon$  is expected to be larger than  $S$ , because HFTs are likely to operate faster communication

networks between venues than other traders.

In this chapter, we examine the impacts of the TSX Quantum XA trading platform upgrade on liquidity fade and quote sniping on other trading venues. Following the findings and predictions in the above literature and the stylized representation of changes in the fragmented trading process in Figure 4.1, we hypothesize an increase in liquidity fade and quote sniping on other trading venues. Furthermore, we predict two resulting effects from higher liquidity fade. Firstly, TSX's market share of trading activity will increase, because liquidity fade on other venues reduces their trading activity. Secondly, adverse selection costs on other venues will decline relative to TSX, because their liquidity suppliers will be better able to avoid multi-venue trades that are typically more informed.

### **4.3 Institutional Details**

Canadian equities trading is fragmented across multiple trading venues. At the time of the TSX Quantum XA upgrade, there were seven lit trading venues for securities listed on the Toronto Stock Exchange (TSX), being the TSX, Alpha, TMX Select, Chi-X, CX2, Omega and Pure Trading. Three dark pools, MatchNow, Instinet and Liquidnet, also offered continuous matching without pre-trade transparency. TSX retains approximately 60% of market share by trading activity, and is operated by TMX Group, which also operates the Alpha and TMX Select venues. Chi-X and CX2 were operated by Chi-X Canada and have since been acquired by Nasdaq.

TSX's trading engine infrastructure was upgraded from the TMX Quantum system to the TMX Quantum Express Accelerated (Quantum XA) system in June 2014, reducing median round-trip order processing latency from 2.3

milliseconds, or 2,300 microseconds, to 26 microseconds.<sup>2</sup> The upgrade was implemented across different stocks in groups constructed by the first letter of the stock code. Stocks beginning with the letters J and Y were migrated on 2 June 2014. The remaining stocks beginning with the letters E through to L and T through to Z were migrated on 9 June 2014. Finally, stocks beginning with letters A through to D and M through to S were migrated on 16 June 2014.<sup>3</sup> The staggered rollout dates provides a series of exogenous shocks that enable strong identification of a causal link between reduced trading platform latency and any observed changes in market efficiency.

The institutional details of both the TSX Quantum XA upgrade and the Canadian equities trading landscape result in an ideal setting to examine whether faster exchange order processing has spillover effects onto other trading venues. Firstly, the 90-fold reduction in latency is extremely large in relative terms while the post-upgrade latency is extremely small in absolute terms, which enables clear observation and analysis of any effects of faster exchange trading speed. More significantly, the pre-upgrade latency exceeds the minimum information transmission and response times between geographically separated Canadian equities trading venues, imposed by speed of light constraints, while the post-upgrade latency is below this threshold. The TSX matching engine is located in a data center in Markham, Toronto whereas the Chi-X and CX2 matching engines are located approximately 30 kilometers away in downtown Toronto. Traveling at the speed of light in a straight line, it takes around 100 microseconds to traverse this distance. Malinova and

---

<sup>2</sup>A comparison of the TMX Quantum and TMX Quantum XA trading platforms and the upgrade's technical details are available at <https://www.tsx.com/resource/en/93>, while the new trading platform's latency is reported at <https://www.newswire.ca/news-releases/tmx-group-readies-launch-of-tmx-quantum-xa-on-toronto-stock-exchange>.

<sup>3</sup>TSX Quantum XA Implementation Update Announcement, available at <https://www.tsx.com/resource/en/1031/2014-017-en.pdf>.

Park (2017) report that Hibernia Networks, a low latency communications provider, estimates the communication latency to be around 400 microseconds in practice. Therefore, the TSX Quantum XA upgrade provides the necessary conditions for increased liquidity fade and quote sniping across venues, if broker SOR systems are not flawlessly calibrated.

Additionally, traditional measures of adverse selection indicate very high levels of informed trading in Canadian equities. The third chapter of this dissertation reports that average realized spreads converge to zero and become negative within a few seconds after trading, due to very high price impacts relative to effective spreads. Approximately half of all trades result in instantaneous adverse selection due to mechanical depletion of order book depth at the best displayed quotes. Price movements immediately after trades are a necessary condition for latency sensitive traders to realize profits from order flow information. Conrad and Wahal (2019) document similar, although less severe, conditions in the United States, where average realized spreads diminish rapidly after trades but remain above zero.

Prior to the TSX Quantum XA upgrade, practitioners raised concerns about its potential to increase quote fade and quote sniping. For example, ITG Canada, an institutional agency broker, forecast that *“this reduction in latency will, in our estimation, lower fill rates for the average smart order router, and create greater opportunity for inter market latency arbitrage ... [because] ... the HFT needs a much smaller latency advantage to get in front of the aggressing order.”*<sup>4</sup> To investigate these predictions, we apply the quote fade metric developed in the third chapter of this dissertation to examine changes in market linkages and the accessibility of liquidity across venues fol-

---

<sup>4</sup>ITG 2014 Canadian Market Structure Forecast, available at <http://itg.com/marketing/2014.CanadianMarketStructureForecast.20140114.pdf>.

lowing the TSX Quantum XA upgrade. We also develop a new metric that examines aggressive competition for order flow to quantify changes in the necessary conditions for quote sniping, which has also been predicted to occur in the theoretical literature (Budish et al., 2015, Menkveld and Zoican, 2017).

## 4.4 Data and Market Efficiency Metrics

### 4.4.1 Data

The data examined in this chapter is from Thomson Reuters Tick History (TRTH), which is provided by the Securities Industry Research Centre of Asia-Pacific (SIRCA). It contains all trade messages with the price, volume, condition code and buyer and seller broker IDs, as well as quote data with the top-of-book bid and ask prices and sizes on each order book update. Messages are stamped with millisecond granularity and correct event sequencing is preserved across the trades and quotes, which are provided within a single file for each venue. Off-market trades, odd lot trades and dark trades are excluded because they do not interact with the displayed limit order book. We also exclude trades larger than \$1 million, even if they are not flagged with an off-market trade condition code.<sup>5</sup>

The trading venues examined are TSX, Alpha, Chi-X and CX2, which together account for around 95% of lit market trading in Canadian equities during the observation period. Omega ATS and TMX Select are excluded due to poor cross-venue timestamp synchronization. Section 4.7 presents the timestamp benchmarking conducted. Canadian securities regulators have also

---

<sup>5</sup>The trade condition codes in the TRTH data may be incomplete, as we are aware of several extremely large trades that did not interact with the limit order book but are not flagged as off-market trades.



identified data quality issues affecting Omega ATS, which makes it impossible to accurately compare its order book activity with concurrent order book activity on other Canadian trading venues.<sup>6</sup> Although several trading venues operate extended trading hours, we constrain our analysis to the TSX's continuous trading session, from 9:30am to 4:00pm.

The observation window is from 31 March 2014 to 15 August 2014, which spans 20 weeks centered on the Quantum XA upgrades on 2, 9 and 16 June 2014. The start of the observation period, in the first quarter of 2014, is constrained by substantial increases in Chi-X market share and corresponding decreases in TSX market share, which significantly precede and are unlikely to be related to the TSX Quantum XA trading platform upgrade. We exclude the quarterly S&P index rebalancing day of 20 June 2014 due to very high trading activity, which is not representative of a typical trading day. We also exclude the United States public holidays of Memorial Day on 26 May 2014 and Independence Day on 4 July 2014, when equities trading activity in Canada is very low. Finally, we exclude 6 August 2014, due to erroneously high market depth on Chi-X Canada across several stocks in the TRTH data. The final observation window consists of 92 days.

The sample of stocks includes the components of the S&P/TSX Composite Index, which contains approximately 250 of the largest stocks listed on the Toronto Stock Exchange, which is the largest equities listing venue in Canada. We exclude stocks that are not members of the TSX Composite Index for the full observation window, as well as CSU and FFH because they have extremely high share prices above \$200. The final sample consists of 233 stocks.

---

<sup>6</sup>Ontario Securities Commission Staff Statement of Allegations in the Matter of Omega Securities Inc., available at [https://www.osc.gov.on.ca/en/Proceedings\\_soa\\_20171116.omega-securities.htm](https://www.osc.gov.on.ca/en/Proceedings_soa_20171116.omega-securities.htm).

#### 4.4.2 Standard Liquidity Metrics

Following the empirical literature, we calculate a range of standard liquidity metrics, which are aggregated at the stock-day level. Firstly, we construct the National Best Bid (NBB) as the highest bid price across all trading venues and the National Best Offer (NBO) as the lowest ask price across all trading venues, at each point in time. We refer to this set as the National Best Bid and Offer (NBBO), which forms the reference price pair for each metric.

Liquidity metrics can be classified as either quoted or traded. Quoted liquidity metrics reflect the displayed liquidity that is immediately available to market participants and are constructed from quote prices and volumes. They are time-weighted throughout each trading day by the duration that each quote update is prevailing, until the next quote update arrives. The quoted liquidity metrics calculated are quoted spread, tick constraint and quoted depth, as well as the proportion of time that TSX is present at the NBBO and its share of total NBBO depth.

The quoted spread measures the absolute difference between the best bid and offer prices. We calculate quoted spreads both in absolute dollar terms, and in relative basis point terms, by dividing the absolute quoted spread by the NBBO midpoint price.

$$QuotedSpread = NBOPrice - NBBPrice \quad (4.1)$$

When the quoted spread is equal to the minimum tick increment at that price level, we consider quoting in the stock to be tick constrained. The minimum tick increment for all stocks in our sample across the full observation period is 1 cent. On a single trading venue in the continuous trading session, the quoted spread will always be a minimum of one tick increment because an

incoming order that narrows the spread to zero will result in a trade at that price. Similar to the United States, Canadian securities regulations prohibit brokers from intentionally submitting orders that lock the NBBO, which occurs if the best bid price on one venue is equal to the best ask price on another venue, for a quoted spread of zero.

Quoted depth is calculated as the sum of all depth displayed across all venues at the NBBO prices, with larger depths indicating that the market is able to sustain larger trades at the NBBO price levels without incurring market impact.

$$QuotedDepth = NBBPrice * NBBVolume + NBOPrice * NBOPVolume \quad (4.2)$$

To examine the price competitiveness of quoting on TSX relative to other trading venues, we calculate the percentage of time that it is present at the NBBO, as the sum of the time that it is quoting at the NBB and the time that it is quoting at the NBO, divided by two.

$$TSXAtNBBO = \frac{1_{TSXBestBidPrice=NBBPrice} + 1_{TSXBestOfferPrice=NBOPrice}}{2} \quad (4.3)$$

To examine the size competitiveness of quoting size on TSX at the NBBO relative to other trading venues, we calculate the percentage of total NBBO depth that is displayed on TSX's order book. The TSX NBBO depth share metric is calculated as the sum of the dollar depth at the best bid price on TSX, conditional on that price being the NBB and the dollar depth at the best ask price on TSX, conditional on that price being the NBO, divided by

the total NBBO depth.

$$TSXNBBO_{Depth} = \frac{TSXDepthAtNBB + TSXDepthAtNBO}{TotalDepthAtNBB + TotalDepthAtNBO} \quad (4.4)$$

Traded liquidity metrics compare trade prices with the NBBO midpoint price, and are turnover-weighted throughout each trading day, to approximate the implicit cost of immediacy that was actually incurred by liquidity demanders. The traded liquidity metrics calculated are effective spread, realized spread and adverse selection. These metrics require assigning trade initiation direction, which can be done with a very high degree of accuracy for TRTH Canadian equities data because the trades and quotes are sourced from the same file and preserve correct message sequencing for each trading venue. Additionally, the order book immediately refreshes after each trade to show liquidity consumption. Trades are assigned as buyer (seller) initiated if the trade price is equal to or greater (less) than the prevailing ask (bid) price on that venue. The greater (less) than condition is necessary because electronic limit order books immediately process large marketable orders that interact with limit orders at multiple price levels, resulting in the potential for trades at multiple prices without interspersed order book updates.

The effective spread measures the implicit transaction cost for each trade that is incurred by the liquidity demander in crossing the bid-ask spread and is calculated by taking the difference between the trade price and the NBBO midpoint price prevailing at the time of the trade,  $NBBOMP_t$ , multiplied by two. Trade direction is +1 for buyer-initiated trades and -1 for seller-initiated trades.

$$EffectiveSpread = 2 * Direction * (TradePrice - NBBOMP_t) \quad (4.5)$$

Adverse selection measures twice the directional change in NBBO midpoint price between the time of the trade and a later reference time, as a proxy for losses incurred by liquidity suppliers due to short-term market movements after the trade. Following Conrad et al. (2015), we compute the adverse selection metric using a reference time of one second after each trade.

$$AdverseSelection = 2 * Direction * (NBBOMP_{t+1sec} - NBBOMP_t) \quad (4.6)$$

Realized spread is calculated as twice the directional difference between the trade price and the NBBO midpoint price at a later reference time and is a proxy for the revenues attributable to liquidity provision, being the effective spread earned minus any adverse selection losses due to NBBO midpoint price movements soon after the trade.

$$\begin{aligned} RealizedSpread &= EffectiveSpread - AdverseSelection \\ &= 2 * Direction * (TradePrice - NBBOMP_{t+1sec}) \end{aligned} \quad (4.7)$$

The above metric specifications for effective spreads, adverse selection and realized spreads are presented in absolute dollar terms. These metrics are also transformed into relative basis point terms, by dividing the metric in dollars by the NBBO midpoint price.

### 4.4.3 Novel Liquidity Metrics

This chapter develops two types of novel liquidity metrics. The first group of metrics examine the competition for simultaneous access to order flow across multiple trading venues, which includes quote fade and measures of potential quote sniping. The second group of metrics quantifies high frequency liquidity replenishment.

We adapt the quote fade metric developed in the third chapter of this dissertation to measure the accessibility of displayed liquidity across fragmented trading venues. Our quote fade metric is similar to those developed by van Kervel (2015) and Malinova and Park (2017). The key difference is that their metrics examine quote cancellations after all trades, whereas we condition on trades that deplete liquidity at a price level. Without depletion and mechanical market impact, there is no need for speed races between liquidity demanders seeking to consume available liquidity before price movements and liquidity suppliers seeking to cancel quotes before they become stale. Price movements are also necessary for informed traders to realize profits from their information.

To construct the quote fade metric, we first benchmark the magnitude of TRTH data timestamp synchronization across our sample, following the approach taken in the third chapter of this dissertation. We rely on securities regulations that prohibit traders from intentionally locking the NBBO (IIROC, 2011) by submitting a bid (offer) at the current best offer (bid) price on another venue, to utilize the distribution of locked NBBO durations as an approximation of timestamp synchronization across venues in our data. Further details are provided in Section 4.7. Figure 4.6 shows that the 90th percentile of the locked NBBO distribution is between 40 milliseconds and 50 milliseconds across almost all trading days for TSX, Alpha, Chi-X and CX2. Therefore,

we take 50 milliseconds as the robust threshold at which to join trades for trade string construction, in contrast with the 30-millisecond threshold used in the third chapter of this dissertation. The benchmarked interval is similar to regulatory requirements that specify a 50-millisecond clock synchronization threshold for trading venues (IIROC, 2016).

Trade strings are constructed by joining all trades in the same direction for each stock across all venues that are separated by less than 50 milliseconds. Trades continue to be added to each string until there are no further trades in the same direction within 50 milliseconds after the last trade in the string. A trade string is categorized as depleting if the NBBO depth on the side of the order book that the trades interact with is fully consumed by the end of the trade string, resulting in an immediate price movement in the direction of the trade. For depleting trade strings, quote fade is calculated as the proportion of displayed liquidity on a venue at the best price level within the trade string that does not result in trades. Trade value can exceed the starting liquidity if trades occurred within a millisecond after new order submissions or the order book was replenished during the trade string, in which case we set a lower bound of zero on the quote fade metric. This metric is aggregated per stock-day by summing across all traded value and all starting quoted value.

$$QuoteFade = 1 - \frac{TradeValue}{Max(StartingLiquidity, TradeValue)} \quad (4.8)$$

Stale quote sniping has been extensively predicted and modeled in the theoretical literature (Budish et al., 2015, Menkveld and Zoican, 2017). Broadly, it describes the scenario where an aggressive HFT liquidity demander senses that another trader intends to submit a marketable order to trade against a limited quantity of limit orders at a given price, and races that trader to con-

sume the scarce liquidity. We develop two metrics to measure outcomes that are necessary, although not sufficient, to demonstrate the existence of stale quote sniping activity.

There are two reasons we are unable to conclude that these metrics definitively identify stale quote sniping by aggressive HFTs. Firstly, the lack of event sequencing and timestamp synchronization across venues means that when two trades are recorded at around the same time, we do not know which trade happened first.<sup>7</sup> Secondly, we do not know whether the trader that was ‘sniped’ by the aggressive HFT actually attempted to trade with but missed out on the limit orders, as the TRTH data does not record individual order submission messages.

We start by constructing trade strings as described above for the quote fade metric. Next, we identify trade strings that have trades on both Alpha and TSX, where the liquidity demanding broker on Alpha is different from the predominant liquidity demanding broker on TSX. Trades with the anonymous broker identifier of “1” are excluded from this analysis. Ideally this metric should be calculated at the client level, but the TRTH trade data is limited to broker identifiers and does not include client identifiers. These trades are classified as potential predatory trades, and their turnover is divided by the total turnover on Alpha to calculate an Alpha potential predatory trading proportion per stock-day.

---

<sup>7</sup>As noted in O’Hara (2015), this limitation affects all empirical analysis conducted by researchers and securities regulators. Regulators have proposed to address this issue by mandating much tighter timestamp synchronization, at 100 microseconds under MiFID II in Europe and 1 millisecond under CAT NMS in the United States, which will assist future empirical analysis.



$$\begin{aligned}
& \textit{PotentialPredatoryTradeOnAlpha} \\
& = 1_{\textit{LiquidityDemanderOnAlpha} \neq \textit{MainLiquidityDemanderOnTSX}}
\end{aligned} \tag{4.9}$$

Our novel metrics for quote fade and potential predatory trading are conceptually similar to those in Malinova and Park (2017) but differ in two important ways. Firstly, Malinova and Park (2017) examines all trade strings, whereas we only examine trade strings in which the best price level was depleted across all trading venues. If the price level is not depleted, high speed competition among traders is unnecessary, because quotes are still available to be consumed or canceled, even for traders who do not participate in the speed race. Additionally, there are no immediate arbitrage profits unless a level of quoted depth is depleted and prices move. Secondly, Malinova and Park (2017) calculate metrics at the trade string level, while we aggregate the metrics per stock-day to streamline the regression analysis with the other liquidity metrics.

To broaden the scope of the potential predatory trading metric, we also calculate a liquidity taker concentration index, which is defined as the Herfindahl Hirschman Index (HHI) of liquidity demanding brokers and measured as the sum of squared market shares of each liquidity demanding broker per trade string. This metric is weighted by the turnover of each trade string and aggregated per stock-day. A lower concentration index indicates more diversity of liquidity demanding brokers competing within a very short duration, and therefore greater potential for predatory trading to have occurred.

$$\begin{aligned}
& \textit{LiquidityTakerConcentrationIndex} \\
& = \sum (\textit{LiquidityDemandingBrokerMarketShare})^2
\end{aligned} \tag{4.10}$$

Foucault et al. (2013) model make and take cycles in which liquidity is supplied and then consumed, after which the process repeats. In this spirit, faster data dissemination alerts traders to the completion of each cycle sooner, facilitating faster liquidity replenishment after quotes have been consumed. We develop a liquidity refresh ratio metric to quantify high speed liquidity replenishment on TSX. For each instance when trades deplete a price level of liquidity and widen the quoted spread, i.e. a take cycle, the liquidity refresh ratio is calculated as the quoted spread 2 milliseconds after the trade divided by the quoted spread immediately before the trade. A 2-millisecond horizon is selected as it is less than the 2.3 millisecond latency of the previous TSX Quantum trading system but greater than the sub-millisecond latency of the new TSX Quantum XA trading system. This metric is equally weighted across stock-days for each TSX price level liquidity depletion event.

$$TSXLiquidityRefreshRatio = \frac{TSXQuotedSpread_{t+2milliseconds}}{TSXQuotedSpread_t} \quad (4.11)$$

Yao and Ye (2018) argue that HFTs compete on speed for queue priority to capture the bid-ask spread, which is evidenced by their increase in proportion of liquidity provision when securities undergo reverse splits that decrease their price and result in larger relative tick sizes. Kirilenko et al. (2017) develop a metric for the futures market that calculates the proportion of the first 100 contracts traded after a price change for which each liquidity supplier type was on the passive side of the trade, to identify the trader types who attain superior queue priority. They note that the average best bid and offer depth in their sample is around 500 contracts. Baldauf and Mollner (2018) apply this metric to the equities market, by calculating a ‘speed index’ measure as

the proportion of the first 100 lots traded at a new higher bid price or lower ask price where each broker was on the liquidity supplying side.

We adapt Baldauf and Mollner (2018)’s metric to measure the proportion of instances a specified broker is the first liquidity supplier after a price movement that sets either a higher bid price or lower ask price. This allows identification of the single fastest liquidity supplier, rather than the largest liquidity supplier who is also relatively fast. In contrast with futures contracts, quoted depths at the top of the order book vary widely across different stocks, so a fixed volume of trades represents different proportions of order book depth across different stocks. For this metric, we exclude trades that have the same broker buying and selling as they may have utilized broker preferencing. In contrast with the standard price-time order matching priority, TSX offers price-broker-time order matching priority to match an incoming marketable order from a broker against its resting orders in the order book ahead of orders that were submitted by other brokers earlier but at the same price.

$$\begin{aligned}
 & \textit{BrokerFrontOfQueueProportion} \\
 &= \frac{\sum_{i=1}^N 1_{\textit{BrokerIsLiquiditySupplierOnFirstTradeAtEnhancedPriceLevel}}}{N}
 \end{aligned} \tag{4.12}$$

#### 4.4.4 Summary Statistics

Table 4.1 presents summary statistics across stock-day observations, including means, medians and standard deviations of the variables prior to and following the staggered TSX Quantum XA upgrade. In this table and all following regression analysis, all variables are double-winsorized at the 1% threshold by stock and then by day. As would be expected of a well-functioning fragmented equities trading environment, quote fade levels are moderate, averaging 8.24%

on TSX, 11.18% on Alpha, 15.52% on Chi-X and 18.92% on CX2 prior to the upgrade event. Quote fade increases across all venues following the upgrade, although the relative and absolute magnitude of increase is smallest on TSX. As the listing venue for large Canadian stocks, TSX has the largest market share of trading volume and NBBO depth, averaging 63.58% and 64.58% respectively prior to the upgrade. It is also present at the NBBO quotes 94.43% of the time. Following the upgrade, TSX's market share of trading volume and time at NBBO increased, while its share of NBBO depth decreased. There is substantial variation in NBBO quoted spreads, with an average of 2.61 cents compared with a median of 1.32 cents, both of which increased following the upgrade event. Prior to the upgrade, effective spreads are 2.29 cents on average, which is slightly narrower than quoted spreads, indicating that trades tend to occur when bid-ask spreads are tighter. Most of the effective spread is subsumed into adverse selection following the trade, rather than realized spread profits attributable to the liquidity supplier. The effective spread does not precisely equal the sum of the realized spread and adverse selection due to winsorization. Applying an exponential transformation to logarithmic variables prior to the upgrade, the average share price in the sample is \$21.54, average NBBO quoted depth is \$70,000, and average daily turnover per stock is \$7.7 million.

## **4.5 Methodology, Results and Discussion**

### **4.5.1 Fleeting Liquidity Across Venues**

This chapter's main contribution is to empirically investigate changes in the accessibility of liquidity across fragmented trading venues in a low latency

Table 4.1: **Summary Statistics**

This table reports descriptive statistics per stock-day across 233 TSX Composite Index component securities from 24 March 2014 to 22 August 2014, which spans a 20-week period centered on the staggered rollout of the TSX Quantum XA trading platform. Quote fade is the proportion of displayed liquidity on each venue at the NBBO that did not result in trades when the entire price level was depleted across all venues. TSX market share of traded volume is presented. TSX time at NBBO is the proportion of time it displays the NBB, plus the proportion of time that it displays the NBO, divided by two. TSX NBBO depth share is the time-weighted proportion of total quoted depth at the NBBO that it displays. Quoted spreads are calculated using the NBBO prices and time weighted across each trading day. Tick constrained is the proportion of time that the NBBO quoted spread is equal to the minimum tick size. Depth is calculated at the NBBO across all venues and time weighted. Effective spread, realized spread and adverse selection are calculated against the NBBO midpoint, the latter two after 1 second. Liquidity refresh ratio is the TSX quoted spread after 2 milliseconds, divided by that at the time of each trade, for trades that move the price. The proportion of instances where a broker associated with HFT DMA activity is at the front of the queue on a new price level is presented. Potential predatory trading on Alpha is the proportion of turnover on Alpha that has a different liquidity demanding broker to the main liquidity demanding broker where the trades occurred within 50 milliseconds. The Herfindahl Hirschman Index of diversity among liquidity takers is presented. Price is the time-weighted NBBO midpoint. Turnover is the dollar value of shares traded. Volatility is the standard deviation of realized one-minute NBBO midpoint returns.

	Pre-Upgrade			Post-Upgrade		
	Mean	Median	Std Dev	Mean	Median	Std Dev
<b>Panel A: Metrics</b>						
Alpha Quote Fade (%)	11.18	8.45	10.42	14.04	10.35	13.49
Chi-X Quote Fade (%)	15.52	12.07	12.76	19.28	15.70	14.58
CX2 Quote Fade (%)	18.92	14.81	18.78	24.92	21.97	19.43
TSX Quote Fade (%)	8.24	6.52	6.43	9.44	7.12	7.85
TSX NBBO Time (%)	94.43	96.88	6.78	95.21	97.11	5.73
TSX NBBO Depth (%)	64.58	64.45	9.78	63.88	63.80	8.41
TSX Market Share (%)	63.58	62.95	10.74	65.26	65.09	9.54
Quoted Spread (cents)	2.61	1.32	2.79	2.68	1.37	2.83
Tick Constrained (%)	60.00	74.06	38.61	58.31	71.05	38.89
Depth (Log \$)	11.15	10.99	0.84	11.25	11.08	0.83
Effective Spread (cents)	2.29	1.31	2.22	2.32	1.35	2.19
Realized Spread (cents)	0.13	-0.07	0.97	0.05	-0.09	0.84
Adverse selection (cents)	2.16	1.59	1.65	2.26	1.65	1.81
TSX Liquidity Refresh Ratio (%)	113.68	111.80	8.59	112.38	109.51	9.63
HFT DMA Front of TSX Queue (%)	13.29	11.19	10.25	16.09	11.63	14.11
Potential Predatory Trading on Alpha (%)	14.00	13.67	8.27	17.25	16.72	9.25
HH Index of Takers	86.76	87.20	5.38	85.64	85.97	5.57
<b>Panel B: Controls</b>						
Price (Log \$)	3.07	3.25	0.90	3.10	3.27	0.89
Turnover (Log \$)	15.85	15.75	1.28	15.88	15.79	1.28
Volatility (basis points)	6.65	5.79	3.54	6.56	5.63	3.48

equities environment following a trading platform speed upgrade on one trading venue. Fleeting liquidity, or quote fade, describes the scenario where quotes on one venue are immediately canceled in response to trades on another venue (Malinova and Park, 2017, van Kervel, 2015).

Figure 4.2 presents the daily level of quote fade on Alpha, Chi-X and CX2 aggregated across all stocks that were upgraded to the TSX Quantum XA trading platform on 9 and 16 June 2014 and equal-weighted per venue. Immediate and substantial increases in quote fade are observed on the day of the upgrade. van Kervel (2015) finds that the incidence of fleeting liquidity decreases when traders use smart order routers (SOR), which send orders to multiple venues simultaneously to efficiently access fragmented liquidity. Our findings validate concerns raised by practitioners that the faster trading platform will increase quote fade for many SORs, especially those that are not perfectly calibrated. Faster order processing and message dissemination increases the amount of time that HFT liquidity suppliers have to observe trading activity on TSX and subsequently cancel their orders on other venues if they determine that depletion of liquidity at a price level, which results in adverse selection, is imminent.

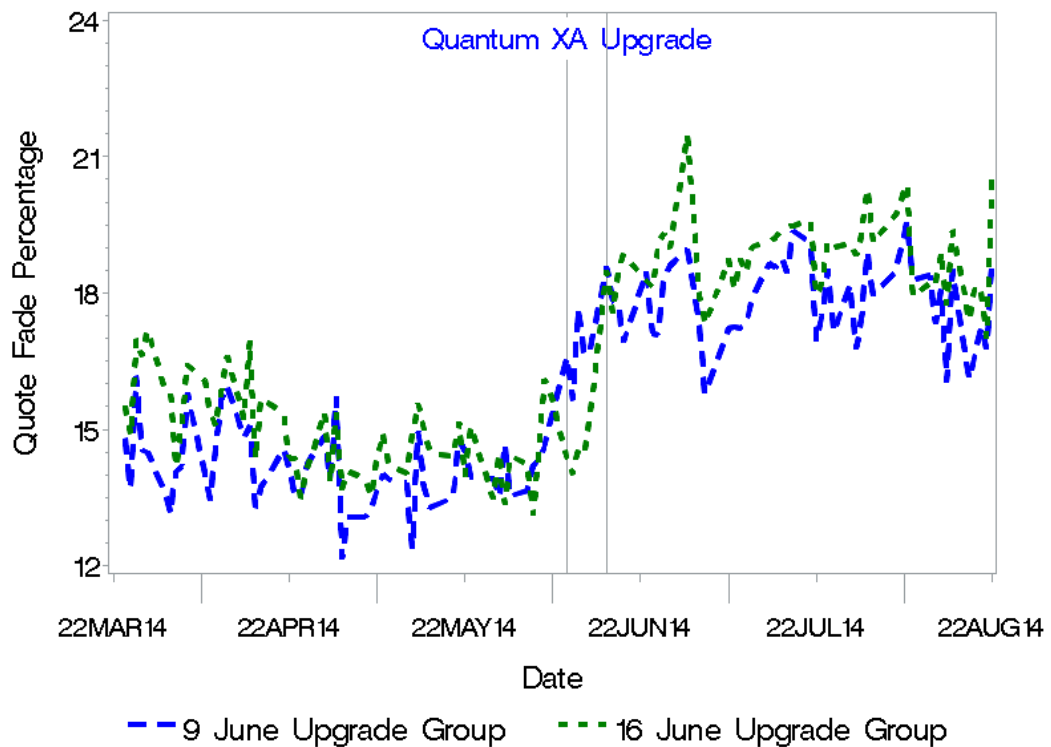
An event study methodology is utilized to formally test the statistical significance of changes in NBBO quote fade following the staggered rollout of the Quantum XA trading platform across stocks, with equations of the form:

$$\begin{aligned}
 QuoteFade_{i,d,v} = & \beta_1 Post_{i,d} + \beta_2 Price_{i,d} + \beta_3 Turnover_{i,d} \\
 & + \beta_4 Volatility_{i,d} + FE_i + \epsilon_{i,d,v}
 \end{aligned} \tag{4.13}$$

where  $QuoteFade_{i,d,v}$  is one minus the volume traded at the best price

Figure 4.2: **Quote Fade by TSX Quantum XA Rollout Groups on Non-TSX Venues**

This figure presents the daily quote fade metric aggregated across Alpha, Chi-X and CX2. Quote fade measures accessibility of liquidity across venues, as a metric of cross-market linkages. Following the second chapter of this dissertation, we construct trade strings by joining all trades in the same direction across all venues separated by less than 50 milliseconds. A trade string is termed depleting if the entire NBBO depth on the side of the market that it interacted with is displaced following the trades. Among all depleting trade strings we calculate the NBBO quote fade metric as the proportion of starting liquidity that did not result in trades. This metric is equal-weighted across stocks. The vertical lines indicate the dates of the staggered Quantum XA trading engine upgrade, which was implemented on 9 and 16 June 2014 depending on the first letter of each security's stock code. The stock group upgraded to TSX Quantum XA on 2 June 2014 is omitted from this figure as it contains only two relatively small stocks and the daily average exhibits much higher variation than the other two groups.



divided by the total starting liquidity at that price, as defined in Equation 4.8, aggregated across all trade strings for stock  $i$  on day  $d$  at venue  $v$ . The key explanatory variables of interest is  $Post_{i,d}$ , an indicator variable equal to one for observations after the TSX Quantum XA upgrade was implemented for stock  $i$  and zero prior. The control variables are  $Price_{i,d}$ , the natural logarithm of the time-weighted NBBO midpoint price,  $Turnover_{i,d}$ , the natural logarithm of on-market turnover, and  $Volatility_{i,d}$ , the realized intraday volatility of one minute NBBO midpoint returns.  $FE_i$  indicates stock fixed effects that control for time-invariant differences in the dependent variable across each stock, and  $\epsilon_{i,d,v}$  is an error term. Standard errors are double-clustered by stock and date.

Table 4.2 reports that following the TSX Quantum XA upgrade, quote fade increased by 6.13% on CX2, 3.61% on Chi-X, 2.82% on Alpha and 1.08% on TSX, which compares with their pre-upgrade means of 18.92%, 15.52%, 11.18% and 8.24% respectively. The coefficient estimates on control variables are qualitatively similar across the venues, with higher quote fade associated with higher share prices and volatility, and lower turnover. Following the trading platform upgrade, it became more difficult for traders seeking to fill marketable orders at a given price level simultaneously across multiple venues. Quote fade is likely to be higher on Chi-X and CX2 than Alpha, because they are located in a different data center to TSX and increased geographical separation results in more opportunities for HFTs to win speed races (Malinova and Park, 2017). The upgrade on TSX increases the information transmission speed of its order book events to other venues, but the change is asymmetric because it does not impact information transmission speeds after trades on other venues. Therefore, the magnitude of quote fade increase on TSX is much lower than on the other venues. We conjecture that quote fade increases slightly on TSX because traders seeking to address substantially higher quote



fade on other venues now send their orders to those venues first, resulting in more opportunities for quote fade on TSX that arise from slight variations in latency.

Table 4.2: **Quote Fade Across All Stocks**

This table reports changes in the inaccessibility of multi-venue liquidity for a sample of 233 TSX Composite Index stocks on all four major Canadian trading venues around the introduction of a trading platform upgrade on TSX, using the regression specification in Equation 4.13. Trade strings are constructed by joining all trades for a stock across venues in the same direction separated by less than 50 milliseconds. Quote fade is the proportion of total displayed liquidity on a venue at the NBB or NBO that did not result in trades when the entire price level was depleted across all venues within a trade string. Post is an indicator for the period after the upgrade, which is 2 June 2014, 9 June 2014 or 16 June 2014, depending on the first letter of the stock code. Control variables include the logarithms of the NBBO midpoint price and turnover and realized 1-minute intraday return volatility. The unit of observation is a stock-day. The sample period covers 20 weeks centered on the middle upgrade date and runs from 24 March 2014 to 22 August 2014. We add a \*/\*\*/\*\* after the t-statistic to indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by both stock and date.

	Alpha	Chi-X	CX2	TSX
Post	2.82 (8.77)***	3.61 (9.54)***	6.13 (11.25)***	1.08 (4.77)***
Price	7.48 (3.53)***	13.80 (4.63)***	-0.99 (-0.22)	7.41 (3.49)***
Turnover	-2.49 (-8.10)***	-4.24 (-12.77)***	-3.47 (-7.36)***	-1.53 (-8.12)***
Volatility	0.44 (7.54)***	0.61 (7.96)***	0.60 (5.87)***	0.23 (5.36)***
Adjusted R2	6.0%	11.5%	4.2%	4.1%
# Obs	21,426	21,409	20,056	21,434

To examine cross-sectional variation in quote fade changes across stocks as a result of the TSX Quantum XA upgrade, we split the stocks into quartiles based on their Herfindhal Hirschman Index (HHI) of market fragmentation by turnover during the pre-upgrade sample period. Stocks where trading is most

fragmented across different venues rely more heavily on cross-market linkages and should experience the largest impacts from the trading platform upgrade. In unreported statistics, we find high correlation between the HHI of market fragmentation by turnover in each stock and the proportion of trading volume in that stock that relies on multi-venue trade strings rather than single-venue trade strings. Additionally, cross-sectional analysis evaluates whether changes in market-wide aggregates are observed across all sub-samples or driven by a particular sub-sample.

Table 4.3 presents the coefficient estimates and t-statistics for the post-upgrade indicator variable from regression models of the quote fade metric separately for each quartile on each trading venue, using the same regression specification as Equation 4.13. For brevity, we do not report the coefficient estimates and t-statistics on control variables. Across all trading venues, quote fade increases the most in the quartile of stocks where trading is most fragmented across venues, and the increases decline monotonically across the quartiles. TSX displays the smallest increase in quote fade for all quartiles. Additionally, quote fade increases across all venues for all quartiles, indicating that the findings are not driven by a small group of stocks. Finally, the economic impact of higher quote fade across the whole market is higher than would be suggested by the market-wide results presented in Table 4.2, because stocks with higher fragmentation across venues also tend to have higher trading activity.

### 4.5.2 Stale Quote Sniping

Budish et al. (2015) and Menkveld and Zoican (2017), among others, predict the existence of aggressive HFTs who “snipe” stale limit orders, which

Table 4.3: **Quote Fade By Stock Fragmentation Quartile**

This table reports changes in the inaccessibility of multi-venue liquidity for a sample of 233 TSX Composite stocks on all four major Canadian trading venues around the introduction of a trading platform upgrade on TSX, using the regression specification in Equation 4.13. Trade strings are constructed by joining all trades for a stock across venues in the same direction separated by less than 30 milliseconds. Quote fade is the proportion of total displayed liquidity on a venue at the NBB or NBO that did not result in trades when the entire price level was depleted across all venues within a trade string. Post is an indicator for the period after the upgrade, which is 2 June 2014, 9 June 2014 or 16 June 2014, depending on the first letter of the stock code, and for brevity we only report coefficient estimates and t-statistics for this variable. The unit of observation is a stock-day. The sample period covers 20 weeks centered on the middle upgrade date and runs from 24 March 2014 to 22 August 2014. Each quartile is examined in a separate regression model, with quartiles constructed based on each stock's Herfindahl index of market fragmentation in the week from 17 March 2014 to 21 March 2014. Quartile 1 includes the most fragmented stocks and quartile 4 includes the least fragmented stocks. We add a \*/\*\*/\*\* after the t-statistic to indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by both stock and date.

	Alpha	Chi-X	CX2	TSX
Most	5.29	6.31	7.46	1.80
Fragmented	(7.85)***	(8.44)***	(7.40)***	(3.87)***
Quartile 2	3.10	3.43	6.84	0.93
	(4.86)***	(5.30)***	(7.21)***	(2.32)**
Quartile 3	2.03	3.03	6.19	0.58
	(4.67)***	(5.01)***	(6.76)***	(2.29)**
Least	0.57	0.98	2.93	0.44
Fragmented	(1.88)*	(2.60)***	(3.92)***	(2.87)***

increases the adverse selection risk faced by liquidity suppliers. Empirically, Kirilenko et al. (2017) find that HFTs frequently consume the last few limit orders at the old price before price changes, which is indicative of sniping activity, while Malinova and Park (2017) observe that HFTs tend to submit more marketable orders in the same direction as, and a few milliseconds after, multi-venue trades from other traders, which can contribute noise to the price discovery process.

We utilize broker identifiers in the TRTH data on individual trades across venues to provide further empirical evidence for the existence of speed races to consume stale quotes before price movements and examine whether this activity increases following the TSX Quantum XA upgrade. We investigate two metrics that describe the necessary conditions for aggressive quote sniping. The first metric is the proportion of trading turnover on Alpha where the liquidity demanding broker is different from the main liquidity demanding broker within that trade string on TSX, defined in Equation 4.9. The second metric is the liquidity taker concentration index, which is broader and measures the propensity for multiple different brokers to be on the marketable side of trades in a high frequency trade string, defined in Equation 4.10. These two metrics quantify the necessary conditions required for the existence of quote sniping. To test for statistically significant changes in aggressive sniping following the trading platform upgrade on TSX, we utilize equations of the form:

$$\begin{aligned} AggressiveSniping_{i,d} = & \beta_1 Post_{i,d} + \beta_2 TickConstraint_{i,d} \\ & + \beta_3 Price_{i,d} + \beta_4 Turnover_{i,d} + \beta_5 Volatility_{i,d} + FE_i + \epsilon_{i,d} \end{aligned} \quad (4.14)$$

where  $AggressiveSniping_{i,d,v}$  is either of the above two quote sniping metrics for stock  $i$  on day  $d$ ,  $TickConstraint_{i,d}$  is the proportion of time for stock  $i$  on day  $d$  that the quoted spread was equal to the minimum tick size of one cent, while other variables are the same as in Equation 4.13. Standard errors are double-clustered by stock and date.

Table 4.4 reports that the liquidity taker concentration index declined by 0.97 after the TSX Quantum XA upgrade. Although the increase in diversity is economically small compared with the pre-upgrade average of 86.76, it is highly statistically significant. The proportion of turnover traded on Alpha that exhibits the necessary conditions for high frequency sniping activity to have occurred increases by 3.18% following the TSX Quantum XA upgrade. The increase is substantial compared with the average of 14% in the pre-upgrade period, which is comparable to the upper bound on this type of activity that Sparrow (2015) estimates at around 17%, using a similar identification approach. To the best of our knowledge, these results provide the first empirical evidence of likely increases in aggressive stale quote sniping across equities trading venues after trading infrastructure becomes faster.

### 4.5.3 Liquidity Provision on TSX

The theoretical literature has offered numerous predictions on the likely impact of faster trading on a single exchange venue, which is examined in this section. We first examine changes in the competitiveness of prices and sizes that TSX quotes at the NBBO. Figure 4.3 presents the proportion of time that TSX was present with quotes at the NBBO, which increases moderately from an average of 94.4% to 95.2% following the Quantum XA upgrade. Visually, the proportion of time that TSX quotes at the NBBO also appears to exhibit

Table 4.4: **High Frequency Competition Among Liquidity Takers Across Venues**

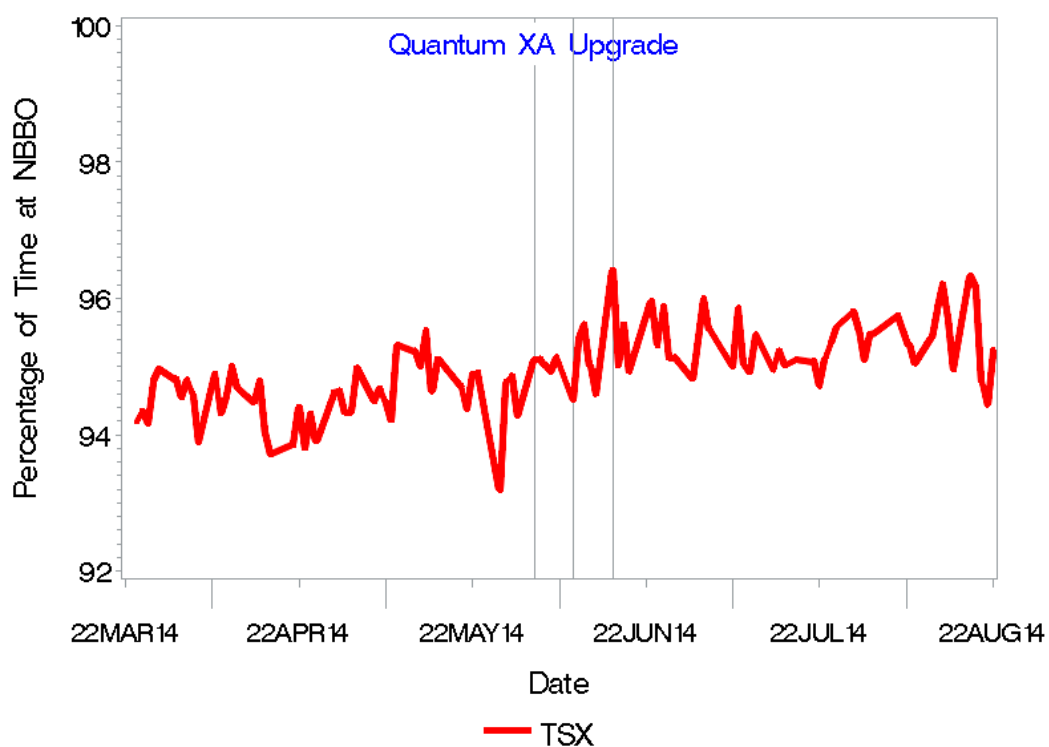
This table reports changes in high frequency competition among liquidity takers across trading venues, for a sample of 233 TSX Composite stocks around the introduction of a trading platform upgrade on TSX, using the regression specification in Equation 4.14. The dependent variable is either the liquidity taker concentration index, defined in Equation 4.10 or the proportion of potential predatory trades on Alpha, defined in Equation 4.9. These two metrics establish necessary, although not sufficient, conditions for the existence of predatory trading activity. Post is an indicator for the period after the upgrade, which is 2 June 2014, 9 June 2014 or 16 June 2014, depending on the first letter of the stock code. Control variables include the proportion of the day over which the quoted spread of the stock was constrained at the minimum tick size, the logarithms of the NBBO midpoint price and turnover and realized 1-minute intraday return volatility. The unit of observation is a stock-day. The sample period covers 20 weeks centered on the middle upgrade date and runs from 24 March 2014 to 22 August 2014. We add a \*/\*\*/\*\* after the t-statistic to indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by both stock and date.

	<b>Liquidity Taker Concentration Index</b>	<b>Potential Predatory Trades on Alpha</b>
Post	-0.97 (-4.08)***	3.18 (9.91)***
Tick	0.01	0.08
Constraint	(1.84)*	(8.81)***
Price	-5.01 (-4.81)***	8.10 (6.11)***
Turnover	-0.80 (-5.06)***	-0.65 (-2.79)***
Volatility	-0.18 (-5.31)***	0.03 (0.83)
Adjusted R2	7.9%	5.6%
# Obs	21,436	21,412

lower volatility and no discernible trend in the post-upgrade period.

Figure 4.3: **TSX Time at NBBO**

This figure presents the proportion of time that TSX quotes were present at the national best bid and offer (NBBO) prices, equal-weighted across stocks. It is calculated by taking the sum of the time that TSX is at the national best bid (NBB) price and the time that it is at the national best offer (NBO) price and dividing by two. The vertical lines indicate the dates of the staggered Quantum XA trading engine upgrade, which was implemented on 2, 9 and 16 June 2014 depending on the first letter of each security's stock code.



As described in Baldauf and Mollner (2018), liquidity suppliers that attain queue priority at the front of the order book after each price change are likely to be the fastest. Figure 4.4 depicts the daily trend in the front-of-queue liquidity provision proportion on TSX of a broker that is associated with HFT direct market access (DMA) service provision. There is an immediate and persistent increase of around 3% following the Quantum XA upgrade, which demonstrates that the fastest liquidity suppliers are able to utilize the trading

engine speed upgrade to enhance their order submission strategies. Interestingly, the proportion increases immediately after the upgrade, but declines within two months, indicating increasing a potential increase in competition among fast HFT liquidity providers.

We conjecture that as market data is disseminated more quickly, the fastest traders are able to observe price level depletion consistently sooner than other traders, and more deterministically be the first to submit limit orders to set a new price level. Under the old TSX Quantum trading engine where market data dissemination was slower, it took longer for traders to become aware of price level depletion so that orders could be submitted to set a new price. The delay introduces a greater degree of randomness in determining who arrives at the front of the queue, which reduces the importance of individual trader speed. For this purpose, Harris (2013) suggests that random speed bumps of short duration could be an effective mechanism to reduce the incentive for speed competition among HFTs.

To formally test for statistically significant changes in liquidity provision on TSX following the trading platform upgrade, we utilize regression specifications of the form:

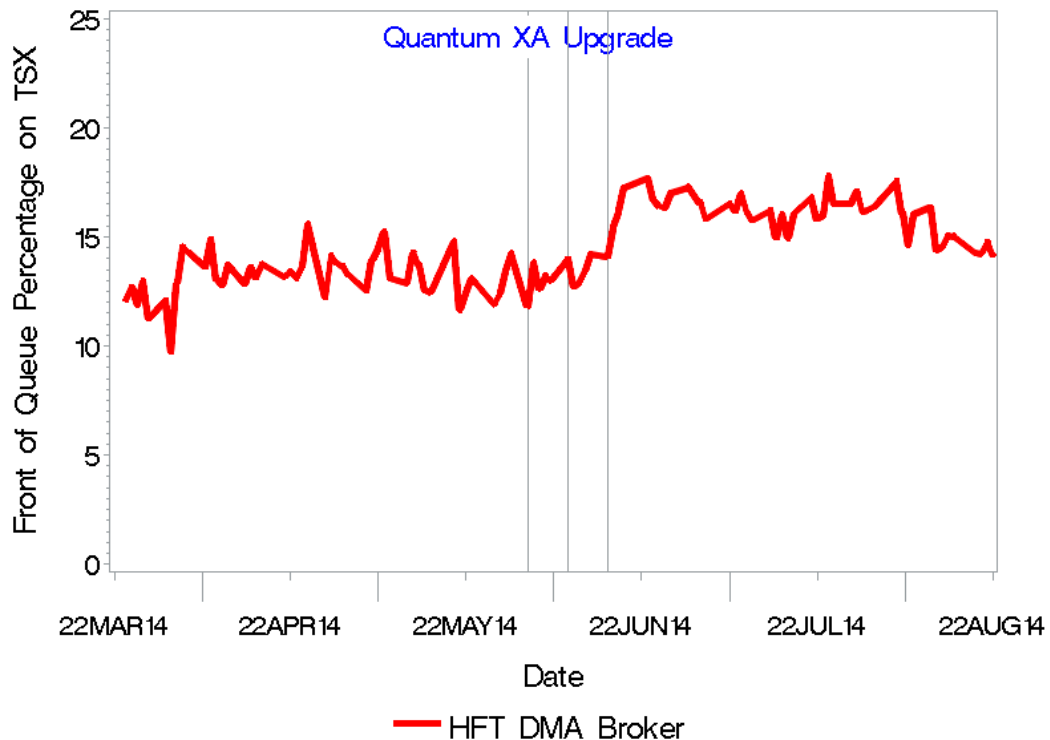
$$\begin{aligned}
 TSXLiquidity_{i,d} = & \beta_1 Post_{i,d} + \beta_2 TickConstraint_{i,d} + \beta_3 Price_{i,d} \\
 & + \beta_4 Turnover_{i,d} + \beta_5 Volatility_{i,d} + FE_i + \epsilon_{i,d}
 \end{aligned} \tag{4.15}$$

Where  $TSXLiquidity_{i,d}$  is TSX's percentage of time at NBBO, percentage of total NBBO depth, liquidity refresh ratio or proportion of trades with HFT DMA broker at the front of the TSX queue for stock  $i$  on day  $d$  as defined in Equation 4.12, and other variables are as defined in Equation 4.14. Standard



Figure 4.4: **HFT DMA Broker at Front of Queue in Liquidity Provision on TSX**

This figure shows the proportion of trades for which a global investment bank that offers direct market access facilities to proprietary traders was at the front of the TSX order book queue on each price step, for trades that were at a superior price to the previous trade. These are buyer-initiated trades at prices lower than the previous trade, and seller-initiated trades at prices higher than the previous trade. This condition ensures that the order was the first to set the quote, rather than being at the front of the queue at that time because forward orders had already been traded earlier. The metric is equal-weighted per observation in each stock, and then equal weighted across all stocks. Trades with the same broker on the buy and sell side of the trade are excluded as broker preferencing functionality on TSX may otherwise confound the results. The vertical lines indicate the dates of the staggered Quantum XA trading engine upgrade, which was implemented on 2, 9 and 16 June 2014 depending on the first letter of each security's stock code.



errors are double-clustered by stock and date.

Table 4.5 reports the results. The proportion of time TSX quotes at the NBBO increases by 0.91%, following the Quantum XA upgrade. Although the statistical significance is marginal, TSX’s proportion of NBBO depth decreases by 0.63%, which is consistent with liquidity suppliers on other venues being willing to quote in greater size on the anticipation of being better able to fade their quotes in response to adverse price movements on TSX, similar to the findings of Malinova and Park (2017).

Menkveld and Zoican (2017) predict that a faster trading platform allows liquidity suppliers to update their quotes more rapidly on incoming news and consequently post narrower quoted spreads. Following the upgrade, the liquidity refresh ratio, which is defined in Equation 4.11 and calculated as the ratio of the average quoted spread 2 milliseconds after a trade relative to the quoted spread at the time of the trade, declines by 1.38%. Quoted spreads on TSX narrow faster after they are widened by large trades that deplete an entire price level of quotes, indicating that liquidity suppliers on TSX utilize the upgrade to replenish orders more rapidly. The HFT-affiliated broker increases the proportion of trades for which it is at the front of the TSX order book queue on each price level by 2.93%. Additionally, the coefficient estimates on the control variables indicate that it attains higher queue priority in stocks with high tick constraint, similar to the findings of Yao and Ye (2018).

#### **4.5.4 TSX Market Share**

We posit that there are two channels via which an exchange speed upgrade may result in that venue attracting additional order flow. Firstly, as predicted by Pagnotta and Philippon (2018), faster trading venues can increase their

Table 4.5: **TSX Liquidity Provision**

This table reports changes in TSX liquidity replenishment for a sample of 233 TSX Composite stocks around the introduction of a trading platform upgrade on TSX, using the regression specification in Equation 4.15. TSX time at NBBO is defined in Equation 4.3. TSX NBBO depth share is defined in Equation 4.4. TSX liquidity refresh ratio is defined in Equation 4.11. The percentage of instances where a global bank that provides direct market access services to proprietary traders sets a new superior price is defined in Equation 4.12. Post is an indicator for the period after the upgrade, which is 2 June 2014, 9 June 2014 or 16 June 2014, depending on the first letter of the stock code. The control variables are natural logarithms of the time-weighted NBBO midpoint price, trading turnover across all venues on that stock-day and realized 1-minute intraday return volatility. The unit of observation is a stock-day. The sample period covers 20 weeks centered on the middle upgrade date and runs from 24 March 2014 to 22 August 2014. We add a \*/\*\*/\*\* after the t-statistic to indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by both stock and date.

	Time at NBBO	NBBO Depth Share	Liquidity Refresh Ratio	HFT DMA at Front of Queue
Post	0.91 (5.52)***	-0.63 (-1.70)*	-1.38 (-6.02)***	2.93 (6.31)***
Tick	0.04	0.06	0.03	0.04
Constraint	(7.36)***	(4.69)***	(2.94)***	(2.82)***
Price	-4.03 (-4.92)***	-2.30 (-1.40)	8.72 (5.89)***	-3.62 (-1.29)
Turnover	0.79 (7.10)***	2.36 (8.07)***	-1.82 (-8.29)***	1.00 (2.77)***
Volatility	-0.12 (-5.39)***	-0.21 (-3.55)***	0.64 (10.85)***	-0.04 (-0.58)
Adjusted R2	4.6%	3.1%	6.3%	3.8%
# Obs	21,436	21,436	21,436	21,433

market share by attracting latency sensitive traders who seek to minimize their search costs and realize gains from trade more rapidly. Secondly, if liquidity suppliers can more successfully fade their quotes on other venues in response to observing a trade on the upgraded venue, the market shares of the other venues will decrease, which in turn increases the relative market share of the upgraded venue.

Figure 4.5 presents the daily time series of TSX's aggregate market share of trading volume, increasing from an average of 64% prior to the upgrade, to 66% over the course of several weeks following the upgrade. We report market shares of trading volumes, rather than turnover values, as trading fees in Canada are calculated based on the number, rather than value, of shares traded.

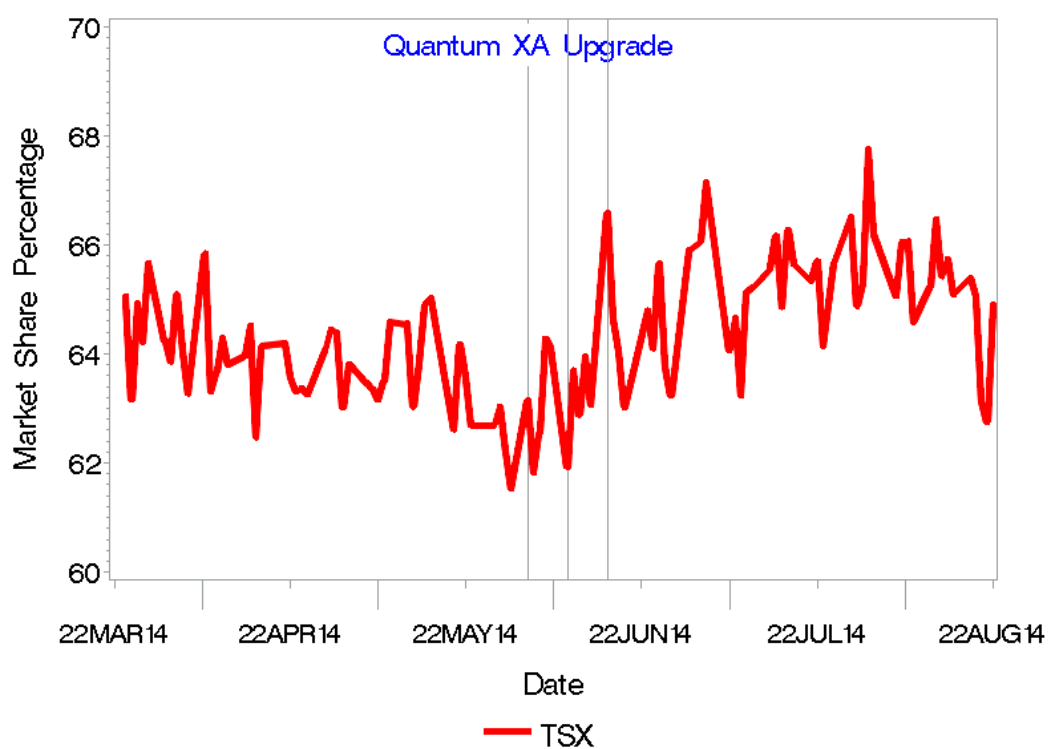
To formally test for statistically significant changes in TSX market share following the trading platform upgrade, we utilize regression specifications of the form:

$$\begin{aligned}
TSXMarketShare_{i,d} = & \beta_1 Post_{i,d} + \beta_2 TickConstraint_{i,d} \\
& + \beta_3 TSXNBBOTime_{i,d} + \beta_4 TSXNBBODepthShare_{i,d} \\
& + \beta_5 Price_{i,d} + \beta_6 Turnover_{i,d} + \beta_7 Volatility_{i,d} + FE_i + \epsilon_{i,d}
\end{aligned} \tag{4.16}$$

where  $TSXMarketShare_{i,d}$  is TSX's share of total on-market volume traded for stock  $i$  on day  $d$ ,  $TSXNBBOTime_{i,d}$  is the proportion of time that TSX was quoting at the NBBO price for stock  $i$  on day  $d$ ,  $TSXNBBODepthShare_{i,d}$  is the proportion of total NBBO depth that is displayed on TSX for stock  $i$  on day  $d$ , and other variables are as defined in Equation 4.15. Standard errors are double-clustered by stock and date.

Figure 4.5: **TSX Market Share**

This figure presents TSX's aggregate market share of traded volumes over the sample period. Market share is presented in terms of volume rather than value as its trading fees are levied based on the number of shares traded, rather than the dollar value traded. The vertical lines indicate the dates of the staggered Quantum XA trading engine upgrade, which was implemented on 2, 9 and 16 June 2014 depending on the first letter of each security's stock code.



In addition to the standard microstructure control variables of share price, trading volume and return volatility, we follow He et al. (2015) to include control variables for the proportion of the stock-day where the quoted spread is constrained by the minimum tick size, TSX's proportion of time quoting at the NBBO and TSX's NBBO depth share. Furthermore, the order protection rule in Canada prohibits marketable orders from being executed on a trading venue when another trading venue is displaying a better price. In practice, this means that trades can only occur on venues that are quoting at the NBBO at the time of the trade. Including the TSX NBBO depth share as a control variable enables disaggregation of changes in market share that are driven by non-marketable order routing decisions and those driven by marketable order routing decisions.

Table 4.6 reports results for the regression specification across the whole market as well as for the same HHI quartiles by trading fragmentation across venues that were constructed in Section 4.1.1. Overall TSX market share increases by 1.73%, with the coefficient estimates on the post-upgrade indicator variable varying between 1.16% and 2.96% across quartiles. The largest increase in TSX market share of 2.96% occurs for the quartile of stocks with the most fragmented trading, which also has the highest increase in quote fade across the other venues, reported in Table 4.3. The second quartile by fragmentation of trading across venues saw the next largest increase in TSX market share of 1.79%. Tick constraint is negatively associated with TSX market share, while share price, total turnover and return volatility are positively correlated with TSX market share. TSX time at NBBO and NBBO depth share have a positive impact on market share, with the latter having a high magnitude and statistical significance, as expected due to the order protection rule. These directionalities are consistent with those found in He et al. (2015),

after accounting for their analysis of market share being for the Chi-X entrant venue, while we examine the main incumbent venue.

#### 4.5.5 Consolidated Quoted and Traded Liquidity

Most of the empirical literature examining the impact of trading platform speed upgrades find improvements in liquidity when latency is reduced from seconds to double digit milliseconds (Conrad et al., 2015, Murray et al., 2016, Riordan and Storkenmaier, 2012) but no changes in liquidity when latency is further reduced below single digit milliseconds (Brogaard et al., 2014a, Murray et al., 2016, Ye et al., 2013). Meanwhile, Foucault and Moinas (2018) postulate that unless traders apply extremely high discount rates, faster realization of gains from trade does not materially increase welfare given the miniscule time scales involved. By examining liquidity changes around the TSX Quantum XA upgrade, which reduces median latency to 26 microseconds, we extend the empirical literature on the impact of faster exchange infrastructure to latencies in the double-digit microseconds. We formally test the statistical significance of changes in consolidated NBBO liquidity metrics following the TSX Quantum XA upgrade with equations of the form:

$$\begin{aligned} LiquidityMetric_{i,d} = & \beta_1 Post_{i,d} + \beta_2 Price_{i,d} \\ & + \beta_3 Turnover_{i,d} + \beta_4 Volatility_{i,d} + FE_i + \epsilon_{i,d} \end{aligned} \quad (4.17)$$

where  $LiquidityMetric_{i,d}$  is one of the consolidated NBBO liquidity metrics for stock  $i$  on day  $d$ ,  $Price_{i,d}$  is the inverse of the time-weighted NBBO midpoint price where the liquidity metric is expressed in basis points, following Hendershott et al. (2011) and the natural logarithm of the time-weighted

**Table 4.6: TSX Market Share Across All Stocks and Per Stock Fragmentation Quartile**

This table reports changes in TSX market share for a sample of 233 TSX Composite stocks around the introduction of a trading platform upgrade on TSX, using the regression specification in Equation 4.16. TSX market share is the proportion of total trading volume across the four largest Canadian equities trading venues that occurred on TSX. Post is an indicator for the period after the upgrade, which is 2, 9, or 16 June 2014, depending on the first letter of the stock code. The control variables are the proportion of time that quoted spreads are constrained by the minimum tick size, TSX time at NBBO, TSX NBBO depth share, the logarithms of the NBBO midpoint price and turnover and realized 1-minute intraday return volatility. The unit of observation is a stock-day. The sample period covers 20 weeks centered on the middle upgrade date and runs from 24 March 2014 to 22 August 2014. The results are presented across all stocks as well as in quartiles that are based on each stock's Herfindahl index of market fragmentation in the week from 17 March 2014 to 21 March 2014, with quartile 1 being the most fragmented stocks and quartile 4 being the least fragmented stocks. We add a \*/\*\*/\*\* after the t-statistic to indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by both stock and date.

	All Stocks	Most Fragmented	Quartile 2	Quartile 3	Least Fragmented
Post	1.73 (9.25)***	2.96 (9.46)***	1.79 (6.20)***	1.16 (4.20)***	1.33 (5.57)***
Tick	-0.08	-0.07	-0.09	-0.09	-0.06
Constraint	(-13.49)***	(-3.64)***	(-7.16)***	(-8.83)***	(-8.68)***
TSX NBBO	0.21	0.34	0.27	0.20	0.15
Time	(10.86)***	(5.90)***	(6.84)***	(6.24)***	(5.18)***
TSX NBBO	0.61	0.52	0.59	0.65	0.65
Depth Share	(54.72)***	(16.60)***	(39.03)***	(34.59)***	(43.54)***
Price	3.02 (3.15)***	-2.07 (-0.97)	3.56 (1.72)*	5.25 (4.94)***	3.02 (2.31)**
Turnover	1.26 (8.14)***	1.43 (3.76)***	1.35 (5.25)***	1.10 (4.19)***	1.12 (4.91)***
Volatility	0.24 (8.19)***	0.17 (2.7)***	0.31 (5.91)***	0.27 (4.38)***	0.24 (5.84)***
Adjusted R2	53.1%	44.4%	51.2%	57.1%	59.9%
# Obs	21,436	5,336	5,428	5,336	5,336



NBBO midpoint price otherwise, and all other variables are as defined in Equation 4.13. Standard errors are double-clustered by stock and date.

Table 4.7 reports change in consolidated NBBO quoted and traded liquidity metrics following the TSX Quantum XA trading platform upgrade. Relative quoted spreads increase by 0.23 basis points after the change, while the change in absolute quoted spreads is economically small and not statistically significant. There is also a 0.92% reduction in the proportion of the trading day that quoted spreads are tick constrained, with marginal statistical significance. Quoted depths increase by 8%. Similar to the findings in Malinova and Park (2017) and consistent with the decrease in TSX's NBBO depth share reported in Table 4.5, greater ability to fade quotes against adverse selection for non-TSX venues increases willingness to post larger quotes. While effective spreads do not change, there is a redistribution of its components. Adverse selection increases by 0.25 basis points and 0.07 cents, consistent with the predictions in Menkveld and Zoican (2017) of faster trading processes resulting in higher adverse selection. Higher adverse selection could also stem from the increases in quote fade reported in Table 4.2, as order book resiliency decreases, and trades are more likely to have market impact. Most of the increase in adverse selection translates into lower realized spreads and therefore lower liquidity supplier profits. In line with previous empirical studies, bid-ask spreads are positively correlated with share prices and return volatility, and negatively correlated with turnover. Overall, the impact of the trading engine upgrade on liquidity is mixed.

Table 4.7: **Consolidated Liquidity Metrics**

This table reports changes in consolidated market-wide liquidity metrics for a sample of 233 TSX Composite stocks around the introduction of the Quantum XA trading platform upgrade on TSX, using the regression specification in Equation 4.17. Tick Constraint is the proportion of each trading day that a stock was quoting at the minimum tick size. Quoted depth is the natural logarithm of the time-weighted dollar value of displayed depth available at the NBBO prices. Quoted spread is the time-weighted average difference between the NBBO prices. Effective spread is the turnover-weighted difference between the trade price and the NBBO midpoint at the time of the trade. Realized spread is the turnover-weighted difference between the trade price and the NBBO midpoint 1 second after the trade. Adverse selection is the directional movement of the NBBO midpoint between the time of the trade and 1 second after the trade. Upgrade is an indicator variable that takes the value 0 before the TSX Quantum XA upgrade and 1 after, which is 2 June 2014, 9 June 2014 or 16 June 2014, depending on the first letter of the stock code. The control variables are the NBBO midpoint price, with an inverse transformation for model specifications where the dependent variable is in basis points and a logarithmic transformation otherwise, the logarithm of turnover across all venues on that stock-day and realized one-minute intraday volatility. The unit of observation is a stock-day. The sample period covers 20 weeks centered on the middle upgrade date and runs from 24 March 2014 to 22 August 2014. We add a \*/\*\*/\*\*\*/ after the t-statistic to indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by both stock and date.

	Tick Constraint	Quoted Depth	Quoted Spread		Effective Spread		Realized Spread		Adverse Selection	
			Cents	Basis Points	Cents	Basis Points	Cents	Basis Points	Cents	Basis Points
Post	-0.92	0.08	0.03	0.23	0.00	0.10	-0.07	-0.18	0.07	0.25
	(-1.79)*	(6.72)***	(0.64)	(2.36)**	(-0.05)	(1.29)	(-2.87)***	(-2.07)**	(2.33)**	(3.03)***
Price	-41.67	-0.20	2.86	79.71	2.13	88.84	-0.28	59.15	2.43	24.17
	(-11.92)***	(-3.28)***	(6.35)***	(15.68)***	(6.81)***	(22.96)***	(-1.71)*	(9.01)***	(11.91)***	(3.64)***
Turnover	8.18	0.28	-0.63	-1.72	-0.47	-1.16	0.04	0.60	-0.50	-1.81
	(15.33)***	(21.07)***	(-10.40)***	(-9.83)***	(-9.02)***	(-7.95)***	(1.69)*	(5.93)***	(-11.75)***	(-16.34)***
Volatility	-1.55	-0.05	0.14	0.37	0.14	0.42	-0.07	-0.47	0.20	0.93
	(-13.60)***	(-18.92)***	(9.22)***	(9.25)***	(9.57)***	(11.05)***	(-12.12)***	(-13.65)***	(12.84)***	(23.84)***
Adjusted R2	18.0%	23.1%	16.4%	26.5%	14.1%	26.2%	4.7%	13.2%	22.0%	26.4%
# Obs	21,436	21,436	21,436	21,436	21,436	21,436	21,436	21,436	21,436	21,436

#### 4.5.6 Liquidity Metrics Across Venues

Following on from the observation of an aggregate increase in adverse selection and reduction in realized spreads, we investigate whether there are changes in the distribution of these liquidity metrics across the trading venues. If liquidity suppliers on other venues are better able to fade their quotes against incoming trades that are likely to impose adverse selection after observing a large market-moving trade on TSX, adverse selection will decline on the venues where quote fade increases. This reasoning is similar to that explored in the third chapter of this dissertation. To examine this hypothesis while controlling for market-wide variation in liquidity across time, we utilize a difference-in-differences regression methodology to examine the changes in liquidity metrics on Alpha, Chi-X and CX2 in excess of changes on TSX, with the following equation:

$$\begin{aligned}
 LiquidityMetric_{i,d,v} = & \beta_1 Post_{i,d} + \beta_2 Alpha_v * Post_{i,d} \\
 & + \beta_3 ChiX_v * Post_{i,d} + \beta_4 CX2_v * Post_{i,d} + \beta_5 Price_{i,d} \\
 & + \beta_6 Turnover_{i,d,v} + \beta_7 Volatility_{i,d} + FE_{i,v} + \epsilon_{i,d,v}
 \end{aligned} \tag{4.18}$$

where  $LiquidityMetric_{i,d,v}$  is the effective spread, realized spread or adverse selection metric for stock  $i$  on day  $d$  at venue  $v$ ,  $Alpha_v$ ,  $ChiX_v$  and  $CX2_v$  are indicator variables equal to one for observations where venue  $v$  is Alpha, Chi-X or CX2, respectively, and zero otherwise,  $Turnover_{i,d,v}$  is the natural logarithm of total on-market trading turnover for stock  $i$  on day  $d$  at venue  $v$ ,  $FE_{i,v}$  indicates stock-venue fixed effects, which control for the time-invariant level of the liquidity metric for each stock on each venue, and  $\epsilon_{i,d,v}$  is an error term. All other variables are as defined in Equation 4.17. Standard errors are

double-clustered by stock and date. Similar to Wang (2018), the treatment indicator variables for Alpha, Chi-X and CX2 are omitted because they are a linear combination of the stock-venue fixed effects, which are included because the treatment and control venues have differing average levels for each liquidity metric.

Table 4.8 reports the results. The baseline observations are TSX stock-days prior to the Quantum XA upgrade. The post-upgrade indicator variable estimates the change in the traded liquidity metric on TSX after the event. The post-upgrade indicator variable is also interacted with indicator variables for each of Alpha, Chi-X and CX2, with the interaction terms estimating the marginal change in the traded liquidity metric on each of these venues after the upgrade, in excess of the change on TSX. The aggregate impact of the upgrade on each venue other than TSX is the sum of the coefficients on the post-upgrade variable and the interaction term of that venue's indicator variable and the post-upgrade variable.

After the Quantum XA upgrade, adverse selection increases by 0.07 cents, realized spreads decrease by 0.06 cents and effective spreads are unchanged on TSX, which is similar to the aggregate analysis across all venues reported in Table 4.7 because TSX is the largest trading venue. Consistent with the hypothesis that increasing quote fade on a trading venue reduces adverse selection for its liquidity suppliers, adverse selection declines by 0.06 cents on Alpha, 0.04 cents on Chi-X and 0.07 cents on CX2, relative to TSX. From Table 4.2, CX2 saw the largest increase in quote fade after the upgrade. Relative adverse selection also declined by 0.16 basis points on Alpha. Across the three smaller trading venues, reductions in adverse selection costs are impounded into narrower effective spreads, leaving the profits retained by liquidity suppliers via realized spreads unchanged. This result indicates a competitive equilibrium

Table 4.8: **Difference-in-Differences Analysis of Traded Liquidity Across Venues**

This table reports changes in effective spreads, realized spreads and adverse selection on Alpha, Chi-X and CX2 relative to the changes on TSX for a sample of 233 TSX Composite stocks around the introduction of the TSX Quantum XA trading platform upgrade, using the regression specification in Equation 4.18. Effective spread is calculated as the difference between the trade price and the NBBO midpoint price at the time of the trade. Adverse selection is calculated as the change in the NBBO midpoint price from the time of the trade to 1 second after the trade in the direction of the liquidity demander. Realized spread is equal to effective spread minus adverse selection. The three metrics are turnover-weighted per stock-day. Post is an indicator for the period after the upgrade, which is 2 June 2014, 9 June 2014 or 16 June 2014, depending on the first letter of the stock code. Alpha, Chi-X and CX2 are indicator variables equal to one for observations on that venue and zero otherwise. The control variables are natural logarithms of the time-weighted NBBO midpoint price, trading volumes across the relevant venues on that stock-day and realized one-minute intraday return volatility. The unit of observation is a stock-day-venue. The sample period covers 20 weeks centered on the middle upgrade date and runs from 31 March 2014 to 15 August 2014. We add a \*/\*\*/\*\* after the t-statistic to indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by both stock and date.

	Effective Spread		Realized Spread		Adverse Selection	
	Cents	Basis Points	Cents	Basis Points	Cents	Basis Points
Post	0.01 (0.30)	0.07 (0.59)	-0.06 (-2.35)**	-0.06 (-0.55)	0.07 (2.34)**	0.15 (1.54)
Alpha * Post	-0.09 (-5.97)***	-0.23 (-2.35)**	-0.04 (-1.94)*	-0.08 (-0.78)	-0.06 (-4.97)***	-0.16 (-2.25)**
Chi-X * Post	-0.05 (-4.65)***	-0.06 (-1.83)*	-0.01 (-0.81)	-0.11 (-1.45)	-0.04 (-3.05)***	0.05 (0.67)
CX2 * Post	-0.13 (-3.14)***	-0.24 (-3.15)***	-0.05 (-1.21)	-0.19 (-1.49)	-0.07 (-2.51)**	-0.05 (-0.44)
Price	2.15 (5.51)***	87.12 (14.61)***	0.26 (1.14)	63.13 (9.57)***	1.84 (9.05)***	22.06 (3.39)***
Turnover	-0.34 (-10.43)***	-0.92 (-10.14)***	-0.15 (-5.74)***	-0.24 (-2.83)***	-0.17 (-9.91)***	-0.64 (-11.54)***
Volatility	0.12 (8.60)***	0.35 (11.49)***	-0.02 (-3.11)***	-0.33 (-7.65)***	0.13 (13.21)***	0.66 (20.51)***
Adjusted R2	7.7%	16.4%	1.1%	5.9%	11.4%	13.7%
# Obs	83,304	83,304	83,304	83,304	83,304	83,304

for liquidity provision across the four major Canadian equities trading venues.

#### **4.5.7 Liquidity Metrics Across Broker Types**

Next, we examine the change in adverse selection for trades on Alpha where a single broker that provides DMA facilities to HFTs was the liquidity supplier, compared with all other trades on Alpha. This is the same broker for which we previously examined order book queue priority on TSX. If fast liquidity suppliers are able to utilize the TSX trading speed upgrade to more quickly cancel their stale quotes on Alpha, we expect their adverse selection costs to decline relative to slower traders. Two reasons explain why Alpha is selected for this analysis. Firstly, the results in Table 4.8 show that Alpha exhibited the most consistent decline in adverse selection relative to TSX following the exchange upgrade, with reductions in both cents and basis points. Secondly, the TRTH data contains broker identifiers for the Alpha trading venue and the venue also offers broker preferencing, which provides an incentive for brokers to broadcast their identity and potentially benefit from matching against their own orders for more advantageous non-marketable order queue priority. Although the TRTH data contains broker identifiers for Chi-X trades, most trades are flagged with the anonymous broker code of “1” as the venue does not offer broker preferencing and there is no incentive for brokers to broadcast their identity. Conversely, CX2 offers broker preferencing functionality but unfortunately the TRTH dataset does not contain broker identifiers for this venue. We formally test for excess changes in adverse selection and profitability of liquidity provision for the HFT-affiliated broker compared with other brokers on Alpha by utilizing a difference-in-differences regression methodology with the following specification:

$$\begin{aligned}
LiquidityMetric_{i,d,b} = & \beta_1 Post_{i,d} + \beta_2 HFTDMA_b * Post_{i,d} \\
& + \beta_3 Price_{i,d} + \beta_4 Turnover_{i,d,b} + \beta_5 Volatility_{i,d} + FE_{i,b} + \epsilon_{i,d,b}
\end{aligned} \tag{4.19}$$

where  $LiquidityMetric_{i,d,b}$  is the effective spread, realized spread or adverse selection on Alpha for stock  $i$  on day  $d$  and broker type  $b$ ,  $HFTDMA_b$  is an indicator variable equal to one for observations where broker type  $b$  is the HFT-affiliated broker and zero otherwise,  $Turnover_{i,d,b}$  is the natural logarithm of total on-market trading turnover on Alpha for stock  $i$  on day  $d$  where broker type  $b$  was the liquidity supplier and  $FE_{i,b}$  indicates stock-broker fixed effects, which control for time-invariant levels of the liquidity metric in each stock trading on Alpha across the HFT-affiliated broker and other brokers, and  $\epsilon_{i,d,b}$  is an error term. All other variables are as defined in Equation 4.18. Standard errors are double-clustered by stock and date. We omit the treatment indicator variable for the HFT-affiliated broker because it is a linear combination of the stock-broker fixed effects, which are included because the average levels of each liquidity metric vary across stocks for the HFT-affiliated broker and the aggregate of other brokers.

Table 4.9 reports the results. Following the TSX Quantum XA upgrade, the HFT-affiliated broker reduces its adverse selection by 0.15 cents or 0.53 basis points, compared with other liquidity suppliers on Alpha. This coefficient estimate is approximately triple the coefficient estimate reported for Alpha in Table 4.8. In fact, excluding this broker, adverse selection on Alpha did not decline at all for other brokers in the post-upgrade period. Also similar to the results reported in Table 4.8, the entire reduction in adverse selection faced by the HFT broker is passed on to liquidity demanders via narrower

effective spreads, which decline by 0.2 cents or 0.62 basis points, relative to trades on Alpha with another liquidity supplying broker. Again, this finding demonstrates efficient competition among liquidity suppliers in the Canadian equities market.

#### 4.5.8 Immediate Impacts on Liquidity Metrics

A key feature of the staggered TSX Quantum XA rollout is that we are able to examine changes in the characteristics of trading for the stocks in each group immediately after implementation, which assists in identifying causality between the upgrade and observed impacts. Figure 4.2 illustrates an immediate increase in quote fade on the event day among stocks in each upgrade group. In this section, we formally test that observation by conducting three difference-in-differences regressions over each of the upgrade dates, using the stocks upgraded on that date as the treatment group, with the stocks upgraded on the other two dates as the control group. Critically, this model setup isolates the impacts that occur within one week of each separate upgrade date. Additionally, any continued drift in effects after the first week following the upgrade will decrease the statistical and economic significance of the coefficient of interest. We use regression specifications of the form:

$$\begin{aligned} LiquidityMetric_{i,d} = & \beta_1 Post_d + \beta_2 (Upgrade_i * Post_d) \\ & + Controls_{i,d} + FE_i + \epsilon_{i,d} \end{aligned} \quad (4.20)$$

where  $LiquidityMetric_{i,d}$  is the liquidity metric of interest for stock  $i$  on day  $d$ ,  $Post_d$  is an indicator variable equal to one for observations after the specified treatment group upgrade date and zero prior,  $Upgrade_i$  is an indicator



**Table 4.9: Difference-in-Differences Analysis of Traded Liquidity for Alpha HFT DMA Broker**

This table reports changes in effective spreads, realized spreads and adverse selection on Alpha for liquidity supplying trades conducted by a broker that provides direct market access (DMA) facilities to high frequency traders (HFTs), relative to the changes on Alpha across all brokers, for a sample of 233 TSX Composite stocks around the introduction of the TSX Quantum XA trading platform upgrade, using the regression specification in Equation 4.19. Effective spread is calculated as the difference between the trade price and the NBBO midpoint price at the time of the trade. Adverse selection is calculated as the change in the NBBO midpoint price from the time of the trade to 1 second after the trade in the direction of the liquidity demander. Realized spread is equal to effective spread minus adverse selection. Post is an indicator for the period after the upgrade, which is 2 June 2014, 9 June 2014 or 16 June 2014, depending on the first letter of the stock code. HFT DMA is an indicator variable equal to one for observations of the broker that is associated with DMA facilities for HFTs and zero for the aggregate Alpha observations. The control variables are natural logarithms of the time-weighted NBBO midpoint price, trading volumes across the relevant broker type on that stock-day and realized one-minute intraday return volatility. The unit of observation is a stock-day-broker. The sample period covers 20 weeks centered on the middle upgrade date and runs from 31 March 2014 to 15 August 2014. We add a \*/\*\*/\*\* after the t-statistic to indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by both stock and date.

	Effective Spread		Realized Spread		Adverse selection	
	Cents	Basis Points	Cents	Basis Points	Cents	Basis Points
Post	-0.06 (-1.63)	-0.09 (-0.90)	-0.10 (-3.82)***	-0.14 (-1.21)	0.03 (0.99)	0.08 (0.80)
HFT DMA	-0.20	-0.62	-0.05	-0.13	-0.15	-0.53
* Post	(-5.50)***	(-6.86)***	(-1.41)	(-1.74)*	(-4.36)***	(-6.54)***
Price	2.00 (5.25)***	97.13 (16.78)***	0.27 (1.17)	70.86 (8.25)***	1.73 (7.88)***	25.88 (3.66)***
Turnover	-0.07 (-2.50)**	-0.36 (-3.37)***	-0.21 (-6.91)***	-0.60 (-5.10)***	0.13 (3.44)***	0.26 (2.61)***
Volatility	0.10 (9.69)***	0.34 (11.01)***	-0.01 (-1.37)	-0.25 (-5.57)***	0.11 (10.71)***	0.61 (16.97)***
Adjusted R2	3.9%	10.4%	1.1%	4.6%	4.0%	7.8%
# Obs	41,852	41,852	41,852	41,852	41,852	41,852

variable equal to one for stocks that are upgraded to TSX Quantum XA on the specified upgrade date and zero for all other stocks,  $Controls_{i,d}$  are the same control variables used for the corresponding liquidity metric panel regression specifications in Equations 4.13, 4.15 and 4.16,  $FE_i$  indicates stock fixed effects, which control for time-invariant levels of the liquidity metric in each stock, and  $\epsilon_{i,d}$  is an error term. Standard errors are double-clustered by stock and date.

The observation window for each upgrade group regression specification spans five days before and after the event date, to avoid overlapping into the previous or next upgrade date. Across the three samples, the treatment stocks were upgraded to Quantum XA on 2, 9 and 16 June 2014 respectively, while stocks not upgraded on that date form the control group. Stock fixed effects are included to control for time-invariant differences in each metric for each stock, as the stocks within each group are not a matched sample. The upgrade group indicator variable is omitted from the regression because it is a linear combination of the individual stock fixed effects.

Table 4.10 reports the results, which for brevity only include the coefficient estimates and t-statistics for statistical significance on the post-upgrade indicator variable and the interaction term between the indicator variables for post-upgrade and upgrade group. The discussion below focuses on the Quantum XA upgrades on 9 and 16 June 2014, as only two relatively small TSX Composite index stocks, JE and YRI, were upgraded on 2 June 2014. All three regression specifications are presented for completeness.

Panel A presents the coefficient estimates for changes in quote fade on each venue. Excluding Alpha on the 9 June 2014 upgrade date, quote fade increases on all venues by between 1.31% and 5.48% for the upgraded stocks in excess of the control stocks, in the week following the upgrade. Quote fade on Chi-X

**Table 4.10: Difference-in-Differences Analysis for Each Upgrade Date**

This table reports results from a difference-in-differences regression analysis of treatment and control stocks over ten-day periods centered on each phase of the staggered Quantum XA upgrade, using the regression specification in Equation 4.20. Quote fade is defined in Equation 4.8. Market share is TSX's proportion of total trading volume. TSX Time at NBBO is defined in Equation 4.3. TSX NBBO depth share is defined in Equation 4.4. Liquidity refresh ratio is defined in Equation 4.11. HFT DMA front of queue is defined in Equation 4.12. Post is an indicator variable equal to zero before the event date and one after. Upgrade is an indicator variable equal to one for the stocks that were upgraded to TSX Quantum XA on the event date and zero otherwise. The control variables used in each regression specification are the same as those in the corresponding event study regressions in Tables 4.2, 4.5 and 4.6. All specifications include stock fixed effects. For brevity, only the coefficients and t-statistics for post and the interaction term between upgrade and post are reported. We add a \*/\*\*/\*\* after the t-statistic to indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by both stock and date.

	<b>2 June 2014 Event</b>		<b>9 June 2014 Event</b>		<b>16 June 2014 Event</b>	
	Post	Upgrade*Post	Post	Upgrade*Post	Post	Upgrade*Post
Panel A: Quote Fade						
Alpha	0.58 (1.31)	5.64 (2.76)***	-0.85 (-1.68)*	0.58 (0.97)	0.38 (0.69)	2.05 (3.51)***
Chi-X	0.03 (0.11)	14.45 (5.43)***	-1.23 (-2.67)***	4.25 (5.98)***	0.16 (0.23)	3.82 (5.74)***
CX2	1.79 (1.87)*	18.82 (4.42)***	-1.48 (-1.48)	5.48 (3.74)***	3.26 (2.51)**	3.07 (1.69)*
TSX	0.54 (1.54)	6.10 (3.01)***	-0.70 (-1.75)*	1.63 (4.06)***	-0.64 (-1.89)*	1.31 (3.99)***
Panel B: Quoting and Trading on TSX						
Market Share	-0.40 (-1.27)	-0.67 (-0.43)	-0.08 (-0.19)	1.29 (3.22)***	0.26 (0.75)	1.40 (3.88)***
Time at NBBO	0.71 (1.50)	-0.22 (-0.48)	0.26 (1.05)	-0.23 (-0.66)	0.60 (1.23)	-0.06 (-0.19)
NBBO Depth	1.25	4.82	-0.59	-0.12	0.39	0.05
Share	(2.00)**	(2.20)**	(-0.89)	(-0.16)	(0.40)	(0.08)
Liquidity	-0.15	-3.69	-0.44	-1.95	1.57	-2.01
Refresh Ratio	(-0.38)	(-6.81)***	(-1.41)	(-6.56)***	(4.15)***	(-5.59)***
HFT DMA	0.00	-5.48	-0.23	2.25	1.45	1.59
Front of Queue	(0.00)	(-2.55)**	(-0.54)	(2.66)***	(1.72)*	(2.30)**

and CX2 increase substantially more than TSX and Alpha, which are located within the same data center. This finding mirrors that in Malinova and Park (2017), which finds that quote fade on a venue decreases after its matching engine was relocated to the same data center as the main trading venue.

Panel B presents the coefficient estimates for changes in quoting and trading activity on TSX. Over the week immediately following the upgrade, TSX's market share of trading volume increases for upgrade stocks in excess of control stocks by 1.29% and 1.40% for stocks upgraded on 9 and 16 June 2014 respectively. Reductions in the liquidity refresh ratio indicate that quoted spreads on TSX are approximately 2% narrower within two milliseconds after trades that are sufficiently large to deplete the liquidity at a price level. The HFT-affiliated broker is at the front of the TSX order book queue 1.59% to 2.25% more frequently within the first week after the trading platform upgrade.

## 4.6 Conclusion

As financial market technology progresses, stock exchanges around the world periodically upgrade their core trading platform infrastructure to reduce the amount of time taken to process orders and disseminate market data. Over the past decade, this latency has decreased by many orders of magnitude, from seconds to milliseconds and now microseconds, which is well beyond the reaction times of most traders. Although the speed of order matching and data dissemination has become much faster, it is not intuitively clear what economically meaningful implications this might have on the organization of securities trading, individual traders or overall market efficiency.

Numerous studies have examined the impact of faster trading infrastructure on a single stock exchange, generally finding that liquidity improves as

exchange latency is reduced to around 10 milliseconds, after which there are no further liquidity improvements with faster exchange technology. This chapter is the first study to examine the impact of increases in trading platform speed in a fragmented trading environment. The potential for spillover effects arising from races among fast traders to transmit orders to one trading venue in response to information ascertained from order flow observed on other trading venues is introduced in this setting.

We find that the Toronto Stock Exchange's Quantum XA trading platform upgrade, which reduces its round-trip order processing latency from 2.3 milliseconds to 26 microseconds, results in an immediate increase in fleeting liquidity. This is concentrated on other trading venues and in stocks where trading is most fragmented across venues. The externality increases TSX's market share by 2%. Adverse selection costs decline for liquidity providers who are better able to fade their quotes and avoid interacting with trades that have market impact. The necessary conditions for quote sniping activity by aggressive HFTs also increases. Consistent with the theoretical literature, we find that liquidity on TSX is replenished sooner after it is consumed, improving its overall quote competitiveness at the NBBO, while an HFT-affiliated broker is at the front of the order book queue more frequently.

This chapter empirically demonstrates that faster trading platforms benefit low latency traders and are a mechanism for trading venues to compete for market share in a high frequency world, adding to the microstructure literature on the arms race for speed. It is the first to document the externalities that enhancements in low latency trading infrastructure on one trading venue can have on other venues and the gains from trade of different types of market participants, with consequences for the fairness of markets. In light of these findings, slightly slowing down market data dissemination could improve investor

outcomes in fast and fragmented equities trading environments. Meanwhile, practitioners have patented systems to address this market friction.<sup>8</sup>

## 4.7 Appendix

### 4.7.1 Timestamp Synchronization Benchmarking

In this section, we follow the methodology developed in the third chapter of this dissertation to benchmark cross-venue timestamp synchronization in our data. Figure 4.6 presents percentiles from the distribution of time intervals during which the national best bid and offer (NBBO) prices constructed from lowest offer price and highest bid price on the specified trading venues indicated a “locked market”, i.e. best bid price equal to best offer price, or a “crossed market”, i.e. best bid price higher than best offer price. The sample is identical to that used in this chapter, being 233 TSX Composite Index component securities from 24 March 2014 to 22 August 2014.

The 90th percentile of the distribution across TSX, Alpha, Chi-X and CX2, presented in the top left panel, is below 50 milliseconds across almost all days in both the pre-upgrade and post-upgrade periods. Therefore, we define the threshold for joining trades to construct trade strings in this chapter at 50 milliseconds, which is slightly wider than the 30 millisecond threshold utilized in the third chapter of this dissertation. Additionally, the stability of each of the distribution percentiles for this quadrant of venues, both before and after the TSX Quantum XA upgrade, validate that the upgrade did not have a material impact on the way that timestamps are recorded in the TRTH data

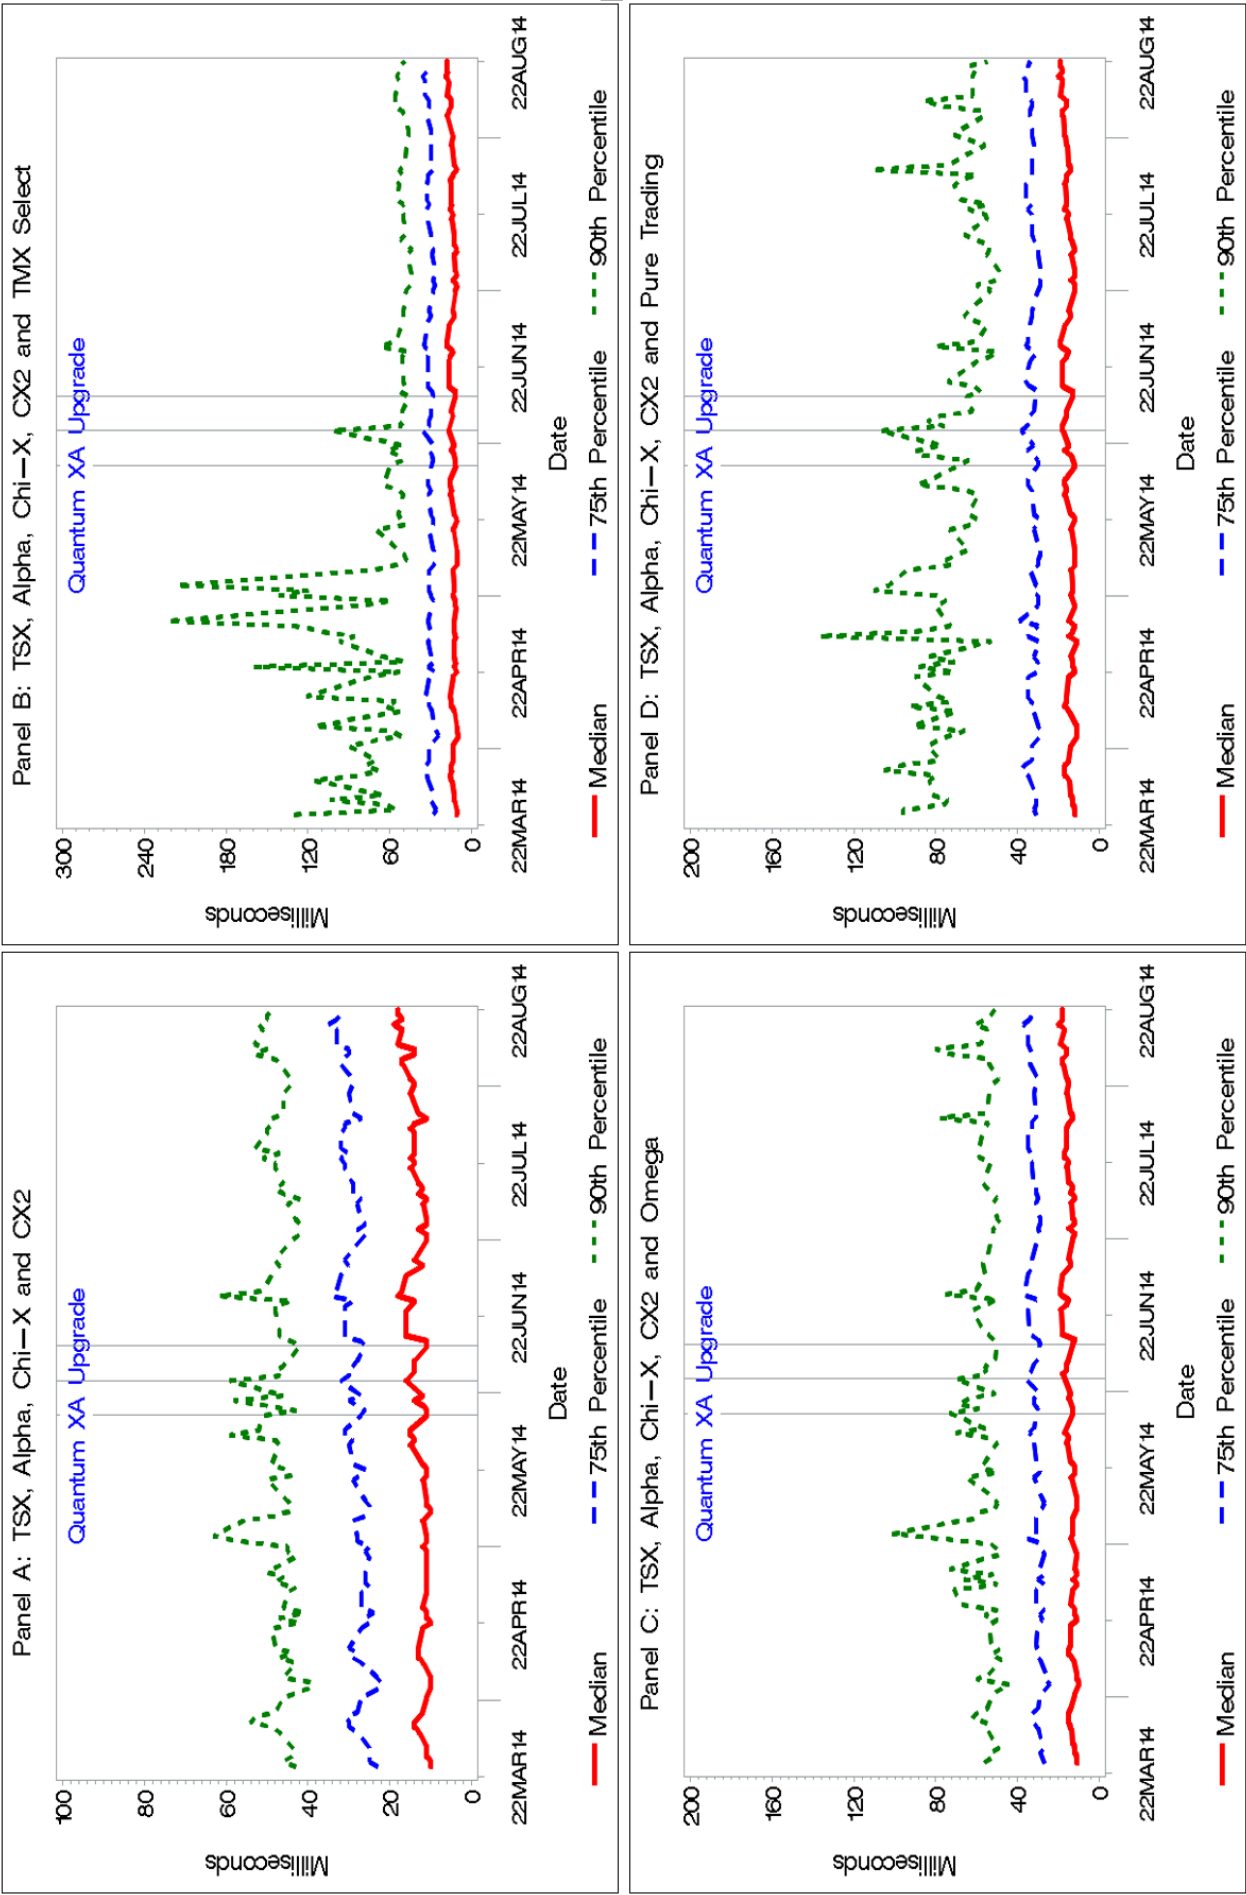
---

<sup>8</sup>Renaissance Technologies, a quantitative hedge fund, has patented a system for synchronized timing of order execution across fragmented trading venues by using co-located atomic clocks at each venue. This mitigates the issue of fleeting liquidity as trading venues become faster, by reducing the margin of timing error in attempting to simultaneously capture liquidity across multiple trading venues. See <https://patents.google.com/patent/US20160035027>.

used for this chapter.

TMX Select, Omega and Pure Trading are excluded from our sample due to large variation in timestamp asynchronicity compared with the four main trading venues.

Figure 4.6: Timestamp Synchronization Benchmarking





# Chapter 5

## Conclusions

This dissertation investigates the nature of competition among stock exchanges and its impacts on liquidity supply in a high frequency equities trading environment. While competition among market participants and the role of high frequency traders has been examined extensively, I add to the small but growing literature that examines the role of stock exchanges in the pursuit of ever faster trading. I show that stock exchanges now use speed as a market design differentiator to compete for various types of order flow in a fragmented trading environment.

Amendments to securities regulations have facilitated the entry of new stock exchanges that provide alternative trading venues for securities listed elsewhere. Competition has generally been found to improve market efficiency, although the channels via which this occurs are less well understood.

The second chapter of this dissertation examines the entry of the Chi-X stock exchange into the previously monopolistic Australian equities trading environment. In anticipation of competition, the incumbent Australian Securities Exchange (ASX) lowered its trading fees, which liquidity suppliers impound into narrower bid-ask spreads. Market efficiency improves following

Chi-X's launch, facilitated by an additional order book queue (Foucault and Menkveld, 2008) and entry of new endogenous liquidity providers (Menkveld, 2013). Most of the liquidity improvements accrue to traders who utilize smart order routing technology to access both exchanges.

To compete for different types of order flow, several stock exchanges have recently introduced speed bumps to delay order processing for some traders or order types, but not others, with IEX in the United States being perhaps the most famous. The theoretical literature has proposed speed bumps as mechanisms to de-emphasize speed in the trading process and mitigate the high frequency trading arms race. Meanwhile, the financial press has described the delays as a way of 'levelling the playing field' among traders with different speed hierarchies. Critically, speed bump design can differ along two dimensions that determine their impacts on trading outcomes: firstly, whether all orders or only specific types of orders are slowed down, and secondly, whether the delay duration is fixed or randomized.

The third chapter of this dissertation examines the introduction of a randomized 1-3 millisecond speed bump on TSX Alpha in Canada, which only applies to liquidity demanders. In an environment where multi-market trades have persistently higher adverse selection than single market trades (Malinova and Park, 2017), the speed bump increases speed differentials among traders and enables fast liquidity suppliers on this venue to avoid order flow driven adverse selection. The profitability of their liquidity provision increases immensely. Without sufficient competition between them, extremely fast traders retain lower adverse selection costs and bid ask spreads do not narrow. Consistent with the theoretical literature, adverse selection increases for slower liquidity suppliers on this venue (Han et al., 2014), and liquidity suppliers on other venues (Biais et al., 2015), which is partially absorbed into lower profits

and partially passed on via wider bid ask spreads. When bestowing systematic speed advantages on some traders over others, stock exchange operators and securities regulators should be cautious in picking the winners and understand the potential impact on the losers.

Motivated by O'Hara (2015)'s assertion that market linkages and fairness are two especially important policy issues for future market microstructure research, the third chapter also develops a novel 'quote fade' metric, to quantify the simultaneous accessibility of liquidity across multiple trading venues, which is indicative of efficient market linkages and fairness for large liquidity demanders. The metric may be constructed using widely available academic datasets, such as Thomson Reuters Tick History. O'Hara (2015) also identifies incorrect order and trade message sequencing and lack of clock synchronization across trading venues as two data quality issues that affect the robustness of empirical research. This chapter develops tools to address these complications.

Utilizing technological innovation, stock exchanges around the world periodically upgrade their trading platforms to increase the speed at which incoming orders are processed. Over the last decade, the time taken to process orders and disseminate market data has decreased tremendously and is now frequently measured in microseconds. Stock exchanges promote these upgrades as key enhancements to their trading venue. While increases in exchange trading speed have been examined extensively on individual stock exchanges, less is known about whether there are spillover effects in fragmented trading environments.

The fourth chapter of this dissertation examines the impact of a trading engine upgrade on the main TSX stock exchange in the competitive Canadian equities trading environment, which reduced latency from 2.3 milliseconds to 26 microseconds. Critically, latency was reduced to below the 100-microsecond minimum data transmission time between the geographically separated trading

venues. While aggregate liquidity did not change substantially, TSX increased its market share and fast liquidity providers were better able to achieve queue priority, which increases the rate at which they realize gains from trade (Foucault and Moinas, 2018). However, quicker trade completion after the upgrade also had the unintended consequence of enabling fast liquidity suppliers on other venues to fade their quotes after observing trades on the upgraded TSX, reducing the accessibility of liquidity across venues and breaking down cross-market linkages. Potential predatory HFT sniping activity also increased. As a corollary of this chapter, and in the spirit of Hu (2018), slowing down market data dissemination may improve natural investor outcomes in a fast and fragmented equities trading environment.

Overall, this dissertation’s findings indicate that given the complexity of trading in fast and fragmented equities markets, nuances in the design of each stock exchange innovation are important for determining their impacts on overall liquidity and welfare for different types of traders. These distinctions include whether participants connect only to the main exchange or to all exchanges in a fragmented trading environment, changes to trading fees, who speed bumps slow down and by how long, and whether exchange order processing and data dissemination speeds exceed the information transmission times between geographically dispersed trading venues. Additionally, while the trading speed innovations examined in chapters three and four are equally available to all traders, few invest in the speed technology required to realize the benefits, raising questions about the social utility and fairness of faster trading in the microsecond environment.

Until very recently, competition among stock exchanges has been restricted to continuous limit order books, which is the market design examined in this dissertation. These venues process orders serially upon arrival on a first-in-

first-out basis. Recent theoretical literature has modeled the potential benefits of alternative market designs and order matching mechanisms (Baldauf and Mollner, 2019). Budish et al. (2015) propose frequent batch auctions as one such alternative, to de-emphasize the socially wasteful arms race for speed in a winner-takes-all contest to consume or cancel stale quotes arising from mechanical latency arbitrage opportunities when prices move in response to symmetric public information. In the fast and fragmented European equities trading environment, periodic auction venues that process orders at discrete time intervals throughout the trading day have recently been launched and operate alongside existing continuous limit order books that trade the same securities. The competitive dynamics between these novel entrant trading venues and incumbent trading venues, and their impacts on market efficiency and welfare for different trader types, are topics for future research.

# Bibliography

AFM, 2016, A case analysis of critiques on high-frequency trading, Technical report, The Netherlands Authority for the Financial Markets.

Ait-Sahalia, Yacine, and Mehmet Sağlam, 2017, High Frequency Market Making: Optimal Quoting, SSRN Scholarly Paper ID 2331613, Social Science Research Network, Rochester, NY.

Amihud, Yakov, 2002, Illiquidity and stock returns: cross-section and time-series effects, *Journal of Financial Markets* 5, 31–56.

Anand, Amber, and Kumar Venkataraman, 2016, Market conditions, fragility, and the economics of market making, *Journal of Financial Economics* 121, 327–349.

Anderson, Lisa, Emad Andrews, Baiju Devani, Michael Mueller, and Adrian Walton, 2018, Speed Segmentation on Exchanges: Competition for Slow Flow, Technical Report 18-3, Bank of Canada.

Anderson, Lisa, Baiju Devani, and Yifan Zhang, 2016, The Hidden Cost: Reference Price Latencies, Trading Review And Analysis, Investment Industry Regulatory Organization of Canada.

Angel, James J., 2014, When Finance Meets Physics: The Impact of the Speed

- of Light on Financial Markets and Their Regulation, *Financial Review* 49, 271–281.
- Aoyagi, Jun, 2019, Strategic Speed Choice by High-Frequency Traders under Speed Bumps, SSRN Scholarly Paper ID 3221990, Social Science Research Network, Rochester, NY.
- Aquilina, Matteo, Sean Foley, Peter O'Neill, and Thomas Ruf, 2016, Asymmetries in Dark Pool Reference Prices, SSRN Scholarly Paper ID 2848120, Social Science Research Network, Rochester, NY.
- Aquilina, Matteo, and Carla Ysusi, 2016, Are High-Frequency Traders Anticipating the Order Flow? Cross-Venue Evidence from the UK Market, *Market Microstructure and Liquidity* 02, 1750005.
- ASIC, 2015, Review of high-frequency trading and dark liquidity, Technical Report 452, Australian Securities and Investments Commission.
- Baldauf, Markus, and Joshua Mollner, 2018, Fast Trading on Fake News: The Role of Speed in Liquidity Provision, SSRN Scholarly Paper ID 2640823, Social Science Research Network, Rochester, NY.
- Baldauf, Markus, and Joshua Mollner, 2019, High-Frequency Trading and Market Performance, SSRN Scholarly Paper ID 2674767, Social Science Research Network, Rochester, NY.
- Baron, Matthew, Jonathan Brogaard, Björn Hagströmer, and Andrei Kirilenko, 2019, Risk and Return in High-Frequency Trading, *Journal of Financial and Quantitative Analysis* 54, 993–1024.
- Battalio, Robert, Shane A. Corwin, and Robert Jennings, 2016, Can Brokers

- Have It All? On the Relation between Make-Take Fees and Limit Order Execution Quality, *The Journal of Finance* 71, 2193–2238.
- Battalio, Robert, Jason Greene, and Robert Jennings, 1997, Do Competing Specialists and Preferencing Dealers Affect Market Quality?, *The Review of Financial Studies* 10, 969–993.
- Bennett, Paul, and Li Wei, 2006, Market structure, fragmentation, and market quality, *Journal of Financial Markets* 9, 49–78.
- Benston, George J., and Robert L. Hagerman, 1974, Determinants of bid-asked spreads in the over-the-counter market, *Journal of Financial Economics* 1, 353–364.
- Bernales, Alejandro, Nicolás Garrido, Satchit Sagade, Marcela Valenzuela, and Christian Westheide, 2018, A Tale of One Exchange and Two Order Books: Effects of Fragmentation in the Absence of Competition, SSRN Scholarly Paper ID 3276548, Social Science Research Network, Rochester, NY.
- Bessembinder, Hendrik, 2000, Tick Size, Spreads, and Liquidity: An Analysis of Nasdaq Securities Trading near Ten Dollars, *Journal of Financial Intermediation* 9, 213–239.
- Bessembinder, Hendrik, 2003, Issues in assessing trade execution costs, *Journal of Financial Markets* 6, 233–257.
- Bessembinder, Hendrik, and Herbert M. Kaufman, 1997, A cross-exchange comparison of execution costs and information flow for NYSE-listed stocks, *Journal of Financial Economics* 46, 293–319.
- Biais, Bruno, Thierry Foucault, and Sophie Moinas, 2015, Equilibrium fast trading, *Journal of Financial Economics* 116, 292–313.



- Boulatov, Alex, and Thomas J. George, 2013, Hidden and Displayed Liquidity in Securities Markets with Informed Liquidity Providers, *The Review of Financial Studies* 26, 2096–2137.
- Brogaard, Jonathan, Björn Hagströmer, Lars Nordén, and Ryan Riordan, 2015, Trading Fast and Slow: Colocation and Liquidity, *The Review of Financial Studies* 28, 3407–3443.
- Brogaard, Jonathan, Terrence Hendershott, Stefan Hunt, and Carla Ysusi, 2014a, High-Frequency Trading and the Execution Costs of Institutional Investors, *Financial Review* 49, 345–369.
- Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan, 2014b, High-Frequency Trading and Price Discovery, *The Review of Financial Studies* 27, 2267–2306.
- Brolley, Michael, and David Cimon, 2018, Order Flow Segmentation, Liquidity and Price Discovery: The Role of Latency Delays, SSRN Scholarly Paper ID 3005738, Social Science Research Network, Rochester, NY.
- Brolley, Michael, and Katya Malinova, 2013, Informed Trading and Maker-Taker Fees in a Low-Latency Limit Order Market, SSRN Scholarly Paper ID 2178102, Social Science Research Network, Rochester, NY.
- Budish, Eric, Peter Cramton, and John Shim, 2015, The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response, *The Quarterly Journal of Economics* 130, 1547–1621.
- Budish, Eric, Robin S Lee, and John J Shim, 2019, Will the Market Fix the Market? A Theory of Stock Exchange Competition and Innovation, Working Paper 25855, National Bureau of Economic Research.

- Buti, Sabrina, Francesco Consonni, Barbara Rindi, Yuanji Wen, and Ingrid M. Werner, 2015, Sub-Penny and Queue-Jumping, SSRN Scholarly Paper ID 2350424, Social Science Research Network, Rochester, NY.
- Carrion, Allen, 2013, Very fast money: High-frequency trading on the NASDAQ, *Journal of Financial Markets* 16, 680–711.
- Chakrabarty, Bidisha, Jianning Huang, and Pankaj K. Jain, 2019, Effects of a Speed Bump on Market Quality and Exchange Competition, SSRN Scholarly Paper ID 3197088, Social Science Research Network, Rochester, NY.
- Chakrabarty, Bidisha, Pankaj K. Jain, Andriy Shkilko, and Konstantin Sokolov, 2014, Unfiltered Market Access and Liquidity: Evidence from the SEC Rule 15c3-5, SSRN Scholarly Paper ID 2328231, Social Science Research Network, Rochester, NY.
- Chakrabarty, Bidisha, Pamela C. Moulton, and Roberto Pascual, 2017, Trading system upgrades and short-sale bans: Uncoupling the effects of technology and regulation, *Journal of Empirical Finance* 43, 74–90.
- Chakravarty, Sugato, and Asani Sarkar, 2002, A model of broker's trading, with applications to order flow internalization, *Review of Financial Economics* 11, 19–36.
- Chlistalla, Michael, and Marco Lutat, 2011, Competition in securities markets: the impact on liquidity, *Financial Markets and Portfolio Management* 25, 149–172.
- Cohen, Kalman J., and Robert M. Conroy, 1990, An Empirical Study of the Effect of Rule 19c-3, *The Journal of Law & Economics* 33, 277–305.

- Comerton-Forde, Carole, and Tālis J. Putniņš, 2015, Dark trading and price discovery, *Journal of Financial Economics* 118, 70–92.
- Conrad, Jennifer, Sunil Wahal, and Jin Xiang, 2015, High-frequency quoting, trading, and the efficiency of prices, *Journal of Financial Economics* 116, 271–291.
- Conrad, Jennifer S., and Sunil Wahal, 2019, The Term Structure of Liquidity Provision, SSRN Scholarly Paper ID 2837111, Social Science Research Network, Rochester, NY.
- Degryse, Hans, Frank de Jong, and Vincent van Kervel, 2015, The Impact of Dark Trading and Visible Fragmentation on Market Quality, *Review of Finance* 19, 1587–1622.
- Ding, Shengwei, John Hanna, and Terrence Hendershott, 2014, How Slow Is the NBBO? A Comparison with Direct Exchange Feeds, *Financial Review* 49, 313–332.
- Domowitz, Ian, Jack Glen, and Ananth Madhavan, 1998, International Cross-Listing and Order Flow Migration: Evidence from an Emerging Market, *The Journal of Finance* 53, 2001–2027.
- Du, Songzi, and Haoxiang Zhu, 2017, What is the Optimal Trading Frequency in Financial Markets?, *The Review of Economic Studies* 84, 1606–1651.
- Easley, David, Nicholas M. Kiefer, and Maureen O’hara, 1996, Cream-Skimming or Profit-Sharing? The Curious Role of Purchased Order Flow, *The Journal of Finance* 51, 811–833.
- Economides, Nicholas, 1996, The economics of networks, *International Journal of Industrial Organization* 14, 673–699.

- Ellis, Katrina, Roni Michaely, and Maureen O'Hara, 2000, The Accuracy of Trade Classification Rules: Evidence from Nasdaq, *Journal of Financial and Quantitative Analysis* 35, 529–551.
- ESMA, 2016, Order duplication and liquidity measurement in EU equity markets, Technical Report 907, European Securities and Markets Authority.
- Foley, Sean, and Tālis J. Putniņš, 2016, Should we be afraid of the dark? Dark trading and market quality, *Journal of Financial Economics* 122, 456–481.
- Fontnouvelle, Patrick De, Raymond P. H. Fishe, and Jeffrey H. Harris, 2003, The Behavior of Bid-Ask Spreads and Volume in Options Markets during the Competition for Listings in 1999, *The Journal of Finance* 58, 2437–2463.
- Foucault, Thierry, Johan Hombert, and Ioanid Roşu, 2016, News Trading and Speed, *The Journal of Finance* 71, 335–382.
- Foucault, Thierry, Ohad Kadan, and Eugene Kandel, 2013, Liquidity Cycles and Make/Take Fees in Electronic Markets, *The Journal of Finance* 68, 299–341.
- Foucault, Thierry, Roman Kozhan, and Wing Wah Tham, 2017, Toxic Arbitrage, *The Review of Financial Studies* 30, 1053–1094.
- Foucault, Thierry, and Albert J. Menkveld, 2008, Competition for Order Flow and Smart Order Routing Systems, *The Journal of Finance* 63, 119–158.
- Foucault, Thierry, and Sophie Moinas, 2018, Is Trading Fast Dangerous?, Technical Report 18-881, Toulouse School of Economics (TSE).
- Gajewski, Jean-François, and Carole Gresse, 2007, Centralised order books versus hybrid order books: A paired comparison of trading costs on NSC

- (Euronext Paris) and SETS (London Stock Exchange), *Journal of Banking & Finance* 31, 2906–2924.
- Goldstein, Michael A., Edith S. Hotchkiss, and Erik R. Sirri, 2007, Transparency and Liquidity: A Controlled Experiment on Corporate Bonds, *The Review of Financial Studies* 20, 235–273.
- Goldstein, Michael A., Amy Kwan, and Richard Philip, 2018, High-Frequency Trading Strategies, SSRN Scholarly Paper ID 2973019, Social Science Research Network, Rochester, NY.
- Goldstein, Michael A., and Edward F. Nelling, 1999, Market Making and Trading in Nasdaq Stocks, *Financial Review* 34, 27–44.
- Gresse, Carole, 2017, Effects of lit and dark market fragmentation on liquidity, *Journal of Financial Markets* 35, 1–20.
- Hagströmer, Björn, and Lars Nordén, 2013, The diversity of high-frequency traders, *Journal of Financial Markets* 16, 741–770.
- Hamilton, James L., 1979, Marketplace Fragmentation, Competition, and the Efficiency of the Stock Exchange, *The Journal of Finance* 34, 171–187.
- Han, Jungsuk, Mariana Khapko, and Albert S. Kyle, 2014, Liquidity with High-Frequency Market Making, SSRN Scholarly Paper ID 2416396, Social Science Research Network, Rochester, NY.
- Harris, Larry, 2013, What to Do about High-Frequency Trading, *Financial Analysts Journal* 69.
- Hasbrouck, Joel, and Gideon Saar, 2013, Low-latency trading, *Journal of Financial Markets* 16, 646–679.

- He, Peng William, Elvis Jarnecic, and Yubo Liu, 2015, The determinants of alternative trading venue market share: Global evidence from the introduction of Chi-X, *Journal of Financial Markets* 22, 27–49.
- Hendershott, Terrence, and Charles M. Jones, 2005, Island Goes Dark: Transparency, Fragmentation, and Regulation, *The Review of Financial Studies* 18, 743–793.
- Hendershott, Terrence, Charles M. Jones, and Albert J. Menkveld, 2011, Does Algorithmic Trading Improve Liquidity?, *The Journal of Finance* 66, 1–33.
- Hendershott, Terrence, and Pamela C. Moulton, 2011, Automation, speed, and stock market quality: The NYSE’s Hybrid, *Journal of Financial Markets* 14, 568–604.
- Hoffmann, Peter, 2014, A dynamic limit order market with fast and slow traders, *Journal of Financial Economics* 113, 156–169.
- Holden, Craig W., and Stacey Jacobsen, 2014, Liquidity Measurement Problems in Fast, Competitive Markets: Expensive and Cheap Solutions, *The Journal of Finance* 69, 1747–1785.
- Hu, Edwin, 2018, Intentional Access Delays, Market Quality, and Price Discovery: Evidence from IEX Becoming an Exchange, SSRN Scholarly Paper ID 3195001, Social Science Research Network, Rochester, NY.
- Huang, Roger D., and Hans R. Stoll, 1996, Dealer versus auction markets: A paired comparison of execution costs on NASDAQ and the NYSE, *Journal of Financial Economics* 41, 313–357.
- Huang, Roger D., and Hans R. Stoll, 1997, The Components of the Bid-Ask

- Spread: A General Approach, *The Review of Financial Studies* 10, 995–1034.
- IIROC, 2011, Notice 11-0043: Guidance on “Locked” and “Crossed” Markets.
- IIROC, 2016, Notice 16-0022: Guidance on Time Synchronization.
- Jain, Pankaj K., Pawan Jain, and Thomas H. McInish, 2016, Does high-frequency trading increase systemic risk?, *Journal of Financial Markets* 31, 1–24.
- Jarnecic, Elvis, and Mark Snape, 2014, The Provision of Liquidity by High-Frequency Participants, *Financial Review* 49, 371–394.
- Jovanovic, Boyan, and Albert J. Menkveld, 2016, Middlemen in Limit Order Markets, SSRN Scholarly Paper ID 1624329, Social Science Research Network, Rochester, NY.
- Kervel, Vincent Van, and Albert J. Menkveld, 2019, High-Frequency Trading around Large Institutional Orders, *The Journal of Finance* 74, 1091–1137.
- Kirilenko, Andrei, Albert S. Kyle, Mehrdad Samadi, and Tugkan Tuzun, 2017, The Flash Crash: High-Frequency Trading in an Electronic Market, *The Journal of Finance* 72, 967–998.
- Kirilenko, Andrei A., and Gui Lamacie, 2015, Latency and Asset Prices, SSRN Scholarly Paper ID 2546567, Social Science Research Network, Rochester, NY.
- Korajczyk, Robert A., and Dermot Murphy, 2019, High-Frequency Market Making to Large Institutional Trades, *The Review of Financial Studies* 32, 1034–1067.

- Kwan, Amy, Ronald Masulis, and Thomas H. McInish, 2015, Trading rules, competition for order flow and market fragmentation, *Journal of Financial Economics* 115, 330–348.
- Larrymore, Norris L., and Albert J. Murphy, 2009, Internalization and Market Quality: An Empirical Investigation, *Journal of Financial Research* 32, 337–363.
- Laughlin, Gregory, Anthony Aguirre, and Joseph Grundfest, 2014, Information Transmission between Financial Markets in Chicago and New York, *Financial Review* 49, 283–312.
- Lee, Charles M. C., and Mark J. Ready, 1991, Inferring Trade Direction from Intraday Data, *The Journal of Finance* 46, 733–746.
- Lee, Tomy, 2019, Latency in fragmented markets, *Review of Economic Dynamics* .
- Lewis, Michael, 2015, *Flash Boys: A Wall Street Revolt*, first edition (W. W. Norton & Company, New York).
- Li, Wei, 2018, High Frequency Trading with Speed Hierarchies, SSRN Scholarly Paper ID 2365121, Social Science Research Network, Rochester, NY.
- Lo, Andrew W., and A. Craig MacKinlay, 1988, Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test, *The Review of Financial Studies* 1, 41–66.
- Macey, Jonathan, and David Swensen, 2017, Recovering the Promise of the Orderly and Fair Stock Exchange, *Faculty Scholarship Series* .
- Maglaras, Costis, Ciamac C. Moallemi, and Hua Zheng, 2015, Optimal Execution in a Limit Order Book and an Associated Microstructure Market



- Impact Model, SSRN Scholarly Paper ID 2610808, Social Science Research Network, Rochester, NY.
- Malceniece, Laura, Kārlis Malcenieks, and Tālis J. Putniņš, 2019, High frequency trading and comovement in financial markets, *Journal of Financial Economics* .
- Malinova, Katya, and Andreas Park, 2015, Subsidizing Liquidity: The Impact of Make/Take Fees on Market Quality, *The Journal of Finance* 70, 509–536.
- Malinova, Katya, and Andreas Park, 2017, Does high frequency trading add noise to prices? 40.
- McInish, Thomas H., and James Upson, 2013, The Quote Exception Rule: Giving High Frequency Traders an Unintended Advantage, *Financial Management* 42, 481–501.
- Menkveld, Albert J., 2013, High frequency trading and the new market makers, *Journal of Financial Markets* 16, 712–740.
- Menkveld, Albert J., 2016, The Economics of High-Frequency Trading: Taking Stock, *Annual Review of Financial Economics* 8, 1–24.
- Menkveld, Albert J., 2018, High-Frequency Trading as Viewed through an Electron Microscope, *Financial Analysts Journal* 74, 24–31.
- Menkveld, Albert J., and Marius A. Zoican, 2017, Need for Speed? Exchange Latency and Liquidity, *The Review of Financial Studies* 30, 1188–1228.
- Murray, Hamish, Thu Phuong Pham, and Harinder Singh, 2016, Latency reduction and market quality: The case of the Australian Stock Exchange, *International Review of Financial Analysis* 46, 257–265.

- O'Hara, Maureen, 2015, High frequency market microstructure, *Journal of Financial Economics* 116, 257–270.
- O'Hara, Maureen, and Mao Ye, 2011, Is market fragmentation harming market quality?, *Journal of Financial Economics* 100, 459–474.
- Pagano, Marco, 1989, Trading Volume and Asset Liquidity, *The Quarterly Journal of Economics* 104, 255–274.
- Pagnotta, Emiliano S., and Thomas Philippon, 2018, Competing on Speed, *Econometrica* 86, 1067–1115.
- Riordan, Ryan, and Andreas Storkenmaier, 2012, Latency, liquidity and price discovery, *Journal of Financial Markets* 15, 416–437.
- Roşu, Ioanid, 2019, Fast and slow informed trading, *Journal of Financial Markets* 43, 1–30.
- Shkilko, Andriy, and Konstantin Sokolov, 2016, Every Cloud Has a Silver Lining: Fast Trading, Microwave Connectivity and Trading Costs, SSRN Scholarly Paper ID 2848562, Social Science Research Network, Rochester, NY.
- van Kervel, Vincent, 2015, Competition for Order Flow with Fast and Slow Traders, *The Review of Financial Studies* 28, 2094–2127.
- Wang, Xin, 2018, Why Do Stock Exchanges Compete on Speed, and How?, SSRN Scholarly Paper ID 3069529, Social Science Research Network, Rochester, NY.
- Yao, Chen, and Mao Ye, 2018, Why Trading Speed Matters: A Tale of Queue Rationing under Price Controls, *The Review of Financial Studies* 31, 2157–2183.

Ye, Mao, Chen Yao, and Jiading Gai, 2013, The Externalities of High Frequency Trading, SSRN Scholarly Paper ID 2066839, Social Science Research Network, Rochester, NY.

Zhang, S. Sarah, 2018, Need for speed: Hard information processing in a high-frequency world, *Journal of Futures Markets* 38, 3–21.

Zhu, Haoxiang, 2014, Do Dark Pools Harm Price Discovery?, *The Review of Financial Studies* 27, 747–789.