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

Haoming Chen
University of New South Wales

Sean Foley
University of Sydney

Michael Goldstein
Babson College

Thomas Ruf[1]
University of New South Wales

September 15, 2017

Abstract

We examine the introduction of an asymmetric, randomized speed bump that exempts certain limit orders, allowing low-latency liquidity providers to avoid order-flow driven adverse selection by reacting to activity on other venues. The speed bump segments order flow and increases profits for fast liquidity providers on that venue at the expense of other liquidity providers and aggregate market quality. The negative effects are concentrated in stocks more exposed to immediate adverse selection ex-ante. Our findings have implications for the speed bump debate and speed differentials more generally, as well as the regulation of market linkages across fragmented trading venues.

Key words: market design, speed bump, market quality, fragmentation, adverse selection

---

[1] Contact Author: t.ruf@unsw.edu.au. We thank the Securities Industry Research Centre of Asia-Pacific (SIRCA) for data. Haoming Chen thanks the Capital Markets Cooperative Research Centre (CMCRC) for funding, and was a visiting researcher at the Ontario Securities Commission (OSC) for part of the time spent on this study. The opinions expressed in this study are our own and do not necessarily reflect the views of the OSC or the CMCRC. We thank Reena Aggarwal, Mike Aitken, Angelo Aspris, Robert Battalio, Alejandro Bernales, Bryan Blake, Michael Brolley, Torstein Braaten, Eric Budish, Sabrina Buti (EFA discussant), Doug Clark, Thierry Foucault, Austin Gerig, Carole Gresse, Edwin Hu, Charles Jones (NBER Discussant), Maureen Jensen, Dan Kessous, Amy Kwan, Katya Malinova, Sophie Moinas, Bernt Odegaard, Andreas Park (NFA Discussant), Richard Phillip, Talis Putnins, Dagfinn Rime, Ryan Riordan, Ioanid Rosu, Jos Schmitt, Andriy Shkilko, Elvira Sojli, Vincent van Kervel, Rene Wells, Pradeep Yadav, Yesha Yadav, Lei Yu (CICF discussant), Bart Yueshen, Haoxiang Zhu, and Marius Zoican as well as seminar participants at the EFA 2017, CICF 2017, Ontario Securities Commission, Financial Conduct Authority UK, ITG MatchNow, Chi-X Canada, Aequitas Neo, the NFA 2016, the National Bureau of Economic Research (NBER) Market Microstructure working group, the Asian Finance Association, BI Norwegian Business School, Universidad de Chile and Catholica University, Chile, for their thoughtful feedback. We are also indebted to the staff at the OSC for sharing their knowledge on Canadian equity market structure. The internet appendix that accompanies this paper may be found at https://goo.gl/3umXjz.

"Speed bumps will now be a competitive tool for exchanges, but each exchange can build its speed bump to target a different audience ... And the SEC's national market system knits all of those exchanges and their disparate speed bumps awkwardly into one market, where everyone is required to route to the exchange showing the best price, no matter how delayed it might be, or for what reason. It's kind of a mess!"

– Matt Levine, BloombergView (8/31/2016)

1 Introduction

Much of the technological innovation in today's financial markets is driven by the incentive of market participants to be faster than their competition. Faster traders capture most of the profits from liquidity provision (Rosu, 2016) and impose adverse selection costs on relatively slower counterparts (Li, 2014; Baldauf and Mollner, 2016), either by picking off stale orders (Aquilina, Foley, O'Neill and Ruf, 2016), or updating limit orders faster than others in response to information revealed on another venue (van Kervel, 2015). This desire for speed motivates technology suppliers to provide faster access to traders (O'Hara, 2015), leading to a perpetual arms race.[2] Over time, the marginal increase in speed has become ever smaller, but what remains constant is that speed provides an advantage to those who possess it over those who do not with the ultimate goal being a reduction in adverse selection costs.

Exchange operators have responded to those needs by introducing innovations such as inverted pricing structures (Maglaras, Moallemi and Zheng, 2015; Tham, Sojli, and Skjeltorp, 2017), dark trading with sub-penny price improvement (Kwan, Masulis and McInish, 2015), co-location (Brogaard, Hagströmer, Nordén and Riordan, 2015) and latency upgrades (Menkveld and Zoican, 2017). Lately, a new tool has emerged in this arena: *speed bumps*. These represent an intentional, systematic delay between an order's arrival at the exchange and that order's entry to the matching engine. Speed bumps can be symmetric (i.e.

[2] This race is nearing its theoretical limit (Angel, 2014), e.g. microwave-based transmissions between Chicago and New York are approaching the speed of light (Laughlin, Aguirre and Grundfest, 2014; Shkilko and Sokolov, 2016). Many argue that this race not only reduces the incentives to collect valuable information, hindering the long-term pricing efficiency of markets and increasing the costs of liquidity provision, but is also socially wasteful due to overinvestment in trading infrastructure (Menkveld, 2014; Budish, Cramton and Shim, 2015; Biais, Foucault and Moinas, 2015).

apply universally to all orders) or asymmetric (i.e. affect only some orders or participants), and can add either a deterministic or randomized delay to processing.

Recent theoretical work (Harris, 2013; Budish, Cramton and Shim, 2015; Baldauf and Mollner, 2015; Brolley and Cimon, 2017) proposes certain structural delays as a means to counter the perpetual arms race. While fully symmetric delays appear to be ineffective, Baldauf and Mollner (2015) suggest that asymmetric, "non-cancellation" delays can be effective against sniping, but point out that additional randomization of delays may impede the efficacy of smart order routers in accessing consolidated liquidity across venues (O'Hara, 2015).[3]

In this paper, we investigate the introduction of an asymmetric, randomized speed bump on the Canadian exchange TSX Alpha on September 21, 2015, which both delays incoming messages for (almost) all market and limit orders by 1-3 milliseconds, but provides the option to pay a higher fee to enter and cancel limit orders *without* experiencing the delay.[4] We conduct a detailed analysis of the behavior of liquidity providers on Alpha, their usage of the delay-exempt order type, and the ability of liquidity demanders to continue to access liquidity on Alpha. We document the consequences for the segmentation of beneficial order flow between venues, changes in transaction costs and overall market quality.

Particularly since speed bump venue IEX, of 'Flash Boys' fame, was approved as a U.S. exchange in June 2016, several competing trading venues have proposed their own version of intentional delays.[5] These designs are remarkably similar along key dimensions. IEX, NYSE American, CHX as well as TSX Alpha

---

[3] Symmetric delays do not alter the incentives of the arms race as traders still compete to be first in line to reach the exchange in order to either pick off (snipe) or cancel outstanding stale limit orders. Symmetric, randomized delays potentially create additional problems as they incentivize traders to submit redundant orders, each essentially representing a lottery ticket for a short delay (Budish et al., 2015).

[4] Exempt from the speed bump are Post-Only limit orders (non-marketable at submission) with a prescribed minimum size that varies by stock and over time (at least 500 shares)

[5] The very first speed bump mechanism was proposed by NASDAQ PHX in 2012, involving a 5 millisecond delay, but was subsequently withdrawn.

in Canada are asymmetric[6] (or discriminatory), slowing down all incoming messages *except for* a specific subset of resting limit orders, the users of which pay the exchange for that privilege.[7],[8] In light of the spate of applications for new speed bumps, empirical evidence is crucial to understand whether these new mechanisms affect liquidity provision and how the elements of current may interact to produce unintended outcomes.

TSX Alpha's speed bump effectively allows low-latency participants1-3ms to observe trading activity on other venues, cancelling standing limit orders prior to any anticipated executions.[9] Orders that access multiple venues likely originate from smart order routers (SOR) of institutional traders and tend to pose a higher risk of adverse selection.[10] Market orders access only Alpha cannot be avoided, but are more likely to be from uninformed participants (e.g. retail orders) than SOR sprays and are thus more profitable on average.[11] The speed bump affords sufficiently fast liquidity providers the opportunity to fade from (likely institutional) SOR sprays which generate immediate adverse selection. The increase in realized spreads from trading (relatively more) with uninformed flow, while avoiding (some) large and informed flow represents the economic value of a millisecond head start.

Our findings for Alpha in the post-event period are consistent with this intuition. About two thirds of initially displayed liquidity on Alpha fades after the introduction of the speed bump without being

---

[6] IEX's speed bump is sometimes considered symmetric as it slows down all lit orders, but exempts dark pegged orders. 80 percent of its trading is dark.

[7] In the case of the updated CHX LEAD proposal, the payment comes in the form of more stringent quoting requirements rather than through a fee differential.

[8] The only exception to this trend that we are aware of is Aequitas in Canada (with a 1-2% market share of lit volume) that specifically slows down HFTs market orders.

[9] 'Last look' is a common, but controversial feature in currency markets, giving brokers a brief time window in which to review and potentially reject a trade if they deem the transaction price to no longer reflect current market conditions.

[10] Orders that access multiple venues tend to be either large or informed or both. Large orders, almost by definition, need to be split across venues to take advantage of all available liquidity. Informed traders may be worried about their trading intentions being revealed to the market too soon and optimally attempt to execute against all available liquidity at once. Liquidity providers would like to avoid either type because both portend future order flow in the same direction, imposing adverse selection. Large orders may be part of an even larger parent order by the same trader, while informed orders predict same direction flow from other traders.

[11] Mechanically, orders that remove liquidity from many venues at the same time are more likely to move prices than orders that remove liquidity from only one venue, particularly if there are multiple venues at the best quote.

executed once market orders start arriving on any venue (up from 15% in the pre-period). For other venues, this figure either remains stable or declines. We show that this dramatic change is due to low latency liquidity providers canceling Post-Only limit orders after trades begin on other venues, but before (most) market orders can reach them.

We also document a change in the composition of traders on Alpha, with institutional trading declining by half, while low price impact order flow more than doubles on Alpha. Using a novel measure – *Segmentation Imbalance* – we show that Alpha is able to capture a larger fraction of trades which do not displace market-wide liquidity at the top of the book, while avoiding those that do. Other exchanges experience a commensurate increase in order flow that instantaneously moves prices. We provide causal evidence that it is the speed bump, or rather its asymmetric application, which allows this segmentation to occur.

Consistent with the ability to capture single-venue trades while avoiding multi-venue trades, we find a significant decline in adverse selection on Alpha (0.6 cents), which drives an increase in realized spreads of over 1 cent, greatly improving the profitability of fast liquidity providers on Alpha. Conversely, we find that other venues see an increase in adverse selection costs (0.3 cents) due to receiving a larger proportion of correlated order flow.[12]

While it is not clear a priori that overall market quality should be affected, we document significant deleterious effects overall, with quoted and effective spreads widening by around 0.2 cents. We construct novel measures of consolidated order book fragility: the fraction of trading that exhausts the top of book; and the fraction of trading that accesses depth behind the NBBO. Both measures increase significantly, implying an increased reluctance of liquidity providers market-wide to absorb demand. While

[12] Overall, these findings are in line with a much earlier literature on the impact of the segmentation of order flow through payment for order flow schemes (Easley, Kiefer and O'Hara, 1996; Chakravarty and Sarkar, 2002) and also with more recent studies on the segmentation of uninformed order flow imposed by dark trading (Zhu, 2014; Comerton-Forde and Putnins, 2015).

improvements in NBBO depth are observed, this increase may be misleading as it is driven by the requirement to post a large minimum number of shares on Alpha, which (protected by the speed bump) are more likely to fade.

Upon further inspection, we find the widening spreads to be exclusive to stocks where trading has a tendency to displace the NBBO: higher priced, less tick-constrained, with thinner quotes, i.e. collectively stocks with a high risk of immediate adverse selection (IAS) for liquidity providers. While quoted and effective spreads in low-IAS stocks do not change, high-IAS stocks experience increases of 0.3-0.45 cents. Realized spreads on Alpha increase by a full 2 cents for these stocks (vs. 0.4 cents).

The very different effects on market quality in the cross-section of stocks should highlight the perils of a 'one-size-fits-all' approaches to market design, but also provides some interesting insights into how the goal of lowering adverse selection costs is achieved by different designs and for a different set of stocks: We confirm and elaborate on recent findings in the literature (Battalio et al., 2016) that inverted venues tend to attract and segment order flow in low-IAS stocks, while being less successful for high-IAS stocks. Alpha's speed bump complements inverted venues by helping liquidity providers avoid adverse selection in the latter group. We show that Alpha by being an inverted venue with speed bump achieves near uniform levels of segmentation of order flow in the cross-section, in contrast to typical purely inverted venues.

To separate the effects of the speed bump from those related to Alpha's switch to an inverted fee structure we compare Alpha to an inverted venue without a speed bump. We utilize broker level data to show that fast liquidity providers profit more from trading on Alpha than CX2, while slower liquidity providers do worse. We attribute this difference to the 1-3 ms head start afforded by the speed bump. We also find that the removal of Alpha from the Order Protection Rule has little effect on traders' routing decisions. Finally, we use two matched U.S. samples to show that our results are isolated to the Canadian market.

In summary, this study provides the first piece of evidence on the impact of intentionally slowing down some participants in a fragmented environment. We document a key insight into what drives fleeting liquidity in today's fast, fragmented markets: participants with speed advantages are able to observe (large) traders actions on other venues, cancelling standing limit orders faster than traders are able to access them. Our evidence suggests that advance knowledge of institutional investors' trading intentions (in a probabilistic sense), even at the millisecond granularity, is valuable: delays as short as one millisecond can enable substantial (costly) information leakage across venues.

In addition to these findings, we also develop a number of new techniques which can be constructed with any public data source to examine market linkages and the accessibility of liquidity in fast, fragmented markets.[13] First, we develop a methodology to benchmark time synchronization across venues, allowing us to aggregate related trades on different venues into trade strings which we use to examine low latency, cross-market liquidity dynamics. Second, we propose two innovative classification schemes to analyze a venue's order flow composition. These identify a) trade strings that cause instantaneous adverse selection costs ("*depleting trade strings*"); and b) trade strings that execute across multiple venues, likely to have originated from a smart order router (henceforth SOR). We further quantify a venue's ability to segment beneficial (from the perspective of the liquidity provider) non-depleting trades away from detrimental depleting trades using a new measure called *Segmentation Imbalance*.

Third, we develop a number of novel measures that describe order fragility during executions. "*Quote Fade"* measures the fraction of liquidity at the NBBO that is cancelled during a trade execution, providing us with an indication of what fraction of posted liquidity is fleeting (van Kervel, 2015). "*Deplete Best"* measures the fraction of trading volume that exhausts liquidity at either the bid or offer (i.e. depleting trades) within 20ms of each transaction, proxying for the likelihood that a liquidity provider will experience immediate adverse selection. Lastly, "*Access Next"* measures the fraction of trading that

[13] A number of recent studies examine the issue of market linkages across fragmented trading venues; see, for example van Kervel (2015), and Malinova and Park (2016).

occurs behind the NBBO. A higher value indicates liquidity demanders had to "walk the book" more frequently, increasing transaction costs and effective spreads.

The remainder of this paper is organized as follows. Section 2 outlines the institutional details of the Canadian trading landscape and the design changes on Alpha. Section 3 describes the data and methodology. Section 4 demonstrates that these design changes lead to a segmentation of order flow and changes in transaction costs across exchanges. Section 5 directly contrasts transaction costs on Alpha with the other major inverted venue CX2. Section 6 presents the cross-sectional analysis. Section 7 investigates the implications for consolidated market quality. Section 8 investigates alternative explanations, while Section 9 concludes and discusses implications for regulators and market participants.

2 Institutional Details

Similar to the United States, Canadian equities trading is fragmented across multiple venues, with six lit and three dark trading venues.[14] Securities are listed on the Toronto Stock Exchange, operated by the TMX Group, which retains approximately 60 percent of market share by trading activity. The TMX Group also operates Alpha and TMX Select, which was decommissioned once the changes on Alpha occurred. In Canada, NASDAQ operates both Chi-X Canada and CX2; notably, CX2 has an inverted maker/taker pricing structure identical to that of new Alpha. Other venues include Omega, Pure Trading, Aequitas Neo, Aequitas Lit and three dedicated continuous dark pools, Match Now, Instinet and Liquidnet. Most lit exchange venues offer price-broker-time priority rather than pure price-time priority as is common in the United States. Across venues, price priority is enforced via "trade-through

[14] In the United States, there are currently over 11 lit markets with publicly displayed limit order books, 44 dark trading venues (without pre-trade transparency) and approximately 200 broker-operated alternative trading systems (ATS) competing for order flow. Non-lit trading accounts for 35 percent of total volume in the United States (Tabb Group and Rosenblatt Securities), but only for 6 percent in Canada.

prohibition." Table A.1 of the appendix reports some relevant characteristics of the four major Canadian markets in our sample such as the prevailing fee structure.

Unlike the United States, internalization of retail order flow in Canada has been significantly constrained. 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 mechanism has prevented the growth of retail internalizing venues common in the United States, which account for around 22 percent of trading there (Kwan et al., 2015).[15] As a result of this regulation, and the subsequent banning of payment-for-order-flow, retail orders remain predominantly on-exchange in Canada.

2.1 The Alpha Speed Bump

Alpha Exchange was launched in 2008 and was acquired by the TMX Group in 2012. On the 21st of September 2015, the trading venue was relaunched as TSX Alpha with several changes, including:

1. A randomized speed bump of 1-3 milliseconds for all orders except Post-Only limit orders.[16]
2. An inverted maker-taker pricing model.
3. Orders on Alpha are no longer subject to the Order Protection Rule.[17]

For brevity, we will sometimes refer to Alpha prior to the introduction of the speed bump as "old" Alpha, and as "new" Alpha from September 21, 2015 onwards.

The market structure facing liquidity demanders after Alpha's relaunch is depicted in the diagram below. Institutional investors who require more liquidity than what is displayed on any single trading venue

[15] In Canada, payment for order flow is prohibited and meaningful price improvement rules apply to trades on dark venues, including regulations designed to ensure orders sent to the United States would also be subject to minimum price improvement regulations. As such, unlike in the United States, internalization is not a common practice.

[16] Post-Only orders require a stock-specific minimum size. In our sample, most stocks require 500 or 700 shares; the maximum is 5,000. A detailed list of minimum Post Only sizes by security is updated monthly and available at http://api.tmxmoney.com/en/research/minpo.csv

[17] Canada's Order Protection Rule is akin to the trade-through prohibition (Rule 611 Reg NMS) in the United States, except that it applies to every price level, not just the NBBO.

typically utilize a SOR to spray marketable orders across multiple trading venues simultaneously, efficiently accessing consolidated quoted depth at the national best bid or offer price (O'Hara and Ye, 2011). Alpha's asymmetric, randomized speed bump for incoming marketable orders *but not limit order entries or cancellations* potentially enables its liquidity suppliers to observe the first legs of any large SOR spray being executed on other venues and cancel their limit orders to avoid adverse selection costs.

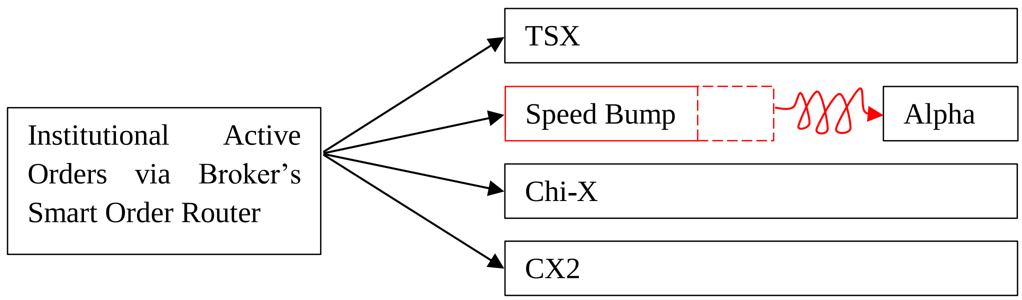

Overall, the optimal trading strategy, from the liquidity takers perspective, may be to send all orders to Alpha when the desired quantity can be filled there alone, but otherwise submit to all venues simultaneously and expect low fill rates on Alpha. Importantly, such concerns are much less relevant for retail traders who demand less than an entire price level.

The creation of the new Alpha resulted in an immediate and significant reduction in market share, from just below 15% to around 4% in the week after the relaunch; however, by December 2015, Alpha's market share had climbed back to around 8-9%. This reduction is consistent with liquidity suppliers being unwilling to pay to post on Alpha, and subsequently providing their liquidity on the remaining venues.

2.2 Comparison with other speed bump designs

Likely the best known speed bump is IEX's "coil in a shoebox" made famous by Michael Lewis' 2014 book 'Flash Boys', but other proposals are quite similar in many regards. IEX, CHX and NYSE American (which launched its speed bump at the end of July 2017) all feature an asymmetric delay of 350 microseconds for all order messages except some clearly defined set of liquidity providing limit orders.

For IEX, the exemption covers so-called dark, discretionary and primary pegged orders which are automatically re-priced by an algorithm that, according to IEX, recognizes 'crumbling quotes' on other markets before they happen.[18] NYSE American also exempts certain types of dark pegged orders. CHX recently updated their LEAD application towards granting the exemption to visible limit orders of approved liquidity providers. TSX Alpha's exemption is very similar to these in that it only exempts visible Post-Only (non-marketable) limit orders of a minimum size.

Also common across these designs is the feature that the benefactors of the exemptions incur a cost compared to the users of other order types: Alpha has an explicit fee difference; IEX effectively has moved towards the same model by waiving fees for lit orders under certain conditions; CHX requires individual approval and minimum liquidity provision.[19] The common asymmetry in favor of liquidity providers is most easily explained by their desire to protect against adverse selection and their willingness to pay for the privilege, increasing exchange revenue in the process.

There are also some differences between Alpha and the U.S. designs, e.g. the removal from order protection as well as the transition to an inverted fee structure. However, we will argue and show in empirical tests that these matter actually surprisingly little on their own, albeit the inverted fee structure is likely helpful in reinforcing the order flow segmentation that the speed bump tries to achieve. Where Alpha and the U.S. speed bump designs deviate most dramatically is the fact that Alpha's delay is not deterministic but randomized between 1 and 3 milliseconds. Consistent with the model of Baldauf and Mollner (2015), we hypothesize that the random nature of the speed bump may render SORs less effective and lead to segmentation of liquidity between venues.

2.3 Switch to inverted fee structure

[18] Some commentators consider IEX to be a symmetric speed bump, which is technically true when one ignores the dark pool of IEX. However, since 80% of IEX's 2% market share originates in the dark, one might debate this classification.

[19] There was little information on the specifics of NYSE American's new fee structure at the time of writing.

To attract retail traders, new Alpha introduces an inverted maker-taker model (also called a "taker-maker" model), providing a rebate to the demander of liquidity, paid for by the liquidity supplier. Without intermarket time priority, liquidity demanders have an incentive to first route marketable orders to venues with the lowest fee or highest rebate (Battalio, Corwin and Jennings, 2016), particularly for brokers that do not pass on the taker rebate to the client. Such a flat fee structure is common for retail brokers in Canada (Brolley and Malinova, 2013).

Prior to Alpha's fee change, three venues (CX2, TMX Select, Omega) already adopted this fee structure, and with TMX Select ceasing to operate, there remain 3 inverted venues in the post period as well. As such we do not expect any significant market-wide impact from the change in Alpha's fee model alone. In the empirical sections of the paper, we find ample evidence in support of this expectation, using a) trader classifications, b) the presence vs. absence of delay-exempt limit orders, c) cross-sectional differences across stocks as well as d) by investigating directly the effect of previous entries of inverted venues.

2.4 Removal from order protection

Since the approval of IEX as an exchange and the subsequent emergence of competing proposals by other exchanges to adopt their own versions of speed bumps, calls in the United States have become louder to re-think one of the central tenets of the 12-year old Reg NMS market framework: the order protection rule (Rule 611) forces brokers to search for the best price across all venues, no matter how delayed by speed bumps or otherwise.[20] In stark contrast, the Canadian regulator OSC explicitly made Alpha's removal from the OPR the basis for the approval of the speed bump. [21]

The most obvious impact from removing a venue from protection means that traders can choose to ignore that venue. However, both in the United States and Canada, brokers are obliged to strive for best

[20] For example, Matt Levine's BloombergView article quoted at the start of the paper. Also, http://tabbforum.com/opinions/when-was-the-last-air-raid-or-why-wait-to-eliminate-the-order-protection-rule

[21] See OSC Bulletin (2015), Volume 38, Issue 16, p. 4045, "Notice of Commission Approval of Proposed Changes to Alpha Exchange Inc."

execution. That is, if they were to ignore a venue with visibly better prices, they would need to explain why this was beneficial for overall execution quality of clients' orders. Conditional on the venue offering competitive quotes it is in clients' and brokers' interest to not categorically ignore it. If on the other hand, the venue is uncompetitive by quoting behind best, then order protection is moot as the venue will not receive orders, protected or not. As long as both sides to a trade find it beneficial to meet on a given venue they will choose to do so, regardless of protection status. The fact that Alpha's market share recovers to 8 percent within a few weeks of the relaunch would indicate that some traders find Alpha still beneficial.

In addition, as we will show later on, new Alpha essentially is never alone at best. Thus, even if the order protection rule was still in force for Alpha, there are few instances where it even can be violated and thus be relevant. We run a test that looks for specific instances where quotes on any venue (not just Alpha) have been traded through and find that the fraction of "trade throughs" on Alpha is not significantly higher than on the other inverted venue, CX-2.

3 Data and Methodology

3.1 Data

The data for this study is sourced from Thomson Reuters Tick History (TRTH), supplied by the Securities Industry Research Centre of Asia Pacific (SIRCA). TRTH provides data for seven Canadian trading venues, namely TSX, Alpha, Chi-X, CX2, TMX Select, Omega and Pure Trading. This encompasses all Canadian trading venues with partial or full pre-trade transparency, except Aequitas NEO and Aequitas Lit, which together account for less than one percent of trading activity. Pure Trading also has a market share of less than one percent, and is dropped from the analysis. Lastly, both TMX Select and Omega currently use a legacy data feed, with time stamp inaccuracies that often exceed 200-300 milliseconds, making it impossible to precisely calculate NBBO prices and volumes, and each account for less than 3 percent of trading volume. Weighing accuracy and quality against completeness, we exclude these two venues as well. This leaves TSX, Alpha, Chi-X and CX2 as the venues of interest in this paper, accounting for more than 90% of all trading in Canada. Our observation period runs from March 23, 2015 to March 18, 2016, covering six months on either side of Alpha's market structure change. We discard all U.S. exchange holidays including the two days where exchanges close early (Friday after Thanksgiving and Christmas Eve) because Canadian markets are much less liquid on those days, as well as August 24, 2015, when a flash crash event in the United States caused abnormally high volatility.

We include all member stocks of the S&P TSX Composite Index at the time of the event. This index contains the 250 largest Canadian equities, representing approximately 70 percent of Canadian market capitalization. We discard four stocks due to delisting, merger activity or stock splits within 3 months of our event so as to not confound the effects of these idiosyncratic events with those of Alpha's relaunch.[22] In addition, 14 stocks are excluded which have a time-weighted midpoint price of $2 or less for the majority of days in either the pre- or post-event period. These stocks are consistently tick-constrained with a natural spread likely well below the 1 cent minimum tick size. Our final sample contains 232 equities.

TRTH provides data for each exchange including the state of the limit order book at each quote update, as well as all trade records. The data fields include exchange, security, date, millisecond time stamp, trade price, trade volume, trade qualifiers, buyer and seller broker ID, as well as the price and size for both the bid and ask. [23] We acquire trades and quotes within the same exchange concurrently to preserve ordering within the same millisecond to enable accurate trade direction classification. We restrict our analysis to the trading hours of the TSX listing market, being 9.30am to 4.00pm. We remove trades marked as off-

---

[22] For any equities that conduct a spin off or stock split in the pre-period, we discard data prior to the corporate action only so as to not add noise to the pre-period averages used in the cross-sectional analysis.
[23] Broker identifiers for buyer and sellers are available for TSX and Alpha, unless the broker chose to remain anonymous and forgo participation in broker preferencing. Chi-X does not offer broker preferencing, but some trades contain broker identifiers. Although CX2 offers broker preferencing, TRTH does not include these identifiers. We thank Chi-X Canada for providing us with public broker identifiers for this venue for the year 2015.

market crossings, odd lot trades or midpoint dark trades. We also remove trades with a value above $2 million, even if they do not have off-market qualifiers.[24]

3.2 Traditional Market Quality Metrics

Our empirical methodology creates one dataset containing the trades on each venue and another dataset containing the national best bid and offer (NBBO) prices and depths. We determine the trade initiator by observing whether the trade happened at the best prevailing bid or offer price on that venue. Our approach assigns trade direction with near certainty and avoids the issues associated with the midpoint or tick tests used in previous studies such as Lee and Ready (1991), Ellis, Michaely and O'Hara (2000), Bessembinder (2003) and Holden and Jacobsen (2014), particularly in the context of fragmented markets. A detailed outline of the full methodology including our attribution of trade direction is provided in Section B.1 of the Internet Appendix.[25] This process creates a file containing exchange, symbol, date, millisecond time stamp, price, volume, trade direction, buyer and seller broker ID for each trade.

We use this file to manually reconstruct the NBBO price and size for each security.[26] The NBBO quoted spread is calculated for each stock (i) and day (d) as the difference between the prevailing national best bid (NBB) and national best offer (NBO) prices and is time-weighted throughout each day. We also calculate the NBBO quoted depth as the total volume quoted at the national best bid and offer prices, updated for each quote (q) across all venues, and measured for the total duration for which that quote prevailed ($Alive_q$).

---

[24] Trade qualifiers in the TRTH data may be incomplete, and we are aware of trades exceeding $100 million in the TRTH data without off-market qualifiers. Trades are recorded from the perspective of the liquidity supplier. Therefore, a trade of $2 million would require the liquidity supplier to have submitted a single limit order for $2 million and the liquidity demander to have also simultaneously submitted a single marketable order larger than $2 million. A frequency distribution of large trade sizes is available upon request.

[25] The internet appendix that accompanies this paper may be found at https://goo.gl/3umXjz.

[26] If the NBBO would be locked or crossed, we take the prevailing quotes on the TSX as being the NBBO. This is due to IIROC's Universal Market Integrity Rules, which stipulate that limit orders that would lock or cross with visible orders on another market are not permitted. In the Reuters data, this occurs for short periods of time due to a lack of clock synchronization across venues. Generally, the venues are synchronized to within 20 milliseconds. Appendix B.2 provides further details on benchmarking of cross-venue clock synchronization.

$$NBBO\ Quoted\ Depth_{i,d} = \frac{\sum_{q=1}^{Q}(NBO\ Depth_q + NBB\ Depth_q) * Alive_q}{\sum_{q=1}^{Q} Alive_q} \quad (1)$$

Additionally, we calculate the proportion of time each venue (*v*) displayed quotes at the NBBO, as well as its share of total NBBO depth.

$$\%\ Time\ at\ NBBO_{i,d,v} = \frac{\sum_{q=1}^{Q}(I_{Venue\ v\ at\ NBB} * Alive_q) + \sum_{q=1}^{Q}(I_{Venue\ v\ at\ NBO} * Alive_q)}{2 * \sum_{q=1}^{Q} Alive_q} \quad (2)$$

$$\%\ Depth\ at\ NBBO_{i,d,v} = \frac{\sum_{q=1}^{Q}\left(\frac{Venue\ NBB\ Depth_q + Venue\ NBO\ Depth_q}{Total\ NBB\ Depth_q + Total\ NBO\ Depth_q} * Alive_q\right)}{\sum_{q=1}^{Q} Alive_q} \quad (3)$$

Effective half-spreads are calculated as the difference between the trade price and the prevailing NBBO midpoint. Similar to Conrad and Wahal (2016), realized spreads are compared to trade prices with the NBBO midpoint that prevailed 1, 5, 10 and 20 seconds after each trade. In the interest of brevity, we report 20 second realized spreads. Adverse selection costs (price impact) are computed as the effective spread minus the realized spread. These metrics are volume-weighted across trades by stock (*i*), day (*d*) and venue (*v*), where applicable, as follows:

$$Effective\ Spread_{i,d,v} = 2 * \frac{\sum_{t=1}^{T}\{D_t * (Price_t - Midpoint_t) * Volume_t\}}{\sum_{t=1}^{T} Volume_t} \quad (4)$$

$$Realized\ Spread_{i,d,v} = 2 * \frac{\sum_{t=1}^{T}\{D_t * (Price_t - Midpoint_{t+20sec}) * Volume_t\}}{\sum_{t=1}^{T} Volume_t} \quad (5)$$

$$Adverse\ Selection_{i,d,v} = Effective\ Spread_{i,d,v} - Realized\ Spread_{i,d,v} \quad (6)$$

Following Malinova and Park (2015), we adjust the effective and realized spreads reported throughout this paper for the fee structure prevailing at each venue. Specifically, in markets with maker-taker pricing, effective spreads are increased by the take fee for a net cost of demanding liquidity, while realized spreads are increased by the maker rebate to reflect the net revenue that accrues to liquidity providers. Conversely, on inverted markets, net effective spreads are lower than gross spreads by the take rebate and realized spreads shrink in line with the make fee that liquidity providers pay to the exchange.

3.3 Construction of High Frequency Trade Strings and Order Routing

Motivated by the importance of linkages between markets highlighted by O'Hara (2015), we investigate the ability of liquidity demanders to access quoted liquidity across venues. To this end, we construct several novel metrics that rely solely on readily available trade and quote data and enable us to estimate the impact of phantom liquidity across venues. We first construct high frequency trade "strings" across venues by grouping together all buyer or seller initiated trades for each security occurring within 30 milliseconds of the last trade in the same direction.[27] While timestamps for any individual exchange may exhibit latency, jitter, caching and lack of cross-venue synchronization, strings of trades that occur over short time intervals are likely related. Section B.3 of the Internet Appendix describes the construction of high frequency trade strings in detail and Table A.1 provides summary statistics for these strings.

In the spirit of van Kervel (2015), we define multi-venue sweep trades as those that are part of a string also containing trades on at least one other venue. These trades likely originate from a SOR spray of a single trader that sought to access the consolidated pools of liquidity across multiple venues. Our methodology results in a median length of multi-venue trade strings of 9-17 ms depending on the number

---

[27] Alpha's speed bump enables quote fade at a 1-3 ms horizon. By bucketing over a 30 ms horizon, we pick up everything caused by the 1-3 ms speed bump, as well as potential additional noise. This noise could be in the form of trades, order additions or cancellations; however, our measured trade strings (even at the 99th percentile) are bursts of activity over several milliseconds, both preceded and followed by many seconds without trades. An examination of the distribution of all trades in our sample suggests that noise is very limited relative to the signal. In particular, Appendix Figure B.1 estimates the 95th percentile of jitter across our venues to be 30 ms. Using this conservative benchmark means we can be sure that trades separated by more than 30 ms of each other are highly unlikely to be contemporaneous trades separated by jitter.

of venues involved. This represents a significant improvement in precision over van Kervel's (2015) decision to bucket trades within 100 ms and is comparable with Malinova and Park (2016).

We further distinguish between trade strings (s) that move prices, i.e. deplete the top of the book ($Deplete_s = 1$), and those that do not ($Deplete_s = 0$). A trade string moves prices when, after its execution, all depth on the opposite side of the NBBO is depleted. More precisely, buyer (seller) initiated trades are considered depleting if they originated from a trade string where the national best bid (or offer) price at the end of the string was higher (or lower) than the best price traded during the string. Trade strings that do not displace the entire NBBO depth are called non-depleting. Note that this classification is not necessarily a proxy for trade size, since orders smaller than pre-trade NBBO depth may also displace a price level a large number of cancelations by liquidity suppliers occur following the execution.

To indicate the relative prevalence of depleting vs. non-depleting trades we define $Deplete\ Best_{i,d,v}$ which measures the fraction of trading volume that depletes the top of book measures for stock (i) on day (d) and venue (v):

$$Deplete\ Best_{i,d,v} = \frac{\sum_{s=1}^{S} Trade\ Volume_{i,d,v,s} \times Deplete_s}{\sum_{s=1}^{S} Trade\ Volume_{i,d,v,s}} \tag{7}$$

O'Hara (2015) suggests that in a high-frequency world, one might consider trades that cause prices to move to be informed in the sense that they impose *instantaneous* adverse selection costs on liquidity providers. Thus, an interpretation of a depleting trade is one with high information content. Traders that are informed (at least about their own orders) cause liquidity providers to withdraw/cancel more liquidity. Our definition is akin to the traditional adverse selection metric; however, we are utilizing a virtually instantaneous horizon of twenty milliseconds, rather than a few minutes (Hendershott, Jones and Menkveld, 2011; Carrion, 2013) or seconds (Conrad, Wahal and Xiang, 2015) after the trade.

Among all trade strings (s) that deplete the NBBO ($Deplete_s = 1$), we calculate *NBBO Quote Fade* as the proportion of starting liquidity at the national best offer (bid) price for buyer (seller) initiated trades that did not result in trades. Recall that the starting liquidity in a depleting trade string can either be consumed or withdrawn. A lower bound of zero is placed on the quote fade metric per trade string to account for the fact that it is not possible for more liquidity to "fade" than exists at the start of the trade string.

$$NBBO\ Quote\ Fade_{i,d,v} = 1 - \frac{\sum_{s=1}^{S} Trade\ Volume_{i,d,v,s} \times Deplete_s}{\sum_{s=1}^{S} Max(Start\ Liquidity_{i,d,v,s}, Trade\ Volume_{i,d,v,s}) \times Deplete_s} \tag{8}$$

Given that orders which deplete the order book generate immediate adverse selection, liquidity providers would ideally avoid such orders, and instead capture those which do not deplete the order book. $Segmentation\ Imbalance_{i,d,v}$ measures the extent to which a venue is able to segment non-depleting orders while avoiding depleting orders, and is calculated as that venue's share of the total non-depleting turnover minus its share of the total depleting turnover, per stock-day. A segmentation imbalance of one (minus one) is recorded when a venue interacts exclusively with trades that do not (do) impose instantaneous adverse selection costs. The metric is symmetric around zero, with a higher segmentation imbalance indicating that the venue was more successful at interacting with trades with lower information content.

$$Segmentation\ Imbalance_{i,d,v} = \frac{Non{-}depleting\ turnover_{i,d,v}}{Non{-}depleting\ turnover_{i,d}} - \frac{Depleting\ turnover_{i,d,v}}{Depleting\ turnover_{i,d}} \tag{9}$$

Finally, we calculate the relative proportion of trades that occurred at the next best price behind the national best bid (offer) for seller (buyer) initiated trades, to measure the tendency for trades to walk the book, and take liquidity from the next level below the best. The metric $Access\ Next_{\ i,d,v}$ captures the lack of top-of-book liquidity where liquidity demanders sought to trade large amounts:

$$Access\ Next_{\ i,d,v} = \frac{\sum_{s=1}^{S} Trade\ volume\ at\ level\ 2\ or\ greater_{\ i,d,v,s}}{\sum_{s=1}^{S} Total\ trade\ volume_{i,d,v,s}} \tag{10}$$

3.4 Summary Statistics

Table 1 presents summary statistics for a number of liquidity metrics, transaction cost measures and stock characteristics based on daily measures for all sample stocks for the full one-year period. We provide the mean, median and the 10th and 90th percentile of each variable separately for the six months before and after the introduction of the speed bump.

We observe an increase in the average quoted spread from 3.47 cents to 3.76 cents in the post period for our sample of 232 TSX equities. *Deplete Best*, the fraction of trading that leads to a change in quotes, is relatively stable across post periods at around 60 percent, while there is an increase in *Access Next*, the fraction of trading that accesses liquidity beyond the top of the book, from 13.2% to 14.6%.

Effective spreads are reported net of exchange fees. We observe increases both on Alpha as well as non-Alpha venues, but adverse selection costs seem to move in different directions. While non-Alpha venues experience an increase in adverse selection from 3.48 cents to 3.74 cents, Alpha experiences a substantial decline, from 2.94 cents to 2.13 cents. Liquidity providers on Alpha could potentially pass on this reduced adverse selection cost through lower quoted spreads, "making" new best prices. However, displaying the best price on Alpha would nullify the advantage of the speed bump, as liquidity providers on Alpha would instead be hit first. Consistent with a "matching" rather than "making" of the best price, we find that new Alpha posts a price at the NBBO only 36% of the time, compared to 60% of the time prior to the speed bump introduction.

Posting at prices equal to (or behind) the NBBO optimizes Alpha's liquidity suppliers' ability to avoid orders which consume the entire level of depth. Consistent with the ability of liquidity providers to provide "phantom" liquidity and fade if desirable, liquidity demanders on new Alpha are able to access only 34% (100% - 66% average quote fade) of the liquidity quoted at the time of order submission, compared to 85% on "old" Alpha. These changes represent a dramatic shift not only from "old" Alpha, but also from the rates of quote fade typically observed across other venues.

< Insert Table 1 Here >

4 Speed Bump Mechanics, Order Flow Segmentation and transaction costs

In this section, we investigate how the introduction of a systematic order processing delay and shift to inverted maker-taker pricing on Alpha affect the routing of informed and uninformed order flow. To understand why the new market design might lead to differential routing among traders with varying information content, we start by analyzing the mechanism by which this segregation occurs, documenting the ability of liquidity suppliers on Alpha to fade against incoming informative trades which execute across multiple venues. We present changes in the market share of active and passive trades by broker type, as a proxy for the level of retail, institutional and proprietary trading. We further examine Alpha's market share of trade strings that either incur or avoid adverse selection costs in comparison to other venues, as well as changes in the realized spread and adverse selection costs for trades across venues.

4.1 Fleeting liquidity and the mechanics of reducing adverse selection costs

Alpha's speed bump of 1 to 3 milliseconds for incoming market orders provides an opportunity for liquidity suppliers to cancel standing limit orders ahead of new marketable orders, particularly after observing trades on other venues. For NBBO-depleting trade strings, we calculate *Quote Fade* on each trading venue by comparing the liquidity visible at the NBBO at the start of the string with the actual volume traded. Absent Quote Fade, all visible liquidity must have been traded.

< Insert Figure 1 Here>

Panel A of Figure 1 presents daily aggregate *NBBO Quote Fade* by trading venue. A rather striking increase in Quote Fade is observed on Alpha immediately after the relaunch, while Quote Fade decreases slightly across TSX, Chi-X and CX2. Given liquidity providers most likely fade when they observe SOR-related executions on other venues, an increase in Quote Fade should enable liquidity providers to avoid order flow which interacts with many venues, whilst capturing order flow interacting with only one venue. Consistent with this intuition, Panel B in Figure 1 shows a marked increase in our measure of *Segmentation Imbalance* for Alpha and a commensurate decline across non-Alpha venues. We test for significant changes in these two liquidity access measures around the speed bump's introduction using Equation 11:

$$LiqAccess_{i,d,v} = \alpha + \delta\ Post_d + \beta_1 Price_{i,d} + \beta_2 Volume_{i,d} + \beta_3 Volatility_{i,d} + \beta_4 VIX_d + \varepsilon_{i,d,v} \quad (11)$$

where $LiqAccess_{i,d,v}$ represents either $Quote\ Fade_{i,d,v}$, the total fraction of starting NBBO liquidity withdrawn across all trade strings or $Segmentation\ Imbalance_{i,d,v}$, the difference between uninformed and informed liquidity shares, on venue *v* for stock *i* on day *d*. The key variable of interest is $Post_d$, an indicator variable equal to one for observations after the introduction of the speed bump and zero prior, which picks up the change in the dependent variable in the post period. The set of control variables includes $Price_{i,d}$, the natural logarithm of the time-weighted NBBO midpoint price, $Volume_{i,d}$ the natural logarithm of aggregate on-market trade volume, $Volatility_{i,d}$ the natural logarithm of the realized intra-day volatility of one minute NBBO midpoint returns, and VIX_d, the daily level of the TSX Composite volatility index.

We use the same set of control variables throughout this study. All variables, except those that are directly observed (such as price and VIX) and those that are naturally bounded, are winsorized daily at the 1 and 99 percent levels. Stock fixed effects are added to all regressions, with standard errors clustered by stock. Results from regression analysis throughout the paper are based on a window of 10 weeks on either side of the event date. For robustness, every test was repeated for the full 12 month sample, with very similar results. We opt for the shorter horizon to avoid capturing confounding effects from other events unrelated to the speed bump. Our analysis is robust to the inclusion of a linear time trend and the exclusion of stock fixed effects. Graphical evidence in our figures is presented for the full 12 months, providing some additional, longer-term perspective.

< Insert Table 2 Here >

Consistent with Figure 1, Table 2 shows that Alpha's Quote Fade jumps by 46% immediately after the introduction of the speed bump. High Quote Fade indicates that quoted liquidity was removed before it could be traded, representing inaccessible liquidity.[28] The ability to fade against the majority of trades that will generate instantaneous adverse selection is the mechanism by which liquidity suppliers on Alpha reduce their interaction with informed trades, minimizing adverse selection costs and increasing realized spreads. This effect is captured by the significant increases in segmentation imbalance on Alpha documented in Table 2, up 4.20 % after the event. Liquidity being removed from the side of the book which is about to become "thin" is consistent with the empirical findings of Goldstein, Kwan and Phillip (2016) that high frequency traders (HFT) primarily supply liquidity on the "thick" side of the book. As a consequence, Alpha becomes unattractive for larger orders that need to access liquidity across multiple venues simultaneously. The random nature of the delay makes it impossible to guarantee consistently low Quote Fade on multiple venues using an SOR.

In contrast, Quote Fade on CX2 decreases almost 3% on the day of the event and declines further over time, indicating that liquidity demanders more aggressively access its displayed limit orders at competitive prices. A high level of accessibility of consolidated market depth across all venues in the pre-event period is consistent with the arguments of O'Hara and Ye (2011) that a trade-through prohibition combined with smart order routing in fragmented markets(virtually) replicates the network advantages of consolidated trading.

---

[28] These effects were not completely un-anticipated. Prior to Alpha's speed bump implementation, several market participants noted that it may result in undesirable consequences. For example, in a submission to the Ontario Securities Commission, Clark (2014) claimed that *"the new Alpha design will allow passive Post Only resting orders the ability to fade should they see trading on another venue"*.

Given several changes were implemented on Alpha simultaneously, it is important to ensure that the increase in Quote Fade on Alpha is driven by the ability to cancel orders without experiencing the speed bump. We conduct several additional tests to ensure this is the case.

In order to separate limit orders which are exempt from the delay (Post Only orders) from those that are not, we make use of the fact that exempt orders must fulfil a minimum order size to be eligible. This minimum is stock-specific and published by the TSX at the beginning of each month as a function of price and liquidity. For the stocks in our sample, this minimum varies between 500 and 5,000 shares. In Table A.5 of the Appendix, we tabulate for each venue the relative share of cancel orders whose size equals (exceeds) the minimum size requirement for that particular stock, separately by minimum size stock groups and for pre and post period.

We observe a highly unusual increase in the cancelation activity of Post-Only orders on Alpha in the post period. For example, in the pre period, for stocks with a minimum of 500 shares, only 21.8% of canceled limit order volume equals 500 shares or more. This increases to over 90% in the post period. The difference at the limit is even stronger: While only 0.6% of pre period cancelations have size of exactly 500, this increases to 62% in the post period. The other venues exhibit no significant changes. This pattern is repeated across all minimum size buckets and clearly demonstrates that the vast majority of limit order volume on Alpha in the post period originates from speed bump exempt Post Only orders. Identical results are obtained in Table A.6 where we repeat this analysis for order entry messages.

Given the majority of limit orders on Alpha are exempt, we repeat the analysis of Quote Fade and segmentation by conditioning on the state of Alpha's order book. We inspect order book data from Alpha in the post period and find that typically when depth is equal or larger than the minimum exemption size it is so in whole multiples of the minimum size. We assume that exempt limit orders are present whenever we observe a quote on Alpha at the best bid or ask with a depth of at least the minimum exemption size. Each trade string is then categorized depending on whether, at the start of the trade, it faces at least one exempt limit order on Alpha or not. In the former case, the market order is subject to the speed bump while (cancellation) messages for the liquidity waiting on Alpha are not (state *Non-Delay*). In the latter case, any cancellation message sent to Alpha is subject to the same delay as the market order. This separation permits potentially two observations per stock-day and venue in the post period. We include a stock-day-venue-state observation if trading volume in that category constitutes at least 10% of total trading volume for that venue to ensure it is representative. We then run the following test:

$$LiqAccess_{i,d,v,s} = \delta\ Post_d + \gamma\ Post_d\ \times NonDelay_{i,d,s} + \Theta' Controls_{i,d} + FE_i + \varepsilon_{i,d,v,s} \qquad (12)$$

where the Post dummy now captures changes relative to the pre period during times when either Alpha was not at best or only non-exempt limit orders were present on Alpha, while $Post\ \times NonDelay$ represents the incremental change (in addition to the Post coefficient) experienced in times when exempt limit orders are present on Alpha.

The differences in Quote Fade behavior and order flow segmentation reported in Table 3 are quite extreme. Compared to Table 2 we can infer that almost all of the increase in Quote Fade is due to cancellations of non-delayed limit orders. For the distributional effects of adverse selection, we see that in the absence of exempt orders, Alpha does not seem to be particularly effective at segmenting order flow. However, with asymmetry present, Alpha is able to capture significant portion of trading activity that does not deplete the top of the book, to the detriment of the two non-inverted markets, Chi-X and TSX.

< Insert Table 3 Here >

4.2 Smart-Order-Routing and the Information Content of Trades

Using the classification schemes derived in Section 3.3, we first split trade strings into two groups: single venue trade strings (likely from small traders); and trade strings accessing more than one venue, likely originating from the SOR of a large, institutional trader (Van Kervel, 2015). Figure 2 shows that this distinction is important, by contrasting realized spreads for these two types of trade strings at several time horizons. We find that the vast majority of the price impact from trading occurs virtually instantaneously for both types of trade strings. However, multi-venue trade strings experience immediate negative realized spreads, while trade strings which only execute on one venue provide liquidity suppliers with positive realized spreads for at least the first 20 seconds. The difference between the negative realized spreads for multi-venue sprays and the positive realized spreads on single venues demonstrates the economic value of avoiding SOR trades to liquidity providers the speed bump provides.

<Insert Figure 2>

In a second classification, we split trade strings into those that deplete the top of book and those that do not. We consider depleting trades more likely to be informed than non-depleting trade strings. The combination of these classifications results in a total of four trade string categories. By definition, depleting trade strings exert immediate adverse selection on the liquidity supplier because quotes change, leading to lower realized spreads. Thus, while the two classifications are related, it is possible that a single-venue trade string leads to a quote change while a multi-venue trade does not.

Figure 3 presents Alpha's trade composition for both depleting/non-depleting trades and those executed with/without a SOR. Small retail orders are likely to be filled on one venue without depleting the NBBO (non-depleting, non-SOR trades). The proportion of non-depleting, non-SOR trades increases dramatically, from 18% on "old" Alpha to 46% on "new" Alpha. Conversely, large institutional trades are likely to exhaust all liquidity available at the NBBO using a SOR. Depleting cross-venue sprays experience a dramatic decline on Alpha, from 46% to 23%. Little movement is observed either for single-venue depleting orders (i.e. large retail orders) or for cross-venue sprays which do not displace the NBBO (i.e. small institutional orders). Given these measures are based on traded liquidity, they demonstrate the ability of liquidity suppliers on Alpha to "fade" away from large institutional orders which access multiple venues, while interacting with a relatively larger proportion of (likely) uninformed retail flow.[29]

< Insert Figure 3 Here>

4.3 Trade-based liquidity metrics on Alpha

To test for statistically significant changes in Alpha's liquidity metrics following the relaunch, we run panel regressions of the form:

$$LiqMetric_{i,d} = \alpha_v + \delta\, Post_d + \beta_1 Price_{i,d} + \beta_2 Volume_{i,d} + \beta_3 Volatility_{i,d} + \beta_4 VIX_d + \varepsilon_{i,d} \quad (13)$$

where $LiqMetric_{i,d}$ is a measure of liquidity on Alpha for stock i on day d and all other variables are as described in Equation 11. Changes in effective spreads, realized spreads and adverse selection costs on Alpha after its relaunch are presented in Table 4. Control variables for price, volume, volatility and VIX have the expected directionality and are statistically significant.

Net-of-fees effective spreads on Alpha increase by 0.58 cents, following the market structure changes. This is despite the fact that, for shares over C$1.00, old Alpha had an active trading fee of 0.18c per share, while active trades under the revised fee structure receive a rebate of 0.10c per share traded, resulting in a substantial decline in explicit trading costs of 0.28c per share. Consistent with Malinova and Park (2015), we document that liquidity suppliers pass on changes in explicit fees, even under inverted maker-taker pricing schemes.

< Insert Table 4 Here >

[29] Section A.1 of the Internet Appendix provides further justification for the inference that these active orders are retail originating. Figure A.1 uses broker ID as a crude proxy for retail brokers, finding the two largest Canadian retail brokers represent around 29% of active orders post-speed bump, up from around 18% in the pre-period.

To explicitly examine the benefits to liquidity suppliers of utilizing the speed bump, we calculate realized spreads by comparing traded prices with NBBO midpoint quotes at both one and 20 seconds after each trade, following Conrad et al. (2015). As shown in Table 5, realized spreads increase 1.10 cents after one second and 1.23 cents after twenty seconds. In relative terms, realized spreads increase 3.71 and 4.53 basis points after one and twenty seconds, respectively. These substantial increases prevail despite the fact that the switch to inverted fees should, ceteris paribus, lower gains to liquidity provision by 0.14c (old make rebate) plus 0.10c (new make fee).

Table 4 also shows that under Alpha's new market structure, adverse selection costs decline by 0.56 cents one second after a trade and 0.69 cents after 20 seconds. In relative terms, price movements away from the liquidity supplier decline 3.09 and 3.9 basis points, one second and 20 seconds after each trade respectively. The increase in the realized spread of trades on Alpha indicates that liquidity suppliers are able to either widen their spreads or avoid adverse selection. The observed decreases in adverse selection costs are slightly larger than the increases in effective spreads, indicating that increased profitability of liquidity provision on Alpha is driven mainly by the ability to avoid correlated order flow.

4.4 Traded Liquidity Metrics on Other Venues

Having established how the design changes to Alpha affected trading costs on Alpha itself, we now examine whether there are any contemporaneous spillover effects onto the other venues, TSX, Chi-X and CX2. Some changes on other venues are likely, particularly given the evidence in Section 4.2 that Alpha's systematic order processing delay for marketable orders enables the segmentation of uninformed order flow. Existing empirical evidence suggests that the segregation of uninformed active orders on dark venues increases the toxicity of the remaining order flow on public lit markets (e.g. Easley et al., 1996; Zhu, 2014, Comerton-Forde and Putnins, 2015).

< Insert Table 5 Here >

We repeat the regressions specified in Equation 13 for the remaining three venues (CX2, Chi-X, and TSX), separately as well as jointly. Table A.2 in the Appendix reports changes in effective spreads, realized spreads and adverse selection costs against the NBBO midpoint, volume-weighted amongst all trades on TSX, Chi-X and CX2 in the same format as Table 5, while Table 5 shows more detailed venue-by-venue results. For brevity, we report results in cents at the 20 second horizon only.

We find that Alpha's improved avoidance of informed trades that sweep multiple venues does indeed increase the toxicity of residual order flow on the other large Canadian trading venues, with adverse selection costs increasing significantly by between 0.21c and 0.31c. In most cases, effective spreads widen less, however, imposing a net cost on liquidity providers. This affects the other inverted venue (CX2) the most, where realized spreads decline a substantial 0.37c, while declines are more modest at 0.12-0.27c on the non-inverted venues.

Alpha's new inverted maker taker pricing and larger quoted depths from minimum Post Only order sizes enable it to aggressively compete with CX2 for active retail flow, reducing the profitability of liquidity provision on that venue. The large reduction on CX2 is consistent with a reduction in the proportion of uninformed (retail) order flow in the aggregate of market orders, likely as a result of a migration to Alpha due to the (mandated) larger quoted depths. We conclude that Alpha's segmentation of order flow increases residual order flow toxicity and imposes negative liquidity externalities on the other trading venues, in particular the competing inverted venue.

5 Direct comparison of inverted venues and speed bumps

Given that in the post period, our sample contains two inverted venues with substantial market share, Alpha and CX2, it seems instructive to directly compare their performance in terms of trading costs. Our goal here is in particular to separate the effects that the switch in fee structure had on Alpha from the effects of the speed bump. To this end, we employ 2 identification strategies, one based on the presence of delay-exempt limit orders, the other based on broker level IDs allowing us to identify fast vs. slow liquidity providers.

5.1 … conditional on the presence of speed bump exempt liquidity

As previously shown in Table 3, the ability of Alpha to segment order flow crucially hinges on the absence vs. presence of speed bump exempt limit orders, i.e. the symmetric vs. asymmetric application of the speed bump. Thus we would expect that trading costs on Alpha and CX2 should be quite similar during times when there is no asymmetry in the application of the speed bump between market and limit orders on Alpha (state *Delay*), while liquidity provision should be vastly more profitable when liquidity providers are exempt (state *non-Delay*). We test our hypothesis in the 10 weeks after the event date using stock-day observations from Alpha and CX2 jointly in the following specification:

$$LiqMetric_{i,d,v,s} = \alpha\ Alpha_v + \gamma\ NonDelay_{i,d,s} + \eta\ NonDelay_{i,d,s} \times Alpha_v + \Theta' Controls_{i,d} + FE_i + \varepsilon_{i,d,v,s}$$

Where *Alpha* is an indicator that is 1 (0) for Alpha (CX2), and *NonDelay* is an indicator that is 1 (0) for that subset of trades that occurs with speed bump exempt limit orders being present (absent). The first 3 columns of Table 6 contain unconditional comparisons of Alpha and CX2, i.e. one observation per stock-day-venue, while the last 3 columns separate trades into 2 subsets as a function of the state of Alpha's order book state (*Delay* vs. *NonDelay*); however, we impose the condition that each observation must represent at least 10% of that stock-day-venue's total trading volume to ensure representativeness.

Table 6 shows that unconditionally, liquidity provision on Alpha experiences realized spreads that are significantly higher than on CX2 by 0.39 cents and lower adverse selection by about 0.18 cents. The difference is made up by higher effective spreads which seems to indicate that LPs on Alpha are less aggressive than on CX2.

Conditioning on symmetry vs. asymmetry of the speed bump, we find quite different results on the next 3 columns of Table 6. The coefficients of Alpha now represent situations where LPs did not have the advantage of being exempt from the speed bump relative to CX2. There is no significant difference between effective and realized spreads between the two venues. Adverse selection is slightly higher on Alpha, and while it is statistically significant, its economic magnitude is rather small with 0.08 cents. In other words, net of fees the two inverted venues have almost identical performance in the absence of speed bump induced asymmetry, confirming the first part of our hypothesis.

Coefficient *NonDelay* represents the incremental effect on both venues when Alpha's order book contains exempt limit orders, while the interaction term reports the additional incremental effect that is Alpha-specific. We find that liquidity provision on either venue is more profitable during times with exempt limit orders present, which suggests that liquidity providers on Alpha are strategic about when to use the exempt order type, but the benefits are about double on Alpha compared to CX2. For example, realized spreads on CX2 improve by 0.55 cents, those on Alpha are higher by 0.55 + 0.52 = 1.07 cents. The interaction term represents the additional value that the speed bump on Alpha provides to its users, beyond the inverted fee structure and beyond strategic considerations of when to post.

< Insert Table 6 Here>

5.2 … conditional on the speed of the passive side

For parts of the sample period we were able to acquire broker level identification data, which allows us to distinguish between fast and slow traders on the passive side of Alpha and CX2. The details of this classification can be found in Appendix A.1. The delay on Alpha is small enough in duration that not every user of limit orders has the means to take advantage of it. Figure A.2 in the Appendix shows that while high frequency traders with direct market access (DMA) provide the majority of liquidity on Alpha in the post period, other trader groups are still present. We assume that most orders that arrive through channels associated with DMA originate from traders fast enough to exploit the exemption (trader type *FastLP*), while at least some of the orders originating through other broker IDs are not (trader type *SlowLP*).[30] Following Rosu (2016), it is reasonable to expect that fast traders harvest higher realized spreads and experience lower adverse selection on either venue. But we also hypothesize that faster traders are incrementally better off along both dimensions on Alpha than on CX2 because of the advantages that the speed bump affords them. To the extent that our identification of differences in reaction time is noisy, it would bias against finding any difference between these two groups that aligns with our hypothesis. We run the following test:

$$LiqMetric_{i,d,v,t} = \alpha\ Alpha_v + \gamma\ FastLP_{i,d,t} + \eta\ FastLP_{i,d,t} \times Alpha_v\ + \Theta' Controls_{i,d} + FE_i + \varepsilon_{i,d,v,t}$$

Where *Alpha* is an indicator that is 1 (0) for Alpha (CX2) and *FastLP* is an indicator that is 1 (0) for that subset of trades with a fast (slow) liquidity provider on the passive side, as defined above. Table 7 repeats the first three columns of Table 6 containing unconditional results for ease of comparison. In the next three columns, we can differentiate between trading costs of slow and fast liquidity providers across the two venues. The coefficients of *Alpha* now represents differences between the 2 venues that slow traders experience. We find that slow users of limit orders seem to do poorly on Alpha compared to CX2 as realized spreads are 0.65 cents lower and adverse selection is 0.35 cents higher. They also tend to trade at lower effective spreads.

FastLP represents the common, incremental effect for fast liquidity providers on both venues, while the interaction term represents the additional incremental effect for fast traders on Alpha. Not surprisingly, fast traders capture more profits and experience lower adverse selection on either venue, but being fast on Alpha is a lot more profitable. The difference between fast and slow traders on CX2 is 0.49 cents, but it is 0.49+1.12 = 1.61 cents on Alpha, again in line with our predictions.

---

[30] We exclude liquidity providers that anonymize their broker ID for this test.

A noteworthy implication of these results is that the speed bump not only segments order flow between venues and redistributes profits from liquidity provision towards Alpha, it also creates winners and losers within Alpha: Fast HFTs and users of exempt limit orders are able to cherry-pick the most favorable trades on Alpha leaving slower traders worse off.[31]

< Insert Table 7 Here>

6 Cross Sectional Variation

(Purely) inverted venues tend to provide liquidity providers with lower adverse selection and higher realized spreads (Battalio et al., 2016).[32] TSX Alpha, being also inverted, seems especially good at this relative to the other major inverted venue in the post period and the evidence in the previous section suggests that it is the asymmetric speed bump feature that causes this difference, as these effects were conditional on the presence of speed bump exempt limit orders and traders fast enough to exploit the exemption.

In this section, we contrast previous market entries of venues with inverted maker/taker fee structures with the relaunch of Alpha as an inverted venue plus speed bump to highlight the different mechanisms at work that lower adverse selection. We develop and then test hypotheses that imply differential outcomes in the cross-section of stocks for pure inverted venues vs. Alpha's speed bump.

6.1 Inverted markets in the cross-section

[31] We do not test this directly, but it seems reasonable to assume that the groups FastLP and NonDelay overlap to a large extent. However, the results between Tables 6 and 7 should differ alone because the latter excludes trades with anonymized broker IDs.
[32] For example, in our pre-event period, inverted venue CX2 has mean and median net realized spreads that are at least 1 cent higher than non-inverted venues and a correspondingly lower adverse selection component.

Recent findings in the literature suggest that typical inverted markets tend to attract stocks that are low priced, tick-constrained, and with large depth at the NBBO (e.g. Battalio et al., 2016; Harris, 2013; Cimon, 2016). These stocks have "deep and dense books" in the sense that every price point is occupied by a large number of limit orders and the price points are only one tick (one cent) apart, making it less likely that the next market order exhausts all liquidity at a given price point, but even if it does, the move will usually be limited to one tick. By contrast, the limit order book for stocks on the opposite end of the spectrum (higher price, unconstrained, thin books) tends to be thin and sparse: a trade is more likely to move quotes and by more than one cent as the next limit order may be several ticks away. We can think of the propensity of trading to move quotes as immediate adverse selection (IAS) and in DepleteBest we have a direct measure of it, in addition to the 3 above-mentioned stock characteristics that serve as proxies.

We confirm our intuition above in unreported results where we find that, prior to the design change on Alpha, DepleteBest is on average around 52% for the former group of stocks, but about 64% for the latter group (defined as having below or above median average values for price, time being tick-constrained, NBBO depth and DepleteBest itself). For brevity, we will henceforth refer to the former (latter) group collectively as low-IAS (high-IAS) stocks.

Because of this, we hypothesize that liquidity providers should be more willing to pay to post quotes in low-IAS stocks on an inverted venue, which has shorter queues and tends to be accessed first due to offering the best price net of fee, knowing the risk of adverse selection is low. And second, we would expect that as a consequence of more frequent liquidity provision in low-IAS stocks on inverted venues, segmentation of non-depleting order flow should be more pronounced than in high-IAS stocks.

To test these conjectures and then highlight differences to TSX Alpha, we investigate how previous market entries by inverted venues affected the willingness of liquidity providers to post and segmentation in the cross-section of stocks compared to TSX Alpha. Prior to Alpha, 3 venues had launched with or switched to a maker/taker pricing scheme between 2011 and 2013. We compile 20-week stock-day panel datasets surrounding each of the 3 previous market entries and filter them in the same way as our main sample. More details can be found in Appendix A.2. We then split the stocks from each of the 4 event samples into two halves according to a total of 4 cross-sectional characteristics related to the risk of IAS: stock price, the proportion of time that the NBBO is tick constrained, the volume of shares quoted at the NBBO (i.e. depth of book), and finally DepleteBest itself.

The difference in IAS risk between these two types of stocks is clearly visible in the willingness of liquidity providers to quote at the NBBO on an inverted venue. Columns 1 to 3 in Table 8 report the average proportion of time that the 3 purely inverted venues quoted at the NBBO in the 10 weeks following their respective market entries. Liquidity providers on purely inverted venues seem to be very reluctant to quote at NBBO for high-IAS stocks, doing so only between 3 and 11 percent of the time; while for low-IAS stocks this figure varies between 20 and 32 percent, i.e. between 2.1 and 7.6 times as much. Casually speaking, the loss caused by immediate adverse selection is the likelihood of exhausting the NBBO times the expected price move conditional on exhausting the NBBO. Since we expect the price move likely to be more than one tick for higher priced stocks, it explains the magnified reluctance to post beyond the pure difference in DepleteBest.

Columns 5 to 7 focus on the ability of the same 3 inverted venues to segment low-IAS order flow from high-IAS order flow, i.e. our measure of Segmentation Imbalance. Given that the data only covers the first 10 weeks of each venue's trading as an inverted venue, the reported numbers may understate the longer term effects, but nevertheless a clear pattern can be observed: Across all three venues, low-IAS stocks are relatively more segmented by (and towards) inverted venues than high-IAS stocks, usually at least twice as much.

By contrast, TSX Alpha in its post-event period "behaves" quite differently. While its liquidity providers are still relatively more willing to quote at NBBO for low-IAS stocks – Alpha is an inverted venue after all – the differential is proportionately much smaller with a ratio of 1.27 to 1.54. In particular, rather than less than 10 percent, Alpha quotes between 26-29 percent of the time at NBBO for high-IAS stocks. Furthermore, there is little difference in Alpha's ability to segment the two sets of stocks. However, the degree of segmentation at around 6-7% is many times larger than what any of the previous inverted venues was capable of.

< Insert Table 8 here >

Given that all 4 exchanges introduced an inverted fee structure, the evidence points towards the speed bump as being the culprit in the differential behavior of TSX Alpha.The cross-sectional differences between Alpha and other inverted venues in terms of quoting behavior and segmentation ability imply knock-on effects for the transaction costs on the other, existing venues: We expect the relaunch of Alpha to effect relatively larger changes in high-IAS stocks, especially compared to purely inverted venues. We test this intuition next.

6.2 Previous market entries of inverted venues

To establish a base line in terms of how the entry of a purely inverted market may affect the distribution of profits from liquidity provision and adverse selection, we investigate the 3 previous entries of inverted markets in more detail using the usual event-study methodology. Against these reference points we can then better judge the effects that Alpha's relaunch had. For each of the 3 previous events, we run tests separately for each of the 4 characteristics, allowing for differences in the change in the post-event period for low vs. high-IAS stocks each time:

$$y_{i,d} = \delta\, Post_d + \gamma\, Post_d \times PreChar_i + \Theta^T Controls_{i,d} + \Phi^T Controls_{i,d} \times PreChar_i + FE_i + \varepsilon_{i,d}$$

Where $y_{i,d,s(c)}$ represents either the segmentation imbalance of all venues except the event venue (Omega, CX2 and TMX Select, respectively), or measures of trading costs (realized spreads and adverse selection) for the event venue (if available in the pre-period) vs. the remaining venues. $PreChar_i$ is a dummy equal to 1 for stocks with a characteristic that corresponds to a higher risk of IAS, i.e. the interaction term represents the differential effect on high-IAS stocks in the post period.The results are reported in tables A.7 through A.9. For brevity, we only report the two coefficients related to the Post dummy from each of the 4 regressions, but all tests include the usual set of controls and fixed effects and even allow for differences in slopes between the 2 subsamples.

Common to all 3 events is that high-IAS experience about half the decline of segmentation that low-IAS stocks do, providing the flip side to the positive segmentation ability of the entrant venue reported in Table 8.

Overall, the existing venues experience modest declines in realized spreads and increases in adverse selection. Some care needs to be taken, however, when interpreting the cross-sectional differences as neither figures in cents nor basis points are directly comparable: high-IAS stocks tend have nominal prices that are on average 2 to 5 times larger than low-IAS stocks in our sample (\$30-\$50 vs. \$8-\$20 depending on the characteristic). As a general pattern, when looking at cent figures, high-IAS stocks do not seem to experience significantly greater declines (increases) in realized spreads (adverse selection) than low-IAS stocks, as the interaction term is often insignificant and of varying signs. Given the difference in price, this suggests that the effects are more pronounced for low-IAS stocks. Basis point figures support this view, as post dummy and interaction term coefficients are consistently of opposite signs and the latter reverses between half and all of the effects on the former.

6.3 Alpha’s speed bump in the cross section

We repeat the cross-sectional event-style analysis that we applied to the previous 3 entries of inverted venues to the relaunch event of TSX Alpha, again splitting the sample into 2 halves along 4 stock characteristics related to the risk of IAS and allowing for differential cross-sectional outcomes in the post event period. Table 9 reports those results.

< Insert Table 9 here >

The last column of Table 8 demonstrated new Alpha's ability to segment order flow rather equally for low vs. high IAS stocks. The second column of Table 9 shows that in order to get there from pre to post period, the *change* in segmentation experienced by the rest of the market was significantly more negative in high-IAS stocks than in low-IAS stocks, at -5.3% compared to -2.9%. In other words, the redistribution of adverse selection risk caused by Alpha's relaunch was much more pronounced in high-IAS stocks. This stands in stark contrast to the previous 3 inverted market entries where the opposite was the case (Tables A.7 to A.9). Also in contrast to the previous events, realized spreads and adverse selection (measured in cents) for the existing venues are more affected in high-IAS stocks. The effect is especially visible in the increasing adverse selection where low-IAS stocks experience economically small increases of at most 0.1 cents, while high IAS stocks see effects of up to 0.47 cents. Realized spreads show a similar pattern but at much smaller economic magnitude. Taken both results together, this strongly suggests that we should expect differential cross-sectional effects for quoted and effective spreads. We will test this implication in Section 7.

6.4 Mechanisms of avoiding adverse selection

We ascribe the differences found in the previous tests between purely inverted venues on the one hand and Alpha, an inverted venue cum speed bump, on the other to the presence of the speed bump. As we showed in Tables 3 and 6, essentially all the effects on spreads are concentrated in times when liquidity on Alpha comes in the form of exempt PO limit orders, i.e. when the top of the book liquidity supplier on Alpha has the ability to observe trading activity on other venues and cancel if he fears adverse selection.

Both mechanisms, inverted venue and speed bump, ultimately allow the same outcome, i.e. to lower adverse selection, but this outcome is achieved in different ways. On the (purely) inverted venue, it is optimal to quote at best when the stack of liquidity "behind it" is large, which serves as an insurance against adverse selection, akin to a level of support or resistance in the parlance of technical analysis. The liquidity provider pays the exchange to jump the queue, but imposes an externality on the liquidity providers on the other venues, who are having to wait longer and facing higher adverse selection cost, because trades tend to first exhaust the liquidity on the inverted venues before moving to non-inverted markets. Thin books make this strategy less profitable, hence liquidity providers tend to quote them less aggressively on inverted venues and inverted venues tend to have lower market shares in stocks with thinner books

By contrast, on Alpha, no thick stack of liquidity is required as insurance against adverse selection, instead the exemption from the delay, for which the liquidity provider pays as well, provides him with a 'last look option', where he can decide to avoid high adverse selection trades bound his way. The fact that Alpha is *also* inverted reinforces the efficiency of the speed bump, as it guarantees that smaller, single-venue (low IAS) orders by fee-conscious brokers are more likely to be routed there. We surmise that if Alpha was not inverted, the speed bump would still allow for 'last look' segmentation, but it would not receive single-venue orders to the extent that it does; thus, segmentation would be less pronounced and its market share would likely be lower.

The avoidance of adverse selection, and thus the externalities imposed, by users of Alpha's speed bump on the rest of the market are more complete when compared to pure inverted venues. While both are first in line to receive low-IAS trades, the purely inverted venue cannot avoid high-IAS trades as efficiently as Alpha can.

7 Effects of speed differentials on consolidated market quality

Introducing speed differentials has the ability to redistribute the impacts of adverse selection. The 1-3ms speed bump on Alpha provides market makers the ability to withdraw liquidity without any delay, whilst slowing liquidity takers. This gives the market makers on Alpha a speed advantage over other participants. They are then able to observe executions on other venues and withdraw liquidity when it is likely they will suffer instantaneous adverse selection. Their inverted maker-taker pricing structure is shown by Batallio et al. (2016) to be particularly attractive to retail brokers. Segmenting retail order flow into a single venue whilst excluding institutional flow is likely to increase market-wide adverse selection, increasing the cost of liquidity. The following section explores these effects, both for the whole market and in the cross-section of securities where the speed bump is expected to be most valuable.

7.1 Theoretical predictions

Many theoretical works have modelled the impact of differential speed between informed traders and market makers. Biais, Foucault and Moinas (2015), Budish, Cramton and Shim (2015), Foucault Humbert and Rosu (2016) and Foucault, Kozhan and Tham (2016) all envision situations in which informed traders are able to respond to new information faster than market makers. These fast, informed traders are then able to adversely select market makers, forcing them to increase their spreads to recover these costs. While our situation involves some market makers increasing their speed, those who are unwilling/unable to do so will also encounter additional adverse selection, requiring them to increase their spreads. Perhaps most closely related to our situation is the work of Hoffman (2014) and Jovanovic and Menkveld (2015), who consider a situation where some market makers become fast, avoiding adverse selection and thus increasing supplied depth. However, the fast market makers impose additional adverse selection on slower market makers, resulting in wider quoted spreads. This situation is very similar to that observed in Canada, with simultaneous increased depth and widened quoted spreads. A recent paper from Brolley and Cimon (2017) models the introduction of a speed bump alongside an incumbent exchange. Consistent with our results, they find reductions in adverse selection on the speed bump venue, and commensurate increases for the non-speed bump venue. Whilst their model predicts a narrowing of spreads on the speed bump venue (due to lower adverse selection) they do not consider that the *order flow itself* could be the information. In such a situation, the speed bump loses its functionality when it ‘makes’ the NBBO.

7.2 Unconditional effects

The previous tests show that Alpha is successful at segmenting non-depleting trades, increasing profits for liquidity providers on Alpha at the expense of the remaining venues. We now examine whether this redistribution of order flow has a measurable impact on consolidated market quality. We test for changes in market quality metrics jointly across all four venues (Alpha, TSX, Chi-X and CX2) using the panel regression:

$$MQ_{i,d} = \delta\, Post_d + \Theta^T Controls_{i,d} + FE_i + \varepsilon_{i,d}$$

where $Market\ Quality_{i,d}$ is a measure of consolidated market quality for stock *i* on day *d*, and all other variables are as previously described.

< Insert Table 10 Here >

Table 10 presents regression results for changes in liquidity metrics across all four trading venues consolidated at the national best bid and offer prices. Quoted spreads increase by a substantial 0.24c in absolute terms and 0.44 basis points in relative terms. These costs are mirrored by overall increases in the cost of accessing liquidity (effective spreads) of 0.17c or 0.16 basis points, similar in magnitude to the impact on individual venue effective spreads reported earlier. With increasing adverse selection, it is unsurprising that liquidity providers on other venues widen their spreads, though in theory liquidity providers on Alpha could quote more aggressively due to lowered adverse selection costs. We find, however that liquidity providers do exactly the opposite. Figure 4 shows that Alpha's time at best (i.e. the fraction of time that quotes on Alpha are at the NBBO) declines dramatically, essentially never quoting alone at the NBBO in the post period.

< Insert Figure 4 Here >

Table 10 documents a significant 17 percent[33] increase in overall depth at the NBBO. This depth increase is in part driven by the increased quoted size on "new" Alpha. While this appears to be beneficial, the large increase in quote fade on Alpha reduces the accessibility of liquidity on Alpha after the introduction of the speed bump.

Table 10 also reports measures of order book fragility. *Deplete Best* represents the proportion of daily trading volume that consumed all depth available on one side of the NBBO. In the post period, this figure rises by a statistically significant 1.87 percent. Similarly, *Access Next* is the proportion of trading volume that 'walked the book', i.e. executing at prices inferior to the pre-trade NBBO. We find that in the post period, an additional 1.74 percent of volume within trade strings was forced to access the next best price. Therefore, although overall displayed market depths increase, trades across all venues were more likely to consume the entire depth available and "walk the book", filling at inferior prices.

7.3 Strategic presence of Liquidity Providers on Alpha

Table 3 provided evidence that the delay-exempt Post Only limit order type was chiefly responsible for the jump in Quote Fade and the segmentation of 'good' order flow towards Alpha. We now identify when liquidity providers on Alpha use the speed bump to post competitive quotes and how this affects transaction costs relative to instances when these traders do not use the speed bump. We analyze the propensity of trading to move quotes, Alpha's market share and liquidity metrics as a function of Alpha's order book state by splitting trades in the post period into two subsets (as done in Table 3). Table 11 reports those results.

The first column shows that the difference in the likelihood of an adverse price change is 11 percent lower in cases where liquidity providers benefit from the speed bump relative to when they do not. As documented in Kwan et al. (2016) it is likely that liquidity providers on Alpha prefer to post competitively on the thick side of the book. Notable also is the large difference in Alpha's market share

[33] Note depth is log(depth), i.e. $\exp(0.16) - 1 \approx 17\%$.

when benefiting from the speed bump 23 vs 3 percent when inactive. Not only do liquidity providers seek out favorable order book conditions ex-ante, they seem to capture a large portion of desirable order flow when they do.

The final 6 columns repeat the analysis for measures of transaction costs, which was conducted unconditionally in Tables 4 and 5, but conditional on the state of Alpha's order book. Notably, realized spreads, i.e. profits to liquidity providers, are on average about 1 cent higher when benefiting from the speed bump, driven mostly by lower adverse selection costs. To avoid sample selection issues, we repeat this test, but require that both categories constitute at least 10 percent of trading volume for either to be included in the test. Our results remain robust to this additional requirement.

< Insert Table 11 Here >

7.4 Cross-sectional effects

We previously documented cross-sectional differences in the *change* in segmentation that the design changes on TSX Alpha caused. Table 9 found that low-IAS stocks experienced increased segmentation of beneficial order flow to Alpha of 2.9%, while for stocks with high-IAS this figure was almost doubled at 5.3%. We also noticed different magnitudes in the effects on realized spreads and adverse selection that seemed to disfavor high-IAS stocks. It thus seems plausible that, if Alpha's speed bump is the cause of overall market deterioration, then overall market quality should be more affected in high-IAS stocks than low-IAS stocks.

To determine the expected impact of introducing an inverted venue, we investigate the impact of the introduction of inverted fee structures on Omega (December 2011), CX2 (May 2013) and TMX Select (November 2013) across all venues for the TSX Composite index. We use the same setup as Table 9, investigating overall market quality changes using 4 different cross-sectional:

$$y_{i,d} = \delta\ Post_d + \gamma\ Post_d \times PreChar_i + \Theta^T Controls_{i,d} + \Phi^T Controls_{i,d} \times PreChar_i + FE_i + \varepsilon_{i,d}$$

where the dependent variables are measures of consolidated quoted and effective spreads as well as order book fragility. $PreChar_i$ is a dummy equal to 1 for high-IAS stocks. The interaction term represents the incremental effect that the dependent variable for high-IAS stocks experienced in the post period relative to low-IAS stocks.

Tables A.10-12 report results for the market entries of Omega, CX-2 and TMX Select respectively. Contrary to the findings of Comerton-Forde, Gregoire and Zhong (2017) that inverted venues increase market quality overall, effects on quoted and effective spreads display little in the way of a consistent pattern, let alone significant effects, both across events as well as within event across different characteristics. Overall, market quality as measured by spreads generally did not deteriorate much in the wake of previous entries of inverted venues and there is also little evidence of a cross-sectional pattern. In terms of NBBO depth, there is some evidence that depth went up for low-IAS stocks but less so for high-IAS stocks. Order book fragility as measured by *Access Next* displays mixed outcomes depending on the event.

Finally, we turn to the relaunch of Alpha for which Table 12 reports results. In contrast to the other events, a clear pattern of significant and economically large effects becomes apparent for quoted and effective spreads. We find no significant effect on either spread variable for low-IAS stocks, but strongly significant increases in both for high-IAS stocks. Quoted spreads widen between 0.36 and 0.49 cents (0.69 and 1.05 basis points) for the latter group relative to the former, while effective spreads widen by slightly lower economic magnitudes. Nominal NBBO depth increases are also more pronounced for high-IAS stocks, which can be explained by the fact that the minimum size requirements for exempt limit orders are likely more binding for the latter group. Finally, increases in order book fragility as measured by *Access Next* are also somewhat higher for high-IAS stocks.

< Insert Table 12 Here >

Prior to Alpha's relaunch, the Canadian market featured 3 inverted venues. With the closure of TMX Select, there remain 3 inverted venues in the post period. As the earlier entries of inverted venues did not cause large changes in market quality, the significant deterioration in market quality suggest the speed bump is driving the causality. This is further strengthened by the consistent nature of the cross-sectional effects with the operation of the speed bump. Overall, the overwhelming majority of tests conducted in this section indicate that despite the number of contemporaneous changes, the primary driver of market-wide impacts is the addition of Alpha's speed bump.

8 Effect of Other Contemporaneous Changes

The introduction of Alpha's speed bump was accompanied by several contemporaneous changes, including inverted maker/taker pricing and the removal of the order protection rule (OPR). Due to the lack of a staggered introduction, it is important to separate the impacts of these contemporaneous events from those caused by the speed bump. This section tests the impact of contemporaneous events to confirm that the observed effects are directly attributable to the speed bump.

8.1 Removal of Alpha from the Order Protection Rule

With Alpha removed from the OPR, it is possible that liquidity demanders exclude Alpha from their routing tables altogether, resulting in low trading activity on that venue. As a result, NBBO liquidity may be present on Alpha but is being "traded through". We know from Figure 4 that quotes on new Alpha are essentially never alone at best in order to take full advantage of the asymmetric delay. Thus, to identify trade through events, we look for instances where Alpha quoted at best jointly with other venues but was ignored, even as transactions on other venues occurred at inferior prices.

To this end, we first collect all trade strings that accessed multiple price levels on any venue. We then define a trade through as an event where a venue had a quote at NBBO at the beginning of the trade and at the end of the trade, but was not accessed during the trade, i.e. no trade occurred on that venue, while trades occurred at a worse price on another venue. For a given stock-day and venue, we compute the ratio of total volume executed at an inferior price, but at most equal to the available liquidity at NBBO on the ignored venue, during a trade through event over the total volume of all trades that accessed multiple price levels and where the venue was at NBBO at the beginning. Thus, we restrict the denominator to trades where Alpha could have been accessed at the best and where we can determine whether it was or was not, but should have been. The numerator represents an estimate of the volume that traded through the venue at an inferior price but could have been filled on the ignored venue at the NBBO. Given that our time stamps are not perfectly synchronized across venues we may mismeasure liquidity available before, during or after a trade. In addition, for venues that do not quote aggressively at the NBBO, the denominator may be rather small for some stock-days. We winsorized at 99 percent to avoid undue outliers. Figure 5 plots the average across stocks over time, separately for each venue.

Two things are noteworthy. First, the larger non-inverted venues have essentially no traded through volume, while the inverted venue CX-2 does have some in both the pre and post period, according to our definition. Alpha's incidence is as low as the other non-inverted venues in the pre-period, but then temporarily spikes to a level similar to CX-2's rate, on average about 0.5%. Within about 2 months though, Alpha's trade through rate has drastically declined.

We thus do not find any persistent and systematically large uptick in trade through incidences on new Alpha, especially when compared to the other inverted venue CX2. Rather it seems likely that the spike in the ratio is driven by a low likelihood of Alpha quoting at the NBBO, especially in the first few weeks after the change.

< Insert Figure 5 Here >

8.2 Global market conditions

Potentially, market quality changes apparent in Canada over the sample period are only reflections of broader, global market trends, unrelated to the speed bump. For instance, financial markets were on average more volatile during the post-event period than prior to the event.

We utilize a difference-in-difference (DID) approach using a sample of U.S. listed securities to examine the market-wide impact of the introduction of the speed bump in Canada. We construct two samples of matched U.S.-Canada stock-pairs. Following Davies and Kim (2009), we first match Canadian stocks to a universe of non-interlisted U.S. equities by market capitalization and nominal price. Second, we match 83 interlisted Canadian stocks in our sample with their corresponding U.S. based listing. Appendix B.4 contains more details on the exact matching methodology. Given the jitter present for the U.S. market, we restrict our analysis to metrics which do not rely on trade strings. Our DID regression has the following design:

$$MQ_{i,d,c} = \theta Post_d \times Canada_{i,c} + \delta Post_d + \beta^T Controls_{i,d,c} + \gamma^T Controls_{i,d,c} \times Canada_{i,c} + FE_{i,c} + \varepsilon_{i,d,c} \quad (17)$$

where $MQ_{i,d,c}$ represents quoted spread, effective spread or NBBO depth, $Post_d$ is a dummy variable which takes a value of one for the post period and 0 otherwise, and $Canada_{i,c}$ is a dummy variable assuming a value of 1 for Canadian stocks and 0 otherwise. The set of control variables is larger in this particular test than previous tests: we start the usual set of controls, but replace the Canadian VIX with the U.S. CBOE volatility index for the U.S. portion of the sample. Next, we account for the fact that spreads in the U.S. may vary to different degrees along stock characteristics by including interaction terms between controls and the Canada dummy. Finally, the U.S. market experiences a slightly negative time trend over the event period, thus we include independent time trend for both countries. Note that because of the inclusion of stock fixed effects, the Canada dummy is subsumed by the fixed effects.

Results in Section 6 showed that spreads are affected differently along 4 cross-sectional stock characteristics. Among these, only stock price carries over to the matched sample because we explicitly used nominal stock price to match. Thus, Table 13 presents the results for both matched samples in 3 different versions: The full samples in Panel A, followed by 2 panels separating stocks below (above) the median price in the pre-event period (based on the Canadian half of each pair).

The Post dummy is generally insignificant and often of negative sign for both quoted and effective spreads across both samples, indicating our results are not driven by global factors around the time of the Alpha relaunch unrelated to the speed bump. We do, however, find significant increases in quoted and effective Canadian spreads in the post period relative to the United States. Mirroring our cross-sectional results from Table 12(?), we find these effects to be strong for the higher priced half of the sample, while completely absent for lower priced stocks.

We also confirm previous results which showed that nominally depth increases substantially in the post period, especially for higher priced stocks, observing e.g. a 11 percent increase (in log terms) relative to cross-listings of the same stocks in the United States.

< Insert Table 13 Here >

Overall, these results indicate that the increases in spreads and depths observed in Canada subsequent to the introduction of the speed bump are not driven by exogenous factors. Together with the results for the Canadian market in Section 5, the increased adverse selection resulting from the segmentation of small, single venue orders on Alpha has resulted in increased overall trading costs for the Canadian market.

9 Discussion and Conclusion

With trading technology approaching the speed of light, speed bumps represent the most recent innovation in the quest to maintain speed differentials. The majority of proposed speed bump designs have one feature in common: discriminatory processing delays which do not apply uniformly to all orders.

We provide the first examination of the market-wide effects of the introduction of a speed bump by a major North-American exchange, TSX Alpha. The asymmetric, randomized nature of this speed bump allows traders to "pay" the exchange for differential speed – that is to exempt their limit order entries and cancellations from the speed bump. Thus, while our results regarding the quality and fairness of markets are relevant to the desirability of speed bumps, they can also be generalized to other situations in which differential speed access can be bought, such as private data feeds, microwave or laser connectivity.

Overall, we find that "new" Alpha is not attractive to all participants, with traded volume decreasing immediately. Using novel trade classification schemes, we show that after the introduction of the speed bump the majority of order flow on Alpha shifts towards trades that do not employ smart-order routers (SOR) nor exhaust the top of the order book. We also show that liquidity suppliers on Alpha who can monitor the market in ultra-high frequency are able to harness the information contained within the order flow, avoiding SOR "sprays" attempting to simultaneously access liquidity on all venues. This significantly reduces liquidity suppliers' adverse selection on Alpha, increasing realized spreads and producing substantial economic benefits in an otherwise low-realized spread environment.[34]

We find that the addition of the speed bump in combination with other changes on Alpha caused market wide effects. We find that Alpha's ability to utilize the speed bump to segment predominantly uninformed, low-impact order flow increases the fraction of informed traders on the remaining venues, increasing quoted and effective spreads on the consolidated market by about half a cent. This increase in

[34] Conrad and Wahal (2016) find that market making in the U.S. has become vastly less profitable, with realized spreads falling from 17 basis points in 2000 to 1.5 basis points in 2015.

spreads is primarily driven by increases in adverse selection, consistent with an increase in the fraction of informed traders. We also observe negative market wide effects for large liquidity demanders, with significant increases in both the fraction of trades consuming all available NBBO depth, and walking the book to achieve execution.

In a competitive market for liquidity provision, the reduction in adverse selection on Alpha would be offset by *tighter* quoted spreads, providing an advantage to traders accessing Alpha. We find the opposite, with Alpha "matching" rather than "making" the NBBO. Alpha's time spent at the NBBO reduces from 60% to 36% immediately after the introduction of the speed bump, consistent with liquidity suppliers' desire to harness information from order flow on other venues – quoting alone at the NBBO removes the value of the speed bump. The proverbial "canary in the goldmine" is not useful if it is alone in the goldmine.

Our empirical results are robust to other changes which occur contemporaneously with the introduction of the speed bump, with none of these able to explain the observed impacts on market quality. Our results are further supported by our cross-sectional analysis, showing that the impacts of the speed bump are most pronounced for stocks infrequently constrained by the minimum tick size, with higher prices, low depth and in which best bid and offer liquidity is frequently exhausted – precisely the type of stocks for which the speed bump is most useful and for which new Alpha is relatively better at segmenting order flow.

We also develop several novel empirical techniques and measures that facilitate our analysis of cross-market linkages and fairness, two particularly important issues for modern regulators and researchers. In particular, we highlight the importance of looking beyond traditional measures of market quality when evaluating market structure changes that involve fragmented order flow and low latency trading, both of which are pervasive features in today's equity markets. We develop techniques to correctly assign trade direction in fragmented markets, benchmark clock synchronization across multiple trading venues and link trades that likely originate from a SOR spray. From these methods, we develop metrics of quote fade, SOR usage, segmentation and order book fragility that empirically validate O'Hara and Ye's (2011) assertion that the combination of a trade-through prohibition and smart order routing in fragmented markets virtually replicate consolidated trading. At the same time, we show that these market linkages are being broken down by Alpha's speed bump, and its ability to segment uninformed order flow.

Speed bumps have been touted as a remedy in the "arms race for speed" by some (Budish et al., 2015; Baldauf and Mollner, 2015), and decried for their unequal access to markets by others. Our results have implications for both the debate surrounding the desirability of speed bumps, and the more general desirability of speed differentials between participants. It seems there are two key choices in constructing Alpha's speed bump which enhance the ability to segment retail order flow: the randomized 1-3 millisecond delay (which disrupts latency detecting smart order routers from synchronizing arrival times across venues, breaking down cross-market linkages) and the asymmetric application of the speed bump, which provides a guaranteed advantage to traders willing to pay for a "de minimis" speed advantage.

Ultimately, not all speed bumps are created equal - the devil is in the details. The speed bumps proposed by IEX, CHX, NYSE and NASDAQ all differ subtly in their construction from that of Alpha. Differences in deterministic versus random delays, and symmetric versus asymmetric applications of such delays will likely result in differing outcomes for market quality, which we leave to future researchers. What we have shown is that these nuances matter, and can generate market wide consequences. Our research highlights that caution is warranted for proposals which lead to the provision of a systematic speed advantage to any class of participant – speed bump or otherwise.

References

Angel, James J., 2014. "When Finance Meets Physics: The Impact of the Speed of Light on Financial Markets and Their Regulation." *Financial Review* 49 (2): 271-281.

Aquilina, Matteo and Foley, Sean and O'Neill, Peter and Ruf, Thomas. "Asymmetries in Dark Pool Reference Prices." FCA Occasional Paper No. 21

Baldauf, Markus and Mollner, Joshua. 2015. "High-Frequency Trading and Market Performance." Working Paper.

Baldauf, Markus and Mollner, Joshua. 2016. "Fast Traders Make a Quick Buck: The Role of Speed in Liquidity Provision." Working Paper.

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 (5): 2193-2238

Biais, Bruno, Thierry Foucault, and Sophie Moinas, 2015. "Equilibrium Fast Trading." *Journal of Financial Economics* 116 (2): 292-313.

Bessembinder, Hendrik. 2003. "Issues in Assessing Trade Execution Costs." *Journal of Financial Markets*, 6 (3): 233–57.

Brogaard, Jonathan, Björn Hagströmer, Lars Nordén, and Ryan Riordan. 2015. "Trading Fast and Slow: Colocation and Liquidity." *Review of Financial Studies*, 28 (12): 3407-3443

Brolley, Michael, and Katya Malinova. 2013. "Informed Trading and Maker-Taker Fees in a Low-Latency Limit Order Market." Working Paper.

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 (4): 1547-1621

Carrion, Allen. 2013. "Very Fast Money: High-Frequency Trading on the NASDAQ." *Journal of Financial Markets*, High-Frequency Trading, 16 (4): 680–711.

Chakravarty, Sugato, and Asani Sarkar. 2002. "A Model of Broker's Trading, with Applications to Order Flow Internalization." *Review of Financial Economics* 11 (1): 19–36.

Cimon, David. 2016. "Broker Routing Decisions in Limit Order Markets." Working Paper.

Clark, Doug. 2014. Letter to the Ontario Securities Commission, sourced at: https://www.osc.gov.on.ca/documents/en/Marketplaces/com_20141208_td-securities.pdf.

Comerton-Forde, Carole, and Tālis J. Putniņš. 2015. "Dark Trading and Price Discovery." *Journal of Financial Economics* 118 (1): 70–92.

Conrad, Jennifer, and Sunil Wahal. 2016. "The Term Structure of Liquidity Provision". Working Paper.

Conrad, Jennifer, Sunil Wahal, and Jin Xiang. 2015. "High-Frequency Quoting, Trading, and the Efficiency of Prices." *Journal of Financial Economics* 116 (2): 271–91.

Davies, Ryan and Sang Soo Kim. 2009. "Using matched samples to test for differences in trade execution costs." *Journal of Financial Markets* 12 (2): 173-202

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 (3): 811–33.

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 (4): 529–51.

Goldstein, Michael A., Amy Kwan, and Richard Phillip. 2016. "High-frequency Trading Strategies and Execution Costs" *Unpublished Manuscript, University of Sydney.*

Harris, Larry. 2013. "What to Do about High-Frequency Trading." *Financial Analysts Journal* 69 (2)

Hendershott, Terrence, Charles M. Jones, and Albert J. Menkveld. 2011. "Does Algorithmic Trading Improve Liquidity?" *The Journal of Finance* 66 (1): 1–33.

Holden, Craig W., and Stacey Jacobsen. 2014. "Liquidity Measurement Problems in Fast, Competitive Markets: Expensive and Cheap Solutions." *The Journal of Finance* 69 (4): 1747–85.

Kwan, Amy, Ronald Masulis, and Thomas H. McInish. 2015. "Trading Rules, Competition for Order Flow and Market Fragmentation." *Journal of Financial Economics* 115 (2): 330–48.

Larrymore, Norris L., and Albert J. Murphy. 2009. "Internalization and Market Quality: An Investigation." *Journal of Financial Research* 32 (3): 337–63.

Laughlin, Gregory, Anthony Aguirre, and Joseph Grundfest. 2014. "Information Transmission between Financial Markets in Chicago and New York." *Financial Review* 49 (2): 283-312.

Lee, Charles M. C., and Mark J. Ready. 1991. "Inferring Trade Direction from Intraday Data." *The Journal of Finance* 46 (2): 733–46.

Li, Wei. 2014. "High Frequency Trading with Speed Hierarchies." Working paper.

Maglaras, Costis, Ciamac C. Moallemi, and Hua Zheng. 2015. "Optimal Execution in a Limit Order Book and an Associated Microstructure Market Impact Model." Working Paper.

Malinova, Katya, and Andreas Park. 2015. "Subsidizing Liquidity: The Impact of Make/Take Fees on Market Quality." *The Journal of Finance* 70 (2): 509–36.

Malinova, Katya, and Andreas Park. 2016. "'Modern Market Makers" *Unpublished Manuscript, University of Toronto*.

Menkveld, Albert J. 2013. "High Frequency Trading and the New Market Makers." *Journal of Financial Markets*, High-Frequency Trading, 16 (4): 712–40.

Menkveld, Albert J. 2014. "High-Frequency Traders and Market Structure." *Financial Review* 49 (2): 333-344.

Menkveld, Albert J., and Marius A. Zoican. 2017. "Need for Speed? Exchange Latency and Liquidity." *Review of Financial Studies*, 30 (2017), 1188-1228

O'Hara, Maureen. 2015. "High Frequency Market Microstructure." *Journal of Financial Economics* 116 (2): 257–70.

O'Hara, Maureen, and Mao Ye. 2011 "Is Market Fragmentation Harming Market Quality?" *Journal of Financial Economics* 100(3): 459–74.

Rosu, Ioanid. 2016. "Fast and slow informed trading." Working Paper.

Shkilko, Andriy, and Sokolov, Konstantin. 2016, "Every Cloud Has a Silver Lining: Fast Trading, Microwave Connectivity and Trading Costs." Working Paper.

Tham, Wing Wah, Elvira Sojli, Johannes A. Skjeltorp. 2017. "Cross-Sided Liquidity Externalities." *Management Science* (online).

van Kervel, Vincent. 2015. "Competition for Order Flow with Fast and Slow Traders." *Review of Financial Studies*, 23 (7), 2094-2127.

Zhu, Haoxiang. 2014. "Do Dark Pools Harm Price Discovery?" *Review of Financial Studies*, 27 (3): 747-789.

Table 1
Summary Statistics

This table reports daily, univariate descriptive statistics across 232 TSX Composite Index component securities 6 months either side of Alpha's relaunch on the 21st of September 2015. *Quoted spread* and *Quoted Depth* are time-weighted and consolidated at the NBBO prices across Alpha, Chi-X, CX2 and TSX. *Deplete Best* is the fraction of trading that is part of a string displacing all NBBO depth. *Access Next* is the fraction of trading that occurs at the next best price behind NBBO. *Effective spreads* (net of fees) are calculated using the prevailing NBBO midpoint. *Adverse Selection* costs are calculated using the NBBO midpoint 20 seconds after the trade. *Time At Best* is the fraction of time a venue is quoting at the NBBO. Among all depleting trade strings, *Quote Fade* is the proportion of total visible starting liquidity on a venue at the NBB or NBO that did not result in trades. *Segmentation Imbalance* is the difference between a venue's share of non-depleting trading volume and its share of depleting trading volume. *Price* is the time-weighted NBBO midpoint. *% Unconstrained* is the fraction of time that the quoted spread was not tick-constrained. *Volume* is the total quantity of on-market traded shares. *Queue Length* is the time-weighted average number of shares at the NBBO. *Realized Volatility* is the standard deviation of one minute NBBO midpoint returns. *VIX* is the market volatility index for the TSX Composite (VIX).

	Pre-event Period				Post-event Period			
	Mean	**P10**	**Median**	**P90**	**Mean**	**P10**	**Median**	**P90**
Panel A: Consolidated Liquidity								
Quoted Spread (cents)	3.47	1.04	1.79	6.63	3.76	1.02	1.92	6.93
Quoted Depth ($'000s)	85.97	24.73	52.22	189.66	85.44	26.70	56.54	181.90
Deplete Best (%)	59.11	44.37	59.68	73.34	60.93	46.46	61.51	74.77
Access Next (%)	13.24	5.57	12.60	21.70	14.58	6.19	13.96	23.69
Panel B: Transaction Costs (cents)								
Eff. Spread on Alpha	2.95	1.22	1.69	5.44	3.46	0.96	1.88	6.60
Eff. Spread on Other	3.07	1.35	1.85	5.47	3.23	1.32	1.83	5.54
Adv. Selection on Alpha	2.94	0.95	1.84	5.43	2.13	0.26	1.16	4.31
Adv. Selection on Other	3.48	1.28	2.24	6.26	3.74	1.28	2.26	6.53
Panel C: Liquidity Provision by venue								
Time At Best on Alpha	60.34	31.22	61.07	88.15	35.82	15.81	35.40	55.43
Time At Best on CX2	36.80	15.93	33.46	62.30	42.30	25.56	42.28	58.54
Time At Best on Chi-X	66.94	34.30	72.62	92.59	68.29	36.45	72.00	95.18
Time At Best on TSX	94.10	85.95	96.28	99.57	95.89	90.21	97.26	99.78
Quote Fade on Alpha	14.89	2.70	11.62	31.60	65.97	32.84	72.48	88.27
Quote Fade on CX-2	24.24	5.26	22.45	43.75	16.78	4.62	15.15	30.16
Quote Fade on Chi-X	20.10	6.26	17.42	37.72	21.82	8.26	20.29	37.15
Quote Fade on TSX	9.51	2.78	7.13	19.25	9.09	2.65	6.60	18.54
Seg. Imbalance on Alpha	2.88	-3.31	2.62	9.54	8.70	1.34	7.78	17.13
Seg. Imbalance on CX-2	7.19	1.31	6.22	14.31	6.47	1.41	5.83	12.41
Seg. Imbalance on Chi-X	1.23	-4.72	0.88	7.42	-0.80	-6.91	-1.01	5.57
Seg. Imbalance on TSX	-11.30	-23.10	-11.53	0.88	-14.37	-26.56	-14.60	-1.81
Panel D: Stock Characteristics								
Price	35.63	6.10	23.51	63.39	32.57	4.32	19.23	59.36
% Unconstrained	0.52	0.04	0.52	0.98	0.53	0.01	0.58	0.99
Volume (000s)	735.51	87.70	362.40	1,901.80	1,040.15	107.30	462.00	2,618.40
Queue Length (000s)	7.63	0.85	2.09	13.24	8.20	1.08	2.65	18.41
Realized Volatility (bps)	9.46	4.40	7.91	16.67	12.50	5.53	10.56	22.06
VIX (%)	15.81	12.54	14.40	24.41	21.88	17.96	21.41	26.43

Table 2
Quote Fade and Segmentation Imbalance by Venue around Alpha Relaunch

This table reports the determinants of post-event changes in *Quote Fade* and *Segmentation Imbalance* separately for each of the 4 major exchanges using the following specification:

$$LiqAccess_{i,d,v} = \delta\ Post_d + \beta_1 Price_{i,d} + \beta_2 Volume_{i,d} + \beta_3 Volatility_{i,d} + \beta_4 VIX_d + FE_i + \varepsilon_{i,d,v}$$

where $LiqAccess_{i,d,v}$ is the *NBBO Quote Fade (Segmentation Imbalance)* for stock *i* on day *d* on venue *v*. $Post_d$ is the indicator for the period after (and including) the introduction of the speed bump, September 21, 2015. Daily stock-level control variables (calculated in log levels) include time-weighted average price, daily volume, realized 1-minute intra-day volatility, and the market volatility index for the TSX Composite (VIX), plus stock fixed effects. Among all depleting trade strings, *Quote Fade* is the proportion of total visible starting liquidity on venue *v* at the NBB or NBO that did not result in trades. *Segmentation Imbalance* is the difference between venue *v*'s share of non-depleting trading volume and its share of depleting trading volume. The sample period covers 10 weeks on either side of the event and runs from July 13, 2015 to November 25, 2015. */**/*** indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by stock.

	Quote Fade				Segmentation Imbalance			
	Alpha	**CX-2**	**Chi-X**	**TSX**	**Alpha**	**CX-2**	**Chi-X**	**TSX**
Post	46.37***	-2.87***	0.48*	0.66***	4.20***	1.03***	-0.91***	-4.32***
	(48.53)	(-9.11)	(1.76)	(4.37)	(17.27)	(6.23)	(-7.14)	(-15.28)
Price	-3.05	5.77***	0.15	1.39*	-1.95**	-0.30	0.69	1.56
	(-1.20)	(3.82)	(0.09)	(1.68)	(-2.09)	(-0.42)	(1.08)	(1.40)
Volume	-0.46	-1.14***	-5.27***	-2.15***	-1.37***	-3.06***	-1.65***	6.08***
	(-0.74)	(-3.86)	(-15.87)	(-10.16)	(-8.02)	(-19.52)	(-12.10)	(22.94)
Volatility	1.49	2.38***	4.03***	2.02***	0.75***	1.43***	1.48***	-3.67***
	(1.62)	(4.59)	(6.99)	(5.99)	(2.90)	(5.22)	(7.47)	(-9.42)
VIX	-2.07***	-0.53	-0.60	-1.61***	-1.54***	0.16	0.16	1.23***
	(-3.67)	(-1.02)	(-1.37)	(-6.61)	(-6.71)	(0.73)	(0.70)	(3.66)
R-squared	0.7152	0.2661	0.6271	0.7453	0.2189	0.2823	0.1972	0.3078
Observations	22,043	21,906	22,157	22,169	22,169	22,169	22,169	22,169

Table 3
Quote Fade and Segmentation Imbalance in the Presence of non-delayed Orders on Alpha

This table reports the determinants of post-event changes in *Quote Fade* and *Segmentation Imbalance* by venue as a function of whether speed bump-exempt limit orders were present at the NBBO on TSX Alpha, separately for each of the 4 major exchanges using the following specification:

$$LiqAccess_{i,d,v,s} = \delta\, Post_d + \gamma\, Post_d \times NonDelay_{i,d,s} + \theta^T Controls_{i,d} + FE_i + \varepsilon_{i,d,v,s}$$

where $LiqAccess_{i,d,v,s}$ is the average *NBBO Quote Fade (Segmentation Imbalance)* for trades in stock *i* on day *d* at venue *v* when Alpha's order book is in state *s*. For every trade in the post period, we observe whether, at the beginning of the trade, the NBBO depth on the passive side of the order book on TSX Alpha exceeds the minimum required size for non-delayed Post-Only limit orders (state *s=Non-Delay*) or not (state *s=Delay*). Trades with Alpha not at the NBBO fall into the latter category. Observations are excluded when trading volume in state *s* represent less than 10 percent of total volume for a given stock-day and venue. $Post_d$ is the indicator for the period after (and including) the introduction of the speed bump, September 21, 2015.

Among all depleting trade strings, *Quote Fade* is the proportion of total visible starting liquidity on venue *v* at the NBB or NBO that did not result in trades. *Segmentation Imbalance* is the difference between venue *v*'s share of non-depleting trading volume and its share of depleting trading volume. Daily stock-level control variables (calculated in log levels) include time-weighted average price, daily volume, realized 1-minute intra-day volatility, and the market volatility index for the TSX Composite (VIX), plus stock fixed effects. The sample period covers 10 weeks on either side of the event. */**/*** indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by stock.

	Quote Fade				Segmentation Imbalance			
	Alpha	**CX-2**	**Chi-X**	**TSX**	**Alpha**	**CX-2**	**Chi-X**	**TSX**
Post	5.86***	-6.77***	0.26	0.35	-0.15	0.21	0.60***	-0.66***
	(4.09)	(-10.97)	(0.67)	(1.27)	(-1.13)	(1.11)	(3.44)	(-2.87)
Post × Non-Delay	64.51***	7.37***	1.41	-0.21	17.80***	1.84***	-7.17***	-12.48***
	(48.33)	(10.96)	(1.64)	(-0.27)	(55.68)	(9.47)	(-27.86)	(-30.98)
Price	-2.78	7.12***	2.98	0.31	-2.58**	0.56	1.31*	0.71
	(-0.77)	(3.59)	(1.16)	(0.17)	(-2.59)	(0.82)	(1.74)	(0.62)
Volume	-3.36***	-2.60***	-9.44***	-6.07***	-1.35***	-2.74***	-1.44***	5.53***
	(-3.66)	(-4.58)	(-14.71)	(-11.10)	(-8.21)	(-16.57)	(-10.97)	(23.55)
Volatility	3.90***	4.59***	9.15***	3.33***	0.95***	1.21***	1.38***	-3.54***
	(2.76)	(4.62)	(7.21)	(3.72)	(3.33)	(4.19)	(6.36)	(-9.72)
VIX	1.37	-0.37	-1.35*	-2.42***	-1.71***	0.73***	-0.15	1.13***
	(1.24)	(-0.38)	(-1.78)	(-5.15)	(-7.16)	(3.26)	(-0.65)	(3.35)
R-squared	0.2322	0.1333	0.2560	0.3992	0.5604	0.2043	0.3031	0.3618
Observations	29,302	29,260	29,470	29,478	29,491	29,491	29,491	29,491

Table 4
Trade-Based Liquidity Metrics on Alpha around Alpha Relaunch

This table reports changes in measures of transaction costs on Alpha for a sample of 232 TSX Composite stocks around the introduction of the speed bump using the following specification:

$$LiqMetric_{i,d,\alpha} = \delta\ Post_d + \theta^T Controls_{i,d} + FE_i + \varepsilon_{i,d,\alpha}$$

where $LiqMetric_{i,d,\alpha}$ is one of several transaction cost metrics on Alpha for stock *i* on day *d*: *Effective spreads* (net of fees) are measured against the prevailing NBBO midpoint, while *Realized spreads* (net of fees) and *Adverse selection* costs use the reference NBBO midpoint 1 second (20 seconds) after the trade. $Post_d$ is the indicator for the period after (and including) the re-launch date, September 21, 2015. Daily stock-level control variables (calculated in log levels) include time-weighted average price, daily volume, realized 1-minute intra-day volatility, and the market volatility index for the TSX Composite (VIX), plus stock fixed effects. Panel A presents metrics in cents while Panel B presents metrics in basis points. The sample period covers 10 weeks on either side of the event and runs from July 13, 2015 to November 25, 2015. */**/*** indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by stock.

	Panel A: Cents					**Panel B: Basis points**				
	Effective	**Realized Spread**		**Adverse Selection**		**Effective**	**Realized Spread**		**Adverse Selection**	
	Spread	**1 s**	**20 s**	**1 s**	**20 s**	**Spread**	**1 s**	**20 s**	**1 s**	**20 s**
Post	0.58***	1.10***	1.23***	-0.56***	-0.69***	0.26	3.71***	4.53***	-3.09***	-3.90***
	(5.80)	(9.22)	(9.63)	(-8.03)	(-8.06)	(1.08)	(18.46)	(20.95)	(-11.81)	(-13.35)
Price	2.47***	0.20	-0.71	2.26***	3.23***	83.28***	66.68***	68.40***	16.43*	14.69
	(4.90)	(0.44)	(-0.89)	(2.78)	(2.80)	(14.27)	(11.68)	(10.52)	(1.96)	(1.51)
Volume	-0.73***	-0.30***	-0.18**	-0.38***	-0.52***	-1.93***	-0.71***	0.06	-1.24***	-2.00***
	(-9.22)	(-4.99)	(-2.39)	(-5.36)	(-6.30)	(-9.32)	(-3.70)	(0.29)	(-6.54)	(-9.45)
Volatility	1.94***	0.43**	-0.28	1.37***	2.17***	6.11***	-1.42***	-4.62***	7.45***	10.78***
	(7.06)	(2.07)	(-1.14)	(9.47)	(9.99)	(16.09)	(-3.45)	(-9.63)	(20.07)	(22.46)
VIX	0.18**	0.07	0.03	0.21	0.28	0.04	0.84***	0.87***	-0.77***	-0.85***
	(2.02)	(0.59)	(0.18)	(1.12)	(1.05)	(0.18)	(3.40)	(3.24)	(-3.10)	(-2.79)
R-squared	0.8407	0.5265	0.3136	0.7194	0.7318	0.8466	0.4525	0.3358	0.5900	0.5760
Observations	22,088	22,088	22,088	22,088	22,088	22,088	22,088	22,088	22,088	22,088

Table 5
Per-Venue Liquidity Metrics on Other Venues around the Alpha Relaunch

This table reports measures of transaction costs separately for non-Alpha venues around the introduction of the speed bump using the following specification:

$$LiqMetric_{i,d,v} = \delta\ Post_d + \theta^T Controls_{i,d} + FE_i + \varepsilon_{i,d,v}$$

where $LiqMetric_{i,d,v}$ is one of several transaction cost metrics for stock *i* on day *d* on venue *v*: *Effective spreads* (net of fees) are measured against the prevailing NBBO midpoint, while *Realized spreads* (net of fees) and *Adverse selection* costs use the reference NBBO midpoint 20 seconds after the trade. $Post_d$ is the indicator for the period after (and including) the re-launch date, September 21, 2015. Daily stock-level control variables (calculated in log levels) include time-weighted average price, daily volume, realized 1-minute intra-day volatility, and the market volatility index for the TSX Composite (VIX), plus stock fixed effects. All measures are in cents. The sample period covers 10 weeks on either side of the event and runs from July 13, 2015 to November 25, 2015. */**/*** indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by stock.

	Effective Spreads			**Realized Spreads**			**Adverse Selection**		
	CX-2	**Chi-X**	**TSX**	**CX-2**	**Chi-X**	**TSX**	**CX-2**	**Chi-X**	**TSX**
Post	-0.10**	0.06	0.15***	-0.37***	-0.27***	-0.12***	0.21***	0.25***	0.31***
	(-2.12)	(0.94)	(3.15)	(-7.59)	(-6.80)	(-5.04)	(4.87)	(3.69)	(4.48)
Price	3.32***	2.95***	2.71***	0.69	-0.73***	-1.22***	2.40***	3.88***	4.06***
	(3.24)	(3.52)	(3.24)	(1.45)	(-3.70)	(-5.48)	(4.92)	(3.69)	(3.71)
Volume	-0.71***	-0.72***	-0.64***	-0.17***	0.09**	0.46***	-0.46***	-0.83***	-1.11***
	(-8.65)	(-8.31)	(-8.62)	(-3.42)	(2.39)	(12.63)	(-8.20)	(-9.69)	(-10.84)
Volatility	1.92***	1.81***	1.55***	-0.27**	-1.19***	-1.99***	2.02***	2.92***	3.59***
	(9.01)	(7.88)	(7.30)	(-2.04)	(-14.75)	(-15.83)	(13.29)	(11.03)	(9.83)
VIX	0.54***	0.46***	0.39**	0.44***	0.35***	0.41***	0.08	0.24	0.06
	(2.73)	(3.04)	(2.44)	(4.47)	(6.17)	(5.52)	(0.71)	(1.14)	(0.28)
R-squared	0.8575	0.8730	0.9174	0.4674	0.2167	0.3960	0.7069	0.7973	0.8845
Observations	22,035	22,159	22,169	22,035	22,159	22,169	22,035	22,159	22,169

Table 6
Liquidity Metrics of Inverted Venues in the Presence of non-delayed Limit Orders on Alpha (Post Period only)

This table directly compares measures of transaction costs on Alpha and CX2 after the introduction of the speed bump (when both have an inverted fee structure) using the following specification:

$$LiqMetric_{i,d,v,s} = \alpha\ Alpha_v + \gamma\ NonDelay_{i,d,s} + \eta\ NonDelay_{i,d,s} \times Alpha_v\ + \theta^T Controls_{i,d} + FE_i + \varepsilon_{i,d,v,s}$$

where $LiqMetric_{i,d,v,s}$ is one of several transaction cost metrics for stock *i* on day *d* on venue *v* when Alpha's order book is in state *s*. For every trade in the post period, we observe whether, at the beginning of the trade, the NBBO depth on the passive side of the order book on TSX Alpha exceeds the minimum required size for non-delayed Post-Only limit orders (state *s=Non-Delay*) or not (state *s=Delay*). Trades with Alpha not at the NBBO fall into the latter category. Observations are excluded when trading volume in state *s* represent less than 10 percent of total volume for a given stock-day and venue. *Effective spreads* (net of fees) are measured against the prevailing NBBO midpoint, while *Realized spreads* (net of fees) and *Adverse selection* costs use the reference NBBO midpoint 20 seconds after the trade. $Alpha_v$ is an indicator equal to 1 (0) for venue TSX Alpha (CX2). Daily stock-level control variables (calculated in log levels) include time-weighted average price, daily volume, realized 1-minute intra-day volatility, and the market volatility index for the TSX Composite (VIX), plus stock fixed effects. All measures are in cents. The sample period covers 10 weeks in the post period of the event, i.e. September 21, 2015 to November 25, 2015. */**/*** indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by stock.

	Effective Spread	Realized Spread	Adverse Selection	Effective Spread	Realized Spread	Adverse Selection
Alpha	0.16***	0.39***	-0.18***	0.05	0.01	0.08***
	(4.82)	(6.16)	(-5.06)	(1.55)	(0.18)	(2.90)
Non-Delay				0.23***	0.55***	-0.33***
				(6.66)	(11.41)	(-11.73)
Non-Delay × Alpha				0.07*	0.52***	-0.44***
				(1.86)	(8.75)	(-10.72)
Price	2.33***	-0.26	2.53***	2.14***	-0.31	2.27***
	(4.37)	(-1.11)	(5.25)	(4.19)	(-1.46)	(5.14)
Volume	-0.60***	-0.11***	-0.47***	-0.54***	-0.11***	-0.40***
	(-9.82)	(-2.62)	(-8.76)	(-10.03)	(-3.16)	(-7.95)
Volatility	1.71***	-0.39***	2.14***	1.54***	-0.39***	1.90***
	(9.79)	(-5.31)	(13.98)	(11.23)	(-6.52)	(14.05)
VIX	0.01	0.08	-0.21	0.05	0.19*	-0.22
	(0.05)	(0.65)	(-1.41)	(0.31)	(1.73)	(-1.60)
R-squared	0.8698	0.3562	0.6864	0.8534	0.3085	0.5878
Observations	21,866	21,866	21,866	40,947	40,947	40,947

Table 7
Fast and Slow Liquidity Providers on "New" Alpha and CX2

This table compares measures of transaction costs between the two major inverted venues (CX2 and Alpha) following Alpha's introduction of a speed bump, while also distinguishing the type of trader that is supplying liquidity, using the following specification:

$$LiqMetric_{i,d,v,t} = \alpha\ Alpha_v + \gamma\ FastLP_{i,d,t} + \eta\ FastLP_{i,d,t} \times Alpha_v\ + \Theta^T Controls_{i,d} + FE_i + \varepsilon_{i,d,v,t}$$

Where $LiqMetric_{i,d,v,t}$ represents a transaction cost metric for stock *i* on day *d* on venue *v* (either CX2 or Alpha) for trader type *t*: *Effective spreads* (net of fees) are measured against the prevailing NBBO midpoint, while *realized spreads* (net of fees) and *adverse selection* costs use the reference NBBO midpoint 20 seconds after the trade. All measures are in cents. $Alpha_v$ is an indicator equal to 1 (0) for venue TSX Alpha (CX2). Using broker-level trade ID, we distinguish liquidity providers that trade through 2 global banks that primarily offer direct market access services to proprietary traders (Type *FastLP*) and those that trade through other brokers (retail, other banks; type *SlowLP*). We exclude trades with anonymized broker ID. Observations are excluded when trading volume of trader type *t* represent less than 10 percent of total volume for a given stock-day and venue.
Daily stock-level control variables (calculated in log levels) include time-weighted average price, daily volume, realized 1-minute intra-day volatility, and the market volatility index for the TSX Composite (VIX), plus stock fixed effects. All measures are in cents. The sample period covers 10 weeks in the post period of the event, i.e. September 21, 2015 to November 25, 2015. */**/*** indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by stock.

	Effective Spread	Realized Spread	Adverse Selection	Effective Spread	Realized Spread	Adverse Selection
Alpha	0.16***	0.39***	-0.18***	-0.37***	-0.64***	0.35***
	(4.82)	(6.16)	(-5.06)	(-8.14)	(-9.10)	(7.14)
FastLP				0.52***	0.49***	0.05
				(12.45)	(7.26)	(1.03)
FastLP x Alpha				0.65***	1.12***	-0.51***
				(9.38)	(12.90)	(-9.58)
Price	2.33***	-0.26	2.53***	2.75***	-0.24	2.50***
	(4.37)	(-1.11)	(5.25)	(3.42)	(-0.69)	(6.71)
Volume	-0.60***	-0.11***	-0.47***	-0.56***	-0.00	-0.50***
	(-9.82)	(-2.62)	(-8.76)	(-9.22)	(-0.05)	(-9.59)
Volatility	1.71***	-0.39***	2.14***	1.73***	-0.50***	2.09***
	(9.79)	(-5.31)	(13.98)	(10.08)	(-6.23)	(15.52)
VIX	0.01	0.08	-0.21	0.39	0.15	-0.05
	(0.05)	(0.65)	(-1.41)	(1.61)	(0.95)	(-0.45)
R-squared	0.8698	0.3562	0.6864	0.8494	0.3091	0.6109
Observations	21,866	21,866	21,866	32,729	32,729	32,729

Table 8
Segmentation Ability of Inverted Venues by Stock Characteristic

This table shows how a) the willingness of market makers to post at the NBBO on inverted venues and b) the ability of inverted venues to segment order flow in the market varies across stock characteristics. We investigate all 4 historical events when a Canadian exchange launched with or switched to an inverted maker/taker fee structure: Omega and CX-2 launched as inverted venues on December 20, 2011 and May 3, 2013, respectively; TMX Select switch to an inverted fee structure on November 1, 2013; Alpha relaunched as inverted venue on September 21, 2015. For each event, we split samples of around 240 TSX Composite stocks into 2 subsets, as a function of whether their average in-sample characteristic is above or below the median value for one of 4 characteristics: *Low vs. High Price* refers to the time-weighted average stock price; *Tick-Constrained vs. Unconstrained* refers to the proportion of time that the quoted spread is constrained by the minimum tick size; *Low vs. High NBBO Depth* refers to the time-weighted average number of shares at the NBBO. *Low vs. High Depletion* refers to the proportion of daily trading volume that displaced the entire depth on one side of the NBBO.

Each panel reports the mean level of both subsamples as well as a simple ratio of these means for a) *Time At NBBO* (i.e. the proportion of time a venue quotes at NBBO) and b) *Segmentation Imbalance* (i.e. the difference between a venue's share of non-depleting trading volume and its share of depleting trading volume) for each venue in the 10 weeks immediately following their respective launches or relaunches as inverted venues.

	Time At NBBO				Segmentation Imbalance			
	Omega Launch	**CX-2 Launch**	**TMX Post Switch**	**ALP Post Relaunch**	**Omega Launch**	**CX-2 Launch**	**TMX Post Switch**	**ALP Post Relaunch**
Panel A:								
Low Price	20.33	25.95	30.71	36.87	0.97	3.14	1.87	6.22
High Price	5.52	12.05	7.47	28.97	0.34	1.85	0.74	6.87
Ratio	3.68	2.15	4.11	1.27	2.85	1.70	2.53	0.91
Panel B:								
Tick-Constrained	22.71	28.18	32.26	39.89	1.11	3.46	1.99	6.94
Unconstrained	3.32	9.96	5.75	26.05	0.21	1.55	0.60	6.21
Ratio	6.84	2.83	5.61	1.53	5.29	2.23	3.32	1.12
Panel C:								
High NBBO Depth	23.04	28.25	32.60	40.02	1.14	3.49	2.02	6.90
Low NBBO Depth	3.02	9.89	5.43	25.94	0.19	1.52	0.57	6.25
Ratio	7.63	2.86	6.00	1.54	6.00	2.30	3.54	1.10
Panel D:								
Low Depletion	20.00	27.18	27.43	38.92	0.99	3.25	1.75	6.73
High Depletion	5.71	11.05	10.61	27.10	0.32	1.76	0.85	6.41
Ratio	3.50	2.46	2.59	1.44	3.09	1.85	2.06	1.05

Table 9

Cross-sectional Differences in Trading Costs after the Alpha Relaunch by Stock Characteristic

This table investigates how quote fade on TSX Alpha, segmentation for all other venues, and transaction costs on Alpha vs. other venues are affected around the introduction of the speed bump on Alpha on September 21, 2015 as a function of cross-sectional characteristics. Stocks are categorized based on one of 4 average pre-event characteristics: *Low vs. High Price* refers to the stock price; *Tick-Constrained vs. Unconstrained* refers to the proportion of time that the quoted spread is constrained by the minimum tick size; *High vs. Low Depth* refers to the time-weighted average number of shares at the NBBO. *Low vs. High Depletion* refers to the proportion of daily trading volume that displaced the entire depth on one side of the NBBO. We run the following test for each characteristic:

$$y_{i,d} = \delta\, Post_d + \gamma\, Post_d \times PreChar_i + \Theta^T Controls_{i,d} + \Phi^T Controls_{i,d} \times PreChar_i + FE_i + \varepsilon_{i,d}$$

where $y_{i,d}$ is one of several measures for stock *i* on day *d*. $Post_d$ is the indicator for the post-event period. $PreChar_i$ is an indicator that is equal to 1 for stocks in the latter half of one of the above characteristics. Among all depleting trade strings, *Quote Fade Alpha* is the proportion of total visible starting liquidity on TSX Alpha at the NBB or NBO that did not result in trades. *Segmentation Imbalance Other* is the difference between the combined share of non-depleting trading volume of the 3 other venues and their share of depleting trading volume. *Realized spreads* and *Adverse Selection* are measured with regards to the NBBO midpoint 20 seconds after the trade, separately for Alpha vs. the other 3 venues. Daily stock-level control variables (calculated in log levels) include time-weighted average price, daily volume, realized 1-minute intra-day volatility, and the market volatility index for the TSX Composite (VIX), plus stock fixed effects. In addition, all control variables are interacted with the $PreChar_i$ dummy. The sample period covers 10 weeks on either side of the event. */**/*** indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by stock.

	Quote Fade	Segment. Imbalance	Realized Spreads in Cents		Adverse Selection in Cents		Realized Spreads in Basis Points		Adverse Selection in Basis Points	
	Alpha	Other	Alpha	Other	Alpha	Other	Alpha	Other	Alpha	Other
Panel A:										
Post	45.84***	-2.91***	0.44***	-0.05***	-0.39***	0.04***	5.05***	-0.76***	-6.03***	0.67***
	(44.14)	(-10.78)	(12.86)	(-2.79)	(-13.67)	(2.71)	(16.85)	(-3.76)	(-12.09)	(2.78)
Post x High Price	1.09	-2.40***	1.49***	-0.12***	-0.58***	0.43***	-1.19***	0.50**	4.19***	-0.10
	(0.58)	(-5.33)	(6.93)	(-2.69)	(-3.82)	(3.42)	(-3.12)	(2.24)	(8.04)	(-0.38)
Panel B:										
Post	46.40***	-3.58***	0.42***	-0.05***	-0.45***	0.04***	4.31***	-0.69***	-5.65***	0.63**
	(48.42)	(-11.78)	(15.80)	(-4.35)	(-20.04)	(3.02)	(14.67)	(-3.91)	(-11.37)	(2.57)
Post x Unconstrained	-0.02	-1.15**	1.56***	-0.10**	-0.48***	0.44***	0.23	0.40*	3.47***	-0.07
	(-0.01)	(-2.43)	(7.30)	(-2.34)	(-3.22)	(3.49)	(0.60)	(1.90)	(6.55)	(-0.27)
Panel C:										
Post	46.84***	-3.83***	0.42***	-0.06***	-0.48***	0.03***	4.24***	-0.78***	-5.68***	0.74***
	(47.56)	(-11.84)	(14.78)	(-5.65)	(-19.50)	(2.99)	(14.36)	(-4.45)	(-11.49)	(2.99)
Post x Low Depth	-0.93	-0.64	1.54***	-0.07*	-0.43***	0.43***	0.34	0.57***	3.52***	-0.25
	(-0.49)	(-1.35)	(7.27)	(-1.72)	(-2.95)	(3.40)	(0.89)	(2.67)	(6.69)	(-0.92)
Panel D:										
Post	44.87***	-3.22***	0.43***	-0.08***	-0.37***	0.10***	4.19***	-0.79***	-5.38***	0.74***
	(39.01)	(-10.54)	(9.64)	(-3.70)	(-11.38)	(3.47)	(13.36)	(-4.08)	(-10.39)	(3.18)
Post x High Depletion	2.90	-1.84***	1.54***	-0.07	-0.61***	0.36***	0.53	0.65***	2.91***	-0.33
	(1.54)	(-3.96)	(7.03)	(-1.47)	(-3.99)	(2.83)	(1.34)	(3.02)	(5.25)	(-1.27)

.

Table 10

Consolidated Market Quality Metrics at NBBO around Alpha Relaunch

This table reports changes in consolidated market quality across all venues around the relaunch of Alpha using the following specification:

$$MQ_{i,d} = \delta\ Post_d + \Theta^T Controls_{i,d} + FE_i + \ \varepsilon_{i,d}$$

where $MQ_{i,d}$ is one of several market quality measures for stock *i* on day *d*: *Quoted spreads* and quoted (log dollar) *Depth* are time-weighted and consolidated at the national best bid and offer prices across Alpha, Chi-X, CX2 and TSX. *Effective spreads* (net of fees) are measured against the prevailing NBBO midpoint. *Deplete Best* represents the proportion of daily trading volume that occurred as part of a trade string that displaced the entire depth on one side of the NBBO. *Access Next* is the proportion of trading volume that occurs at any price behind NBBO. $Post_d$ is the indicator for the period after (and including) the re-launch date, September 21, 2015. Daily stock-level control variables (calculated in log levels) include time-weighted average price, daily volume, realized 1-minute intra-day volatility, and the market volatility index for the TSX Composite (VIX), plus stock fixed effects. The sample period covers 10 weeks on either side of the event and runs from July 13, 2015 to November 25, 2015. */**/*** indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by stock.

	Quoted Spreads		**Eff. Spreads**		**Depth**	**Deplete**	**Access**
	cents	**bps**	**cents**	**bps**		**Best**	**Next**
Post	0.24***	0.44***	0.17***	0.16*	0.16***	1.87***	1.74***
	(3.77)	(3.06)	(3.33)	(1.75)	(19.85)	(8.63)	(12.94)
Price	3.59***	79.39***	2.77***	107.18***	0.19**	12.83***	6.07***
	(3.39)	(26.10)	(3.34)	(39.19)	(2.43)	(10.16)	(7.74)
Volume	-1.05***	-3.62***	-0.67***	-1.79***	0.25***	-5.66***	-0.53***
	(-10.64)	(-14.36)	(-8.81)	(-10.31)	(21.71)	(-23.00)	(-3.25)
Volatility	1.82***	5.89***	1.62***	5.06***	-0.38***	12.81***	5.56***
	(7.65)	(15.92)	(7.49)	(18.38)	(-20.69)	(33.55)	(20.36)
VIX	0.74***	1.37***	0.39**	0.21	-0.15***	0.88**	-0.26
	(3.96)	(7.61)	(2.49)	(1.56)	(-9.02)	(2.43)	(-1.03)
R-squared	0.9354	0.9267	0.9171	0.9566	0.9419	0.6004	0.5816
Observations	22,169	22,169	22,169	22,169	22,169	22,169	22,169

Table 11
Consolidated Market Quality in the Presence of non-delayed Limit Orders on Alpha

This table reports changes in consolidated market quality across all venues as a function of whether speed bump exempt Post Only limit orders were present at the NBBO on Alpha, using the following specification:

$$MQ_{i,d,s} = \delta\, Post_d + \gamma\, Post_d \times NonDelay_{i,d,s} + \theta^T Controls_{i,d} + FE_i + \varepsilon_{i,d,s}$$

where $MQ_{i,d,s}$ is one of several market quality measures in stock *i* on day *d* when Alpha's order book is in state *s*. For every trade in the post period, we observe whether, at the beginning of the trade, the NBBO depth on the passive side of the order book on Alpha exceeds the minimum required size for non-delayed Post-Only limit orders (state *Non-Delay*) or not (state *Delay*). Trades with Alpha not at the NBBO fall into the latter category as well. Observations are excluded when trading volume in subset *s* represent less than 10 percent of total volume for a given stock-day and venue. $Post_d$ is the indicator for the period after (and including) the introduction of the speed bump, September 21, 2015.

Quoted spreads and quoted (log dollar) *Depth* are measured at the start of a trade and volume-weighted based on the consolidated national best bid and offer prices across Alpha, Chi-X, CX2 and TSX. *Effective spreads* (net of fees) are measured against the prevailing NBBO midpoint. *Deplete Best* represents the proportion of daily trading volume that occurred as part of a trade string that displaced the entire depth on one side of the NBBO. *Access Next* is the proportion of trading volume that occurs at any price behind NBBO. Daily stock-level control variables (calculated in log levels) include time-weighted average price, daily volume, realized 1-minute intra-day volatility, and the market volatility index for the TSX Composite (VIX), plus stock fixed effects. The sample period covers 10 weeks on either side of the event and runs from July 13, 2015 to November 25, 2015. */**/*** indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by stock.

	Quoted Spreads		Effective Spreads		Depth	Deplete	Access
	Cents	bps	Cents	bps		Best	Next
Post	0.00	-0.15	0.03	0.05	0.16***	3.75***	2.20***
	(0.02)	(-1.56)	(0.97)	(0.60)	(19.57)	(16.58)	(15.77)
Post × Non-Delay	0.50***	1.86***	0.34***	0.54***	0.02***	-10.80***	-2.65***
	(11.67)	(14.75)	(8.64)	(3.33)	(8.40)	(-28.75)	(-15.97)
Price	1.77***	-11.50***	1.80***	104.60***	0.18**	13.01***	5.71***
	(4.02)	(-7.41)	(4.35)	(36.10)	(2.39)	(11.23)	(7.52)
Volume	-0.59***	-2.22***	-0.55***	-1.57***	0.24***	-4.24***	-0.37**
	(-10.38)	(-14.89)	(-9.63)	(-9.50)	(20.17)	(-14.87)	(-2.11)
Volatility	1.24***	4.32***	1.31***	4.93***	-0.37***	11.33***	5.55***
	(10.50)	(16.49)	(10.97)	(18.43)	(-18.61)	(26.93)	(20.01)
VIX	0.20***	0.15	0.18***	0.24*	-0.16***	1.08***	-0.43*
	(3.03)	(0.70)	(2.78)	(1.67)	(-9.52)	(2.96)	(-1.70)
R-squared	0.9087	0.9399	0.9084	0.9452	0.9432	0.5229	0.4873
Observations	29,497	29,497	29,497	29,497	29,497	29,497	29,497

Table 12

Cross-sectional Differences in Consolidated Market Quality after the Alpha Relaunch by Stock Characteristic

This table investigates how consolidated measures of market quality are affected around the introduction of the speed bump on Alpha on September 21, 2015 as a function of cross-sectional characteristics. Stocks are categorized based on one of 4 average pre-event characteristics: *Low vs. High Price* refers to the stock price; *Tick-Constrained vs. Unconstrained* refers to the proportion of time that the quoted spread is constrained by the minimum tick size; *Thin vs. Thick Book* refers to the time-weighted average number of shares at the NBBO. *Low vs. High Depletion* refers to the proportion of daily trading volume that displaced the entire depth on one side of the NBBO. We run the following test for each subset:

$$y_{i,d} = \delta\ Post_d + \gamma\ Post_d \times PreChar_i + \Theta^T Controls_{i,d} + \Phi^T Controls_{i,d} \times PreChar_i + FE_i + \varepsilon_{i,d}$$

where $y_{i,d}$ is one of several measures for stock *i* on day *d*. $Post_d$ is the indicator for the post-event period. $PreChar_i$ is an indicator that is equal to 1 for stocks in the second half of one of the above characteristics. Quoted spreads and quoted (log dollar) Depth are time-weighted and consolidated at the NBBO. Effective spreads are measured against the prevailing NBBO midpoint. *Access Next* is the proportion of trading volume that occurs behind the NBBO. Daily stock-level control variables (calculated in log levels) include time-weighted average price, daily volume, realized 1-minute intra-day volatility, and the market volatility index for the TSX Composite (VIX), plus stock fixed effects. In addition, all control variables are interacted with the $PreChar_i$ dummy. The sample period covers 10 weeks on either side of the event. */**/*** indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by stock.

	Quoted Spreads		Effective Spreads		Depth	Access
	cents	bps	cents	bps		Next
Panel A:						
Post	0.02	0.01	0.01	-0.17	0.13***	1.73***
	(0.66)	(0.06)	(0.39)	(-1.18)	(13.04)	(8.85)
Post x High Price	0.39***	0.75***	0.28***	0.59***	0.04**	0.05
	(3.65)	(2.62)	(3.22)	(3.32)	(2.55)	(0.18)
Panel B:						
Post	-0.00	0.00	-0.01	-0.18	0.12***	1.37***
	(-0.08)	(0.00)	(-0.70)	(-1.62)	(11.79)	(8.39)
Post x Unconstrained	0.44***	0.79***	0.31***	0.59***	0.07***	0.74***
	(4.13)	(2.85)	(3.58)	(3.57)	(4.84)	(2.84)
Panel C:						
Post	-0.03*	-0.16	-0.02*	-0.24**	0.12***	1.37***
	(-1.69)	(-0.76)	(-1.96)	(-2.14)	(11.65)	(8.32)
Post x Low Depth	0.49***	1.05***	0.33***	0.67***	0.07***	0.75***
	(4.58)	(3.72)	(3.81)	(4.05)	(4.70)	(2.89)
Panel D:						
Post	0.06	0.07	0.02	-0.15	0.12***	1.53***
	(1.55)	(0.34)	(0.98)	(-1.23)	(11.89)	(9.78)
Post x High Depletion	0.36***	0.69**	0.28***	0.60***	0.07***	0.43
	(3.20)	(2.44)	(3.13)	(3.54)	(4.82)	(1.63)

Table 13
Market Quality Comparison with Matched and Interlisted U.S. Equities around the Alpha Relaunch

This table reports how measures of consolidated market quality experience different changes after Alpha's relaunch i) between a sample of 235 TSX Composite stocks and 235 matched equities in the U.S. and ii) between a subset of 83 interlisted TSX composite stocks and their respective U.S. exchange listings. Results are reported for A) the full sample, B) stocks with pre-period average price below the median, C) above the median using the following specification:

$$MQ_{i,d,c} = \theta Post_d \times Canada_{i,c} + \delta Post_d + \Theta^T Controls_{i,d,c} + \Phi^T Controls_{i,d,c} \times Canada_{i,c} + FE_{i,c} + \varepsilon_{i,d,c}$$

where $MQ_{i,d,c}$ is one of several market quality measures for stock *i* on day *d* in country *c*: *Quoted spreads* (in cents) and quoted *Depth* are time-weighted and consolidated at the NBBO. *Effective spreads* (in cents) are measured against the prevailing NBBO midpoint. $Post_d$ is the indicator for the period after the introduction of the speed bump, September 21, 2015. $Canada_{i,c}$ is a stock-level indicator equal to 1 for the Canadian listings of the sample. Daily stock-level control variables (calculated in log levels) include time-weighted average price, daily volume, realized 1-minute intra-day volatility, the market volatility index for the TSX Composite (VIX), and a linear time trend. In addition, all control variables are interacted with the country dummy. Individual stock fixed effects subsume the country dummy. The sample period covers 10 weeks on either side of the event. */**/*** indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by stock.

	Matched Sample			Interlisted Stock Pairs		
	Quoted Spreads	**Effective Spreads**	**NBBO Depth**	**Quoted Spreads**	**Effective Spreads**	**NBBO Depth**
<u>Panel A: Full Sample</u>						
Post x Canada	0.25*	0.23**	0.04***	0.25	0.22**	0.08***
	(1.90)	(2.37)	(2.62)	(1.38)	(2.00)	(4.84)
Post	-0.12	-0.06	0.02*	-0.23	-0.14	0.08***
	(-1.64)	(-1.57)	(1.82)	(-1.20)	(-1.11)	(6.08)
Observations	44,150	44,150	44,150	15,602	15,602	15,602
<u>Panel B: Low Price</u>						
Post x Canada	-0.10	-0.01	0.01	-0.04***	0.00	0.07***
	(-1.22)	(-0.15)	(0.22)	(-2.96)	(0.12)	(3.53)
Post	0.06	0.00	0.05**	-0.01	-0.01	0.09***
	(1.10)	(0.06)	(2.53)	(-0.93)	(-0.93)	(4.18)
Observations	22,154	22,154	22,154	7,332	7,332	7,332
<u>Panel C: High Price</u>						
Post x Canada	0.82**	0.62**	0.14***	0.59	0.45*	0.11***
	(2.03)	(2.16)	(5.44)	(1.59)	(2.01)	(4.19)
Post	-0.20	-0.12	-0.02	-0.44	-0.27	0.07***
	(-0.85)	(-0.54)	(-1.05)	(-1.19)	(-1.09)	(4.02)
Observations	21,996	21,996	21,996	8,270	8,270	8,270

Figure 1

We construct trade strings by joining all trades in the same direction separated by less than 30 milliseconds. A full explanation of the construction of trade strings can be found in Internet Appendix B. A trade string is called depleting when the entire NBBO depth is displaced following the trade. Among all depleting trade strings we calculate the NBBO *Quote Fade* as the proportion of starting liquidity that did not result in trades. *Segmentation Imbalance* is the difference between venue *v*'s share of non-depleting trading volume and its share of depleting trading volume.

Panel A: NBBO Quote Fade By Venue

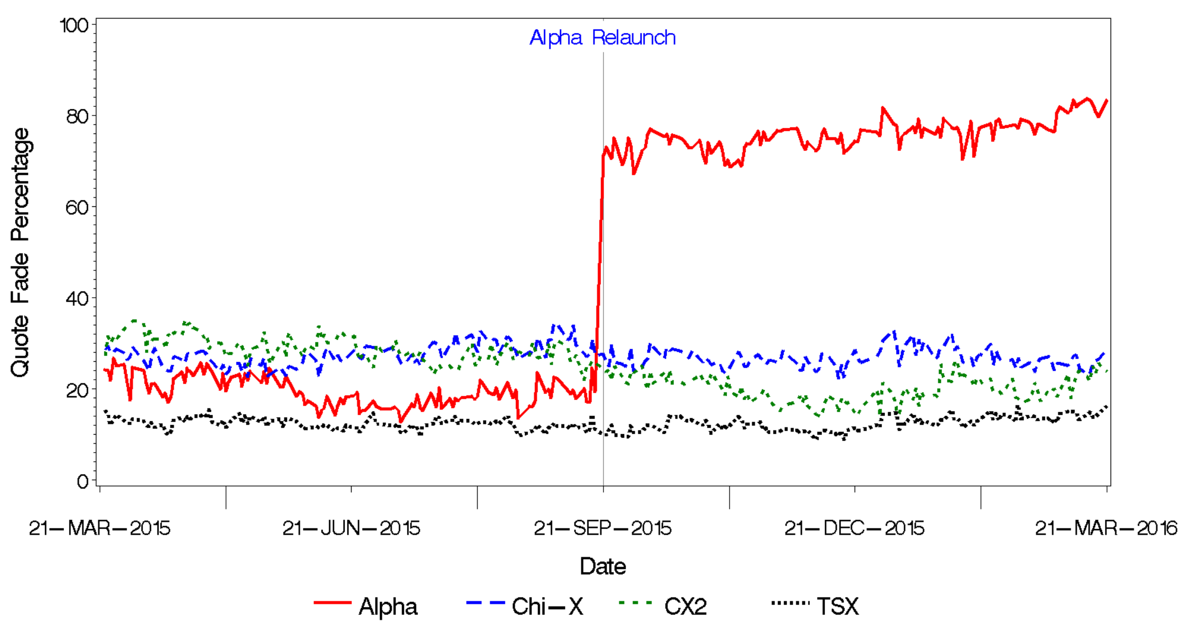

Panel B: Segmentation Imbalance by Venue

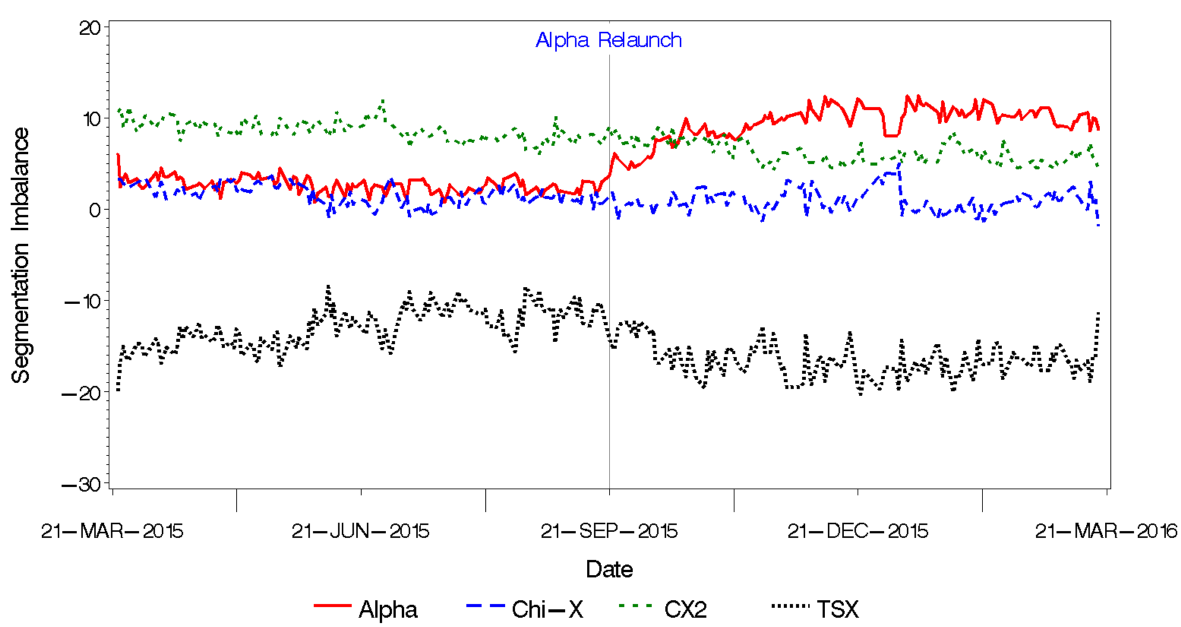

Figure 2
Realized Spread within One Minute by Number of Venues Accessed

This figure presents average realized spread associated with trade strings over 100 milliseconds, 1, 5, 10, 20 and 60 seconds. Trade strings are split into those which access only a single venue, and those which access multiple venues. Trade strings are defined as series of trades which execute within 30 milliseconds of each other. A full explanation of the construction of trade strings can be found in Internet Appendix B.

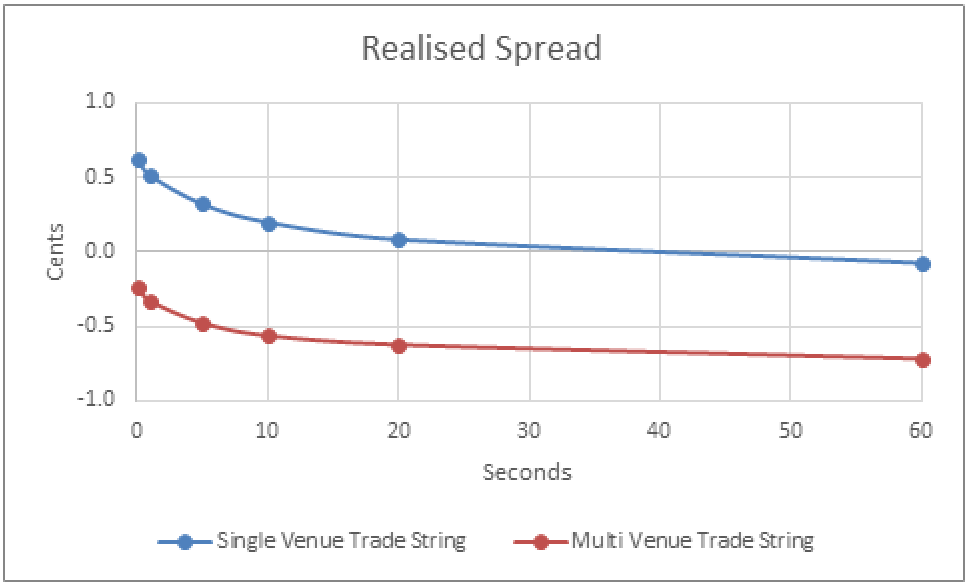

Figure 3
Trading Volume Composition by Trade String Type on Alpha

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 separated by less than 30 milliseconds. We distinguish between trade strings that leave the top level of quoted depth at the NBBO depleted vs. those that do not (undepleted). Smart order router (SOR) strings are those that execute on multiple venues.

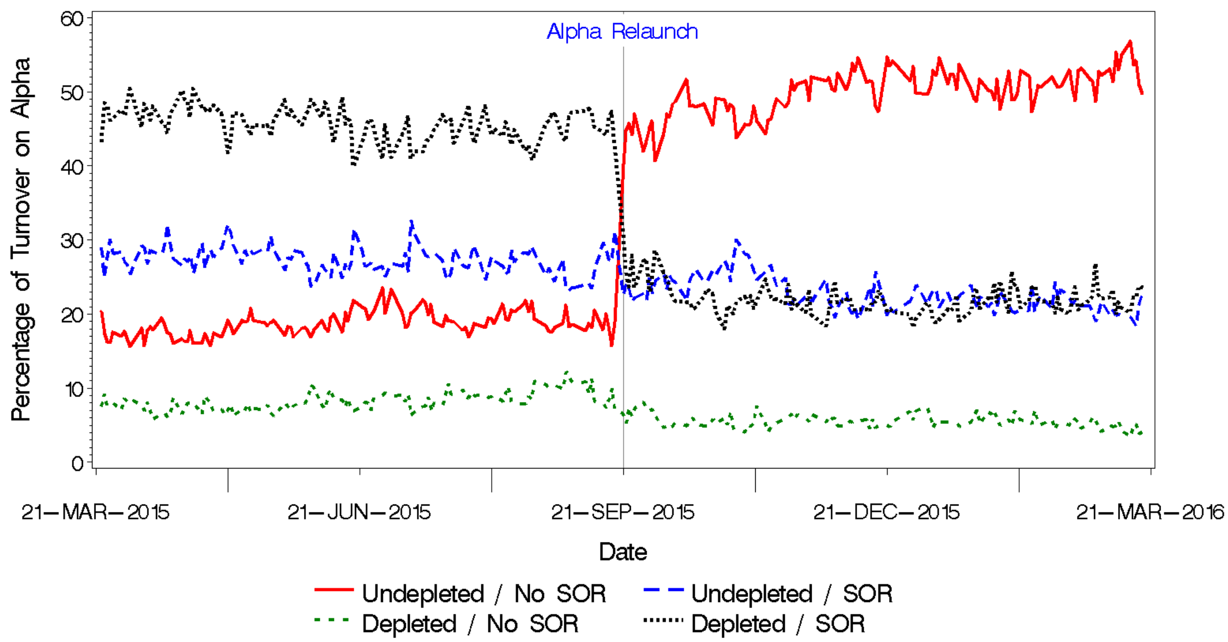

Figure 4
Aggressiveness of Quoting on Alpha

This figure presents the level of aggressiveness exhibited by quotes on venue Alpha over time, time-weighted during the day, equal-weighted across stocks and averaged between bid and ask quotes. A venue's best bid or ask can be one of three things: a) alone at the NBBO, with all other venues at inferior price points; b) at the NBBO jointly with other venues; c) behind the NBBO at an inferior price point.

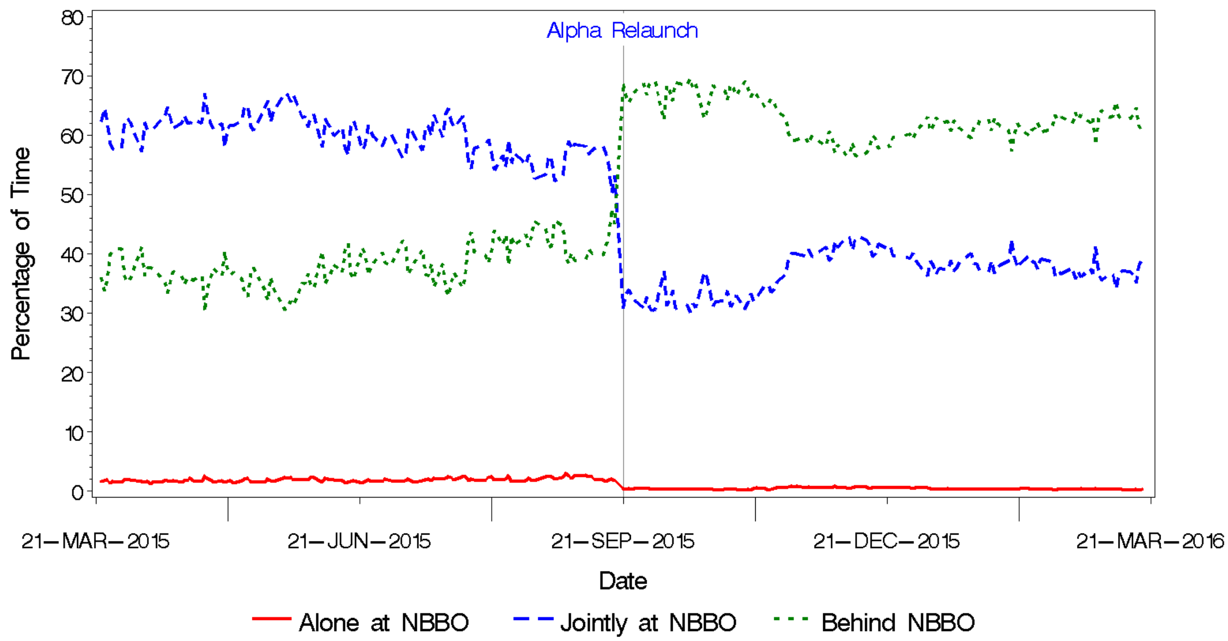

Figure 5
Incidence of "Trade Throughs" by Venue

To identify trade through events, we first identify all trade strings across any set of venues that accessed multiple price levels. We then define a trade through as the event where a venue had a quote at best at the beginning of the trade and at the end of the trade but was not accessed, i.e. no trade executed on that venue, while trades occurred at a worse price on another venue. Every stock-day and for every venue, we compute the ratio of volume executed at inferior prices, capped by the liquidity displayed at the NBBO on the traded-through venue at the start of the trade through event, over the total volume of all trades that accessed multiple price levels and where the traded-through venue was at best at the beginning. In other words, we compute the proportion of trading that could have traded at better prices on the traded-through venue during all multi-price level trades. This figure displays the average level of this ratio over time by venue.