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

Haoming Chen
University of New South Wales

Sean Foley[1]
University of Sydney

Michael Goldstein
Babson College

Thomas Ruf
University of New South Wales

November 17, 2016

Abstract

We examine the introduction of a speed bump by an existing exchange, which provides certain participants with guaranteed speed advantages. A selective order processing delay for marketable orders on TSX Alpha allows low-latency liquidity providers to avoid adverse selection through their ability to react to activity on other venues. These changes increase profits for liquidity providers on TSX Alpha but negatively impact aggregate liquidity: market-wide costs for liquidity demanders increase, with liquidity suppliers' profits reduced across remaining venues. Our findings have implications for the speed bump debate in the United States, speed differentials more generally, as well as the regulation of market linkages across fragmented trading venues.

[1] Contact Author: sean.foley@sydney.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, Thierry Foucault, Carole Gresse, Maureen Jensen, Petko Kalev (FIRN Discussant), Dan Kessous, Amy Kwan, Katya Malinova, Sophie Moinas, Bernt Odegaard, Andreas Park (NFA Discussant), Richard Phillip, Talis Putnins, Dagfinn Rimes, Ryan Riordan, Ioanid Rosu, Jos Schmitt, Andriy Shkilko, Elvira Sojli, Wing Wah Tham, Vincent van Kervel, Rene Wells, Pradeep Yadav, Yesha Yadav, and Haoxiang Zhu as well as seminar participants at the Ontario Securities Commission, Financial Conduct Authority UK, ITG MatchNow, Chi-X Canada, Aequitas Neo, the Northern Finance Association, the Asian Finance Association, 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.

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 are able to capture most of the profits of liquidity provision (Rosu, 2016) and impose adverse selection costs on relatively slower counterparts (Li, 2014), either by picking off stale orders, or updating limit orders faster than others in response to new market conditions or information revealed on another venue (van Kervel, 2015). This in turn motivates trading venues to compete on offering the fastest access possible to traders (O'Hara, 2015).

It is arguably the relative, not the absolute level of speed that matters, leading to a perpetual arms race for speed. 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 (Hoffman, 2014; Menkveld, 2014; Budish, Cramton and Shim, 2015; and Biais, Foucault and Moinas, 2015). Marginal increases in speed have become increasingly expensive as technology advances, through innovations such as colocation (Brogaard, Hagströmer, Nordén, and Riordan, 2015), signal transmission technology (Laughlin, Aguirre and Grundfest, 2014) and beyond.

Surprisingly, the latest development in the arms race for speed may actually be an intentional slowdown: speed bumps. Speed bumps and other such mechanisms are billed as an attempt to mitigate the advantages that the investment in faster trading technology provides. Among others, Harris (2013) and Budish, Cramton and Shim (2015) suggest small systematic delays could mitigate the structural advantages enjoyed by low latency traders. Such suggestions have recently gained, as evidenced by the U.S. Securities and Exchange Commission's (SEC) approval on June 17, 2016 of IEX's application for registration as a national securities exchange, with orders subject to a speed bump. A similar proposal by Chicago Stock Exchange is also currently under consideration.

An interesting feature often overlooked in the debate about speed bumps, however, is the fact that it need not uniformly apply to all messages and all investors. In the case of IEX, un-displayed (dark) pegged orders are re-priced by the exchange without delay, whilst orders submitted, amended or cancelled by participants are required to traverse the 350 microsecond speed bump. Thus, speed bumps have the potential to create (or add to) speed differentials between exchange venues, as well as between market participants, rather than alleviate them. This intentional speed differential is potentially valuable in fragmented markets, notably to traders with the ability to observe and react to order flow on *other* venues almost instantaneously. Whether this latest iteration in the relative speed race improves market outcomes and fairness *on balance* is an open question, which we address in this study.

We examine the effects of a recent introduction of a speed bump by the Canadian exchange venue TSX Alpha, which provides systematic speed advantages to some liquidity providers. The speed differential is created through the combination of two unique features: a randomized speed bump (between 1-3 milliseconds); and the ability for traders to pay higher fees for the right to enter and cancel limit orders *without* experiencing the delay. This effectively provides liquidity suppliers on TSX Alpha a guaranteed 1-3 millisecond window in which to cancel standing limit orders before any incoming marketable orders can access them. Why is this valuable to liquidity suppliers?

<Insert Figure 1>

Figure 1 plots realized spreads attained by liquidity providers in Canada prior to the introduction of the speed bump, divided into trade strings that access multiple exchanges vs. those that execute on one exchange only. It shows that multi-venue trade strings experience immediate and declining negative realized spreads, while trade strings which only execute on one venue benefit liquidity suppliers with positive realized spreads. As we will later argue, multi-venue trades, which account for 53 percent of trade volume in Canada, likely originate from institutional traders employing smart order routing technology, while single-venue trades are relatively more likely to stem from less informed retail traders.

Intuitively, the Alpha speed bump allows liquidity suppliers to avoid the loss-creating multi-venue order flow as they can observe executions on other venues, while still retaining the time to respond. At the same time, they are able to remain in the market for the lucrative single-venue order flow. The difference between the two realized spreads visible in Figure 1, then, is the economic value of being provided a millisecond in which to avoid order flow in the presence of immediate, adverse price impact.

Our findings for Alpha in the post-event period are consistent with this intuition. In short, we observe a remarkable increase in quote fade on Alpha from around 14% to about 60% of displayed liquidity, which shows that liquidity providers are using the speed bump to cancel many of their limit orders after trades begin on other venues, but before market orders can reach them. When we investigate how the composition of trades changes after the event, we find that likely institutional trading declines by half while the proportion of order flow with low price impact more than doubles, confirming that liquidity providers on Alpha are able to avoid interacting with orders that adversely select them. Lastly, we document that realized spreads on Alpha increase as predicted by our intuition from Figure 1, while the adverse selection component of the spread decreases. Both of these changes greatly improve the profitability of liquidity providers on Alpha.

Our evidence suggests that advance knowledge of institutional investors' trading intentions, even at the millisecond granularity, is valuable (see also van Kervel and Menkveld, 2016). We find that durations as short as one millisecond can allow for substantial information leakage of trading intentions. Our paper documents a key insight into the mechanism driving 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 the original trader is able to access them.

As part of this analysis, we propose a number of new techniques to examine market linkages and the accessibility of liquidity in fast, fragmented markets.[2] First, we develop a methodology to benchmark time synchronization across venues which is robust to any market structure, which allows us to aggregate related trades on different venues into trade strings – a tool which allows for the examination of low latency cross-market liquidity dynamics. This allows us to test whether, in the presence of a speed bump, liquidity faded after executions on other venues. It also allows us to determine what proportion of pre-trade displayed liquidity on each venue participants were actually able to access.

Second, we propose two innovative classification schemes that identify a) trade strings that cause instantaneous adverse selection costs on liquidity providers, termed depleting trade strings, and b) trade strings that include trades across multiple venues and are likely to have originated from a smart order

[2] A number of recent studies examine the issue of market linkages across fragmented trading venues (van Kervel, 2015; Malinova and Park, 2016; FCA, 2016; ESMA, 2016; AFM, 2016).

router (henceforth SORs). These proxies can be constructed with any public data feed and allow us to identify which types of traders continue to successfully interact with liquidity on Alpha, and which types of trades are no longer represented. We find a significant increase in trade strings on which do not displace the NBBO price level and do not originate from a SOR (access only one venue in the string). We find a significant reduction in trades that result in the removal of all liquidity at the price level and that originate from a SOR (access multiple venues). This evidence is further corroborated by rough proxies based on retail and high-frequency trader (HFT) broker IDs, with a significant increase in the proportion of passive liquidity supplying orders originating from HFT brokers, and the majority of aggressive, liquidity taking orders coming from retail brokers.

We find that the benefits afforded to high frequency traders on Alpha impose negative externalities on other traders and adversely affect market liquidity as a whole. We document that transaction costs increase significantly, with market wide quoted spreads increasing by 0.66 basis points, and effective spreads on non-Alpha venues increasing by 0.46 basis points on average. This is due to the adverse selection component of spreads on the remaining venues increasing by 0.67 basis points, consistent with an increase in the proportion of informed traders on those venues. Realized spreads on non-Alpha venues similarly reduce by about 0.19 basis points. In that sense, our findings are in line with a much earlier literature on the segmentation of order flow through payment for order flow schemes (for example Easley et al., 1996 and Chakravarty and Sarkar, 2002), but also with more recent studies on the segmentation of uninformed order flow imposed by dark trading (Zhu, 2014, Comerton-Forde and Putnins, 2015).

To summarize, our analysis shows that the speed bump implemented by TSX Alpha benefits one particular group of market participants, low-latency liquidity providers on Alpha, while institutional traders as well as liquidity providers on other venues face higher trading costs due to increased toxicity.

Taken as a whole, we provide the first piece of evidence on the impact of intentionally slowing down some participants relative to others in a fragmented environment. The speed bump introduced by TSX Alpha allows high frequency market participants to profit from the order flow information on other venues, removing liquidity at times they believe the risk of adverse selection to be high. This behavior contributes significantly to the problem of fleeting liquidity, and is responsible for increased market-wide transactions costs.

The remainder of this paper is organized as follows. Section 2 discusses information arising from the order flow in fast fragmented markets. Section 3 outlines the institutional details of the Canadian trading landscape, and in particular the newly implemented design changes on Alpha. Section 4 describes the data and methodology. Section 5 demonstrates these design changes lead to a segmentation of order flow across exchange venues. Section 6 assesses the impact on the market quality of other Canadian trading venues, while Section 7 concludes and discusses implications for regulators.

2. Information in Fast, Fragmented Markets

As noted by Cardella et al. (2014), Goldstein, Kumar, and Graves (2014) and others, there has been a recent evolution in markets towards more computerized trading, resulting in faster and faster markets. What famously started with carrier pigeons has evolved through the telegraph to telephone, co-location, fiber-optic cables and microwave towers: Laughlin, Aguirre, and Grundfest (2014) demonstrate that the placement of microwave towers between Chicago and New York are resulting in trade response functions approaching the speed of light. These changes have increased in both cost and technological complexity and, as Angel (2014) notes, raise interesting issues for financial markets and their regulation, and have contributed to an "arms race" for speed.[3] 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 don't. These technologies, along with innovations introduced by stock exchange operators, such as inverted pricing, dark trading with sub-penny price improvement, discretionary pegged orders and speed bumps are used by liquidity suppliers to attempt to avoid "toxic" order flow.

Within individual exchanges, price-time priority is typically enforced. Between venues price priority is frequently enforced,[4] with participants able to select any venue when multiple venues display the best price. Battalio et al. (2015) documents venues with inverted maker-taker pricing schemes experience lower adverse selection and higher realized spreads than conventional maker-taker pricing schemes. Without intermarket time priority, liquidity demanders maximize their welfare by first routing marketable orders to venues with the lowest fee (or highest rebate). Malagaras, Moallemi and Zheng (2015) argue that trading venues with inverted fee structures tend to interact with a larger proportion of small trades, which are less likely to impose instantaneous adverse selection costs.

Dark trading without pre-trade transparency is also a common feature of modern equities trading. Numerous exchanges offer limit orders with no pre-trade transparency and sub-penny price improvement as a functionality integrated in their continuous limit order books.[5] Zhu (2014) and Comerton-Forde and Putnins (2015) suggest that order flow that migrates to dark venues is more likely to be uninformed (and hence balanced in nature). Dark orders experience similar benefits to inverted markets since they are likely to interact with a larger proportion of uninformed trades which are less likely to impose the instantaneous adverse selection cost of sweeping the entire price level.[6]

Co-location provides another example of an innovation with the potential to allow fast liquidity providers to adjust their quotes to avoid adverse selection. Unsurprisingly Brogaard, Hagstromer, Norden and Riordan (2015) find that the fastest co-location services are utilized by low latency market makers. Studies examining the impact of co-location on market quality have found improvements in bid-ask

---

[3] Haldane (2012) warns that these competitions are often winner takes all, and can be socially deleterious. Menkveld (2014) suggests such an arms race could be "socially wasteful". Biais, Foucault, and Moinas (2015), suggest that it may be better to have a market for fast traders and one for slow traders.

[4] Commonly referred to as the trade-through prohibition in Canada and the U.S.

[5] These venues include TSX, Chi-X Canada, Chi-X Australia and many others. IEX Group and NYSE have also separately developed undisplayed discretionary pegged limit orders. These specialised order types continually monitor limit order book imbalance to avoid trading immediately prior to a "crumbling quote", where instantaneous adverse selection costs would have been incurred.

[6] Numerous studies have found that small percentages of dark trades occurring against NBBO reference prices are stale by a few milliseconds. Aquilina, O'Neill, Foley and Ruf (2016) find that this issue affects at least 3.54% of trades in the United Kingdom, IIROC (2016) find that this issue affects at least 4% of dark trades in Canada and ASIC (2015) find that less than 1% of dark trades in Australia are affected. IEX introduced a discriminatory order processing delay in the United States, where all incoming participant order entries, amendments and cancellations are delayed by 350 microseconds, whilst the exchange re-prices un-displayed pegged orders against the midpoint of the national best bid and offer prices without a delay. The asymmetry of IEX's speed bump ensures that participants using this specialized dark order type are protected from transacting at stale reference prices

spreads (Boehmer, Fong and Wu, 2014) or increases in depth (Gai, Yao and Ye, 2013) consistent with a reduction in adverse selection.

Trading participants have also invested heavily in technology to compete in a winner-takes-all arms race to transmit and process order flow information the fastest (Budish et al., 2016). With fiber optic cables, microwave towers (Shkilko and Solokov, 2016) and laser beams linking geographically dispersed trading centers, these investments allow traders to harness the order flow information on one venue for their trading strategies on other venues, rather than analyze stock-specific or macroeconomic news. Participants able to observe price movements on one venue can successfully avoid (or impose) instantaneous adverse selection costs across other venues, maximizing trading profits. Such strategies have been shown to extend to creating "smokescreens" by quote-stuffing, to slow the information processing capacity of other traders (Egginton, Van Ness and Van Ness, 2016).

Our paper examines the latest incarnation of methods to avoid adverse selection – speed bumps. Market operators have recently started introducing discriminatory systematic order processing delays to provide some participants with guaranteed latency advantages. This issue is at the heart of the HFT arms race, as small latency advantages are only relevant for low latency participants. In the context of a speed bump, market participants pay the trading venue operator a higher fee in exchange for guaranteed latency advantages, rather than having to invest in new infrastructure. However, the outcomes are identical – some participants are able to use the order book information on one venue for their trading strategy on another venue. The speed bump's advantage over other mechanisms to avoid adverse selection is that it is able to segment *all* incoming order flow, rather than accessing order flow with lower aggregate toxicity.

3. Institutional Details

Similar to the United States, Canadian equities trading is fragmented across multiple venues, with six lit venues and three dark trading venues.[7] Securities are listed on the Toronto Stock Exchange, operated by the TMX Group and retains approximately 60 percent market share of trading activity. The TMX Group also operates Alpha and TMX Select (which was decommissioned once the "new" Alpha was launched), whilst NASDAQ operates both Chi-X Canada and CX2. Other venues include Omega, Pure Trading, Aequitas Neo, Aequitas Lit and three dedicated continuous dark pools, Match Now, Instinet and Liquidnet.

Unlike the U.S., internalization of retail order flow in Canada has been significantly constrained. Brokers wishing to internalize trades of less than 5,000 shares were required to provide one full tick of price improvement, or a half tick when the bid ask spread is one tick wide.[8] This mechanism prevented the growth of retail internalizing venues such as those that exist in the United States, which account for

[7] The Canadian market is comparatively unfragmented when compared with the United States. 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 U.S. (Tabb Group and Rosenblatt Securities), but only 6 percent in Canada.

[8] For further details of this change, see Larrymore and Murphy (2009).

around 22 percent of trading (Kwan, Masulis and McInish, 2015).[9] As a result of this regulation, and the subsequent banning of payment-for-order-flow, retail orders remain predominantly on-exchange in Canada.

3.1 The Alpha Speed Bump

Alpha Exchange was launched in 2008 and was merged with 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 for all non-post only orders of between 1-3 milliseconds.
2. Minimum size requirements for post-only orders,[10] typically 5 board lots per quote.[11]
3. Inverted maker-taker pricing model.
4. Orders on Alpha are no longer subject to the Order Protection Rule.

Prior to Alpha's speed bump implementation, several market participants noted that it may result in undesirable consequences. For example, TD Securities[12] argued that *"the introduction of speed bumps on both Alpha and Aequitas will slow down the operation of smart order routers … aggravating quote fade across all marketplaces"* and ITG Canada[13] claimed that *"the new Alpha design will allow passive post only resting orders the ability to fade should they see trading on another venue"*. These concerns are depicted in the diagram below. Institutional investors who require more liquidity than what is displayed on any single trading venue may utilize a SOR to simultaneously spray marketable orders across multiple trading venues, efficiently accessing consolidated quoted depth at the national best bid or offer price. Alpha's randomized speed bump enable its' liquidity suppliers to observe the first legs of any large SOR spray being executed on other venues, and have 1-3 milliseconds to cancel their limit orders and avoid adverse selection costs, should they deem those orders informed.

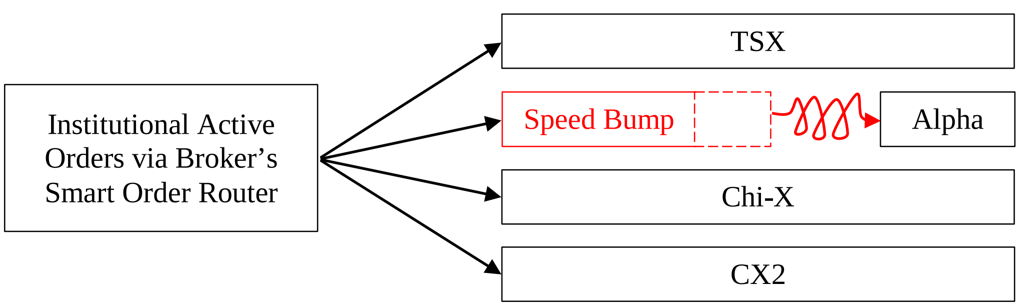

The application of the speed bump to incoming marketable orders gives rise to some interesting market dynamics. As reported by van Kervel, (2015), many liquidity suppliers duplicate their quotes across

[9] 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 U.S. would also be subject to minimum price improvement regulations. As such, unlike in the U.S., internalization is not a common practice.
[10] Post-only orders are limit orders which will cancel if they become marketable, ensuring they can only provide liquidity to the market.
[11] Minimum post only volumes for each security are available at http://api.tmxmoney.com/en/research/minpo.csv
[12] https://www.osc.gov.on.ca/documents/en/Marketplaces/com_20141208_td-securities.pdf
[13] https://www.osc.gov.on.ca/documents/en/Marketplaces/com_20141208_itg-canada-corp.pdf

multiple venues. This enables them to maximize the probability of execution, but also necessitates that liquidity demanders enter orders across a variety of venues in order to access all available liquidity. This duplication of orders allows liquidity suppliers the opportunity to remove duplicate orders subsequent to the first execution, leading to what many term "phantom liquidity". The introduction of a speed bump for incoming marketable orders *but not limit order entries or cancellations* allows liquidity suppliers who are able to monitor and respond to changes in the market in under 1-3 milliseconds[14] to cancel their standing limit orders subsequent to observing trades in other venues. Such conduct makes it unattractive for traders using SORs for large orders to include Alpha in their routing table, as the speed bump provides an opportunity for liquidity suppliers on Alpha to remove orders after observing trades on other venues, particularly if all available liquidity at a particular price level has been removed. [15]

A trader may be tempted to put Alpha *first* on the routing table, and to route to other venues once the order has resolved. While this strategy may provide superior access to limit orders on Alpha, the randomized delay of 1-3 milliseconds provides uncertainty about when to send the remainder of the order. This random delay provides an opportunity for fast liquidity suppliers to pull their limit orders from non-Alpha venues, especially if all available liquidity at a price level on Alpha has been removed. In such a situation, the optimal trading strategy may be to send *all* orders to Alpha when the desired quantity can be filled, and send *none* of the order to Alpha otherwise. This ability to "fade" away from large orders makes the "new" Alpha an undesirable venue for institutional traders. Importantly, such concerns are much less relevant for retail traders who are less likely to demand an entire price level.

To further attract retail traders, Alpha has employed an inverted maker-taker model. The maker-taker pricing model has been used to reward the provision of lit market liquidity in the United States[16] and Canada since 2005.[17] Since 2011, the proliferation of alternative trading venues in Canada led many venues to adopt inverted maker-taker pricing (such as CX2, TMX Select and Omega). Inverted maker-taker pricing provides a rebate to the demander of liquidity, which is paid for by the liquidity supplier. On the 21st of September (when the "new" Alpha was launched as an inverted maker-taker market) the existing TMX Select inverted market was decommissioned. Table 1 provides an explanation of the current fee structure of each of the major Canadian markets.

< Insert Table 1 Here >

Alpha's provision of an inverted maker-taker structure encourages fee-sensitive brokers to route aggressive orders to their venue, particularly if the taker rebate is not passed through to the client (such as when a flat fee is levied regardless of maker-taker rebates). This flat fee structure is common for retail brokers in Canada, as noted in Brolley and Malinova (2013).

[14] Given co-location and Menkveld's (2013) upper bound estimate of 1.67milliseconds round trip latency it seems likely fast participants are able to cancel orders within the speed bump duration.
[15] In the United States, order protection rules protect displayed quotations at the best bid or best offer from being traded through at other venues. In Canada, this protection extends to all levels of the order book, not just the top.
[16] For more detailed explanation of the usage of maker-taker fee structures see Battalio, Shkilko and Van Ness (2016).
[17] For a more detailed explanation of the introduction of maker-taker to Canada see Malinova and Park (2015).

The “new” Alpha was also removed from the order protection rule,[18] which requires any incoming marketable order to be sent to the venue displaying the best price prior to accessing liquidity on any other market at an inferior price. A condition of Alpha’s regulatory approval was that it would not be a “protected marketplace”, owing to the randomized delay, which would make it impractical for marketable orders to have to execute at prices quoted with a speed bump.[19]

The minimum passive post only volume requirement on new Alpha (typically 5 board lots per quote for large securities) is also attractive to retail investors who prefer to execute active orders in one trade - with rebates if possible. The requirement that liquidity suppliers post a minimum size ensures that most average size retail orders can be completed, while the speed bump ensures that this minimum size requirement does not expose the liquidity supplier to orders with larger adverse selection costs.

The decommissioning of TMX Select, which used inverted maker-taker pricing to target active retail traders, would have resulted in active retail volume being redistributed amongst other trading venues, potentially reducing the toxicity of aggregate order flow. As such, any observed liquidity deterioration would need to overcome this redistributive effect on the consolidated Canadian equities market.

The creation of the “new” Alpha resulted in an immediate and significant reduction in market share, from just below 15% to around 4%, as shown in Figure 2. In recent months, Alpha’s market share has climbed back towards 10 percent. While a portion of this decline is attributable to a number of smaller securities ceasing to quote at all on Alpha, the reduction among the remaining equities is consistent with liquidity suppliers being unwilling to pay to post on Alpha, and consequently providing their liquidity on the remaining venues.

<Insert Figure 2 Here >

4. Data and Methodology

4.1 Data

The data for this study was sourced from Thomson Reuters Tick History (TRTH), supplied by the Securities Industry Research Centre of Asia Pacific (SIRCA). Data for seven Canadian trading venues is available from TRTH, 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.[20] 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 can exceed 200-300 milliseconds, making it impossible to precisely calculate NBBO prices and volumes. Weighing data accuracy and quality against sample completeness, we exclude these two venues as well[21]. This leaves TSX, Alpha, Chi-X and CX2 as the venues of interest in this paper. Our observation period runs

[18] Analagous to the “trade through” prohibition in the US.
[19] https://www.osc.gov.on.ca/en/Marketplaces_alpha-exchange_20150421_noa-proposed-changes.htm
[20] Aequitas Lit and Neo combined accounted for less than 1% of total on-market trading in TSX listed securities during our sample according to IIROC’s Report of market share by marketplace.
[21] TMX Select and Omega each account for less than 3 percent of trading volume.

from the 13th of July 2015 to the 27th of November 2015, accounting for ten weeks on either side of Alpha's market structure changes.[22] Our universe of securities spans all 236 securities which remain in the S&P TSX Composite Index for the duration of our sample.

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,[23] as well as the price and size for both the bid and ask. We download trades and quotes within the same exchange concurrently to preserve ordering within the same millisecond to enable accurate trade direction classification. Although several venues operate extended trading hours, we restrict our analysis to the trading hours of the TSX listing market, being 9.30am to 4.00pm. We remove trades whose qualifiers identify them as off-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]

4.2 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 assign trade initiation direction based on 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 et al. (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 A 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]

From this file, we construct the national best bid and offer (NBBO) prices and sizes. 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,

---

[22] We exclude the 26th of November, a NYSE trading holiday, the 27th of November, a partial NYSE trading holiday, the 21st of October, during which extreme volatility occurred in Canadian equities, and the 24th of August, a U.S. stock market "flash crash".

[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. Although CX2 offers broker preferencing, the data does not include these identifiers. Chi-X does not offer broker preferencing, but some trades contain broker identifiers.

[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 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 provides further details on benchmarking of cross-venue clock synchronization.

updated for each quote (q) across all venues, and measured for the total duration for which that quote prevailed ($Alive_q$).

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

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

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

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

Effective half-spreads are calculated as the difference between the trade price and the prevailing NBBO midpoint. Realized spreads compare trade prices with the NBBO midpoint twenty seconds after the trade. Similar to Conrad et al. (2015), we calculate realized spreads at intervals of one, five, ten and twenty seconds after each trade. For brevity, we report this metric after twenty seconds as our primary result. Price impacts are computed as the effective spread minus the realized spread. Following Malinova and Park (2015) in markets with maker-taker pricing, effective spreads may be increased by the taker fee for a net cost of demanding liquidity, whilst realized spreads may be reduced by the maker rebate for a net revenue attributable to liquidity provision. Per trade (t), these metrics are volume weighted.

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

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

$$Price\ Impact_{i,d} = Effective\ Spread_{i,d} - Realized\ Spread_{i,d} \tag{6}$$

4.3. Construction of High Frequency Trade Strings

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 new metrics that rely solely on readily available trade and quote data and are able to estimate the impact of phantom liquidity across venues.

Building on the measurement of arbitrage opportunities across geographically separated markets in Budish et. al. (2015), we evaluate the accuracy of time stamp synchronicity across venues by calculating the duration of locked/crossed markets. Figure 3 presents the distribution of the duration of locked/crossed markets in our sample. As the order protection rules in the U.S. and Canada prohibit the entry of an order which would lock or cross the market, the observance of any such period is mostly driven by non-synchronicity in the timestamps. We use this feature to characterize the maximum observed latency in our data where 30 milliseconds correspond roughly to the 95th-percentile in the distribution. Thus, by concatenating trades occurring within 30 milliseconds of each other according to their database time stamp, we are able to capture the vast majority of trades that occurred in close proximity to each other in real time.[27] Section B of the Internet Appendix provides a more detailed outline of the methodology.[28]

<Insert Figure 3 Here >

We use this 30 millisecond or 95% confidence interval to construct high frequency trade "strings" by grouping together all buyer or seller initiated trades for each security that occur within 30 milliseconds of the last trade in the same direction. While timestamps for any individual trade may exhibit latency, jitter, caching and lack of cross-venue synchronization, strings of trades that occur over short time intervals are likely related.[29] If there are multiple trades within a string, they may have originated from a SOR spray by a single participant or active competition for order flow by multiple participants.

For each trade string, we snapshot the state of the limit order book across each venue 1 millisecond before the start of the first trade, since order book updates are produced to show trades consuming liquidity. We also snapshot the limit order books across all venues 20 milliseconds after the end of the last trade, to allow sufficient time for the venues with slower clocks to update their order books to reflect the information of the last trade. Since this is less than the 30 milliseconds required to group trades, neither snapshot overlaps into the previous or the next trade string for the same security. Buyer initiated trade strings are compared with changes in the offer prices and sizes, while seller initiated trade strings are compared with changes in the bid prices and sizes, on each venue. For trades that occurred at the best price within each string (generally the prevailing NBBO price at the start of the string) we record the trade price, start time and end time, as well as recording the trade volume, start price, start volume, end price and end volume on each trading venue. Only trades occurring at the best prices within each string are analyzed, to enable trade attribution to the consumption of visible liquidity at each venue's best bid or offer price.

<Insert Table 2 Here >

[27] This is in line with the U.S. Securities and Exchange Commission's proposed Rule 613 Consolidated Audit Trail National Market System (CAT NMS), which will require clock synchronization for each trading venue to be within 50 milliseconds of Coordinated Universal Time; See
http://www.catnmsplan.com/web/groups/catnms/@catnms/documents/appsupportdocs/p571933.pdf

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

[29] Section C of the Internet Appendix describes the construction of high frequency trade strings in detail.

Table 2 provides summary statistics on the frequency, size, and duration of the trade strings we construct during the entire sample period. We separate trade strings into four groups according to the number of venues on which at least one related execution occurs. Generally, trade strings on one venue are more frequent but smaller, while trade strings that cover all four venues are much less frequent but larger. For example, while slightly over three-quarters of trade strings are single trade, single venue events, they account for slightly less than half of all trading by volume. Trade strings that execute across all four venues account for only 1.21% of the total number of trade strings, but represent over 6.36% of traded value due to their larger size. Consistent with the use of smart order routers by institutional traders, the frequency with which a trade string leaves the order book depleted increases monotonically in the number of venues accessed, from 32 percent for single-venue trade strings to 70 percent for four-venue strings.

While our definition of trade strings theoretically allows day-long trade strings (as long as each consecutive trade follows less than 30 milliseconds after the previous trade), in practice the duration of trade strings is rather short. The median length for multi-venue trade strings is between 9 and 17 milliseconds. Even at the 90th percentile, durations increase modestly to between 21 and 36 milliseconds respectively for 1-4 venues. Overall, the median length of a trade string executed across multiple venues is 11 milliseconds in our sample, which is comparable to the analysis of liquidity provision using regulatory data from IIROC in Malinova and Park (2016), and significantly smaller than the 100 millisecond snapshots taken by van Kervel (2015).[30] Instead, the results in Table 2 are more consistent with the time horizons in which high frequency traders are known to operate; see, for example, Hasbrouck and Saar (2013).[31]

4.4. Multi-Venue Trade Strings and Depletion of Top of Book Quotes

In parts of our analysis, we distinguish between trade strings that do move prices, i.e. deplete the top of the book, and those that do not. 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 called 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 an order smaller than pre-trade NBBO depth can also displace an entire price level if it leads to a large number of cancelations by liquidity suppliers during or immediately following the execution.

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

[30] Malovina and Park (2016) groups trades originating from a (unique) SOR conditional on being separated by less than 5 milliseconds between each trade, and less than 9 milliseconds in total for the full trade string. This finer grid is facilitated by the use of regulatory data from IIROC, which exhibits less cross-market jitter.

[31] We extend the methodological underpinnings of van Kervel (2015) from 100millisecond buckets to continuous time. Given that Bessembinder (2003) finds that trades tend to occur immediately after order book cancellations in the opposite direction, our new method avoids the potential endogeneity which could be generated by using 100 millisecond buckets (where order book changes *before* each trade could be associated with the trade itself).

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 et al., 2011; Carrion, 2013) or seconds (Conrad et al., 2015) after the trade. Figure 4 provides an example of the logic applied to constructing trade strings for the purpose of our metrics.

<Insert Figure 4 Here >

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. This allows us to divide trade string into four separate categories: depleting vs. non-depleting on the one hand; and multi-venue sweep orders vs. single-venue orders on the other.

For depleting trade strings (s), we calculate the **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 of 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.

$$NBBO\ Quote\ Fade_{i,d,v} = 1 - \frac{\sum_{s=1}^{S} Trade\ Volume_{i,d,v,s}}{\sum_{s=1}^{S} Max(Start\ Liquidity_{i,d,v,s}, Trade\ Volume_{i,d,v,s})} \tag{7}$$

Finally, within each trade string we calculate the relative proportion of trades that occurred at the next best price behind the national best bid (offer) price 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. This metric, called **"Take Next"**, captures the sufficiency of top-of-book liquidity where liquidity demanders sought to trade large amounts.

$$Take\ Next_{\ i,d,v} = \frac{\sum_{s=1}^{S} Trade\ volume\ at\ level\ 2\ or\ greater_{\ i,d,v,s}}{\sum_{s=1}^{S} Total\ trade\ volume_{i,d,v,s}} \tag{8}$$

4.5. Summary Statistics

For each liquidity metric and control variable formulated, Table 3 presents summary statistics for the ten weeks before and after Alpha's relaunch, along with the difference in means and t-statistics from a univariate test of statistical significance. In the post-relaunch period, quoted spreads averaged 3.67 cents whilst quoted depths averaged $92,690 at the national best bid and offer prices. Average share prices declined slightly from $31.84 to $29.65. Daily traded volume per security averaged slightly less than 1 million shares. Realized one minute intraday volatility decreased slightly by 3%.Trades originating from strings that displaced the entire NBBO depth accounted for 60% of volume, and 13% of trading volume "walked the book", with trades in a string occurring at prices behind the national best bid or offer price.

Effective spreads on Alpha increased from 2.86 cents to 3.48 cents, despite the adverse selection component of the spread decreasing from 3.09 cents to 2.17 cents. Liquidity suppliers on Alpha could choose to 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 suppliers on Alpha would instead be hit first. Consistent with a “matching” rather than “making” of the best price, new Alpha posted a price equal to the NBBO 34% of the time, compared to 59% 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 this ability to provide “phantom” liquidity, traders on new Alpha were able to access only 40% of the liquidity quoted at the time they tried to trade, compared to 94% on the old Alpha. Commensurate with the levels of significance described in Table 3, these changes represent dramatic shifts not only from the “old” Alpha, but also from the status quo enjoyed on the other measured markets.

< Insert Table 3 Here >

5. System Delay, Pricing Change, and Order Flow Segmentation

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 motivate why the new market design might lead to differential routing among trades 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 orders after observing large trades on other venues. Then, 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 also examine Alpha’s market share of trade strings that incur and avoid adverse selection costs. Finally, we analyze changes in realized spreads and adverse selection costs for trades on Alpha.

5.1. Fleeting liquidity and the mechanics of reducing adverse selection costs

Alpha’s speed bump of 1 to 3 milliseconds against incoming market orders provides an opportunity for liquidity suppliers to cancel their standing limit orders ahead of new marketable orders, particularly after observing large trades on other venues. For NBBO-depleting trade strings, we calculate quote fade on each trading venue by comparing the visible liquidity at the start of the string with the actual volume traded. If there is no quote fade, all visible liquidity results in trades. Our analysis of cross-venue liquidity access is at the NBBO only. Our analysis of cross-venue liquidity access is at the NBBO only. Analysis at a single specified price level allows us to attribute the consumption of liquidity by incoming active orders to passive limit orders visible immediately prior to the trades. The Canadian market is characterized by particularly low relative minimum tick sizes. Figure 5 presents the average fraction of trade strings which consume an entire level of liquidity per stock across our sample period. The majority of stocks have relative minimum tick sizes of 1-10 basis points, and between 50-80% of trade strings consume all available liquidity at the NBBO. Those stocks with wider relative minimum ticks of between 10-100 basis points experience significantly lower levels of depleting trade strings. The advantage of the Alpha speed bump to liquidity suppliers is likely much lower for these stocks due to the much lower levels of instantaneous adverse selection.

< Insert Figure 5 Here >

Figure 6 presents daily aggregate NBBO quote fade per trading venue, calculated as the total trade value among all depleting trade strings, divided by the total dollar value of visible liquidity available at the national best bid (offer) price at the start of seller (buyer) initiated trade strings. A sharp increase in quote fade is observed on Alpha immediately after the relaunch, whilst quote fade decreases slightly across TSX, Chi-X and CX2. We formally test for statistically significant changes in NBBO quote fade with equations of the form

$$Quote\ Fade_{i,d,v} = Post_d + Price_{i,d} + Turnover_{i,d,v} + Volatility_{i,d} + Depth_{i,d} + NBBO\ Depth\ Share_{i,d,v} + FE_i + e_{i,d} \quad (9)$$

where $Fill\ Rate_{i,d}$ is the total trade volume divided by the total starting liquidity among all trade strings, at the NBBO on venue v for stock i on day d, $Post_d$ is an indicator variable equal to one for observations after the 21st of September 2015 and zero prior, $Price_{i,d}$ is the natural logarithm of the time-weighted NBBO midpoint price, $Turnover_{i,d,v}$ is the natural logarithm of on-market trade turnover on each venue, $Volatility_{i,d}$ is the standard deviation of one minute NBBO midpoint returns, $Depth_{i,d}$ is the natural logarithm of the time-weighted consolidated depth at the national best bid and offer prices, $NBBO\ Depth\ Share_{i,d,v}$ is the percentage of consolidated depth at the national best bid and offer prices quoted by each venue, FE_i indicates stock fixed effects and $e_{i,d}$ is an error term. Observations are winsorized at the 1% level per day.

< Insert Figure 6 Here>

Table 4 reports that Alpha's quote fade increased by 44%. High quote fade indicates that quoted liquidity available at the start of a trade string after which the NBBO changes was removed before it could be traded against, representing increased quote fade. The ability to fade against the majority of trades that will incur 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. Liquidity being removed from the side of the book which is about to be very "thin" is consistent with the empirical findings of Goldstein, Kwan and Phillip (2016) that HFT liquidity suppliers primarily supply liquidity on the "thick" side of the book. As a consequence, Alpha becomes unattractive for larger parent orders that need to access consolidated pools of liquidity across multiple venues simultaneously. The random nature of the delay makes it impossible to guarantee consistently low quote fade on multiple venues by a SOR. In contrast, quote fade on CX2 decreases 3%, 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 without significant speed differentiation (virtually) replicates the network advantages of consolidated trading.

< Insert Table 4 Here >

A potential limitation of the quote fade metric is that it only considers buyer (seller) initiated trade strings where all liquidity at the national best offer (bid) price was depleted across all venues by the end of the string. With Alpha removed from the Order Protection Rule at the same time as the introduction of the speed bump, a possible outcome is that liquidity demanders seeking to trade large volumes and requiring access to the consolidated liquidity across multiple venues might exclude Alpha from their routing tables completely, resulting in very low trading activity. Given that Alpha was also removed from the Order Protection Rule, it is possible that liquidity at the BBO existed and was simply "traded through". To test this, we examine what happens to quoted liquidity on Alpha 20 milliseconds after the end of the trade string. Liquidity resting on Alpha has three options following trades on other venues: also trade, fade before being executed, or remain (and likely be "traded through"). Figure 7 presents four area charts which display the eventual outcome of liquidity resting on Alpha. Following the relaunch with randomized speed bump, we observe a large increase in the proportion of starting liquidity that fades during trade strings that deplete the available liquidity, on either two or three of the analyzed non-Alpha venues (TSX, Chi-X and CX2). In these cases, a very limited amount of liquidity "stays" on the Alpha order book. The remaining liquidity fades, indicating that whether or not the liquidity demander sent an order to Alpha, they were unlikely to have been able to access that liquidity. A similar, though smaller, increase in liquidity fade is observed when one of the other three trading venues is depleted of liquidity. The fleeting nature of orders on Alpha against trades that deplete liquidity across multiple trading venues is independent of whether liquidity demanders seeking to trade large volumes continue to route their marketable orders to Alpha.

< Insert Figure 7 Here >

5.2. Market share of active and passive trades by broker account

Figures 5 and 6 present further (albeit noisy) evidence that the composition of traders on Alpha changes after the re-launch. Figure 8 presents changes in the proportion of aggressive market orders by broker type. Through conversations with industry participants and regulators we gathered that two domestic banks capture the majority of retail order flow in Canada, while two global investment banks capture a large portion of low-latency trading (HFT) through their direct market access (DMA). The order flow of other banks and broker-dealers contains a mixture of various client types.

We find that the two retail banks' share of aggressive orders increases significantly from 18% to 29%. These retail orders, on average, are unlikely to need to execute quantities larger than the 5 board lot minimum Alpha enforces. Further, consistent with Battalio et al. (2015) the rebates offered by Alpha for aggressive orders will be attractive to the typical retail broker who does not pass this rebate on to their customer.[32] Figure 9 presents changes in passive market share by broker type. Here, the combined passive market share of the two banks with a large HFT presence increases from 19% to 48%. This is consistent with the idea that the main benefit provided by a speed bump requires the (sophisticated) ability to continuously monitor the market in high speed. Non-specialized firms (such as our "other" category) may be unwilling to invest in such sophisticated technology, removing any advantage to posting on new Alpha. These trends suggest that on the "new" Alpha the main active order flow is derived from uninformed retail participants, while liquidity is provided by sophisticated low latency proprietary traders.

< Insert Figure 8 Here >

< Insert Figure 9 Here >

5.3. Smart-Order-Routing and the Information Content of Trades

The existing empirical literature (e.g. Hendershott et al., 2011) calculates realized spreads and adverse selection five minutes after each trade. Carrion (2013) decreases the post-trade interval to one minute, whilst Conrad et al (2015) further decreases the delay to one second. Our approach of constructing trade strings to gauge the information content of each trade is equivalent to a snapshot *twenty milliseconds* after the end of each string of related trades. In Figure 1 we find that the vast majority of price impacts after a trade occur virtually instantaneously, since adverse selection costs often result from trades displacing all available depth at the NBBO, moving the midpoint price.

We utilize the characteristics of each trade string to create proxies for both the information content of the trade and whether the trade originated from a SOR. As described in Section 4.4, trades which deplete the NBBO (either through trades or cancellations) are "informed" (at least about their own trading intentions), while trade strings which retain liquidity at the NBBO subsequent to the trade are less informed. Utilizing a similar definition to van Kervel (2015), we identify trade strings which access more than one venue as originating from a SOR, with those that access only one venue assumed not to have used a SOR.

Figure 10 presents Alpha's trade composition for both depleting and non-depleting trades and those that do or do not use a SOR. Small retail orders are likely to be fully filled on one venue without depleting the NBBO. As such, they would be categorized as non-depleting, non-SOR trades. The proportion of non-depleting, non-SOR trades increases dramatically after the speed bump, from 18% to 46%. Conversely, large institutional trades are likely to use a SOR and to exhaust all liquidity available at the NBBO. Depleting cross-venue sprays experience a dramatic decline, from 46% to 23%. Little movement is

[32] Brolley and Malinova (2013) note that the majority of Canadian retail brokers charge a flat fee, retaining any exchange rebates.

observed in depleting orders which access only one venue (i.e. large retail orders) nor in cross-venue sprays which do not displace an entire NBBO level (i.e. small institutional orders). Given these measures are based on traded (as opposed to quoted) 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.

< Insert Figure 10 Here>

5.4. Trade-based liquidity metrics

To test for statistically significant changes in Alpha's market quality following the relaunch, we utilize equations of the form

$$y_{i,d} = Post_d + Price_{i,d} + Turnover_{i,d} + Volatility_{i,d} + FE_i + e_{i,d} \quad (10)$$

where $y_{i,d}$ is a measure of market quality for stock *i* on day d, $Post_d$ is an indicator variable equal to one for observations after the 21st of September 2015 and zero prior, $Price_{i,d}$ is either the natural logarithm or inverse of the time-weighted NBBO midpoint price, $Turnover_{i,d}$ is the natural logarithm of on-market trade turnover on Alpha, $Volatility_{i,d}$ is the standard deviation of one minute NBBO midpoint returns, FE_i indicates stock fixed effects and $e_{i,d}$ is an error term. Observations are winsorized daily at the 1% level.

Changes in effective spreads, realized spreads and adverse selection costs on Alpha after its relaunch are presented in Table 5. Effective spreads on Alpha increase 0.66 cents, or 1.95 basis points, following the market structure changes. Control variables for price, volume and volatility have the expected directionality and are statistically significant. Old Alpha had an active trading fee of 0.18c per share,[33] with active trades under the revised fee structure instead receiving a rebate of 0.10c per share traded. This resulted in the fee for active orders declining by 0.28c per share traded, slightly smaller than the 0.33c increase observed in the effective half-spread. Consistent with Malinova and Park (2015), we document that liquidity suppliers pass on changes in explicit trading fees, even in markets transitioning to inverted maker-taker pricing schemes. Applying the net-of-fees implicit transaction cost analysis of Malinova and Park (2015), the increase in exchange fees of 0.28c is significantly lower than the 0.58c reduction in adverse selection experienced through access to the speed bump.

< Insert Table 5 Here >

To explicitly examine the benefits to liquidity suppliers of utilizing the speed bump, we follow Conrad et al. (2015) in calculating realized spreads by comparing traded prices with NBBO midpoint quotes at

[33] For shares priced more than C$1.00.

intervals of 1, 5, 10 and 20 seconds after each trade. As shown in Table 5, realized spreads increase 1.24 cents after one second and 1.40 cents after twenty seconds. In relative terms, realized spreads increase 6.29 and 7.52 basis points after one and twenty seconds, respectively. Alpha previously had a passive trading rebate of 0.14c per share traded, for shares priced above $1. For these shares, passive trades under the revised fee structure paid a fee of 0.10c per share during the observation period.[34] The fee for adding liquidity increased 0.24c per share traded, which is substantially smaller than the 0.70c increase in realized half-spread 20 seconds after the trade. Multiplying by trading volumes, net-of-fees profits attributable to liquidity provision on Alpha increase by approximately C$1.48 million per month, suggesting that liquidity suppliers on Alpha benefit from the change. Figure 11 presents average net-of-fees realized half-spreads across each of the major Canadian trading venues. Significant increases in the realized spread earned on Alpha are immediately evident. This is matched by a slight decline on CX2, consistent with a reduction in the aggregate proportion of uninformed order flow arriving at this alternate inverted venue.

< Insert Figure 11 Here >

Table 5 also shows how these changes affected adverse selection. Adverse selection costs measure the directional change in the NBBO midpoint price after a trade. Under Alpha's new market structure, we observe a decline in adverse selection costs of 0.58 cents 1 second after a trade and 0.72 cents after 20 seconds. In relative terms, price movements away from the liquidity supplier decline 4.31 and 5.53 basis points, 1 second and 20 seconds after each trade respectively. Figure 12 presents average adverse selection costs across each of the major Canadian trading venues over the sample period and confirms the decline on Alpha, while other venues, most notably CX2, see their adverse selection costs rise slightly. 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 toxic order flow.

< Insert Figure 12 Here >

Given the simultaneous nature of the introduction of the speed bump and a shift to inverted maker taker pricing, it is difficult to isolate the effects of each of these changes. To provide evidence on how participants are able to utilize the speed bump to their advantage, Table 6 compares the realized spread and adverse selection costs on the two inverted markets in the post period – Alpha and CX2. Any differences observed between the two venues are likely attributable to the existence of the speed bump on Alpha. We further decompose our analysis of liquidity provision on Alpha by trader type, comparing HFT, non-HFT and anonymous participants. Each regression contains all trades on CX2 and trades on

[34] From the 1st of December 2015, passive trading fees will increase to 0.16c per share for post only orders and 0.14c for non- post only orders.

Alpha attributable to only one of each of the 3 classes of traders. The regression specification is of the form:

$$y_{i,d,v} = Alpha_{i,d} + Price_{i,d} + Turnover_{i,d} + Volatility_{i,d} + FE_i + e_{i,d} \quad (11)$$

where $Alpha_{i,d}$ indicates trades on Alpha for the specified class of trader (HFT, non-HFT and Anonymous). All other variables take the same meaning as in Equation 10.

Table 6 presents the findings of this comparison. HFT participants show a clear ability to earn higher realized spreads than the average trader on CX2 by reducing their adverse selection costs. Traders may choose to anonymize their broker ID, in which case it is not possible for us to determine their class. Interestingly, anonymous traders (who may also be HFT) on Alpha show a much greater ability to earn higher spreads and avoid adverse selection. The difference observed for HFT (and anonymous) traders is likely attributable to the existence of the speed bump, which allows fast participants on Alpha to avoid interacting with the (relatively) more toxic multi-market sweep orders. Non-HFT traders on Alpha do not exhibit a statistically significant ability to reduce their adverse selection costs below those of CX2. If anything, non-HFT participants choosing to post liquidity on Alpha earn significantly lower realized spreads than the average trader on CX2. This could be evidence of additional adverse selection imposed on these traders by the HFT participants' rapid removal of liquidity. An alternate explanation is that non-HFT participants capture lower realized spreads than the average on CX2, which contains both HFT and non-HFT liquidity providers.[35]

Overall, our evidence indicates that it is primarily HFT participants who are able to utilize the differential speed advantage provided by the speed bump to increase their realized spread and reduce their adverse selection. They do so by harnessing the information contained within the order flow on other markets, reducing their interaction with costly multi-market sweeps. This strategy is much more successful with the speed bump on Alpha than on a comparable inverted market.

6. Impact on Market Quality for Other Trading Venues

Now that we have established how these changes affected trading on Alpha, we turn to see if these changes affected trading on other venues. Some changes on other venues are likely, particularly since Section 4 establishes that Alpha's systematic order processing delay against marketable orders enables the segmentation of uninformed order flow. In this section, we address the question of whether order flow segmentation increases adverse selection on TSX, Chi-X and CX2, the other large Canadian trading venues. The existing literature 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). We also analyze the impact on consolidated market quality at the national best bid and offer prices.

[35] Unfortunately TRTH does not carry broker IDs for CX2, prohibiting a like-for-like analysis.

We test for changes in market quality metrics across all four venues (Alpha, TSX, Chi-X and CX2) as well as traded liquidity metrics on the three non-Alpha venues (just TSX, Chi-X, and CX2). In each case, we utilize equations of the form

$$y_{i,d} = Post_d + Price_{i,d} + Turnover_{i,d} + Volatility_{i,d} + FE_i + e_{i,d} \quad (12)$$

where $y_{i,d}$ is a measure of consolidated market quality for stock i on day d, $Post_d$ is an indicator variable equal to one for observations after the 21st of September 2015 and zero prior, $Price_{i,d}$ is either the natural logarithm or inverse of the time-weighted NBBO midpoint price, $Turnover_{i,d}$ is the natural logarithm of total on-market trade turnover across either the four venues or three venues, $Volatility_{i,d}$ is the standard deviation of one minute NBBO midpoint returns, FE_i indicates stock fixed effects and $e_{i,d}$ is an error term. Observations are winsorized at the 1% level per day.

< Insert Table 7 Here >

6.1. Impact on Consolidated NBBO Liquidity

Table 7 presents the results of the regression in equation 12 for changes in liquidity metrics across all four trading venues consolidated at the national best bid and offer prices. Quoted spreads increase 0.35 cents in absolute terms and 0.66 basis points in relative terms, while quoted depth increases 0.13; all of these results are statistically significant at the 1% level. These results suggest that the change in Alpha is associated with notable and important decrease in overall market quality as measured by the NBBO quoted spread across all markets, with a smaller but positive increase in depth at the NBBO across all markets.

Why did these changes occur, and are they due to Alpha? Figure 13 indicates that some of the increase in quoted spreads is due to reduced competition from Alpha. Consistent with the increase in Alpha's effective spread reported in Table 6, Figure 13 illustrates that Alpha's proportion of time quoting at the NBBO decreases substantially with the introduction of the speed bump, with a slight increase observed across the other venues. Figure 14 shows that Alpha's share of total NBBO quoted depth increases over time, ultimately increasing by over 70% compared to pre-speed bump levels by the end of our sample. While this may seem like a substantial increase, recall Figure 6 shows that the accessibility of orders at NBBO on Alpha declines sharply with the introduction of the speed bump. Thus, while the minimum size of post only orders on Alpha is effective at increasing quoted liquidity, the speed bump allows this liquidity to fade before being accessed.

In addition, in Table 7, we also look at the resilience of the order book and the necessity of liquidity takers to access limit orders outside of the NBBO. "Take First" represents the proportion of daily trading volume by stock that consumed all depth available on one side of the NBBO. In the post period, this quantity increases by about 2%.

"Take 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 1.6% of volume within trade strings was forced to access the next best price levels to finalize execution. 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.

< Insert Figure 13 Here >

< Insert Figure 14 Here >

6.2. Traded Liquidity Metrics on Other Venues

Alpha’s relative avoidance of informed trades that sweep multiple venues and impose adverse selection costs may increase the toxicity of residual order flow on the other large Canadian trading venues. Table 8 examines changes in effective spreads, realized spreads and adverse selection costs against the NBBO midpoint, volume-weighted amongst trades on TSX, Chi-X and CX2. All control variables have the expected directionality and are statistically significant. After Alpha’s relaunch, effective spreads increase 0.27 cents in absolute terms, or 0.46 basis points in relative terms, both of which are significant at the 1% level. Multiplying by trading volumes, the cost of demanding liquidity increases by $6.12 million per month. Effective spreads increase by a smaller magnitude than quoted spreads, potentially due to the concurrent increase in market depths resulting in competition between liquidity suppliers.

Similar to Conrad et al. (2015), we calculate a range of realized spreads and adverse selection costs from 1 second to 20 seconds after each trade. For brevity, we report results after 20 seconds as our base specification. Realized spreads decline 0.06 cents, signaling a reduction in profits attributable to liquidity provision, but it is only weakly statistically significant (at the 10% level) and not statistically significant if measured in basis points. Multiplying by traded volume, liquidity provider profitability decreases by $1.36 million per month. Although effective spreads widen, the narrowing in realized spreads result from a sharp increase in adverse selection costs of 0.38 cents, or 0.67 basis points. Since adverse price movements after each trade are a proxy for order flow toxicity, we conclude that Alpha’s segmentation of order flow increases residual order flow toxicity and imposes negative liquidity externalities on other trading venues.

< Insert Table 8 Here >

Next, we separately examine traded liquidity metrics on each venue against the national best bid and offer midpoint, to identify where the largest impact of Alpha’s order flow segmentation occurs. Table 9 presents regression results for changes in effective spreads, as well as realized spreads and adverse selection after 20 seconds, separately for TSX, Chi-X and CX2. Effective spreads increase 0.24c on TSX and 0.29c on Chi-X, consistent with the observed widening in quoted spreads at the national best bid and offer prices. No significant change in effective spreads occurred on CX2, potentially due to its relatively low proportion of time quoting at the NBBO. Adverse selection costs increase 0.36c on TSX and Chi-X, and 0.29c on CX2. With Alpha capturing a larger proportion of the uninformed order flow, flow toxicity on all other venues increases at the same time as order book resiliency (i.e. the likelihood of a trade not removing all available depth at the top of the book) declines. Realized spreads decline 0.07c, 0.10c and 0.21c on TSX, Chi-X and CX2 respectively. Alpha's new inverted maker taker pricing and larger quoted depths from minimum post only order sizes enable it to compete with CX2 for active retail order flow, substantially 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 hitting that market, likely as a result of a migration to Alpha encouraged by the (mandated) larger quoted depths.

< Insert Table 9 Here >

6.3. Consolidated Liquidity Metrics by Nominal Stock Price

Alpha's speed bump provides an opportunity for liquidity suppliers to avoid large trades that execute across multiple venues simultaneously, displacing all available depth at the best price level, resulting in immediate adverse selection costs for liquidity suppliers. For stocks with a higher nominal price there are relatively more ticks on the price grid within a given percentage distance from the mid quote and naturally liquidity supply is distributed over more price points than for an otherwise comparable stock with a smaller price. In other words, the quoted depth at each tick is thinner for high price stocks, making the book less resilient and thus increasing the value of the ability to fade, i.e. to not interact with order flow that will move the price. Hence, in the cross-section the introduction of Alpha's speed bump should have a larger impact on the consolidated market quality for higher priced stocks.

We formally test this intuition by grouping stocks into deciles of 24 each and repeating the regression analysis of consolidated market quality metrics conducted in the previous section separately for each group. Separate analysis by deciles further serves as a robustness test, demonstrating that changes in market quality are not driven by a small subset of securities in the sample.

To conserve space, Table 10 reports only the coefficients and t-statistics of the post-launch dummy and omits those for the standard controls. Average stock price ranges from C$2.52 in decile 1 to C$143 in decile 10. Quoted spreads increase the most for high-price stocks, by about 1c for the top 2 deciles, while they move much less for all other deciles.

Adverse selection costs show a somewhat monotonic pattern with increases being concentrated again among high priced stocks. Effective spreads of trades on TSX, Chi-X and CX2 calculated against the prevailing NBBO midpoint also significantly widen across the higher deciles, but do not change by much for the other half of the sample. As a consequence, to a large extent realized spreads do not change across deciles. Higher adverse selection costs faced by liquidity suppliers are being passed on to liquidity demanders in the form of higher quoted and effective spreads, with no net impact on the trading profits attributable to liquidity provision.

These results indicate that our findings are robust across various subsets of stocks, with stocks with a high nominal price experiencing the highest market-wide impact of liquidity provider segmentation.

< Insert Table 10 Here >

7. Conclusion

Counterintuitively, speed bumps represent the most recent innovation in the quest for ever faster trading. With the SEC's recent approval of IEX as the first national securities exchange with a speed bump, studying this new feature of market design has become a question of great practical relevance for market participants and regulators alike. We provide the first examination of the market-wide effects from the introduction of a speed bump by a major North-American exchange.

All speed bump designs currently approved or under consideration in both the U.S. and Canada have one feature in common: the processing delay does not apply to all orders uniformly. The speed bump considered in this study is discriminatory in the sense that it allows traders to "pay" to exempt their limit order entries and cancellations from the speed bump. Thus, our results regarding the quality and fairness of markets are relevant not only to the recent argument surrounding the desirability of speed bumps, but also to any situation in which differential access to speed can be bought. This may include (and to some extent, explain) the prevalence of new deep-sea cables (such as the Hibernia Express cable from New York to London completed in 2015) and microwave towers which now crisscross territories with numerous exchanges (such as the proposed microwave link spanning from London to Frankfurt).

Conrad and Wahal (2016) find that realized spreads in the U.S. have fallen dramatically, from 17 basis points in 2000 to 1.5 basis points in 2015. As a result, market making has become vastly less profitable over this period. At the same time, competition between venues has increased with the fragmentation of liquidity. Against this backdrop, speed bumps have become the next step in the competition for market share between exchanges. Offering some participants a more profitable trading environment relative to competing venues provides a valuable "innovation". With HFTs representing a large and concentrated group of market participants, catering to their needs, (i.e. helping them increase realized spreads), makes sense for an individual venue operator. As we show in this paper, the consequences for other market participants (and market quality overall) need not be positive.

We find that the new speed bump (combined with an inverted maker/taker pricing) is not attractive to all participants, with traded volume on Alpha immediately decreasing. Using novel identification schemes at the trade string level, we show that after the introduction of the speed bump the majority of liquidity consuming trade volume likely shifts to retail traders who typically do not deplete the order book and do not simultaneously interact with multiple trading venues.

Further, we show that the liquidity suppliers on Alpha are predominantly electronic traders who can monitor the market in ultra-high frequency. We show that these low-latency electronic liquidity suppliers are able to harness the information contained within the order flow fragmented across other venues, to avoid trading with large (likely institutional) orders emanating from a SOR which attempt to simultaneously access liquidity on all venues. This results in significantly reduced adverse selection for liquidity suppliers on Alpha, increasing realized spreads and producing substantial economic benefits in an environment where realized spreads are otherwise very close to zero.

The segmentation of predominantly retail order flow to Alpha increases the fraction of informed traders on the remaining venues. Overall, we find significant increases in quoted and effective spreads on the consolidated market of around half a basis point. Consistent with an increase in the fraction of informed traders, this increase in spreads is primarily driven by increases in adverse selection. We also observe negative market wide effects for liquidity demanders, with significant increases in both the fraction of trade strings that consume the entire level of NBBO depth, and the fraction of trade strings that need to walk the book to achieve their desired quantity.

It may at first appear that the reduction in adverse selection costs on Alpha (in a competitive market for liquidity provision) would be offset by tighter spreads, providing an advantage to traders accessing Alpha. However, we do not observe this in the data. Instead, we see "matching" rather than "making" of the NBBO. Alpha quotes at the NBBO only 30% of the time after the introduction of the speed bump, almost half as frequently as the 60% they quoted prior to the introduction of the speed bump. This behavior is consistent with liquidity suppliers' usage of the speed bump to harness information from order flow on other venues. If Alpha was quoting (alone) at the NBBO, the speed bump would lose some (all) of its value: Due to order protection rules, liquidity demanders could route to Alpha first, negating the advantage that comes from observing order flow on other venues prior to Alpha being hit.

A secondary contribution of this study is our methodological innovations. We develop several novel empirical techniques that enable the analysis of cross-market linkages and fairness, which O'Hara (2015) argues are two particularly important issues for both modern regulators and researchers. 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. To this end, we propose techniques to correctly assign trade direction in fragmented markets using trades and quotes emanating from a single data feed, benchmark clock synchronization across multiple trading venues using prohibitions against locked/crossed markets and join trades that likely originate from a SOR spray. From these methods, we develop metrics that empirically validate the conjecture of O'Hara and Ye (2011), that the trade-through prohibition and smart order routing, when combined in fragmented markets, virtually replicate the network advantages of consolidated trading. We show that these market linkages have been circumvented by Alpha's speed bump and its ability to segregate uninformed order flow.

Globally, speed bumps are being discussed as one potential remedy to the "arms race for speed" on the one hand, and unequal access to markets by certain participants on the other. Our results have implications for both the debate surrounding the introduction and desirability of speed bumps, as well as the more general desirability of speed differentials between participants. It seems there are two key components in Alpha's ability to segment retail order flow: the randomized 1-3 millisecond delay (which disrupts smart order routers such as RBC's "THOR" from synchronizing arrival times across venues, breaking down cross-market linkages) and the nature of the application of the speed bump, which provides a guaranteed advantage to traders willing to pay for a "de-minimis" speed advantage.

Of course, different speed bumps may be implemented differently, with differing effects. The IEX speed bump, for example, differs in many ways from that of Alpha, with e.g. a fixed 350 microsecond delay, but also grants preferential treatment to un-displayed pegged orders that automatically re-price against the midpoint of the NBBO. We leave the analysis of the impact of that particular market design change, as well as others, to future research. In general, however, our research suggests 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

Ait-Sahalia, Yacine, and Mehmet Saglam. 2016. "High Frequency Market Making." Working Paper.

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.

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*, Forthcoming

Battalio, Robert, Andriy Shkilko, and Robert Van Ness. "To Pay or be Paid? The Impact of Taker Fees and Order Flow Inducements on Trading Costs in U.S. Options Markets." *Journal of Financial and Quantitative Analysis*, Forthcoming.

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.

Boehmer, Ekkehart, Kingsley Y. L. Fong, and Juan (Julie) Wu. 2015. "International Evidence on Algorithmic Trading." Working Paper.

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

Cardella, Laura, Jia Hao, Ivalina Kalcheva, and Yung-Yu Ma. 2014, "Computerization of the Equity, Foreign Exchange, Derivatives, and Fixed-Income Markets." *Financial Review* 49 (2): 231-243.

Carrion, Allen. "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.

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.

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.

Eggington, Jared F., Bonnie F. Van Ness, and Robert A. Van Ness. 2016. "Quote Stuffing" *Financial Management*, Forthcoming.

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.

Foucault, Thierry, Ohad Kadan, and Eugene Kandel. "Liquidity Cycles and Make/Take Fees in Electronic Markets." *The Journal of Finance* 68(1): 299–341.

Goldstein, Michael A., Pavitra Kumar, and Frank C. Groves. 2014. "Computerized and High Frequency Trading." *Financial Review* 49 (2): 177-202.

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)

Hasbrouck, Joel, and Gideon Saar. 2009. "Technology and Liquidity Provision: The Blurring of Traditional Definitions." *Journal of Financial Markets* 12 (2): 143–72.

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.
Jovanovic, Boyan, and Albert J. Menkveld. 2016. “Middlemen in Limit Order Markets.” Working Paper.
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.
Kyle, Albert S. 1985. “Continuous Auctions and Insider Trading.” *Econometrica* 53 (6): 1315–35.
Larrymore, Norris L., and Albert J. Murphy. 2009. “Internalization and Market Quality: An Investigation.” *Journal of Financial Research* 32 (3): 337–63.
Latza, Torben, Ian W. Marsh, and Richard Payne. 2014. “Fast Aggressive Trading.” Working Paper.
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*.
Malinova, Katya, Andreas Park, and Ryan Riordan. 2016. “Taxing High Frequency Market Making: Who Pays the Bill”. Working Paper.
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. 2016. “The Economics of High-Frequency Trading: Taking Stock.” Working Paper.
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 Solokov, Konstantin. 2016, “Every Cloud Has a Silver Lining: Fast Trading, Microwave Connectivity and Trading Costs.” Working Paper.
van Kervel, Vincent. 2015. “Competition for Order Flow with Fast and Slow Traders.” *Review of Financial Studies*, 23 (7), 2094-2127.
van Kervel, Vincent, and Albert J. Menkveld. 2016. “High-Frequency Trading around Large Institutional Orders.” Working Paper.
Ye, Mao, Chen Yao, and Jiading Gai. 2013. “The Externalities of High Frequency Trading.” Working Paper.
Zhu, Haoxiang. 2014. “Do Dark Pools Harm Price Discovery?” *Review of Financial Studies*, 27 (3): 747-789.

Table 1
Specifications of Major Canadian Lit Trading Venues

This table presents institutional details for each of the major Canadian lit trading venues, including trading fees, order protection rule status, speed bump status and continuous trading hours. Negative trading fees, i.e. rebates, are enclosed in parentheses.

	New Alpha	Old Alpha	TSX[36]	Chi-X	CX2
Taker Fee (above $1)	(0.0010)	0.0018	0.0030 for interlisted 0.0023 for non-interlisted	0.0028	(0.0010)
Maker Fee (above $1)	0.0016 for post only, otherwise 0.0014[37]	(0.0014)	(0.0026) for interlisted (0.0019) for non-interlisted	(0.0024)	0.0014
Speed Bump	1 – 3 milliseconds randomized[38]	No	No	No	No
OPR Protected	No	Yes	Yes	Yes	Yes
Continuous Trading Hours	8:00am – 5:00pm	9:30am – 4:00pm	9:30am – 4:00pm	8:30am – 5:00pm	8:30am – 5:00pm
Average Daily Volume[39]	14,812,413	27,724,226	152,553,868	39,564,726	15,876,833

[36] At the start of each month, TSX updates a list of securities for which the interlisted trading fees apply during that month, available at http://www.tsx.com/resource/en/1130/tsx-symbols-subject-to-applicable-interlisted-trading-fees.csv
[37] New Alpha offers a discounted maker fee of 0.0010 for both post only and non- post only until the 1st of December 2015
[38] Alpha's speed bump applies to all orders except those designated as post-only, which are unable to remove liquidity
[39] Average daily trading volume of on-market lit trades in TSX Composite Index component securities

Table 2
Summary Statistics for Trade Strings

This table presents summary statistics for the trade strings, by the number of venues on which trades were executed in each string. Proportion of Count and Proportion of Traded Value present the percentage of trade strings under each group, by frequency and value of turnover. Frequency is the number of trade strings observed. Means and medians are presented for traded value and traded volume. Proportion of trade strings that deplete all of the available NBBO are presented by frequency and turnover. Mean and median starting depths at NBBO across all venues tend to be higher prior to trade strings that accessed more venues. The 10th, 25th, 50th, 75th and 90th percentiles of the distributions for trade string length, time between the end of the previous trade string to the current trade string, and time between the end of the current trade string and the start of the next trade string are presented in seconds.

Number of Venues on Which Trade String Executed	**1**	**2**	**3**	**4**
Frequency (#)	26,142,010	6,022,150	1,973,604	418,438
Mean Traded Value ($)	6,763	18,752	32,938	57,626
Median Traded Value ($)	3,419	11,142	20,775	36,225
Mean Traded Volume (#)	295	985	2,235	5,596
Median Traded Volume (#)	100	500	1,100	2,200
Mean Starting Depth ($)	44,830	40,839	54,325	80,390
Median Starting Depth ($)	17,140	21,395	32,400	50,096
Proportion of Count (%)	75.65%	17.43%	5.71%	1.21%
Proportion of Traded Value (%)	46.67%	29.81%	17.16%	6.36%
Proportion Depleted by Count (%)	32.24%	56.04%	63.71%	70.34%
Proportion Depleted by Traded Value (%)	48.90%	64.91%	69.90%	74.97%
P10 String Length (sec)	0.000	0.000	0.004	0.008
P25 String Length (sec)	0.000	0.003	0.009	0.011
P50 String Length (sec)	0.000	0.009	0.013	0.017
P75 String Length (sec)	0.000	0.014	0.020	0.024
P90 String Length (sec)	0.000	0.021	0.030	0.036
P10 Space Before (sec)	0.181	0.280	0.506	0.963
P25 Space Before (sec)	1.363	1.958	2.997	4.653
P50 Space Before (sec)	6.668	8.871	11.989	16.145
P75 Space Before (sec)	24.749	31.000	39.629	51.329
P90 Space Before (sec)	71.570	85.410	105.419	131.722
P10 Space After (sec)	0.250	0.139	0.109	0.094
P25 Space After (sec)	1.695	1.180	0.965	0.659
P50 Space After (sec)	7.680	6.619	6.180	5.611
P75 Space After (sec)	27.501	25.373	24.817	24.696
P90 Space After (sec)	77.736	73.157	73.450	76.320

Table 3
Summary Statistics

This table reports univariate descriptive statistics across the 247 TSX Composite Index component securities. The first and second observation periods include the ten weeks prior to and following Alpha Exchange's relaunch on the 21st of September 2015. Quoted spreads and quoted depths are time-weighted and presented at the national best bid and offer prices across Alpha, Chi-X, CX2 and TSX. Depleting (or Take First) trades are those that were part of a string that displaced the entire NBBO depth, where strings are constructed by grouping trades in the same direction separated by less than 30 milliseconds. Take Next trades are those that occur at the next best price behind NBBO within each trade string. Time at NBBO is the proportion of time from 9:30am to 4:00pm that each venue is quoting at the NBB plus the proportion of time quoting at the NBO, divided by two. Depth at NBBO is the proportion of total dollar depth at the NBBO that is quoted by each venue. Metrics are presented separately for Alpha and Chi-X, CX2 and TSX. Effective spreads are calculated against the prevailing NBBO midpoint. Realized spreads are calculated against the NBBO midpoint twenty seconds after the trade. For all depleting trades on each venue, the NBBO quote fade is the proportion of the total visible liquidity at NBB or NBO at the start of the trade string that did not result in trades. Price is the time-weighted NBBO midpoint. Volume is the total quantity of on-market trades. Volatility is the standard deviation of one minute NBBO midpoint returns.

	13 JUL 2015 – 18 SEP 2015			21 SEP 2015 – 27 NOV 2015			Change	T Stat
	Mean	Median	Std. Dev.	Mean	Median	Std. Dev.		
A: Consolidated Liquidity								
NBBO Quoted Spread (cents)	3.58	1.83	5.33	3.67	1.83	5.89	0.09	1.49
NBBO Quoted Depth ($'000s)	81.76	49.64	85.97	92.69	57.33	91.60	10.93	7.08
Depleting/Take First Trades (%)	58.84	59.83	12.31	59.59	60.50	11.73	0.75	2.32
Take Next Trades (%)	11.78	11.19	5.92	12.98	12.42	6.30	1.20	7.89
B: Transaction Costs								
Alpha Effective Spread (cents)	2.86	1.49	4.13	3.48	1.84	5.57	0.62	7.67
Other Effective Spread (cents)	2.92	1.62	4.12	2.94	1.57	4.52	0.02	0.36
Alpha Adverse Selection (cents)	3.09	1.88	4.25	2.17	1.16	3.62	-0.92	-9.40
Other Adverse Selection (cents)	3.65	2.28	4.78	3.66	2.16	5.64	0.02	0.22
C: Percentages at NBBO								
Alpha Time (%)	59.08	59.51	21.95	34.18	31.66	17.59	-24.90	-22.09
Chi-X Time (%)	64.87	68.70	24.13	68.60	71.77	22.29	3.73	5.13
CX2 Time (%)	38.48	35.81	17.75	44.88	44.78	14.55	6.40	8.00
TSX Time (%)	94.23	96.35	6.32	96.14	97.65	4.37	1.91	9.16
Alpha Depth (%)	13.84	13.17	5.72	15.86	14.63	9.49	2.02	2.54
Chi-X Depth (%)	16.61	16.32	6.17	16.89	16.92	5.93	0.28	0.87
CX2 Depth (%)	7.30	6.37	4.15	7.50	6.89	3.46	0.20	0.73
TSX Depth (%)	62.17	62.03	8 .07	59.64	59.46	9.59	-2.53	-4.02
Alpha Quote Fade (%)	14.15	11.1	11.35	60.22	67.16	25.96	46.08	40.92
Chi-X Quote Fade (%)	21.37	19.08	12.81	21.25	19.44	11.74	-0.12	-0.39
CX2 Quote Fade (%)	20.22	18.7	13.83	16.88	15.41	11.03	-3.34	-7.86
TSX Quote Fade (%)	8.82	6.42	7.6	8.93	6.68	7.6	0.11	0.64
D: Control Variables								
Price ($)	31.84	21.21	42.15	29.65	20.26	33.93	-2.19	-2.97
Volume (millions)	0.90	0.41	1.23	0.97	0.46	1.33	0.07	2.19
Volatility (basis points)	11.84	9.69	6.96	11.44	9.61	6.35	-0.40	-1.24

Table 4

Quote Fade at the National Best Bid and Offer Prices Relative to the Pre-Relaunch Period

This table reports coefficient estimates for the determinants of the NBBO quote fade by market for each of Alpha, Chi-X, CX2 and TSX for TSX Composite Index securities, after Alpha's relaunch relative to previous levels using the following specification:

$$Quote\ Fade_{i,d,v} = Post_d + Price_{i,d} + Turnover_{i,d,v} + Volatility_{i,d} + Depth_{i,d} + NBBO\ Depth\ Share_{i,d,v} + FE_i + e_{i,d}$$

where the NBBO quote fade for stock *i* on day *d* at venue *v* is expressed as the sum of an indicator variable for the post-relaunch period, control variables for price, volume, volatility, total NBBO quoted depth, each venue's NBBO depth share, and a stock specific mean. We construct trade strings by joining all trades in the same direction separated by less than 30 milliseconds. 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. The pre-relaunch period runs from the 13th of July 2015 to the 18th of September 2015 and the post-relaunch period from the 21st of September 2015 to the 27th of November 2015. */**/*** indicate statistical significance at the 90%/95%/99% levels, respectively. We double cluster standard errors by stock and date.

	Alpha	Chi-X	CX2	TSX
$Post_d$	43.79	0.00	-2.99	0.16
	(48.76)***	(0.00)	(-7.03)***	(0.86)
$Price_{i,d}$	-0.96	0.50	3.92	1.92
	(-0.39)	(0.30)	(2.93)***	(1.92)*
$Turnover_{i,d,v}$	-0.44	-5.47	-0.71	-2.09
	(-0.73)	(-15.47)***	(-2.15)**	(-9.01)***
$Volatility_{i,d}$	0.17	0.36	0.12	0.16
	(2.03)**	(7.59)***	(2.94)***	(5.09)***
$Depth_{i,d}$	6.19	2.01	-0.30	0.76
	(6.79)***	(3.82)***	(-0.51)	(2.60)***
$Depth\ Share_{i,d,v}$	83.51	17.95	11.74	-4.84
	(14.9)***	(5.99)***	(2.44)**	(-6.41)***
Adjusted R^2	76.2%	63.9%	28.0%	75.4%
# Obs	21,827	21,936	21,682	21,948

Table 5

Trade-Based Liquidity Metrics on Alpha Relative to the Pre-Relaunch Period

This table reports coefficient estimates for measures of transactions costs for TSX Composite Index securities traded on Alpha around the relaunch of the venue using the following specification:

$$y_{i,d} = Post_d + Price_{i,d} + Turnover_{i,d} + Volatility_{i,d} + FE_i + e_{i,d}$$

each liquidity metric for stock *i* on day *d* is expressed as the sum of an indicator variable for the post-relaunch period, and control variables for price, volume and volatility, a stock specific mean and an error term. Effective spreads are measured against the prevailing NBBO midpoint, while realized spreads and adverse selection costs use the reference NBBO midpoint 20 seconds after the trade. Panel A presents metrics in cents whilst panel B presents metrics in basis points. The pre-relaunch period runs from the 13th of July 2015 to the 18th of September 2015 and the post-relaunch period from the 21st of September 2015 to the 27th of November 2015. */**/*** indicate statistical significance at the 90%/95%/99% levels, respectively. We double cluster standard errors by stock and date.

Panel A: In Cents

	Effective Spread	Realized Spread 1 second	Realized Spread 20 seconds	Adverse Selection 1 second	Adverse Selection 20 seconds
$Post_d$	0.66	1.24	1.40	-0.58	-0.72
	(6.34)***	(10.57)***	(10.99)***	(-9.69)***	(-8.14)***
$Price_{i,d}$	2.11	-0.01	-0.97	2.57	3.59
	(2.59)***	(-0.02)	(-1.39)	(3.85)***	(3.82)***
$Turnover_{i,d}$	-0.20	-0.09	-0.07	-0.03	-0.06
	(-3.88)***	(-1.55)	(-1.00)	(-0.51)	(-0.61)
$Volatility_{i,d}$	0.11	0.03	-0.01	0.07	0.11
	(5.48)***	(2.00)**	(-0.58)	(7.99)***	(8.58)***
Adjusted R^2	6.1%	9.8%	13.1%	8.4%	8.7%
# Obs	21,870	21,870	21,870	21,870	21,870

Panel B: In Basis Points

	Effective Spread	Realized Spread 1 second	Realized Spread 20 seconds	Adverse Selection 1 second	Adverse Selection 20 seconds
$Post_d$	1.95	6.29	7.52	-4.31	-5.53
	(10.71)***	(14.98)***	(15.14)***	(-11.04)***	(-11.70)***
$Price_{i,d}$	89.43	72.04	69.87	9.79	11.04
	(42.89)***	(8.78)***	(7.67)***	(1.36)	(1.38)
$Turnover_{i,d}$	-0.61	0.15	0.40	-0.80	-1.03
	(-4.91)***	(0.85)	(1.87)*	(-5.26)***	(-5.28)***
$Volatility_{i,d}$	0.34	-0.26	-0.47	0.60	0.83
	(12.71)***	(-4.86)***	(-7.59)***	(14.56)***	(15.06)***
Adjusted R^2	33.2%	27.5%	27.7%	19.7%	22.5%
# Obs	21,870	21,870	21,870	21,870	21,870

Table 6
Realized Spreads and Price Impacts of Various Participant Types on Alpha Relative to CX2

This table reports coefficient estimates for measures of transactions costs for TSX Composite Index securities on CX2 and for various broker accounts supplying liquidity on Alpha after the latter's relaunch using the following specification:

$$y_{i,d} = Alpha_{i,d} + Price_{i,d} + Turnover_{i,d} + Volatility_{i,d} + FE_i + e_{i,d}$$

each liquidity metric for stock *i* on day *d* is expressed as the sum of an indicator variable for observations on Alpha, and control variables for price, volume and volatility, a stock specific mean and an error term. Realized spreads and adverse selection costs use the reference NBBO midpoint 100 milliseconds after the trade. For three groups of broker accounts that supply liquidity on Alpha, these metrics are compared with CX2, an alternative trading venue that offers a similar inverted fee structure. HFT DMA consists of two global banks that offer direct market access services to proprietary traders. Anonymous consists of all participants that chose not to broadcast their broker number, forgoing the opportunity to participate in broker preferencing. All other brokers are grouped as other. The post-relaunch period runs from the 21st of September 2015 to the 27th of November 2015. */**/*** indicate statistical significance at the 90%/95%/99% levels, respectively. We double cluster standard errors by stock and date.

	Adverse Selection			Realized Spread		
	HFT DMA	Anonymous	Other	HFT DMA	Anonymous	Other
$Alpha_v$	-0.16	-0.50	-0.04	0.27	0.45	-0.23
	(-5.62)***	(-8.90)***	(-0.85)	(7.09)***	(3.42)***	(-4.73)***
$Price_{i,d,v}$	0.93	1.17	1.09	0.55	0.11	0.20
	(6.15)***	(9.86)***	(9.67)***	(3.15)***	(0.47)	(1.29)
$Turnover_{i,d,v}$	0.07	-0.12	-0.05	-0.12	0.20	0.09
	(3.79)***	(-8.49)***	(-2.31)**	(-4.83)***	(7.24)***	(4.97)***
$Volatility_{i,d,v}$	0.04	0.05	0.05	0.02	0.00	0.01
	(9.25)***	(13.36)***	(10.82)***	(4.74)***	(0.00)	(1.13)
Adjusted R^2	46.2%	43.1%	49.5%	15.7%	5.9%	4.9%
# Obs	21,235	18,124	21,284	21,235	18,124	21,284

Table 7

Consolidated Liquidity Metrics at NBBO Relative to the Pre-Relaunch Period

This table reports coefficient estimates for measures of liquidity for TSX Composite Index securities traded on TSX, Alpha, Chi-X and CX2 around the relaunch of the venue using the following specification:

$$y_{i,d} = Post_d + Price_{i,d} + Turnover_{i,d} + Volatility_{i,d} + FE_i + e_{i,d}$$

each liquidity metric for stock *i* on day *d* is expressed as the sum of an indicator variable for the post-relaunch period, and control variables for price, volume and volatility, a stock specific mean and an error term. Quoted spreads and quoted depths are time-weighted and presented at the national best bid and offer prices across Alpha, Chi-X, CX2 and TSX. "Take First" 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. "Take Next" is the proportion of trading volume that occurs at any price behind NBBO. The pre-relaunch period runs from the 13th of July 2015 to the 18th of September 2015 and the post-relaunch period from the 21st of September 2015 to the 27th of November 2015. */**/*** indicate statistical significance at the 90%/95%/99% levels, respectively. We double cluster standard errors by stock and date.

	Quoted Spread		Quoted Depth	Take First	Take Next
	Cents	Basis Points			
$Post_d$	0.35	0.66	0.13	1.93	1.60
	(4.05)***	(3.90)***	(8.98)***	(6.45)***	(10.70)***
$Price_{i,d}$	3.15	85.56	0.33	11.51	5.58
	(2.99)***	(32.01)***	(4.58)***	(8.38)***	(9.85)***
$Turnover_{i,d}$	-0.96	-3.17	0.24	-4.74	-0.38
	(-9.57)***	(-14.13)***	(16.31)***	(-16.32)***	(-2.53)**
$Volatility_{i,d}$	0.13	0.43	-0.03	0.83	0.32
	(6.74)***	(14.29)***	(-17.61)***	(17.04)***	(16.08)***
Adjusted R^2	10.6%	47.5%	32.4%	11.6%	8.8%
# Obs	21,948	21,948	21,948	21,948	21,948

Table 8
Consolidated Liquidity Metrics across Other Venues Relative to the Pre-Relaunch Period

This table reports coefficient estimates for measures of transactions costs for TSX Composite Index securities traded across all venues apart from Alpha (i.e. TSX, Chi-X and CX2) around the relaunch of the venue using the following specification:

$$y_{i,d} = Post_d + Price_{i,d} + Turnover_{i,d} + Volatility_{i,d} + FE_i + e_{i,d}$$

each liquidity metric for stock *i* on day *d* is expressed as the sum of an indicator variable for the post-relaunch period, and control variables for price, volume and volatility, a stock specific mean and an error term. Effective spreads are measured against the prevailing NBBO midpoint, while realized spreads and adverse selection costs use the reference NBBO midpoint 20 seconds after the trade. The pre-relaunch period runs from the 13th of July 2015 to the 18th of September 2015 and the post-relaunch period from the 21st of September 2015 to the 27th of November 2015. */**/*** indicate statistical significance at the 90%/95%/99% levels, respectively. We double cluster standard errors by stock and date.

	Effective Spread		Realized Spread		Adverse Selection	
	Cents	Basis Points	Cents	Basis Points	Cents	Basis Points
$Post_d$	0.27	0.46	-0.06	-0.19	0.38	0.67
	(3.83)***	(4.54)***	(-1.83)*	(-1.21)	(3.82)***	(4.48)***
$Price_{i,d}$	2.63	89.47	-1.06	37.12	3.63	49.44
	(3.13)***	(39.11)***	(-6.24)***	(5.78)***	(3.33)***	(11.86)***
$Turnover_{i,d}$	-0.59	-1.58	0.16	1.49	-0.78	-3.22
	(-7.68)***	(-9.11)***	(4.72)***	(8.25)***	(-7.66)***	(-17.82)***
$Volatility_{i,d}$	0.11	0.36	-0.10	-0.66	0.22	1.06
	(6.43)***	(14.06)***	(-11.27)***	(-17.78)***	(6.99)***	(25.39)***
Adjusted R^2	8.9%	49.3%	7.8%	21.5%	11.8%	40.5%
# Obs	21,948	21,948	21,948	21,948	21,948	21,948

Table 9
Per-Venue Liquidity Metrics on Other Venues Relative to the Pre-Relaunch Period

This table reports coefficient estimates for measures of transactions costs (measured in cents) for TSX Composite Index securities by venue traded, across TSX, Chi-X and CX2 around the relaunch of the venue using the following specification:

$$y_{i,d,v} = Post_d + Price_{i,d,v} + Turnover_{i,d,v} + Volatility_{i,d,v} + FE_i + e_{i,d}$$

each liquidity metric for stock *i* on day *d* is expressed as the sum of an indicator variable for the post-relaunch period, and control variables for price, volume and volatility, a stock specific mean and an error term. Effective spreads are measured against the prevailing NBBO midpoint, while realized spreads and adverse selection costs use the reference NBBO midpoint 20 seconds after the trade. The pre-relaunch period runs from the 13th of July 2015 to the 18th of September 2015 and the post-relaunch period from the 21st of September 2015 to the 27th of November 2015. */**/*** indicate statistical significance at the 90%/95%/99% levels, respectively. We double cluster standard errors by stock and date.

	Effective Spread			Realized Spread			Adverse Selection		
	TSX	Chi-X	CX2	TSX	Chi-X	CX2	TSX	Chi-X	CX2
$Post_d$	0.24	0.29	0.13	-0.07	-0.10	-0.21	0.36	0.36	0.29
	(3.59)***	(3.50)***	(1.64)	(-2.16)**	(-1.96)**	(-3.14)***	(3.80)***	(3.94)***	(4.46)***
$Price_{i,d}$	2.61	2.80	2.94	-1.27	-0.53	0.50	3.79	3.64	2.59
	(3.18)***	(3.26)***	(3.19)***	(-5.90)***	(-1.96)**	(1.24)	(3.41)***	(3.73)***	(5.87)***
$Turnover_{i,d,v}$	-0.52	-0.49	-0.49	0.25	-0.06	-0.28	-0.81	-0.44	-0.20
	(-7.39)***	(-8.54)***	(-7.25)***	(7.26)***	(-1.87)*	(-3.82)***	(-8.40)***	(-6.48)***	(-3.22)***
$Volatility_{i,d}$	0.11	0.11	0.11	-0.11	-0.06	-0.01	0.22	0.17	0.11
	(6.18)***	(6.81)***	(7.49)***	(-12.05)***	(-8.82)***	(-0.77)	(7.19)***	(8.26)***	(9.90)***
Adjusted R^2	8.7%	6.2%	5.1%	7.9%	1.6%	1.6%	11.9%	6.9%	4.9%
# Obs	21,948	21,939	21,818	21,948	21,939	21,818	21,948	21,939	21,818

Table 10
Consolidated Liquidity Metrics on Other Venues Relative to Pre-Relaunch Period
By Average Share Price

This table reports coefficient estimates and t-statistics for the post-relaunch indicator variable for changes in consolidated market quality metrics (measures in cents) across deciles of TSX Composite Index securities. Deciles are constructed from each stock's average time-weighted midpoint price in the ten weeks prior to Alpha's relaunch. The regression specification used is as follows:

$$y_{i,d} = Post_d + Price_{i,d} + Turnover_{i,d} + Volatility_{i,d} + FE_i + e_{i,d}$$

each liquidity metric for stock *i* on day *d* is expressed as the sum of an indicator variable for the post-relaunch period, and control variables for price, volume and volatility, a stock specific mean and an error term. Quoted spreads are consolidated across TSX, Alpha, Chi-X and CX2. Effective spreads, realized spreads and adverse selection are consolidated across TSX, Chi-X and CX2. The pre-relaunch period runs from the 13th of July 2015 to the 18th of September 2015 and the post-relaunch period from the 21st of September 2015 to the 27th of November 2015. */**/*** indicate statistical significance at the 90%/95%/99% levels, respectively. We double cluster standard errors by stock and date.

Decile	Average Price	Quoted Spread	Effective Spread	Realized Spread	Adverse Selection
10	143.51	**1.27**	**1.06**	**-0.33**	**1.91**
		(2.40)**	(2.15)**	(-2.06)**	(2.57)**
9	51.01	**0.71**	**0.49**	**0.00**	**0.47**
		(2.95)***	(2.92)***	(-0.04)	(4.76)***
8	38.56	**0.17**	**0.10**	**0.02**	**0.13**
		(1.86)*	(1.48)	(0.26)	(2.77)***
7	30.86	**0.21**	**0.13**	**-0.02**	**0.16**
		(2.69)***	(2.25)**	(-0.32)	(1.8)*
6	24.12	**0.28**	**0.17**	**-0.06**	**0.23**
		(2.84)***	(2.27)**	(-0.93)	(3.8)***
5	18.30	**0.18**	**0.17**	**0.02**	**0.14**
		(2.15)**	(2.44)**	(0.41)	(2.86)***
4	13.47	**0.10**	**0.06**	**0.00**	**0.06**
		(2.2)**	(2.1)**	(-0.01)	(1.65)*
3	9.87	**0.01**	**0.01**	**0.03**	**-0.02**
		(0.13)	(0.45)	(1.05)	(-0.5)
2	6.22	**-0.03**	**0.00**	**0.00**	**0.01**
		(-1.79)*	(0.33)	(-0.03)	(0.23)
1	2.52	**0.00**	**-0.01**	**-0.06**	**0.05**
		(-1.27)	(-1.04)	(-3.89)***	(2.63)***

Figure 1
Realized Spread Within One Minute by Number of Venues Accessed

This figure presents the realized spread associated with trades over 100 milliseconds, 1, 5, 10, 20 and 60 seconds. These are split into trade strings which access only a single venue, and trade strings which access multiple venues. Trade strings are defined as series of trades which execute within 30 milliseconds of each other. This means trade strings will be separated by at least 30 milliseconds of no-trade. A full explanation of the construction of trade strings is available in Appendix C.

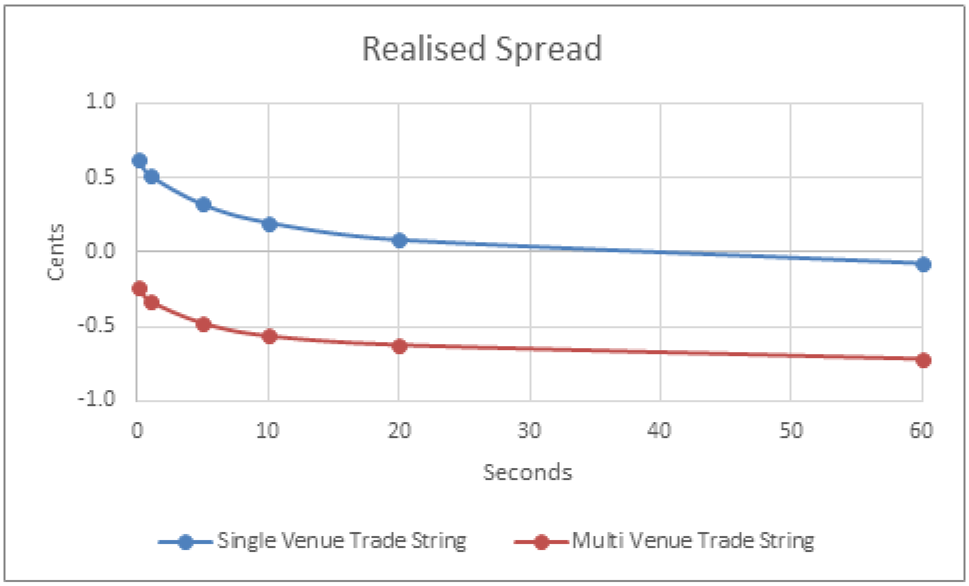

Figure 2
On-Market Volume Share per Venue

This figure presents each venue's market share of total daily on-market lit trading volume in TSX Composite Index securities. We present market share of volume, rather than dollar turnover, since trading fees in Canada are a fixed price per share instead of a fixed percentage of dollar value traded.

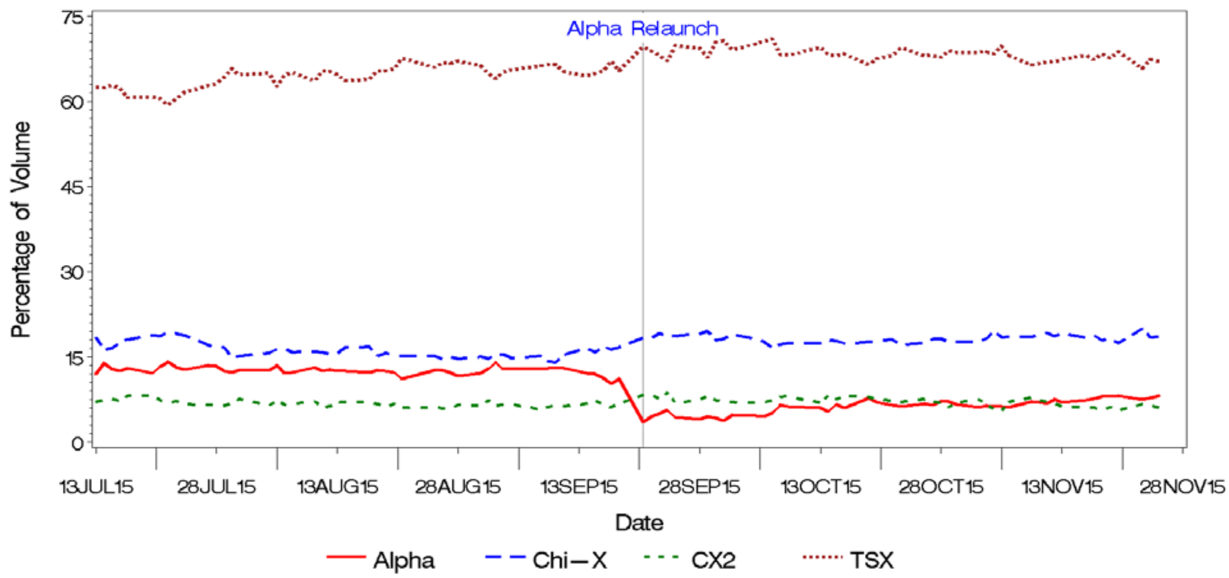

Figure 3

Duration of locked/crossed markets

This figure presents a histogram of the duration of periods of locked/crossed markets using potentially asynchronous time-stamps across venues. Time stamps in TRTH are reported to the nearest millisecond, thus a locked period of 0 milliseconds means that quotes across markets changed in a consistent way within the same 1 millisecond period and so on. The 95th percentile of locked/crossed durations is marked with a dashed line at 30 milliseconds.

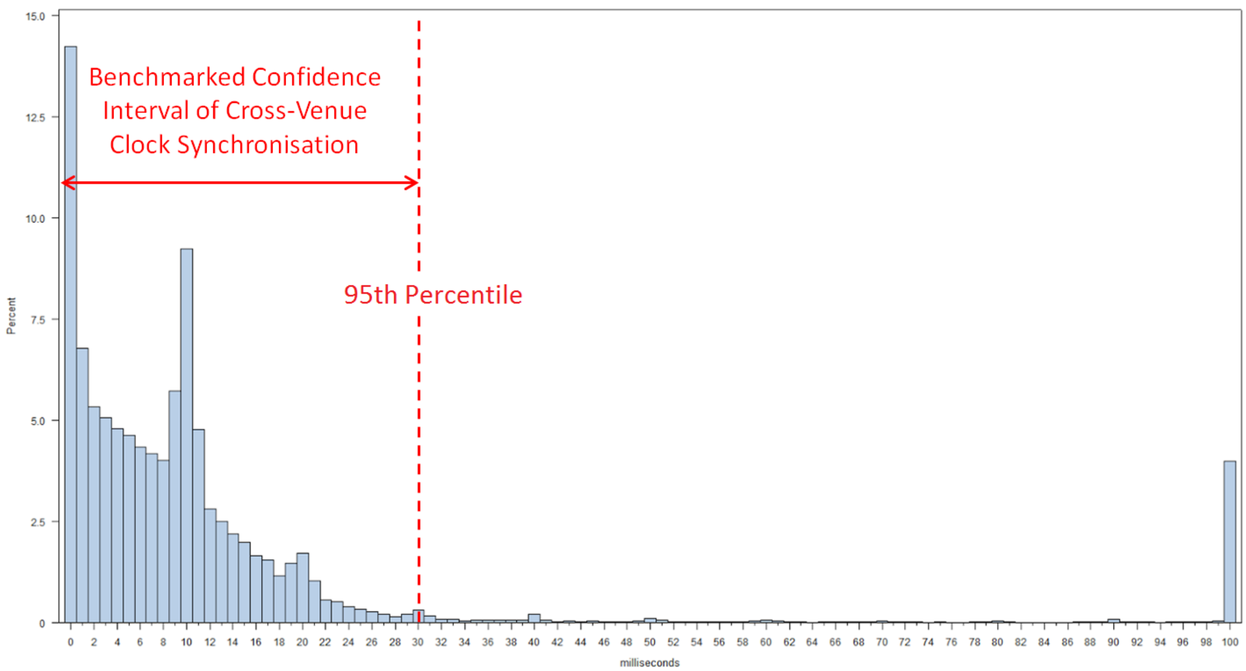

Figure 4

Example of Trade String Construction for Quote Fade

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

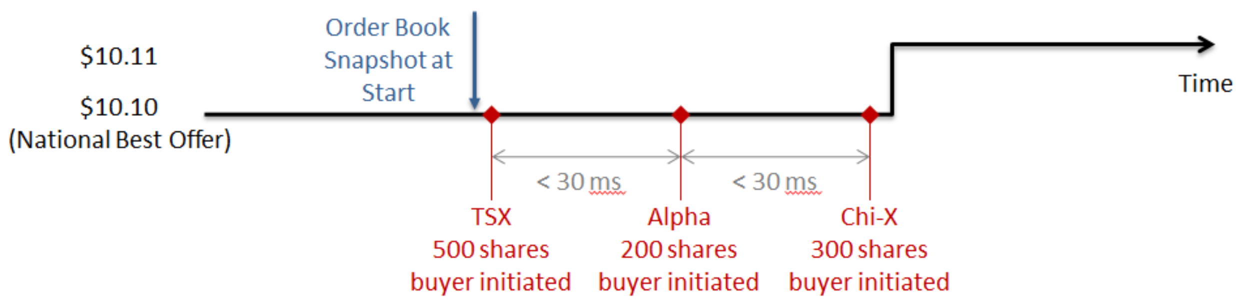

Figure 5
Fraction of Depleting Trade Strings per Stock

This figure presents the fraction of trades by stock which consume (i.e. deplete) an entire level of liquidity across the duration of our sample. Stocks are separated by their relative minimum tick size, which is the average of the daily minimum tick size divided by the time-weighted quoted mid-point. The horizontal axis is in log scale due to the significant variation in the relative minimum tick size.

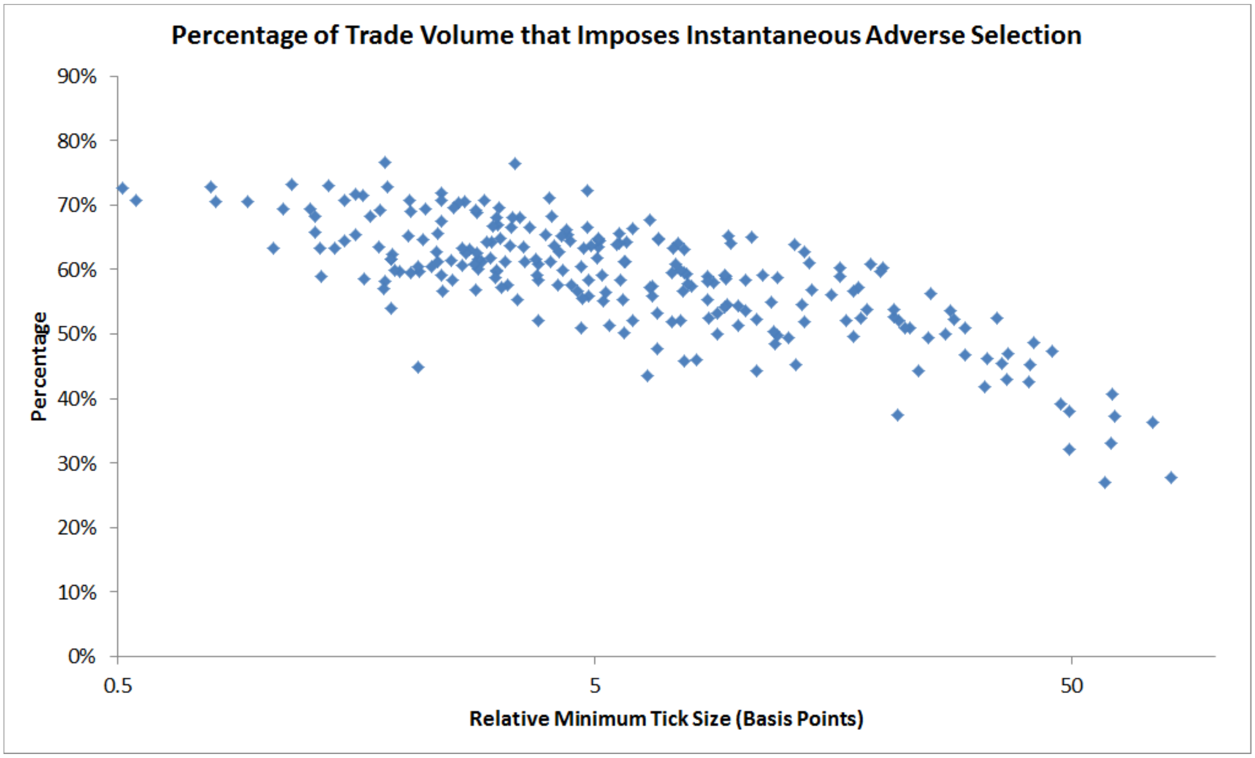

Figure 6
NBBO Quote Fade per Venue

This figure presents the aggregate quote fade within each market for trade strings that deplete an entire level of quoted depth at the NBBO. It measures the proportion of visible liquidity that active traders were unable to access. This metric is restricted by a lower bound of zero per trade string.

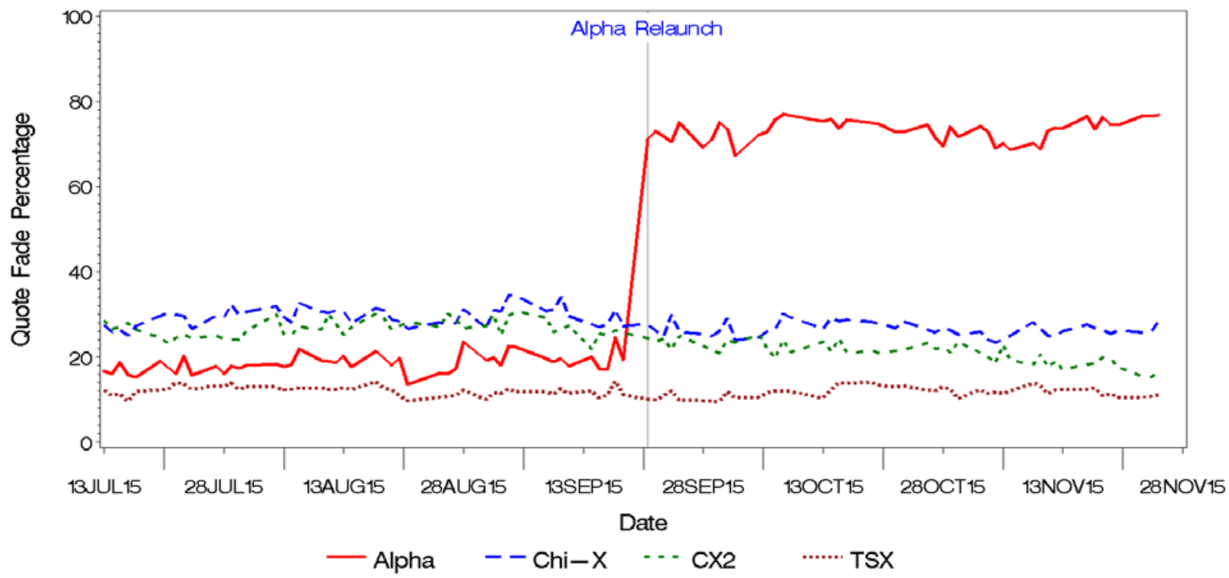

Figure 7
Percentage of Displayed Liquidity on Alpha that Stayed, Traded and Faded

This quadrant of figures presents the proportion of displayed liquidity on Alpha at the national best offer (bid) price at the start of buyer (seller) initiated trade strings that resulted in trades, stayed in the order book, or were faded from the order book conditional on the number of venues accessed during each trade string. Trade strings are constructed by grouping all trades separated by less than 30 milliseconds. The starting liquidity snapshot is taken immediately prior to the first trade in the string and the ending liquidity snapshot is taken 20 milliseconds after the last trade in the string.

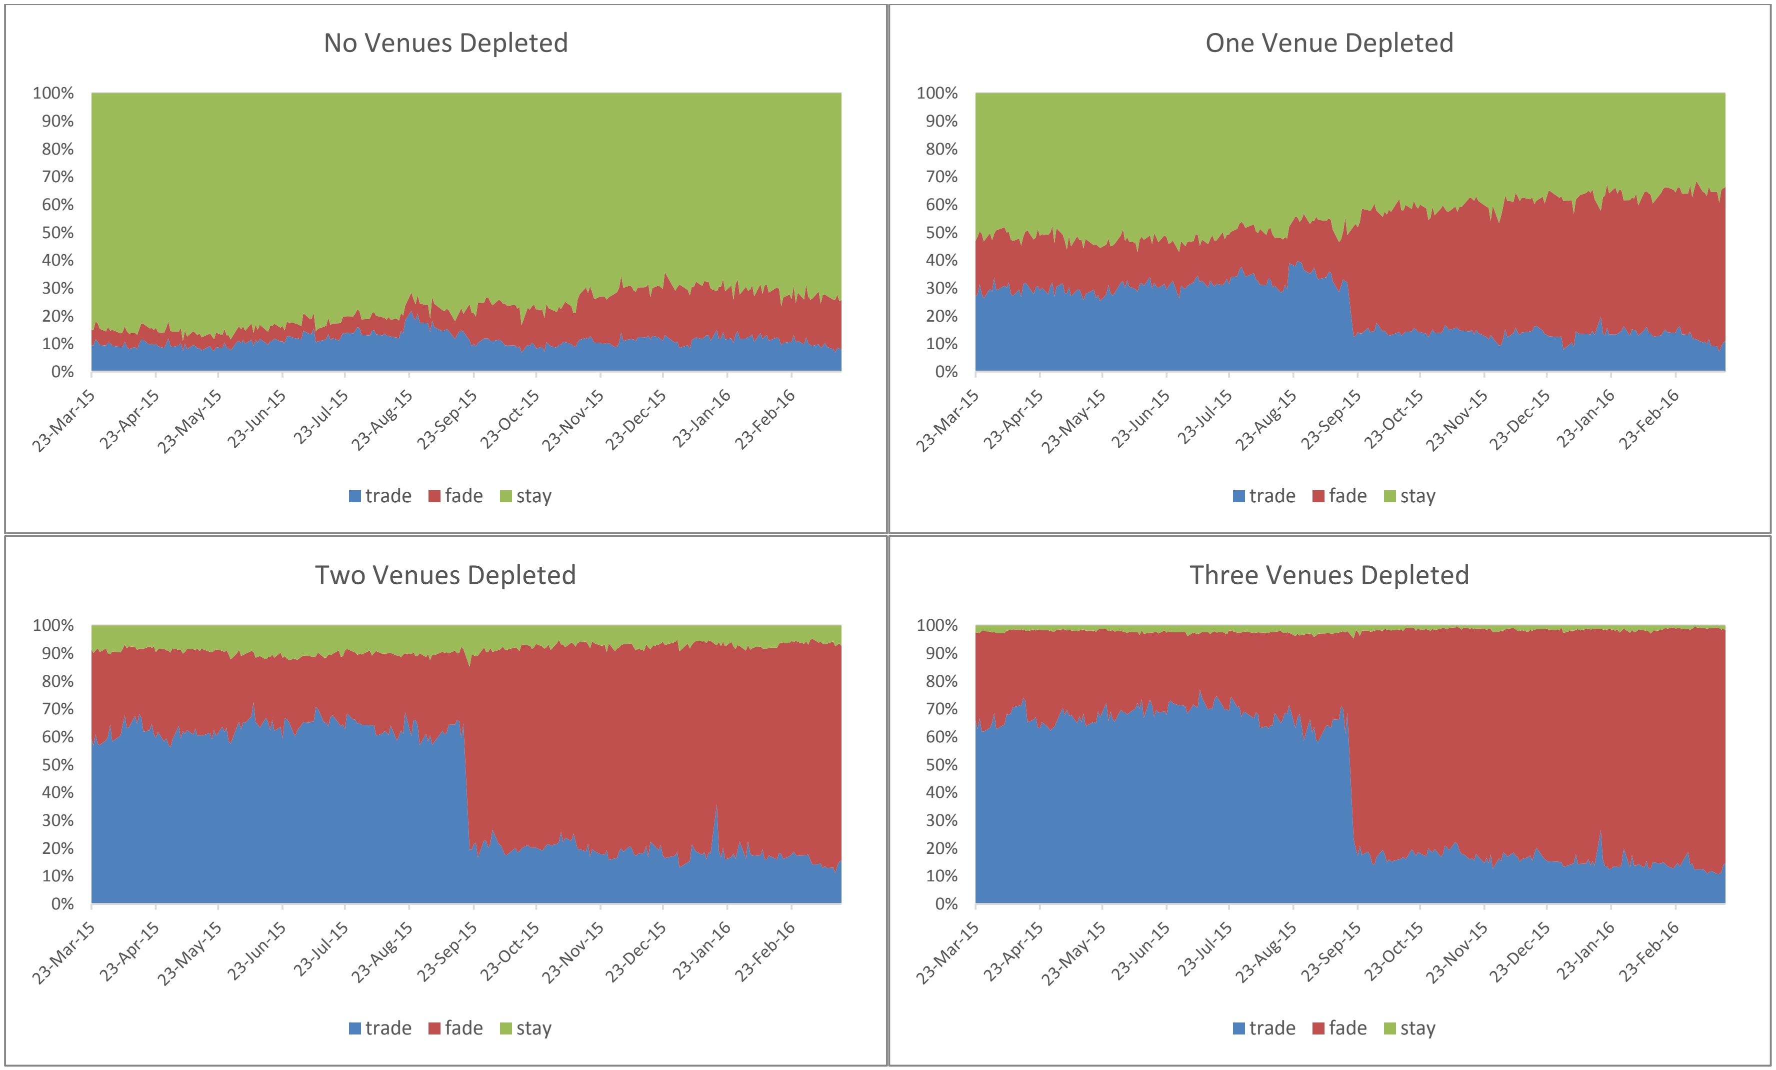

Figure 8
Active Market Share by Broker Type on Alpha

This figure presents Alpha's market share of active trade turnover by broker type. Retail consists of two local Canadian banks that are known to constitute a large proportion of retail broking activity.

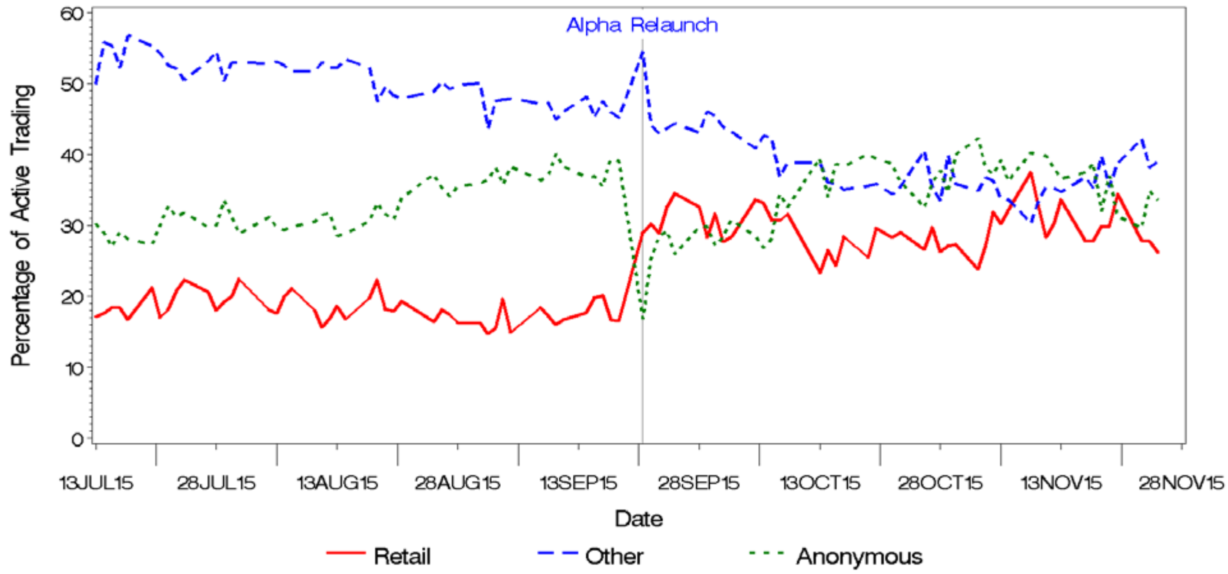

Figure 9
Passive Market Share by Broker Type on Alpha

This figure presents Alpha's market share of passive trade turnover by broker type. HFT DMA consists of two global investment banks that offer direct market access services to proprietary trading firms that act as low latency market makers.

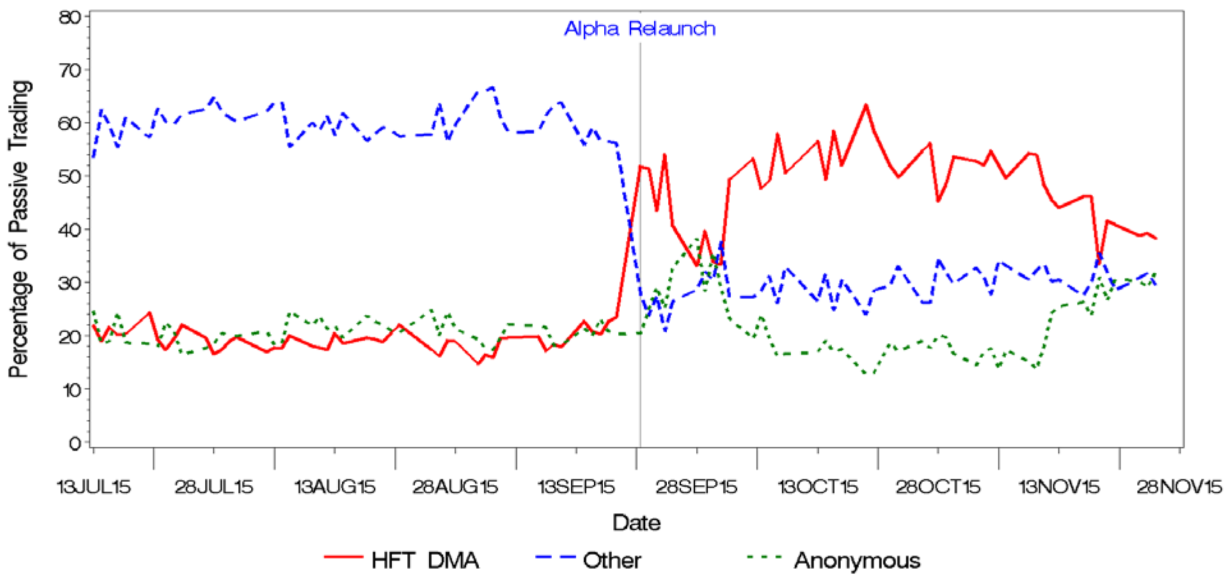

Figure 10
Trading Volume Composition by Trade String Type on Alpha

This figure presents a decomposition of Alpha's on-market turnover by trade string type. We distinguish between trade strings that deplete the top level of quoted depth at the NBBO vs. those that do not. Smart order router (SOR) strings are those that execute on multiple venues.

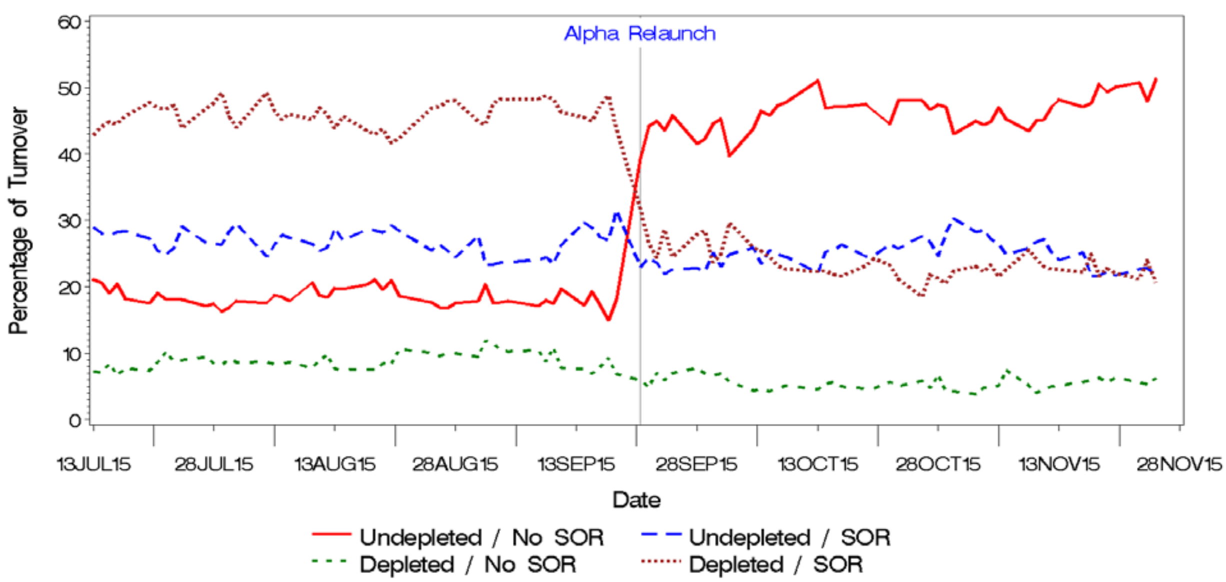

Figure 11
Net-of-Fees Realized Half-Spreads per Venue

This figure presents the volume-weighted average realized spreads of trades against the midpoint of the national best bid and offer prices twenty second after the trade, adjusted by the venue's passive trading fee or rebate. The net-of-fees realized spread proxies for the liquidity supplier's trading profits.

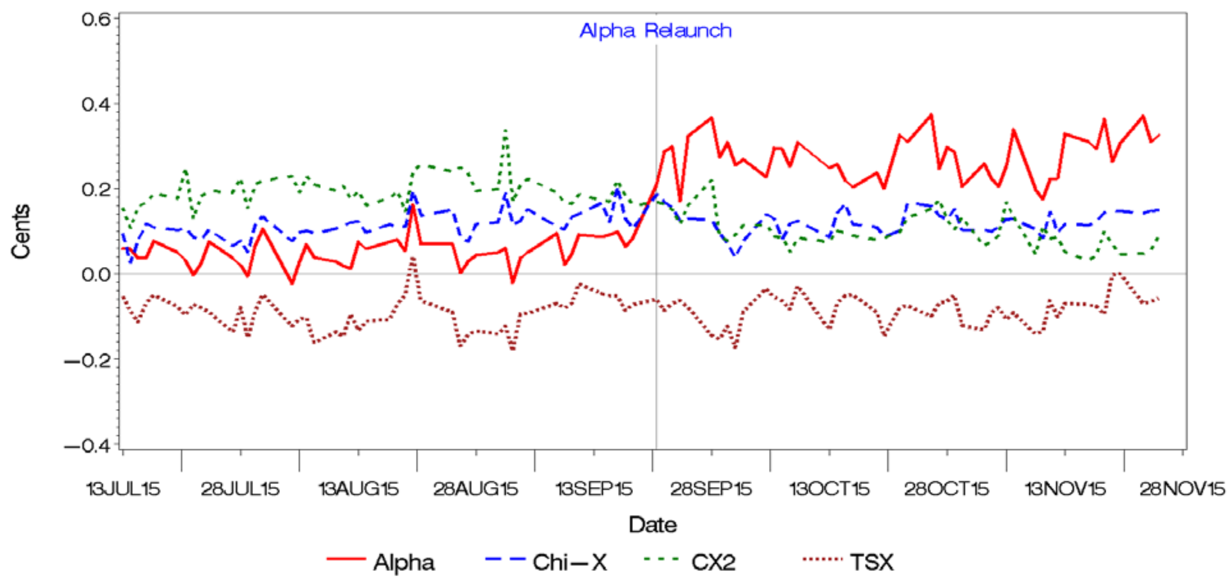

Figure 12
Adverse Selection Costs per Venue

This figure presents the volume-weighted average adverse selection costs of trades, measured as the directional change in midpoint of the national best bid and offer prices from immediately before the trade occurred to twenty seconds after the trade. This metric gauges the price impact as a result of the trade.

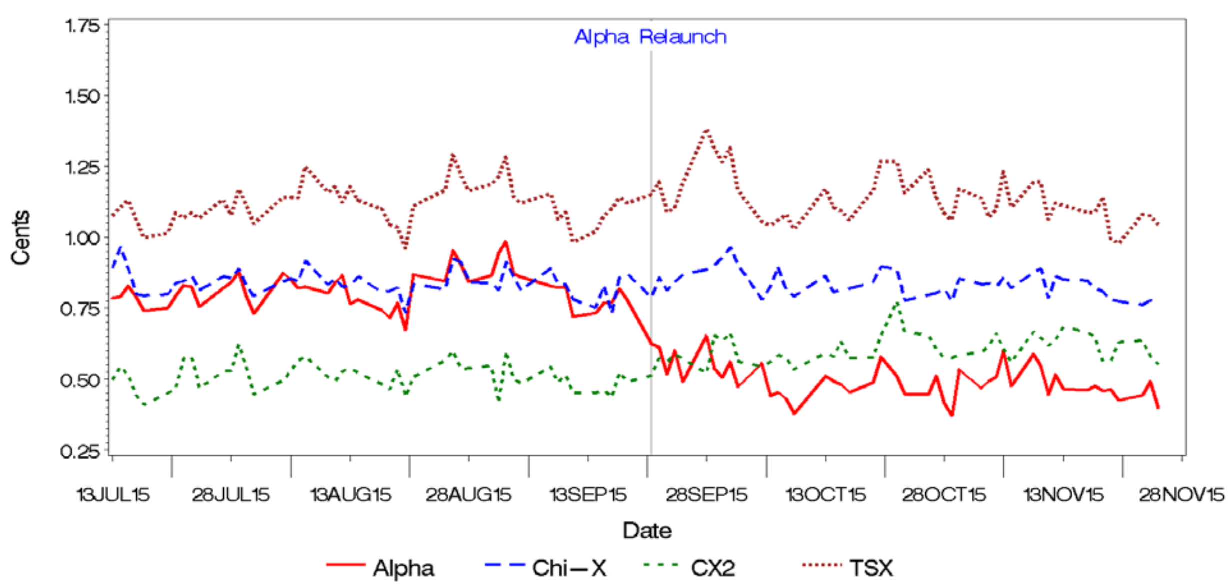

Figure 13
Percentage of Time Quoting at NBBO per Venue

This figure presents the average proportion of time each venue was quoting at the national best bid and offer prices, equal-weighted per security. A large decrease in the proportion of time the relaunched Alpha venue posts competitive quotes occurs.

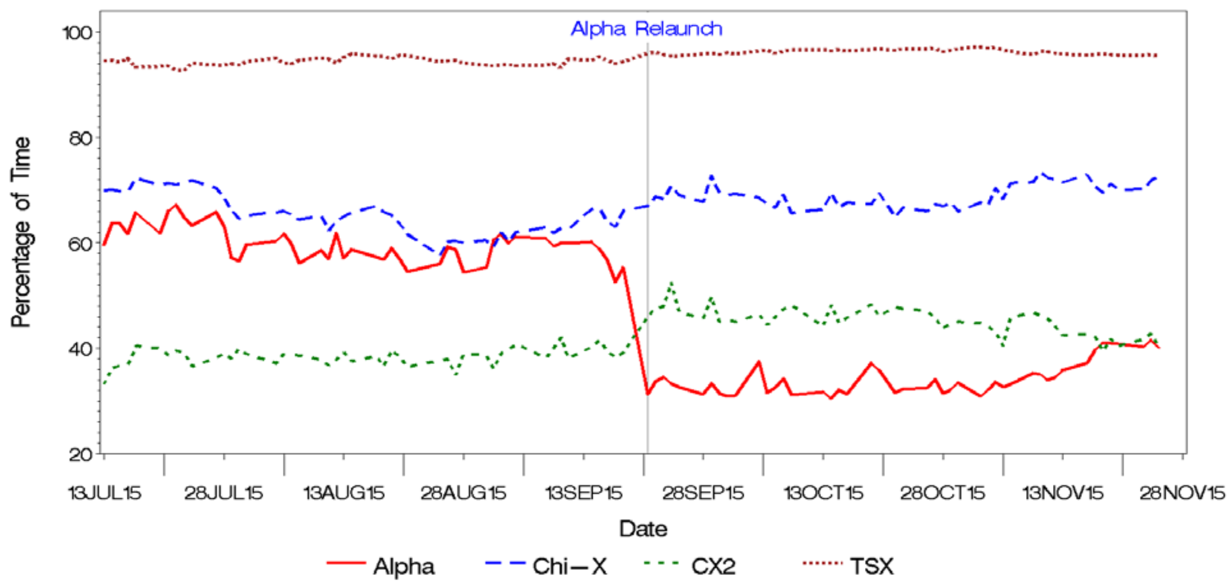

Figure 14
Percentage of Total Depth Quoted at NBBO per Venue

This figure presents the cumulative percentage change in the proportion of total dollar depth each venue quoted at the national best bid and offer prices, aggregated across all securities, over the sample period. Proportions are normalized to 100% as of the event date, 21st September 2015.