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

Haoming Chen
University of New South Wales

Sean Foley
University of Sydney

Michael Goldstein
Babson College

Thomas Ruf[1]
University of New South Wales

November 18, 2017

Abstract

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

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

[1] Contact Author: t.ruf@unsw.edu.au. We thank the Securities Industry Research Centre of Asia-Pacific (SIRCA) for data. Haoming Chen thanks the Capital Markets Cooperative Research Centre (CMCRC) for funding, and was a visiting researcher at the Ontario Securities Commission (OSC) for part of the time spent on this study. The opinions expressed in this study are our own and do not necessarily reflect the views of the OSC or the CMCRC. We thank Reena Aggarwal, Mike Aitken, Angelo Aspris, Robert Battalio, Alejandro Bernales, Bryan Blake, Michael Brolley, Torstein Braaten, Eric Budish, Sabrina Buti (EFA discussant), Doug Clark, Thierry Foucault, Austin Gerig, Carole Gresse, Edwin Hu, Charles Jones (NBER Discussant), Maureen Jensen, Dan Kessous, Amy Kwan, Katya Malinova, Sophie Moinas, Bernt Odegaard, Andreas Park (NFA Discussant), Richard Phillip, Talis Putnins, Dagfinn Rime, Ryan Riordan, Ioanid Rosu, Jos Schmitt, Andriy Shkilko, Elvira Sojli, Vincent van Kervel, Rene Wells, Pradeep Yadav, Yesha Yadav, Lei Yu (CICF discussant), Bart Yueshen, Haoxiang Zhu, and Marius Zoican as well as seminar participants at the EFA 2017, CICF 2017, Ontario Securities Commission, Financial Conduct Authority UK, ITG MatchNow, Chi-X Canada, Aequitas Neo, the NFA 2016, the National Bureau of Economic Research (NBER) Market Microstructure working group, the Asian Finance Association, BI Norwegian Business School, Universidad de Chile and Catholica University, Chile, for their thoughtful feedback. We are also indebted to the staff at the OSC for sharing their knowledge on Canadian equity market structure. The internet appendix that accompanies this paper may be found at https://goo.gl/Sg5qz2.

1 Introduction

The desire for speed has led technology suppliers to provide faster access to traders, leading to a perpetual arms race. Over time, the marginal increase in speed has become ever smaller, but what remains constant is that speed provides an advantage to those who possess it over those who do not. Faster traders capture most of the profits from liquidity provision and impose adverse selection costs on relatively slower counterparts.[2] Exchange operators have responded with innovations such as co-location, microwave towers and latency upgrades as communications technology approaches the speed of light.[3]

Given these speeds, it has become easier to *slow* participants in order to speed others up. Several competing trading venues have proposed a number of speed bumps with a variety of features. Such delays may apply to all orders equally (symmetrically), or, as is the case for IEX, NYSE American, and CHX in the United States and TSX Alpha in Canada, almost all incoming messages are slowed – *except for* a specific subset of resting limit orders (asymmetric) which pay additional fees.[4] Further, the duration of such speed bumps can be either deterministic or random. Brolley and Cimon (2017) predict increased trading costs on non-speed bump venues for asymmetric speed bump. Baldauf and Mollner (2015) suggest that asymmetric, "non-cancellation" delays can be effective against sniping, but also point out that randomization of delays impedes the efficacy of smart order routers in accessing consolidated liquidity across venues. It is unclear whether these new mechanisms affect overall liquidity provision and how the elements of each mechanism may interact to produce unintended consequences.

In this paper, we investigate the introduction of an asymmetric, randomized speed bump to the Canadian exchange TSX Alpha on September 21, 2015. This speed bump delays (almost) all incoming messages

---

[2] See Rosu (2016 and Baron et al. (2017) for how faster traders capture profits from liquidity provision. Li (2014) and Baldauf and Mollner (2016) show that faster traders impose adverse selection costs.

[3] See Brogaard, Hagströmer, Nordén and Riordan (2015) for co-location, Laughlin, Aguirre, and Grundfest, (2014) and Shkilko and Sokolov (2016) for a discussion of how microwave towers have affected fast trading between the Chicago future pits and the New York equity markets, Menkveld and Zoican (2017) for a discussion of latency upgrades, and Angel (2014) and Laughlin, Aguirre, and Grundfest (2014) for speed of light comparisons.

[4] IEX's speed bump is sometimes considered symmetric as it slows down all lit orders, but exempts dark pegged orders. 80 percent of its trading is dark according to a 2017 report. In the case of the updated CHX LEAD proposal, the payment comes in the form of more stringent quoting requirements rather than through a fee differential.

for 1-3 milliseconds, but provides the option to pay a higher fee to enter and cancel limit orders *without* experiencing the delay.[5] We provide the first empirical evidence on the impact of intentionally slowing down *some* participants in a fragmented environment and document a key insight into what drives fleeting liquidity in today's fast, fragmented markets: participants with speed advantages are able to observe (large) traders' actions on other venues, cancelling standing limit orders faster than traders are able to access them. Our evidence suggests that advance knowledge – even for a millisecond – of institutional investors' trading intentions (in a probabilistic sense) is valuable: delays as short as one millisecond can enable substantial (costly) information leakage across venues.

We start with a detailed analysis of the behavior of liquidity providers on Alpha, including their usage of the delay-exempt order type and the ability of liquidity demanders to access liquidity across venues. Using several newly developed metrics, we show that low-latency liquidity providers on Alpha use the delay-exempt limit orders to "fade away" from incoming market orders, i.e. cancel their orders on Alpha, after observing executions on other venues.[6] We also observe a redistribution of information content and adverse selection contained in trading volume, with Alpha capturing a larger fraction of low adverse selection volume at the expense of the rest of the market. Large liquidity demanders are less successful on Alpha and execute relatively more on the remaining venues, while smaller orders shift towards Alpha (e.g. driven by the incentives of retail brokers to collect the rebates). As a consequence, the profitability of liquidity provision on Alpha vastly improves, while realized spreads on the remaining venues fall, consistent with the theoretical predictions of Baldauf and Mollner (2016).

Battalio, Corwin and Jennings (2016) suggest inverted and non-inverted venues differ in terms of adverse selection and realized spreads. We use the presence or absence of delay-exempt liquidity on Alpha to

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

[6] Such a 'last look' is a common, but controversial feature in currency markets, giving brokers a brief time window in which to review and potentially reject a trade if they deem the transaction price to no longer reflect market conditions. See Rock (1990) for a model of the last-mover advantage and the ability to choose whether or not to trade for equity markets.

disentangle the effects of Alpha's addition of a speed bump and its switch to inverted pricing. Our strongest test is a direct comparison of the two inverted venues in our sample, with Alpha and CX-2 having near identical fee structures and market shares in the post period. The stark differences observed on Alpha relative to CX-2 are exclusive to times when delay-exempt liquidity (which can take advantage of the speed bump) is present on Alpha, while performance metrics are near identical otherwise. Not all market participants can take advantage of the speed bump, as it requires sub-millisecond reaction times. While fast liquidity providers tend to be more profitable everywhere (Rosu 2016), on Alpha they achieve realized spreads more than a full cent higher than they do on CX-2. Slower traders, on the other hand, do worse on Alpha than on CX-2, suggesting a redistribution of adverse selection not only *across* venues but also across speed tiers *within* venues (Shkilko and Sokolov, 2017).

We also compare Alpha's relaunch to previous entries by inverted venues (without speed bump) into the Canadian market and highlight the different mechanism at play between purely inverted venues and Alpha's speed bump mechanism as a function of stock characteristics. The redistribution of adverse selection away from Alpha (and profitability towards Alpha) is heavily concentrated in stocks where trading has a tendency to displace the NBBO: higher priced, less tick-constrained, with thinner quotes, i.e. collectively stocks with a high risk of immediate adverse selection (IAS) for liquidity providers. In contrast, the market entries of other inverted venues previously mostly affected the opposite spectrum, i.e. stocks with a low risk of immediate adverse selection, both in terms of market share as well as the ability to segment order flow, in line with Battalio et al. (2016). These differences suggest Alpha's speed bump complements the inverted fee structure by helping liquidity providers avoid adverse selection also in high-IAS stocks. We also investigate alternative drivers of our results, and find that the removal of Alpha from the Order Protection Rule has little effect on traders' routing decisions. Using two matched U.S. samples, we also show that our results not driven by changes in global market conditions.

Finally, we investigate potential changes in consolidated market quality. Facing increased adverse selection, liquidity providers widen their spreads on other venues, but liquidity providers on Alpha do not narrow spreads in response because being delay-exempt is not useful when being alone at the NBBO.[7] We again confirm that almost all effects are concentrated in high-IAS stocks in the case of Alpha, while earlier entries of inverted venues did not exhibit this cross-sectional pattern.

Our study also makes several methodological contributions by developing new techniques which can be constructed with any public data source to examine market linkages and the accessibility of liquidity in fast, fragmented markets. First, we devise a methodology to aggregate related trades on different venues into trade strings which we use to examine low latency, cross-market liquidity dynamics, extending the metric devised by Malinova and Park (2016) to public data. Second, we propose two innovative classification schemes to analyze a venue's order flow composition. These identify a) trade strings that cause instantaneous adverse selection costs ("*depleting trade strings*"); and b) trade strings that execute across multiple venues, likely to have originated from a smart order router (conceptually similar to the work of van Kervel, 2015). We further quantify a venue's ability to segment beneficial (from the perspective of the liquidity provider) non-depleting trades away from detrimental depleting trades using a new measure called *Segmentation Imbalance*. Third, we employ several novel measures that describe order book fragility during executions. "*Quote Fade"* measures the fraction of liquidity at the NBBO that is cancelled during a trade execution, "*Deplete Best"* measures the fraction of trading volume that exhausts liquidity at either the bid or offer (i.e. depleting trades), and "*Access Next"* measures the fraction of trading that occurs behind the NBBO, indicating liquidity demanders had to "walk the book" increasing transaction costs and effective spreads.

Our results demonstrate that differential speed is valuable to participants, particularly in circumstances where they are able to react to the actions of other traders. Uniquely, the introduction of the randomized speed bump enables fast liquidity suppliers to 'fade' away from orders which consume liquidity across multiple venues, allowing them to profit from the information contained in the order flow itself.

---

[7] In fact, even when an inverted venue only matches a best quote from another venue, it represents the best price after considering fees and rebates and will often be routed to first.

Regulators and market participants should be cautious in approving any mechanism, speedbump or otherwise, which facilitates such conduct, as our results show that it increases total transaction costs and reduces the resiliency of the order book.

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

2 Institutional Details

Similar to the United States, Canadian equities trading is fragmented across multiple venues, with six lit and three dark trading venues.[8] Securities are listed on the Toronto Stock Exchange, operated by the TMX Group, which retains approximately 60 percent of market share by trading activity. The TMX Group also operates Alpha and TMX Select, which was decommissioned once the changes on Alpha occurred. In Canada, NASDAQ operates both Chi-X Canada and CX-2; notably, CX-2 has an inverted maker/taker pricing structure identical to that of new Alpha. Other venues include Omega, Pure Trading, Aequitas Neo, Aequitas Lit and three dedicated continuous dark pools, Match Now, Instinet and Liquidnet. Most lit exchange venues offer price-broker-time priority rather than pure price-time priority as is common in the United States. Across venues, price priority is enforced via "trade-through prohibition". Table A.1 of the

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

appendix reports some relevant characteristics of the four major Canadian markets in our sample such as the prevailing fee structures.

Unlike the United States, internalization of retail order flow in Canada has been significantly constrained. Brokers wishing to internalize trades of less than 5,000 shares are required to provide one full tick of price improvement, or a half tick when the bid ask spread is one tick wide (Larrymore and Murphy, 2009). This mechanism has prevented the growth of retail internalizing venues common in the United States, which account for around 22 percent of trading there (Kwan et al., 2015).[9] As a result of this regulation, and the subsequent banning of payment-for-order-flow, retail orders, which are generally considered to exhibit lower adverse selection and therefore greater profitability for the counterparty, remain predominantly on-exchange in Canada.

2.1 The Alpha Speed Bump

On September 21, 2015, venue Alpha was relaunched as TSX Alpha with the following market design modifications:

1. A randomized speed bump of 1-3 milliseconds for all orders except Post-Only limit orders.[10]
2. An inverted maker-taker pricing model.
3. Orders on Alpha are no longer subject to the Order Protection Rule.[11]

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

---

[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 United States would also be subject to minimum price improvement regulations. As such, unlike in the United States, internalization is not a common practice.

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

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

The market structure facing liquidity demanders after Alpha's relaunch is depicted in the diagram below. Institutional investors who require more liquidity than what is displayed on any single trading venue typically utilize a smart order router (SOR) to spray marketable orders across multiple trading venues simultaneously, efficiently accessing consolidated quoted depth at the national best bid or offer price (O'Hara and Ye, 2011). Alpha's asymmetric, randomized speed bump for incoming marketable orders *but not delay-exempt limit order entries or cancellations* potentially enables its liquidity suppliers to observe the first legs of any large SOR spray being executed on other venues and cancel their limit orders to avoid adverse selection costs.

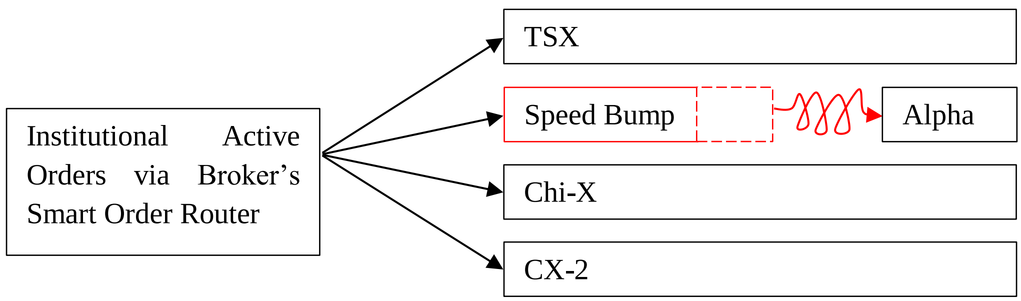

As a consequence, the optimal trading strategy, from the liquidity taker's perspective, may be to send all orders to Alpha when the desired quantity can be filled there alone, but otherwise submit to all venues simultaneously and expect low fill rates on Alpha. Importantly, such concerns are much less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

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

2.2 Comparison with other speed bump designs

Alpha's speed bump design is both similar and dissimilar to other existing and proposed speed bumps in the U.S. The speed bumps on IEX, NYSE American and CHX's proposal all feature an asymmetric delay of 350 microseconds for all liquidity consuming orders but only some liquidity supplying orders. The IEX and NYSE American exemption covers so-called dark pegged orders which are automatically re-priced by reference to the national best bid and offer price.[12] Under CHX's proposal an exemption is provided to the entry and cancellation of visible limit orders from approved liquidity providers.[13] TSX Alpha's exemption is very similar to these in that it only exempts the entry and cancellation of visible Post-Only (non-marketable) limit orders of a minimum size.

Also common across these designs is the feature that the beneficiaries of the exemptions incur a cost compared to the users of other order types: Alpha, IEX and NYSE American have an explicit fee difference; while CHX requires individual approval and minimum liquidity provision.[14] The common asymmetry in favor of liquidity providers is sensible – increases in order processing costs (exchange fees) which are more than recovered through reduced adverse selection increase both liquidity supplier and exchange profits.

There are also some differences between Alpha and the U.S. designs. The largest difference in speed bump design is that Alpha's delay is not deterministic but randomized. Consistent with the model of Baldauf and Mollner (2015), we hypothesize that the random nature of the speed bump may render SORs less effective, leading to the segmentation of liquidity between venues. Other differences are not differences in the speed bump *per se* but are changes that occurred on Alpha concurrent with the speed bump implementation, such as the removal from order protection and the transition to an inverted fee

---

[12] Some consider IEX to be a symmetric speed bump, which is technically true for their lit orders. However, since 80% of IEX's 2% market share originates in the dark, one might debate this classification.
[13] CHX proposal was approved by SEC staff on October 18, 2017, but put on hold by SEC commission on October 26, 2017 to be decided at a later date. See https://www.reuters.com/article/us-sec-chicagostockexchange/u-s-sec-puts-chicago-stock-exchange-speed-bump-on-hold-idUSKBN1CV3EC
[14] Alpha charges a higher fee for non-marketable orders that are exempt from the speed bump. IEX and NYSE American charge a fee for liquidity supplying dark orders, in contrast with no fee on lit orders (which do not benefit from the speed bump), and rebates received on liquidity supplying dark orders on other exchanges.

structure. As shown below, our empirical tests suggest that the change to an inverted fee structure is likely helpful in reinforcing the order flow segmentation that the speed bump tries to achieve. Interestingly, our empirical tests suggest the other two changes matter surprisingly little on their own, and at times on their own provide opposite results of what we find for Alpha, suggesting that they are not driving the results. We now turn to each of these two changes.

2.2.1 Switch to inverted fee structure

To attract retail traders, new Alpha introduces an inverted maker-taker model (also called a "taker-maker" model), providing a rebate to the demander of liquidity, paid for by the liquidity supplier. Since fragmented markets do not have intermarket time priority, liquidity demanders have an incentive to first route marketable orders to venues with the lowest fee or highest rebate (Battalio, Corwin and Jennings, 2016), particularly for brokers that do not pass on the taker rebate to the client. As Brolley and Malinova (2013) note, such a flat fee structure is common for retail brokers in Canada.

Prior to Alpha's fee change, three venues (CX-2, TMX Select, Omega) already adopted an inverted fee structure; with TMX Select ceasing to operate, there remain three inverted venues in the post period as well. As such, we do not expect any significant market-wide impact from the change in Alpha's fee model alone. Using trader classifications, the presence vs. absence of delay-exempt limit orders, cross-sectional differences across stocks, and investigating directly the effect of previous entries of inverted venues without speed bumps, we present below ample empirical evidence that Alpha's joining the other three previously existing inverted markets had little market-wide impact and is not driving the results.

2.2.2 Removal from order protection

The Canadian regulator made Alpha's removal from the order protection rule (OPR) a pre-condition for the approval of the speed bump.[15] Removing a venue's protected status means that traders can choose to

[15] See OSC Bulletin (2015), Volume 38, Issue 16, p. 4045, "Notice of Commission Approval of Proposed Changes to Alpha Exchange Inc." Many U.S. participants have called for reforms of the order protection rule (Rule 611) which forces brokers to search for the best price across all venues, irrespective of delays - speed bumps or otherwise – since the IEX speed bump was approved; see, for example, http://tabbforum.com/opinions/when-was-the-last-air-raid-or-why-wait-to-eliminate-the-order-protection-rule.

avoid routing there. However, as in the United States, brokers in Canada are obliged to strive for best execution: if brokers were to ignore a venue (such as Alpha) with visibly better prices, they would need to explain why this was beneficial for overall execution quality of clients' orders. Thus, regardless of whether the venue is protected, it is in both the clients' and brokers' interest to not categorically ignore a venue, particularly when the venue is offering competitive quotes. If on the other hand, the venue is uncompetitive by quoting behind best, then order protection is moot: the venue will not receive orders, protected or not, due to the best execution requirements. Thus, as long as both sides to a trade find it beneficial to meet on a given venue they will choose to do so, regardless of protection status. While Alpha's market share fell from around 15% to 4% with the relaunch, it quickly recovered to 6-7 percent within a few weeks, indicating that some traders find the "new" Alpha beneficial.

Our later empirical analyses also shows that "new" Alpha is almost never alone at the best, making its' inclusion or exclusion from the order protection rule something of a moot point. We specifically test for instances where quotes on any venue (not just Alpha) have been traded through and find that the fraction of "trade throughs" on new Alpha (which is not subject to the order protection rule) is comparable to the other major inverted venue (which is subject to the order protection rule), indicating that the removal of the order protection rule in itself had little effect on marketable order routing decisions and the incidence of "trade throughs".

3 Data and Methodology

3.1 Data

The data for this study is sourced from Thomson Reuters Tick History (TRTH), supplied by the Securities Industry Research Centre of Asia Pacific (SIRCA). TRTH provides data for each exchange including the state of the limit order book at each quote update, as well as all trade records. The data fields include exchange, security, date, millisecond time stamp, trade price, trade volume, trade qualifiers, buyer and seller broker ID, as well as the price and size at the bid and ask.[16] We remove trades marked 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.[17] We analyze data the for six months around the introduction of the speed bump for the four major Canadian trading venues, TSX, Alpha, Chi-X and CX-2 which together account for more than 90% of all trading.[18]

We include all stocks which remain members of the S&P TSX Composite Index across our sample period which trade above $2 for the majority of days and do not experience M&A or stock split events. This results in a final sample of 232 out of the 250 candidate securities.

We construct a variety of liquidity and order book fragility measures to analyze the impact of the speed bump. To achieve this, we use intraday trade and quote data to manually reconstruct the NBBO price and quantity for each stock-day.[19] A detailed outline of the full methodology including our approach to the attribution of trade direction is provided in Section B.1 of the Internet Appendix,[20] while section B.3 provides technical details on the metric construction.

3.2 Market Quality Metrics

---

[16] Broker identifiers for buyer and sellers are available for TSX and Alpha, unless the broker chose to remain anonymous and forgo participation in broker preferencing. Chi-X does not offer broker preferencing, but some trades contain broker identifiers. Although CX-2 offers broker preferencing, TRTH does not include these identifiers. We thank Chi-X Canada for providing us with public broker identifiers for this venue for the year 2015.

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

[18] Section B.2 of the Appendix provides further description of the sample selection and construction. US public holidays as well as a US flash crash on August 24th, 2015 are excluded from our analysis.

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

[20] The internet appendix that accompanies this paper may be found at https://goo.gl/Sg5qz2.

We measure consolidated liquidity provision against the NBBO using quoted, effective and realized spreads, adverse selection and total market depth. Quoted spreads indicate the average return required by liquidity suppliers for a small round-trip trade at the best quotes. Effective spreads reflect the cost of immediate liquidity demand, taking into consideration the size of the trade. Realized spreads measure the rewards to liquidity suppliers over a 20 second interval, while adverse selection represents the permanent price impact induced by each trade. Depth is the average quoted dollar value available to be bought and sold at the NBBO and is time-weighted across all venues. We also measure the contribution of individual venues to pre-trade liquidity by calculating both the percentage of time a venue is at the NBBO and the percentage of displayed NBBO liquidity.

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

3.3 Construction of High Frequency Trade Strings and Order Routing

Motivated by the importance of linkages between markets highlighted by O'Hara (2015), we construct several novel metrics that rely solely on publicly available trade and quote data and enable us to estimate the ability of liquidity demanders to access quoted liquidity across venues.

These metrics are constructed at both the trade-string level, and at the venue level. We investigate the fraction of trade strings which exhausts available liquidity at the BBO (Deplete Best) creating a new 'best price', and the fraction that 'walks the book', consuming depth at levels past the best (Access Next) as part of their execution. These metrics are not necessarily proxies for trade size, since both executions and cancellations result in the depletion of liquidity.

To construct these measures, we first aggregate high frequency trade "strings" across venues which access liquidity in the same direction at the same time, similar to Malinova and Park (2016). Section B.3 of the Internet Appendix describes the construction of high frequency trade strings in detail and Table A.1 provides summary statistics for these strings.

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

At the venue level, we are interested in the extent to which posted liquidity 'fades' before execution (Quote Fade) and how successful a venue is at segmenting order flow with low instantaneous adverse selection while avoiding order flow with high instantaneous adverse selection (Segmentation Imbalance).

We calculate *Quote Fade* as the proportion of starting liquidity at the NBBO that did not result in a trade. In a depleting trade string, this indicates that the liquidity was withdrawn before it could be executed. Given that orders which deplete the order book generate immediate adverse selection, liquidity providers would ideally avoid such orders, and instead capture those which do not deplete the order book. *Segmentation Imbalance* measures the extent to which a venue is able to access non-depleting orders while avoiding depleting orders, and is calculated as that venue's share of the total non-depleting turnover minus its share of the total depleting turnover, per stock-day. A segmentation imbalance of one (minus one) indicates a venue interacts exclusively with trades that do not (do) impose instantaneous adverse selection costs. The metric is symmetric around zero, with a more positive segmentation imbalance indicating that the venue was more successful at interacting with trades with lower information content.

3.4 Summary Statistics

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

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

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

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

< Insert Table 1 Here >

4 Speed Bump Mechanics and Order Flow Segmentation

In this section, we investigate the effects of TSX Alpha’s implementation of a systematic order processing delay on each of the four major Canadian venues. We start by quantifying the drastically increased ability of liquidity providers on Alpha to “fade away” from incoming market orders compared to other venues and examine how this ability changes the relative composition and information content of trading volume across venues. We verify that this ability is strictly linked to the speed bump rather than the inverted fee structure.

4.1 Fleeting liquidity and the mechanics of reducing adverse selection costs

Alpha’s speed bump of 1 to 3 milliseconds for incoming market orders provides an opportunity for liquidity suppliers to cancel standing limit orders ahead of new marketable orders, particularly after observing trades on other venues. For NBBO-depleting trade strings, we calculate *Quote Fade* for each trading venue as the difference between the liquidity visible at the NBBO at the start of the string and the actual volume traded. A *Quote Fade* measure of zero means that all visible liquidity has been traded.

< Insert Figure 1 Here>

Panel A of Figure 1 presents daily average *Quote Fade* across all stocks in the sample by trading venue. A significant increase in *Quote Fade* is visible on Alpha immediately after the relaunch, while *Quote Fade* decreases slightly across the remaining venues. Given liquidity providers most likely fade when they observe SOR-related executions on other venues, an increase in *Quote Fade* suggests that liquidity providers avoid order flow which interacts with many venues, whilst capturing order flow interacting with only one venue. Consistent with this intuition, Panel B of Figure 1 shows a marked increase in our measure of *Segmentation Imbalance* for Alpha and a commensurate decline across non-Alpha venues. We more formally test the significance of changes in these two liquidity access measures as follows:

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

where $LiqAccess_{i,d,v}$ represents either $Quote\,Fade_{i,d,v}$, the fraction of starting NBBO liquidity withdrawn during NBBO-depleting trade strings, or $Segmentation\,Imbalance_{i,d,v}$, the difference between the uninformed and informed trading volume shares, on venue *v* for stock *i* on day d. The key variable of interest is $Post_d$, an indicator variable equal to one for observations after the introduction of the speed bump and zero prior, which picks up the change in the dependent variable in the post period. The set of control variables includes $Price_{i,d}$, the natural logarithm of the time-weighted NBBO midpoint price, $Volume_{i,d}$ the natural logarithm of aggregate on-market trade volume, $Volatility_{i,d}$ the natural logarithm of the realized intra-day volatility of one minute NBBO midpoint returns, and VIX_d, the daily level of the TSX Composite volatility index.[21]

< Insert Table 2 Here >

Columns marked with (1) in Table 2 contain the results for this test. Consistent with Figure 1, Table 2 shows that Alpha's Quote Fade jumps by 46% on average immediately after the introduction of the speed bump while changes for other venues are minor by comparison. High Quote Fade indicates that quoted

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

liquidity was removed before it could be traded, representing inaccessible liquidity.[22] As a consequence, Alpha becomes unattractive for larger orders that need to access liquidity across multiple venues simultaneously. A high level of accessibility of consolidated market depth across all venues in the pre-event period is consistent with the arguments of O'Hara and Ye (2011) that a trade-through prohibition combined with smart order routing in fragmented markets (virtually) replicates the network advantages of consolidated trading. However, in the post period, the random nature of the delay makes it impossible to guarantee consistently low Quote Fade on multiple venues using an SOR. The corresponding effect in Panel B, an increase in order flow segmentation of 4.2% suggests that cancelations are not random but selective against market orders that are high in adverse selection. The avoidance of depleting trade string by liquidity providers on Alpha largely happens at the expense of liquidity suppliers on the largest venue TSX, who experience a commensurate deterioration in segmentation imbalance.

To confirm that the documented spike in Quote Fade is specifically due to liquidity providers canceling speed bump-exempt limit orders, we repeat the analysis conditional on the presence of limit orders exempt from the speed bump. We identify exempt limit orders through their minimum order size requirement which is both stock-specific and time-varying, determined by the TSX at the beginning of each month as a function of price and liquidity. For the stocks in our sample, this minimum varies between 500 and 5,000 shares.

In Table A.5 of the Appendix, we tabulate for each venue the relative share of cancelled orders whose size at least equals (equals exactly) the minimum size requirement for that particular stock, separately by minimum size stock groups for both the pre and post periods. We observe a significant increase in the cancelation activity of Post-Only orders on Alpha in the post period. For example, in the pre period, for

---

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

stocks with a minimum of 500 shares, only 21.8% of canceled limit order volume is due to cancelations of 500 shares or more. This increases to over 90% in the post period. The difference *exactly at* the limit is even stronger: While only 0.6% of pre-period cancelations have a size of exactly 500, this increases to 62% in the post period. The other venues exhibit no significant changes in cancellations. This pattern is repeated across all minimum size buckets and clearly demonstrates that the vast majority of order cancelations on Alpha in the post period originate from speed bump exempt Post Only orders.[23]

In the following analysis, we assume that exempt orders are present whenever we observe depth on Alpha of at least the minimum exemption size. Each trade string is then categorized depending on whether, at the start of the trade, it faces exempt limit orders on Alpha or not. When exempt orders are present, market orders are subject to the speed bump, while cancellation messages for resting liquidity on Alpha are not (state *Non-Delay*). When depth is below the minimum post-only size, any cancellation message sent to Alpha is subject to the same delay as a market order (state *Delay*). We aggregate trade strings by state in the post-period which results in up to two observations per stock-venue-day and venue. We include a stock- venue-day-state observation if trading volume in that category constitutes at least 10% of total trading volume for that venue to ensure it is representative. We then run the following test:

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

where the Post dummy now captures changes relative to the pre-period during times when either Alpha was not at best or only non-exempt limit orders were present on Alpha, while $Post \times NonDelay$ represents the incremental change (in addition to the Post coefficient) experienced in times when exempt limit orders are present on Alpha. These results are reported in columns marked (2) in Table 2.

The differences in Alpha's Quote Fade behavior and order flow segmentation reported between specification (1) and (2) are economically large. The Quote Fade of exempt limit orders is 64% higher than for non-exempt limit orders while segmentation is 17% higher. In the absence of exempt orders,

[23] Identical results are obtained in Table A.6 where we repeat this analysis for order entry messages.

Alpha does not seem to be particularly effective at segmenting order flow (compared to its pre-period). However, in the presence of an asymmetric delay, Alpha is able to capture a significant portion of trading activity that does not deplete the top of the book, to the detriment of the two non-inverted markets, Chi-X and TSX.

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

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

<Insert Figure 2>

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

Figure 3 presents Alpha's 2-by-2 trade composition for both depleting/non-depleting trades and those executed with/without a SOR. Small orders are likely to be filled on one venue without depleting the NBBO (non-depleting, non-SOR trades). The proportion of non-depleting, non-SOR trades increases dramatically, from 18% on "old" Alpha to 46% on "new" Alpha. Conversely, large (institutional) trades are likely to exhaust all liquidity available at the NBBO using a SOR. The incidence of depleting, multi-venue sprays is halved on new Alpha, from 46% to 23%. Little movement is observed either for single-venue depleting orders (i.e. large retail orders) or for multi-venue sprays which do not displace the NBBO (i.e. small institutional orders). Given these measures are based on traded liquidity, they demonstrate the ability of liquidity suppliers on Alpha to "fade" away from large institutional orders which access multiple venues, while interacting with a relatively larger proportion of (likely) uninformed (retail) flow.[24]

< Insert Figure 3 Here>

4.3 Alternative causes

The previous tests clearly show that the dramatic changes in accessibility and segmentation of "new" Alpha are driven by speed bump-exempt limit orders. Here, we show that these results are not driven by the other contemporaneous changes to Alpha, notably its switch to inverted pricing or its removal from the order protection rule may also contribute.

With Alpha no longer subject to the order protection rule, large liquidity demanders may cease to route to Alpha as part of their multi-venue order submissions, which would give liquidity providers plenty of time to cancel both exempt and non-exempt limit orders. However, we do not find any evidence to support this conjecture. Instead, we find that exempt orders are more than 11 times as likely to be cancelled: the 5.8%

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

increase in quote fade for non-exempt orders found in Table 2 is modest compared to the 70% increase for exempt orders and the entire increase in segmentation occurs in the presence of exempt limit orders only. After the change, non-exempt limit orders on Alpha are still being traded against and are no better at avoiding adverse selection than they were in the pre-period. Second, while the fraction of SOR trades does decline on Alpha, it still represents close to 50% of traded volume in the post period on average (Figure 3). This is a strong indication that Alpha, despite being no longer quote protected, still receives a substantial amount of market orders as part of multi-venue SOR sprays.

It is also possible liquidity providers on Alpha would exhibit similar behavior in the post period even if the delay-exempt limit orders were not available to them. To explore this possibility, we compare "new" Alpha to the other (purely) inverted venue CX-2, having identical fee structures and market shares. Regular limit orders are still available on Alpha and liquidity providers electing to use exempt limit orders incur both an explicit cost (in the form of a higher exchange fee) and an implicit cost (through higher inventory risk due to the higher minimum order size). Thus, their choice of order type is a direct reflection of the benefits of the speed bump. We formally test this below.

4.4 Direct comparison of inverted venues

In the post period, our sample contains two inverted venues with substantial market share and near identical fee structures, Alpha and CX-2. In the following test, we compare their respective ability to fade liquidity and to segment "good" order flow in the post period side-by-side:

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

where, as before, $LiqAccess_{i,d,v,s}$ represents either Quote Fade or Segmentation Imbalance. $Alpha_v$ is an indicator for the venue, equal to 1 (0) for venue TSX Alpha (CX-2). *NonDelay* is an indicator for the presence of delay-exempt liquidity on Alpha. Table 3 presents those results.

In columns marked (1), the coefficient for Alpha represents the unconditional difference in the levels of quote fade and segmentation imbalance between the two inverted venues. On average, Alpha fades almost an additional 50% of its pre-trade liquidity compared to CX-2, while the difference in segmentation is relatively small unconditionally. Specification (2) conditions on the presence of delay-exempt limit orders on Alpha. Even in the absence of the speed bump ,Alpha is almost 5.7% more likely to fade than CX-2, but in its presence that difference grows by an additional 58.7%, confirming the outstanding ability of users of speed bump exempt limit orders to fade on Alpha under identical circumstances as those faced by liquidity providers on CX-2.

More revealing is the comparison of segmentation imbalance. In the absence of the speed bump, Alpha is actually worse than CX-2 at segmenting non-depleting order flow. This may be due to a separating equilibrium, with fast liquidity providers who can take advantage of the speed bump no longer using regular limit orders, leaving only slower traders on CX-2 (Rosu, 2016). We confirm this intuition in Section 5.2 where we explicitly investigate the performance of fast and slow liquidity providers. In the presence of the speed bump on Alpha, CX-2 does 2.5% better at segmenting order flow than in their absence. This suggests that the users of exempt limit orders are selective as to when they post this type of liquidity (likely driven by the larger inventory risk that comes with the larger order size). Finally, the interaction term *NonDelay* × *Alpha* contains the strongest evidence that the speed bump is the reason for the success of Alpha in segmenting order flow. Relative to CX-2, with identical order book conditions and ex-ante adverse selection risk, liquidity providers on Alpha are able to access 15% less toxic order flow than liquidity providers on CX-2.

Taken together, these findings suggest that it is the speed bump on Alpha which plays a pivotal role in allowing traders to avoid adverse selection from larger institutional sweep orders by removing liquidity when they observe contemporaneous executions on other venues. Neither the existence of an inverted fee structure nor the removal of the order protection rule appear to drive these findings, suggesting that although the changes are implemented simultaneously the speed bump is the critical component in the mechanism which facilitates improved realised spreads.

5 Transaction costs across individual venues

We test for changes in transaction costs across the four major venues in our sample in the wake of Alpha's design change as follows:

$$LiqMetric_{i,d,v} = \delta\ Post_d + \Theta^T Controls_{i,d} + FE_i + \varepsilon_{i,d,v} \qquad (13)$$

where $LiqMetric_{i,d}$ is a measure of liquidity on venue *v* for stock *i* on day *d* and all other variables are as described in Equation 11. Changes in effective spreads, realized spreads and adverse selection costs on Alpha after its relaunch are presented in the first column of each block in Table 4. Control variables for price, volume, volatility and VIX have the expected directionality and are statistically significant. Net-of-fees effective spreads on Alpha increase by 0.58 cents on average, following the market structure changes. This is despite the fact that shares over C$1.00 experienced a decline in explicit trading costs of 0.28c per share.[25] Consistent with Malinova and Park (2015), we document that liquidity suppliers pass on changes in explicit fees, even under inverted maker-taker pricing schemes.

< Insert Table 4 Here >

We calculate realized spreads by comparing traded prices with NBBO midpoint quotes 20 seconds after each trade, following Conrad et al. (2015). Realized spreads on Alpha increase 1.23 cents after 20 seconds. This substantial increase prevails despite the fact that the switch to inverted fees should, *ceteris paribus*, lower gains to liquidity provision by 0.14c (old make rebate) plus 0.10c (new make fee).

Table 4 also shows that under Alpha's new market structure, adverse selection costs decline by 0.69 cents. The increase in the realized spread of trades on Alpha indicates that liquidity suppliers are able to either widen their spreads or avoid adverse selection. The observed decreases in adverse selection costs are slightly larger than the increases in effective spreads, indicating that increased profitability of liquidity provision on Alpha is driven mainly by the ability to avoid correlated order flow.

---

[25] Active fees decreased from 0.10c per share on 'old' Alpha to a rebate of 0.18c per share on 'new' Alpha.

What are the effects on the remaining venues? In light of the evidence in Section 4 where we found that Alpha's composition of trading moves towards less informed and less depleting at the expense of its competition, we expect some negative spillover effects into the profitability of liquidity provision on the remaining venues. This expectation is in line with existing empirical evidence suggesting 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).

In the remaining columns of Table 4, we find that Alpha's improved avoidance of informed trades that sweep multiple venues does indeed increase the toxicity of residual order flow on the other large Canadian trading venues, with adverse selection costs increasing significantly by between 0.21c and 0.31c. In most cases, effective spreads widen less, however, imposing a net cost on liquidity providers. This affects the other inverted venue (CX-2) the most, where realized spreads decline a substantial 0.37c, while declines are more modest at 0.12-0.27c on the non-inverted venues.

Alpha's new inverted maker taker pricing and larger quoted depths from minimum Post Only order sizes enable it to aggressively compete with CX-2 for active retail flow, reducing the profitability of liquidity provision on that venue. The large reduction on CX-2 is consistent with a reduction in the proportion of uninformed (retail) order flow in the aggregate of market orders, likely as a result of a migration to Alpha due to the (mandated) larger quoted depths. We conclude that Alpha's segmentation of order flow increases residual order flow toxicity and imposes negative liquidity externalities on the other trading venues, in particular the competing inverted venue.

It is of course possible that some portion of the increase in realized spreads and(decrease in adverse selection is due to Alpha's switch to an inverted pricing structure (Battalio et al., 2016), and not the speed bump alone. To isolate these two channels, we directly compare the performance of Alpha with the other inverted venue CX-2 in the post period while employing two identification strategies for the usefulness of the speed bump: The first one is based on the presence of delay-exempt limit orders, which we already used in Table 3 to understand the incremental ability of the speed bump to effect segmentation. The second identification is based on broker-level IDs of trades allowing us to separate liquidity providers fast enough to exploit the speed bump from those who are likely not.

5.1 Comparison of inverted venues by liquidity type

As previously shown in Tables 2 and 3, the ability of Alpha to segment order flow crucially hinges on the absence vs. presence of speed bump exempt limit orders, i.e. the symmetric vs. asymmetric application of the speed bump. We expect that trading costs on Alpha and CX-2 should be quite similar during times when there is no asymmetry in the application of the speed bump between market and limit orders on Alpha (state *Delay*), while liquidity provision on Alpha should be vastly more profitable when liquidity providers are exempt from the delay (state *non-Delay*). We test our hypothesis in the 10 weeks after the event date using stock-day observations from Alpha and CX-2 jointly in the following specification:

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

Where *Alpha* is an indicator that is 1 (0) for Alpha (CX-2), and *NonDelay* is an indicator that is 1 (0) for that subset of trades that occurs with speed bump exempt limit orders being present (absent). Columns marked (1) in Table 5 contain unconditional comparisons of Alpha and CX-2, i.e. one observation per stock-day-venue, while columns marked (2) separate trades into 2 subsets as a function of the state of Alpha's order book state (*Delay* vs. *NonDelay*); however, we impose the condition that each observation must represent at least 10% of that stock-day-venue's total trading volume to ensure representativeness.

Table 5 shows that unconditionally (columns marked by (1)), liquidity provision on Alpha experiences realized spreads that are significantly higher than on CX-2 by 0.39 cents and lower adverse selection by about 0.18 cents. The difference is made up by higher effective spreads which seems to indicate that LPs on Alpha are less aggressive than on CX-2. However, when we condition on symmetry vs. asymmetry of the speed bump, we find stark differences in columns marked (2) in Table 5. The coefficients of Alpha now represent situations where liquidity providers did not have the advantage of being exempt from the speed bump relative to CX-2. There is no significant difference for effective and realized spreads between the two venues. Adverse selection is slightly higher on Alpha, and while it is statistically significant, its economic magnitude is rather small with 0.08 cents. In other words, net of fees the two inverted venues have almost identical performance in the absence of speed bump induced asymmetry, confirming the first part of our hypothesis.

The coefficient of *NonDelay* represents the incremental effect on both venues when Alpha's order book contains exempt limit orders, while the interaction term reports the additional effect that is Alpha-specific. We find that liquidity provision on either venue is more profitable during times with exempt limit orders present, which suggests that liquidity providers that employ the exempt order type on Alpha are strategic about when to quote at best, but the benefits are about double on Alpha compared to CX-2. For example, realized spreads on CX-2 improve by 0.55 cents, those on Alpha by about double with 0.55 + 0.52 = 1.07 cents. This interaction term represents the additional value that the speed bump on Alpha provides to its users, beyond the inverted fee structure and beyond strategic considerations of when to post.

< Insert Table 5 Here>

5.2 Comparison of inverted venues by trader type

For parts of the sample period we are able to acquire broker level identification data, which we use to distinguish between fast and slow traders on the passive side of trades on Alpha and CX-2. The details of this classification can be found in Appendix A.1. The delay on Alpha is small enough in duration that not every user of limit orders has the means to take advantage of it. Figure A.2 in the Appendix shows that while high frequency traders with direct market access (DMA) provide the majority of liquidity on Alpha in the post period, other trader groups are still present. We assume that most orders that arrive through channels associated with DMA originate from traders fast enough to exploit the exemption (trader type *FastLP*), while at least some of the orders originating through other broker IDs are not (trader type

SlowLP).[26] Following Rosu (2016), it is reasonable to expect that fast traders harvest higher realized spreads and experience lower adverse selection on either venue. But we also hypothesize that faster traders are incrementally better off along both dimensions on Alpha than on CX-2 because of the advantages that the speed bump affords them. To the extent that our identification of differences in reaction time is noisy, it would bias against finding any difference between these two groups that aligns with our hypothesis. We run the following test:

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

Where *Alpha* is an indicator that is 1 (0) for Alpha (CX-2) and *FastLP* is an indicator that is 1 (0) for that subset of trades with a fast (slow) liquidity provider on the passive side, as defined above. These results are reported in columns marked as (3) in Table 5, where we differentiate between trading costs of slow and fast liquidity providers across the two venues. The coefficient of *Alpha* now represents differences between the 2 venues that slow traders experience. We find that slow users of limit orders seem to do poorly on Alpha compared to CX-2 as realized spreads are 0.65 cents lower and adverse selection is 0.35 cents higher. They also tend to trade at lower effective spreads.

FastLP represents the common, incremental effect for fast liquidity providers on both venues, while the interaction term represents the additional, incremental effect for fast traders on Alpha. Not surprisingly, fast traders capture more profits and experience lower adverse selection on either venue (Rosu, 2016), but being fast on Alpha is more profitable still. The difference between fast and slow traders on CX-2 is 0.49 cents, but it is 0.49+1.12 = 1.61 cents on Alpha, again in line with our predictions.

Overall, both comparisons provide evidence for the effectiveness of the speed bump to alter transaction costs in favor of liquidity providers on Alpha. A second noteworthy implication of these results is that the speed bump not only segments order flow between venues and redistributes profits from liquidity

---

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

provision towards Alpha, it also creates winners and losers within Alpha: Fast HFTs and users of exempt limit orders are able to cherry-pick the most favorable trades on Alpha leaving slower traders worse off.[27]

6 Cross Sectional Variation

Battalio et al. (2016) suggest that inverted venues tend to provide liquidity providers with lower adverse selection and higher realized spreads.[28] As shown in the previous section, in the post period the new inverted Alpha provides even lower adverse selection and higher realized spreads than CX-2, which we attributed to the asymmetric randomized speed bump on Alpha using several identification strategies. The relaunch of Alphas as an inverted venue did not represent the first entry of an inverted venue into the Canadian trading landscape. In part to help further separate the effects of a new inverted entrant and the addition of the speed bump, in this section we contrast previous market entries of venues with inverted maker/taker fee structures (henceforth *purely* inverted venues) with the relaunch of Alpha as an inverted venue *plus speed bump* to highlight the different mechanisms at work that lower adverse selection. We develop and then test hypotheses that imply differential outcomes in the cross-section of stocks for pure inverted venues vs. TSX Alpha.

6.1 Inverted markets in the cross-section: A historical comparison

Recent findings in the literature suggest that typical inverted markets tend to attract stocks that are low priced, tick-constrained, and with large depth at the NBBO (e.g. Harris, 2013; Battalio et al., 2016; Cimon, 2016). Such stocks tend to have "deep and dense books", making it less likely that the next market order exhausts all liquidity at a given price point, and even if the market order does exhaust

---

[27] We do not test this directly, but it seems reasonable to assume that the groups *NonDelay* and *FastLP* overlap to a large extent. However, the results between the two tests should differ somewhat because the latter excludes trades with anonymized broker IDs.

[28] For example, in our pre-event period, inverted venue CX-2 has mean and median net realized spreads that are at least 1 cent higher than non-inverted venues and a correspondingly lower adverse selection component.

liquidity at a given price point, the subsequent move in quotes will usually be limited to one tick.[29] By contrast, the limit order book for stocks on the opposite end of the spectrum (higher price, unconstrained, small depth) tends to be thin and sparse: a trade is more likely to move quotes and by more than one cent as the next limit order may be several ticks away.

We can think of the propensity of trading to move quotes as risk of immediate adverse selection (IAS) and in *DepleteBest* we already possess a direct measure of the likelihood of a liquidity provider experiencing IAS. In unreported results, we find that prior to the design change on Alpha, *DepleteBest* is on average around 52% for the former group of stocks (i.e. low priced, tick-constrained, high depth), but about 64% for the latter (high priced, unconstrained, low depth). Henceforth, we will refer to the first group stocks collectively as low-IAS stocks and the latter group as high-IAS stocks. We exploit these cross-sectional differences in IAS (as measured by the above characteristics) in subsequent empirical tests.

Because of the difference in IAS, we first hypothesize that liquidity providers should be relatively more willing to pay to post quotes in low-IAS stocks on an inverted venue, where queues are shorter and tend to be accessed first due to offering the best price net of fee while the risk of adverse selection is low. We also expect that as a consequence of more frequent liquidity provision in low-IAS stocks on inverted venues, segmentation of non-depleting order flow should be more pronounced than for high-IAS stocks.

To test these conjectures and then highlight differences to TSX Alpha, we investigate how previous market entries by inverted venues affected the willingness of liquidity providers to post and the resulting segmentation in the cross-section of stocks. Prior to Alpha, three venues had launched with or switched to a maker/taker pricing scheme between 2011 and 2013. We compile 20-week stock-day panel datasets surrounding each of the three previous market entries and filter them in the same way as our main sample surrounding the relaunch of Alpha. More details can be found in Appendix A.2. We then split the stocks

[29] "Deep and dense books" generally have every price point occupied by a large number of limit orders and the price points are only one tick apart.

from each of the four event samples (including Alpha's relaunch) into two halves according to a total of four cross-sectional characteristics related to the risk of IAS: stock price, the proportion of time that the NBBO is tick constrained, the volume of shares quoted at the NBBO (i.e. depth of book), and *DepleteBest*.

The difference in IAS risk between these two groups of stocks is clearly visible in the willingness of liquidity providers to quote at the NBBO on an inverted venue. Columns 1 to 3 in Table 6 report the average proportion of time that the three purely inverted venues quoted at the NBBO in the 10 weeks following their respective market entries. Liquidity providers on purely inverted venues seem to be very reluctant to quote at NBBO for high-IAS stocks, doing so only between 3% and 11% of the time, while for low-IAS stocks they are more aggressive and quote at best between 20% and 32% of the time, i.e. between 2.1 and 7.6 times as much. Casually speaking, the expected loss from immediate adverse selection is the likelihood of exhausting the NBBO times the expected price move conditional on exhausting the NBBO. Since we expect the price move to be more than one tick for higher priced stocks, it explains the magnified reluctance to post beyond the simple difference in *DepleteBest*.

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

By contrast, TSX Alpha in its post-event period "behaves" quite differently, as is evident in columns 4 and 8. While liquidity providers on Alpha are still relatively more willing to quote at NBBO for low-IAS stocks – Alpha is an inverted venue after all – the differential is proportionately much smaller with ratios of 1.27 to 1.54. In particular, rather than less than 10%, Alpha quotes between 26-29% of the time at

NBBO for high-IAS stocks. Finally, there is strikingly little difference in Alpha's ability to segment the two sets of stocks. However, the level of segmentation at around 6-7% is many times larger than what any of the previous inverted venues was capable of, a result suggestive of the power of the speed bump to magnify segmentation on an inverted venue with speed bump.

< Insert Table 6 here >

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

6.2 The effect of entries of inverted venues on the rest of the market

Table 6 reports the average *level* of segmentation that the newly inverted venues are able to attain in the 10 weeks after their respective entries/relaunches. We now focus on how segmentation (and adverse selection costs) is altered on the remaining venues by a new venue's entry relative to the pre-period, when that venue was either not inverted or not present at all. We investigate post-event *changes* in the cross-section of stocks along the four dimensions that relate to the risk of immediate adverse selection (IAS), introduced in Section 6.1. For each of the four historical entries of inverted venues (Omega, CX-2, TMX Select and TSX Alpha), we run tests separately for segmentation imbalance and adverse selection, allowing for differences in the change in the post-event period for low vs. high-IAS stocks each time:

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

where $y_{i,d,s(c)}$ represents either the segmentation imbalance of all venues except the event venue (Omega, CX-2, TMX Select or TSX Alpha, respectively) or the adverse selection costs imposed on the remaining venues. $HighIAS_i$ is a dummy equal to 1 for stocks with a characteristic that corresponds to a higher risk of IAS (high price, unconstrained tick, low-depth, or high depletion rate). Therefore, the interaction term represents the differential effect on high-IAS stocks in the post period. Table 7 reports those results. For brevity, we only report the two coefficients related to the Post dummy from each of the four regressions, but all tests include the set of daily stock-level control variables (price, volume, realized intra-day volatility, and market volatility) and fixed effects as in previous tables. In addition, we interact all controls with the $HighIAS_i$ dummy to allow for different effects between the two subsamples.

The first three columns of Table 7 demonstrate that consistently across events and across cross-sectional splits, non-event venues experience a greater loss of non-depleting order flow in low-IAS stocks than in high-IAS stocks, as the interaction term is consistently positive and reduces the negative effects observed for low-IAS stocks. In columns 5 to 7, we observe a corresponding increase in adverse selection for low-IAS stocks, while the effects on high-IAS are less clear. The latter group of stocks is on average higher priced and generally spreads should be somewhat higher as well (though not to the same degree). Finding insignificant or even lower effects in some cases should be seen as evidence that the effects are concentrated in low-IAS stocks, consistent with our prediction.

The last column of Table 6 demonstrated new Alpha's ability to segment order flow rather equally for low vs. high IAS stocks. Column 4 of Table 7 shows that from pre- to post-period, the *change* in segmentation experienced by the rest of the market as a consequence is significantly more negative in high-IAS stocks than in low-IAS stocks, at -5.3% (-2.91% + -2.40%) compared to -2.91%. This result for Alpha is many times higher than the effects of any of the three previous entrants had on high-IAS stocks.

Further recall that high-IAS stocks have by definition less trading that does not move quotes (100%-64%=36% vs. 100%-52%=48% for low-IAS). In other words, the redistribution of adverse selection risk caused by Alpha's relaunch is further amplified in high-IAS stocks because there is less non-depleting order flow to begin with. This stands in stark contrast to the previous three inverted market entries where the opposite was the case.

Not surprisingly then, Alpha's entry affects the adverse selection costs for the remaining venues vastly more in high-IAS stocks. While we observe increases of around 0.04 to 0.10 cents for low-IAS stocks, which is comparable to the effects observed in the previous three events, high-IAS stocks experience increases in adverse selection of almost half a cent across all cross-sectional splits.

In the same manner, in Appendix Table A.7, we look at realized spreads of the remaining venues across these four historical events, and by split samples in the cross-section of stocks. Mirroring our findings from Table 7, we observe declines in realized spreads on the non-event venues that are more pronounced in low-IAS stocks for the three previous entries of purely inverted venues. By contrast, TSX Alpha imposes higher reductions on its competing venues for high-IAS stocks. In terms of magnitude, however, it is striking that the reduction in realized spreads is significantly smaller than the observed increases in adverse selections. This would suggest that liquidity providers are able to partially compensate through widening effective spreads. This will be the focus of Section 7.

6.3 Mechanisms of avoiding adverse selection

We ascribe the differences found in the previous tests between purely inverted venues on the one hand and Alpha, an inverted venue cum speed bump, on the other to the presence of the speed bump. As we showed in Tables 2 and 4, essentially all the effects on segmentation and transaction costs are concentrated in times when liquidity provision on Alpha comes in the form of exempt PO limit orders, i.e. when the top of the book liquidity supplier on Alpha has the ability to observe trading activity on other venues and cancel if she fears adverse selection.

Both mechanisms, inverted venue and speed bump, ultimately allow the same outcome, i.e. to lower adverse selection, but this outcome is achieved in different ways. On a (purely) inverted venue, it is optimal to quote at best when the stack of liquidity "behind it" is large (Goldstein, Kwan and Richards, 2017), which serves as an insurance against adverse selection, akin to a level of support or resistance in the parlance of technical analysis. The liquidity provider pays the exchange to jump the queue, but imposes an externality on liquidity providers on the other venues, who are having to wait longer and facing higher adverse selection cost, as trades tend to first exhaust the liquidity on the inverted venues before moving to non-inverted markets. Thin books make this strategy less profitable, hence liquidity providers tend to quote less aggressively in those stocks on inverted venues, and inverted venues tend to have lower market shares in stocks with thinner books.

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

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

7 Effects on consolidated market quality

Speed differentials redistribute adverse selection. The randomized 1-3ms speed bump on Alpha provides market makers with the ability to withdraw liquidity without any delay but delays liquidity takers, giving market makers on Alpha a speed advantage over other participants. Since the market makers on Alpha are able to observe executions on other venues, market makers on Alpha (but not on other venues) can then choose to withdraw liquidity when it is likely they will suffer instantaneous adverse selection. This differential effect will redistribute adverse selection away from the fast market makers on Alpha (who can jump out of the way) to slower liquidity providers on Alpha and on other exchanges.

In addition, inverted maker-taker pricing structure such as that on the new Alpha is shown by Battalio et al. (2016) to be particularly attractive to retail brokers, which choose where to send presumably uninformed (and therefore valuable) retail order flow. Consolidating retail order flow primarily onto a single venue while also creating structures that make the venue less attractive to presumably informed institutional flow is likely to increase market-wide adverse selection, increasing the cost of liquidity. This increase in market-wide adverse selection primarily comes from the increase in information that is created by the speed bump, in that it now allows some market participants (the fast market makers on Alpha) "inside" information on market-wide order flow by providing the market makers a look at the trades on other exchanges. This increase in information is the primary benefit in the "last mover advantage" described in Rock (1990) and others. The following section explores these effects, both for the whole market and in the cross-section of securities where the speed bump is expected to be most valuable.

7.1 Theoretical predictions

Many theoretical works have modelled the impact of differential speed between informed traders and market makers. Biais, Foucault and Moinas (2015), Budish, Cramton and Shim (2015), Foucault Humbert and Rosu (2016) and Foucault, Kozhan and Tham (2016) all envision situations in which informed traders are able to respond to new information faster than market makers. These fast, informed traders are then able to adversely select market makers, forcing them to increase their spreads to recover these costs. While our situation involves some market makers increasing their speed, those who are unwilling/unable to do so will also encounter additional adverse selection, requiring them to increase their spreads. Perhaps most closely related to our situation is the work of Hoffman (2014) and Jovanovic and Menkveld (2016), who consider a situation where some market makers become fast, avoiding adverse selection and thus increasing supplied depth. However, the fast market makers impose additional adverse selection on slower market makers, resulting in wider quoted spreads. This situation is very similar to that observed in Canada, with simultaneous increased depth and widened quoted spreads.[30] A recent paper by Brolley and Cimon (2017) models the introduction of a speed bump alongside an incumbent exchange. Consistent with our results, they find reductions in adverse selection on the speed bump venue, and commensurate increases for the non-speed bump venue. While their model predicts a narrowing of spreads on the speed bump venue (due to lower adverse selection) they do not consider that the *order flow itself* could be the information. In such a situation, the speed bump loses its functionality when it 'makes' the NBBO.

7.2 Unconditional effects

The previous tests show that Alpha is successful at segmenting non-depleting trades, increasing profits for fast liquidity providers on Alpha at the expense of slower traders on Alpha and other liquidity providers on the remaining venues. We now examine whether this redistribution of order flow has a measurable impact on consolidated market quality. We test for changes in market quality metrics jointly across all four venues (Alpha, TSX, Chi-X and CX-2) using the panel regression:

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

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

< Insert Table 8 Here >

The very top panel of Table 8 presents full sample regression results for changes in liquidity metrics across all four trading venues consolidated at the national best bid and offer prices. Quoted spreads increase by a substantial 0.24c in absolute terms. These costs are mirrored by overall increases in the cost

[30] Symmetric deterministic delays do not alter the incentives of the arms race as traders still compete to be first in line to reach the exchange in order to either pick off (snipe) or cancel outstanding stale limit orders. Symmetric, randomized delays potentially create additional problems as they incentivize traders to submit redundant orders, each essentially representing a lottery ticket for a short delay (Budish et al., 2015). Overall, these findings are in line with a much earlier literature on the impact of the segmentation of order flow through payment for order flow schemes (Easley, Kiefer and O'Hara, 1996; Chakravarty and Sarkar, 2002) and also with more recent studies on the segmentation of uninformed order flow imposed by dark trading (Zhu, 2014; Comerton-Forde and Putnins, 2015).

of accessing liquidity (effective spreads) of 0.17c, similar in magnitude to the impact on individual venue effective spreads reported earlier. With increasing adverse selection, it is unsurprising that liquidity providers on other venues widen their spreads, though in theory liquidity providers on Alpha could quote more aggressively due to lowered adverse selection costs. We find, however that liquidity providers do exactly the opposite. Figure 4 shows that Alpha's time at best (i.e. the fraction of time that quotes on Alpha are at the NBBO) declines dramatically, essentially never quoting alone at the NBBO in the post period.

< Insert Figure 4 Here >

Table 8 also documents a significant 17 percent[31] increase in overall depth at the NBBO. This depth increase is in part driven by the increased quoted size on "new" Alpha. While this appears to be beneficial, the large increase in quote fade on Alpha reduces the accessibility of liquidity on Alpha after the introduction of the speed bump.

Table 8 further reports one of our measures of order book fragility. *Access Next* is the proportion of trading volume that 'walked the book', i.e. executing at prices inferior to the pre-trade NBBO. We find that in the post period, an additional 1.7 percent of volume within trade strings was forced to access the next best price. Therefore, although overall displayed market depths increase, trades across all venues were more likely to consume the entire depth available and "walk the book", filling at inferior prices.

7.3 Cross-sectional effects

In Section 6, we document stark cross-sectional differences in the *change* in segmentation caused by the design changes on TSX Alpha. Table 7 reports an 2.9% increase in the segmentation of beneficial (non-depleting) order flow towards Alpha for low-IAS stocks, while for stocks with high-IAS this figure was almost doubled at 5.3%. We also observe different magnitudes in the changes in adverse selection costs that disfavor high-IAS stocks. It thus seems plausible that, if Alpha's speed bump is the cause of overall

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

market deterioration, then overall market quality should be more affected in high-IAS stocks than low-IAS stocks.

To determine the expected impact on market quality from the introduction of an inverted venue (without a speed bump), we turn again to previous introductions of inverted fee structures (Omega in December 2011; CX-2 in May 2013; TMX Select in November 2013). We employ the same setup as before, using 4 cross-sectional sample splits related to measures of immediate adverse selection (IAS), this time investigating overall market quality changes:

$$MQ_{i,d} = (\delta + \gamma\ HighIAS_i)Post_d + (\Theta^T + \Phi^T\ HighIAS_i)\ Controls_{i,d} + FE_i + \varepsilon_{i,d}$$

where $MQ_{i,d}$ represents consolidated quoted and effective spreads, realized spreads, adverse selection, depth as well as order book fragility. $HighIAS_i$ is a dummy equal to 1 for high-IAS stocks. The interaction term represents the incremental effect that the dependent variable for high-IAS stocks experienced in the post period relative to low-IAS stocks.

Appendix Tables A.9-A.11 report results for the market entries of Omega, CX-2 and TMX Select respectively. Contrary to the findings of Comerton-Forde, Gregoire and Zhong (2017) that inverted venues increase market quality overall, across those three events we do not recognize any consistent pattern: The entry of Omega appears to modestly lower quoted and effective spreads for low-IAS stocks, whereas in the case of CX-2 spreads and adverse selection increase modestly across both subsamples. The switch of TMX Select does not change quoted and effective spreads by much, but increases adverse selection. In terms of NBBO depth, there is some evidence that depth went up for low-IAS stocks but less so for high-IAS stocks. Order book fragility deteriorates in the first two cases, but improves in the last. Overall, market quality as measured by spreads generally did not deteriorate a lot in the wake of previous entries of inverted venues and there is also little evidence of a cross-sectional pattern.

With that in mind, we turn to the relaunch of Alpha, results for which are reported in the bottom 4 panels of Table 8. In contrast to the other events, a consistent cross-section pattern of significant and economically large effects becomes apparent for quoted and effective spreads. We find no significant effect on either spread variable for low-IAS stocks, but strongly significant and large increases in both for high-IAS stocks. Quoted spreads widen between 0.36 and 0.49 cents for the latter group relative to the former, while effective spreads widen by slightly lower economic magnitudes. Nominal NBBO depth increases are also more pronounced for high-IAS stocks, which can be explained by the fact that the minimum size requirements for exempt limit orders are likely more binding for the latter group. Finally, increases in order book fragility as measured by *Access Next* are also somewhat higher for high-IAS stocks.

< Insert Table 10 Here >

Prior to Alpha's relaunch, the Canadian market featured three inverted venues. With the closure of TMX Select, three inverted venues remain also in the post period. As the earlier entries of inverted venues do not appear to cause significant changes in market quality, the significant deterioration in market quality suggests a causal effect of the speed bump. This argument is further strengthened by the fact that the deterioration is consistently concentrated in high-IAS stocks, which aligns with our understanding of the operation of the speed bump. Overall, all tests conducted in this section indicate that despite several contemporaneous changes, the primary driver of market-wide impacts is the addition of Alpha's speed bump.

7.4 Strategic presence of Liquidity Providers on Alpha

Table 2 provided evidence that the delay-exempt Post Only limit order type was chiefly responsible for the jump in Quote Fade and the segmentation of 'good' order flow towards Alpha. We now investigate under which market conditions liquidity providers on Alpha are using Post Only orders to aggressively provide liquidity compared to regular limit orders, and how this affects market shares as well as consolidated market outcomes. To do so, we use the following specification:

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

where $MQ_{i,d,s}$ represents a) our two measures of order book fragility that we interpret here as measures of current market conditions (*DepleteBest* and *AccessNext*); b) by-venue market shares; c) consolidated transaction costs measures across all venues (quoted, effective, realized spreads and adverse selection costs). As in Tables 2 and 4, we categorize all trades in the post period to occur either during state *Non-Delay*, when at the start of the trade, the quoted depth at best on the passive side of Alpha exceeded the minimum size that is required for liquidity providers to be exempt from the speed bump, or state *Delay* otherwise. We also require that state s represents at least 10% of that stock-day's aggregate trading volume to avoid giving undue weight to observation based on only a few trades. The focus in this analysis is on the interaction term $Post_d\ \times NonDelay_{i,d,s}$, which represents the differential in the outcome variable between the two states in the post-period.

We report the results of this test in Table 9. In addition to full sample results, Table 9 also reports two subsample analyses for stocks with a price that was below (above) the median stock price in the pre-period. The first two columns of Panel A in Table 9 show that liquidity providers using PO limit orders on Alpha are selective as to when they post the (relatively large) minimum liquidity: the risk of immediate adverse selection is almost 11% lower and the likelihood that a given trade walks the book is 2.6% lower compared to situations where liquidity providers do not post at best or use regular limit orders. This is again consistent with the results of Kwan et al. (2016) who find that HFTs prefer to post competitively on the thick side of the book.

The next four columns report the differential market shares for each of the four venues. Naturally, Alpha's market share is a lot higher, by 17%, when its book contains large Post Only limit orders at best; at the same time the other inverted venue CX-2 also trades a lot more, to the detriment of venue TSX, whose market share is diminished by over 21% during those times. Thus, not only do Alpha's liquidity providers seek out favorable order book conditions ex-ante, they also seem to capture a large portion of desirable order flow when they do.

The final four columns repeat the analysis for measures of transaction costs or market quality, which was conducted unconditionally in the top panel of Table 8, but now conditional on the state of Alpha's order book. Given the difference in market conditions, it is not surprising to see that market-wide adverse selection costs are lower by on average 0.4 cents and realized spreads are higher by on average 0.77 cents during times when Alpha's liquidity providers are aggressive. This naturally requires effective spreads to be wider, at +0.34 cents. Quoted spreads are wider still at 0.5 cents. There are two possible interpretations: the first is that Alpha's liquidity providers are more aggressive during times when quoted spreads are higher; the second is a composition effect whereby Alpha's liquidity provider are more aggressive in stocks with generally higher spreads, for example those with higher prices. To distinguish between these two explanations, we repeat the full sample analysis, but impose a balanced sample condition in addition to the 10% minimum volume filter. That is, in the post-period, a stock-day with trades in state *Non-Delay* is only included in the analysis when the same stock-day also includes an observation from the other state, and vice versa. Results (omitted for brevity) are virtually unchanged and support the first explanation: Users of exempt limit orders on Alpha are very strategic about when they provide liquidity, cream skimming the order flow with the lowest IAS risk and capturing more of it.

The cross-section split shown in Panels B and C shows that all aforementioned effects – order book fragility, market share and market quality differentials – are more pronounced in high priced stocks relative to low priced stocks, which is in line with previous cross-sectional tests. Similar results obtain along the other cross-sectional dimensions related to IAS.

< Insert Table 9 Here >

8 Effect of Other Contemporaneous Changes

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

8.1 Removal of Alpha from the Order Protection Rule

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

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

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

We thus do not find any persistent and systematically large uptick in trade through incidences on new Alpha, especially when compared to the other inverted venue CX-2. Rather, it seems likely that the spike in the ratio is driven by a low likelihood of Alpha quoting at the NBBO, especially in the first few weeks after the change, resulting in a small denominator of our ratio.

< Insert Figure 5 Here >

8.2 Global market conditions

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

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

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

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

Results in Section 6 showed that spreads are affected differently along 4 cross-sectional stock characteristics. Given data quality constraint, we restrict the analysis to a sample split according to size, which is easily implemented in our matched sample because we explicitly used nominal stock price to match. Table 10 presents the results for both matched samples in 3 different versions: The full samples in Panel A, followed by 2 panels separating stocks below (above) the median price in the pre-event period (based on the Canadian half of each pair).

The Post dummy is generally insignificant and often of negative sign for both quoted and effective spreads across both samples, indicating our results are not driven by global factors around the time of the Alpha relaunch unrelated to the speed bump. We do, however, find significant increases in quoted and effective Canadian spreads in the post period relative to the United States. Mirroring our cross-sectional results from Table 8, we find these effects to be strong for the higher priced half of the sample, while completely absent for lower priced stocks.

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

< Insert Table 10 Here >

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

9 Discussion and Conclusion

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

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

Overall, we find that "new" Alpha is not attractive to all participants, with traded volume decreasing immediately. Using novel trade classification schemes, we show that after the introduction of the speed bump the majority of order flow on Alpha shifts towards trades that do not employ smart-order routers (SOR) nor exhaust the top of the order book. We also show that liquidity suppliers on Alpha who can monitor the market in ultra-high frequency are able to harness the information contained within the order flow, avoiding SOR "sprays" attempting to simultaneously access liquidity on all venues. This significantly reduces liquidity suppliers' adverse selection on Alpha, increasing realized spreads and producing substantial economic benefits in an otherwise low-realized spread environment.[32]

We find that the addition of the speed bump in combination with other changes on Alpha caused market wide effects. We find that Alpha's ability to utilize the speed bump to segment predominantly uninformed, low-impact order flow increases the fraction of informed traders on the remaining venues, increasing quoted and effective spreads on the consolidated market by about half a cent. This increase in spreads is primarily driven by increases in adverse selection, consistent with an increase in the fraction of informed traders. We also observe negative market wide effects for large liquidity demanders, with significant increases in both the fraction of trades consuming all available NBBO depth, and walking the book to achieve execution.

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

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

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

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

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

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

References

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

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

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

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

Baron, Matthew, Jonathan Brogaard, Björn Hagströmer, and Andrei Kirilenko. 2017. "Risk and Return in High-Frequency Trading." Working paper.

Battalio, Robert, Shane A. Corwin, and Robert Jennings. 2016. "Can Brokers Have It All? On the Relation between Make-Take Fees and Limit Order Execution Quality." *The Journal of Finance* 71 (5): 2193-2238

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

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

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

Brolley, Michael, and David Cimon. 2017. "Order Flow Segmentation, Liquidity and Price Discovery: The Role of Latency Delays." Working Paper.

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

Budish, Eric, Peter Cramton, and John Shim. 2015. "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response." *The Quarterly Journal of Economics*, 130 (4): 1547-1621

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

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

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

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

Comerton-Forde, Carole, Vincent Gregoire, and Zhuo Zhong. 2017. "Inverted Fee Venues and Market Quality." Working paper.

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

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

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

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

Easley, David, Nicholas M. Kiefer, and Maureen O'Hara. 1996. "Cream-Skimming or Profit-Sharing? The Curious Role of Purchased Order Flow." *The Journal of Finance* 51 (3): 811–33.

Ellis, Katrina, Roni Michaely, and Maureen O'Hara. 2000. "The Accuracy of Trade Classification Rules: Evidence from Nasdaq." *Journal of Financial and Quantitative Analysis* 35 (4): 529–51.

Foucault, Thierry, Johan Hombert, and Ioanid Roşu. 2016. "News trading and speed." *The Journal of Finance* 71(1): 335-382.

Foucault, Thierry, Roman Kozhan, and Wing Wah Tham. 2017. "Toxic arbitrage." *Review of Financial Studies*, 30(4): 1053-1094.

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

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

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

Hoffmann, Peter. 2014. "A Dynamic Limit Order Market with Fast and Slow Traders." *Journal of Financial Economics* 113 (1):156–69.

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.

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

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

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

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

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

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

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

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

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

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

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

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

Rock, Kevin. 1990. "The specialist's order book and price anomalies." Unpublished manuscript, Harvard University.

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

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

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

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

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

Table 1
Summary Statistics

This table reports daily, univariate descriptive statistics across 232 TSX Composite Index component securities 6 months either side of Alpha's relaunch on the 21st of September 2015. *Quoted spread* and *Quoted Depth* are time-weighted and consolidated at the NBBO prices across Alpha, Chi-X, CX-2 and TSX. *Deplete Best* is the fraction of trading that is part of a string displacing all NBBO depth. *Access Next* is the fraction of trading that occurs at the next best price behind NBBO. *Effective spreads* (net of fees) are calculated using the prevailing NBBO midpoint. *Adverse Selection* costs are calculated using the NBBO midpoint 20 seconds after the trade. *Time At Best* is the fraction of time a venue is quoting at the NBBO. Among all depleting trade strings, *Quote Fade* is the proportion of total visible starting liquidity on a venue at the NBB or NBO that did not result in trades. *Segmentation Imbalance* is the difference between a venue's share of non-depleting trading volume and its share of depleting trading volume. *Price* is the time-weighted NBBO midpoint. *% Unconstrained* is the fraction of time that the quoted spread was not tick-constrained. *Volume* is the total quantity of on-market traded shares. *Queue Length* is the time-weighted average number of shares at the NBBO. *Realized Volatility* is the standard deviation of one minute NBBO midpoint returns. *VIX* is the market volatility index for the TSX Composite (VIX).

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

Table 2
Quote Fade and Segmentation Imbalance with/without speed bump

This table investigates changes in *Quote Fade* and *Segmentation Imbalance* by venue using the following 2 specifications:

$$(1)\ LiqAccess_{i,d,v} = \delta\ Post_d + \theta^T Controls_{i,d} + FE_i + \varepsilon_{i,d,v,s}$$

$$(2)\ LiqAccess_{i,d,v,s} = \delta\ Post_d + \gamma\ Post_d\ \times NonDelay_{i,d,s} + \theta^T Controls_{i,d} + FE_i + \varepsilon_{i,d,v,s}$$

Where $LiqAccess_{i,d,v}$ ($LiqAccess_{i,d,v,s}$) is the average *NBBO Quote Fade* or *Segmentation Imbalance* for trades in stock *i* on day *d* at venue *v* (and when Alpha's order book is in state *s*).

For specification (2), at the start of every trade in the post period only, we observe whether TSX Alpha is posting liquidity on the passive side at best and in excess of the minimum require size for non-delayed Post-Only limit orders (state *s=Non-Delay*) or not (state *s=Delay*), where the former describes a situation where liquidity providers on Alpha post competitively and can react without delay to new information. Stock-day-venue-state observations are excluded when trading volume in state *s* represent less than 10 percent of that venue's volume on a given stock-day. By contrast, in specification (1) all trades for a given stock-day and venue are aggregated into one observation. Trades in the pre-period are never split in either specification.

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

Panel A: Quote Fade By Venue

	Alpha		CX-2		Chi-X		TSX	
	(1)	**(2)**	**(1)**	**(2)**	**(1)**	**(2)**	**(1)**	**(2)**
Post	46.37***	5.86***	-2.87***	-6.77***	0.48*	0.26	0.66***	0.35
	(48.53)	(4.09)	(-9.11)	(-10.97)	(1.76)	(0.67)	(4.37)	(1.27)
Post × Non-Delay		64.51***		7.37***		1.41		-0.21
		(48.33)		(10.96)		(1.64)		(-0.27)
Price	-3.05	-2.78	5.77***	7.12***	0.15	2.98	1.39*	0.31
	(-1.20)	(-0.77)	(3.82)	(3.59)	(0.09)	(1.16)	(1.68)	(0.17)
Volume	-0.46	-3.36***	-1.14***	-2.60***	-5.27***	-9.44***	-2.15***	-6.07***
	(-0.74)	(-3.66)	(-3.86)	(-4.58)	(-15.87)	(-14.71)	(-10.16)	(-11.10)
Volatility	1.49	3.90***	2.38***	4.59***	4.03***	9.15***	2.02***	3.33***
	(1.62)	(2.76)	(4.59)	(4.62)	(6.99)	(7.21)	(5.99)	(3.72)
VIX	-2.07***	1.37	-0.53	-0.37	-0.60	-1.35*	-1.61***	-2.42***
	(-3.67)	(1.24)	(-1.02)	(-0.38)	(-1.37)	(-1.78)	(-6.61)	(-5.15)
Stock F.E.	YES	YES	YES	YES	YES	YES	YES	YES
R-squared	0.7152	0.2322	0.2661	0.1333	0.6271	0.2560	0.7453	0.3992
Observations	22,043	29,302	21,906	29,260	22,157	29,470	22,169	29,478

Panel B: Segmentation Imbalance By Venue

	Alpha		CX-2		Chi-X		TSX	
	(1)	**(2)**	**(1)**	**(2)**	**(1)**	**(2)**	**(1)**	**(2)**
Post	4.20***	-0.15	1.03***	0.21	-0.91***	0.60***	-4.32***	-0.66***
	(17.27)	(-1.13)	(6.23)	(1.11)	(-7.14)	(3.44)	(-15.28)	(-2.87)
Post × Non-Delay		17.80***		1.84***		-7.17***		-12.48***
		(55.68)		(9.47)		(-27.86)		(-30.98)
Price	-1.95**	-2.58**	-0.30	0.56	0.69	1.31*	1.56	0.71
	(-2.09)	(-2.59)	(-0.42)	(0.82)	(1.08)	(1.74)	(1.40)	(0.62)
Volume	-1.37***	-1.35***	-3.06***	-2.74***	-1.65***	-1.44***	6.08***	5.53***
	(-8.02)	(-8.21)	(-19.52)	(-16.57)	(-12.10)	(-10.97)	(22.94)	(23.55)
Volatility	0.75***	0.95***	1.43***	1.21***	1.48***	1.38***	-3.67***	-3.54***
	(2.90)	(3.33)	(5.22)	(4.19)	(7.47)	(6.36)	(-9.42)	(-9.72)
VIX	-1.54***	-1.71***	0.16	0.73***	0.16	-0.15	1.23***	1.13***
	(-6.71)	(-7.16)	(0.73)	(3.26)	(0.70)	(-0.65)	(3.66)	(3.35)
Stock F.E.	YES	YES	YES	YES	YES	YES	YES	YES
R-squared	0.2189	0.5604	0.2823	0.2043	0.1972	0.3031	0.3078	0.3618
Observations	22,169	29,491	22,169	29,491	22,169	29,491	22,169	29,491

Table 3

Post Period Only: Quote Fade and Segmentation Imbalance on inverted venues

This table investigates differences in *Quote Fade* and *Segmentation Imbalance* between the two major inverted venues (CX2 and Alpha) following Alpha's introduction of a speed bump, using the following 2 specifications:

$$(1)\ LiqAccess_{i,d,v} = \alpha\ Alpha_v + \Theta^T Controls_{i,d} + FE_i + \varepsilon_{i,d,v}$$

$$(2)\ LiqAccess_{i,d,v,s} = \alpha\ Alpha_v + \gamma\ NonDelay_{i,d,s} + \eta\ NonDelay_{i,d,s} \times Alpha_v + \Theta^T Controls_{i,d} + FE_i + \varepsilon_{i,d,v,s}$$

Where $LiqAccess_{i,d,v}$ ($LiqAccess_{i,d,v,s}$) is the average *NBBO Quote Fade* or *Segmentation Imbalance* for trades in stock *i* on day *d* at venue *v* (and when Alpha's order book is in state *s*). $Alpha_v$ is an indicator for the venue, equal to 1 (0) for venue TSX Alpha (CX-2). For specification (2), at the start of every trade, we observe whether TSX Alpha is posting liquidity on the passive side at best and in excess of the minimum require size for non-delayed Post-Only limit orders (state *s=Non-Delay*) or not (state *s=Delay*), where the former describes a situation where liquidity providers on Alpha post competitively and can react without delay to new information. Stock-day-venue-state observations are excluded when trading volume in state *s* represent less than 10 percent of that venue's volume on a given stock-day. In specification (1) all trades for a given stock-day and venue are aggregated into one observation.

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

	Quote Fade		**Segmentation Imbalance**	
	(1)	**(2)**	**(1)**	**(2)**
Alpha	49.18***	5.73***	-1.01***	-4.65***
	(46.74)	(4.69)	(-4.81)	(-33.27)
Non-Delay		5.91***		2.54***
		(8.04)		(13.61)
Non-Delay × Alpha		58.71***		15.64***
		(43.43)		(34.38)
Price	2.13	6.37**	-0.15	0.13
	(0.83)	(1.99)	(-0.27)	(0.19)
Volume	-0.94*	-2.18***	-2.83***	-2.44***
	(-1.69)	(-2.79)	(-22.88)	(-15.95)
Volatility	0.48	2.56**	1.58***	1.43***
	(0.41)	(2.23)	(7.77)	(6.04)
VIX	-8.38***	7.89***	-3.88***	-0.68
	(-4.10)	(2.61)	(-10.16)	(-1.34)
R-squared	0.5889	0.2507	0.2097	0.3688
Observations	21,804	40,424	21,966	41,014

Table 4
Trade-Based Liquidity Metrics By Venue around Alpha Relaunch

This table reports changes in measures of transaction costs for a sample of 232 TSX Composite stocks on all 4 major trading venues separately around the introduction of the speed bump using the following specification:

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

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

	Effective Spreads				Realized Spreads				Adverse Selection			
	Alpha	**CX-2**	**Chi-X**	**TSX**	**Alpha**	**CX-2**	**Chi-X**	**TSX**	**Alpha**	**CX-2**	**Chi-X**	**TSX**
Post	0.58***	-0.10**	0.06	0.15***	1.23***	-0.37***	-0.27***	-0.12***	-0.69***	0.21***	0.25***	0.31***
	(5.80)	(-2.12)	(0.94)	(3.15)	(9.63)	(-7.59)	(-6.80)	(-5.04)	(-8.06)	(4.87)	(3.69)	(4.48)
Price	2.47***	3.32***	2.95***	2.71***	-0.71	0.69	-0.73***	-1.22***	3.23***	2.40***	3.88***	4.06***
	(4.90)	(3.24)	(3.52)	(3.24)	(-0.89)	(1.45)	(-3.70)	(-5.48)	(2.80)	(4.92)	(3.69)	(3.71)
Volume	-0.73***	-0.71***	-0.72***	-0.64***	-0.18**	-0.17***	0.09**	0.46***	-0.52***	-0.46***	-0.83***	-1.11***
	(-9.22)	(-8.65)	(-8.31)	(-8.62)	(-2.39)	(-3.42)	(2.39)	(12.63)	(-6.30)	(-8.20)	(-9.69)	(-10.84)
Volatility	1.94***	1.92***	1.81***	1.55***	-0.28	-0.27**	-1.19***	-1.99***	2.17***	2.02***	2.92***	3.59***
	(7.06)	(9.01)	(7.88)	(7.30)	(-1.14)	(-2.04)	(-14.75)	(-15.83)	(9.99)	(13.29)	(11.03)	(9.83)
VIX	0.18**	0.54***	0.46***	0.39**	0.03	0.44***	0.35***	0.41***	0.28	0.08	0.24	0.06
	(2.02)	(2.73)	(3.04)	(2.44)	(0.18)	(4.47)	(6.17)	(5.52)	(1.05)	(0.71)	(1.14)	(0.28)
Stock F.E.	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
R-squared	0.8407	0.8575	0.8730	0.9174	0.3136	0.4674	0.2167	0.3960	0.7318	0.7069	0.7973	0.8845
Observations	22,088	22,035	22,159	22,169	22,088	22,035	22,159	22,169	22,088	22,035	22,159	22,169

Table 5

Post Period Only: Liquidity Metrics of Inverted Venues, Non-Delayed Orders and Fast v. Slow Traders

This table compares measures of transaction costs between the two major inverted venues (CX-2 and Alpha) following Alpha's introduction of a speed bump. We alternatively classify trades according to whether speed bump exempt liquidity is present on Alpha at the start of the trade, or what type of trader is supplying liquidity, using the following 3 specifications:

$$(1)\ LiqMetric_{i,d,v} = \alpha\ Alpha_v + \theta^T Controls_{i,d} + FE_i + \varepsilon_{i,d,v}$$

$$(2)\ LiqMetric_{i,d,v,s} = \alpha\ Alpha_v + \gamma\ NonDelay_{i,d,s} + \eta\ NonDelay_{i,d,s} \times Alpha_v + \theta^T Controls_{i,d} + FE_i + \varepsilon_{i,d,v,s}$$

$$(3)\ LiqMetric_{i,d,v,t} = \alpha\ Alpha_v + \gamma\ FastLP_{i,d,t} + \eta\ FastLP_{i,d,t} \times Alpha_v + \theta^T Controls_{i,d} + FE_i + \varepsilon_{i,d,v,t}$$

where $LiqMetric_{i,d,v}$ ($LiqMetric_{i,d,v,s}$; $LiqMetric_{i,d,v,t}$) is one of several transaction cost metrics for stock *i* on day *d* on venue *v* (when Alpha's order book is in state *s*; or when the liquidity supplying trader is of type *t*). *Effective spreads* (net of fees) are measured against the prevailing NBBO midpoint, while *Realized spreads* (net of fees) and *Adverse selection* costs use the reference NBBO midpoint 20 seconds after the trade. $Alpha_v$ is an indicator equal to 1 (0) for venue TSX Alpha (CX-2). For specification (2), at the start of every trade, we observe whether TSX Alpha is posting liquidity on the passive side at best and in excess of the minimum require size for non-delayed Post-Only limit orders (state *s=Non-Delay*) or not (state *s=Delay*). For specification (3), using broker-level trade ID, we distinguish liquidity providers that trade through 2 global banks that primarily offer direct market access services to proprietary traders (Type *FastLP*) and those that trade through other brokers (retail, other banks; type *SlowLP*). We exclude trades with anonymized broker ID. Observations are excluded when trading volume in state *s* (or of trader type *t*) represent less than 10 percent of that venue *v*'s volume for a given stock-day. Daily stock-level control variables (calculated in log levels) include time-weighted average price, daily volume, realized 1-minute intra-day volatility, and the market volatility index for the TSX Composite (VIX), plus stock fixed effects. All measures are in cents. The sample period covers 10 weeks in the post period of the event, i.e. September 21, 2015 to November 25, 2015. */**/*** indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by stock.

	Effective Spreads			**Realized Spreads**			**Adverse Selection**		
	(1)	**(2)**	**(3)**	**(1)**	**(2)**	**(3)**	**(1)**	**(2)**	**(3)**
Alpha	0.16***	0.05	-0.37***	0.39***	0.01	-0.64***	-0.18***	0.08***	0.35***
	(4.82)	(1.55)	(-8.14)	(6.16)	(0.18)	(-9.10)	(-5.06)	(2.90)	(7.14)
Non-Delay		0.23***			0.55***			-0.33***	
		(6.66)			(11.41)			(-11.73)	
Non-Delay × Alpha		0.07*			0.52***			-0.44***	
		(1.86)			(8.75)			(-10.72)	
Fast			0.52***			0.49***			0.05
			(12.45)			(7.26)			(1.03)
Fast x Alpha			0.65***			1.12***			-0.51***
			(9.38)			(12.90)			(-9.58)
Price	2.33***	2.14***	2.75***	-0.26	-0.31	-0.24	2.53***	2.27***	2.50***
	(4.37)	(4.19)	(3.42)	(-1.11)	(-1.46)	(-0.69)	(5.25)	(5.14)	(6.71)
Volume	-0.60***	-0.54***	-0.56***	-0.11***	-0.11***	-0.00	-0.47***	-0.40***	-0.50***
	(-9.82)	(-10.03)	(-9.22)	(-2.62)	(-3.16)	(-0.05)	(-8.76)	(-7.95)	(-9.59)
Volatility	1.71***	1.54***	1.73***	-0.39***	-0.39***	-0.50***	2.14***	1.90***	2.09***
	(9.79)	(11.23)	(10.08)	(-5.31)	(-6.52)	(-6.23)	(13.98)	(14.05)	(15.52)
VIX	0.01	0.05	0.39	0.08	0.19*	0.15	-0.21	-0.22	-0.05
	(0.05)	(0.31)	(1.61)	(0.65)	(1.73)	(0.95)	(-1.41)	(-1.60)	(-0.45)
Stock F.E.	YES	YES	YES	YES	YES	YES	YES	YES	YES
R-squared	0.8698	0.8534	0.8494	0.3562	0.3085	0.3091	0.6864	0.5878	0.6109
Observations	21,866	40,947	32,729	21,866	40,947	32,729	21,866	40,947	32,729

Table 6

Historical comparison of inverted market entries: Quote Aggressiveness and Segmentation Ability by Stock Characteristic

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

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

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

Table 7
Historical comparison of inverted market entries: Effects on the existing venues by Stock Characteristic

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

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

where $y_{i,d}$ is segmentation imbalance or adverse selection across all remaining venues (i.e. not the event venue) for stock i on day d. $Post_d$ is the indicator for the post-event period. $HighIAS_i$ is an indicator that is equal to 1 for stocks in the half of the sample associated with a higher risk of immediate adverse selection using one of the above characteristics. *Segmentation Imbalance* is the difference between the combined share of non-depleting trading volume of all other venues and their share of depleting trading volume. *Adverse Selection* is measured with regards to the NBBO midpoint 20 seconds after the trade. Daily stock-level control variables (calculated in log levels, omitted for brevity) include time-weighted average price, daily volume, realized 1-minute intra-day volatility, and the market volatility index for the TSX Composite (VIX), plus stock fixed effects. In addition, all control variables are interacted with the $HighIAS_i$ dummy. The sample period covers 10 weeks on either side of the event. */**/*** indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by stock.

	Segmentation Imbalance for all venues except:				Adverse Selection for all venues except:			
	Omega	**CX-2**	**TMX**	**Alpha**	**Omega**	**CX-2**	**TMX**	**Alpha**
Panel A:								
Post	-0.84***	-2.76***	-0.52***	-2.91***	0.04	0.11***	0.07***	0.04***
	(-6.65)	(-13.81)	(-7.79)	(-10.78)	(1.25)	(7.49)	(3.75)	(2.71)
Post x High Price	0.47***	0.84***	0.14*	-2.40***	-0.01	-0.02	-0.10**	0.43***
	(3.39)	(3.57)	(1.67)	(-5.33)	(-0.11)	(-0.35)	(-2.09)	(3.42)
Panel B:								
Post	-0.96***	-3.10***	-0.58***	-3.58***	0.02	0.09***	0.05***	0.04***
	(-7.61)	(-16.42)	(-9.44)	(-11.78)	(0.76)	(8.19)	(5.63)	(3.02)
Post x Unconstrained	0.71***	1.54***	0.24***	-1.15**	0.07	0.09*	-0.09**	0.44***
	(5.26)	(7.09)	(3.06)	(-2.43)	(1.03)	(1.85)	(-1.99)	(3.49)
Panel C:								
Post	-0.99***	-3.09***	-0.57***	-3.83***	0.01	0.09***	0.04***	0.03***
	(-7.77)	(-16.00)	(-9.42)	(-11.84)	(0.44)	(6.70)	(4.77)	(2.99)
Post x Low Depth	0.76***	1.56***	0.23***	-0.64	0.09	0.08	-0.08*	0.43***
	(5.59)	(7.23)	(2.90)	(-1.35)	(1.26)	(1.62)	(-1.85)	(3.40)
Panel D:								
Post	-0.78***	-2.58***	-0.53***	-3.22***	0.09	0.09***	0.08***	0.10***
	(-6.79)	(-13.00)	(-8.58)	(-10.54)	(1.55)	(4.89)	(3.98)	(3.47)
Post x High Depletion	0.37***	0.45*	0.10	-1.84***	-0.07	0.15***	-0.04	0.36***
	(2.69)	(1.82)	(1.21)	(-3.96)	(-0.96)	(3.21)	(-0.77)	(2.83)

Table 8

Consolidated Market Quality after Alpha Relaunch by Stock Characteristic

This table investigates how consolidated measures of market quality are affected around the introduction of the speed bump on Alpha on September 21, 2015 as a function of cross-sectional characteristics. Stocks are categorized based on one of 4 pre-event characteristics: *Low vs. High Price* refers to the stock price; *Tick-Constrained vs. Unconstrained* refers to the proportion of time that the quoted spread is constrained by the minimum tick size; *Thin vs. Thick Book* refers to the time-weighted average number of shares at the NBBO. *Low vs. High Depletion* refers to the proportion of daily trading volume that displaced the entire depth on one side of the NBBO. We run the following test using each of the 4 characteristics, in additional to an unconditional test without interaction terms:

$$MQ_{i,d} = (\delta + \gamma\ HighIAS_i)Post_d + (\Theta^T + \Phi^T\ HighIAS_i)\ Controls_{i,d} + FE_i + \varepsilon_{i,d}$$

where $MQ_{i,d}$ is one of several market quality measures for stock *i* on day *d*: *Quoted spreads* and quoted (log dollar) *Depth* are time-weighted and consolidated at the national best bid and offer prices across Alpha, Chi-X, CX-2 and TSX. *Effective spreads* (net of fees) are measured against the prevailing NBBO midpoint. *Realized spreads* and *Adverse Selection* are measured with regards to the NBBO midpoint 20 seconds after the trade. *Access Next* is the proportion of trading volume that occurs at any price behind NBBO. $Post_d$ is the indicator for the post-event period. $HighIAS_i$ is an indicator that is equal to 1 for stocks in the half of the sample associated with a higher risk of immediate adverse selection using one of the above characteristics (these terms are excluded in the full sample test). Daily stock-level control variables (calculated in log levels) include time-weighted average price, daily volume, realized 1-minute intra-day volatility, and the market volatility index for the TSX Composite (VIX), plus stock fixed effects. In addition, all control variables are interacted with the $HighIAS_i$ dummy. For brevity, all control terms are omitted from the table. The sample period covers 10 weeks on either side of the event and runs from July 13, 2015 to November 25, 2015. */**/*** indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by stock.

	Quoted Spreads	Effective Spreads	Realized Spreads	Adverse Selection	Depth	Access Next
Full Sample:						
Post	0.24***	0.17***	-0.09***	0.29***	0.16***	1.74***
	(3.77)	(3.33)	(-3.55)	(4.10)	(19.85)	(12.94)
Sample Split 1:						
Post	0.02	0.01	-0.04**	0.05***	0.13***	1.73***
	(0.66)	(0.39)	(-2.26)	(2.69)	(13.04)	(8.85)
Post x High Price	0.39***	0.28***	-0.07	0.41***	0.04**	0.05
	(3.65)	(3.22)	(-1.57)	(3.37)	(2.55)	(0.18)
Sample Split 2:						
Post	-0.00	-0.01	-0.03***	0.03**	0.12***	1.37***
	(-0.08)	(-0.70)	(-3.01)	(2.45)	(11.79)	(8.39)
Post x Unconstrained	0.44***	0.31***	-0.06	0.43***	0.07***	0.74***
	(4.13)	(3.58)	(-1.42)	(3.55)	(4.84)	(2.84)
Sample Split 3:						
Post	-0.03*	-0.02*	-0.04***	0.03**	0.12***	1.37***
	(-1.69)	(-1.96)	(-3.99)	(2.18)	(11.65)	(8.32)
Post x Low Depth	0.49***	0.33***	-0.04	0.43***	0.07***	0.75***
	(4.58)	(3.81)	(-0.92)	(3.51)	(4.70)	(2.89)
Sample Split 4:						
Post	0.06	0.02	-0.07***	0.10***	0.12***	1.53***
	(1.55)	(0.98)	(-3.39)	(3.18)	(11.89)	(9.78)
Post x High Depletion	0.36***	0.28***	-0.01	0.34***	0.07***	0.43
	(3.20)	(3.13)	(-0.14)	(2.72)	(4.82)	(1.63)

Table 9
Strategic Liquidity Provision and Consolidated Market Quality in the cross-section

This table investigates how market conditions, by-venue market shares and consolidated market quality differ in the post period as a function of whether speed bump exempt Post Only limit orders were present at the NBBO on Alpha or not, using the following specification:

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

where $MQ_{i,d,s}$ is one of several measures of market conditions, market shares and market quality in stock *i* on day *d* when Alpha's order book is in state *s*. At the start of every trade in the post period, we observe whether TSX Alpha is posting liquidity on the passive side at best and in excess of the minimum require size for non-delayed Post-Only limit orders (state *s=Non-Delay*) or not (state *s=Delay*). Observations are excluded when trading volume in state *s* or represent less than 10 percent of that venue *v*'s volume for a given stock-day. No split occurs in the pre-period. $Post_d$ is the indicator for the period after (and including) the introduction of the speed bump, September 21, 2015. In addition to full sample results reported in Panel A, we also divide the sample in 2 equal parts using the average pre period price; below (above) median price stocks are analysed in Panel B (Panel C).

Deplete Best represents the proportion of trading volume in state *s* that occurred as part of a trade string that displaced the entire depth on one side of the NBBO. *Access Next* is the proportion of trading volume in state *s* that occurs at any price behind NBBO. Market shares in percent and computed within state *s* across all venues. *Quoted spreads* are measured at the start of a trade and volume-weighted based on the consolidated national best bid and offer prices across Alpha, Chi-X, CX-2 and TSX. *Effective spreads* (net of fees) are measured against the prevailing NBBO midpoint. *Realized spreads* and *Adverse Selection* are measured with regards to the NBBO midpoint 20 seconds after the trade. Daily stock-level control variables (calculated in log levels, omitted for brevity) include time-weighted average price, daily volume, realized 1-minute intra-day volatility, and the market volatility index for the TSX Composite (VIX), plus stock fixed effects. The sample period covers 10 weeks on either side of the event and runs from July 13, 2015 to November 25, 2015. */**/*** indicate statistical significance at the 90/95/99 percent levels, respectively. Standard errors are clustered by stock.

	Deplete Best	Access Next	Market Share By State: Alpha	CX-2	Chi-X	TSX	Quoted Spreads	Effective Spreads	Realized Spreads	Adverse Selection
Panel A: Full Sample										
Post	3.75***	2.20***	-10.05***	1.06***	2.90***	6.09***	0.00	0.03	-0.17***	0.17***
	(16.58)	(15.77)	(-42.02)	(5.61)	(19.64)	(23.06)	(0.02)	(0.97)	(-7.41)	(4.52)
Post × Non-Delay	-10.80***	-2.65***	17.64***	5.51***	-1.56***	-21.59***	0.50***	0.34***	0.77***	-0.40***
	(-28.75)	(-15.97)	(50.35)	(35.22)	(-11.38)	(-63.94)	(11.67)	(8.64)	(13.76)	(-14.76)
Panel B: Below-median price										
Post	4.08***	1.98***	-10.54***	1.38***	3.56***	5.61***	-0.02	0.00	-0.08***	0.09***
	(12.67)	(10.03)	(-31.89)	(5.73)	(16.34)	(13.72)	(-1.19)	(0.14)	(-5.16)	(5.10)
Post × Non-Delay	-8.63***	-1.79***	15.99***	5.12***	-2.17***	-18.93***	0.18***	0.05**	0.36***	-0.30***
	(-15.73)	(-9.27)	(38.58)	(25.25)	(-10.74)	(-46.62)	(7.79)	(2.19)	(14.86)	(-17.21)
Panel C: Above-median price										
Post	3.50***	2.42***	-9.65***	0.79***	2.30***	6.56***	0.00	0.04	-0.25***	0.23***
	(11.10)	(12.35)	(-28.69)	(2.76)	(12.44)	(19.11)	(0.10)	(0.77)	(-5.99)	(3.59)
Post × Non-Delay	-13.29***	-3.60***	19.56***	5.92***	-0.92***	-24.57***	0.85***	0.65***	1.23***	-0.53***
	(-34.58)	(-14.33)	(35.85)	(25.34)	(-5.92)	(-62.57)	(11.31)	(9.71)	(12.18)	(-11.30)

Table 10
Market Quality Comparison with Matched and Interlisted U.S. Equities around the Alpha Relaunch

This table reports how measures of consolidated market quality experience different changes after Alpha's relaunch i) between a sample of 232 TSX Composite stocks and 232 matched equities in the U.S. and ii) between a subset of 83 interlisted TSX composite stocks and their respective U.S. exchange listings. Results are reported for A) the full sample, B) stocks with pre-period average price below the median, C) above the median using the following specification:

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

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

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

Figure 1

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

Panel A: NBBO Quote Fade By Venue

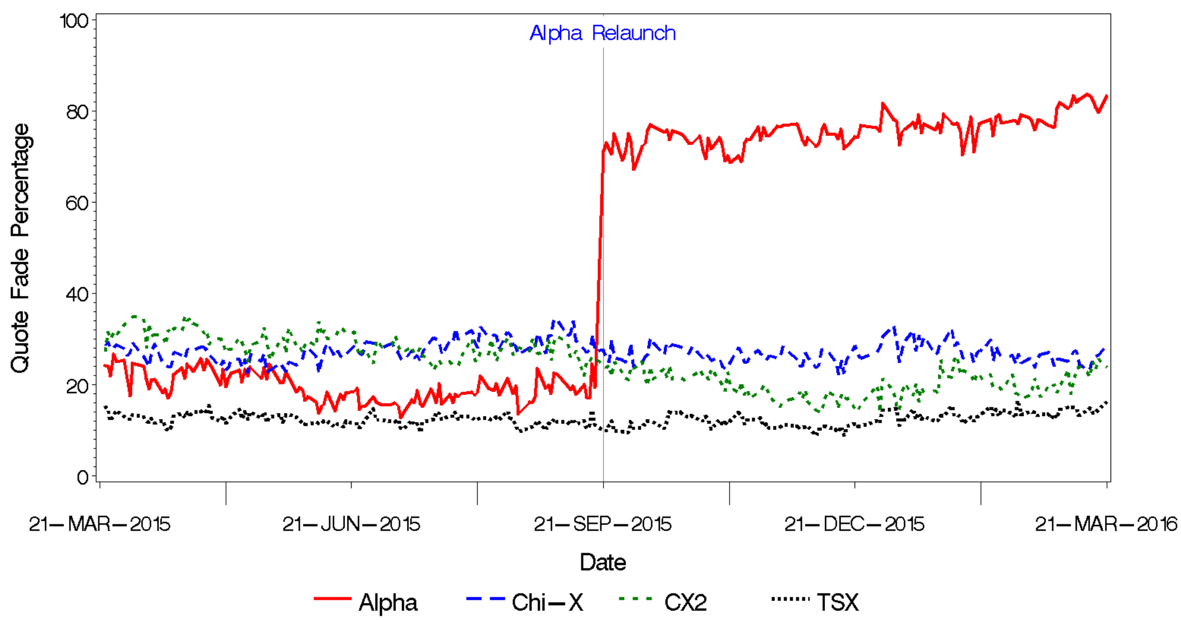

Panel B: Segmentation Imbalance by Venue

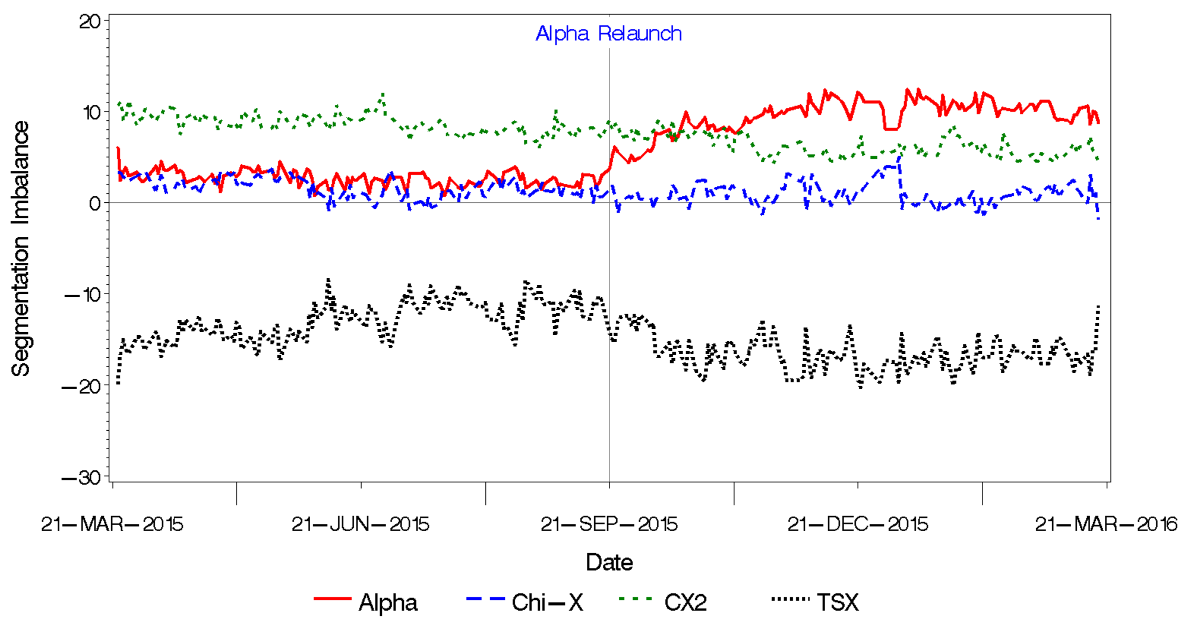

Figure 2
Realized Spread within One Minute by Number of Venues Accessed

This figure presents average realized spread associated with trade strings over 100 milliseconds, 1, 5, 10, 20 and 60 seconds. Trade strings are split into those which access only a single venue, and those which access multiple venues. Trade strings are defined as series of trades in the same direction which execute within a 30 millisecond window. A full explanation of the construction of trade strings can be found in Internet Appendix B.4.

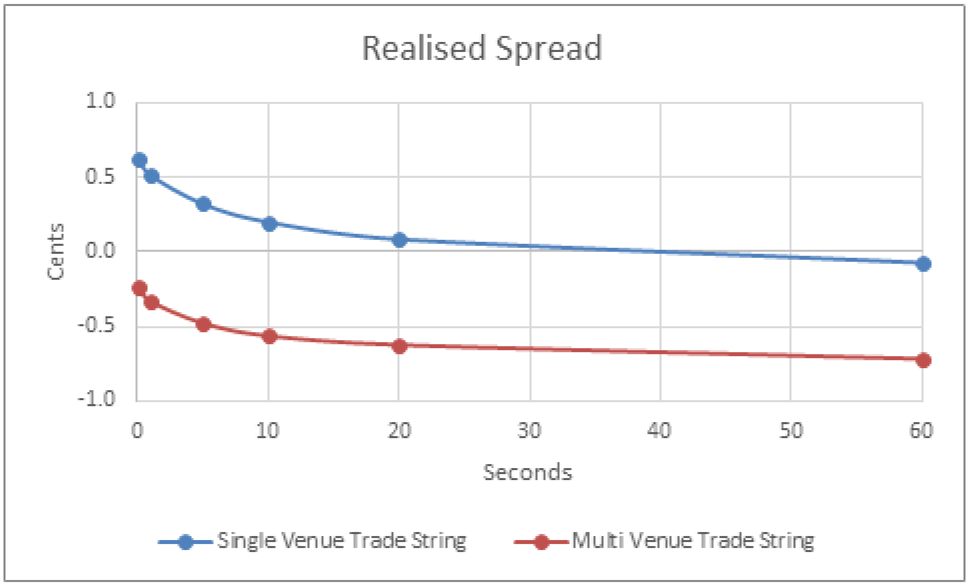

Figure 3
Trading Volume Composition by Trade String Type on Alpha

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

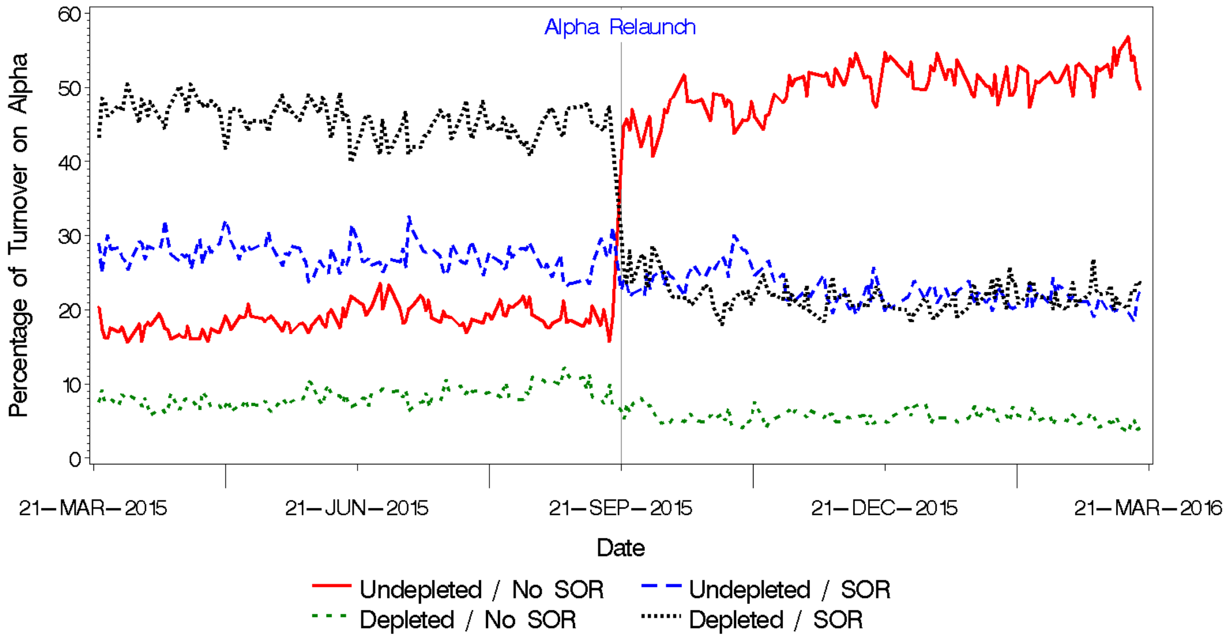

Figure 4
Aggressiveness of Quoting on Alpha

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

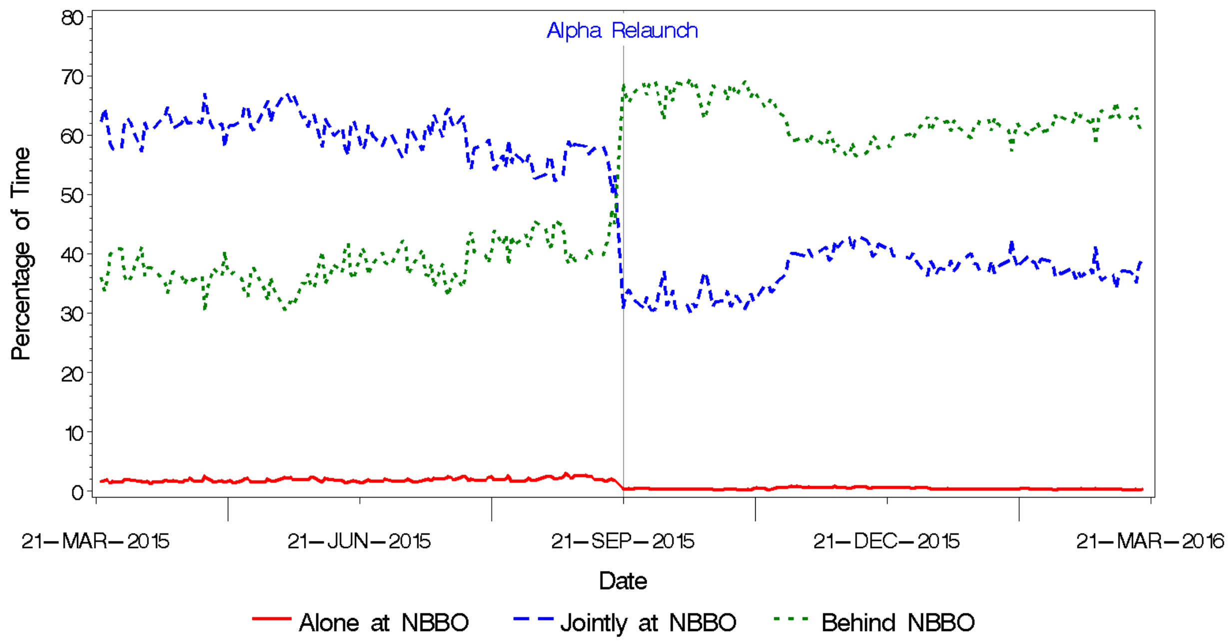

Figure 5
Incidence of "Trade Throughs" by Venue

To identify trade through events, we first identify all trade strings that accessed multiple price levels. We then define a trade through as the event where one venue had a quote at best at the beginning of a trade string and at the end of the trade string but no shares were traded on that venue, while trades occurred at a worse price on another venue. Every stock-day and for every venue, we compute the ratio of volume executed at inferior prices, capped by the liquidity displayed at the NBBO on the traded-through venue at the start of the trade through event, over the total volume of all trades that accessed multiple price levels and where the traded-through venue was at best at the beginning. In other words, we compute the proportion of trading that could have traded at better prices on the traded-through venue during all multi-price level trades. This figure displays the average level of this ratio over time by venue.