Fast Traders Make a Quick Buck: The Role of Speed in Liquidity Provision*

Markus Baldauf† Joshua Mollner‡

December 31, 2016

Abstract

We study the consequences of information arrival for market outcomes when both high-frequency and slower traders provide liquidity. We present a model that predicts faster traders achieve a relative increase in profits obtained from liquidity provision following a news event through (*i*) avoiding adverse selection by canceling mispriced quotes, and (*ii*) winning the race to post updated quotes. We find strong support for these model predictions using data from the Toronto Stock Exchange. The identification strategy is based on an unanticipated news event in which the Twitter feed of the Associated Press falsely reported a terrorist attack.

JEL classification: G14, G15
Keywords: high-frequency trading, liquidity provision, speed heterogeneity

*We thank Timothy Bresnahan, Darrell Duffie, Liran Einav, Matthew Gentzkow, Joseph Grundfest, Laurie Hodrick, Jonathan Levin, Paul Milgrom, Monika Piazzesi, Peter Reiss, Bernhard Silli, an anonymous referee, and seminar participants at Stanford for useful comments. We also thank various industry experts for sharing their insights, as well as NASDAQ OMX and TMX Datalinx for providing the data and answering our questions. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number OCI-1053575. We acknowledge financial support by the Kohlhagen Fellowship Fund and the Kapnick Fellowship Program through grants to the Stanford Institute for Economic Policy Research. A previous version of this paper was entitled "The Private Value of Speed: High-Frequency Trading and the Supply of Liquidity."

†Sauder School of Business, University of British Columbia. E-mail: baldauf@mail.ubc.ca.

‡Kellogg School of Management, Northwestern University. E-mail: joshua.mollner@kellogg.northwestern.edu.

1 Introduction

In financial markets, liquidity provision is the activity of intermediating trade between agents who arrive to the market at different times. Typically, this is accomplished by posting non-marketable limit orders and then trading passively. Over the last two decades, liquidity provision has been reshaped by a general shift toward democratization. Liquidity provision had once been almost exclusively been done by specialists. However, that began to change with the advent of electronic markets. Today, direct market access allows traders to have very fine control over their limit orders, effectively enabling anyone to become a liquidity provider.

Electronic markets also engendered the rise of high-frequency trading, whereby a large portion of trading has become automated and geared toward progressively faster speeds. High-frequency trading has fundamentally reshaped financial markets over the last decade, and according to some estimates, it has grown to account for 49 percent of trading volume in the United States (TABB Group, 2014). Many of these high-frequency traders have taken advantage of the aforementioned democratization to become liquidity providers themselves.[1] These fundamental changes in the nature of liquidity provision raise a number of important questions, including "what are the advantages of faster liquidity providers?" and "what are their externalities upon slower liquidity providers?"

We answer these questions by providing a detailed and cleanly identified analysis of the strategic advantages that accrue to faster liquidity providers. In order to protect themselves from arbitrageurs, liquidity providers must constantly monitor the market and modify their quotes quickly in response to new information. Informative signals can be derived from a variety of sources including news about current events or price movements of related securities. It is intuitive that faster liquidity providers have an advantage in responding to news. In this paper we quantify this advantage and identify two channels through which it is derived. First, faster liquidity providers participate in fewer passive trades immediately after information arrives, when trades tend to be less profitable for the passive side. They achieve this by using their speed to cancel mispriced quotes before they can be exploited. Second, faster liquidity providers participate in relatively more trades in the subsequent period, when trades tend to be more profitable for the passive side. They achieve this by winning the race to post quotes at updated prices and thus obtaining better queue position.

[1]For example, Virtu Financial and KCG, two prominent high-frequency trading firms, describe themselves respectively as "the global leader in electronic market making" (Virtu, 2016), and "one of the largest independent market makers in U.S. equities" (KCG, 2016). Similarly, several academic studies have documented that many high-frequency traders are extensively involved in liquidity provision. (See, for example, the literature cited in the first paragraph of Section 2.)

To fix ideas, we first present a modeling framework, which is a straightforward application of Budish, Cramton and Shim (2015). In the model, traders may choose to upgrade their trading speed. While the upgrade is costly, faster traders are better able to respond to new information about the value of the asset. In equilibrium, both fast and slow traders provide liquidity. The model generates a set of hypotheses about how liquidity providers of differing speeds react to news. In testing them with data from the Toronto Stock Exchange (TSX), we exploit identification through an unanticipated and exogenous news event. On April 23, 2013 the Associated Press Twitter account, having been hacked by the Syrian Electronic Army, falsely announced a successful terrorist attack on the White House, which triggered an appreciable market response. Our analysis focuses on how liquidity providers of differing speeds reacted to this "news."

Our data identifies the participating organizations involved in each trade. We construct a measure of relative speed for each participating organization, which is based on its behavior during the morning before the event. Each participating organization may represent several clients, which makes this a moderately coarse level of aggregation. However, our empirical approach takes advantage of the existence of substantial variation in the composition of clientèle across across participating organizations. A "faster" participating organization is one whose clients tend to be fast, and a "slower" participating organization is one whose clients tend to be slow. Consistent with the model, we find empirical evidence of a bimodal speed distribution among traders who provide liquidity. We then investigate the consequences of this event for liquidity providers along this distribution.

As expected, speed conveys a benefit to liquidity providers in this event. We estimate that a faster participating organization achieves a relative increase in profits earned from passive trading of \$0.79 per second in the moments following the event. The size of this effect is economically significant, and is on the order of magnitude of the revenue that an average liquidity provider derives from passive trading. We then isolate the channels responsible for this effect. First, we find that faster participating organizations take the passive side of relatively fewer trades immediately after the event, and that these trades tend to be less profitable for the passive side. Second, we show that faster participating organizations take the passive side of relatively more trades in the subsequent period, and that these trades tend to be more profitable for the passive side. Specifically, a faster participating organization takes the passive side of the trades of 11.55 fewer lots per second immediately after the event and of 46.29 additional lots per second during the subsequent period, controlling for the level of passive trading prior to the event.[2] Furthermore, a lot traded immediately after

[2] In this paper, we measure traded volume in terms of board lots. For most securities, the TSX defines a board lot as 100 shares. However, for securities trading at less than \$0.10 per share, a board lot is 1000

the event yields an additional loss to passive side of \$0.02 per lot traded compared to trades made before the news arrived, whereas a lot traded during the subsequent period yields an additional profit of \$0.02.

This analysis highlights a particular externality that high-frequency traders exert upon slower traders. Faster traders effectively "cream-skim" liquidity provision, participating relatively less often when trades tend to be less profitable for the passive side and participating relatively more often when trades tend to be more profitable for the passive side. As a result, the liquidity provision of slower traders tends to be at worse terms. In particular, faster liquidity providers are better able to cancel their mispriced quotes when information arrives, which shifts the adverse selection burden onto their slower counterparts—or, in the language of the industry, leaves them to be "overrun."

Certain participating organizations on TSX are labeled designated market makers, a role intended to ensure that liquidity will be provided, even during market stress. One might expect this contractual obligation to affect the nature of liquidity provision. Interestingly, after controlling for trading speed, we find no systematic evidence of this.

The remainder of the paper is organized as follows. Section 2 discusses related literature. Section 3 presents a model of liquidity provision. Section 4 describes the news event we use for identification. Section 5 describes the data and defines the variables used in the analysis. Section 6 presents our main results, which support the predictions of theory. Section 7 discusses our findings and concludes. The appendix contains proofs, technical details, robustness checks, and additional analysis.

2 Related Literature

This paper is connected to several strands of the growing empirical literature on high-frequency trading. One branch of this literature has investigated how high-frequency traders take advantage of their superior technologies. Menkveld (2013) studies the behavior of a high-frequency trading firm after Chi-X opens a new exchange in Europe. The firm is primarily a provider of liquidity, and his analysis suggest that its technological advantages may be crucial for managing inventory across markets. Brogaard, Hendershott and Riordan (2014a) show that when high-frequency traders take liquidity, they profit by trading in the direction of permanent price changes and in the opposite direction of transitory pricing errors. This suggests that these traders may be obtaining an informational advantage, which they monetize by taking liquidity. Baron, Brogaard, Hagströmer and Kirilenko (2016) find that successful high-frequency traders use their technology to succeed in winner-take-all

shares, and for securities trading between \$0.10 and \$1.00 per share, a board lot is 500 shares.

races to take liquidity. Hendershott and Riordan (2013) find that high-frequency traders use their technology to monitor the market closely and to tailor their strategies to current market conditions, taking liquidity when it is plentiful and providing it when it is scarce. Carrion (2013) argues that high-frequency traders possess the ability to time the market. In this paper, we identify two advantages that high-frequency traders obtain from their speed: they are able to avoid trading at stale prices by canceling mispriced quotes, and they are able to trade more at updated prices by winning the race for queue position.

We also find that high-frequency liquidity providers impose a negative externality on their slower counterparts: because of their superior ability to cancel mispriced quotes, faster traders can shift the adverse selection burden onto slower ones. This finding complements a number of other papers, which also consider the externalities that high-frequency traders exert on other classes of traders. Narrow spreads are a positive externality for small, retail traders, and a large empirical literature has established a link between high-frequency trading and narrower bid-ask spreads. Evidence has been found across a wide variety of contexts, including Chi-X and Euronext (Menkveld, 2013); NYSE (Hendershott, Jones and Menkveld, 2011); NASDAQ OMX Stockholm (Brogaard, Hagströmer, Nordén and Riordan, 2015); NASDAQ (Hasbrouck and Saar, 2013); the Tokyo Stock Exchange (Conrad, Wahal and Xiang, 2015); and TSX (Malinova, Park and Riordan, 2016). A great deal of research has also linked high-frequency trading to more efficient prices (Carrion, 2013; Hendershott, Jones and Menkveld, 2011; Brogaard, Hendershott and Riordan, 2014a; Chaboud, Chiquoine, Hjalmarsson and Vega, 2014). In contrast, a host of other papers find that high-frequency trading exerts negative externalities, either in the form of increased institutional trading costs (Tong, 2015; Korajczyk and Murphy, 2016); slowing the markets via quote stuffing (Egginton, Van Ness and Van Ness, 2016; Gai, Yao and Ye, 2013); or by exacerbating high-volatility events (Easley, Lopez de Prado and O'Hara, 2011; Kirilenko, Kyle, Samadi and Tuzun, forthcoming; Huh, 2014).[3]

This paper is also related to the theoretical literature on trading in limit order books, seminal members of which include Copeland and Galai (1983), Glosten and Milgrom (1985), Glosten (1994), and Foucault (1999). In addition there is a more recent theoretical literature on high-frequency trading, which includes Foucault, Hombert and Roşu (2016), Roşu (2016), Biais, Foucault and Moinas (2015), and Baldauf and Mollner (2015). Particularly related is Budish, Cramton and Shim (2015), whose theoretical framework we adopt for our model.

[3] Although the discussion in this paragraph reflects the conclusions of the bulk of the literature, there are some exceptions. For example, Menkveld and Zoican (forthcoming) associate high-frequency trading with wider spreads; Brogaard, Hendershott, Hunt and Ysusi (2014b) find no systematic link between high-frequency trading and institutional trading costs; and Hasbrouck (2015) rejects quote stuffing and instead explains fleeting orders with Edgeworth cycles.

In all these models, markets are made by either a single dealer, or by a homogeneous group of them. We, however, are interested in what transpires when liquidity providers have heterogeneous speeds. More closely related to our paper in that respect is Hoffmann (2014), who extends the Foucault (1999) model to include traders of heterogeneous speeds. In his model, like ours, two effects occur following information events: (*i*) fast traders provide relatively less liquidity at stale prices than their slower counterparts, and (*ii*) fast traders provide relatively more liquidity at updated prices. However, our model is able to make finer predictions about when these two effects occur. In line with the data, our model predicts the first effect to manifest in the moments following the information event and the second effect to manifest in the subsequent period.

3 Model

In this section we present a theoretical framework, which generates several testable implications about how fast and slow liquidity providers are affected by the arrival of public information. In the model, "fast" traders obtain this information sooner, and they represent high-frequency traders. "Slow" traders do not obtain this information until after some lag, and they represent other traders.

The model is almost precisely that of Budish, Cramton and Shim (2015), with the only relevant point of divergence being in respect to the selection of equilibrium. They focus on an equilibrium in which only fast traders provide liquidity. We instead identify and study a broader class of equilibria in which liquidity is provided by fast traders, *together with varying numbers of slow traders.* While this is only a small extension of their approach, it is useful because it allows us to match the heterogeneity in trading speed among liquidity providers observed in the data and to model the implications of that heterogeneity.

3.1 Environment

There is a single security, which is traded in a limit order book by two types of agents: investors and market makers. Time is continuous.

Security. There is a single security whose fundamental value at time t is v_t. Trading begins at $t = 0$, at which point the fundamental value v_0 is common knowledge. From then on v_t evolves as a compound Poisson jump process with arrival rate $\lambda_{jump} \in \mathbb{R}_+$ and jump distribution F_{jump}. We assume F_{jump} has finite bounded support and is symmetric with mean zero.

Limit order book. Shares of the security are traded in a limit order book. Orders are processed serially in the order received. In the event two orders are received simultaneously, the tie is broken at random (i.e. by random latency). We assume that order flow is non-anonymous and publicly observed in real time.

In what follows, we refer to four types of orders. A *limit order* consists of a buy/sell designation, a limit price, the maximum quantity desired to transact, and a time until when the order stays in force. An *immediate-or-cancel (IOC) order* is a limit order with a time in force of zero. A *market order* may be thought of as an IOC order with a limit price of positive or negative infinity. Traders may also submit *cancel orders*, which instruct the exchange to remove previously placed limit orders from the book.

Investors. The first type of traders are investors, who arrive at the Poisson rate $\lambda_{invest} \in \mathbb{R}_+$. Investors are non-strategic. Upon arrival, they submit market orders to buy or sell one share of the security, each with probability one half.

Market makers. The remaining traders are market makers. There are an infinite number of them, each with the objective of maximizing profits. Market makers costlessly observe jumps in v_t with a lag of $\delta > 0$, but by paying to be fast, a market maker may observe these jumps without lag. The flow rate cost of being fast is c_{speed}.

3.2 Equilibrium

In this section, we identify a class of stationary equilibria in which the bid-ask spread, s, is defined implicitly as the solution to the following equation, where J denotes a random variable drawn from F_{jump}.[4]

$$\lambda_{jump} \cdot \Pr\left(|J| > \frac{s}{2}\right) \cdot \mathbb{E}\left(|J| - \frac{s}{2} \;\middle|\; |J| > \frac{s}{2}\right) = \lambda_{invest} \cdot \frac{s}{2}. \tag{1}$$

All market makers earn zero profits in the equilibria we identify, and in fact equation (1) can be interpreted as a flow rate zero profit condition for a slow market maker who is currently quoting the best bid and the best ask. It equates the revenues from trade with investors (RHS) to the costs of adverse selection from trade with fast market makers (LHS).[5]

[4] A solution to (1) exists as a consequence of the intermediate value theorem.

[5] Consider a slow market maker who owns the best bid and the best ask. The arrival rate of investors is λ_{invest}. Conditional on the arrival of an investor, she earns $\frac{s}{2}$. The arrival rate of jumps in v_t is λ_{jump}. Conditional on a jump arriving, she will lose $\Pr\left(|J| > \frac{s}{2}\right) \cdot \mathbb{E}\left(|J| - \frac{s}{2} \;\middle|\; |J| > \frac{s}{2}\right)$ to fast investors who "snipe" her stale quotes.

At any time t, the action space of a market maker includes (i) a choice of whether to pay to be fast, and (ii) whether to submit any limit orders or cancel orders. Market makers therefore have a wide set of strategies available to them. However, in what follows, we will make special reference to the three particular strategies defined below.

Strategy 1: Never pay to be fast. At time zero, submit a limit order to buy one share at $b_0 = v_0 - \frac{s}{2}$ and a limit order to sell one share at $a_0 = v_0 + \frac{s}{2}$. If either order does not receive the first queue position, then cancel it immediately. If at any time t one of the standing quotes is filled by an investor, then submit a "replacement" limit order, which is canceled immediately if it does not receive the first queue position. If at any time t, other market makers submit quotes around a new price level, then cancel any standing quotes.

Strategy 2: Pay to be fast at all times. At time zero, submit a limit order to buy one share at $b_0 = v_0 - \frac{s}{2}$ and a limit order to sell one share at $a_0 = v_0 + \frac{s}{2}$. If either order does not receive the first queue position, then cancel it immediately. If at any time t one of the standing quotes is filled by an investor, then submit a "replacement" limit order, which is canceled immediately if it does not receive the first queue position. If at any time t a jump in v_t occurs, then (i) cancel any active quotes; (ii) if $v_{t^+} > a_{t^-}$, then submit an IOC order to buy one share at price a_{t^-}; (iii) if $v_{t^+} < b_{t^-}$, then submit an IOC order to sell one share at price b_{t^-}; (iv) submit a limit order to buy one share at $b_{t^+} = v_{t^+} - \frac{s}{2}$ and a limit order to sell one share at $a_{t^+} = v_{t^+} + \frac{s}{2}$. If either of the latter orders does not receive the first queue position, then cancel it immediately.

Strategy 3: Never pay to be fast. Never submit any orders.

We are now prepared to state the following result concerning the existence of stationary equilibria in which both fast (Strategy 2) and slow (Strategy 1) market makers engage in liquidity provision. For convenience, we do not restrict the number of market makers who use a particular strategy to be a natural number.[6] Finally, as a point of clarification, although market makers use only the three aforementioned strategies on path in these equilibria, potential deviations are not restricted in that way.[7]

Proposition 1. *For any $N_{slow} \in \mathbb{R}_+$, it is an equilibrium for N_{slow} market makers to use Strategy 1, $N_{fast} = \frac{\lambda_{invest}}{c_{speed}} \cdot \frac{s}{2}$ market makers to use Strategy 2, and the remaining market makers to use Strategy 3.*

[6]Thereby we follow the long tradition of entry models, e.g. Panzar and Rosse (1987). Budish, Cramton and Shim (2015) also use this simplification.

[7]In other words, these are equilibria of the full game, as opposed to only equilibria of an auxiliary game in which market makers are restricted to Strategies 1, 2, and 3.

Proof. In Appendix A.1.

This result identifies a class of equilibria in which any number of slow market makers may be active in liquidity provision. The analogue of the equilibrium identified by Budish, Cramton and Shim (2015) is the special case in which no slow market makers are active (i.e. $N_{slow} = 0$). While this class of equilibria can be thought of as generalizing the Budish, Cramton and Shim (2015) equilibrium in a somewhat natural way, still more equilibria exist.

Notice that in these equilibria, the active market makers (those who use Strategies 1 and 2) often cancel limit orders immediately after placing them. This provides an equilibrium justification for fleeting orders, whose prevalence has been documented empirically (Hasbrouck and Saar, 2009; Gai, Yao and Ye, 2013). In this model, fleeting orders are neither the result of nefarious trading practices, nor do they destabilize the market. Rather, they are a natural consequence of competition among liquidity providers. Fleeting orders are also an equilibrium feature in Baruch and Glosten (2013) and Hasbrouck (2015).

3.3 Testable Implications

The model generates testable hypotheses about how fast and slow liquidity providers are affected by the arrival of a piece of information. In describing these hypotheses, we consider three time periods around the information event: (*i*) the equilibrium steady state, which we refer to as the *pre period*, (*ii*) the instant of the jump, which we refer to as the *jump period*, and (*iii*) the time interval immediately after the jump, which we refer to as the *post period*. The pre period serves as a benchmark for the analysis against which the jump period and the post period are compared.

We denote the number of passive trades that a trader $i \in \{fast, slow\}$ makes during time period $t \in \{pre, jump, post\}$ by $q_{i,t}$. We also use $Q_t = N_{fast} \cdot q_{fast,t} + N_{slow} \cdot q_{slow,t}$ to denote the total volume traded during time period t. Similarly, we use Π_t to denote the total trading profits (gross of speed costs) earned from passive trading during time period t.

Proposition 2. *In any of the equilibria described in Proposition 1,*

(i) $\frac{q_{slow,jump}}{q_{fast,jump}} \geq \frac{q_{slow,pre}}{q_{fast,pre}}$,

(ii) $\frac{q_{slow,post}}{q_{fast,post}} \leq \frac{q_{slow,pre}}{q_{fast,pre}}$,

(iii) $\frac{\Pi_{jump}}{Q_{jump}} < \frac{\Pi_{pre}}{Q_{pre}}$, *and*

(iv) $\frac{\Pi_{post}}{Q_{post}} \geq \frac{\Pi_{pre}}{Q_{pre}}$.

Proof. In Appendix A.2.

In words, the model generates the following predictions: (*i*) the number of passive trades increases from the pre period to the jump period by relatively less for fast traders than for slow traders, (*ii*) the number of passive trades increases from the pre period to the post period by relatively more for fast traders than for slow traders, (*iii*) passive-side profits per trade are lower in the jump period than in the pre period, and (*iv*) passive-side profits per trade are higher in the post period than in the pre period.

More succinctly, fast traders have two advantages over their slower counterparts in liquidity provision. First, they participate in fewer passive trades during the jump period, when trades tend to be less profitable for the passive side. Second, they participate in more trades during the post period, when trades tend to be more profitable for the passive side.

The intuition for these results is as follows. Quotes become mispriced during the jump period, and trades made in this period are initiated by traders who are exploiting these mispriced quotes. Trades during the jump period therefore tend to be less profitable for the passive side (i.e. part (*iii*) of the proposition). Due to their speed, fast traders have a better chance of canceling misplaced quotes and are therefore better at avoiding this adverse selection. They therefore trade relatively fewer shares during the jump period (i.e. part (*i*) of the proposition). Fast traders can also use their speed to better position themselves in the queue as the market settles at a new level. They therefore trade relatively more shares during the post period (i.e. part (*ii*) of the proposition). Finally, since the mix of liquidity provision in the post period shifts toward fast traders, who are better at avoiding adverse selection, trades tend to be more profitable for the passive side in the post period (i.e. part (*iv*) of the proposition).

4 Event Description

On April 23, 2013, at 13:07:50 EDT,[8] the Associated Press Twitter account, having fallen victim to a cyber attack, issued the following post:

> Breaking: Two Explosions in the White House and Barack Obama is injured (@AP).

The hacker collective Syrian Electronic Army (SEA) eventually claimed responsibility for the attack and was subsequently investigated by the FBI Cyber Division for it.

In the minutes following the tweet, the market responded with a drastic decline on average, which has sometimes been referred to since as the "Hack Crash" or the "Hash Crash." We illustrate the magnitude of the market's response using market data on the

[8] In all that follows, times will be expressed in EDT.

iShares Core S&P 500 Index ETF (CAD-Hedged), which has the ticker symbol XSP. This ETF tracks a version of the S&P 500 index that is hedged for exposure to the CAD-USD exchange rate. This security is chosen because it is the largest ETF tracking the S&P 500 index that is traded at TSX, and therefore an appropriate summary of the market. As Figure 3(a) shows, its mid price dropped by 95 bps in just over two minutes following the tweet, from 18.085 CAD to 17.915 CAD.[9] It recovered equally quickly thereafter, following announcements that the Twitter post was a hoax.

Two tweets are considered to have given the first definite confirmation of a hoax. The first of these tweets was sent at 13:10:02 by Associated Press writer Sam Hananel and read

> Please Ignore AP Tweet on explosions, we've been hacked (@SamHananelAP).

The second was sent at 13:10:06 by Evan McMorris-Santoro, a White House correspondent who retweeted the original message, appending

> from here in the WH basement, this acct seems hacked (@EvanMcSan).

In addition to the price movement triggered by the tweet, bid-ask spreads were also affected. As Figure 1(b) shows, the spread of XSP increases immediately after the the tweet was posted, likely in response to the additional uncertainty that had been injected into the market. Prior to the tweet, the spread was nearly always at the one cent minimum, but was occasionally higher afterwards.

[9] In all that follows, money will be expressed in CAD unless otherwise noted.

Figure 1: Market response to event, S&P 500 Index ETF TSX:XSP

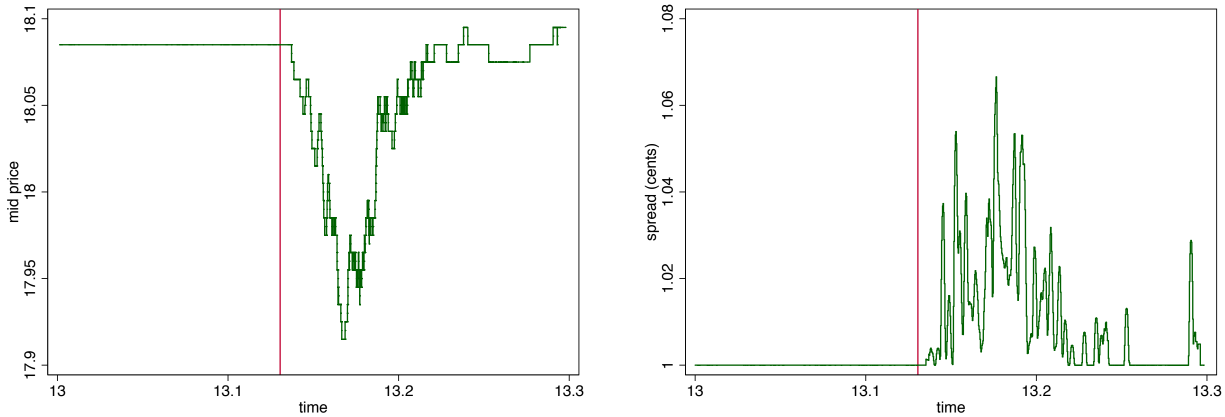

(a) Mid price

(b) Spread

In both panels, the horizontal axis is time on April 23, 2013 measured in hours after midnight EDT. The vertical line at 13:07:50 marks the published time of the tweet. The left panel plots, for each second, the time-weighted average mid price of XSP on TSX. For the right panel, we compute the spread of XSP on TSX measured in cents, in every second, where we take a simple average if multiple spreads prevailed in a given second. The right panel then plots a lowess over these observations using a 72 second bandwidth.

Our analysis focuses on the decline phase, the time period prior to the trough. This phase is relatively clean: the initial tweet was exogenous, unanticipated, and at a time that can be pinpointed. However, the recovery phase is much less straightforward to analyze: two conflicting pieces of information were in play during this phase because some traders were still reacting on the basis of the initial tweet.

This event provides a unique opportunity for studying how the algorithms of fast and slow liquidity providers respond to the sorts of information events that normally occur hundreds or thousands of times throughout the trading day. The event possesses five advantageous features. First, the event was arguably exogenous. Second, the timing of the event can be pinpointed ex post—which is advantageous econometrically—yet, like most information events, its timing was unknown ex ante. Third, the resulting price fluctuations were neither too small nor too large. Smaller fluctuations might have been difficult to detect in the data. On the other hand, larger fluctuations would likely have triggered either circuit breakers at the exchange or the safety thresholds of the trading algorithms in play, resulting in behavior that was not representative of how traders respond to typical information events. Fourth, the event was short. A longer event would have given humans more time to intervene, which would also make the event unrepresentative. An fifth advantage of this event is that it had a very different impact across stocks, affecting some negatively, some positively, and leaving some prices flat. This heterogeneity allows us to test a richer range of hypotheses.

There are also a number of other events to which the literature has paid attention. Many papers have studied how markets absorb news around scheduled macro policy announcements such as Federal Open Market Committee (FOMC) announcements. While the timing of these events can be pinpointed ex post their timing is also known ex ante, which gives rise to anomalies surrounding these events and makes them less suitable for our purposes.[10] On the other hand, a number of other papers study drastic events such as the 2010 Flash Crash, which saw a collapse of historic proportions and subsequent rebound of the entire United States market over the course of about a half hour. This event is less suitable for our purposes because it is endogenous, relatively lengthy, and resulted in price swings that were extremely large.

5 Data

We obtain trade and quote data from the Toronto Stock Exchange (TSX) for the day of April 23, 2013. This section introduces this dataset and also defines several variables that are used in the analysis.

5.1 Toronto Stock Exchange Dataset

TSX provides us with a complete record of trades and a nearly complete record of quotes for all securities that were traded on April 23, 2013.[11] For each trade, we observe the quantity traded, the transaction price, and identifiers for the buying and selling participating organizations (if they choose to remain non-anonymous). Each quote includes the bid and the ask prices, as well as the number of shares that can be traded at these prices (top-coded at 999 lots). Trades and quotes are recorded chronologically and time-stamped to the second.

An advantage of using Canadian data is that, given the dominance of TSX in the Canadian market, cross-exchange strategies are less important than in American markets.[12] Nevertheless, many of the same phenomena we document using TSX data would also be true for American markets. In Appendix E, we repeat a portion of the analyses of this paper on

[10] In particular, since liquidity providers can anticipate the timing of these announcements, they can protect themselves from adverse selection by withdrawing their quotes just before the information arrives. Evidence of liquidity providers doing exactly this is documented by Fleming and Remolona (1999) and Balduzzi, Elton and Green (2001) for the treasury market, by Kavajecz (1999) for equities, and more recently by Rosa (2016) for the E-Mini. All establish that spreads widen and depth decreases in anticipation of these announcements.

[11] The dataset occasionally fails to report a change in the quote for a security in the event that another change follows shortly thereafter.

[12] In April 2013, TSX accounted for 59% of value traded and 47% of volume (IIROC, 2015). This implies a much higher concentration than the American market where in April 2013, the leading exchange, NASDAQ, accounted for only 18% of value traded and 16% of volume (BATS, 2016).

NASDAQ data.

5.1.1 Participating Organizations

On April 23, 2013, 89 of the participating organizations on TSX identified themselves as a party in at least one trade. TSX does allow participating organizations to conceal their identities and trade anonymously.[13] For the purposes of the analysis that follows, we treat "anonymous" as its own participating organization. This yields a total of 90 participating organizations.

In Table 11, we provide summary statistics pertaining to the activity of those participating organizations on April 23, 2013. Apparent from this table is the fact that the distribution of trader size is heavily skewed. CIBC World Markets, Inc. is the most active player, accounting for roughly one quarter of trading activity. Other large players include TD Securities Inc., Dundee Securities Ltd., RBC Capital Markets, and National Bank Financial Inc.

Note that each participating organization may represent multiple clients. For the large banks mentioned above, these clients may include retail traders, institutional traders, as well as high-frequency trading firms whose direct market access they sponsor. Unfortunately, our data do not allow us to observe the activities of each separate client. However, our empirical approach takes advantage of variation in the composition of clientèle across participating organizations. To interpret our findings, we assume that a trader's choice of a participating organization is not systematically related to the tweet. We return to discuss this issue in more detail in Section 7.1.

5.1.2 Securities

Our sample consists of all securities that were either quoted or traded on TSX on April 23, 2013. This includes stocks and exchange traded funds, as well as warrants, debentures, and investment trusts.[14] In total, our sample contains 2367 securities.

Summary statistics pertaining to the traded volume and the bid-ask spread of these securities are provided in Table 12.

5.2 Variables Used in Analysis

This section provides details concerning the definitions of variables that we construct for use in the analysis. Notation is summarized in Table 13.

[13] In the TSX data, we observe the participating organization for 86.45 percent of the lots traded in the pre period (where that calculation counts each lot twice: once per side). The corresponding figures for the jump and post periods are 82.74 percent and 88.28 percent.

[14] In the event that a company issues stock in different share classes, we treat each as a separate security.

5.2.1 Time Periods

We define three time periods, the empirical analogues of those defined in Section 3. The pre period comprises the morning of the trading day until the news event. The jump period begins just prior to the moment when the information first becomes available and ends once fast traders have had an opportunity to react. The post period follows the jump period, and it ends slightly before the tweet was reported to be a hoax by Sam Hananel and Evan McMorris-Santoro.

Specifically, we define the jump period to begin at 13:07:40 and to end at 13:08:10.[15] The pre period covers all time from the opening of markets at 9:30:00 until the beginning of the jump period. The post period covers all time from the end of the jump period until 13:10:00.

Formally, for a second t, we define indicators pre_t, jump_t, and post_t equal to one if t lies in either the pre, jump, or post period, respectively. Additionally, for second t, we define $(\text{jump} + \text{post})_t$ as an indicator equal to one if t lies either in the jump or post period. Similarly, for the trade of a particular lot l, we define indicators pre_l, jump_l, and post_l equal to one if the trade occurred in either the pre, jump, or post period, respectively. We define the indicator $(\text{jump}+\text{post})_l$ equal to one if a lot was traded in either the jump or post period.

Our main specifications involve discrete time periods because they are the empirical analogues of the time periods in the model. However, in Appendix C.2, we verify the robustness of our findings using specifications in which time enters linearly.

5.2.2 Directional Response Classification

For each security, we calculate the prevailing mid price at the beginning of the jump period, as well as the prevailing mid price at the end of the post period. Depending on how these mid prices change, we classify the security as either *down* (in the case of a negative price change), *flat* (in the case of no price change), or *up* (in the case of a positive price change).[16]

In some cases, one or both of these mid prices does not exist, which is the case if there are either no resting limit orders to buy or no resting limit orders to sell. We do not classify such securities into any of these three categories.

489 securities have a negative price change. 1405 securities have no price change. 219 securities have a positive price change. The remaining 254 securities cannot be classified.

[15] There is some ambiguity regarding the alignment of timestamps used by TSX and Twitter. We err on the side of beginning the jump period too early, in order to guarantee that this period contains the moment at which the information became available. If the jump period includes trades that belong in the pre period, then this would merely cause attenuation in our results.

[16] Mining stocks, for example, tended to be positively affected during this episode, while banking and finance stocks tended to be negatively affected.

5.2.3 Passive-Side Profits

The TSX data do not distinguish which participant in a trade was the passive party (i.e. the participating organization whose order had been resting in the limit order book) and which was the aggressive party (i.e. the participating organization whose order triggers the trade). In order to study the profits derived from liquidity provision, we must first perform this classification. We do this using a modified version of the Lee and Ready (1991) algorithm. The primary difference is that we do not rely upon statistical assumptions to classify trades taking place within the spread. See Appendix B for a more detailed discussion of this classification.

For all traded lots that we classify either as passive buys or as passive sells, we compute the profits earned by the passive side of the trade. For the trade of lot l, let $x(l)$ denote the security that is traded, let p_l denote the price at which the trade occurs, and let $||l||$ denote the number of shares in the lot. We calculate profits using marked-to-market accounting, where inventory is valued at $m_{x(l),13:10:00}$, the mid price at the end of the post period. Because we are interested in understanding how traders react to and profit from real information, we mark profits out to a point in time before the market learned that the information contained in the tweet had been false. Specifically, we calculate passive-side profits π_l from the trade of lot l as:

$$\pi_l = \begin{cases} (m_{x(l),13:10:00} - p_l) \cdot ||l|| & \text{if } l \text{ was a passive buy} \\ (p_l - m_{x(l),13:10:00}) \cdot ||l|| & \text{if } l \text{ was a passive sell} \end{cases}$$

We use $\pi_{i,t}$ to denote participating organization i's profits (in dollars) from passive trades in second t. We also use $q_{i,t}$ to denote the number of lots passively traded by participating organization i in second t.

5.2.4 Speed Index

While we do not directly observe trading speeds, we can infer relative speeds from the liquidity provision behavior of participating organizations. To motivate our measure, consider how traders would react to positive news about the value of a security. Traders currently offering to sell would attempt to cancel their mispriced quotes, and other traders might attempt to trade at the stale ask. More important for our purposes, traders would also race to establish queue position at a new bid, which would form a new level of the limit order book. Faster liquidity providers are more likely to win these races, and they therefore facilitate a larger fraction of any subsequent trades that occur at the new bid.[17] The same is true for ask-side

[17] In contrast, depth would typically have already existed at the new ask. Subsequent ask-side trades are therefore less informative about speed.

trades that follow negative news.

Motivated by this, we define our speed index based on a participating organization's share of liquidity provision at the new level following these types of episodes. First, we identify all episodes in which the best bid ticks up or the best ask ticks down on any TSX security during the pre period. Second, we identify the first 100 lots traded at the new bid in the case of a bid increase, or the first 100 lots traded at the new ask in the case of an ask decrease. (In case the bid or ask changed again before 100 lots were traded, then we use the actual number of lots traded at the new level.) Third, we calculate, for each participating organization, the share of these lots in which they were the passive party. Finally, we create the variable speedIndex by applying min-max normalization to these shares.

Figure 2 plots, by the speed index of the participating organization on the passive side, the distribution of lots traded in the pre period across all TSX securities. The distribution of speed is bimodal, with the slowest participating organizations (i.e. speedIndex ≤ 0.2) and the fastest participating organizations (i.e. speedIndex ≥ 0.8) accounting for 46.1 percent and 38.1 percent of liquidity provision, respectively.[18] This distribution is in line with other estimates of high-frequency trading activity.[19]

The interquartile range of this distribution is 0.70. We interpret this range as the difference in speedIndex between a slower liquidity provider and a faster liquidity provider. In the analysis that follows, we evaluate the magnitude of our estimates using this range.

[18] However, faster participating organizations are more active on a per-participant basis. Of the 90 participating organizations on TSX, 86 have speedIndex ≤ 0.2, and 2 have speedIndex ≥ 0.8.

[19] In the United States, high-frequency trading is estimated to have accounted for 49 percent of volume in 2013 (TABB Group, 2014).

Figure 2: Histogram of trades by passive party speed

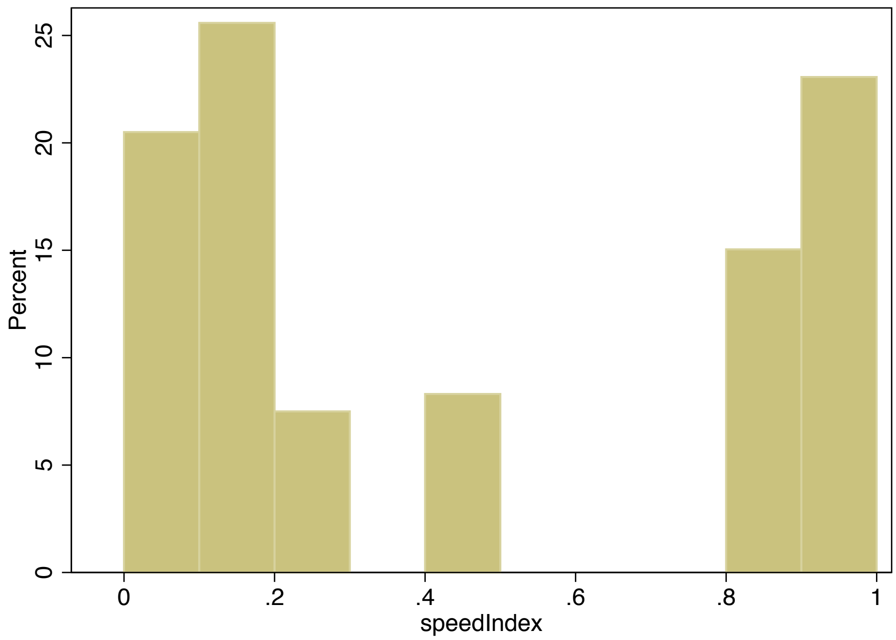

Histogram of lots traded on TSX between 9:30:00 and 13:07:40 on April 23, 2013, plotted by the speed index of the passive party, which is computed as described in Section 5.2.4.

To establish that our results are robust to alternative measurements of speed, in Appendix C.1 we re-run our main specifications using variants of the above definition. Moreover, in Appendix D, we demonstrate that our speed index correlates well with other measures of speed that are already established in the literature.

5.2.5 Designated Market Makers

While any participating organization can serve as a liquidity provider by submitting limit orders that reside in the order book, TSX identifies certain participating organizations as designated market makers.[20] The responsibilities of designated market makers are geared toward ensuring liquidity of the market. In particular, market makers are required to ensure that, for their securities of responsibility, the time-weighted average spread is below some agreed upon cutoff (TMX Group, 2012a). For participating organization i, we define DMM_i as an indicator equal to one if TSX lists i as a designated market maker.

[20] We obtain the list of designated market makers from TMX Group (2013).

6 Main Results

Faster liquidity providers have an advantage in this event. The model predicts two sources for their additional profits. First, faster traders are relatively better able to avoid providing liquidity in the jump period, when trades tend to be less profitable for the passive side. Second, they are relatively better able to position themselves to provide liquidity in the post period, when trades tend to be more profitable for the passive side. In this section, we document strong support for each of these predictions.

Section 6.1 begins by quantifying the additional profits that accrue to faster liquidity providers. To further understand the source of these additional profits, we test for the specific predictions of the model. Section 6.2 finds support for the predictions concerning the number of passive trades made by faster and slower participating organizations. Section 6.3 finds support for the predictions concerning the profitability of trades for the passive party.

6.1 Passive-Side Profits per Second

We start by establishing that faster traders earn relatively more profits from passive trading in the aftermath of the tweet. To that end, we estimate the following equation, where the dependent variable $\pi_{i,t}$ is the profits (in dollars) earned by participating organization i from its passive trades in second t.

$$\begin{aligned} \pi_{i,t} = \alpha_i + \beta_1 \cdot (\text{jump} + \text{post})_t + \beta_2 \cdot \text{speedIndex}_i \times (\text{jump} + \text{post})_t \\ + \beta_3 \cdot \text{DMM}_i \times (\text{jump} + \text{post})_t + \varepsilon_{i,t} \end{aligned} \tag{2}$$

This specification includes participating organization fixed effects, which allow us to control for participating organization behavior on that trading day (for example, whether the participating organization tended to be primarily buying or selling). We compute profits with respect to four different samples of securities. First, we compute them using all securities traded on TSX. These results are displayed in column (1) of Table 1. We then re-estimate the relationship, computing profits using only TSX securities whose prices are negatively affected, unaffected, and positively affected by the tweet. Those results are reported in columns (2)–(4) of the table.

Table 1: Passive-side profits over speed and time (in dollars per second)

	(1) All	(2) Down	(3) Flat	(4) Up
jump + post	-0.0327	-0.0165	0.00443	-0.0202
	(0.0246)	(0.0215)	(0.00304)	(0.0123)
speedIndex × (jump + post)	1.130^{+}	0.448**	-0.0362	0.721**
	(0.364)	(0.196)	(0.0390)	(0.316)
DMM × (jump + post)	0.166	0.0131	0.00190	0.148
	(0.194)	(0.0757)	(0.00884)	(0.143)
Observations	1174889	1174889	1174889	1174889
Participating organization FE	YES	YES	YES	YES

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, $^{+}$ $p < 0.005$, $^{++}$ $p < 0.001$
Coefficients are from OLS regressions with standard errors clustered by participating organization in parentheses. The dependent variables, profits of a participating organization from passive trades in a given second, are computed using marked-to-market accounting as described in Section 5.2.3, and are measured in dollars. An observation is a (participating organization, second) pair. The sample includes all seconds in the interval [09:30:00,13:10:00]. In column (1), we include all TSX securities. In columns (2)–(4), we include only TSX securities whose mid prices decrease, remain constant, or increase between 13:07:40 and 13:10:00, respectively. The variable DMM is an indicator for whether the participating organization is a designated market maker on TSX. The variable speedIndex is computed as described in Section 5.2.4. The variable (jump + post) is an indicator for the time period [13:07:40,13:10:00].

Table 1 demonstrates that faster liquidity providers have an advantage in how they respond to the arrival of information during this event. We would expect $\text{speedIndex}_i \times (\text{jump} + \text{post})_t$ to enter positively (i.e. $\beta_2 > 0$) for samples including securities whose prices are affected by the tweet. As expected, we find positive and significant estimates in columns (1), (2), and (4). The magnitude in column (1) can be interpreted as follows. An increase in speedIndex of 0.70 (an amount equal to the quantity-weighted interquartile range of this variable) is associated with an additional increase in profits per second from the pre period to the post period of \$0.79 per second. Moreover, as expected, we do not find the effect to be present for the flat securities on which column (3) is estimated.

In order to put this estimate in perspective, we compute an upper bound on the profits that a typical participating organization might obtain from liquidity provision in a typical second. For each trade taking place in the pre period that we were able to classify as either a passive buy or a passive sell (*cf.* Appendix B), we compute the half-spread and maker rebate earned by the passive party. Then for each second in the pre period, we compute aggregate revenue from liquidity provision as the sum of the half-spreads and rebates earned by passive parties in that second.[21] This is an upper bound on the aggregate profits from liquidity provision, since it does not account for the costs of adverse selection (i.e. the fact that passive parties tend to lose money on their inventory). In the median pre period second, the aggregate revenue from liquidity provision was \$63.89. Dividing by the 90 participating organizations in our sample, we conclude that a typical participating organization might obtain no more than \$0.71 from liquidity provision in a typical second. This is the same order of magnitude as the effect that we estimated above. In other words, speed is a vital aspect of liquidity provision: during information events, the additional profits from liquidity provision that accrue to faster participating organizations are roughly equal to an upper bound on the baseline profits from that activity.

In Sections 6.2 and 6.3, we delve deeper into this relationship, examining the channels through which faster liquidity providers obtain the additional profits quantified here. In doing so, we find support for the predictions of the model: faster participating organizations make fewer passive trades during the jump period, when trades tend to be less profitable for the passive side, and they make more passive trades during the post period, when trades tend to be more profitable for the passive side.

6.2 Quantity of Passive Trading

The model makes two predictions about how faster and slower participating organizations differ in terms of the number of passive trades they make in the time before, during, and after information arrives. Proposition 2(*i*) predicts that the change from the pre period to the jump period in quantity passively traded should be relatively lower for faster participating organizations than for slower ones. Intuitively, faster traders are able to use their advantage to avoid trade in the jump period by canceling mispriced quotes before they can be exploited. In contrast, Proposition 2(*ii*) predicts that the change from the pre period to the *post* period in quantity passively traded should be relatively *higher* for faster participating organizations than for slower ones. Intuitively, faster traders are able to use their advantage to better position themselves in the queue at the new prices levels.

To test whether these predictions are borne out in the data, we estimate the following equation, where the dependent variable $q_{i,t}$ is the number of lots passively traded by

[21] When the trade price is \$1 or more, the maker rebate is \$0.0031 per share. (There is also an \$0.0001 paid if the participating organization is in a certain rewards tier, which we do not have the data to account for.) When the trade price is below \$1, the maker rebate is \$0.0001 (TMX Group, 2012b).

participating organization i in second t.

$$\begin{aligned} q_{i,t} = \alpha_i &+ \beta_1 \cdot \text{jump}_t + \beta_2 \cdot \text{post}_t \\ &+ \beta_3 \cdot \text{speedIndex}_i \times \text{jump}_t + \beta_4 \cdot \text{speedIndex}_i \times \text{post}_t \\ &+ \beta_5 \cdot \text{DMM}_i \times \text{jump}_t + \beta_6 \cdot \text{DMM}_i \times \text{post}_t + \varepsilon_{i,t} \end{aligned} \tag{3}$$

This specification includes participating organization fixed effects, which allow us to control for participating organization behavior on that trading day (for example, whether the participating organization tended to be relatively active or inactive). We compute the number of lots passively traded with respect to four different samples of securities. First, we compute them using all securities traded on TSX. These results are displayed in column (1) of Table 2. We then re-estimate the relationship, computing lots passively traded using TSX securities whose prices are negatively affected, unaffected, and positively affected by the tweet. Those results are reported in columns (2)–(4) of the table.

Table 2: Quantity of passive trading over speed and time (in lots per second)

	(1) All	(2) Down	(3) Flat	(4) Up
jump	-0.00794	-0.0712	-0.0506	0.114^{*}
	(0.0932)	(0.0447)	(0.0357)	(0.0679)
post	-0.143	-0.130	-0.0769	0.0466
	(0.575)	(0.416)	(0.0469)	(0.166)
speedIndex × jump	-16.51^{++}	-4.036^{++}	-4.614^{++}	-7.331^{+}
	(2.489)	(0.980)	(0.734)	(2.411)
speedIndex × post	66.25^{++}	50.13^{++}	1.406	14.29^{++}
	(8.267)	(9.069)	(1.467)	(1.789)
DMM × jump	-0.357	-0.177	-0.226	0.0796
	(0.417)	(0.307)	(0.142)	(0.197)
DMM × post	2.737	3.687	-0.600^{**}	-0.268
	(5.540)	(4.944)	(0.282)	(0.779)
Observations	1174889	1174889	1174889	1174889
Participating organization FE	YES	YES	YES	YES

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, $^{+}$ $p < 0.005$, $^{++}$ $p < 0.001$
Coefficients are from OLS regressions with standard errors clustered by participating organization in parentheses. The dependent variables are the number of lots passively traded by a participating organization in a given second. An observation is a (participating organization, second) pair. The sample includes all seconds in the interval [09:30:00,13:10:00]. Column (1) considers the number of lots across all TSX securities. Columns (2)–(4) considers only trades on TSX securities whose mid prices decrease, remain constant, or increase between 13:07:40 and 13:10:00, respectively. The variable DMM is an indicator for whether the participating organization is a designated market maker on TSX. The variable speedIndex is computed as described in Section 5.2.4. The variables jump and post are indicators for the time periods [13:07:40,13:08:10) and [13:08:10,13:10:00], respectively.

The implication of Proposition 2(i) is that $\text{speedIndex}_i \times \text{jump}_t$ should enter negatively (i.e. $\beta_3 < 0$) for samples including securities whose prices are affected by the tweet. In words, the change from the pre period to the jump period in lots passively traded should be relatively lower for faster participating organizations than for slower ones. As Table 2 shows, this prediction is generally borne out in the data. We would expect negative estimates in columns (1), (2), and (4). The estimates are indeed negative, and they are highly significant. The magnitude of the estimate in column (1) can be interpreted as follows: an increase in speedIndex of 0.70 (an amount equal to the quantity-weighted interquartile range of this variable) is associated with 11.55 fewer lots passively traded per second in the jump period.

The flat securities that compose the sample on which column (3) is estimated do not experience a jump. Under a literal interpretation of the model, we would therefore not expect $\text{speedIndex}_i \times \text{jump}_t$ to enter significantly in that column. Nevertheless, we do find it to be negative and significant. An explanation for this is that the model assumes market makers are perfectly able to analyze the content of information. They then cancel quotes if and only if the price will move. However, we conjecture that in a richer and more realistic model in which market makers could only imperfectly analyze the content of information, a fast market maker would think that the price *might* move and respond by canceling his quote, even if the price does not move *ex-post*. Such a model would reconcile the negative estimate in column (3).

The implication of Proposition 2(*ii*) is that $\text{speedIndex}_i \times \text{post}_t$ should enter positively (i.e. $\beta_4 > 0$) for samples including securities whose prices are affected by the tweet. In words, the change from the pre period to the post period in lots passively traded should be relatively higher for faster participating organizations than for slower ones. As Table 2 shows, this prediction is also borne out in the data. We would expect positive estimates in columns (1), (2), and (4). The estimates are indeed positive, and they are highly significant. The magnitude of the estimate in column (1) can be interpreted as follows: an increase in speedIndex of 0.70 (an amount equal to the quantity-weighted interquartile range of this variable) is associated with 46.29 more lots passively traded per second in the post period. Moreover, as expected, we do not find the effect to be present for the flat securities on which column (3) is estimated.

A further implication suggested by these predictions is that liquidity provision should become more concentrated in both the jump period and the post period, relative to the pre period. Whereas both slower and faster traders provide liquidity in the pre period, Proposition 2(*i*) predicts that liquidity provision will be predominantly done by slower traders in the jump period, causing concentration to increase. Similarly, Proposition 2(*ii*) predicts that liquidity provision will be predominantly done by faster traders in the jump period, also causing concentration to increase. In Appendix C.3, we present evidence that this prediction is also supported by the data.

It is also interesting to note that we do not obtain a statistically significant estimate of either the coefficient of $\text{DMM}_i \times \text{jump}_t$ (i.e. β_5) or the coefficient of $\text{DMM}_i \times \text{post}_t$ (i.e. β_6). That is, after controlling for speed, whether or not a participating organization is a designated market maker does not seem to affect the way in which it provides liquidity. Despite their additional obligations to provide liquidity, designated market makers neither appear more likely to be "overrun," trading larger quantities during the jump period, nor appear more likely to establish depth at new levels after a price change, and thus trading larger quantities during the post period. This could be the case either because (*i*) designated market makers fail in their obligation to provide liquidity during this episode, or (*ii*) other traders provide enough liquidity that designated market makers do not need to step in. We therefore fail to find evidence that the designated market maker designation plays a vital role in the exchange ecosystem.

The insignificance of the designated market maker label corroborates the findings of Korajczyk and Murphy (2016), who also study liquidity provision on TSX. In contrast, Anand and Venkataraman (2016), who also study liquidity provision on TSX, do find significance. One explanation for why their findings differ from ours is that they do not separately control for trader speed. Another explanation is that they have more power, which stems from a longer sample as well as the ability to identify traders at a less aggregated level.

6.3 Passive-Side Profits per Lot

The model makes two predictions about the profitability of trades for the passive party in the time before, during, and after information arrives. Proposition 2(*iii*) predicts that passive-side profits per trade should be lower in the jump period than in the pre period. Intuitively, liquidity providers face the highest amount of adverse selection in the jump period, as other traders race to exploit their mispriced quotes. In contrast, Proposition 2(*iv*) predicts that passive-side profits per trade should be higher in the post period than in the pre period. Intuitively, most liquidity in the post period is provided by faster traders who are better equipped to protect themselves from adverse selection, and consequently, liquidity providers as a whole face the lowest amount of adverse selection in this period.

To test whether these predictions are borne out in the data, we estimate the following equation, where the dependent variable π_l is the profits earned by the passive side of the trade of a lot l, and $x(l)$ denotes the security that is traded.

$$\pi_l = \alpha_{x(l)} + \beta_1 \cdot \text{jump}_l + \beta_2 \cdot \text{post}_l + \varepsilon_l \tag{4}$$

This specification includes security fixed effects, which allow us to control for security-specific behavior on that trading day (for example, whether passive parties tended to be primarily selling or primarily buying that security). We estimate this equation on four different samples. First, we estimate it for all traded lots that we classify as either passive buys or passive sells for all securities traded on TSX. These results are displayed in column (1) of Table 3. We then re-estimate the relationship, focusing on traded lots for TSX securities whose prices are negatively affected, unaffected, and positively affected by the tweet. Those results are reported in columns (2)–(4) of the table.

Table 3: Passive-side profits over time (in dollars per lot)

	(1) All	(2) Down	(3) Flat	(4) Up
jump	-0.0191*	-0.0310**	0.0148	0.0110
	(0.0111)	(0.0132)	(0.0125)	(0.0212)
post	0.0243**	0.0193	-0.00692	0.0422
	(0.0115)	(0.0141)	(0.00644)	(0.0329)
Observations	1579556	700153	401443	454773
Security FE	YES	YES	YES	YES

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, + $p < 0.005$, ++ $p < 0.001$
Coefficients are from OLS regressions with standard errors clustered by security in parentheses. The dependent variables, profits from a passive trade, are computed using marked-to-market accounting as described in Section 5.2.3, and is measured in dollars. An observation is a lot traded. The sample includes all trades in the time period [09:30:00,13:10:00] that could be classified as either passive buys or passive sells. In column (1), we include trades for all TSX securities. In columns (2)–(4), we include only TSX securities whose mid prices decrease, remain constant, or increase between 13:07:40 and 13:10:00, respectively. The variables jump and post are indicators for the time periods [13:07:40,13:08:10) and [13:08:10,13:10:00], respectively.

The implication of Proposition 2(*iii*) is that jump_l should enter negatively (i.e. $\beta_1 < 0$) for samples including securities whose prices are affected by the tweet. In words, trades made in the jump period should be less profitable for the passive side than trades made in the pre period. The implication of Proposition 2(*iv*) is that post_l should enter positively (i.e. $\beta_2 > 0$) for samples including securities whose prices are affected by the tweet. In words, trades made in the post period should be more profitable for the passive side than trades made in the pre period.

As Table 3 shows, these predictions are generally borne out in the data. Evidence in favor of these predictions comes through most strongly in column (1), which is estimated on data from all securities on TSX. As column (1) demonstrates, on TSX, the passive side fares \$0.0191 worse per lot traded in the jump period than in the pre period and \$0.0243 better in the post period than in the pre period.

In Appendix E, we repeat the analysis of this section on data from NASDAQ. This is possible because this analysis does not require any knowledge of the identities of participating organizations, which we do not observe for NASDAQ in any meaningful sense. The analyses of the previous sections, however, do require knowledge of these identities, and therefore cannot be repeated with our NASDAQ data.

7 Discussion

This section begins by discussing two potential concerns with our empirical approach. We then conclude.

7.1 Potential Concerns

One potential concern with our empirical approach stems from the fact that we observe identities only at the level of the participating organization, which is a moderately coarse level of aggregation. Furthermore, participating organizations have the ability to conceal their identity, trading anonymously. These issues would undermine our empirical approach if it were the case that either (*i*) the decision of a given trader about which participating organization to use for market access is correlated with market conditions, or (*ii*) the decision of a given trader about whether to trade anonymously is correlated with market conditions. To interpret our results as we wish, we must assume that neither of these forces are at work over the course of our sample. Since decisions about which participating organization to use and whether to trade anonymously seem likely to be fairly fixed over the short-run, we view this assumption as not unreasonable.

Given the assumption that the decisions about which participating organization to use and whether to trade anonymously are orthogonal to the arrival of the tweet, we can interpret our results as informative about the advantages of speed. The empirical approach relies upon the existence of natural variation in the speed of the clientèle of participating organizations (and in the propensity of different traders to remain anonymous). If anything, our results are likely attenuated relative to what we would find if we could observe traders at a less aggregated level.

A second potential concern deals with the interpretation of our speed index variable. It is possible to envision scenarios in which a slower trader is the first to establish a new bid after a price increase or to establish a new ask after a price decrease.[22] However, such scenarios seem unlikely to bias our results systematically in any direction. Moreover, in Appendix D, we show that our speed index variable correlates well with other proxies of speed that have been used in the literature.

[22] We thank an anonymous referee for suggesting two such scenarios. First, a large institution might submit a very large order to buy. If that order is sufficiently large, then it might both take out the old ask and establish a new bid. Second, a retail trader might place an order to buy, thinking it was a marketable limit order. However, if the ask were to rise before the order reaches the exchange, then the limit order might become non-marketable, in which case it would instead establish a new bid.

7.2 Conclusion

The race toward faster trading technology is not new to financial markets. What is different now, however, is the ability of fast traders to make vastly more trading decisions than a human trader and to do so within the blink of an eye. Combined with the introduction of open limit order books, this raises important questions about how liquidity provision has been reshaped by high-frequency trading. This paper informs the debate about how high-frequency liquidity providers behave and whether their presence imposes a negative externality on other market participants.

We propose a new measure of trading speed, which is based on the behavior of participating organizations in response to price changes. We document a considerable amount of heterogeneity in this measure of speed. As expected, faster liquidity providers are better equipped to cope with information shocks. We identify and quantify two channels through which their advantage is realized. First, faster liquidity providers are able to mitigate adverse selection by canceling mispriced quotes before they can be exploited. This entails a negative externality on slower liquidity providers, who must then shoulder a greater portion of the adverse selection risk. Second, faster liquidity providers are more likely to win the race to post quotes at updated prices and thus obtain better queue position. This also entails a negative externality on slower liquidity providers, who are left with worse queue position. Both channels are economically and statistically significant.

Our analysis leverages a unique event in order to obtain clean identification and insight into these issues. Our findings demonstrate that faster liquidity providers do indeed have an advantage in how they react to news, which indicates that equity trading has been reshaped by the recent introduction of faster trading technologies and the limit order book.

A Proofs

A.1 Proof of Proposition 1

Proof. We first establish that all market makers earn zero profits in equilibrium. The second step is to show that no market maker possesses a deviation that would yield positive profits.

We first compute the profits from trades occurring on the bid-side of the book that accrue to market makers using Strategy 1. If the market maker does not own the best bid, then her flow profits are zero. If the market maker does own the best bid, then we compute her profits as follows. The arrival rate of investors who wish to sell is $\frac{\lambda_{invest}}{2}$. Conditional on the arrival of an investor who wishes to sell, she earns $\frac{s}{2}$. The arrival rate of downward jumps in v_t is $\frac{\lambda_{jump}}{2}$. Conditional on a downward jump arriving, this market maker loses $\Pr\left(|J| > \frac{s}{2}\right) \cdot \mathbb{E}\left(|J| - \frac{s}{2} \;\middle|\; |J| > \frac{s}{2}\right)$. Therefore if she owns the best bid, then her flow profits from trades occurring at the bid are

$$\frac{\lambda_{invest}}{2} \cdot \frac{s}{2} - \frac{\lambda_{jump}}{2} \cdot \Pr\left(|J| > \frac{s}{2}\right) \cdot \mathbb{E}\left(|J| - \frac{s}{2} \;\middle|\; |J| > \frac{s}{2}\right) = 0.$$

By symmetry, market makers using Strategy 1 also earn zero profits from trades occurring on the ask-side of the book, for total equilibrium flow profits of zero.

We next compute the profits from trades occurring on the bid-side of the book that accrue to market makers using Strategy 2. If the market maker does not own the best bid, then we compute her profits as follows. The arrival rate of downward jumps in v_t is $\frac{\lambda_{jump}}{2}$. Conditional on a downward jump arriving, she earns $\Pr\left(|J| > \frac{s}{2}\right) \cdot \mathbb{E}\left(|J| - \frac{s}{2} \;\middle|\; |J| > \frac{s}{2}\right)$ if her order is processed by the exchange before the cancel/IOC orders of the other fast market makers, which occurs with probability $\frac{1}{N_{fast}}$. Therefore if she does not own the best bid, then her flow profits from trades occurring at the bid are

$$\frac{\lambda_{jump}}{2} \cdot \frac{1}{N_{fast}} \cdot \Pr\left(|J| > \frac{s}{2}\right) \cdot \mathbb{E}\left(|J| - \frac{s}{2} \;\middle|\; |J| > \frac{s}{2}\right) = \frac{c_{speed}}{2}.$$

If the market maker does own the best bid, then we compute her profits as follows. The arrival rate of investors who wish to sell is $\frac{\lambda_{invest}}{2}$. Conditional on the arrival of an investor who wishes to sell, she earns $\frac{s}{2}$. The arrival rate of downward jumps in v_t is $\frac{\lambda_{jump}}{2}$. Conditional on a downward jump arriving, she loses $\Pr\left(|J| > \frac{s}{2}\right) \cdot \mathbb{E}\left(|J| - \frac{s}{2} \;\middle|\; |J| > \frac{s}{2}\right)$ unless her order is processed by the exchange before the IOC orders of the other fast market makers; her order is processed first with probability $\frac{1}{N_{fast}}$. Therefore, if she does own the best bid, then her flow profits from trades occurring at the bid are

$$\frac{\lambda_{invest}}{2} \cdot \frac{s}{2} - \frac{\lambda_{jump}}{2} \cdot \frac{N_{fast} - 1}{N_{fast}} \cdot \Pr\left(|J| > \frac{s}{2}\right) \cdot \mathbb{E}\left(|J| - \frac{s}{2} \;\middle|\; |J| > \frac{s}{2}\right) = \frac{c_{speed}}{2}.$$

By symmetry, market makers using Strategy 2 also earn flow profits of $\frac{c_{speed}}{2}$ from trades occurring on the ask-side of the book. Because such market makers also pay c_{speed} to be fast, their total equilibrium flow profits are zero.

Finally, the total equilibrium flow profits of market makers who use Strategy 3 and do not trade are also obviously zero.

That no market maker has a profitable deviation is established through essentially the same arguments as in Budish, Cramton and Shim (2015). In short, it is not profitable to deviate by quoting a narrower spread, since such a quote would yield negative profits. It is also not profitable to deviate by quoting a wider spread, since such a quote would never receive the first queue position. □

A.2 Proof of Proposition 2

Proof. We let

$$\alpha = \frac{\lambda_{invest}}{2\lambda_{jump} + \lambda_{invest}} \cdot \frac{N_{slow}}{N_{fast} + N_{slow}}$$

denote the steady state probability that a slow market maker quotes at the best bid (also the best ask) in the equilibrium described in Proposition 1.[23]

[23] A slow trader quotes at the best bid only if both (i) the most recent investor with a need to buy arrived more recently than the most recent jump, and (ii) a slow market maker won the race that occurs among all active market makers to replace the quote at the bid after the arrival of that investor.

The expected number of passive trades that fast and slow traders make are as follows.[24]

$$
\begin{aligned}
q_{slow,pre} &= \frac{\alpha}{N_{slow}} \cdot \left[\lambda_{invest} + \lambda_{jump} \cdot \Pr\left(|J| > \frac{s}{2}\right)\right] \\
q_{fast,pre} &= \frac{(1-\alpha)}{N_{fast}} \cdot \left[\lambda_{invest} + \lambda_{jump} \cdot \Pr\left(|J| > \frac{s}{2}\right) \cdot \frac{N_{fast}-1}{N_{fast}}\right] \\
q_{slow,jump} &= \frac{\alpha}{N_{slow}} \cdot \Pr\left(|J| > \frac{s}{2}\right) \\
q_{fast,jump} &= \frac{(1-\alpha)}{N_{fast}} \cdot \Pr\left(|J| > \frac{s}{2}\right) \cdot \frac{N_{fast}-1}{N_{fast}} \\
q_{slow,post} &= 0 \\
q_{fast,post} &= \frac{1}{N_{fast}} \left[\lambda_{invest} + \lambda_{jump} \cdot \Pr\left(|J| > \frac{s}{2}\right) \cdot \frac{N_{fast}-1}{N_{fast}}\right]
\end{aligned}
$$

We therefore have

$$
\begin{aligned}
Q_{pre} &= \lambda_{invest} + \lambda_{jump} \cdot \Pr\left(|J| > \frac{s}{2}\right)\left(1 - \frac{1-\alpha}{N_{fast}}\right) \\
Q_{jump} &= \Pr\left(|J| > \frac{s}{2}\right)\left(1 - \frac{1-\alpha}{N_{fast}}\right) \\
Q_{post} &= \lambda_{invest} + \lambda_{jump} \cdot \Pr\left(|J| > \frac{s}{2}\right) \cdot \frac{N_{fast}-1}{N_{fast}}
\end{aligned}
$$

We also have

$$
\begin{aligned}
\Pi_{pre} &= \lambda_{invest} \cdot \frac{s}{2} - \lambda_{jump} \cdot \Pr\left(|J| > \frac{s}{2}\right)\left(1 - \frac{1-\alpha}{N_{fast}}\right) \cdot \mathbb{E}\left(|J| - \frac{s}{2} \;\middle|\; |J| > \frac{s}{2}\right) \\
\Pi_{jump} &= -\Pr\left(|J| > \frac{s}{2}\right)\left(1 - \frac{1-\alpha}{N_{fast}}\right) \cdot \mathbb{E}\left(|J| - \frac{s}{2} \;\middle|\; |J| > \frac{s}{2}\right) \\
\Pi_{post} &= \lambda_{invest} \cdot \frac{s}{2} - \lambda_{jump} \cdot \Pr\left(|J| > \frac{s}{2}\right) \cdot \frac{N_{fast}-1}{N_{fast}} \cdot \mathbb{E}\left(|J| - \frac{s}{2} \;\middle|\; |J| > \frac{s}{2}\right)
\end{aligned}
$$

The claims follow from these expressions. □

B Classification of Trades

It is crucial for our analysis to know whether the passive party was the buyer or the seller in a given trade. Since we cannot observe this directly from the TSX data, we must make

[24] For the purposes of these expressions, we use flow rates for the pre and post period, but simply quantities for the jump period. Nevertheless, units cancel when we compute the ratios that appear in the proposition.

the classification ourselves.

Specifically, we classify as passive buys all trades taking place at a price at or below the most recently published bid. Conversely, we classify as passive sells all trades taking place at a price at or above the most recently published ask. However, a small number of trades take place at prices within the most recently published bid-ask spread.[25] We therefore classify most—but not all—trades occurring on TSX.[26]

Our classification differs from the classic Lee and Ready (1991) algorithm in two ways. First, we compare the price of a trade to the most recently published quote, as opposed to the most recently published quote as of five seconds previous. This is to account for the faster speeds in modern electronic markets. Second, we do not classify trades that occur within the spread on the basis of educated guesses about the direction of those trades. Rather, we prefer to be conservative and classify only when we can do so with confidence.

Note also that it is possible for the same participating organization to take both sides of a trade.[27]

C Robustness Checks

We conduct three types of robustness checks. The first considers variants of the speed index defined in Section 5.2.4. The second considers an alternate specification of time. Third, we test the model's predictions about quantities passively traded using an alternative strategy based upon the concentration of liquidity provision.

C.1 Alternative Speed Indices

Here, we re-estimate our main results using two variants of the speed index defined in Section 5.2.4. The speed index calculated previously was based on the share of liquidity provision in the first 100 lots traded at the new bid, in the case of a bid increase, or the first 100 lots traded at the new ask, in the case of an ask decrease.

[25] Trades may take place within the most recently published bid-ask spread for a number of reasons, including (*i*) the most recently published bid-ask spread may not be the prevailing bid-ask spread, which is a possibility due to the incomplete nature of the TSX quote data; or (*ii*) the passive party was using a dark midpoint order, a special type of order occasionally used on TSX, which can be thought of as a limit order whose price moves in lockstep with the mid price and whose quantity is not displayed.

[26] We classify 93.0 percent of the 1,660,735 lots traded in the pre period as either passive buys or passive sells. We also classify 97.9 percent of the 1144 lots traded in the jump period and 98.6 percent of the 49,603 lots traded in the post period.

[27] It is true that in those cases, we would know who must have been the passive party, even if the trade takes place within the spread. However, we do not make use of that information, so as not to treat those trades differently. In other words, we leave as unclassified *all* trades that occur within the most recently published bid-ask spread.

We proceed similarly to define two alternative speed indices, speedIndex(10) and speedIndex (1000), which are based on the first 10 or 1000 lots traded, respectively. In this section, we demonstrate that our results are robust to using either speedIndex(10) or speedIndex (1000) as a measure of speed.

C.1.1 Passive-Side Profits per Second

As a robustness check for the results on passive-side profits presented in Table 1, we repeat the analysis using these alternative speed indices. Panels A and B of Table 4 are constructed using speedIndex(10) and speedIndex(1000), respectively, but are otherwise perfect analogues of Table 1. The results are indeed robust to these alternate speed indices.

Table 4: Passive-side profits over speed and time (alternative speed index, in dollars per second)

	(1) All	(2) Down	(3) Flat	(4) Up
A. speedIndex(10)				
jump + post	-0.0328	-0.0178	0.00447	-0.0191
	(0.0241)	(0.0220)	(0.00313)	(0.0117)
speedIndex(10) × (jump + post)	1.193^{+}	0.504^{+}	-0.0389	0.732**
	(0.357)	(0.167)	(0.0393)	(0.326)
DMM × (jump + post)	0.176	0.0162	0.00159	0.155
	(0.197)	(0.0757)	(0.00871)	(0.146)
Observations	1174889	1174889	1174889	1174889
Participating organization FE	YES	YES	YES	YES
B. speedIndex(1000)				
jump + post	-0.0324	-0.0160	0.00418	-0.0201
	(0.0245)	(0.0213)	(0.00297)	(0.0122)
speedIndex(1000) × (jump + post)	1.117^{+}	0.435**	-0.0301	0.714**
	(0.373)	(0.214)	(0.0404)	(0.313)
DMM × (jump + post)	0.161	0.0117	0.00177	0.145
	(0.193)	(0.0752)	(0.00895)	(0.143)
Observations	1174889	1174889	1174889	1174889
Participating organization FE	YES	YES	YES	YES

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, $^{+}$ $p < 0.005$, $^{++}$ $p < 0.001$
Coefficients are from OLS regressions with standard errors clustered by participating organization in parentheses. The dependent variables, profits of a participating organization from passive trades in a given second, are computed using marked-to-market accounting as described in Section 5.2.3, and are measured in dollars. An observation is a (participating organization, second) pair. The sample includes all seconds in the interval [09:30:00,13:10:00]. In column (1), we include all TSX securities. In columns (2)–(4), we include only TSX securities whose mid prices decrease, remain constant, or increase between 13:07:40 and 13:10:00, respectively. The variable DMM is an indicator for whether the participating organization is a designated market maker on TSX. The variables speedIndex(10) and speedIndex(1000) are computed as described in Section 5.2.4, but are instead based on, respectively, the first 10 and 1000 lots traded at the new bid or new ask. The variable (jump + post) is an indicator for the time period [13:07:40,13:10:00].

C.1.2 Quantity of Passive Trading

As a robustness check for the results on quantity of passive trading presented in Table 2, we repeat the analysis using these alternative speed indicies. Panels A and B of Table 5 are constructed using speedIndex(10) and speedIndex(1000), respectively, but are otherwise perfect analogues of Table 2. The results are indeed robust to these alternate speed indices.

Table 5: Quantity of passive trading over speed and time (alternative speed index, in lots per second)

	(1) All	(2) Down	(3) Flat	(4) Up
A. speedIndex(10)				
jump	-0.0180	-0.0752*	-0.0579	0.116*
	(0.0961)	(0.0446)	(0.0378)	(0.0679)
post	-0.109	-0.104	-0.0794*	0.0574
	(0.613)	(0.436)	(0.0459)	(0.179)
speedIndex(10) × jump	-17.15++	-4.156++	-4.683++	-7.773+
	(2.254)	(0.923)	(0.864)	(2.352)
speedIndex(10) × post	68.94++	52.17++	1.540	14.79++
	(7.094)	(8.313)	(1.484)	(1.764)
DMM × jump	-0.516	-0.216	-0.274*	0.0141
	(0.440)	(0.316)	(0.160)	(0.160)
DMM × post	3.369	4.165	-0.589**	-0.129
	(5.529)	(4.957)	(0.277)	(0.758)
Observations	1174889	1174889	1174889	1174889
Participating organization FE	YES	YES	YES	YES
B. speedIndex(1000)				
jump	0.00507	-0.0649	-0.0477	0.117*
	(0.0902)	(0.0443)	(0.0341)	(0.0681)
post	-0.224	-0.212	-0.0722	0.0425
	(0.615)	(0.449)	(0.0466)	(0.173)
speedIndex(1000) × jump	-16.71++	-4.157++	-4.652++	-7.360+
	(2.338)	(0.954)	(0.756)	(2.359)
speedIndex(1000) × post	67.69++	51.69++	1.291	14.30++
	(8.819)	(9.336)	(1.481)	(1.924)
DMM × jump	-0.277	-0.154	-0.204	0.113
	(0.378)	(0.294)	(0.139)	(0.206)
DMM × post	2.385	3.399	-0.601**	-0.330
	(5.268)	(4.709)	(0.282)	(0.753)
Observations	1174889	1174889	1174889	1174889
Participating organization FE	YES	YES	YES	YES

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, $^{+}$ $p < 0.005$, $^{++}$ $p < 0.001$.

Coefficients are from OLS regressions with standard errors clustered by participating organization in parentheses. The dependent variables are the number of lots passively traded by a participating organization in a given second. An observation is a (participating organization ID, second) pair. The sample includes all seconds in the interval [09:30:00,13:10:00]. Column (1) considers the number of lots across all TSX securities. Columns (2)–(4) considers only trades on TSX securities whose mid prices decrease, remain constant, or increase between 13:07:40 and 13:10:00, respectively. The variable DMM is an indicator for whether the participating organization is a designated market maker on TSX. The variables speedIndex(10) and speedIndex(1000) are computed as described in Section 5.2.4, but are instead based on, respectively, the first 10 and 1000 lots traded at the new bid or new ask. The variables jump and post are indicators for the time periods [13:07:40,13:08:10) and [13:08:10,13:10:00], respectively.

C.2 Linear Time Trend

Here, we re-estimate our main results using an alternate specification of time. Previously we had divided the aftermath of the tweet into two discrete time periods: the jump period and the post period. Here, we instead treat the aftermath of the tweet as a single discrete time period where we also allow for a linear time trend.

For the purposes of this analysis, we use the following variables. For a particular second t, we, as before, define $(\text{jump} + \text{post})_t$ to be an indicator equal to one when $t \in [13{:}07{:}40, 13{:}10{:}00]$. We also define sec_t as the larger of either zero or the number of seconds after 13:07:40. Similarly, for a particular lot traded l, we define $(\text{jump} + \text{post})_l$ as an indicator equal to one when the trade occurred in the interval $[13{:}07{:}40, 13{:}10{:}00]$. We also define sec_l as the larger of either zero and the number of seconds after 13:07:40 that the trade occurred.

C.2.1 Quantity of Passive Trading

As a robustness check for the results on the quantity of passive trading presented in Table 2, we repeat the analysis using this alternate specification of time. To do so, we estimate the following analogue of equation (3):

$$\begin{aligned} q_{i,t} = \alpha_i &+ \beta_1 \cdot (\text{jump} + \text{post})_t + \beta_2 \cdot (\text{jump} + \text{post})_t \times \text{sec}_t \\ &+ \beta_3 \cdot \text{speedIndex}_i \times (\text{jump} + \text{post})_t + \beta_4 \cdot \text{speedIndex}_i \times (\text{jump} + \text{post})_t \times \text{sec}_t \\ &+ \beta_5 \cdot \text{DMM}_i \times (\text{jump} + \text{post})_t + \beta_6 \cdot \text{DMM}_i \times (\text{jump} + \text{post})_t \times \text{sec}_t + \varepsilon_{i,t} \end{aligned}$$

We compute quantities with respect to four different samples of securities. First, we compute them using all securities traded on TSX. These results are displayed in column (1) of Table 6. We then re-estimate the relationship, computing quantities using TSX securities whose prices are negatively affected, unaffected, and positively affected by the tweet. Those results are reported in columns (2)–(4) of the table.

The estimates are similar in flavor to those of our baseline specification. They indicate that faster participating organizations are the passive party of relatively fewer trades immediately after the jump, but relatively more trades in the subsequent period. In particular, we find negative estimates of the coefficient on the "intercept term," $\text{speedIndex}_i \times (\text{jump} + \text{post})_t$, (i.e. $\beta_3 < 0$), and positive estimates of the coefficient on the "slope term," $\text{speedIndex}_i \times (\text{jump} + \text{post})_t \times \text{sec}_t$, (i.e. $\beta_4 > 0$).

Table 6: Quantity of passive trading over speed and time (linear time trend, in lots per second)

	(1) All	(2) Down	(3) Flat	(4) Up
jump + post	-0.150	-0.0376	-0.0124	-0.111
	(0.241)	(0.210)	(0.0567)	(0.0955)
(jump + post) × sec	0.000495	-0.00113	-0.000838	0.00243
	(0.00938)	(0.00735)	(0.000874)	(0.00272)
speedIndex × (jump + post)	-29.21^{++}	-19.47**	-7.516**	-1.604
	(7.822)	(8.349)	(3.174)	(2.717)
speedIndex × (jump + post) × sec	1.111^{++}	0.828^{++}	0.109*	0.162^{++}
	(0.190)	(0.214)	(0.0576)	(0.0338)
DMM × (jump + post)	-1.575	-1.110	-0.112	-0.301
	(2.736)	(2.095)	(0.264)	(0.718)
DMM × (jump + post) × sec	0.0521	0.0567	-0.00582	0.00150
	(0.1000)	(0.0845)	(0.00508)	(0.0160)
Observations	1174889	1174889	1174889	1174889
Participating organization FE	YES	YES	YES	YES

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, $^{+}$ $p < 0.005$, $^{++}$ $p < 0.001$
Coefficients are from OLS regressions with standard errors clustered by participating organization in parentheses. The dependent variables are number of lots passively traded by a participating organization in a given second. An observation is a (participating organization, second) pair. The sample includes all seconds in the interval [09:30:00,13:10:00]. Column (1) considers the number of lots across all TSX securities. Columns (2)–(4) considers only trades on TSX securities whose mid prices decrease, remain constant, or increase between 13:07:40 and 13:10:00, respectively. The variable DMM is an indicator for whether the participating organization is a designated market maker on TSX. The variable speedIndex is computed as described in Section 5.2.4. The variable (jump + post) is an indicator for the time period [13:07:40,13:10:00]. The variable sec is equal to the number of seconds after 13:07:40 if that is positive, and is zero otherwise.

C.2.2 Passive-Side Profits per Lot

As a robustness check for the results on profitability of passive trades presented in Table 3, we repeat the analysis using this alternate specification of time. To do so, we estimate the following analogue of equation (4):

$$\pi_l = \alpha_{x(l)} + \beta_1 \cdot (\text{jump} + \text{post})_l + \beta_2 \cdot (\text{jump} + \text{post})_l \times \text{sec}_l + \varepsilon_l$$

We estimate this equation on four different samples. First, we estimate it for all traded lots that we classify as either passive buys or passive sells for all securities traded on TSX. These results are displayed in column (1) of Table 7. We then re-estimate the relationship, focusing on traded lots for TSX securities whose prices are negatively affected, unaffected, and positively affected by the tweet. Those results are reported in columns (2)–(4) of the table.

The implications of Proposition 2(*iii*) and (*iv*) are that $\beta_1 < 0$ and $\beta_2 > 0$ for samples including securities whose prices are affected by the tweet. As Table 7 illustrates, neither of these predictions are borne out. A potential explanation for why we fail to find support for these predictions is that the linearity imposed by this specification may be too strong an assumption.

Table 7: Passive-side profits over time (linear time trend, in dollars per lot)

	(1) All	(2) Down	(3) Flat	(4) Up
jump + post	0.0861*	0.0465*	0.00989	0.188
	(0.0483)	(0.0242)	(0.00853)	(0.128)
(jump + post) × sec	-0.000650	-0.000287	-0.000160**	-0.00163
	(0.000424)	(0.000246)	(0.0000704)	(0.00113)
Observations	1579556	700153	401443	454773
Security FE	YES	YES	YES	YES

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, + $p < 0.005$, ++ $p < 0.001$
Coefficients are from OLS regressions with standard errors clustered by security in parentheses. The dependent variables, profits from a passive trade, are computed using marked-to-market accounting as described in Section 5.2.3, and is measured in dollars. An observation is a lot traded. The sample includes all trades in the time period [09:30:00,13:10:00] that could be classified as either passive buys or passive sells. In column (1), we include trades for all TSX securities. In columns (2)–(4), we include only TSX securities whose mid prices decrease, remain constant, or increase between 13:07:40 and 13:10:00, respectively. The variable (jump + post) is an indicator for the time period [13:07:40,13:10:00]. The variable sec is equal to the number of seconds after 13:07:40 if that is positive, and is zero otherwise.

C.3 Market Concentration

A further implication suggested by Proposition 2(*i*) and (*ii*) is that liquidity provision should become more concentrated in both the jump period and the post period, relative to the pre period. Whereas both faster and slower traders provide liquidity in the pre period, Proposition 2(*i*) predicts that liquidity provision will be predominantly done by slower traders in the jump period and therefore more concentrated. Similarly, Proposition 2(*ii*) predicts that liquidity provision will be predominantly done by faster traders in the jump period, also resulting in more concentrated liquidity provision.

In this section we investigate the extent to which the concentration of liquidity provision changed in response to the news event. To that end, we calculate for each stock x and each of two time periods t (the pre period and the combination of the jump and post periods), the Herfindahl-Hirschman Index (HHI) for passive trading of stock x during period t as follows:

$$\mathrm{HHI}_{x,t} = \sum_i \left(\frac{q_{x,i,t}}{\sum_i q_{x,i,t}} \right)^2,$$

where $q_{x,i,t}$ denotes the number of lots of security x traded in second t (that we can classify as either passive buys or passive sells) for which participating organization i was the passive party.

To ensure that we are comparing like with like, for each stock, we restrict the number of trades on which we base the market concentration to be the same in each period. When it is necessary to ration trades conducted in the pre period, we retain the latest trades. When it is necessary to ration trades conducted in the (jump + post) period, we retain the earliest trades.[28]

We then estimate the following equation on four different samples:

$$\mathrm{HHI}_{x,t} = \alpha_x + \beta \cdot (\text{jump} + \text{post})_t + \varepsilon_{x,t} \tag{5}$$

We first estimate (5) on all TSX securities. These results are displayed in column (1) of Table 8. We then re-estimate the relationship, focusing on TSX securities whose prices are negatively affected, unaffected, and positively affected by the tweet. Those results are reported in columns (2)–(4) of the table.

[28] Note that some securities were not traded in either the jump or post period, and other securities were not traded in the pre period. In these cases, we are unable to compute the HHI, and they fall out of the sample.

Table 8: HHI of liquidity provision over time

	(1) All	(2) Down	(3) Flat	(4) Up
jump + post	0.0397^{+}	0.0535***	0.0387	0.0253
	(0.0133)	(0.0191)	(0.0324)	(0.0269)
Observations	2163	637	1090	263
Security FE	YES	YES	YES	YES

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, $^{+}$ $p < 0.005$, $^{++}$ $p < 0.001$
Coefficients are from OLS regressions with standard errors in parentheses. The dependent variables, the Herfindahl-Hirschman Index (HHI) of liquidity provision for a security during a time period, are computed as in the text. An observation is a (security, time period) pair. The sample includes two time periods: pre, [9:30:00,13:07:40), and (jump + post), [13:07:40,13:10:00]. In column (1), we include all TSX securities. In columns (2)–(4), we include only TSX securities whose mid prices decrease, remain constant, or increase between 13:07:40 and 13:10:00, respectively.

The implication of Proposition 2(*i*) and (*ii*) is that $(\text{jump} + \text{post})_t$ should enter negatively (i.e. $\beta > 0$) for samples including securities whose prices are affected by the tweet. In words, liquidity provision should be more concentrated after the tweet than before the tweet. Table 8 produces the results of this analysis and shows that the concentration of liquidity provision increases significantly. Column (1) reveals that the HHI increased on average by 0.0303 after the tweet, relative to an average pre period concentration of 0.585.

D Other Proxies for Speed

The speed index that we use for this paper, which is described in Section 5.2.4, is new to the literature. However, other academic studies have used alternative proxies for speed. In this appendix, we argue that our speed index correlates well with many of these alternative proxies.

D.1 Alternative Proxies

Perhaps the most common proxies for speed are the three used by Kirilenko, Kyle, Samadi and Tuzun (forthcoming), as well as much follow-on literature, to identify high-frequency traders. The first of these proxies is volume, where large values are associated with high-frequency traders. The second is the end-of-day net position as a fraction of total trading volume, where small values are associated with high-frequency traders. The third involves a measure of intraday inventory, where small values are again associated with high-frequency traders.

Another proxy for speed is average trade size, where small values are associated with high-frequency traders. This is used by Weller (2016), for example.

D.2 Variable Construction

In this appendix, we argue that our speed index correlates well in our data with the four alternative proxies for speed mentioned above. To that end, we construct variables to represent each of these three proxies.

First, we construct speedProxy(1) to capture volume. For each participating organization, we compute the number of lots traded on April 23, 2013 in all securities. We then apply min-max normalization.

Next, we construct speedProxy(2) to capture end-of-day net position as a fraction of total trading volume. For each participating organization, we compute the number of lots traded on April 23, 2013 in all securities. We also compute the sum across securities of the absolute value of their net position in each security at the end of the trading day (also in lots, and assuming no initial inventory). We then take the ratio of those two numbers, apply min-max normalization, and subtract the resulting value from one (so that larger values of the proxy correspond to faster participating organizations).

We also construct speedProxy(3) to capture the notion of intraday inventory used by Kirilenko, Kyle, Samadi and Tuzun (forthcoming). For each participating organization, each security, and each minute, we compute the squared deviation of the net position held in that minute from the net position at the end of the day. We then sum across minutes and securities. We then take the square root of that sum and divide by the number of lots traded on that day in all securities. Finally, we apply min-max normalization and subtract the resulting value from one (so that larger values of the proxy correspond to faster participating organizations).

Lastly, we construct speedProxy(4) to capture average trade size. For each participating organization, we compute the average number of lots involved in their trades on April 23, 2013 in all securities. We then apply min-max normalization and subtract the resulting value from one (so that larger values of the proxy correspond to faster participating organizations).

D.3 Correlation

Finally, we calculate the correlations among all pairs of the five variables of interest: our speedIndex, as well as the four alternative proxies for speed: speedProxy(1), speedProxy(2), speedProxy(3), and speedProxy(4). This correlation matrix is reported in Table 9. We find that speedIndex correlates positively with all four proxies. Moreover, the strength of this correlation is on the order of the strength of the correlations between the proxies that are already established in the literature.

Table 9: Correlation of proxies for speed

	speedIndex	speedProxy			
		(1)	(2)	(3)	(4)
speedIndex	1.000				
speedProxy(1)	0.977	1.000			
speedProxy(2)	0.376	0.404	1.000		
speedProxy(3)	0.257	0.255	0.586	1.000	
speedProxy(4)	0.061	0.046	0.176	0.557	1.000

This table shows the correlation of five measures of trading speed over the participating organizations ($N = 90$). speedIndex is defined in Section 5.2.4. speedProxy(1), speedProxy(2), and speedProxy(3) are based on volume, end-of-day net position, and the variation of intraday inventory as in Kirilenko, Kyle, Samadi and Tuzun (forthcoming). speedProxy(4) is based on average trade size as in Weller (2016).

E External Validity: NASDAQ

While the main analysis of this paper is conducted on Canadian data from TSX, markets in the United States were similarly affected. To demonstrate this, we establish in this appendix that many of the same features we document for TSX are also true in NASDAQ data. One caveat is that the NASDAQ data do not contain meaningful participating organization identifiers, so it is not possible for us to repeat all analyses on it. However the analyses that can be repeated yield results in line with the conclusions derived from TSX data.

E.1 Data

We use the NASDAQ ITCH TotalView data stream, which provides order-level information about all actions that affect the limit order book on NASDAQ.[29] Each observation contains a time stamp, an order number, and details about price, volume, and order type. The data are a historical record of the real-time data feed that participating organizations use to track orders and trades. Messages are recorded chronologically and time-stamped to the nanosecond.

[29] On April 23, 2013 the message feed contained 296.8 million messages. These messages relate to (*i*) the state of the order book (e.g. additions of limit orders and cancellations), (*ii*) trades, (*iii*) the open and closing crosses, as well as (*iv*) various system messages.

Of these data, we use only observations pertaining to a single security: the SPDR S&P 500 ETF Trust (SPY), which tracks the S&P 500. This security is a natural one to focus on because of its significance as the world's largest ETF and because it summarizes the market as a whole.

The NASDAQ data and TSX data differ along several dimensions. On one hand, the NASDAQ data include a more detailed account of the actions affecting the limit order book. On the other hand, we are able to observe which participating organization was on the passive side of a trade with a much higher frequency on TSX than on NASDAQ.[30]

E.2 Event Description

The reaction of the Canadian market to the tweet was described in Section 4. Here, we demonstrate that American markets responded in the same way, using NASDAQ data on SPY.

As Figure 3(a) shows, the price of SPY dropped by 97.4bps in just over two minutes following the tweet.[31] It recovered equally quickly thereafter, following an announcement that the Twitter post was a hoax. This price series is very similar to the corresponding price series for XSP, depicted in Figure 1(a).

As Figure 3(b) shows, the spread of SPY increases immediately after the posting of the tweet, ultimately reaching a peak almost twice its typical value before the tweet. This series is also very similar to its counterpart for XSP, depicted in Figure 1(b).

[30] In the NASDAQ data for SPY, we would be able to infer the participating organization on the passive side of a trade for at most 0.86 percent of the lots traded in the data.

[31] The magnitude of this decline is similar to the security's average daily price movement of 85bps, where the calculation of the latter figure is based on the distribution of $(High - Low)/Close$ for all trading days of the period January 2013 – July 2013.

Figure 3: Market response to event, NASDAQ:SPY

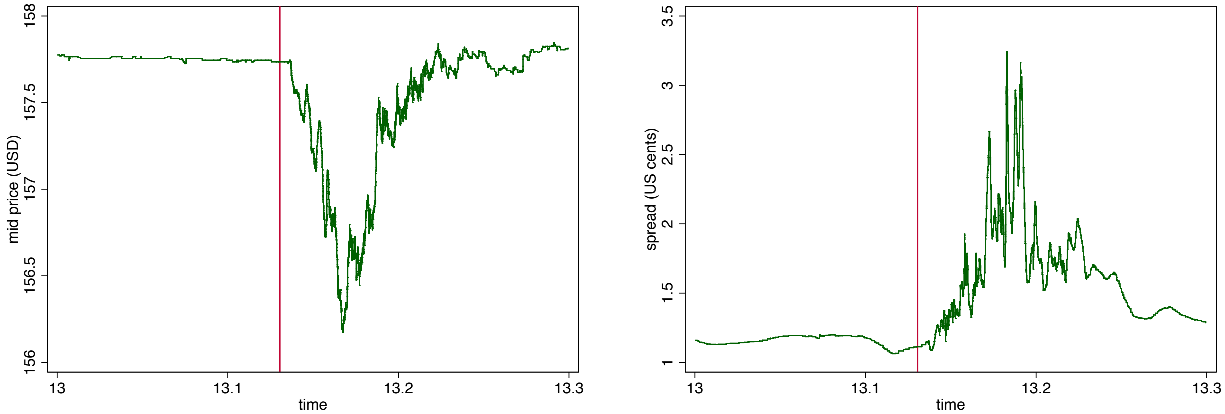

(a) Mid price

(b) Spread

In both panels, the horizontal axis is time on April 23, 2013 measured in hours after midnight EDT. The vertical line at 13:07:50 marks the published time of the tweet. The left panel plots, for each second, the time-weighted average mid price of SPDR S&P 500 ETF Trust (SPY) on NASDAQ, measured in USD. For the right panel, we compute the spread of SPY on NASDAQ measured in US cents, in every second, where we take a time-weighted average if multiple spreads prevailed in a given second. The right panel then plots a lowess over these observations using a 72 second bandwidth.

Finally, we are also able to use the detailed order-level nature of the NASDAQ data to study how order flows responded to the tweet. As Figure 4 demonstrates, the tweet triggers a drastic increase in trades. Immediately following the tweet, contracts were traded at a rate of 2.5 million per minute on NASDAQ, a 37-fold increase over the rate at which contracts were traded in the hours leading up the event. The pattern for order flow at the bid and ask is similar.

Figure 4: Trades and order flow at the bid and ask, NASDAQ:SPY

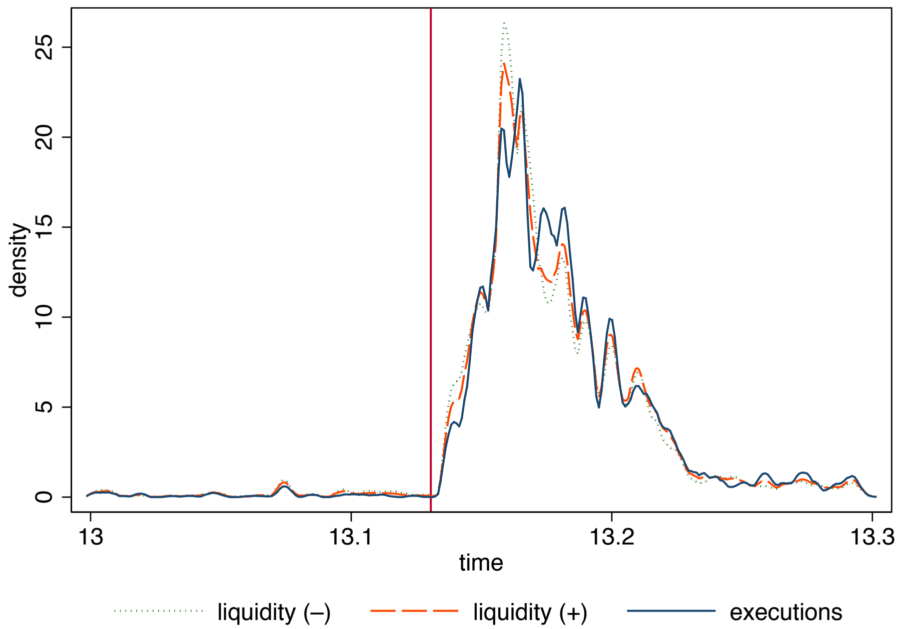

The horizontal axis is time on April 23, 2013 measured in hours after midnight EDT. The vertical line at 13:07:50 marks the published time of the tweet. The graph plots kernel densities (using an Epanechnikov kernel function with a six second bandwidth) for three types of aggregate share-weighted flow variables: (*i*) liquidity (–): deletions, cancellations, or modifications that remove liquidity from the bid or ask, (*ii*) liquidity (+): additions and modifications that add liquidity at the bid or ask, and (*iii*) trades.

E.3 Passive-Side Profits per Lot

Because the NASDAQ dataset does not contain meaningful participating organization identifiers, we are unable to use it to investigate either whether faster traders obtain relatively higher profits from passive trades in the aftermath of the tweet (*cf.* Section 6.1), or whether faster traders are the passive side of relatively fewer trades in the jump period and relatively more in the post period (*cf.* Section 6.2). We can, however, investigate whether trades in the jump period are relatively less profitable for the passive side and whether trades in the post period are relatively more profitable for the passive side (*cf.* Section 6.3). And indeed, we do find evidence of these same phenomena transpiring on NASDAQ.

To test this formally, we estimate the following equation, where the dependent variable π_l is the passive-side profits from the trade of lot l, measured in USD.

$$\pi_l = \beta_0 + \beta_1 \cdot \mathrm{jump}_l + \beta_2 \cdot \mathrm{post}_l + \varepsilon_l$$

This is the analogue of the estimation performed in Section 6.3.

Column (1) of Table 10 contains the results of estimating this equation on the NASDAQ data for SPY. Theory predicts $\beta_1 < 0$ and $\beta_2 > 0$. Both predictions are borne out in the data, and in fact the effect appears to be even stronger for SPY than than it was for TSX. We find that for SPY, the passive party fares 0.529 USD worse per lot traded in the jump period than in the pre period, and 0.390 USD better in the post period than in the pre period.

As a robustness check, we also estimate the relationship on an alternative specification that includes a linear time trend rather than discrete time periods.

$$\pi_l = \beta_0 + \beta_1 \cdot (\text{jump} + \text{post})_l + \beta_2 \cdot (\text{jump} + \text{post})_l \times \text{sec}_l + \varepsilon_l$$

This is the analogue of the estimation performed in Appendix C.2.2.

Column (2) of Table 10 contains the results of estimating this equation on the NASDAQ data for SPY. Theory predicts $\beta_1 < 0$ and $\beta_2 > 0$. The latter prediction is borne out. However, we do not find support for the former prediction, perhaps because the assumption of linearity is too strong.

Table 10: Passive-side profits over time (NASDAQ:SPY, in dollars per lot)

	(1)	(2)
jump	-0.529^{++}	
	(0.140)	
post	0.390^{++}	
	(0.0441)	
jump + post		0.0113
		(0.0484)
(jump + post) × sec		0.00370^{++}
		(0.000175)
Constant	-0.509^{++}	-0.509^{++}
	(0.0439)	(0.0439)
Observations	40409	40409

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, $^{+}$ $p < 0.005$, $^{++}$ $p < 0.001$
Coefficients are from OLS regressions with robust standard errors in parentheses. The dependent variables, profits from a passive trade, are computed using marked-to-market accounting as described in Section 5.2.3, and is measured in US dollars. An observation is a lot traded. The sample includes all trades in the time period [09:30:00,13:10:00]. The variables jump and post are indicators for the time periods [13:07:40,13:08:10) and [13:08:10,13:10:00], respectively. The variable (jump + post) is an indicator for the time period [13:07:40,13:10:00]. The variable sec is equal to the number of seconds after 13:07:40 if that is positive, and is zero otherwise.

F Tables of Variables

Table 11: Participating Organizations *(N=90)*

	pre period			jump period			post period		
	p25	p50	p75	p25	p50	p75	p25	p50	p75
All trades									
lots traded	678.25	3338.37	18762.99	0	0	11	0	13.5	144
number of trades	84	545	3545	0	0	6	0	4.5	56
value traded (CAD thousands)	711.29	5333.73	42661.44	0	0	17.11	0	16.3	237.65
Passive trades									
lots traded	212	1603.49	8563.70	0	0	4	0	4.75	100.34
number of trades	26	176.5	1780	0	0	3	0	2	32
value traded (CAD thousands)	189.63	2060.68	11780.77	0	0	6.18	0	9.27	171.76
Aggressive trades									
lots traded	269.56	1462.78	10741.76	0	0	1	0	.01	37
number of trades	64	348	1780	0	0	1	0	.5	17
value traded (CAD thousands)	350.73	1890.65	11977.11	0	0	1.26	0	.09	56.04

Summary statistics for all participating organizations on TSX who identified as a party in at least one trade on April 23, 2013. The "anonymous" identifier is treated as a separate participating organization. The pre period is the time interval [9:30:00,13:07:40). The jump period is the time interval [13:07:40,13:08:10). The post period is the time interval [13:08:10,13:10:00]. For every period and every participating organization, we compute the number of lots traded, the number of trades made, and the value traded (in thousands of dollars). In the first set of rows, these statistics are computed with respect to all trades made in a period by a participating organization. For second and third sets of rows, these statistics are computed with respect to all trades in a period for which the participating organization was classified as the passive party or the active party, respectively. Not all trades are classified (*cf.* Appendix B). In each case, we compute the 25th, 50th, and 75th percentiles of the values of the statistic over the set of participating organizations.

Table 12: Securities

	pre period			jump period			post period		
	p25	p50	p75	p25	p50	p75	p25	p50	p75
All securities (N=2,367)									
lots traded	0	32.64	213.79	0	0	0	0	0	0
spread (bps)	23.14	69.10	330.58	21	63.19	289.86	21.67	64.56	295.57
spread (cents)	2	4	14	1.5	4	13	1.5	4	14
Securities with negative price change (N=489)									
lots traded	6.85	84	883.99	0	0	0	0	.5	21
spread (bps)	11.05	21.86	55.81	9.95	22.08	51.39	11.27	23.49	53.91
spread (cents)	2	3	6	1	3	6	1	3	6
Securities with no price change (N=1,405)									
lots traded	0	25.25	130	0	0	0	0	0	0
spread (bps)	37.99	113.37	454.55	31.95	103.84	409.15	32.3	103.84	416.67
spread (cents)	2	5	15	1	4	14	1	4	14
Securities with positive price change (N=219)									
lots traded	5	76.3	905.58	0	0	0	0	0	12
spread (bps)	18.66	33.28	76.82	16.46	34.23	80.32	16.43	33.3	74.14
spread (cents)	2	4	9	1	4	8	2	4	9

Summary statistics for all securities traded or quoted on TSX on April 23, 2013. The pre period is the time interval [9:30:00,13:07:40). The jump period is the time interval [13:07:40,13:08:10). The post period is the time interval [13:08:10,13:10:00]. For every period and every security, we compute the number of lots traded and two measures of the spread. The spread is first computed as a fraction relative to the prevailing mid price and expressed in basis points. It is second computed in raw form and expressed in cents. In both cases, the spread of a security in a period is the median taken over all spreads prevailing at the end of each second in that period (where an undefined spread is treated as infinite). The first set of rows corresponds to all securities that were either traded or quoted on TSX on April 23, 2013. Later sets of rows correspond to all securities that responded to the tweet with a negative price change, with no price change, or with a positive price change, respectively. The price change of a security x is determined by the sign of $m_{x,13:10:00} - m_{x,13:07:40}$, which can be undefined if either of those mid prices fail to exist. For each period and statistic, we compute the 25th, 50th, and 75th percentiles of the values of the statistic for the period over the relevant set of securities.

Table 13: Variables used in estimation

Name	Description
Dependent Variables	
$\pi_{i,t}$	profits from passive trades of participating organization i in second t (dollars)
$q_{i,t}$	number of lots passively traded by participating organization i in second t
π_l	passive-side profits from the trade of lot l (dollars)
$\text{HHI}_{x,t}$	HHI of liquidity provision for security x in second t
Independent Variables	
speedIndex_i	speed index of participating organization i (using first 100 lots traded at new level)
$\text{speedIndex}(10)_i$	speed index of participating organization i (using first 10 lots traded at new level)
$\text{speedIndex}(1000)_i$	speed index of participating organization i (using first 1000 lots traded at new level)
DMM_i	indicator for participating organization i being a designated market maker
jump_t	indicator for second t being in [13:07:40,13:08:10)
jump_l	indicator for trade of lot l being in [13:07:40,13:08:10)
post_t	indicator for second t being in [13:08:10,13:10:00]
post_l	indicator for trade of lot l being in [13:08:10,13:10:00]
$(\text{jump} + \text{post})_t$	indicator for second t being in [13:07:40,13:10:00]
$(\text{jump} + \text{post})_l$	indicator for trade of lot l being in either [13:07:40,13:10:00]
sec_t	larger of zero and number of seconds that second t is after 13:07:40
sec_l	larger of zero and number of seconds that trade of lot l is after 13:07:40

We index participating organizations with the subscript i. We index seconds with the subscript t. We index traded lots with the subscript l. We index securities with the subscript x.

References

Anand, Amber and Kumar Venkataraman, "Market Conditions, Fragility and the Economics of Market Making," *Journal of Financial Economics*, 2016, *121* (2), 327–349. http://ssrn.com/abstract=2179259.

The Associated Press (@AP), "Breaking: Two Explosions in the White House and Barack Obama is injured," 2013. 23 April 2013, 13:07:50. Tweet.

Baldauf, Markus and Joshua Mollner, "High-Frequency Trading and Market Performance," *Working Paper*, 2015. http://ssrn.com/abstract=2674767.

Balduzzi, Pierluigi, Edwin J Elton, and T Clifton Green, "Economic News and Bond Prices: Evidence From the U.S. Treasury Market," *Journal of Financial and Quantitative Analysis*, 2001, *36* (4), 523–543.

Baron, Matthew, Jonathan Brogaard, Björn Hagströmer, and Andrei Kirilenko, "Risk and Return in High Frequency Trading," *Working Paper*, 2016. http://ssrn.com/abstract=2433118.

Baruch, Shmuel and Lawrence Glosten, "Fleeting Orders," *Working Paper*, 2013. http://ssrn.com/abstract=2278457.

BATS Global Markets, Inc., "Historical Market Volume Data," http://batstrading.com/market_data/market_volume_history 2016. Accessed: July 20, 2016.

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

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

_ , Terrence Hendershott, and Ryan Riordan, "High-Frequency Trading and Price Discovery," *Review of Financial Studies*, 2014, *27* (8), 2267–2306.

_ , _ , Stefan Hunt, and Carla Ysusi, "High-Frequency Trading and the Execution Costs of Institutional Investors," *Financial Review*, 2014, *49* (2), 345–369.

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

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

Chaboud, Alain, Ben Chiquoine, Erik Hjalmarsson, and Clara Vega, "Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market," *The Journal of Finance*, 2014, *69* (5), 2045–2084.

Conrad, Jennifer, Sunil Wahal, and Jin Xiang, "High-Frequency Quoting, Trading, and the Efficiency of Prices," *Journal of Financial Economics*, 2015, *116* (2), 271–291.

Copeland, Thomas E. and Dan Galai, "Information Effects on the Bid-Ask Spread," *The Journal of Finance*, 1983, *38* (5), 1457–1469.

Easley, David, Marcos Lopez de Prado, and Maureen O'Hara, "The Microstructure of the 'Flash Crash': Flow Toxicity, Liquidity Crashes and the Probability of Informed Trading," *The Journal of Portfolio Management*, 2011, *37* (2), 118–128.

Egginton, Jared F., Bonnie F. Van Ness, and Robert A. Van Ness, "Quote Stuffing," *Working Paper*, 2016. http://ssrn.com/abstract=1958281.

Fleming, Michael J and Eli M Remolona, “Price Formation and Liquidity in the U.S. Treasury Market: The Response to Public Information,” *The Journal of Finance*, 1999, *54* (5), 1901–1915.

Foucault, Thierry, “Order Flow Composition and Trading Costs in a Dynamic Limit Order Market,” *Journal of Financial Markets*, 1999, *2* (2), 99–134.

_ , Johan Hombert, and Ioanid Roşu, “News Trading and Speed,” *The Journal of Finance*, 2016, *71* (1), 335–382.

Gai, Jiading, Chen Yao, and Mao Ye, “The Externalities of High Frequency Trading,” *Working Paper*, 2013. http://ssrn.com/abstract=2066839.

Glosten, Lawrence R., “Is the Electronic Open Limit Order Book Inevitable?,” *The Journal of Finance*, 1994, *49* (4), 1127–1161.

_ and Paul R. Milgrom, “Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders,” *Journal of Financial Economics*, 1985, *14* (1), 71–100.

Hananel, Sam (@SamHananelAP), “Please Ignore AP Tweet on explosions, we’ve been hacked.,” https://twitter.com/SamHananelAP/status/326744786444886017 2013. 23 April 2013, 13:10:02. Tweet.

Hasbrouck, Joel, “High Frequency Quoting: Short-Term Volatility in Bids and Offers,” *Working Paper*, 2015. http://ssrn.com/abstract=2237499.

_ and Gideon Saar, “Technology and Liquidity Provision: The Blurring of Traditional Definitions,” *Journal of Financial Markets*, 2009, *12* (2), 143–172.

_ and _ , “Low-Latency Trading,” *Journal of Financial Markets*, 2013, *16* (4), 646–679.

Hendershott, Terrence and Ryan Riordan, “Algorithmic Trading and the Market for Liquidity,” *Journal of Financial and Quantitative Analysis*, 8 2013, *48*, 1001–1024.

_ , Charles M. Jones, and Albert J. Menkveld, “Does Algorithmic Trading Improve Liquidity?,” *The Journal of Finance*, 2011, *66* (1), 1–33.

Hoffmann, Peter, “A Dynamic Limit Order Market with Fast and Slow Traders,” *Journal of Financial Economics*, 2014, *113* (1), 156–169.

Huh, Yesol, “Machines vs. Machines: High Frequency Trading and Hard Information,” *Working Paper*, 2014. http://ssrn.com/abstract=2436048.

Investment Industry Regulatory Organization of Canada, "Report of Marketshare by Marketplace (Historical 2007-2014)," http://www.iiroc.ca/industry/marketmonitoringanalysis/Documents/MarketplaceStatisticsReportHistorical_en.pdf 2015. Accessed: 14 Jan. 2016.

Kavajecz, Kenneth A., "A Specialist's Quoted Depth and the Limit Order Book," *The Journal of Finance*, 1999, *54* (2), 747–771.

KCG Holdings, "Exchange-Registered Market Making," https://www.kcg.com/market-making/exchange-registered-market-making 2016. Accessed: 12 Jan. 2016.

Kirilenko, Andrei A., Albert S. Kyle, Mehrdad Samadi, and Tugkan Tuzun, "The Flash Crash: The Impact of High Frequency Trading on an Electronic Market," *Journal of Finance*, forthcoming. http://ssrn.com/abstract=1686004.

Korajczyk, Robert A. and Dermot Murphy, "High Frequency Market Making to Large Institutional Trades," *Working Paper*, 2016. http://ssrn.com/abstract=2567016.

Lee, Charles and Mark J Ready, "Inferring Trade Direction from Intraday Data," *The Journal of Finance*, 1991, *46* (2), 733–746.

Malinova, Katya, Andreas Park, and Ryan Riordan, "Taxing High Frequency Market Making: Who Pays the Bill?," *Working Paper*, 2016. http://ssrn.com/abstract=2183806.

McMorris-Santoro, Evan (@EvanMcSan), "from here in the WH basement, this acct seems hacked RT @AP: Breaking: Two Explosions in the White House and Barack Obama is injured," https://twitter.com/EvanMcSan/status/326744804560084992 2013. 23 April 2013, 13:10:06. Tweet.

Menkveld, Albert J., "High Frequency Trading and the *New Market* Makers," *Journal of Financial Markets*, 2013, *16* (4), 712–740.

_ **and Marius A. Zoican**, "Need for Speed? Exchange Latency and Liquidity," *Review of Financial Studies*, forthcoming. http://ssrn.com/abstract=2442690.

Panzar, John C. and James N. Rosse, "Testing For 'Monopoly' Equilibrium," *The Journal of Industrial Economics*, 1987, *35* (4), 443–456.

Rosa, Carlo, "Walking on Thin Ice: Market Quality Around FOMC Announcements," *Economics Letters*, 2016, *138*, 5–8.

Roşu, Ioanid, "Fast and Slow Informed Trading," *Working Paper*, 2016. http://ssrn.com/abstract=1859265.

TABB Group, "High-Frequency Trading: An Important Conversation," http://tabbforum.com/opinions/high-frequency-trading-an-important-conversation 2014. Accessed: 12 Jan. 2016.

TMX Group, "Policy 4-604 Responsibilities of Market Makers," *TSX Rulebook*, 2012. http://tmx.complinet.com/en/display/display_main.html?rbid=2073&record_id=737, Accessed: 14 Jan. 2016.

_ , "TSX Markets Fee Schedule," http://www.tmx.com/en/trading/fee_schedule/ 2012. Accessed: 21 Oct. 2013.

_ , "TSX Market Makers List," http://www.tmx.com/en/trading/products_services/market_makers_list.html 2013. Accessed: 21 Oct. 2013.

Tong, Lin, "A Blessing or a Curse? The Impact of High Frequency Trading on Institutional Investors," *Working Paper*, 2015. http://ssrn.com/abstract=2330053.

Virtu Financial, "Company Overview," http://www.virtu.com/company/ 2016. Accessed: 12 Jan. 2016.

Weller, Brian, "Efficient Prices at Any Cost: Does Algorithmic Trading Deter Information Acquisition?," *Working Paper*, 2016. http://ssrn.com/abstract=2662254.