Every cloud has a silver lining: Fast trading, microwave connectivity and trading costs

Andriy Shkilko*
Wilfrid Laurier University

Konstantin Sokolov
Wilfrid Laurier University

First version: March 2016
This version: February 2017

Abstract: The modern marketplace is characterized by speed differentials, whereby some traders are faster than others. How do these differentials affect liquidity? To answer this question, we study a series of exogenous weather-related episodes that temporarily remove speed advantages of the fastest traders by disrupting their microwave networks. During these episodes, adverse selection declines accompanied by lower trading costs, reduced volatility and an overall strengthening of liquidity supply. The results are confirmed in an event-study setting, whereby a new business model adopted by one of the technology providers reduces speed differentials among traders, resulting in liquidity improvements.

JEL classification: G10; G14

*Correspondence: ashkilko@wlu.ca (A. Shkilko). We thank Jim Angel, Matt Barron, Robert Battalio, Michael Brolley, Eric Budish, Jean-Edouard Colliard, Amy Edwards (NBER discussant), Sean Foley, Thierry Foucault, Austin Gerig, Michael Goldstein, Björn Hagströmer, Terrence Hendershott, Sophie Moinas (Central Bank Workshop discussant), Peter O'Neill, Andreas Park, Lasse Pedersen, Fabricio Perez, Barbara Rindi, Matt Ringgenberg, Ryan Riordan, Eric Stockland, Tugkan Tuzun, Brian Weller, René Wells, Bart Yueshen, Yelin Zhang, Haoxiang Zhu, and the audiences at the Annual Central Bank Workshop on the Microstructure of Financial Markets, NBER Market Microstructure meeting, Northern Finance Association annual meeting, Paris Finance Meeting, the University of South Australia, University of Sydney and Wilfrid Laurier University for insightful comments. Jiacheng Zhou has provided excellent research assistance. Stuart Hinson from the NOAA National Climatic Data Center and Jireh Ray from the CME have provided detailed guidance on the use of data.

Every cloud has a silver lining:
Fast trading, microwave connectivity and trading costs

Abstract: The modern marketplace is characterized by speed differentials, whereby some traders are faster than others. How do these differentials affect liquidity? To answer this question, we study a series of exogenous weather-related episodes that temporarily remove speed advantages of the fastest traders by disrupting their microwave networks. During these episodes, adverse selection declines accompanied by lower trading costs, reduced volatility and an overall strengthening of liquidity supply. The results are confirmed in an event-study setting, whereby a new business model adopted by one of the technology providers reduces speed differentials among traders, resulting in liquidity improvements.

1. Introduction

Competition on relative speed is a defining characteristic of modern markets, where trading firms spend generously to gain sub-second speed advantages over their rivals. Speed-improving technology is expensive and sometimes only available to a select few, leading to speed differentials. A rich theory literature suggests two possible effects of such differentials on liquidity.[1] On the one hand, being faster may allow liquidity providers to avoid adverse selection and to manage inventory more efficiently. As a result, liquidity may improve. Alternatively, the differentials may allow some traders to pick off stale limit orders, impairing liquidity. To shed light on these possibilities, we examine a multi-year time series of exogenous shocks to speed differentials. The results show that when the differentials exist liquidity is impaired.

We examine information transmission between financial markets in Chicago and New York, where signals are sent via two channels: a fiber-optic cable and several microwave networks. Microwave networks are about 30% faster than cable, and have two important characteristics. First, in 2011-2012 (the first two years of our four-year sample period) they are only accessible by a select group of traders. Second, precipitation (i.e., rain and snow) disrupts them. The first characteristic creates a two-tiered market, where some traders are faster than others. The second characteristic intermittently eliminates the speed advantage of the fastest tier. We show that when the microwave networks are functional, the speed advantage is used primarily to pick off stale limit orders in the course of latency arbitrage. When precipitation eliminates the speed advantage, adverse selection and trading costs decline by more than 7%, while volatility declines by about 6%. Aggressiveness of limit orders increases, pointing to more robust liquidity supply.

[1] See Hoffmann (2014), Biais, Foucault and Moinas (2015), Foucault, Hombert and Roşu (2016), Foucault, Kozhan and Tham (2016), Menkveld and Zoican (2016), Aït-Sahalia and Sağlam (2017).

On balance, the results suggest that microwave users prefer to take liquidity rather than supply it. Many theory models recognize this preference, yet some empirical studies find that fast traders often trade via limit orders.[2] There are two possible explanations for liquidity taking in our setting. First, the execution probability of limit orders is relatively low, especially for the short-lived latency arbitrage opportunities. Second, in many liquid stocks spreads are narrow and order queues are long, further reducing execution chances. Corroborating the latter explanation, the largest reductions in adverse selection during microwave disruptions occur in assets with narrow spreads.

In addition to posting new orders, speed advantages should allow traders to timely cancel stale limit orders. The data however suggest that an average limit order trader does not possess such advantages, likely because the microwave connections are available only to a select few in 2011-2012. Notably, the status quo changes in winter of 2012-2013, when a technology provider McKay Brothers democratizes microwave transmissions. Instead of selling microwave bandwidth that select traders use to outpace others, the firm begins to use its network to transmit the latest price updates and sell them to anyone on a subscription basis. As a result, the speed advantages previously enjoyed by a select few are diminished. We find that once information transmission is democratized in this manner precipitation stops having an effect on trading costs, suggesting an elimination of the speed differential between an average liquidity taker and an average supplier. Furthermore, democratization leads to a one-time reduction in adverse selection and trading costs.

Our results point to a negative relation between speed differentials and liquidity. As such, they provide a complementary perspective to that of Brogaard, Hagströmer, Nordén and Riordan

[2] See for instance, O'Hara (2015), Yao and Ye (2015), Brogaard, Hendershott and Riordan (2016), Chordia, Green and Kottimukkalur (2016).

(2015), who show that in the Swedish market speed differentials resulting from colocation are mainly sought by market makers and therefore benefit liquidity. The authors suggest that although colocation is most attractive to market makers, technology that increases information transmission speeds between markets may be sought by other traders, such as latency arbitragers, leading to negative liquidity effects. Our study also corroborates the findings of Baron, Brogaard, Hagströmer and Kirilenko (2016) and Foucault, Kozhan and Tham (2016), who suggest that modern arbitragers often use marketable orders, thus increasing order flow toxicity and impairing liquidity.

The financial economics literature has previously explored the effects of weather on trader behavior, but these effects have been mainly ascribed to investor mood. Although we examine a different weather-induced regularity, a technological one, it is important that we address the possibility that our results come from slower information processing attributed to weather-induced moods of traders in Chicago and New York (deHaan, Madsen and Piotroski, 2015). To do so, we show that our results are robust to focusing exclusively on precipitation in Ohio, a state that hosts all microwave network paths yet has a relatively low concentration of financial firms. We also confirm the robustness of the results to various sample selection procedures and to alternative precipitation variables.

Our contribution to the literature is as follows. First, we shed new light on the predictions of theory models that examine speed differentials and provide empirical evidence on the models' insights into (i) order choices of the fastest traders and (ii) the liquidity suppliers' response to lower adverse selection risk. Second, we offer evidence complementary to existing empirical research that examines the relation between speed differentials and liquidity. Some market participants claim that faster markets are unconditionally better; our results suggest that the benefits are conditional on how speed advancements are used. Finally, we describe a new approach to measuring exogenous variation in relative speed in modern markets that, to our knowledge, has not been previously examined.[3]

The remainder of the paper is as follows. Section 2 discusses the history and physics of information transmission, the state of the trading speed literature, and latency arbitrage between the futures and equity markets. Section 3 describes the data and sample. Section 4 discusses the main empirical tests. Section 5 reports robustness tests. Section 6 concludes.

2. Institutional background and related literature

2.1. History and physics of information transmission between Chicago and New York

In the world of ultra-fast trading, the physics of signal transmission plays an important role. The most common way to transmit information over long distances is via a fiber-optic cable. The first such cable between Chicago and New York was laid in the mid-1980s; however, its path was not optimal for ultra-fast communications. The cable was placed along the existing rail lines, making multiple detours from a straight line, going south to Pittsburgh and thereby exceeding the straight-line distance between Chicago and New York by about 300 miles. Realizing potential latency reduction from a more linear setup, a technology company Spread Networks laid another cable in 2010. The new cable had significantly fewer detours, went through the Appalachian Mountains and shaved valuable milliseconds off the signal transmission time.

Although fiber is a very fast transmission medium, it is not the fastest. Because microwaves travel faster through air than photons do through fiber, a network of microwave

[3] Koudijs (2015, 2016) uses adverse weather events to study information transmission between London and Amsterdam in the 18th century.

towers placed in a straight line can shave additional milliseconds off the signal transmission time. At the time of this study, microwave networks advertise round-trip information transmission speeds that are about 30% faster than their fiber-optic competitors. Specifically, during the sample period, microwave networks transmit information from Chicago to New York in about 4.5 ms, whereas fiber transmission takes about 6.5 ms.

Although faster than cable, microwaves have a disadvantage – they are relatively easily disrupted. Among the disruptors are rain droplets and snowflakes, especially when rainfall/snowfall is substantial. During such disruptions, traders who use microwave links lose their speed advantage and must either stop trading or transition from microwave to fiber transmissions. Industry insiders suggest that mainly the former happens; certain strategies are switched off when firms realize that their speed advantage is temporarily lost. The switch is automatic and does not require human involvement.

The first microwave network that linked Chicago and New York was operational at the end of 2010, with several additional networks built in 2011 and 2012. During this period, access to microwave transmission speeds was limited to a small group of trading firms, because the Federal Communications Commission restricted the number of network licenses citing airwave congestion. As such, the 2011-2012 period provides us with a unique opportunity to examine a two-tiered marketplace where some traders have access to the fastest speeds and others do not. Our results linking precipitation episodes to lower adverse selection and trading costs come from this period.

2.2. Information transmission speed and market quality

The speed-related effects have been extensively modeled in recent literature; Menkveld (2016) provides a comprehensive review. Bernales (2014), Biais, Foucault and Moinas (2015), Foucault, Hombert and Roşu (2016) and Foucault, Kozhan and Tham (2016) model a market where speed differentials result in fast traders' generating adverse selection for slower limit order traders. Limit order traders in turn seek higher compensation for providing liquidity, thereby increasing liquidity costs for all market participants. Budish, Cramton and Shim (2015) and Menkveld and Zoican (2016) show that, even in absence of speed differentials between the fast traders, sequential order processing and increases in exchange engine speeds lead to adverse selection of liquidity providers, increasing transaction costs.

Hoffmann (2014) and Jovanovic and Menkveld (2015) show that when some market makers become fast they avoid being adversely selected and increase liquidity supply. In Hoffmann (2014) however, slower market makers become more exposed to adverse selection and widen their quotes. Depending on the relative size and competitiveness of the two groups, speeding up of select market makers may have both positive and negative consequences. Bongaerts, Kong and Van Achter (2016) show that both liquidity takers and liquidity makers will engage in speed competition. Roşu (2015) and Du and Zhu (2016) suggest that when some traders are faster than others, volatility may increase.

2.3. Information flow between futures and equities

We focus on information transmission between Chicago and New York. In the U.S., most futures contracts trade on the Chicago Mercantile Exchange (CME), particularly in its data center in Aurora, IL. Meanwhile, equities mainly trade at data centers that are located in New Jersey, close to New York City. During our sample period, the NYSE data center is in Mahwah, NJ; Nasdaq data center is in Carteret, NJ; BATS is in Weehawken, NJ; and Direct Edge is in Secaucus, NJ. To continue with academic tradition, throughout the paper we refer to the two locales as Chicago and New York.

Information transmission between the two market centers is driven by fast arbitrageurs. Our data show that when microwave technology allows these arbitrageurs to speed up, both price impacts and trading costs increase. This result may appear counterintuitive because arbitrageurs are often viewed as liquidity providers who enhance market efficiency. Several theory models suggest that arbitrageurs may respond to supply and demand shocks faster and more effectively than traditional market makers thereby improving liquidity (Holden, 1995; Gromb and Vayanos, 2002, 2010). Guided by the insights of Grossman and Stiglitz (1980), these models assume that arbitrageurs are passive and provide liquidity when it is required by noise traders.

Recent theory relaxes this assumption and allows arbitrageurs to demand liquidity when it is profitable. Foucault, Kozhan and Tham (2016) model a market in which arbitrageurs are faster than market makers. When arbitrageurs trade to enforce the law of one price, they often expose market makers to adverse selection risk. As in Copeland and Galai (1983), market makers require compensation for the risk of being adversely selected, and liquidity becomes more expensive. Foucault, Kozhan and Tham (2016) conclude that although arbitrage makes prices more efficient, it may hurt liquidity. This conclusion echoes the result in Roll, Schwartz and Subrahmanyam (2007), who find that arbitrage opportunities Granger-cause illiquidity.

3. Data and sample

Our main analysis is based on the millisecond DTAQ data. The sample period spans four years, from January 2011 through December 2014. The first two years (2011-2012) are characterised by limited access to microwave technology. The latter period (2013-2014) captures the time after the technology was democratized.

To achieve the fastest speeds, microwave networks follow paths that are as straight as possible and therefore rather similar. For illustration, Figure 1 reports tower locations of three select networks connecting Chicago to the New York data centers. The data on tower locations are obtained from the Federal Communications Commission (https://www.fcc.gov). Going east from the CME data center, the networks pass through Illinois, Indiana, Ohio, western Pennsylvania and then split in eastern Pennsylvania, with the southern branches going to Nasdaq's data center in Carteret and the northern branches going to the NYSE in Mahwah. To avoid clutter, Figure 1 maps three microwave networks; FCC data show that all networks follow similar paths.

[Figure 1]

3.1. Precipitation data

We obtain precipitation data from the National Oceanic and Atmospheric Administration (http://www.noaa.gov). The data contain precipitation statistics collected by weather stations across the U.S., in 15-minute intervals. The data also contain precise station locations. The stations report in local time, so for stations in Illinois and northwestern Indiana located in the Central time zone we add one hour to report times to match DTAQ time stamps. A standard piece of equipment at every station is a precipitation tank equipped with an automatic gauge that measures accumulated precipitation. We focus on data collected by 83 stations located along the Chicago-New York corridor (Figure 2). In the robustness section, we examine station samples of different sizes.

[Figure 2]

We note that although it may only rain over Indiana or Ohio, the entire microwave network will be disrupted. A relatively narrow weather front like the one in Figure 3 will result in weather stations located within the front reporting high levels of precipitation. In the meantime, stations located outside the front will report no precipitation. To capture relatively narrow bands of intense precipitation, our main independent variable $PRECIP$ is computed as the sum of precipitation amounts reported by all stations. We examine alternative specifications in the robustness section.

[Figure 3]

Statistics reported in Panel A of Table 1 indicate that an average 15-minute sampling interval sees 0.155 mm of precipitation. The distribution is rather skewed, with a median of 0.07, indicating that periods of low precipitation are occasionally interrupted by significant rain or snow. We note that microwave networks are only disrupted when precipitation is substantial. We therefore focus on high levels of precipitation and compute two additional metrics, $PRECIP1$ and $PRECIP2$, that capture intervals when precipitation is 0.5 and 1 standard deviations above the mean. The two groups contain, respectively, 17% and 10.5% of all intervals, and $PRECIP1$ and $PRECIP2$ events last on average 54 and 49 minutes. As such, significant precipitation is observed rather frequently but ends quickly, forming a time series with sufficient variability.

[Table 1]

3.2 Asset samples

The importance of information flows between the futures markets in Chicago and the equity markets in New York is well recognized in the literature. Some studies find that futures markets lead price discovery (Kawaller, Koch, and Koch, 1987; Chan, 1992). Others suggest that information may flow both ways (Chan, Chan and Karolyi, 1991; Roll, Schwartz and Subrahmanyam, 2007). Hasbrouck (2003) shows that the direction of information flow depends on the futures trading activity; for the most active contracts, futures dominate price discovery.

Given that the most active futures contracts track baskets of securities, our focus in the equity market is on the ETFs. As long as price discovery via futures is non-trivial, the speed of information transmission between Chicago and New York should matter for trading costs in ETFs. In a later section, we examine the direction of price discovery between the two markets in more detail.

We use millisecond DTAQ data for a sample of 100 most actively traded ETFs. Among these, 50 ETFs track U.S. equity indexes; 22 – international indexes; 20 – corporate or treasury interest rate indexes; 4 – metals (i.e., gold and silver); 1 – a real estate portfolio; and 3 – other assets (Panel B of Table 1). Many ETFs in our sample track the same baskets of securities as the CME futures contracts (e.g., the QQQ ETF and the E-mini Nasdaq-100 futures). Others track baskets similar to those of major CME contracts. For example, the iShares Russell 1000 ETF does not have a corresponding CME futures contract; however, a portion of price discovery in this ETF comes from futures on other indexes such as the S&P 500.[4]

3.3. DTAQ data and summary statistics

Following Holden and Jacobsen (2014), we combine the DTAQ NBBO and Quote files to obtain the complete NBBO record and merge the resulting dataset with the Trade file. We sign trades using the Lee and Ready (1991) algorithm and exclude the first and the last five minutes of each trading day to avoid the influence of the opening and closing procedures. Panel A of Table 2 reports market activity statistics. Precipitation data are in 15-minute intervals, and we aggregate the statistics accordingly. An average ETF has 5,305 NBBO updates every 15 minutes, equivalent to about 6 updates per second. In addition, this ETF trades 500 times every 15 minutes, for a total volume of 190,522 shares.

[4] The CME delisted E-mini Russell 1000 futures contract in 2007 and relisted it in 2015.

[Table 2]

3.4. Precipitation and information transmission speed

Our empirical tests are based on the premise that precipitation disrupts microwave networks and slows down the fastest traders. In Figure 4, we illustrate such disruptions. The figure reports the number of equity trades that follow a futures trade when precipitation is zero or very low ($PRECIP2$<0, orange bars), and when it is high ($PRECIP2$>1, blue bars). The number of trades is standardized for each asset to allow for cross-sectional comparability. To reduce serial correlation effects, we focus on the standalone futures trades, those not preceded by another futures or equity trade in the previous 100 milliseconds. Futures trades come from a 2012 CME dataset that contains four E-mini contracts: S&P 500, S&P MidCap 400, Nasdaq 100 and Financial Sector Select. Intraday data for these contracts are sold by the CME as a bundle. ETF trades are from DTAQ.

[Figure 4]

Figure 4 shows that when precipitation is low, trading activity in the equity market picks up 5 ms after a futures trade. Meanwhile, when precipitation is heavy, equity trading begins 7 ms after a futures trade. During the sample period, microwave networks transmit information from Chicago to New York in about 4.5 ms, whereas fiber transmission takes about 6.5 ms. As such, the results corroborate the notion that precipitation slows traders by the 2-ms difference between the microwave and fiber speeds.

Until mid-2015, U.S. exchanges are not required to synchronize their clocks (Bartlett and McCrary, 2016). During our sample period, the CME clock lags DTAQ by about one millisecond, and we adjust for this lag. Without the adjustment, it appears that equity trading during periods of zero precipitation picks up 4 ms after a futures trade, which is not possible as it would require microwave speed to be equal to or faster than the speed of light. Importantly, the adjustment does not affect the evidence of the 2-ms speed advantage of the fastest traders. This is because the adjustment affects both $PRECIP2<0$ and $PRECIP2>1$ equally, and therefore the difference between them remains the same.

3.5. Picking-off risk

Recent literature suggests that fast informed traders often trade via limit orders. Brogaard, Hendershott and Riordan (2016) show that limit orders submitted by fast traders play a significant role in price discovery. O'Hara (2015) also suggests that fast informed traders often prefer limit to marketable orders. Both studies however point out that most traders do not resort to one order type exclusively, but rather use them interchangeably depending on the circumstances.

One of such circumstances is the constraint introduced by the minimum tick size. A binding tick size provides a strong incentive for fast traders to use marketable orders. Assume that a fast trader learns that an asset is underpriced. She wants to buy, but if the tick size is binding she cannot raise the outstanding bid without locking or crossing the market. Given these considerations, and if her signal is sufficiently strong, she may choose to consume liquidity (pick off the outstanding ask quote) despite having to pay the spread. As such, picking-off risk may be higher in assets with binding tick sizes.

The very active ETFs in our sample are quite liquid and therefore are likely to be constrained by the minimum tick size. Panel B of Table 2 shows that the average NBBO is 1.9 cents, with a median of 1.2 cents. Given these constraints, trade-related price discovery and the associated picking-off risk may be important. In subsequent tests, we subdivide assets into two categories: most and least constrained. To do so, we divide the assets into terciles according to the average NBBO. An average (median) NBBO in the first tercile is 1.0 (1.0) cent, whereas it is 3.6 (1.8) cents in the third tercile. We define the first tercile as the most tick constrained and the third tercile as the least constrained.

To further examine the issue of picking-off risk, we compute two metrics. First, we estimate a share of price discovery attributable to trades. Second, we compute trade price impacts. The former metric follows Hasbrouck's (1991 a,b) and decomposes the efficient price variance into the trade-related and trade-unrelated components. The details of this calculation are in the Appendix. Panel B of Table 2 shows that the trade-related component is 29.6%. As such, new information is incorporated into prices through trades rather frequently, and therefore concerns with the picking-off risk are warranted.

Our second proxy for the picking-off risk is the conventional price impact metric, computed on a round-trip basis as twice the signed difference between the midquote at a time after the trade and the midquote at the time of the trade: $PRIMP_t = 2q_t\left(mid_{t+\gamma} - mid_t\right)$, where q_t is the Lee and Ready (1991) trade direction indicator, mid_t is the midquote computed as $(NBBO\ Ask_t + NBBO\ Bid_t)/2$, and γ indicates the time elapsed since the trade. Recent research uses γs of just a few seconds. For instance, O'Hara (2015) suggests that 5- to 15-second intervals may be the most useful, whereas Conrad, Wahal and Xiang (2015) use price impacts up to 20 seconds.

To check if intervals of these lengths are practical in our setting, Figure 5 traces price impacts for 60 seconds after a trade. The results clarify our understanding of price dynamics on two levels. First, the data show that price impacts are greater than zero, corroborating the earlier assertion that non-trivial amounts of information are incorporated into prices through trades.

Second, although a significant proportion of information is incorporated into quotes within a second after the trade, incorporation continues at a slower pace up to 60 seconds. In the remainder of the tests, we focus on 15-second intervals, with robustness checks examining intervals between 1 and 60 seconds.

[Figure 5]

It may not be immediately obvious that there is enough adverse selection in ETFs to warrant non-zero price impacts. We suggest that as long as sufficient amounts of macro information are present, price impacts in ETFs may be quite sizeable. In a study that examines a recent sample of large U.S. equities, Chakrabarty, Jain, Shkilko and Sokolov (2016) report price impacts that are 35% of the effective spread. In Table 2, the ETF price impacts are 31% of the effective spread. As such, adverse selection is a non-trivial component of ETF trading costs and is comparable to the levels found in equities.

3.6. Trading costs and liquidity provider revenues

Table 2 also reports liquidity costs and liquidity provider revenues proxied by effective spreads, ESP_t, and realized spreads, RSP_t. ESP is computed as twice the signed difference between the prevailing midquote and the trade price, p_t: $ESP_t = 2q_t(p_t - mid_t)$. RSP_t is computed as the difference between the effective spread and the price impact. We volume-weight effective and realized spreads. The average (median) effective spread is 1.9 (1.0) cents and the average (median) realized spread is 1.3 (0.7) cents.

4. Empirical findings

4.1. Connectivity disruptions and picking-off risk

When the microwave networks are fully functional, their users have a speed advantage.

Theory models make several assumptions as to how this advantage may be used. Some models assume that the fastest traders can better manage adverse selection risk. Others suggest that speed advantages are used to generate such risk by picking off slower traders. In this section, we aim to better understand which of these assumptions prevails.

If speed advantages allow fast traders to pick off outstanding limit orders, connectivity disruptions should result in lower price impacts. Alternatively, if fast connections are used to incorporate the latest information into quotes, the disruptions may be accompanied by larger price impacts. Certainly, it is possible that both explanations have merit, and our data allow us to gauge which of them prevails. We focus on the 2011-2012 period when the microwave networks allowed for speed differentials among traders. The post-democratization period (2013-2014) is examined in a later section. Chung and Chuwonganant (2014) and Malinova, Park and Riordan (2014) argue that VIX is a first-order determinant of trading activity and liquidity, and we use their insight in a regression setup as follows:

$$DEPVAR_{it} = \alpha_0 + \beta_1 PRECIP_t + \beta_2 VIX_t + \varepsilon_{it}, \tag{1}$$

where $DEPVAR$ is the price impact; $PRECIP$ is total precipitation in the Chicago-New York corridor; and VIX is the intraday volatility index proxied by the iPath S&P500 VIX ST Futures ETF that tracks VIX. As discussed earlier, we also use $PRECIP1$ and $PRECIP2$ to identify the most significant precipitation events. All asset-specific variables are standardized (by demeaning and scaling by the standard deviation for each stock), so regression models control for asset fixed effects. Additionally, the standard errors are double-clustered along the asset and time dimensions.

Table 3 shows that price impacts decline during network disruptions. Significant amounts of precipitation captured by $PRECIP2$ are associated with a 0.047 standard deviations, or 7.1%, decline in price impacts (Panel A). It therefore appears that microwave users prefer marketable orders to limit orders.

[Table 3]

Per our earlier suggestion, marketable orders may be the only choice when the tick size is binding. If this is so, microwave network disruptions will have a larger effect on price impacts in the most constrained assets. The results in Panel B are consistent with this expectation. Price impacts in the most constrained ETFs decline by 0.051 standard deviations, whereas they decline by only 0.039 standard deviations in the least constrained ETFs. It therefore appears that fast traders use more limit orders when the tick size allows. This said, even in the least constrained ETFs liquidity taking is preferred by microwave users as evident from the decline in price impacts.

4.2. Trading costs and liquidity provider revenues

To the extent that liquidity providers use speed advantages to avoid being adversely selected, microwave connectivity disruptions should cause them to widen spreads. If however speed advantages are used mainly to pick off standing orders, precipitation should result in reduced adverse selection, and spreads may narrow. The tests discussed earlier provide support to the picking-off story, so we expect trading costs to decline during precipitation episodes. We however note that it is not clear how quickly liquidity providers adjust to lower adverse selection, and if they adjust at all given that precipitation episodes are relatively short. Easley and O'Hara (1992) describe price adjustment as gradual learning. In their model, market makers do not immediately know if informed traders are active, but learn over time. Whether such learning happens in our setting is an empirical question.

In Table 3, we report eq. 1 coefficient estimates for effective and realized spreads.

Effective spreads decline by 0.043 standard deviations, or 7.2%, during heavy precipitation episodes ($PRECIP2$ in Panel A). Expectedly, this result is more pronounced for the least constrained assets (Panel B), where the spreads have room to decline. The results are consistent with predictions of the models that emphasize the picking-off risk and are also informative about the speed of adjustment to changing levels of adverse selection. Specifically, the length of an average precipitation episode appears sufficient for liquidity providers to adjust. It is however unclear if this adjustment reflects an equilibrium. In a later section, we report results based on an exogenous shock that resulted in the long-term reduction in speed differentials. This shock further improves our understanding of the equilibrium effects.

We note that there is an alternative, mechanical, explanation to the decline in effective spreads. Aggressively priced limit orders may be added to the book equally often when the microwaves are up and when they are down. Since such orders are consumed less when the arbitrage strategies are switched off, they are more readily available to the remaining liquidity takers, who obtain better prices. In later sections, we examine trades and limit order book data to show that both explanations have merit.

Realized spreads also decline, by 0.021 standard deviations, or 5.3%, during $PRECIP2$ events. As such, network disruptions not only reduce liquidity costs, but also reduce liquidity provider revenues. Similarly to the effective spread result, there are two possible explanations. On the one hand, liquidity supply may become more robust when picking-off risk is reduced. Chakrabarty, Jain, Shkilko and Sokolov (2016) show that when the adverse selection risk declines liquidity providers reposition orders from the deeper layers of the book to the inside. On the other hand, the effect may be due to the minimum tick size. When microwaves allow arbitrageurs to pick off inside quotes, the spread increases by at least one tick. Given the coarseness of this change, the resulting realized spreads may be unduly large.

4.3. Trading activity and volatility

The literature often assumes that lower trading costs attract additional trading interest and therefore result in higher trading volume. In our setting, this assumption will not necessarily hold. This is because aside from lower costs, network disruptions lead to a reduction in the number of picking-off opportunities and consequently the volume generated by latency arbitrage. The regression results in Panel A of Table 4 are consistent with this notion. The number of trades declines by 0.072 standard deviations during $PRECIP2$ events. Trading volume also declines; by 0.042 standard deviations.

[Table 4]

The finance literature has not yet come to a consensus on the relation between electronic trading and volatility. While some studies report that the relation is negative (Hasbrouck and Saar, 2013; Brogaard, Hendershott and Riordan, 2014), others find it to be positive (Boehmer, Fong and Wu, 2015). Closest to our setting, a theory model by Roşu (2015) suggests that as fast traders pick off market makers' quotes, volatility may increase. Du and Zhu (2016) also show that when some traders are faster than others, liquidity shocks result in greater volatility. Our results are consistent with these insights; volatility, which we define as the difference between the high and low prices during an interval scaled by the average price, declines by 0.118 standard deviations, or 5.8%, during $PRECIP2$ events.

In assets with wider spreads new information may be incorporated into prices through both marketable and aggressive limit orders. Meanwhile, in assets whose spreads are constrained by the minimum tick size, fast traders must rely on marketable orders. Naturally, these considerations should affect changes in trading activity during network disruptions. Panel B shows that the number of trades and trading volume decline in the most constrained assets, yet remain unchanged in the least constrained assets. These results corroborate two of our earlier conjectures: (i) fast traders use fewer marketable orders in the least constrained ETFs, and (ii) trading volume may increase in response to lower trading costs compensating for some of the volume lost when the arbitrage strategies are switched off.

Overall, the results suggest that even though lower spreads may attract additional trading interest, trading volume generated by this interest is smaller or equal to the lost arbitrage volume. One possibility is that the disruptions are not long enough or not sufficiently predictable for additional trading interest to emerge. A trading strategy that is highly sensitive to transaction costs may not be viable in a high cost environment, even if high cost periods are occasionally interrupted by low cost periods. This said, an extended period of lower spreads may make the strategy viable, thus generating new trading interest. In a later section, we examine this possibility by studying an event that resulted in a long-lasting loss of speed advantage by the network users.

4.4. Limit order aggressiveness

Earlier, we suggest that the reduction in trading costs may have two explanations: (i) strengthening of liquidity supply and (ii) fewer realizations of picking-off opportunities. The declines in price impacts and trades reported earlier corroborate the latter explanation. In this section, we use ITCH limit order book data to examine the former one.

We ask if the proportion of aggressively priced limit orders increases during microwave disruptions. We compute two metrics: the number of orders that (i) match the prevailing NBBO and (ii) the number of orders that match or improve the NBBO. We scale both metrics by the number of total order submissions. Eq. 1 coefficients reported in Table 5 suggest that the number of aggressively priced limit orders increases during precipitation episodes. For orders that match the NBBO, the increase is seen in the full sample as well as the most and the least constrained subsamples (Panel A). For orders that match or improve the NBBO, the coefficients are insignificant for the ETFs with narrow spreads. This result is expected given that improving the best quotes in such assets is largely impossible once they narrow during precipitation episodes.

[Table 5]

4.5. Information asymmetry in the futures market

Do speed differentials also affect the picking-off risk in futures contracts? On the one hand, fast traders may carry information both from futures to equities and in the opposite direction. On the other hand, prior research suggests that futures provide the lion's share of price discovery in index instruments. If so, using limited microwave bandwidth to transmit information from ETFs to futures may be wasteful. If this is the case, speed differentials may not have much of an effect on the CME.

To examine this issue, we use the 2012 CME data described in Section 3.4 and compute information shares as in Hasbrouck (1995) for the four futures-ETF pairs. The details of the methodology are in the Appendix. The results are consistent with earlier studies, in that price discovery occurs mainly on the CME; the CME information shares are in the [0.64; 0.82] range. Second, in Panel A of Table 6 we examine the price impacts in futures during precipitation episodes and find no relation.

[Table 6]

Statistical power is a concern in the analysis above given the small size of the cross-section. To address this concern, in Panel B we replicate the results for the sample of four corresponding ETFs. Similarly to the main sample, precipitation is associated with reductions in price impacts.

4.6. Democratization of MWN access

In late December 2012 – early January 2013, a microwave technology provider McKay Brothers, disrupted the business model used by the microwave firms. Instead of selling bandwidth on its network to select traders, McKay Brothers began selling information transmitted by the network to everyone who was willing to pay a nominal fee. Subscribers to this service obtained access to an affordable and non-exclusive channel of information transmission that was among the fastest in the industry. The offer was soon replicated by other providers, and the market for microwave transmissions was democratized. Put differently, microwave users lost the speed advantage they had in 2011-2012.

Democratization of access to microwave transmission speeds may lead to two outcomes. First, the relation between precipitation and market quality observed in 2011-2012 may diminish because access to superior speeds is no longer limited to a small group of traders. Second, democratization may result in market quality changes similar to those observed during precipitation events. In this sense, precipitation episodes in 2011-2012 may be viewed as periods of short-term democratization, whereas the 2012-2013 event may be viewed as long-term democratization.

In Table 7, we report the coefficients of the $PRECIP2$ variable obtained from estimating eq. 1 during the post-democratization period. The results confirm expectations. Precipitation episodes no longer have an effect on price impacts, effective spreads, realized spreads, volatility and trading activity. The change is observed for the full sample and for the most and least constrained subsamples. As such, democratization appears to be a significant market disruptor.

[Table 7]

Given the significance of democratization, the liquidity effects associated with the loss of the speed advantage may reappear around the event. To examine if this is the case, we estimate an event study regression that compares market quality and activity variables in a three-month pre-event window (September – November 2012) and a three-month post-event window (February – April 2013). We exclude December 2012 and January 2013 to allow for a transition period, however the results are similar when these months are included. The regression is set up as follows:

$$DEPVAR_{it} = \alpha_0 + \gamma_t + \beta_1 POST_t + \beta_2 VIX_t + \varepsilon_{it}, \qquad (2)$$

where $DEPVAR_{it}$ is one of the following variables (price impacts, effective spreads, realized spreads, the number of trades, traded volume, volatility and stock price) in asset i on day t, γ_t denotes a time trend, $POST$ is a dummy variable that equals to one in February-April 2013, and VIX is the volatility index. All variables are standardized. The model controls for asset fixed effects, and the standard errors (in parentheses) are double-clustered along the asset and time dimensions.

The main variable of interest in eq. 2 is $POST$ as it captures the difference between the pre- and post-democratization periods. Consistent with expectations, the results for the full sample in Table 8 indicate that price impacts, effective spreads, realized spreads and volatility decline post-democratization. We must however note an important caveat. Democratization was a single event that affected all assets at the same time. As such, we are unable to eliminate the possibility that the results are driven by a confounding event(s) unrelated to democratization. Although we are unaware of any such events, the event study results should be interpreted with due caution.

[Table 8]

A notable difference between the event study and the panel findings discussed earlier comes from the trade-related variables: the number of trades and volume. Recall that in the full sample these variables decline during precipitation events (Table 4), likely because the incidences of latency arbitrage are reduced. In Table 8, these variables do not change post-democratization in the full sample, corroborating our earlier suggestion that lower trading costs may encourage new trading interest over long periods of time. The new interest offsets the loss of arbitrage volume. The results pointing to new trading interest are observed even in the sample of the most constrained ETFs where trade and volume losses during precipitation-related disruptions are the largest. Although Table 8 shows a declining number of trades in the most constrained assets, the change in volume is statistically insignificant, suggesting that new trading interest arises even in these assets.

The market structure literature often suggests that lower trading costs translate into higher stock prices (e.g., Amihud and Mendelson, 1986; Easley, Hendershott and Ramadorai, 2014). If so, liquidity improvements caused by democratization may result in price appreciations. Table 8 indicates that such appreciations indeed occur.

Given the private benefits from exclusive network access that accrued to the select few trading firms in 2011-2012, democratization is a curious case of the market fixing itself. It is however not immediately clear why McKay Brothers chose to disrupt the status quo. Cespa and Foucault (2014) note that it is in a data provider’s best interest to restrict dissemination of pricing data only to select traders. Their model shows that if information is accessible to many, it is less valuable to the few who may be willing to pay a premium for the exclusive use. It is our understanding that McKay Brothers was driven by the following two motives. First, the firm believed that even if others were to replicate its offering, there would be enough pricing information for everyone to transmit given the large numbers of traded instruments and the low bandwidth of microwave links. Second, the firm argued that growing its customer base was more profitable than providing restricted access to a small group of clients. As long as the firm maintained its latency advantage, it expected to always retain its customers. Many other connectivity providers launched similar offerings in the months after democratization, possibly driving down the price of the service. This said, McKay Brothers continues to offer the service to this day and has expanded it to several continents.

5. Robustness

A rich literature examines the effects of weather on the behavior of market participants and finds that poor weather is associated with investor pessimism, which is reflected in stock returns (Hirshleifer and Shumway, 2003). The pessimism affects even the sophisticated investors (Goetzmann, Kim, Kumar and Wang, 2015). Furthermore, deHaan, Madsen and Piotroski (2015) show that pessimistic moods induced by poor weather often delay equilibrium price adjustments. As such, the reduction in adverse selection during precipitation episodes may be attributed (at least in part) to slower price discovery caused by the poor weather in Chicago and/or New York rather than to the microwave disruptions.

To examine this possibility, we recalculate the $PRECIP2$ variable to capture periods when the networks are disrupted, yet the moods of traders in Chicago and New York are not affected. Specifically, we compute $PRECIP2$ that satisfies the following two conditions: (i) only weather stations in Ohio indicate high levels of precipitation, and (ii) weather stations in the western and eastern parts of the Chicago-New York corridor indicate near-zero precipitation. We then re-estimate eq. 1 for the 2011-2012 sample and report the results in the *mood control* specification in Panel A of Table 9. The effects are consistent with those reported in the earlier tables. As such, trader moods do not seem to be the source of our findings.

[Table 9]

Our sample of weather stations is selected to capture the area closely surrounding the microwave paths. As with any such selection procedure, it is important to show that the results are not driven by the specific set of stations. The *mood control* specification takes the first step in this direction by restricting the sample to the Ohio stations. In two additional Table 9 specifications, we show that using information from an expanded area surrounding the MWN paths leads to similar conclusions, while precipitation in the placebo area over Colorado, Utah and Wyoming (far removed from the Chicago-New York corridor) has no effect on the variables of interest.

Information asymmetry, trading costs and trading activity vary throughout the day. For instance, effective spreads follow an intraday J-pattern, with wider spreads in the morning that become narrower in the afternoon (Figure 6). Notably, intraday precipitation too follows a reverse J-pattern, with precipitation amounts being lower in the morning hours. Since the results in the previous section point to a negative relation between precipitation and spreads, we need to establish that the findings are not due to these intraday patterns.

[Figure 6]

We examine this possibility in two additional specifications in Panel A of Table 9. First, we focus on the afternoon period, when spreads and precipitation are relatively flat. Our results hold for every variable of interest. Second, the results continue to hold when we add intraday fixed effects to eq. 1. As such, the relations between precipitation and spreads observed in the earlier sections are independent of intraday patterns.

Recall that the $PRECIP$ variable estimates total precipitation in the Chicago-New York corridor. This variable is well-suited to capture periods of high precipitation over small areas, but may occasionally acquire high values if relatively minor precipitation extends over the entire corridor. This possibility is the reason for our focus on $PRECIP2$ that captures very high precipitation totals not likely to be achieved through anything other than significant precipitation. To provide another alternative to $PRECIP$, in Panel B of Table 9 we report the results using the average precipitation per station, $MPRECIP$, and its variations, $MPRECIP1$ and $MPRECIP2$, that capture periods when average precipitation is 0.5 and 1 standard deviations above the mean. We note that although these variables mitigate the abovementioned concern, they potentially reduce our ability to detect relatively narrow bands of strong precipitation, especially those accompanied by near-zero precipitation in the rest of the corridor (Figure 3). Corroborating this reasoning, the results for $MPRECIP$ are weaker than those reported earlier for $PRECIP$, yet the results for $MPRECIP1$ and $MPRECIP2$ are equally as strong as those for their counterparts computed using total precipitation.

In the main analysis, we compute the effective spreads and their components on a volume-weighted basis. As such, large trades have a stronger effect on the estimates than small trades. To shed more light on the effects of network disruptions on small trades, in Table 10 we report eq. 1 regression results for the equally-weighted variables (specifications $EW\_$). The results reported earlier hold.

[Table 10]

The results for the volume-weighted effective spreads and their components reported in earlier tables use raw dollar metrics. Naturally, raw spreads may vary in the price of the asset. Although our regressions account for the overall price levels by using asset fixed effects, intraday price changes remain unaccounted for. The $VWP\_$ specifications in Table 10 address this issue using effective spreads, price impacts and realized spreads scaled by the midquote at the time of trade. The results corroborate those reported in the earlier tables.

In the previous sections, we discuss the effects of network disruptions on effective spreads. We also show that the effects differ between the assets most and least constrained by the minimum tick size. In Table 11, we estimate eq. 1 for two additional variables – the quoted NBBO spread and quoted depth. Whereas effective spreads capture the realized trading costs, the quoted spreads summarize liquidity that is available at all times. As long as investors choose to trade when costs are low, effective spreads may not be fully indicative of changes in available liquidity. Table 11 shows that quoted spreads decline when the networks are down across all sample groups. The coefficients repeat the patterns reported for effective spreads, with quoted spreads declining more for the least constrained ETFs, in which more price improvement is possible.

[Table 11]

Table 11 also reports the results for quoted NBBO depth, which increases during network disruptions, but only for the most constrained ETFs. This finding is consistent with the earlier discussion. In the most constrained ETFs, there is not always room to improve the spread, so liquidity improvements are reflected in quoted depth. There is certainly an alternative mechanical explanation in that less depth is consumed by arbitrage strategies during precipitation episodes, therefore allowing for larger depth averages. Notably, depth does not increase in the least constrained ETFs.

6. Conclusions

This study examines the effects of speed differentials on liquidity. During our sample period, microwave networks located between Chicago and New York allow for the fastest information transmission and are only available to select trading firms. When it rains or snows in the area between the two cities, the networks are disrupted because rain droplets and snowflakes block the microwave paths. With the networks temporarily down, information transmission falls back onto the fiber-optic cable – a more reliable, yet slower transmission medium – effectively eliminating the speed advantages of the fastest traders. We show that when this happens, adverse selection and trading costs decline. This result is consistent with predictions of theory models that show that speed differentials among traders may be associated with lower liquidity.

Our results also shed new light on traders' order choices. Recent research suggests that informed fast traders may prefer to trade via limit orders. Our results confirm that this is the case, yet this preference varies in the cross-section. Specifically, in assets with binding tick sizes, trading on short-lived information through limit orders is difficult due to long queues. In such assets, traders prefer marketable orders.

Our results are confirmed in an event-study setting. In winter of 2012-2013, a technology firm, McKay Brothers, democratized microwave transmissions by introducing a new business model. Instead of selling bandwidth on its network, the firm began selling information on both sides of the Chicago-New York corridor. This event had positive liquidity consequences similar to precipitation-related network disruptions.

The technological race continues to drive spending in the trading industry. Recent examples include new data transmission towers to connect the U.K. and European markets. The towers will be among the tallest structures in the U.K. and will rival the Eiffel Tower. They will provide trading firms with a completely unobstructed optical and radio line of sight, never previously offered in Europe, increasing signal transmission speed. In the meantime, traders in the U.S. have been switching from microwave transmissions to more reliable, yet costly, laser links. Our findings shed light on the possible consequences of these developments.

References

Aït-Sahalia, Yacine, and Mehmet Sağlam, 2017, High frequency market making: Optimal quoting, Working paper.

Amihud, Yakov, and Haim Mendelson, 1986, Asset pricing with liquidity risk, Journal of Financial Economics 77, 375-410.

Baron, Matthew, Jonathan Brogaard, Björn Hagströmer and Andrei Kirilenko, 2016, Risk and return in high-frequency trading, Working paper.

Bartlett, Robert, and Justin McCrary, 2016, How rigged are stock markets? Evidence from microsecond timestamps, Working paper.

Bernales, Alejandro, 2014, Algorithmic and high frequency trading in dynamic limit order markets, Working paper.

Biais, Bruno, Thierry Foucault, and Sophie Moinas, 2015, Equilibrium fast trading, Journal of Financial Economics 116, 292-313.

Boehmer, Ekkehart, Kingsley Fong, and Julie Wu, 2015, International evidence on algorithmic trading, Working paper.

Bongaerts, Dion, Lingtian Kong, and Mark Van Achter, 2016, Trading speed competition: Can the arms race go too far? Working paper.

Brogaard, Jonathan, Björn Hagströmer, Lars Nordén, and Ryan Riordan, 2015, Trading fast and slow: Colocation and liquidity, Review of Financial Studies 28, 3407-3443.

Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan, 2014, High-frequency trading and price discovery, Review of Financial Studies 27, 2267-2306.

Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan, 2016, Price discovery without trading: Evidence from limit orders, Working paper.

Budish, Eric, Peter Cramton, and John Shim, 2015, The high-frequency trading arms race: Frequent batch auctions as a market design response, Quarterly Journal of Economics 130, 1547-1621.

Carrion, Allen, 2013, Very fast money: High-frequency trading on the NASDAQ, Journal of Financial Markets 16, 680-711.

Cespa, Giovanni, and Thierry Foucault, 2014, Sale of price information by exchanges: Does it promote price discovery? Management Science 60, 148-165.

Chakrabarty, Bidisha, Pankaj Jain, Andriy Shkilko, and Konstantin Sokolov, 2016, Speed of market access and market quality: Evidence from the SEC naked access ban, Working paper.

Chordia, Tarun, Clifton Green, and Badrinath Kottimukkalur, 2016, Do high frequency traders need to be regulated? Evidence from trading on macroeconomic announcements, Working paper.

Conrad, Jennifer, Sunil Wahal, and Jin Xiang, 2015, High-frequency quoting, trading, and the efficiency of prices, Journal of Financial Economics 116, 271-291.

Copeland, Thomas, and Dan Galai, 1983, Information effects on the bid-ask spread, Journal of Finance 38, 1457-1469.

Chan, Kalok, K.C. Chan, and Andrew Karolyi, 1991, Intraday volatility in the stock index and stock index futures markets, Review of Financial Studies 4, 657-684.

Chan, Kalok, 1992, A further analysis of the lead-lag relationship between the cash market and stock index futures market, Review of Financial Studies 5, 123-151.

Chung, Kee, and Chairat Chuwonganant, 2014, Uncertainty, market structure, and liquidity, Journal of Financial Economics 113, 476-499.

deHaan, Ed, Joshua Madsen, and Joseph Piotroski, 2015, Do weather-induced moods affect the processing of earnings news? Working paper.

Du, Songzi, and Haoxiang Zhu, 2016, What is the optimal trading frequency in financial markets? Review of Economic Studies, forthcoming.

Easley, David, Terrence Hendershott and Tarun Ramadorai, 2014, Leveling the trading field, Journal of Financial Markets 17, 65-93.

Easley, David, and Maureen O'Hara, 1992, Time and process of security price adjustment, Journal of Finance 47, 577-605.

Foucault, Thierry, Johan Hombert, and Ioanid Roşu, 2016, News trading and speed, Journal of Finance 71, 335-382.

Foucault, Thierry, Roman Kozhan, and Wing Wah Tham, 2016, Toxic arbitrage, Review of Financial Studies, forthcoming.

Goetzmann, William, Dasol Kim, Alok Kumar, and Qin Wang, 2015, Weather-induced mood, institutional investors, and stock returns, Working paper.

Gromb, Denis, and Dimitri Vayanos, 2002, Equilibrium and welfare in markets with financially constrained arbitrageurs, Journal of Financial Economics 66, 361-407.

Gromb, Denis, and Dimitri Vayanos, 2010, Limits of arbitrage: The state of the theory, Annual Review of Financial Economics 2, 251-275.

Grossman, Sanford, and Joseph Stiglitz, 1980, On the impossibility of informationally efficient markets, American Economic Review 70, 393-408.

Hagströmer, Björn, and Lars Nordén, 2013, The diversity of high-frequency traders, Journal of Financial Markets 16, 741-770.

Hasbrouck, Joel, 1991a, Measuring the information content of stock trades, Journal of Finance 46, 179-207.

Hasbrouck, Joel, 1991b, The summary informativeness of stock trades: An economic analysis, Review of Financial Studies 4, 571-595.

Hasbrouck, Joel, 1995, One security, many markets: Determining the contributions to price discovery, Journal of Finance 50, 1175-1199.

Hasbrouck, Joel, 2003, Intraday price formation in the U.S. equity index markets, Journal of Finance 58, 2375-2399.

Hasbrouck, Joel, and Gideon Saar, 2013, Low-latency trading, Journal of Financial Markets 16, 646-679.

Hendershott, Terrence, and Ryan Riordan, 2013, Algorithmic trading and the market for liquidity, Journal of Financial and Quantitative Analysis 48, 1001-1024.

Hirshleifer, David, and Tyler Shumway, 2003, Good day sunshine: Stock returns and the weather, Journal of Finance 58, 1009-1032.

Hoffmann, Peter, 2014, A dynamic limit order market with fast and slow traders, Journal of Financial Economics 113, 156-169.

Holden, Craig, 1995, Index arbitrage as cross-sectional market making, Journal of Futures Markets 15, 423-455.

Holden, Craig, and Stacey Jacobsen, 2014, Liquidity measurement problems in fast, competitive markets: Expensive and cheap solutions, Journal of Finance 69, 1747-1785.

Jovanovic, Boyan, and Albert Menkveld, 2015, Middlemen in limit order markets, Working paper.

Kawaller, Ira, Paul Koch, and Timothy Koch, 1987, The temporal price relationship between S&P 500 futures and the S&P 500 index, Journal of Finance 5, 1309-1329.

Koudijs, Peter, 2015, Those who know most: Insider trading in 18th century Amsterdan, Journal of Political Economy 123, 1356-1409.

Koudijs, Peter, 2016, The boats that did not sail: Asset price volatility in a natural experiment, Journal of Finance 71, 1185-1226.

Lee, Charles, and Mark Ready, 1991, Inferring trade direction from intraday data, Journal of Finance 46, 733–746.

Malinova, Katya, Andreas Park, and Ryan Riordan, 2014, Do retail traders suffer from high frequency traders, Working paper.

Menkveld, Albert, 2016, The economics of high-frequency trading: Taking stock, Annual Review of Financial Economics, forthcoming.

Menkveld, Albert, and Marius Zoican, 2016, Need for speed? Exchange latency and liquidity, Review of Financial Studies, forthcoming.

O'Hara, Maureen, 2015, High frequency market microstructure, Journal of Financial Economics 116, 257-270.

Roll, Richard, Eduardo Schwartz, and Avanidhar Subrahmanyam, 2007, Liquidity and the law of one price: The case of the futures-cash basis, Journal of Finance 62, 2201-2234.

Roşu, Ioanid, 2015, Fast and slow informed trading, Working paper.

Yao, Chen, and Mao Ye, 2015, Why trading speed matters: A tale of queue rationing under price controls, Working paper.

Table 1: Descriptive statistics

The table reports descriptive statistics for precipitation and for the sample of 100 ETFs. In Panel A, $PRECIP$ is the variable that captures total precipitation recorded by the weather stations along the Chicago-New York corridor. Along with precipitation statistics (in mm per a 15-minute sampling interval), we report the percent share of intervals with $PRECIP$ greater than 0.5 standard deviations ($PRECIP1$) and with $PRECIP$ greater than 1 standard deviation ($PRECIP2$). Finally, we report the length of an average period with consecutive $PRECIP1$ and $PRECIP2$ as well as the percent share of days with episodes of $PRECIP1$ or $PRECIP2$. Panel B classifies 100 sample ETFs into categories according to the underlying asset basket.

Panel A: Precipitation	
$PRECIP$, mm/interval	
mean	0.155
median	0.070
std. dev.	0.218
% intervals with $PRECIP1$	17.0
% intervals with $PRECIP2$	10.5
length $PRECIP1$, min	54.2
length $PRECIP2$, min	49.1
Panel B: ETF sample	
Equities	
US index	50
International index	22
Interest rate products	20
Metals	4
Real estate	1
Other	3

Table 2. Market activity statistics

The table contains summary statistics for the sample of 100 ETFs. Statistics are derived from the millisecond DTAQ data and aggregated into 15-minute intervals to match precipitation data. Volatility is defined as the difference between the high and low price in a 15-minute interval scaled by the average price. Trade price discovery is the percentage of efficient price variance that may be attributed to trades (Hasbrouck, 1991). The National Best Bid and Offer (NBBO) is defined as the difference between the lowest offer quote and the highest bid quote across all markets. We divide the assets into terciles by their average NBBO. Assets with the smallest (largest) NBBOs are considered the most (least) tick-constrained. Price impact is defined as twice the signed difference between the NBBO midquote 15 seconds after the trade and the midquote at the time of the trade. Effective spread is twice the signed difference between the trade price and the corresponding midquote. Realized spread is the difference between the effective spread and the corresponding price impact.

	mean	std. dev.	25%	median	75%
Panel A: Activity statistics					
# NBBO updates	5,305	8,169	608	2,470	6,953
# trades	500	1,233	39	113	432
volume, sh.	190,522	485,076	13,438	32,171	120,444
price, $	71.69	36.67	42.08	69.61	92.06
trade size, sh.	448	852	246	311	425
volatility, %	0.154	0.076	0.111	0.158	0.195
Panel B: Trading cost statistics					
NBBO, $	0.019	0.024	0.010	0.012	0.019
most constrained	0.010	0.003	0.010	0.010	0.010
least constrained	0.036	0.037	0.018	0.028	0.042
trade price disc., %	0.296	0.159	0.194	0.261	0.350
price impact, $	0.006	0.009	0.000	0.004	0.009
effective spread, $	0.019	0.032	0.010	0.011	0.018
realized spread, $	0.013	0.033	0.002	0.007	0.014

Table 3. Microwave connectivity and trading costs

The table contains coefficient estimates from the following panel regression:

$$DEPVAR_{it} = \alpha_0 + \beta_1 PRECIP_t + \beta_2 VIX_t + \varepsilon_{it},$$

where $DEPVAR_{it}$ is one of the following variables: the price impact, $PIMP$, the effective spread, ESP, or the realized spread, RSP, in asset i; $PRECIP$ is total precipitation in the Chicago-New York corridor; and VIX is the volatility index. We also use $PRECIP1$ and $PRECIP2$ to identify the most significant precipitation events. All variables are standardized (by demeaning and scaling by the standard deviation for each stock) and as such the regression models control for asset fixed effects, and the standard errors (in parentheses) are double-clustered along the asset and time dimensions. The data are from the 2011-2012 period. Panel A examines the full sample of 100 ETFs, and Panel B examines the assets for which the minimum tick size is the most (least) binding. For this test, we separate the assets into terciles by their average NBBO on the previous day. Assets with the smallest (largest) NBBOs are considered the most (least) tick-constrained. Asterisks ***, ** and * denote statistical significance at the 1%, 5% and 10% levels.

	$PIMP$			ESP			RSP		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Full sample									
$PRECIP$	-.010***			-.010***			-.005**		
	(.004)			(.003)			(.002)		
$PRECIP1$		-.035***			-.041***			-.024***	
		(.012)			(.010)			(.007)	
$PRECIP2$			-.047***			-.043***			-.021***
			(.013)			(.011)			(.008)
VIX	.035***	.035***	.035***	.057***	.058***	.057***	.036***	.036***	.036***
	(.009)	(.009)	(.009)	(.008)	(.008)	(.008)	(.006)	(.006)	(.006)

Panel B: Effects of $PRECIP2$ for assets that are the most (least) constrained by the minimum tick size (large sample)

	$PIMP$		ESP		RSP	
	most	least	most	least	most	least
$PRECIP2$	-.051***	-.039***	-.023***	-.079***	-.006	-.058***
	(.017)	(.010)	(.008)	(.020)	(.007)	(.017)

Table 4. Microwave connectivity, trading activity and volatility

The table contains coefficient estimates from the following panel regression:

$$DEPVAR_{it} = \alpha_0 + \beta_1 PRECIP_t + \beta_2 VIX_t + \varepsilon_{it},$$

where $DEPVAR_{it}$ is one of the following three variables (the number of trades, traded volume, or volatility) in asset i during a 15-minute interval t; $PRECIP$ is total precipitation in the Chicago-New York corridor; and VIX is the volatility index. We also use $PRECIP1$ and $PRECIP2$ to identify the most significant precipitation events. All variables are standardized and as such the regression models control for asset fixed effects, and the standard errors (in parentheses) are double-clustered along the asset and time dimensions. The data are from the 2011-2012 period. Panel A examines the full sample of 100 ETFs, Panel B examines the assets for which the minimum tick size is the least (most) binding. For this test, we separate the assets into terciles by their average NBBO on the previous day. Assets with the smallest (largest) NBBOs are considered the most (least) tick-constrained. Asterisks ***, ** and * denote statistical significance at the 1%, 5% and 10% levels.

	trades			volume			volatility		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Full sample									
$PRECIP$	-.010***			-.012***			-.025**		
	(.006)			(.004)			(.010)		
$PRECIP1$		-.070***			-.044***			-.103***	
		(.020)			(.013)			(.032)	
$PRECIP2$			-.072***			-.042***			-.118***
			(.023)			(.015)			(.036)
VIX	.079***	.079***	.079***	.049***	.050***	.049***	.185***	.186***	.185***
	(.015)	(.015)	(.015)	(.009)	(.009)	(.009)	(.024)	(.024)	(.024)

Panel B: Effects of $PRECIP2$ for assets that are the most (least) constrained by the minimum tick size (large sample)

	trades		volume		volatility	
	most	least	most	least	most	least
$PRECIP2$	-.111***	-.015	-.064***	-.010	-.119***	-.109***
	(.034)	(.021)	(.025)	(.013)	(.038)	(.032)

Table 5. Limit order aggressiveness: ITCH sample

The table contains coefficient estimates β_1 from the following panel regression estimated using ITCH limit order book data:

$$DEPVAR_{it} = \alpha_0 + \beta_1 PRECIP_t + \beta_2 VIX_t + \varepsilon_{it},$$

where $DEPVAR_{it}$ is one of the following two variables: (i) the number of limit orders that match the NBBO in ETF i during period t (Panel A), and (ii) the number of orders that match or improve the NBBO (Panel B). Both variables are scaled by the number of order submissions. $PRECIP$ is total precipitation in the Chicago-New York corridor; and VIX is the volatility index. We also use $PRECIP1$ and $PRECIP2$ to identify the most significant precipitation events. All variables are standardized and as such the regression models control for asset fixed effects, and the standard errors (in parentheses) are double-clustered along the asset and time dimensions. Asterisks ***, ** and * denote statistical significance at the 1%, 5% and 10% levels. Note: results reported here are derived from 2012; we are in the process of adding the results from 2011.

	Panel A: NBBO match			Panel B: NBBO match or improve		
	full sample	most constr.	least constr.	full sample	most constr.	least constr.
	(1)	(2)	(3)	(4)	(5)	(6)
$PRECIP$	.017***	.025***	.038***	.006	.008	.041***
	(.004)	(.004)	(.006)	(.004)	(.005)	(.006)
$PRECIP1$	.054***	.080***	.095***	.028***	.006	.107***
	(.007)	(.008)	(.011)	(.006)	(.010)	(.010)
$PRECIP2$	.040***	.068***	.113***	.063***	.010	.126***
	(.012)	(.013)	(.013)	(.012)	(.015)	(.013)

Table 6. Microwave connectivity and information asymmetries in futures and equities: small sample

The table contains coefficient estimates from the following panel regression:

$$PIMP_{it} = \alpha_0 + \beta_1 PRECIP_t + \beta_2 VIX_t + \varepsilon_{it},$$

where $PIMP_{it}$ is the price impact; $PRECIP$ is total precipitation in the Chicago-New York corridor; and VIX is the volatility index. We also use $PRECIP1$ and $PRECIP2$ to identify the most significant precipitation events. All variables are standardized, and as such the regression models control for asset fixed effects, and the standard errors (in parentheses) are time-clustered. The futures sample (Panel A) is from 2012 and includes four e-mini contracts: S&P 500, S&P MidCap 400, Nasdaq 100 and Financial Sector Select. The equities sample (Panel B) is from the same time period and includes the four corresponding ETFs. Asterisks ***, ** and * denote statistical significance at the 1%, 5% and 10% levels.

	Panel A: futures			Panel B: equities		
	(1)	(2)	(3)	(4)	(5)	(6)
$PRECIP$	-.004			-.037***		
	(.009)			(.006)		
$PRECIP1$		-.003			-.078***	
		(.022)			(.018)	
$PRECIP2$			-.014			-.079***
			(.031)			(.022)

Table 7. Post-democratization period

The table reports the β_1 coefficient estimates from the following panel regression:

$$DEPVAR_{it} = \alpha_0 + \beta_1 PRECIP2_t + \beta_2 VIX_t + \varepsilon_{it},$$

where $DEPVAR_{it}$ is one of the following four variables (price impacts, effective spreads, realized spreads, number of trades, traded volume, or volatility) in asset i during a 15-minute interval t; $PRECIP2$ is a dummy variable that captures periods when total precipitation in the Chicago-New York corridor exceeds one standard deviation; and VIX is the volatility index. All variables are standardized and as such the regression models control for asset fixed effects, and the standard errors (in parentheses) are double-clustered along the asset and time dimensions. The data are from the 2013-2014 period. We examine three groups of assets: (i) 100 ETFs in the full sample, and (ii/iii) the ETF terciles for which the tick size is the most (least) binding. Asterisks ***, ** and * denote statistical significance at the 1%, 5% and 10% levels.

	$PIMP$	ESP	RSP	trades	volume	volatility
	(1)	(2)	(3)	(4)	(5)	(6)
full sample	.007	.001	-.003	.027	.007	.016
	(.013)	(.012)	(.009)	(.018)	(.012)	(.033)
most constr.	-.016	.003	.010	.024	.009	.001
	(.015)	(.008)	(.007)	(.023)	(.018)	(.031)
least constr.	.014	.002	-.006	.031	.005	.023
	(.012)	(.021)	(.016)	(.021)	(.009)	(.032)

Table 8. Democratization: Event study

The event window spans the months of September 2012 to April 2013. In this window, the months of September, October and November capture the period prior to the democratization, and the months of February, March and April capture the post-democratization period. We report the coefficient estimates β_1 from the following panel regression:

$$DEPVAR_{it} = \alpha_0 + \gamma_t + \beta_1 POST_t + \beta_2 VIX_t + \varepsilon_{it},$$

where $DEPVAR_{it}$ is one of the following variables (price impacts, effective spreads, realized spreads, number of trades, traded volume, volatility or stock price) in asset i on day t; γ_t denotes a time trend, $POST$ is a dummy variable that equals to one in February-April 2013; and VIX is the volatility index. All variables are standardized and as such the regression models control for asset fixed effects, and the standard errors are double-clustered along the asset and time dimensions. We examine three groups of assets: (i) 100 ETFs in the full sample, and (ii/iii) the ETF terciles for which the tick size is the most (least) binding. Asterisks ***, ** and * denote statistical significance at the 1%, 5% and 10% levels.

	$PIMP$	ESP	RSP	trades	volume	volatility	price
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
full sample	-.484***	-.711***	-.546***	-.071	.045	-.836***	.319***
	(.122)	(.127)	(.125)	(.129)	(.125)	(.244)	(.083)
most constr.	-.590***	-.454**	-.190	-.475***	-.174	-1.09***	.496***
	(.185)	(.179)	(.184)	(.171)	(.177)	(.254)	(.135)
least constr.	-.448***	-.965***	-.905***	.084	.095	-.542**	.284**
	(.100)	(.181)	(.178)	(.120)	(.123)	(.222)	(.129)

Table 9. Robustness: alternative sampling and regression setup

The table reports the β_1 coefficient estimates from the following panel regression:

$$DEPVAR_{it} = \alpha_0 + \beta_1 PRECIPx_t + \beta_2 VIX_t + \varepsilon_{it},$$

where $DEPVAR_{it}$ is one of the following variables (price impacts, effective spreads, realized spreads, number of trades, traded volume, or volatility) in asset i during a 15-minute interval t; $PRECIP2$ is a dummy variable that captures periods when total precipitation in the Chicago-New York corridor exceeds one standard deviation (Panel A); and VIX is the volatility index. We examine several specifications of the model. The *mood control* specification restricts precipitation episodes to those occurring in Ohio when precipitation is near-zero in the eastern and western parts of the Chicago-New York corridor. The *expanded area* specification uses additional weather stations, forming a wider area around the corridor. The *placebo area* specification uses data from the weather stations located in Colorado, Utah and Wyoming, away from the corridor. The *afternoon only* specification uses data between noon and the market close. The *intraday FE* specification adds intraday fixed effects. Finally, Panel B replaces total precipitation across all stations with the average precipitation per station $MPRECIP$, and its two variations, $MPRECIP1$ and $MPRECIP2$, which are dummies that capture episodes when the average precipitation is more than 0.5 and 1 standard deviation removed from the mean. All variables are standardized and as such the regression models control for asset fixed effects, and the standard errors (in parentheses) are double-clustered along the asset and time dimensions. The data are from the 2011-2012 period, and we examine the full sample of 100 ETFs. Asterisks ***, ** and * denote statistical significance at the 1%, 5% and 10% levels.

	$PIMP$	ESP	RSP	trades	volume	volatility
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: $PRECIPx = PRECIP2$						
mood control	-.060***	-.061***	-.026***	-.094***	-.053***	-.166***
	(.013)	(.012)	(.009)	(.024)	(.016)	(.035)
expanded area	-.034***	-.040***	-.020**	-.055**	-.032**	-.087**
	(.013)	(.012)	(.008)	(.023)	(.015)	(.039)
placebo area	.006	-.012	-.001	.015	.003	-.036
	(.016)	(.025)	(.019)	(.024)	(.016)	(.038)
afternoon only	-.061***	-.063***	-.028***	-.080***	-.048***	-.147***
	(.015)	(.014)	(.010)	(.026)	(.018)	(.040)
intraday FE	-.054***	-.060***	-.028***	-.067***	-.043***	-.141***
	(.012)	(.012)	(.008)	(.021)	(.014)	(.035)
Panel B: $PRECIPx \in \{MPRECIP, MPRECIP1, MPRECIP2\}$						
$MPRECIP$	-.007*	-.006*	-.003	-.014**	-.009**	-.013
	(.004)	(.004)	(.002)	(.006)	(.004)	(.011)
$MPRECIP1$	-.024**	-.027**	-.013*	-.051***	-.030**	-.057*
	(.012)	(.011)	(.007)	(.020)	(.013)	(.034)
$MPRECIP2$	-.043***	-.039***	-.014*	-.067***	-.037***	-.096***
	(.012)	(.011)	(.008)	(.021)	(.014)	(.035)

Table 10. Robustness: alternative variables of interest

The table reports the β_1 coefficient estimates from the following panel regression:

$$DEPVAR_{it} = \alpha_0 + \beta_1 PRECIP2_t + \beta_2 VIX_t + \varepsilon_{it},$$

where $DEPVAR_{it}$ is one of the following three variables: price impacts ($PRIMP$), effective spreads (ESP) and realized spreads (RSP). Each variable is computed as equally-weighted ($EW\_$) or volume-weighted scaled by the corresponding midquote ($VWP\_$); $PRECIP2$ is a dummy variable that captures periods when total precipitation in the Chicago-New York corridor exceeds one standard deviation; and VIX is the volatility index. All variables are standardized and as such the regression models control for asset fixed effects, and the standard errors (in parentheses) are double-clustered along the asset and time dimensions. The data are from the 2011-2012 period. We examine three groups of assets: (i) 100 ETFs in the full sample, and (ii/iii) the ETF terciles for which the tick size is the most (least) binding. Asterisks ***, ** and * denote statistical significance at the 1%, 5% and 10% levels.

	EW_PIMP	VWP_PIMP	EW_ESP	VWP_ESP	EW_RSP	VWP_RSP
	(1)	(2)	(3)	(4)	(5)	(6)
full sample	-.064***	-.059***	-.089***	-.067***	-.008	-.032***
	(.018)	(.014)	(.021)	(.015)	(.012)	(.010)
most constr.	-.079***	-.069***	-.086***	-.035***	.030*	-.005
	(.026)	(.020)	(.020)	(.010)	(.018)	(.008)
least constr.	-.046***	-.047***	-.105***	-.109***	-.067***	-.077***
	(.014)	(.011)	(.030)	(.027)	(.026)	(.022)

Table 11. Quoted spread and inside depth

The table reports the β_1 coefficient estimates from the following panel regression:

$$DEPVAR_{it} = \alpha_0 + \beta_1 PRECIP2_t + \beta_2 VIX_t + \varepsilon_{it},$$

where $DEPVAR_{it}$ is one of the following four variables (NBBO spread or NBBO inside depth) in asset i during a 15-minute interval t; $PRECIP2$ is a dummy variable that captures periods when total precipitation in the Chicago-New York corridor exceeds one standard deviation; and VIX is the volatility index. All variables are standardized and as such the regression models control for asset fixed effects, and the standard errors (in parentheses) are double-clustered along the asset and time dimensions. The data are from the 2011-2012 period. We examine three groups of assets: (i) 100 ETFs in the full sample, and (ii/iii) the ETF terciles for which the tick size is the most (least) binding. Asterisks ***, ** and * denote statistical significance at the 1%, 5% and 10% levels.

	QSP	$DEPTH$
	(1)	(2)
full sample	-.065***	.014
	(.020)	(.029)
most constr.	-.026**	.118***
	(.012)	(.045)
least constr.	-.105***	-.013
	(.034)	(.034)

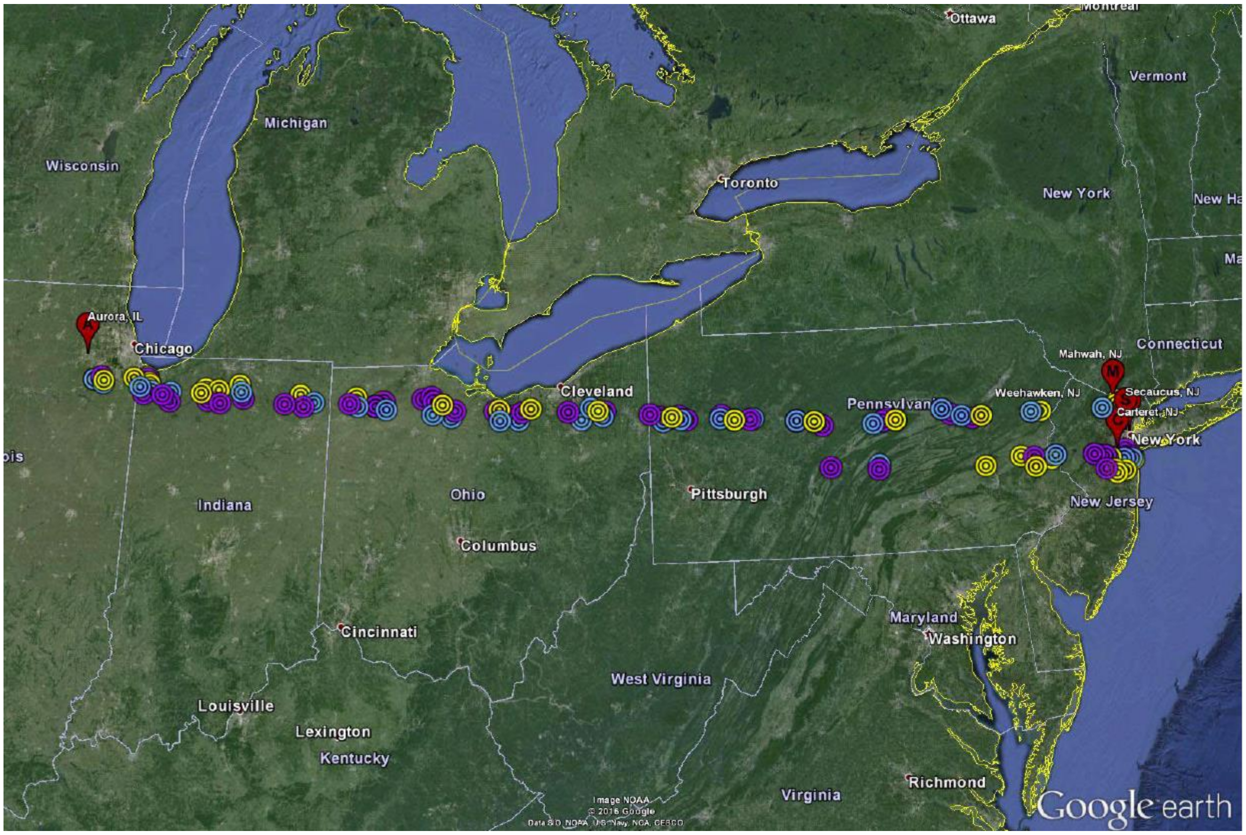

Figure 1. Microwave network paths

The figure maps tower locations of three microwave networks (blue, yellow and purple icons) obtained from the Federal Communications Commission. There are more than three microwave networks between Chicago and New York during our sample period; however, we plot only three to avoid clutter. The remaining networks follow very similar paths. The red markers indicate locations of the CME's data center in Aurora, IL (marker A); the NYSE data center in Mahwah, NJ (marker M); Nasdaq data center in Carteret, NJ (marker C); BATS data center in Weehawken, NJ (marker W); and Direct Edge data center in Secaucus, NJ (marker S).

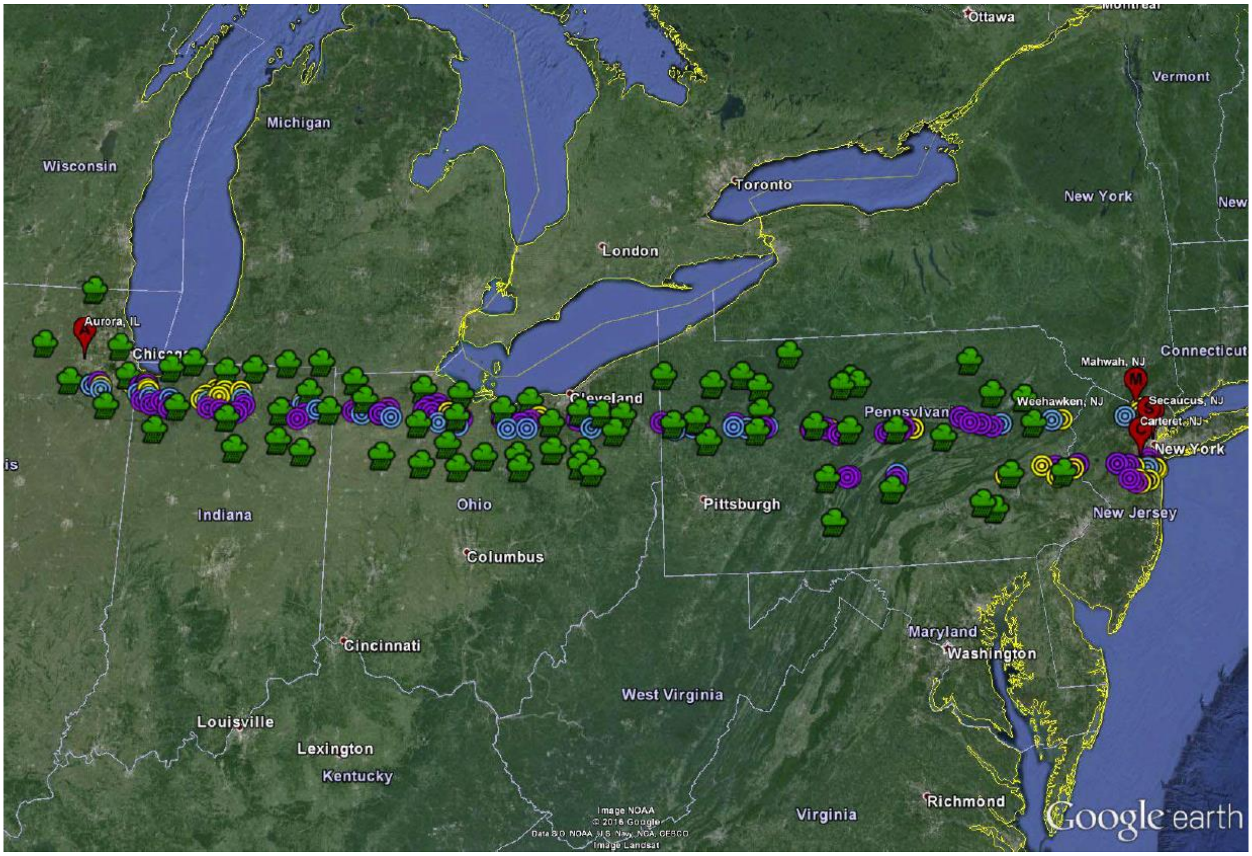

Figure 2. Locations of microwave networks and weather stations

The figure maps the weather stations (green icons) located near the microwave network paths. Station data are obtained from the National Oceanic and Atmospheric Administration. The red markers indicate locations of the CME's data center in Aurora, IL (marker A); the NYSE data center in Mahwah, NJ (marker M); Nasdaq data center in Carteret, NJ (marker C); BATS data center in Weehawken, NJ (marker W); and Direct Edge data center in Secaucus, NJ (marker S).

Figure 3. A typical weather front

As a weather front moves over the microwave paths, it disrupts data transmission forcing trading firms to fall back on the fiber-optic cable.

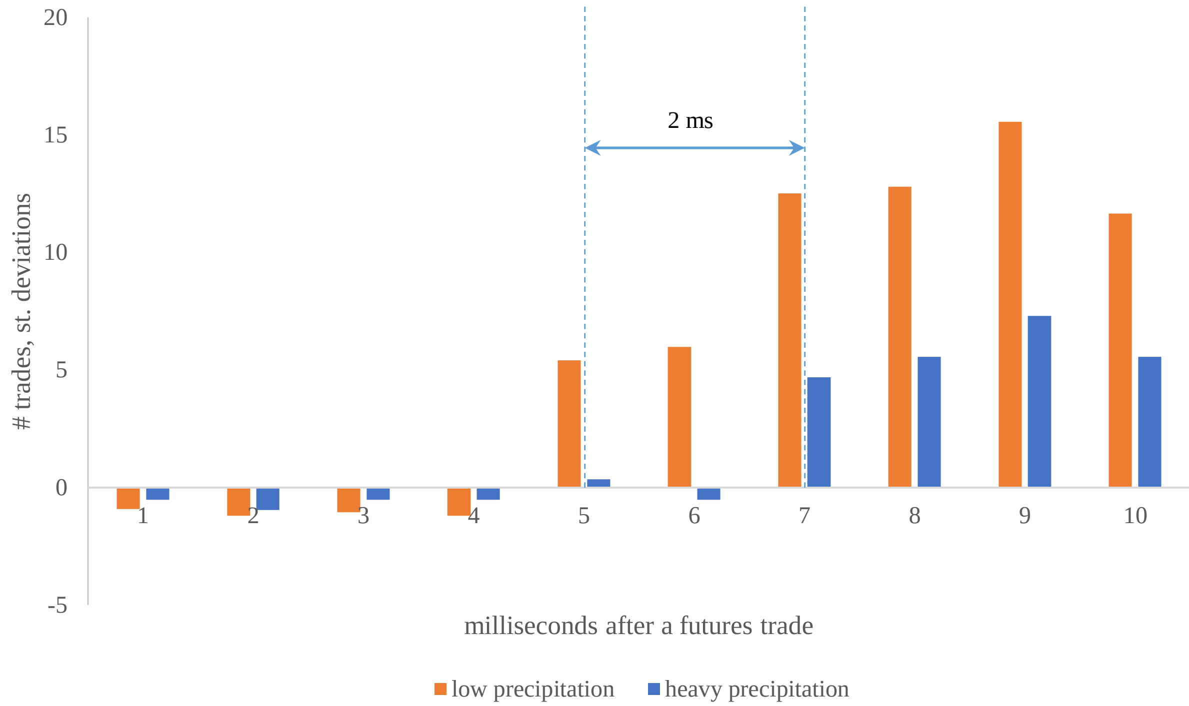

Figure 4. Equity trades after a futures trade during low and heavy precipitation episodes

The figure reports a timeline of equity trades that follow a futures trade. Orange bars represent periods of zero or very low precipitation ($PRECIP2$<0), and blue bars represent periods of heavy precipitation ($PRECIP2$>1) when the microwave networks are disrupted. The number of trades is standardized on an asset by asset basis to allow for cross-sectional comparability. We focus on the standalone futures trades (t = 0), those not preceded by another futures or equity trade in the previous 100 milliseconds. Note: light covers the distance from Chicago to New York in 4 milliseconds (ms), the microwave signal covers this distance in about 4.5 ms, and the same signal takes 6.5 ms to cover the distance through fiber. During our sample period, the CME clock lags DTAQ by about one millisecond, and we adjust for this lag. The 2-ms advantage of the fastest traders is evident even without the adjustment.

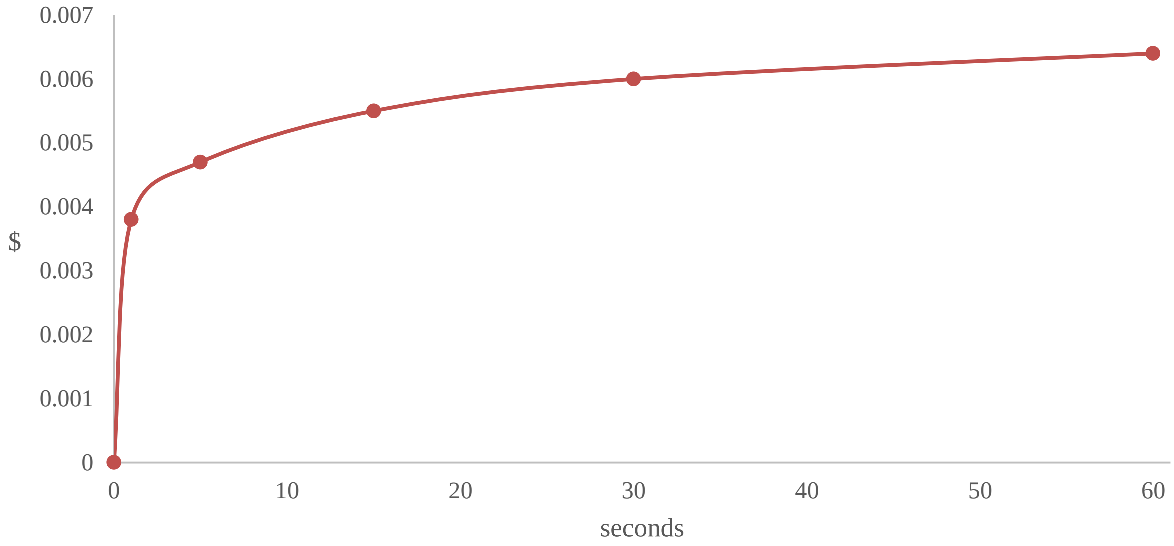

Figure 5. Price impacts

The figure reports price impacts computed as the signed difference between a midquote at a certain time after the trade and the midquote at the time of the trade: $PRIMP_t = 2q_t\left(mid_{t+\gamma} - mid_t\right)$, where q_t is the trade direction indicator, mid_t is the midquote computed as $(NBBO\ Ask_t + NBBO\ Bid_t)/2$, and γ indicates the time elapsed since the trade, with $\gamma \in \{1s, 5s, 15s, 30s, 60s\}$.

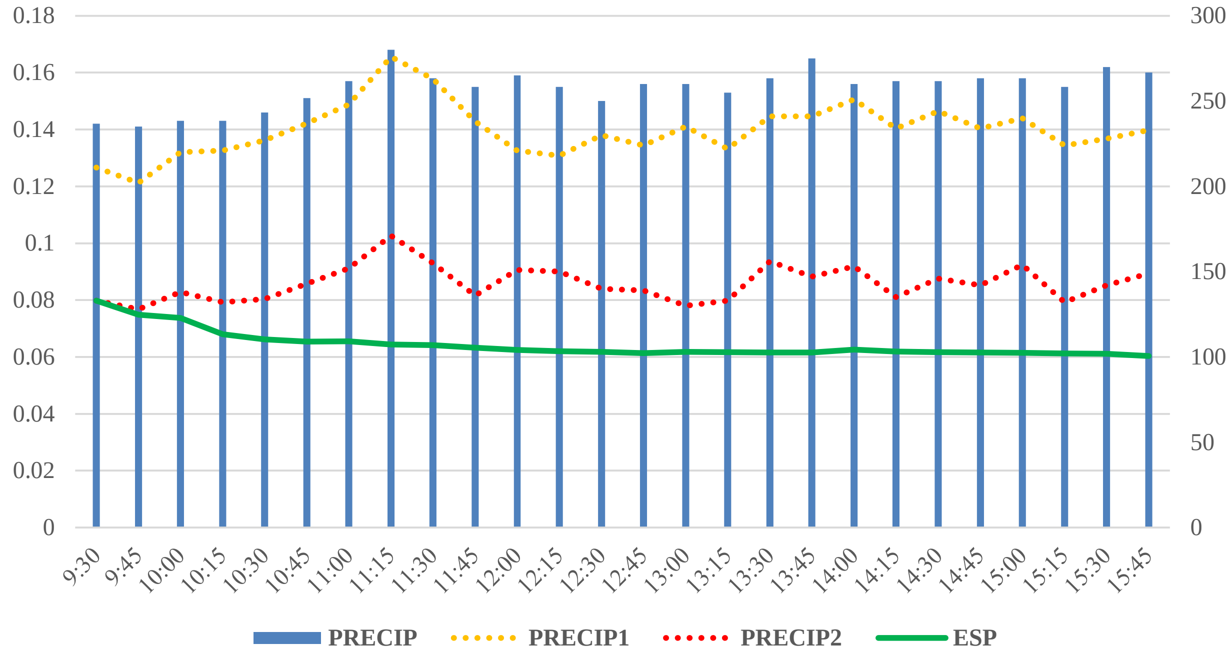

Figure 6. Intraday patterns

The figure reports intraday patterns for $PRECIP$ (in mm average per intraday period, left axis), $PRECIP1$ and $PRECIP2$ (both in number of occasions per intraday period, right axis), and ESP (scaled by 10000 for display purposes, right axis).

Appendix to "Every cloud has a silver lining: Fast trading, microwave connectivity and trading costs" by A. Shkilko and K. Sokolov

A1. Price discovery via trades and quotes

To examine price discovery via trades and quotes, we use the methodology described in Hasbrouck (1991 a,b) to decompose the efficient price variance into the trade-related and trade-unrelated components. We begin with an assumption that the observed midquotes p_t follow a random walk with two components:

$$p_t = m_t + s_t,$$

where m_t is the efficient price (the expectation of price conditioned on all available information at time t), and s_t is a deviation of the price from the efficient price. We then estimate the VAR with ten lags as follows:

$$r_t = a_1 r_{t-1} + a_2 r_{t-2} + \cdots + b_0 q_t + b_1 q_{t-1} + b_2 q_{t-2} + \cdots + v_{r,t}$$

$$q_t = c_1 r_{t-1} + c_2 r_{t-2} + \cdots + d_1 q_{t-1} + d_2 q_{t-2} + \cdots + v_{q,t},$$

where r_t is the difference in log-midquotes, and q_t is a vector of three trade-related variables, including a trade direction indicator, signed volume and signed square root of volume. The VAR is then converted into the VMA model:

$$r_t = a_1^* v_{r,t-1} + a_2^* v_{r,t-2} + \cdots + b_0^* v_{q,t} + b_1^* v_{q,t-1} + b_2^* v_{q,t-2} + \cdots$$

$$q_t = c_1^* v_{r,t-1} + c_2^* v_{r,t-2} + \cdots + d_1^* v_{q,t-1} + d_2^* v_{q,t-2} + \cdots,$$

and the total variance of the random walk component is given by:

$$\sigma_w^2 = (1 + \textstyle\sum_{i=1}^{\infty} a_i^*)^2 \sigma_r^2 + (\sum_{i=0}^{\infty} b_i^*)\Omega(\sum_{i=0}^{\infty} b_i^{*\prime}),$$

where the first term corresponds to the trade-unrelated component of the efficient price innovation, and the second term corresponds to the trade-related component of this innovation. The model is estimated in event time, with ts indexing every new midquote.

A2. Information share estimation

To compute information shares using the methodology in Hasbrouck (1995), we first estimate the following vector error correction model (VECM) for each futures-ETF pair:

$$\Delta p_{f,t} = \alpha_1 \Delta p_{f,t-1} + \cdots + \alpha_k \Delta p_{f,t-k} + \beta_0 \Delta p_{e,t-1} + \cdots + \beta_k \Delta p_{e,t-k} + g_1 (p_{f,t-1} - p_{e,t-1} - \mu) + u_{f,t}$$
$$\Delta p_{e,t} = \gamma_1 \Delta p_{f,t-1} + \cdots + \gamma_k \Delta p_{f,t-k} + \delta_0 \Delta p_{e,t-1} + \cdots + \delta_k \Delta p_{e,t-k} + g_2 (p_{f,t-1} - p_{e,t-1} - \mu) + u_{e,t},$$

where $\Delta p_{f,t}$ ($\Delta p_{e,t}$) is the difference between the current and lagged prices of the futures (ETF), and μ is the mean difference between the price of the futures and the ETF.

In the second step, we obtain the VMA representation of the above model:

$$\Delta p_{f,t} = a_0 u_{f,t} + \cdots + a_k u_{f,t-k} + b_0 u_{e,t} + \cdots + b_k u_{e,t-k}$$
$$\Delta p_{e,t} = c_0 u_{f,t} + \cdots + c_k u_{f,t-k} + d_0 u_{e,t} + \cdots + d_k u_{e,t-k}$$

and add the coefficients $A = \sum_0^k a_i$ and $B = \sum_0^k b_i$. Next we obtain the covariance matrix of the residuals:

$$\Omega = \begin{bmatrix} \sigma_f^2 & \sigma_{ef} \\ \sigma_{ef} & \sigma_e^2 \end{bmatrix},$$

and finally, the information share (IS) of the futures market is calculated as:

$$IS_f = \frac{A^2 \sigma_f^2}{\sigma_W^2}, \text{ where } \sigma_W^2 = \begin{bmatrix} A \\ B \end{bmatrix}' \begin{bmatrix} \sigma_f^2 & \sigma_{ef} \\ \sigma_{ef} & \sigma_e^2 \end{bmatrix} \begin{bmatrix} A \\ B \end{bmatrix}.$$

Since some price innovations happen in both markets within the same millisecond, $\sigma_{ef} \neq 0$. To address this, we follow Hasbrouck (1995) and orthogonalize Ω. Orthogonalization maximises (minimises) the variance of the futures market and gives the upper (lower) bound of the true variance.