The Rogue Algorithm and its Discontents: Evidence from a Major Trading Glitch

Mehmet Sağlam
Lindner College of Business
University of Cincinnati
email: mehmet.saglam@uc.edu

Preliminary Draft: June 1, 2015

Abstract

I examine the impact of an exogenous trading glitch in a high-frequency market-making firm on institutional trading costs. At the first look, the trading glitch does not appear to affect institutional investors as it leads to dramatic increases in volume without any information content but controlling for various stock- and order-level characteristics, I find that executing a large order on a glitch-affected stock incurs substantially higher costs on the event day. Moreover, the cost increase is persistent up to one week roughly with the same additional cost magnifying the total economic costs. These findings can be interpreted as negative externalities of algorithmic trading which has important policy implications.

Keywords: Algorithmic Trading, Institutional Trading Costs, Rogue Algorithm, Negative Externality

JEL Classification: G10.

1. Introduction

Institutional investors such as mutual funds, pension funds, and hedge funds are the largest holder of publicly traded stocks in the United States and their trading accounts for the majority of daily trading volume (Boehmer and Kelley, 2009). There is ample evidence that their trading facilitates price discovery process by identifying under- and over-priced securities (Diether et al., 2009). Due to their large-size trades, the magnitude of their trading costs heavily affect the resulting investment performance (Perold, 1988). In order to control these execution costs, they usually work with sophisticated execution desks in large brokerage firms. Given this significance, institutional trading costs have been a major area of research in the overall study of market liquidity.

The recent fundamental transformation in securities trading industry through algorithmic trading has been heavily debated through the perspective of institutional investors. A particular subset of algorithmic trading, high frequency trading, has been the main driver of the discussions, partly due to its continuous coverage in the mainstream media. Jones (2013) defines this category of algorithmic trading as "the use of computer algorithm to make decisions about order submissions and cancellations," and provides three main categories for the broad HFT strategies: market-making, relative-value trading and directional news trading.

Although the link between HFT and institutional trading cost has not been heavily studied, there are numerous empirical studies documenting that HFT activity is positively correlated with improvement in market quality measures [1]. These findings are easily justified from the economics of liquidity provision. Given the increased competition in liquidity suppliers, HFT stands out as a more efficient competitor with better inventory management which consequently leads to smaller spreads and price inefficiency [2]. On the other hand, Tong (2013) finds that increase in HFT activity is actually positively associated with institutional trading costs and underscores that the cost increase is more pronounced when the high-frequency traders follow a directional trading strategy. This additional cost could be, for example, due to the ability of the high-frequency trader to anticipate the execution of a large order from an institutional investor and trade in front of him.

In this paper, I study the impact of a major rogue algorithm originating from a large high-

[1] See e.g., Hendershott et al. (2011) and Brogaard et al. (2014).
[2] See e.g., Aït-Sahalia and Saglam (2013) for a theoretical analysis.

frequency market-making firm on institutional trading costs. This exogenous shock occurred on August 1, 2012 during the first 30 minutes of trading and caused numerous erroneous trades on a set of NYSE-listed stocks. Using public trade data, I identify the set of stocks affected by the trading glitch and test whether it is costlier to trade the affected stocks on the same day using a proprietary execution data from a large investment bank.

The trading glitch led to dramatic increases in number of trades and since they were erroneous, they cannot contain much information about the fundamental value of the stock. If these erroneous trades did not have any adversary effect on the remaining market participants, one would expect that the institutional trading costs should remain unaffected after the end of the glitch[3]. Surprisingly, I find that executions on affected stocks had higher trading costs on August 1, 2012 as measured by two popular metrics for institutional trading costs, implementation shortfall and VWAP slippage. Furthermore, the cost increase for affected executions is economically substantial. While the median execution has an implementation shortfall (resp. VWAP slippage) of 3 bps (resp. 1 bp), affected executions suffer a cost increase of 18 bps (resp. 7 bps)[4]. My findings are robust to different identification strategies of the affected group or alternative explanations.

More importantly, I find that the impact of the trading glitch is not short-lived but actually persists for another week with roughly the same additional cost. Hence, the total cost of the trading glitch is substantially higher than expected due to this persistence. A back-of-the-envelope calculation shows that the total cost to institutional traders is on the order of $100 million. These results suggest that the rogue algorithm might have exerted a negative externality on other market participants to process the flow of erroneous trades to determine an efficient price.

The rest of the paper is organized as follows: In Section 2, I provide a brief background information on the trading glitch. Section 3 introduces the metrics for institutional trading costs, while Section 4 provides a detailed information about the data sources. Section 5 introduces the main methodology of the analysis and discusses main findings. Section 6 examine the robustness of our results with respect to different identification strategies and Section 7 investigates whether the cost increase is persistent. We conclude in Section 8.

[3]During the glitch period, one would even expect to see decrease in trading costs due to the higher volume.

[4]The most conservative estimate for the cost increase is 11 bps (resp. 3 bps) which is still substantial economically.

2. Background Information on the Trading Glitch

At 10:38 a.m. on August 1, 2012, Wall Street Journal published one of the first news piece regarding the trading glitch citing the abnormal volume increase in several stocks due to Knight Capital and quoting a note from NYSE Euronext about an initial review of 148 symbols for trading irregularities between 09:30 a.m. and 10:15 a.m.[5]. According to the company statement on August 2, 2012, Knight Capital announced that the company lost $440 million due to a technology issue at the open of the trading at the NYSE on August 1, 2012 related to an installation of a trading software that resulted in Knight sending numerous erroneous orders in NYSE-listed stocks.

NYSE cancelled only very small portion of the erroneous trades due to the regulations. In six stocks where there were price swings of 30% or more, the trades were cancelled. Trading was also halted in several stocks due to tripping the circuit breakers. It is reported that the size of the position built during the trading glitch was on the order of billions [6]. During this two-day period, share price of Knight Capital plummeted from $10.33 to $2.58 and forced the company to search for a buyer for the following months.

Knight Capital was the designated market-maker for about 400 NYSE-listed stocks. The software update that Knight implemented was due to the launch of a new NYSE proposal to attract more retail investors as the market volume has been consistently low. The trading irregularities were first noticed by large price movements with unexpected high volume in NYSE-listed stocks. For example, Table 1 reports 8 NYSE-listed securities having larger share volume than SPY, the most liquid ETF for S&P 500, during the initial 30 minutes of the trading. Specifically, without any link to a corporate news, Juniper Networks (NYSE:JNPR) realized roughly 20 million share volume in this 30-minute period whereas its daily 20-day average share volume is only 8 million. Figure 1 and 2 illustrate the number of trades and the corresponding size of the trades occurred in each minute of the trading days on July 31, 2012 and August 1, 2012. It is very apparent from the figures that there is a large volume increase between 9:32 a.m. and 10:00am after which both number of trades and corresponding share volume drop drastically. This volume pattern can easily determine any affected stock from the trading glitch. However, we also note that compared to the

[5]"Fat Finger Trade? Knight Looking Into Stock-Trading Irregularities," *Wall Street Journal*, August 1, 2012.
[6]"SEC Nixed Knight's Plea for a Do-Over," *Wall Street Journal*, August 6, 2012.

previous day, there is also episodes of high volume in the afternoon session suggesting that Knight Capital might have traded later in the day to offload any inventory built during the trading glitch. For this reason, total daily volume on August 1, 2012, compared to historical averages may be a good proxy to identify the affected stocks as well.

3. Two Cost Metrics for Institutional Trading

In this section, I introduce two most frequently used cost measures for institutional trading: implementation shortfall and volume-weighted average price (VWAP) slippage.

3.1. Implementation Shortfall

As a first measure, I use implementation shortfall (IS) as introduced by Perold (1988). This approach is based on comparing the weighted average price of the actual trades to a benchmark price observed prior to the execution period. Formally, implementation shortfall of the kth execution in our data is given by

$$\text{Implementation Shortfall}_k = \text{sgn}\left(Q_k\right) \frac{P_k^{\text{avg}} - P_{k,0}}{P_{k,0}}, \tag{1}$$

where P_k^{avg} is the value-weighted execution price of the parent order and $P_{k,0}$ is the mid-quote price of the security (arrival price) when the parent order starts being executed.

3.2. VWAP Slippage

Implementation shortfall considers an ex-ante static benchmark and does not consider any permanent price changes during the execution period. For this reason, one popular ex-post benchmark for average execution price is the volume-weighted average price during the trading interval. Formally, VWAP in a particular trading interval period is defined as follows. Suppose that there are N trades during this period. Given the sequence of trades at prices $P_1, \ldots, P_N$ with corresponding quantities $V_1, \ldots, V_N$, VWAP is given by

$$\bar{P} = \frac{\sum_{i=1}^{N} P_i V_i}{\sum_{i=1}^{N} V_i}. \tag{2}$$

Using this definition, VWAP slippage for the kth execution in our data equals

$$\text{VWAP Slippage}_k = \operatorname{sgn}(Q_k) \frac{P_k^{\text{avg}} - \bar{P}_k}{\bar{P}_k}, \tag{3}$$

where $\bar{P}_k$ is the realized VWAP over the kth execution period.

4. Data Sources

I compile the data from several sources. Stock returns, volume, outstanding shares and prices come from the Center for Research in Security Prices (CRSP). Earnings announcement dates come from I/B/E/S. Intraday transactions data (trades and quotes) come from the Trade and Quote (TAQ) database. The proprietary data on institutional large trades is obtained from an algorithmic trading desk of a large investment bank. In the next section, I describe this dataset.

4.1. Execution Data

I use a novel proprietary execution data from the historical order databases of a large investment bank providing algorithmic trading services ("The Bank"). The orders originate from a diverse pool of investors, such as institutional portfolio managers, quantitative investment funds, internal trading desks and retail customers. The Bank offers a large selection of algorithms to match the investors' trading styles and expectations. My dataset consists of two frequently used algorithms, the volume weighted average price (VWAP) and the percentage of volume (POV). The VWAP algorithm is designed to achieve an average execution price that is as close as possible to the volume weighted average price over the trading interval. Similarly, the main objective of the POV algorithm is to have constant participation rate in the market within the trading interval.

This proprietary dataset provides a rich set of attributes. For each order, this data contains trade- and stock-level statistics. Trade-level statistics include order size, direction of the order (buy or sell), order start and end times, participation rate (the ratio of order size to the total volume during the trading interval), average execution price, proportional bid-ask spread and mid-quote volatility based on the duration of the execution. Stock-level information includes average daily volume, proportional bid-offer spread and mid-volatility of the stock on the trading day along with their rolling averages over the last 20 trading days prior to the execution.

The traded asset universe includes all S&P 500 stocks with an execution duration greater than 5 minutes but no longer than 6.5 hours, the duration of a regular trading day. All executions occur between January 2012 and December 2012, inclusive. All orders have been fully filled without intermediate replacements or cancellations.

Our sample consists of 39,570 executions coming from 18,357 buy and 21,213 sell orders. The trading algorithms used are 29,027 VWAP and 10,543 POV. There are 158.3 orders per trading day on average. The highest number of executions on a single stock is 330 which corresponds to 0.83% of all executions. Table 5 provides additional summary statistics for our complete execution data.

[Insert Table 5 here]

There is a wide range of participation rates across executions with an average (median) of 5.24% (0.89%). As a fraction of daily volume, most of the orders are less than 1%. This is expected as my dataset contains the most liquid names from S&P 500 Index. Average and median implementation shortfall is 3 bps. The mean duration of the executions is a little above than 2.5 hours. Finally, the average and median percentage return realized during an execution is roughly zero.

[Insert Table 6 here]

Table 6 provides the summary statistics for executions occurred on August 1, 2012. There are 662 parent orders executed by The Bank on that day which is substantially higher than the daily average of 158.3. I observe that when compared with the full sample the order sizes are smaller relative to both interval and daily volume. Consequently, the average values for implementation shortfall and execution duration are also smaller.

5. Did Erroneous Trades Increase Institutional Trading Cost?

5.1. Identification of Affected Stocks

From available public data, it is not possible to identify all the affected stocks from the trading glitch with certainty. However, according to a company statement on August 2, 2012, the erroneous orders were limited to NYSE stocks. Beside this broad restriction, numerous stocks had a tremendous increase in daily trading volume without any link to a specific corporate news as emphasized in Section 2. This abnormal change in daily volume can help us identify the affected stocks.

Given that the abnormal trades have dropped significantly after 10:00 a.m., I focus on the number of distinct trades occurring in the first 30 minutes of the trading day. For all the available S&P 500 stocks available in my dataset, I compute the number of trades in this interval. As a benchmark security, I also compute the number of trades for the most active S&P 500 ETF, SPY. Table 2 shows the list of top 50 stocks which has the highest number of trades and compares it with number of trades observed in SPY and total number of trades observed on the previous trading day, July 31, 2012, for the same stock. I find that 35 securities have actually more trades than SPY has. Not surprisingly, all of these stocks are NYSE-listed. I also verify that all of these stocks have the similar abnormal volume profile as depicted with JNPR in Figures 1 and 2. Furthermore, all 35 stocks also appear in the initial list of affected stocks that Wall Street Journal reported on Auugust 1, 2012.

Given this evidence, I consider this list of stocks as being affected by the trading glitch. In my execution dataset, 30 out of 35 securities are executed on the day of the trading glitch and I have 49 different parent executions implemented on the affected group. In Section 6, I also discuss different identification strategies to determine the group of affected stocks.

5.2. Determinants of Trading Costs

Broad literature analyzing the variation in trading costs find that relative order size, market capitalization, asset volatility, share turnover, and bid-offer spread are the main determinants (see e.g., Domowitz et al. (2001)). Based on this primary list, I use the following order- and stock-level characteristics as control variables in the main analysis:

- **Order Size.** In order to capture variations related to the size of the order, I use participation rate (the ratio of order size to interval volume); fraction of daily volume (the ratio of order size to daily volume) and relative order size (the ratio of order size to average daily volume over the past month).

- **Firm Size.** Since the size of the executed stock may affect the available liquidity, I use the logarithm of market capitalization to control for the variation in firm size.

- **Share Volume and Turnover.** In order to control for variations in volume and turnover, I use relative daily volume (the ratio of daily volume to average daily volume over the past month), interval turnover (the ratio of interval volume to number of shares outstanding), daily turnover (the ratio of daily volume to number of shares outstanding), and average turnover (the ratio of average daily volume over the past month to number of shares outstanding).

- **Volatility.** In order to control for differences in volatility, I use interval volatility (mid-quote volatility during the interval expressed in annualized percentage), daily volatility (mid-quote volatility during the trading day expressed in annualized percentage), and average volatility (average of the daily volatilities over the past month).

- **Bid-offer spread.** In order to control for differences in spread values, I use interval spread (average bid-offer spread during the interval expressed in basis points), daily spread (average bid-offer spread during the trading day expressed in basis points), and average spread (average of the daily spread values over the past month).

- **Execution Duration.** I use execution duration expressed as a fraction of total trading hours to control for differences in urgency.

- **Market and Stock Returns.** In order to control for extreme past and contemporaneous price movements, I use absolute value of the interval, daily and prior-day returns for the stock and the S&P 500 index.

- **Earnings Announcement Days.** On earnings announcement days, firms may also have increase in trading volume due to the mere desire of opening or closing position in anticipation of good or bad quarterly results. Therefore, I also include a control variable, IsEarningsDay, which takes a value of 1 if the executed stock has an earnings announcement on the execution day.

These various stock- and order-level characteristics serve as the main control variables for all models explaining institutional trading costs.

5.3. Methodology

In order to formally test the impact of erroneous trades on execution costs, I run an OLS regression model with a treatment dummy for executions on the stocks hit by the erroneous trades while controlling for execution-level characteristics, and using calendar day dummies. Let $t(k)$ map the kth execution to the jth trading day of 2012, with $j = 1 \dots 250$. Formally, I run the following OLS regression model with

$$C_k = aD_k + bX_k + \sum_{j=1}^{250} d_j \mathbb{I}_{\{t(k)=j\}} + \epsilon_k \tag{4}$$

where C_k is either implementation shortfall or VWAP slippage, D_k is a dummy variable if the executed stock is in the affected group as described in Section 5.1 and X_k are execution-level control variables specified in 5.2.

We are concerned with heteroskedasticity, contemporaneous correlation across stocks, and autocorrelation within each stock and adjust our standard errors by clustering on calendar day and stock throughout the analysis as suggested by Petersen (2009).

5.4. Results

Using the complete dataset, we run the model specified in Equation 4 using implementation shortfall and VWAP slippage as our cost measures.

[Insert Table 7 here]

In the first column, Table 7 illustrates whether the affected executions due to the trading glitch have higher implementation shortfall after controlling for execution characteristics and calendar day dummies. The main variable of interest is of course the dummy variable for the group of affected executions and I find that it is highly statistically significant. The estimated cost increase due to the trading glitch is approximately 18 bps. This increase is economically substantial as the median (75% percentile) implementation shortfall in our dataset is only 3 bps (24 bps) as depicted in Table 5. We can also interpret the estimated coefficient by comparing it with the most important execution-level parameter, participation rate. 18 bps of cost increase roughly corresponds to an additional 30% increase in participation rate. For example, all else equal a glitch-affected execution with a participation rate of 10% is expected to have a similar implementation shortfall compared to an unaffected execution with a participation rate of 40%.

In the second column, Table 7 also illustrates the estimation results using our second cost measure with VWAP slippage. Here I regress VWAP slippage on the affected dummy, control variables, and calendar day dummies using the complete dataset. I obtain qualitatively very similar results for this cost measure as well. First, the dummy variable is highly statistically significant. The estimated cost increase due to the trading glitch in this case is approximately 6.7 bps. Since the median (75% percentile) VWAP slippage in my dataset is 1.2 bps (3.1 bps), the estimated cost increase is indeed economically substantial. I observe that in all four of the considered cutoff values for abnormal volume, the dummy for our affected group is highly statistically significant.

Apart from the affected group dummy, participation rate, fraction of daily volume, daily bid-offer spread, and the market volatility proxy (measured by the absolute value of the market return on the trading day) are the statistically significant variables for both cost metrics. I find that execution costs are positively correlated with participation rate, fraction of daily volume, daily bid-offer spread, and the market volatility proxy.

Using these estimates for the cost increase, a back of the envelope calculation can be computed for the total welfare loss for institutional traders. Using July 2012 data, I find that average dollar volume per day on the affected set of stocks is roughly 20 billion. Given that majority of the daily trades are due to institutional investors, we can roughly attribute 10 billion of dollar volume to institutional investors. Thus, using implementation shortfall (VWAP Slippage) as our cost measure, the total loss due to the trading glitch is approximately \$18 (\$7) million.

6. Robustness Tests

In this section, we assess the robustness of our results with respect to alternative hypotheses and another identification of the affected group due to the trading glitch.

6.1. Identification with Abnormal Volume

As a second identification method, I use a very simple criteria to distinguish the affected stocks. For each NYSE-listed traded stock in my execution data, I compare the daily share volume on August 1, to its past 20-day average daily share volume. Formally, let abnormal volume ratio defined by

$$\Delta V^i = \frac{V^i_{t^*}}{V^i_{\mathsf{avg}}},$$

where $V^i_{t^*}$ corresponds to the daily volume of the ith stock on August 1, 2012 and V^i_{avg} denotes the average daily volume from trailing 20 trading days for the ith stock. If ΔV^i is greater than 2 and the stock is NYSE-listed, I label the ith stock to be *affected*. Stocks listed in other exchanges are directly excluded to be in the affected group due to the aforementioned company statement. I find that 31 stocks in my dataset satisfy this criteria and 47 executions are implemented on this set of stocks. It is worthwhile to note that 18 out of 31 stocks also appeared in our affected group identified in Section 5.1.

Figure 3 illustrates the number of securities that satisfies this criteria between June 2012 and August 2012. On most of the days, there are only a few stocks with this abnormal volume whereas on August 1, 2012, this number is at record 31. Table 8 displays this list of stocks with their corresponding abnormal volume ratio.

Table 8 illustrates the effect of the trading glitch using this identification strategy. I observe that my earlier results remain largely the same both qualitatively and quantitatively. As measured by implementation shortfall (VWAP slippage), affected executions realize an additional 23 bps (6 bps) on average.

6.2. Identification with NYSE Initial List of Affected Stocks

Various media outlets reported on August 1, 2012 that NYSE Euronext issued a note to traders about the trading glitch in the morning stating that it started an initial review of 148 symbols for trading irregularities between 09:30 a.m. and 10:15 a.m. Wall Street Journal published an initial list of these stocks at 12:41 p.m. even though there are actually 140 securities listed in the news article [7].

Table 4 displays the list of stocks which also appears on the executions on August 1, 2012 in my dataset. Since this is an initial list that NYSE reported right after the news of trading glitch, it seems to be not perfect. For each security in the list, Table 4 compares the number of trades in the first 30 minutes on the day of the trading glitch to the average number of trades in the same trading interval during the past week. For some stocks, average number of trades are much higher than what is observed on August 1, 2012, suggesting that they may not be exposed to irregular trading activity. All 42 securities are executed on the day of the trading glitch with a total number of 73 orders.

Table 9 illustrates the regression results when I use this set of securities as affected by the trading glitch. I observe that my earlier results remain largely the same. As measured by implementation shortfall (VWAP slippage), affected executions realize approximately an additional 11 bps (3 bps) on average. This additional cost estimates are economically substantial as the median implementation shortfall (VWAP slippage) is only 3 bps (1.2 bps) in the dataset.

6.3. Alternative Hypotheses

One possible alternative hypothesis is that the affected group of stocks might have been hit by the trading glitch because of an earlier shock to their trading costs. In this case, our methodology would incorrectly attribute the cost increase due to the trading glitch. In order to investigate this potential concern, I look at the costs of executing affected group of stocks before the occurrence of the trading glitch. For this purpose, I create a dummy variable, IsAffectedBefore, which takes a value of 1 if an execution happened within 5 business days before the trading glitch and the executed stock is in the affected group.

[Insert Table 10 here]

Table 10 illustrates the regression results with this new dummy variable using both cost measures. I find that there is no earlier shock in execution costs leading to the trading glitch and my

[7] "Trading Snafu: The List of the 148 Affected Stocks," *Wall Street Journal*, August 1, 2012. The news article does not mention a specific source. Through a personal communication, the reporter of the story told that the source is probably NYSE Euronext.

original conclusion regarding the cost increase due to the trading glitch remains unchanged with the same estimate.

7. Persistence in Cost Increase

In the previous sections, I provided strong evidence that executions exposed to the trading glitch on August 1, 2012 are subject to substantially higher cost. It would be interesting to investigate whether this cost increase is persistent as such a conclusion would drastically increase the impact of the rogue algorithm.

I will study whether there is any persistence in the cost increase using the available executions on the affected group in the following week (5 trading days). Using the affected 35 stocks from Section 5.1, there are 60 executions implemented on 35 different stocks in the upcoming week. I create another dummy variable, IsAffectedNextWeek, which takes a value of 1 if an execution happened within 5 business days after the trading glitch and the executed stock is in the affected group.

Table 11 illustrates the regression results with this new dummy variable using both cost measures. I find that for both cost measures, executions occurring in the following week is exposed to roughly the same cost increase. As measured by implementation shortfall (VWAP slippage), executions in the affected group realize on average an additional 15 bps (6 bps) in the following week. I also do not observe any change in the cost increase estimates for the day of the trading glitch.

These results highlight that the impact of the rogue algorithm may not be isolated to a single trading day of occurrence. In Section 5.4, I computed the approximate total dollar cost of institutional investors on August 1, 2012 as $18 million. Due to the persistence in the following week (additional 5 more trading days), the total cost of the trading glitch easily exceeds $100 million. This estimate is actually pretty conservative as given the commonality in liquidity, the total cost might be even more larger considering the spillover effects on other unaffected stocks.

8. Conclusion

Using an exogenous rogue algorithm of Knight Capital, a high-frequency market-making firm, leading to dramatic increases in number of trades for affected securities, I investigate whether institutional trading costs are negatively impacted by this seemingly unrelated shock. Given that the trading glitch results in tremendous volume, it is not clear why executions on the affected group of stocks can be more costlier after controlling for various stock- and order-level characteristics.

I first identify affected stocks using the abnormal number of trades observed due to the trading glitch. 35 stocks belonging to S&P 500 index, had actually more number of trades than the most actively traded security in equity markets, S&P 500 ETF, SPY. Using these stocks as the affected group, I first observe that executions on this affected group resulted in substantially higher costs as measured by implementation shortfall and VWAP slippage. Using two another identification strategies of the affected group, the estimate for the cost increase largely remain the same. Furthermore, the cost increase is also economically significant with an additional five-fold increase in both cost metrics for an average execution.

More importantly, the cost increase is not limited to the single event day. I find that all the affected executions in the following week suffer roughly the same cost increase underscoring the persistent impact of the shock. A back-of-the-envelope calculation shows that the total cost of the trading glitch to institutional traders is on the order of 100 million.

References

Yacine Aït-Sahalia and Mehmet Saglam. High frequency traders: Taking advantage of speed. Technical report, National Bureau of Economic Research, 2013.

Ekkehart Boehmer and Eric K Kelley. Institutional investors and the informational efficiency of prices. *Review of Financial Studies*, 22(9):3563–3594, 2009.

Jonathan Brogaard, Terrence Hendershott, and Ryan Riordan. High-frequency trading and price discovery. *Review of Financial Studies*, 27(8):2267–2306, 2014.

Karl B Diether, Kuan-Hui Lee, and Ingrid M Werner. Short-sale strategies and return predictability. *Review of Financial Studies*, 22(2):575–607, 2009.

Ian Domowitz, Jack Glen, and Ananth Madhavan. Liquidity, volatility and equity trading costs across countries and over time. *International Finance*, 4(2):221–255, 2001.

Terrence Hendershott, Charles M Jones, and Albert J Menkveld. Does algorithmic trading improve liquidity? *The Journal of Finance*, 66(1):1–33, 2011.

Charles M Jones. What do we know about high-frequency trading? *Columbia Business School Research Paper*, (13-11), 2013.

Andre F Perold. The implementation shortfall: Paper versus reality. *The Journal of Portfolio Management*, 14(3):4–9, 1988.

Mitchell A Petersen. Estimating standard errors in finance panel data sets: Comparing approaches. *Review of financial studies*, 22(1):435–480, 2009.

Lin Tong. A blessing or a curse? the impact of high frequency trading on institutional investors. *SSRN Working Paper*, 2013.

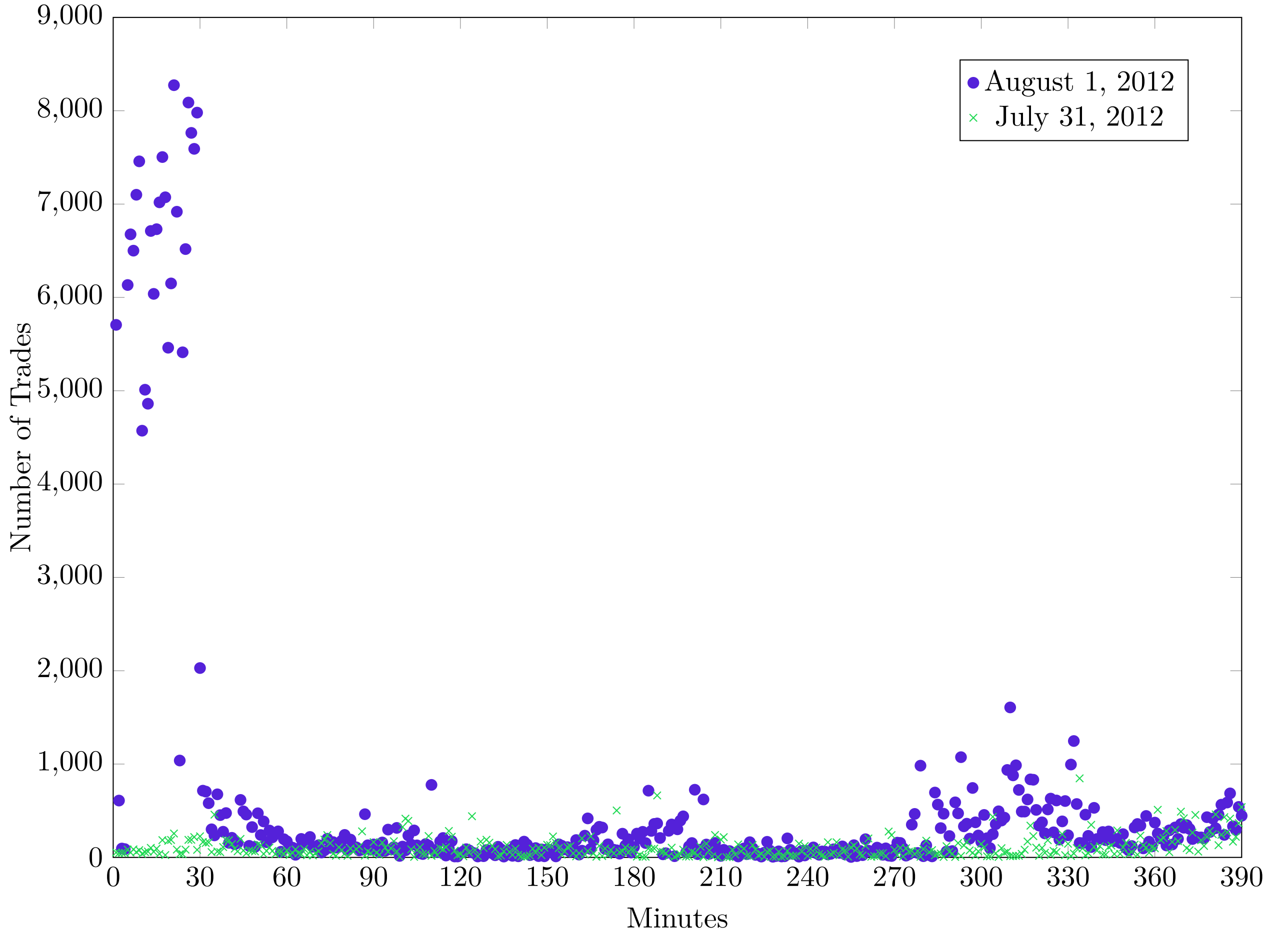

Figure 1: An illustration of the impact of the trading glitch by comparing the number of trades aggregated over each minute on NYSE:JNPR on July 31, 2012 and August 1, 2012.

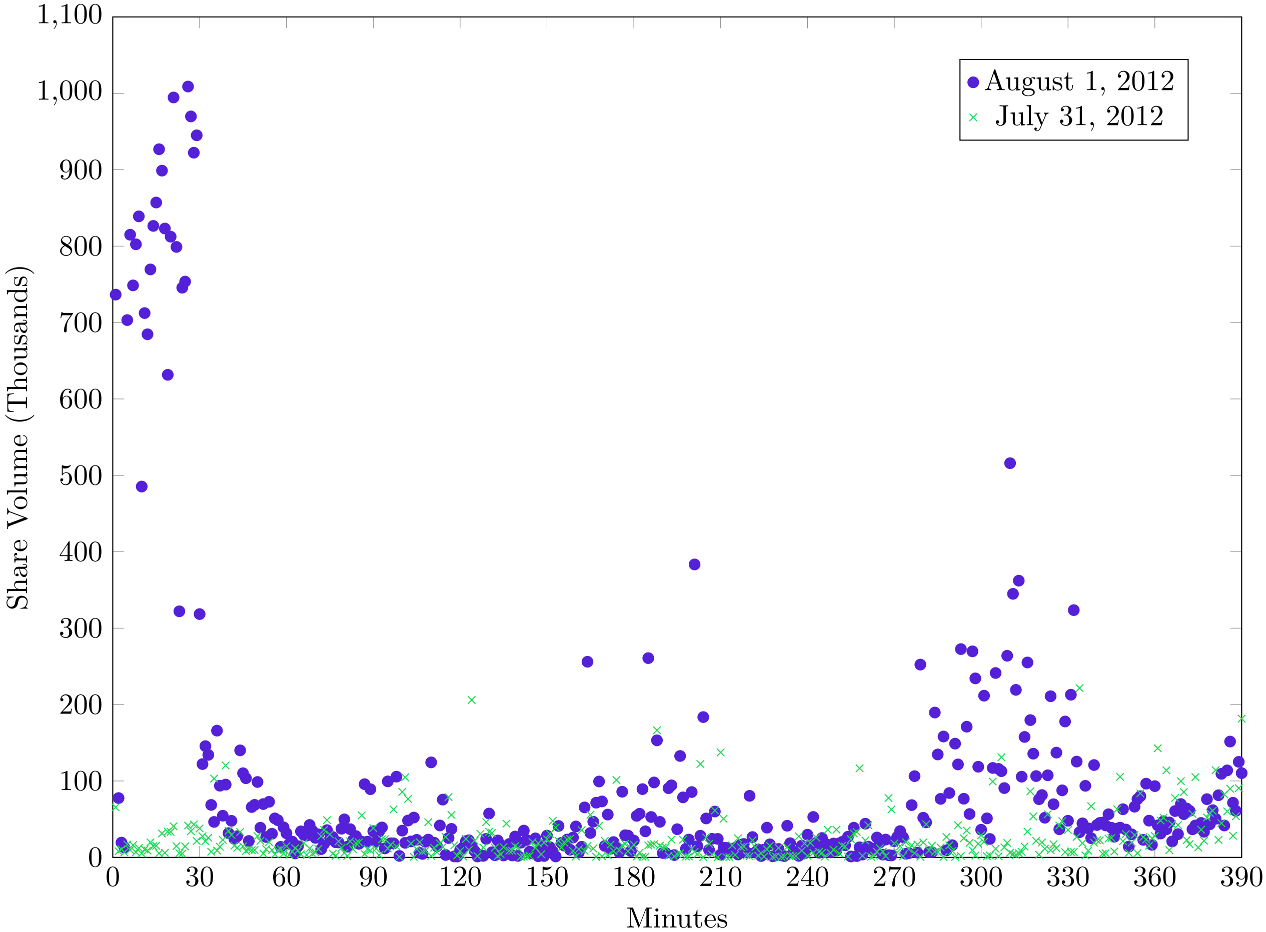

Figure 2: An illustration of the impact of the trading glitch by comparing the share volume aggregated over each minute on NYSE:JNPR on July 31, 2012 and August 1, 2012.

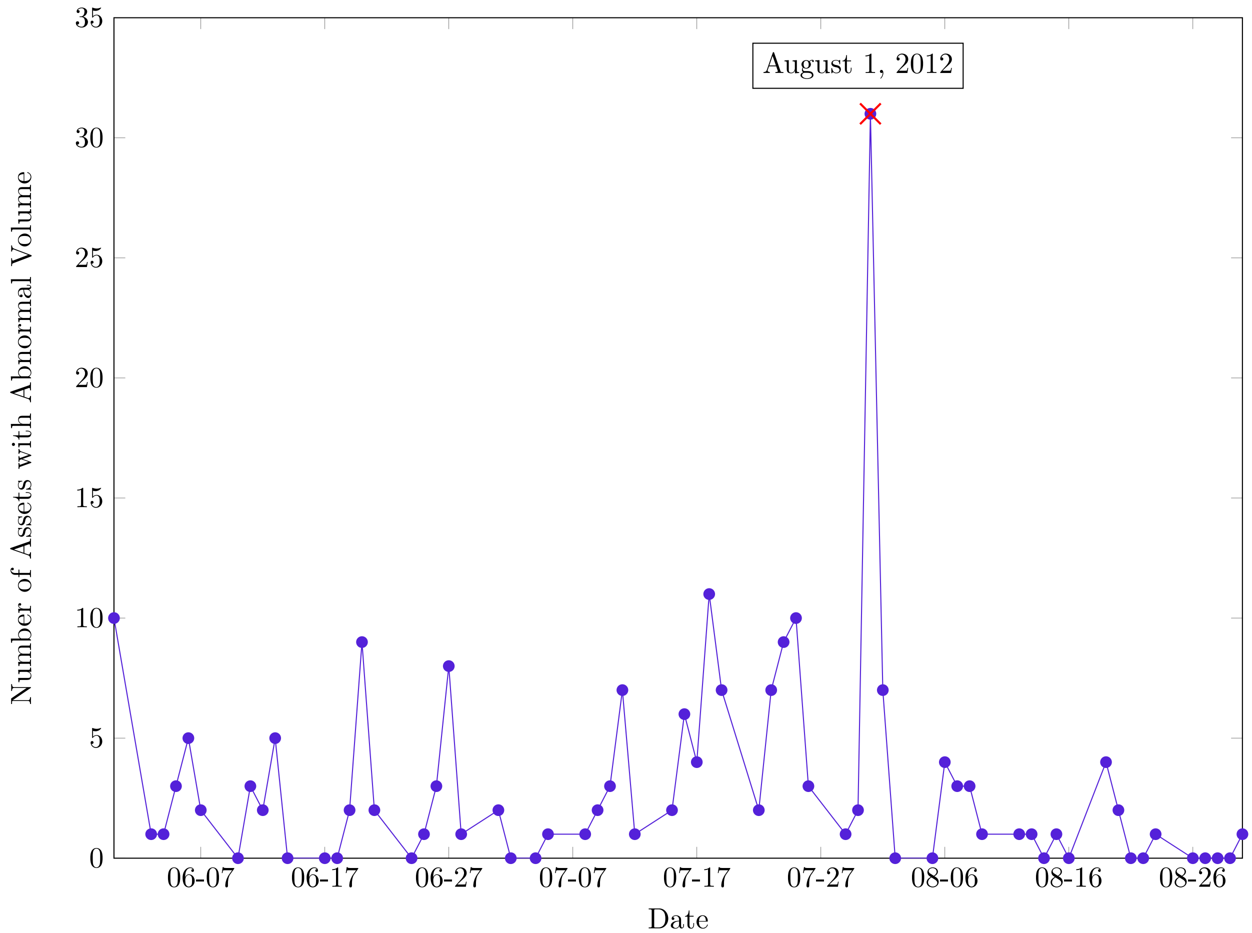

Figure 3: For every trading day between June 2012 and August 2012, this figure plots the number of securities in my execution data that have at least 100% increase in share volume compared to its average daily volume over the past month.

Table 1: 8 NYSE-listed securities had larger share volume than SPY, the most liquid ETF for S&P 500, during the initial 30 minutes of the trading.

Rank	Ticker	Share Volume
1	BAC	22,818,515
2	JNPR	20,963,198
3	F	19,913,150
4	GLW	14,213,554
5	C	13,222,612
6	AA	12,183,727
7	LSI	12,089,144
8	WFC	10,846,746
	SPY	10,842,391

Table 2: List of top 50 stocks which had the highest number of trades during the first 30 minutes of trading on August 1, 2012 compared with total number of trades observed on the previous trading day, July 31, 2012, for the same stock.

Rank	Symbol	Number of Trades between 9:30 and 10:00 on August 1, 2012	Number of Trades on July 31, 2012
1	JNPR	169118	40761
2	BAC	124200	86903
3	F	109102	70393
4	AA	102624	33242
5	LOW	94248	90262
6	C	91345	87162
7	LSI	81165	34372
8	GLW	80270	33678
9	WFC	71743	63523
10	PFE	67624	137883
11	BBY	60243	29313
12	EXC	58305	23653
13	GNW	57730	25242
14	AMT	56804	19056
15	WLP	55165	28018
16	DE	43236	21420
17	ANR	42413	54556
18	GT	41142	70202
19	HOG	40322	16259
20	DD	39901	23922
21	COH	39740	157866
22	SWN	39640	28182
23	PEP	38740	28566
24	JWN	37611	15880
25	BRK.B	37000	14625
26	MPC	36460	26508
27	AGN	36442	45871
28	ANF	35809	29762
29	CVS	34571	26820
30	AMD	34326	26948
31	AFL	33838	14045
32	VZ	33336	45183
33	DOW	32304	39277
34	HUM	31782	72389
35	GME	31572	8340
	SPY	**31341**	**267312**
36	FRX	26714	6546
37	NEM	20080	26443
38	T	17893	103633
39	CMCSA	17151	71820
40	WAG	15099	40258
41	AAPL	14208	96085
42	HPQ	12314	53872
43	PSX	12144	30345
44	FTR	11237	69763
45	VLO	10754	97558
46	GE	10727	81671
47	LUV	10278	26421
48	EA	9768	37867
49	AVP	9751	31330
50	SBUX	9635	72354

Table 3: List of stocks which had an abnormal volume ratio of greater than 2 when the share volume on August 1, 2012 is compared with the average daily volume over the past month.

Rank	Symbol	Abnormal Volume Ratio
1	HOG	6.53
2	JNPR	4.88
3	HSP	4.67
4	FRX	4.47
5	HUM	4.45
6	AGN	4.39
7	DNB	4.26
8	LH	4.17
9	AFL	3.87
10	JWN	3.86
11	AMT	3.78
12	GME	3.60
13	ANF	3.41
14	MPC	3.02
15	ALL	2.97
16	EXC	2.95
17	DE	2.92
18	CBG	2.82
19	MA	2.75
20	NYX	2.71
21	COH	2.58
22	DVN	2.43
23	QEP	2.43
24	NU	2.41
25	DD	2.39
26	PEP	2.27
27	GLW	2.25
28	LOW	2.19
29	PXD	2.17
30	TWX	2.04
31	AVP	2.02

Table 4: List of the 42 S&P 500 stocks that are on the initial list of affected stocks on WSJ article, "Trading Snafu: The List of the 148 Affected Stocks," *Wall Street Journal*, August 1, 2012. I compare the number of trades observed during the first 30 minutes on August 1, 2012 to average number of trades observed during the first 30 minutes over the past month.

Symbol	Number of Trades between 9:30 and 10:00 on August 1, 2012	Average Number of Trades between 9:30 and 10:00 over the past week
JNPR	169118	4446
BAC	124200	2050
F	109102	1712
AA	102624	4736
LOW	94248	1533
C	91345	2318
GLW	80270	3262
WFC	71743	15964
PFE	67624	4001
BBY	60243	15045
EXC	58305	8555
AMT	56804	5481
WLP	55165	9386
DE	43236	3974
HOG	40322	2592
DD	39901	2976
COH	39740	5847
SWN	39640	1579
PEP	38740	10444
JWN	37611	638
MPC	36460	12571
AGN	36442	1510
ANF	35809	5308
CVS	34571	930
AFL	33838	1047
VZ	33336	3269
DOW	32304	9398
HUM	31782	1286
GME	31572	3301
FRX	26714	4988
NEM	20080	2770
T	17893	1363
WAG	15099	5357
GE	10727	4430
LUV	10278	12873
CAT	8190	2778
CHK	6823	11738
GIS	4640	5492
AXP	4291	3755
AIG	4163	11279
WPI	3267	4836
LEN	3018	1139

Table 5: Summary statistics for the main attributes in our complete execution data. Average daily volatilities are computed using the previous 20 trading days before the execution date.

Statistic	Mean	St. Dev.	Min	1st Qu.	Median	3rd Qu.	Max
Participation Rate (%)	5.24	7.81	<0.01	0.13	0.89	8.91	100.00
Fraction of Daily Volume (%)	0.65	1.25	<0.01	0.04	0.22	0.71	33.27
Implementation Shortfall (%)	0.03	0.62	-10.06	-0.18	0.03	0.24	9.97
Average Daily Volatility (%)	1.21	0.43	0.13	0.91	1.14	1.42	7.34
Fraction of Trading Time (%)	41.18	40.54	1.28	4.84	18.24	93.30	100.00
Interval Return (%)	0.01	0.83	-18.10	-0.29	0.00	0.31	11.09

Table 6: Summary statistics for the main attributes of the execution data restricted to 662 executions implemented on August 1, 2012. Average daily volatilities are computed using the previous 20 trading days before the execution date.

Statistic	Mean	St. Dev.	Min	1st Qu.	Median	3rd Qu.	Max
Participation Rate (%)	1.27	3.18	0.01	0.17	0.51	1.05	34.97
Fraction of Daily Volume (%)	0.08	0.29	0.001	0.01	0.02	0.03	3.29
Implementation Shortfall (%)	0.001	0.33	−1.81	−0.11	−0.004	0.09	4.07
Average Daily Volatility (%)	1.31	0.44	0.61	0.99	1.24	1.54	3.33
Fraction of Trading Time (%)	11.56	21.56	1.33	3.38	3.72	15.34	100.00
Interval Return (%)	−0.10	0.44	−3.48	−0.24	−0.06	0.08	4.46

Table 7: Results of regressing implementation shortfall and VWAP slippage on the glitch-affected dummy and various control variables and calendar day dummies. Standard errors are given in parentheses and are adjusted by clustering on calendar day and stock as suggested by Petersen (2009).

	Dependent variable:	
	Implementation Shortfall	VWAP Slippage
IsAffected	17.42^{**}	6.69^{***}
	(8.08)	(1.62)
Participation Rate	61.92^{***}	25.28^{***}
	(7.07)	(2.65)
Fraction of Daily Volume	102.66^{**}	12.34^{**}
	(43.57)	(5.30)
Relative Order Size	3.21	1.72
	(4.09)	(1.23)
Relative Daily Volume	2.22	−0.72
	(1.36)	(0.50)
IsEarningsDay	8.15^{**}	−0.51
	(3.73)	(0.79)
Market Capitalization	−0.23	0.52^{***}
	(0.44)	(0.14)
Interval Turnover	−0.29	0.23^{***}
	(0.32)	(0.08)
Day Turnover	-368.13^{**}	−10.05
	(150.36)	(52.92)
Average Turnover	0.44^{***}	−0.002
	(0.15)	(0.05)
Interval Volatility	214.23^{*}	−10.05
	(117.26)	(28.56)
Day Volatility	-704.39^{**}	−62.65
	(315.53)	(52.72)
Average Volatility	353.48	−31.08
	(330.38)	(68.25)
Interval Spread	−0.13	0.12^{**}
	(0.18)	(0.05)
Day Spread	0.10^{***}	0.02^{***}
	(0.02)	(0.01)
Average Spread	0.01	0.09
	(0.11)	(0.05)
Execution Duration Fraction	3.75	−0.02
	(4.13)	(1.32)
abs(Day Return)	217.08^{**}	9.85
	(98.74)	(16.01)
abs(Interval Return)	0.02	0.01
	(0.03)	(0.01)
abs(Market Return)	$2{,}297.62^{***}$	$1{,}751.05^{***}$
	(735.01)	(103.17)
abs(Prior Market Return)	−421.78	$-1{,}636.55^{***}$
	(503.25)	(101.58)
abs(Prior Day Return)	13.99	−5.84
	(48.16)	(9.80)
N	39,068	39,068
R^2	4.1%	7.5%
Adjusted R^2	3.5%	6.9%

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 8: Results of regressing implementation shortfall and VWAP slippage on the glitch-affected dummy and various control variables and calendar day dummies. I identify the affected group using abnormal volume ratios of greater than 2. Standard errors are given in parentheses and are adjusted by clustering on calendar day and stock as suggested by Petersen (2009).

	Dependent variable:	
	Implementation Shortfall	VWAP Slippage
IsAffected	23.41^{***}	6.34^{***}
	(6.68)	(1.62)
Participation Rate	61.99^{***}	25.29^{***}
	(7.06)	(2.65)
Fraction of Daily Volume	102.55^{**}	12.39^{**}
	(43.50)	(5.30)
Relative Order Size	3.33	1.75
	(4.10)	(1.23)
Relative Daily Volume	2.13	−0.74
	(1.36)	(0.51)
IsEarningsDay	7.96^{**}	−0.58
	(3.76)	(0.79)
Market Capitalization	−0.21	0.52^{***}
	(0.44)	(0.14)
Interval Turnover	−0.29	0.23^{***}
	(0.32)	(0.08)
Day Turnover	-364.15^{**}	−9.41
	(150.69)	(53.02)
Average Turnover	0.44^{***}	−0.002
	(0.15)	(0.05)
Interval Volatility	211.75^{*}	−11.81
	(118.06)	(28.84)
Day Volatility	-713.55^{**}	−57.40
	(302.53)	(52.91)
Average Volatility	366.66	−33.03
	(323.98)	(68.10)
Interval Spread	−0.13	0.13^{**}
	(0.18)	(0.05)
Day Spread	0.10^{***}	0.02^{***}
	(0.02)	(0.01)
Average Spread	0.01	0.09
	(0.11)	(0.05)
Execution Duration Fraction	3.64	−0.04
	(4.13)	(1.32)
abs(Day Return)	215.69^{**}	8.94
	(98.20)	(16.16)
abs(Interval Return)	0.02	0.01
	(0.03)	(0.01)
abs(Market Return)	$2{,}290.76^{***}$	$1{,}757.93^{***}$
	(730.56)	(102.05)
abs(Prior Market Return)	−412.43	$-1{,}640.88^{***}$
	(501.26)	(100.33)
abs(Prior Day Return)	13.01	−6.10
	(48.20)	(9.86)
N	39,068	39,068
R^2	4.1%	7.5%
Adjusted R^2	3.5%	6.9%

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 9: Results of regressing implementation shortfall and VWAP slippage on the glitch-affected dummy and various control variables and calendar day dummies. I identify the affected group using the initial list of affected stocks published by WSJ. Standard errors are given in parentheses and are adjusted by clustering on calendar day and stock as suggested by Petersen (2009).

	Dependent variable:	
	Implementation Shortfall	VWAP Slippage
IsAffected	11.27^{*}	3.20^{***}
	(6.57)	(1.25)
Participation Rate	61.89^{***}	25.27^{***}
	(7.07)	(2.65)
Fraction of Daily Volume	102.80^{**}	12.45^{**}
	(43.58)	(5.31)
Relative Order Size	3.18	1.71
	(4.09)	(1.24)
Relative Daily Volume	2.23	−0.72
	(1.36)	(0.50)
IsEarningsDay	8.12^{**}	−0.54
	(3.74)	(0.79)
Market Capitalization	−0.23	0.52^{***}
	(0.44)	(0.14)
Interval Turnover	−0.29	0.23^{***}
	(0.33)	(0.08)
Day Turnover	-368.32^{**}	−10.49
	(150.55)	(52.95)
Average Turnover	0.44^{***}	−0.001
	(0.15)	(0.05)
Interval Volatility	211.19^{*}	−11.87
	(117.76)	(28.85)
Day Volatility	-687.34^{**}	−50.99
	(319.95)	(55.26)
Average Volatility	341.03	−39.48
	(331.73)	(67.30)
Interval Spread	−0.13	0.13^{**}
	(0.18)	(0.05)
Day Spread	0.10^{***}	0.02^{***}
	(0.02)	(0.01)
Average Spread	0.01	0.09
	(0.11)	(0.05)
Execution Duration Fraction	3.77	−0.01
	(4.13)	(1.32)
abs(Day Return)	215.72^{**}	8.99
	(98.75)	(16.15)
abs(Interval Return)	0.02	0.01
	(0.03)	(0.01)
abs(Market Return)	$2{,}316.21^{***}$	$1{,}764.03^{***}$
	(737.51)	(103.04)
abs(Prior Market Return)	−436.38	$-1{,}646.75^{***}$
	(505.00)	(100.96)
abs(Prior Day Return)	14.19	−5.78
	(48.15)	(9.78)
N	39,068	39,068
R^2	4.1%	7.5%
Adjusted R^2	3.4%	6.8%

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 10: Results of regressing implementation shortfall and VWAP slippage on the glitch-affected dummy and various control variables and calendar day dummies. I create a dummy variable, IsAffectedBefore, to control for any potential cost increase before the trading glitch. Standard errors are given in parentheses and are adjusted by clustering on calendar day and stock as suggested by Petersen (2009).

	Dependent variable:	
	Implementation Shortfall	VWAP Slippage
IsAffected	17.43^{**}	6.69^{***}
	(8.09)	(1.62)
IsAffectedBefore	-10.21	-0.50
	(7.02)	(1.38)
Participation Rate	61.83^{***}	25.28^{***}
	(7.06)	(2.65)
Fraction of Daily Volume	102.83^{**}	12.35^{**}
	(43.58)	(5.30)
Relative Order Size	3.19	1.72
	(4.09)	(1.23)
Relative Daily Volume	2.23	-0.72
	(1.36)	(0.50)
IsEarningsDay	8.10^{**}	-0.51
	(3.72)	(0.79)
Market Capitalization	-0.22	0.52^{***}
	(0.44)	(0.14)
Interval Turnover	-0.29	0.23^{***}
	(0.32)	(0.08)
Day Turnover	-368.41^{**}	-10.07
	(150.35)	(52.92)
Average Turnover	0.44^{***}	-0.002
	(0.15)	(0.05)
Interval Volatility	213.92^{*}	-10.07
	(117.24)	(28.56)
Day Volatility	-705.26^{**}	-62.69
	(315.62)	(52.73)
Average Volatility	357.35	-30.89
	(330.39)	(68.25)
Interval Spread	-0.13	0.12^{**}
	(0.18)	(0.05)
Day Spread	0.10^{***}	0.02^{***}
	(0.02)	(0.01)
Average Spread	0.01	0.09
	(0.11)	(0.05)
Execution Duration Fraction	3.75	-0.02
	(4.13)	(1.32)
abs(Day Return)	217.40^{**}	9.87
	(98.71)	(16.00)
abs(Interval Return)	0.02	0.01
	(0.03)	(0.01)
abs(Market Return)	$2{,}291.53^{***}$	$1{,}750.75^{***}$
	(734.68)	(103.00)
abs(Prior Market Return)	-418.44	$-1{,}636.39^{***}$
	(502.89)	(101.52)
abs(Prior Day Return)	14.12	-5.83
	(48.05)	(9.80)
N	39,068	39,068
R^2	4.1%	7.5%
Adjusted R^2	3.5%	6.9%

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 11: Results of regressing implementation shortfall and VWAP slippage on the glitch-affected dummy and various control variables and calendar day dummies. I create a dummy variable, IsAffectedNextWeek, to account for any persistent cost increase in the following week. Standard errors are given in parentheses and are adjusted by clustering on calendar day and stock as suggested by Petersen (2009).

	Dependent variable:	
	Implementation Shortfall	VWAP Slippage
IsAffected	16.94^{**}	6.56^{***}
	(8.03)	(1.60)
IsAffectedNextWeek	15.00^{**}	5.95^{***}
	(6.18)	(1.69)
Participation Rate	62.13^{***}	25.34^{***}
	(7.06)	(2.65)
Fraction of Daily Volume	101.82^{**}	11.73^{**}
	(43.41)	(5.33)
Relative Order Size	3.27	1.72
	(4.09)	(1.23)
Relative Daily Volume	2.20	−0.72
	(1.35)	(0.50)
IsEarningsDay	8.09^{**}	−0.52
	(3.73)	(0.80)
Market Capitalization	−0.23	0.51^{***}
	(0.44)	(0.14)
Interval Turnover	−0.30	0.23^{***}
	(0.33)	(0.08)
Day Turnover	-368.73^{**}	−11.04
	(150.42)	(52.78)
Average Turnover	0.45^{***}	−0.001
	(0.15)	(0.05)
Interval Volatility	213.16^{*}	−8.62
	(117.16)	(28.48)
Day Volatility	-694.24^{**}	−60.68
	(313.91)	(52.25)
Average Volatility	334.15	−34.51
	(329.98)	(67.96)
Interval Spread	−0.12	0.12^{**}
	(0.17)	(0.05)
Day Spread	0.10^{***}	0.02^{***}
	(0.02)	(0.01)
Average Spread	0.01	0.09^{*}
	(0.11)	(0.06)
Execution Duration Fraction	3.73	0.01
	(4.13)	(1.32)
abs(Day Return)	215.69^{**}	9.76
	(98.73)	(16.00)
abs(Interval Return)	0.02	0.01
	(0.03)	(0.01)
abs(Market Return)	$2{,}322.35^{***}$	$1{,}755.62^{***}$
	(735.25)	(102.94)
abs(Prior Market Return)	−440.35	$-1{,}640.36^{***}$
	(503.17)	(101.35)
N	39,068	39,068
R^2	4.1%	7.5%
Adjusted R^2	3.5%	6.9%

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$