

No. 636

Jorge Gonçalves, Roman Kräussl and Vladimir Levin

Do “Speed Bumps” Prevent Accidents in Financial Markets?

The CFS Working Paper Series

presents ongoing research on selected topics in the fields of money, banking and finance. The papers are circulated to encourage discussion and comment. Any opinions expressed in CFS Working Papers are those of the author(s) and not of the CFS.

The Center for Financial Studies, located in Goethe University Frankfurt's House of Finance, conducts independent and internationally oriented research in important areas of Finance. It serves as a forum for dialogue between academia, policy-making institutions and the financial industry. It offers a platform for top-level fundamental research as well as applied research relevant for the financial sector in Europe. CFS is funded by the non-profit-organization Gesellschaft für Kapitalmarktforschung e.V. (GfK). Established in 1967 and closely affiliated with the University of Frankfurt, it provides a strong link between the financial community and academia. GfK members comprise major players in Germany's financial industry. The funding institutions do not give prior review to CFS publications, nor do they necessarily share the views expressed therein.

Do “Speed Bumps” Prevent Accidents in Financial Markets?*

Jorge Gonçalves, Roman Kräussl and Vladimir Levin[†]

This version - July 2019

Abstract

Is it true that speed bumps level the playing field, make financial markets more stable and reduce negative externalities of high-frequency trading (HFT) firms? We examine how the implementation of a particular speed bump – Midpoint Extended Life order (M-ELO) on Nasdaq impacted financial markets stability in terms of occurrences of mini-flash crashes in individual securities. We use high-frequency order book message data around the implementation date and apply difference-in-differences analysis to estimate the average treatment effect of the speed bump on market stability and liquidity provision. The results suggest that the introduction of the M-ELO decreases the average number of crashes on Nasdaq compared to other exchanges by 4.7%. Liquidity provision by HFT firms also improves. These findings imply that technology-based solutions by exchanges are feasible alternatives to regulatory intervention towards safer markets.

JEL classification: C21, G14, G18

Key-words: Mini-flash crash; speed bump; midpoint extended life order.

*We are grateful to Nasdaq Global Data Products for the data. All errors are our own.

[†]Jorge Gonçalves (jorge.goncalves@uni.lu) is affiliated with the University of Luxembourg, the Department of Engineering at the University of Cambridge. Roman Kräussl (roman.kraussl@uni.lu) is affiliated with the Luxembourg School of Finance, the Hoover Institution at Stanford University, the Center for Alternative Investments at Goizueta Business School, Emory University in Atlanta, and the Center for Financial Studies in Frankfurt am Main. Vladimir Levin (vladimir.levin@uni.lu) is affiliated with the Luxembourg School of Finance.

1 Introduction

Rule 611 of Regulation National Market System (NMS) was designed to establish intermarket price protection in financial markets. It ensures that market participants have immediate, fair, and efficient access to exchange quotes in all automated markets. Despite obvious advantages of Rule 611, academic literature [Budish et al., 2015, Pagnotta and Philippon, 2018] argues that it results in fragmented markets characterized by high speed but low allocative efficiency. To tackle with negative externalities coming from such high-frequency market structure Biais et al. [2014] and Biais et al. [2015] propose the models of “slow market”. The important note here is that the authors do not propose to ban all high-frequency trading but to create a segment of slow-friendly markets and to leave room for the investment in the fast trading technology. Financial markets seem to find the technological solution to deemphasize speed externalities. Recently exchanges began to add intentional access delays to their feeds to improve market quality and even price discovery [Hu, 2018]. In the industry, these delays are called “speed bumps” because their functioning resembles traffic speed bumps. The focus of this paper is the effect of speed bumps on market stability. We find that the introduction of a speed bump favors overall market stability by increasing liquidity provision and reducing the occurrence of extreme price movements such as mini-flash crashes.

The financial markets’ speed bump mainly responds to the speed of high-frequency trading (HFT) firms, who represent today a considerable part of all U.S. financial market participants. According to Easley et al. [2012] and Meyer et al. [2018], HFT firms account for more than 55% of the volume on equity markets and approaching 85% of the volume on futures markets. One of the common roles of HFT firms is providing liquidity for other market participants [Hasbrouck and Saar, 2013, Menkveld, 2013, Conrad et al., 2015], but there is a growing literature documenting that during periods of uncertainty they may start demanding liquidity instead [Chung and Chuwonganant, 2014]. Moreover, HFT firms are blamed for using their speed advantage to trade against less sophisticated retail traders. A number of manipulative strategies [Brunnermeier and Pedersen, 2005, Cumming et al., 2011] is known to be used by the HFT firms in order to scalp profits from “slow” traders.

Previous research suggests HFT activity worsen market stability by provoking extreme price movements also known as flash crashes. [Golub et al. \[2012\]](#) study the increase in the number of mini-flash crashes in individual securities for the period of 2006-2011 and suggest those crashes are caused by HFT. [Leal et al. \[2016\]](#) build an agent-based model to study how the interplay between low- and high-frequency trading affects asset price dynamics. They find that the presence of HFT increases market volatility and plays a fundamental role in the generation of flash crashes. Later [Kirilenko et al. \[2017\]](#) examine the structure of the E-mini S&P 500 stock index futures market on May 6, 2010¹ and report that HFT indeed exacerbated market volatility.

[Biais et al. \[2015\]](#) argues that investment in fast trading technology helps financial institutions to cope with market fragmentation, and improves social welfare but simultaneously creates adverse selection. One way to mitigate the adverse consequences of fast traders is to create “slow-only” markets and to impose taxes on investments in fast trading technology. Instead of regulatory interventions, we consider the new kind of speed competitions among the exchanges. This competition works in the opposite direction – slowing down. Our paper discusses how exchanges bring to life one of the proposals of [Biais et al. \[2015\]](#) to create a “slow” market. This paper adds to the recent literature on regulation and market design by studying the effect of intentional access delays and minimum resting times on stability of equity markets.

Since the year of 2016 exchanges started to introduce their own technology-based solutions to protect interests of long-term investors and reduce the HFT firms domination as well as the adverse selection of participants who do not have privileged high-speed access to market information. Those solutions were implemented in the form of latency delays and are known now as speed bumps. The Investors Exchange (IEX) applied the first such measure: it introduced a 350 microseconds delay to all incoming and outgoing correspondence. [Hu \[2018\]](#) studies improvements in market functioning around the period when IEX becomes a national securities exchange. He documents a positive impact of the speed bump on the market quality in terms of tighter spreads and improved liquidity. The increasing market

¹The date of the most notorious Flash Crash in US markets.

share of IEX suggests that protected by speed bump markets are attractive for participants. Some other exchanges later implemented the IEX’s approach and launched their own versions of speed bumps². For example, Nasdaq created a new order type for a submission called Midpoint Extended Life Order (M-ELO) which is targeted to long-term investors and will make HFT firms to abstain from using it.

We focus on the M-ELO implemented by Nasdaq as a “voluntary” speed bump. The M-ELO is a pegged to the midprice order type which is also not displayed but trade reported like any other order, without any new or special indication. It provides a valuable tool to investors seeking to find liquidity in size, with the midpoint of the National Best Bid and Offer (NBBO). Anonymity and confidentiality of the M-ELO are the critical tools in preventing potentially predatory counterparties from determining intention and using that information to generate short-term profits at the expense of longer-term investors. This order type does not interact with other orders on the Nasdaq market that does not meet the requirements of a M-ELO, particularly the proposed holding period of half of a second. Nothing, however, obliges market participants from canceling or changing the order thus giving to them the freedom to risk manage their exposure. Only if the holding period requirement was met the M-ELO will become eligible for execution against an opposite side M-ELO order which fulfilled all the requirements as well.

This paper studies whether the M-ELO introduction improved financial markets’ stability by reducing HFT externalities. We apply the difference-in-differences analysis for the case of Nasdaq and investigate the effects of HFT around March 12, 2018, when the M-ELO was put into operation. As a proxy of market stability, we use the occurrence of mini-flash crashes in individual securities. We also assess the effects of this new order type on liquidity provision using such proxies as cancel-to-trade and trade-to-order ratios. We find robust positive effects of the M-ELO introduction for both market stability and liquidity provision, thus supporting previous literature that speed bumps are effective tools for deemphasizing speed externalities without regulatory intervention.

The remainder of the paper is organized as follows. Section 2 discusses data sources and

²For a comprehensive description of speed bumps see [Woodward \[2018\]](#) and [Hu \[2018\]](#)

describes the identification strategy of mini-flash crashes. It continues with summary statistics and outlines the model to be estimated. Section 3 presents our empirical results and discusses various robustness checks. Finally, Section 4 concludes.

2 Data and Methodology

2.1 Data

The Order Book Message data come from Nasdaq historical ITCH and span 42 trading days: from February 12, 2018, to April 13, 2018. This period covers the introduction of the M-ELO on Nasdaq on March 12, 2018. For each order message, the data set contains its order type, which can be one of the following: order added, order deleted, order canceled, order executed, order replaced. The data allow us to directly observe liquidity provision on each depth level of the limit order book at any period of time. An illustration of the data is provided in the Appendix.

The limitations of this data are straightforward and similar to what previous research encountered [Carrion, 2013, Brogaard et al., 2014, O’Hara et al., 2014, Brogaard et al., 2018]. We do not observe individual HFT activity as well as trading activity on other venues. Trades on Nasdaq account for on average 40% of trading activity for Nasdaq listed stocks and about 15% for stocks with other primary listing venues. Despite the high fragmentation nature of financial markets, we share the belief of Brogaard et al. [2018] that liquidity transfers to other venues are unlikely due to the short time period of interest and overall similar liquidity provision rules among exchanges. Securities characteristics to serve as covariates in our analysis were obtained from the SEC’s MIDAS Market Structure Metrics database.

2.2 Mini-Flash Crash Identification

We identify mini-flash crashes in NBBO midprice. When HFT firms search for hidden orders inside the spread they submit and immediately cancel limit orders forcing bid and ask prices

to fluctuate rapidly without any trades occurring. Using changes in midprice we are able to reduce this microstructure noise created by the strong coupling of the best bid and best ask prices. We use the order book messages to reconstruct the limit order book at the time of every quote update which happens during the continuous trading period and obtain the midprice as an average of NBBO prices. The data set consists of over 8,000 equities and ETFs listed on Nasdaq, NYSE, NYSE American, NYSE Arca and BATS exchanges. Medium and small stocks by market capitalization are not traded frequently enough. We follow considerations of similar analyses of [Andersen et al. \[2001\]](#) and [Brogaard et al. \[2018\]](#) and preserve the sufficiently big number of observations by focusing our analysis on the largest stocks.

To identify mini-flash crashes we use the following approach. The trading day for a stock is split into n intervals of length t which we call identification intervals. The length of each interval should satisfy: (i) $t \leq 5$ minutes and (ii) the interval contains 50 trades on average. If both these conditions cannot hold simultaneously we consider the stock not to be liquid enough for our analysis. Thus, depending on the trading intensity of a stock the length of the identification interval can vary in the range $(0, 300]$ seconds. Then, we label an interval $[(i - 1) \cdot t, i \cdot t]$ as one containing a mini-flash crash if the following conditions hold: (i) the maximum absolute mid price return of a stock within this interval exceeds 0.8%, (ii) the price volatility within $[(i - 1) \cdot t, i \cdot t]$ is at least 10 times greater than in the previous period $[(i - 2) \cdot t, (i - 1) \cdot t]$, (iii) the price reverts at least 33% after the crash. To account for the price reversal we calculate the cumulative return from the crash starting point to its maximum absolute value. The time length of this event we call the duration of the mini-flash crash within interval i , or just τ_i . The starting time of a crash in the identification interval is chosen in the way that it will lead to the maximum possible absolute cumulative return. Thus, t is the upper bound for the crash duration ($\tau_i \leq t$). After calculating τ_i we calculate which fraction of the cumulative return had the stock recovered in the next τ_i -long time window. Therefore, we document a reversal if the price retraced at least 33% of the initial change.

This identification technique is in line with previous works of [Golub et al. \[2012\]](#) and [Johnson](#)

et al. [2013] where authors do not do a sharp distinction between crashes and spikes and just require the price to move fast and severe. Similar to the approach of Bellia et al. [2018] we study only those mini-flash crashes which possess a transitory dynamics and do not exhibit a permanent effect only. In our approach, we add a second dimension to the previous identification techniques in order to capture mini-flash crashes better. This dimension is represented by price volatility. In fact, in our approach, the interval is labeled as a mini-flash crash if both return and price volatility are far in the right tail of their distributions. In the Appendix, we address possible limitations of our approach by repeating our analysis using a non-parametric test for “drift burst” detection in stock returns proposed by Aït-Sahalia et al. [2009].

2.3 Summary Statistics

In our identification, we do not consider crashes detected within first and last 10 minutes of the continuous trading period since opening and closing auctions processes are governed by different rules and the price is usually extremely volatile near those times [Barclay and Hendershott, 2003]. Table 1 reports descriptive statistics related to the main features of mini-flash crashes for the sample of NASDAQ listed stocks (Panel A) and those which are listed elsewhere (Panel B).

Table 1 around here

The average absolute return during the mini-flash crash is around 2.2% on NASDAQ and about 1.6% on other exchanges. The duration of crashes is about 10% lower in mean and in quartiles for NASDAQ listed stocks, being on average 86.7 seconds. There is no sizable change between exchanges in terms of price volatilities during the crash days, as well as in terms of reversals. The huge difference is, however, in market activity around the time of a mini-flash crash: the average number of trades during the crash period is about 148 for NASDAQ listed stocks and 236 for other ones. The dollar trading volume difference is even higher: non-NASDAQ listed stocks are traded for about \$1.37 million which is more than two times higher than on NASDAQ listed ones for which the volume is almost \$640 thousand.

Figure 1 presents the example of the mini-flash crash in Procter&Gamble (PG) on March 21, 2018, identified by our approach. Panel A spans opening auction, continuous trading period and closing auction. This example illustrates the typical dynamics of midprice during a mini-flash crash. At 2 p.m. the price experienced a rapid massive spike of a magnitude about 1%, but within the next five minutes the price dropped more than 1.4% down and eventually returned back to the region of previous daily consolidation. Panel B in which each dot is a trade is a zoomed representation of the identified crash. The crash did not trigger the circuit breaker, even though the initial spike accounted for more than 46 standard deviations of price time series.

Figure 1 around here

This mini-flash crash in the price of PG has the following properties: duration of the crash is 26.2 seconds (the initial spike), cumulative return during the crash is 0.98%, the spike move accounted for more than 46 standard deviations of the previous price dynamics, the dollar volume traded during the crash exceeds \$633 thousand.

Table 2 reports summary statistics on market capitalization and some HFT activity proxies for stocks listed on Nasdaq and on other exchanges (for companies listed on NYSE, NYSE Arca, NYSE American and BATS Z the aggregated statistics is reported). Panel A gives the broad overview of the companies' characteristics and provides a comparison between developing and high tech companies listed on Nasdaq and well established high market capitalization companies who have their primary listings on NYSE Group of exchanges, (BATS Z and NYSE American specialize mostly on small companies and constitute less than 0.2% in our sample since we focus on large capitalization companies). Panel B provides a similar comparison between Nasdaq and non-Nasdaq listed securities for which we observed mini-flash crashes with price recoveries.

Table 2 around here

A cancel-to-trade ratio is one of the proxies of liquidity providing activity. This is the ratio of the number of order cancellations to the number of executions for a particular security. Market makers operate mostly through limit orders and cancel many of them in order to

quickly adjust to the current situation on the market. According to cancel-to-trade ratio, there is almost double market making activity on other exchanges than on Nasdaq with levels of 18.32 and 9.95 respectively. Another proxy is trade-to-order ratio which shows what fraction of the initial orders was eventually executed during the continuous trading period. The average ratio for Nasdaq listed securities suggests that about 16% of orders were traded on average during a day which is almost twice as high as the ratio for non-Nasdaq listed companies. The odd-to-trade ratio, which is represented by the fraction of odd trades executed to the number of all trades, documents similar phenomenon. About 41% of trades executed on Nasdaq were not equal a multiple of a round lot versus 33% for non-Nasdaq listed stocks. The average hidden rate which is the ratio of non-displayed executions to the total number of trades is comparable between the two groups and a little higher for Nasdaq listed companies being equal to 15% versus 11% on other exchanges.

When investigating the stocks for which we identify mini-flash crashes (Panel B) we document the substantial difference in the market capitalization of companies that are prone to crashes and the companies from the general sample. For instance, the median capitalization of a company who at least once experienced a mini-flash crash during the period in question is about \$0.98 bn., against the median of \$12.43 bn. of the total sample of Nasdaq companies, selected for the study. In Panel B the difference between exchanges in terms of market capitalization becomes more severe, and so the median capitalization on non-Nasdaq listed companies is equal to \$4.39 bn which is four times the corresponding median capitalization on Nasdaq. Other characteristics as well became relatively more distant: the average cancel-to-trade ratio on NYSE exchanges is almost three times higher than on Nasdaq (14 versus 5.13), trade-to-order ratio for Nasdaq stocks is now twice as high as for NYSE exchanges, while Hidden rate and odd-to-trade rate increased only slightly.

In general, we observe different features of stock characteristics across the primary listing exchange as well as across the type of a crash in several aspects. Mini-Flash crashes with subsequent price recovery present objects of main interest for us, therefore the sample of 865 detected crashes will be used to perform the analysis of the effects of the M-ELO on the market stability.

2.4 Difference-in-Differences Model

The introduction of a M-ELO creates the conditions close to the natural experiment, which allow evaluating its effect on the negative externalities of HFT characterized by the number of times the market experience a mini-flash crash in a single security. We compare stocks listed on Nasdaq and on the other exchanges: NYSE, NYSE American, NYSE Arca and BATS. Treatment takes place on the implementation day of March 12, 2018. The possibility to submit M-ELOs was granted to all securities, thus we define the “treatment group” as the set of symbols having primary listings on Nasdaq and the “control group” as the set of symbols listed on other exchanges. Treatment group symbols are most likely to be impacted by the M-ELO whereas control group stocks are affected by the new order type to much lower extent due to a relatively small share of trading on other exchanges.

To identify the treatment effect we use a difference-in-differences (DID) design. It is beneficial due to the following reasons. First, the timing of treatment is exogenous with respect to stock characteristics and the same for all sample. Second, since the treatment group assignment is based on the primary listing exchange the outcome of whether the stock falls into treatment or control groups should not be affected by unobservable factors. We run the following regression to estimate the Average Treatment Effect on the Treated group (ATET):

$$Y_{it} = \beta_0 + \beta_1 \lambda_t + \beta_2 Exchange_i + \beta_3 D_{it} + \beta_4 X_{it} + \epsilon_{it} \quad (1)$$

where Y_{it} is a measure of interest (either the average number of mini-flash crashes, price volatility, dollar trading volume, etc.), λ_t is the day of observation dummy variable, $Exchange_i$ is an indicator which equals one if the company is primarily listed on Nasdaq and zero otherwise, $D_{it} \equiv Exchange_i \times d_t$ is a dummy variable which equals one for treatment units in the post-treatment period ($Exchange_i = 1, d_t = 1$) and is zero otherwise, X_{it} is a vector of individual-specific covariates. In this case, the estimate of β_3 coefficient represents the difference in the average result of actual treatment and a counter-factual effect on treatment group if there would not be any treatment, so this is indeed the average treatment effect on the treated group (ATET).

The identification assumption one need to make in order to infer the treatment effect can be represented below:

$$E(\epsilon_{it} | Exchange_i, d_t, \lambda_t, X_{it}) = 0.$$

It is commonly known as the parallel trend assumption. To support this assumption consider Figure 2 which represents trends in the occurrences of mini-flash crashes experienced across all stocks listed on Nasdaq and on other exchanges around the implementation day. Before March 12 the number of crashes observed display almost equal downward trends, although trend estimates in Nasdaq’s case are noisier. There is a decrease in the average number of crashes per stock happening on the Nasdaq after the change. This mainly supports the hypothesis that financial markets where the interests of long term investors are protected results in a safer environment. Moreover, we observe the change in trend for the other exchanges. Specifically, we notice that the average number of mini-flash crashes stopped decreasing and leveled up at around 0.025.

Figure 2 around here

As a set of control variables in model (1) we use exogenous variables that lead to differential trends and that are not influenced by the treatment. The inclusion of control variables has positive and negative aspects, even when these additional variables do not lead to a violation of the DID assumptions. Among positive aspects, covariates could help to detect effect heterogeneity for Nasdaq listed and other securities separately although the average frequency of mini-flash crashes in both types of securities experiences the same trends for their potential outcomes. On the negative side, every additional variable makes the common support assumption more difficult to fulfill. We use time-constant covariates like the rank of market capitalization, turnover rank, volatility rank and price rank. To assure the validity of the parallel trend assumption a number of diagnostics and robustness checks are provided in Section 3.3.

3 Discussion of Results

In this section, we show that the introduction of the M-ELO which was designed to leverage the competition between low- and high-frequency traders also resulted in a more stable market environment in terms of extreme price movements. The average number of mini-flash crashes decreases by 4.7% after the introduction of the M-ELO as well as the average dollar volume that is traded during the crash periods. The overall volatility increases slightly, but this effect is coupled with the volatility increase prior to the crash which can to some degree serve as an indication of the crash. We also show that the liquidity provision increases after the M-ELO becomes available for market participants.

3.1 Mini-Flash Crashes

Table 3 reports estimates of the treatment effects of introducing the M-ELO on different features characterizing mini-flash crashes. Column 1 reports regression results for model (1), where the dependent variable is the average number of mini-flash crashes. In general we observe more crashes on Nasdaq than on other exchanges: the average number of crashes on Nasdaq is 5.8% higher than on other exchanges. After the treatment, however, we document a decline in the number of mini-flash crashes by 4.7%. Column 2 reports estimation results for the case when the time dummy is just an indicator of a treatment period. In this case, we also document a decrease in the average number of mini-flash crashes as a result of the treatment. The coefficient is greater in absolute value and suggests the average number of crashes decreased by 5.3%. This is mainly due to the fact that the general downward trend is not taken into account in this setting. In both cases, the estimates of the covariates are significant and in line with the economic intuition. Higher the market capitalization and turnover of the company leads to a smaller number of flash crashes since there more often will be limit orders of market participants willing to buy or sell the stock at an attractive price. In contrast, if a stock is generally more volatile the spreads are set wider by the market makers and the overall liquidity worsens which explains more frequent crashes.

Columns 3–5 report the estimates of the model (1) where the dependent variables are

respectively daily volatility of the price, price volatility prior to the outbreak of the crash and dollar volume traded during the crash. Both volatilities are measured in percentage terms of the price of the security. The estimates suggest that the overall volatility of Nasdaq listed securities slightly increased by 0.137% after the introduction of the new order type. The volatility of the price just before the crash raises due to the treatment as well by 0.112%. This can have two possible explanations: firstly, after the treatment took place the sample of stocks which are still experiencing mini-flash crashes may include the most volatile one, and, secondly, the rise in volatility due to the treatment might suggest an increased predictability power of stock price volatility for the mini-flash crash. The dollar volume of trades during the crash times reduced due to the treatment by approximately \$287 thousand. For the other parameters like the duration of a crash, the percentage change, number of trades and recovery ratio we do not find a significant impact of the introduction of the M-ELO.

Table 3 around here

3.2 Liquidity Provision

Table 4 reports the estimates of the treatment effect of the M-ELO introduction on such measures of liquidity provision and market making activity as cancel-to-trade, trade-to-order, odd-to-trade and the hidden rate. We document a positive and significant effect of the treatment on the liquidity provision on Nasdaq. The average cancel-to-trade ratio on Nasdaq after the introduction of M-ELO increased by 1.32, taking into account the fact that on NASDAQ it is on average 4.19 points lower than on other exchanges. This supports the hypothesis that M-ELO favors enhanced market making activity due to decreasing the degree of adverse selection. Similarly, Column 2 documents a decrease in the trade-to-order ratio on Nasdaq after the M-ELO became available. The ratio for Nasdaq due to the M-ELO implementation decreases by 0.003. The introduction of this protected order type brings more protection to the long term reduces the price impact of big trades and result in an improved liquidity provision. Since M-ELO reduces price and market impact of trades, market participants with long term horizons are not forced to split their orders in order to receive better execution price.

In contrast, odd-to-trade ratio (see Column 3) increases due to the treatment by about 0.016 points. Early research [[Ritter, 1988](#), [Lakonishok and Maberly, 1990](#)] suggests that this ratio correlates with the fraction of retail investors mostly due to the budget constraints associated with submitting a round lot. In the era of algorithmic trading, however, odd lots are used also by HFT firms as was shown by [O’Hara et al. \[2014\]](#) and [Upson and Johnson \[2017\]](#). Thus, the statement that the increase of odd-to-trade ratio suggests retail investors are more willing to participate in price discovery should be made with caution. Column 4 reports the estimation result for hidden rate being the dependent variable. The coefficient on the treatment effect is not significant in this case, suggesting that the M-ELO does not impact the overall willingness of market participants to disclose their intentions.

Table 4 around here

3.3 Robustness

The crucial assumption of the model (1) is the presence of parallel trends in treatment and control groups prior to the treatment. We suggest that violations of the parallel trend assumption are not the case in our study. First of all, the assumption may not hold due to self-selection. In our case, the self-selection is not possible since the treatment affected only NASDAQ trading stocks the majority of which are listed on NASDAQ as well. In a similar way, targeting of the treatment group is also ruled out. We also rule out long-term effects which can harm reliability. In the long run, many events can happen that potentially confound the effect of the treatment, for example, regulatory changes, seasonality in overall financial market trading activity and other factors. However, since we analyze a relatively short time period around the implementation date we substantially reduce the occurrence of long-term effects.

On the other hand, we select stocks for the analysis according to their trading activity to maintain a high level of liquidity in our sample. This resulted in slightly different samples for each day because stocks who exceeded the threshold of 3,900 trades on a particular day do not necessarily exceed it on some other days. Of course, some highly liquid stocks can be

found in the sample all the times while some stocks might be considered for one day only. There are only 10 stocks with full presence of 42 days in our sample and as much as 307 of those who just appeared once. This fact of compositional differences across time is a source of a potential problem in identification since the change in the sample may confound the DID estimate because the “effect” could be attributable to change in population.

To provide a robustness check for this issue we construct a panel with time and fixed effects. Panel data analysis is a relevant robustness check because observations of individual stock for several periods reduce the variance compared to repeated random selections. We can account for this by estimating the panel data version of the model (1). The “within” transformation will allow us to estimate the coefficient β_3 . We compare the estimation of the panel data model to our main results using both balanced and unbalanced panels. In the balanced panel we include only 152 stocks with a full history of 35 days including the date of the M-ELO introduction.

Table 5 reports the estimation results of panel data models in Columns 1–4. In Columns 1 and 3 we report the results for an unbalanced panel, while in Columns 2 and 4 the results for a balanced subset ($N = 152$, $T = 35$) are provided. The estimation results of unbalanced panel suggest that treatment effect is still significant at 5% level and the estimates are close to what we have obtained previously. For the strictly balanced panel, however, the treatment effect is no longer significant. This is mainly due to the considerably strict sub-setting of the sample: out of 5320 observations we have only 82 crash events out of which only 17 hypothetically related to the treatment effect. The dimensionality of the data is crucial and we suppose to find significant effects encompassing more data.

Table 5 around here

To further strengthen our results in Table 5 we also report two types of placebo tests: placebo test using previous periods and alternative groups. For the first type of test, we exclude around one-third of the latest available data. Thus, we assume our data ends on March 25 and also assume the treatment happens at a different day: in our case, the middle point of the data is now February 26. Then we re-estimate model (1). The results are presented in Column 5. In this case, the coefficient of the treatment effect does not statistically differ

from zero which supports the validity of our initial problem setup. Next, Column 6 shows results for the placebo test using alternative groups. We re-coded some of the control groups (stocks which are not listed on Nasdaq) as treated also maintaining the initial ratio of the treated group of about 45% of the sample. Then we re-estimated the model (1) with the placebo-treated units and without actually treated units. The results in Column 6 again supports our initial setting: the treatment effect for alternative groups is also not significant at any reasonable level. This makes us confident about the statistical power of the main results reported in Table 3.

4 Conclusion

We provide a novel evidence on market stability and liquidity provision due to the implementation of technology-based solutions to reduce HFT negative externalities. Market participants require regulators to review the Regulation NMS and consider the possibility of eliminating Rule 611 which results in excessive market fragmentation. While there is little progress in this direction financial sector creates alternative solutions in the form of speed bumps to level the playing field and to give slower non-HFT firms a chance to compete with the HFT firms that otherwise would outrun them.

The introduction of the M-ELO provides a flexible solution to deemphasize HFT firms speed advantages and at the same time leaves a possibility to manage the risk of open positions. Figure 3 documents that the number of M-ELO orders is growing since its inception. As this type of protecting order also leads to a safer market environment in terms of mini-flash crashes occurrences the new Nasdaq's structure appears to work as a "slow market". Beyond the scope of this paper, more research can help generalize or qualify the findings. For example, it will be relevant to know whether and how the HFT evolved in response to the speed bumps. Also, it is important to understand how the introduction of new venues with different rules and speed deemphasizing solutions impacts the liquidity provision by the HFT firms.

References

- Y. Aït-Sahalia, J. Jacod, et al. Testing for jumps in a discretely observed process. *Annals of Statistics*, 37(1):184–222, 2009.
- T. G. Andersen, T. Bollerslev, F. X. Diebold, and H. Ebens. The distribution of realized stock return volatility. *Journal of Financial Economics*, 61(1):43–76, 2001.
- M. J. Barclay and T. Hendershott. Price discovery and trading after hours. *Review of Financial Studies*, 16(4):1041–1073, 2003.
- M. Bellia, K. Christensen, A. Kolokolov, L. Pelizzon, and R. Reno. High-frequency trading during flash crashes: Walk of fame or hall of shame. Technical report, 2018.
- B. Biais, T. Foucault, et al. HFT and market quality. *Bankers, Markets & Investors*, 128(1): 5–19, 2014.
- B. Biais, T. Foucault, and S. Moinas. Equilibrium fast trading. *Journal of Financial Economics*, 116(2):292–313, 2015.
- J. Brogaard, T. Hendershott, and R. Riordan. High-frequency trading and price discovery. *Review of Financial Studies*, 27(8):2267–2306, 2014.
- J. Brogaard, A. Carrion, T. Moyaert, R. Riordan, A. Shkilko, and K. Sokolov. High frequency trading and extreme price movements. *Journal of Financial Economics*, 128(2):253–265, 2018.
- M. K. Brunnermeier and L. H. Pedersen. Predatory trading. *Journal of Finance*, 60(4): 1825–1863, 2005.
- E. Budish, P. Cramton, and J. Shim. The high-frequency trading arms race: Frequent batch auctions as a market design response. *Quarterly Journal of Economics*, 130(4):1547–1621, 2015.
- A. Carrion. Very fast money: High-frequency trading on the NASDAQ. *Journal of Financial Markets*, 16(4):680–711, 2013.
- K. H. Chung and C. Chuwonganant. Uncertainty, market structure, and liquidity. *Journal of Financial Economics*, 113(3):476–499, 2014.
- J. Conrad, S. Wahal, and J. Xiang. High-frequency quoting, trading, and the efficiency of prices. *Journal of Financial Economics*, 116(2):271–291, 2015.
- D. Cumming, S. Johan, and D. Li. Exchange trading rules and stock market liquidity. *Journal of Financial Economics*, 99(3):651–671, 2011.
- D. Easley, M. M. López de Prado, and M. O’Hara. Flow toxicity and liquidity in a high-frequency world. *Review of Financial Studies*, 25(5):1457–1493, 2012.

- A. Golub, J. Keane, and S.-H. Poon. High frequency trading and mini-flash crashes. *Working paper*, 2012.
- J. Hasbrouck and G. Saar. Low-latency trading. *Journal of Financial Markets*, 16(4):646–679, 2013.
- E. Hu. Intentional access delays, market quality, and price discovery: Evidence from IEX becoming an exchange. *DERA Working Paper*, 2018.
- N. Johnson, G. Zhao, E. Hunsader, H. Qi, N. Johnson, J. Meng, and B. Tivnan. Abrupt rise of new machine ecology beyond human response time. *Scientific reports*, 3:2627, 2013.
- A. Kirilenko, A. S. Kyle, M. Samadi, and T. Tuzun. The Flash Crash: High-frequency trading in an electronic market. *Journal of Finance*, 72(3):967–998, 2017.
- J. Lakonishok and E. Maberly. The weekend effect: Trading patterns of individual and institutional investors. *Journal of Finance*, 45(1):231–243, 1990.
- S. J. Leal, M. Napoletano, A. Roventini, and G. Fagiolo. Rock around the clock: An agent-based model of low- and high-frequency trading. *Journal of Evolutionary Economics*, 26(1):49–76, 2016.
- A. J. Menkveld. High frequency trading and the new market makers. *Journal of Financial Markets*, 16(4):712–740, 2013.
- G. Meyer, N. Bullock, and J. Renninson. How high-frequency trading hit a speed bump. *Financial Times*, Jan 2018. URL <https://www.ft.com/content/d81f96ea-d43c-11e7-a303-9060cb1e5f44>.
- M. O’Hara, C. Yao, and M. Ye. What’s not there: Odd lots and market data. *Journal of Finance*, 69(5):2199–2236, 2014.
- E. S. Pagnotta and T. Philippon. Competing on speed. *Econometrica*, 86(3):1067–1115, 2018.
- J. R. Ritter. The buying and selling behavior of individual investors at the turn of the year. *Journal of Finance*, 43(3):701–717, 1988.
- J. Upson and H. Johnson. Are odd-lot orders informed? *Financial Review*, 52(1):37–67, 2017.
- M. Woodward. Bumping up the competition: The influence of IEX’s speed bump in US financial market. *Working paper*, 2018.

Tables and Figures

Table 1: Summary Statistics for Mini-Flash Crashes affecting individual stocks for the period February 12, 2018 – April 13, 2018.

The table reports summary statistics for the sample of mini-flash crashes for stocks with NASDAQ as a listing exchange (Panel A) and for all other stocks (Panel B). σ_i/σ_{i-1} shows how much greater was the price volatility during the crash interval compared to the volatility in the previous period, $|return_i|$ is the maximum absolute cumulative return during the crash, $duration$ is the time in seconds from the start of the mini-flash crash to the time point of maximum absolute cumulative return, $recovery$ shows how much does the price recover after the crash during time equal to the duration of the crash. σ_{day} represents daily price volatility of a security during the day of a mini-flash crash, while σ_{prior} reports price volatility for the period from the opening to the outbreak of the crash, $q(\cdot)$ reports quantiles.

	mean	std	min	max	q(0.25)	q(0.5)	q(0.75)
Panel A: Mini-Flash Crashes for NASDAQ listed stocks (604 events)							
σ_i/σ_{i-1}	17.77	15.80	10.01	243.55	11.32	13.49	18.69
$ return_i $, %	2.19	1.97	0.80	18.46	1.05	1.49	2.54
duration, s	86.71	61.39	0.01	274.51	37.12	71.38	124.38
recovery, %	56.42	20.14	33.33	233.33	41.87	52.38	65.03
Numb. trades	148.12	133.59	3.00	1,125.00	67.00	112.00	178.25
Volume \$'000	639.47	1,173.00	0.51	10,252.61	88.95	258.04	680.07
σ_{day} , %	3.88	4.08	0.15	31.53	1.51	2.47	5.16
σ_{prior} , %	3.13	3.32	0.12	24.36	1.09	1.99	4.12
Panel B: Mini-Flash Crashes for non-NASDAQ listed stocks (261 events)							
σ_i/σ_{i-1}	24.75	61.27	10.05	793.69	11.65	13.95	20.45
$ return_i $, %	1.65	1.28	0.80	11.52	0.93	1.20	1.78
duration, s	96.34	60.14	0.02	273.90	45.54	86.30	138.56
recovery, %	55.75	23.72	33.19	205.88	39.25	49.45	64.89
Numb. trades	235.79	228.17	29.00	1,800.00	91.00	162.00	302.00
Volume \$'000	1,368.61	2,581.30	16.50	25,978.07	234.11	588.02	1,562.25
σ_{day} , %	1.78	2.31	0.25	29.44	0.79	1.19	2.06
σ_{prior} , %	1.22	1.02	0.05	5.50	0.52	0.91	1.65

Table 2: Summary Statistics for Liquidity Provision.

The table reports summary statistics of liquidity provision proxies for the sample of most liquid securities on Nasdaq and on other exchanges (Panel A) and for the sample of stocks which experienced a mini-flash crash (Panel B); $q(\cdot)$ reports quantiles.

	mean	std	min	max	q(0.25)	q(0.5)	q(0.75)
Panel A: All securities							
NASDAQ listed securities (611 stocks)							
Cancel/Trade	9.95	7.85	0.62	123.57	4.83	8.23	12.61
Trade/Order	0.16	0.10	0.01	0.97	0.09	0.13	0.20
Odd/Trade	0.41	0.15	0.08	0.93	0.28	0.39	0.51
Hidden Rate	0.15	0.08	0.02	0.52	0.09	0.12	0.19
Market Cap. (\$'B)	64.70	149.20	0.00	922.05	3.89	12.43	49.91
non-NASDAQ listed securities (722 stocks)							
Cancel/Trade	18.32	13.46	1.99	155.89	9.48	15.13	22.70
Trade/Order	0.09	0.06	0.01	0.52	0.05	0.07	0.11
Odd/Trade	0.33	0.14	0.04	0.85	0.21	0.30	0.42
Hidden Rate	0.11	0.07	0.01	0.70	0.07	0.10	0.14
Market Cap. (\$'B)	55.68	77.56	0.00	412.10	8.92	22.19	59.66
Panel B: Securities with Mini-Flash Crashes							
NASDAQ listed securities (235 stocks)							
Cancel/Trade	5.13	5.27	0.62	61.54	2.53	3.90	5.97
Trade/Order	0.27	0.13	0.02	0.97	0.17	0.25	0.35
Odd/Trade	0.35	0.12	0.13	0.79	0.26	0.32	0.44
Hidden Rate	0.19	0.08	0.04	0.52	0.14	0.18	0.24
Market Cap. (\$'B)	4.21	13.17	0.01	159.86	0.18	0.98	2.67
non-NASDAQ listed securities (168 stocks)							
Cancel/Trade	14.00	12.56	2.44	89.37	5.35	10.21	16.81
Trade/Order	0.13	0.09	0.01	0.39	0.07	0.11	0.20
Odd/Trade	0.30	0.12	0.05	0.77	0.20	0.28	0.38
Hidden Rate	0.17	0.10	0.02	0.70	0.09	0.14	0.21
Market Cap. (\$'B)	13.67	28.25	0.03	273.41	1.76	4.39	13.67

Table 3: Mini-Flash Crash related features.

The table reports estimated coefficients from the following regression:

$$Y_{it} = \beta_0 + \beta_1 \lambda_t + \beta_2 Exchange_i + \beta_3 D_{it} + \beta_4 X_{it} + \epsilon_{it},$$

estimated on the sample of liquid stocks traded on NASDAQ for the period around introduction of the M-ELO from February 12 to April 13, 2018. Y_{it} is the average number of mini-flash crashes (Column 1 and 2), average volatility of the price during the day with a crash (Column 3), price volatility prior to the crash (Column 4) and dollar volume traded during the crash (Column 5). *Post Period* is a dummy variable indicating the post treatment period (when the M-ELO was launched), *Exchange* is a dummy variable which is equal to one if the stock is primarily listed on Nasdaq and zero otherwise, $D_{it} = Exchange_i \times d_t$ is the average treatment effect of the M-ELO introduction on NASDAQ. Regressions are estimated with stock fixed effects and time effects. Standard errors are in parentheses.

	<i>Dependent variable:</i>				
	Avg. Number of Crashes	Volatility, %	Volatility prior, %	Volume \$'M	
	(1)	(2)	(3)	(4)	(5)
Post Period (d_t)		0.003 (0.015)			
Exchange	0.058*** (0.015)	0.058*** (0.015)	0.092** (0.041)	0.117*** (0.034)	0.466*** (0.094)
Exchange $\times d_t$	-0.047** (0.022)	-0.053** (0.022)	0.137** (0.061)	0.112** (0.051)	-0.287** (0.140)
Market Cap Rank	-0.020*** (0.005)	-0.020*** (0.005)	-0.305*** (0.011)	-0.231*** (0.009)	-0.297*** (0.027)
Turnover Rank	-0.012*** (0.004)	-0.011*** (0.004)	0.195*** (0.007)	0.160*** (0.006)	-0.118*** (0.023)
Volatility Rank	0.086*** (0.004)	0.084*** (0.004)			-0.091*** (0.021)
Price Rank	0.001 (0.004)	0.0005 (0.004)	0.030*** (0.010)	0.011 (0.008)	0.206*** (0.023)
Days' dummies	<i>Yes</i>	<i>No</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Observations	14,565	14,565	5,915	5,915	5,915
R ²	0.095	0.090	0.304	0.291	0.085
F Statistic	32.431***	206.382***	59.571***	56.143***	12.390***

Note:

*p<0.1; **p<0.05; ***p<0.01

Table 4: Proxies for market making activity.

The table reports estimated coefficients from the following regression:

$$Y_{it} = \beta_0 + \beta_1 \lambda_t + \beta_2 Exchange_i + \beta_3 D_{it} + \beta_4 X_{it} + \epsilon_{it},$$

estimated on the sample of liquid stocks traded on NASDAQ for the period around introduction of the M-ELO from February 12 to April 13, 2018. Y_{it} is the cancel-to-trade ratio (Column 1), trade-to-order ratio (Column 2), odd-to-trade ratio (Column 3) and hidden rate (Column 4). $Exchange$ is a dummy variable which is equal to one if the stock is primarily listed on Nasdaq and zero otherwise, $D_{it} = Exchange_i \times d_t$ is the average treatment effect of the M-ELO introduction on NASDAQ. Regressions are estimated with stock fixed effects and time effects. Standard errors are in parentheses.

	<i>Dependent variable:</i>			
	Cancel/Trade	Trade/Order	Odd/Trade	Hidden Rate
	(1)	(2)	(3)	(4)
Exchange	-4.191*** (0.190)	0.025*** (0.001)	0.054*** (0.002)	0.012*** (0.001)
Exchange \times d_t	1.317*** (0.284)	-0.003* (0.002)	0.016*** (0.003)	-0.0003 (0.002)
Market Cap Rank	-0.628*** (0.059)	-0.015*** (0.0003)	-0.025*** (0.001)	-0.007*** (0.0004)
Turnover Rank	-1.398*** (0.047)	0.006*** (0.0003)	-0.0001 (0.001)	0.0003 (0.0003)
Volatility Rank	-1.061*** (0.047)	0.011*** (0.0003)	0.005*** (0.001)	0.010*** (0.0003)
Price Rank	-2.681*** (0.048)	0.022*** (0.0003)	0.066*** (0.001)	0.013*** (0.0003)
Constant	58.970*** (0.528)	-0.036*** (0.003)	0.071*** (0.006)	0.034*** (0.004)
Time dummies	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Observations	14,565	14,565	14,565	14,565
R ²	0.540	0.519	0.654	0.218
F Statistic	362.602***	333.759***	584.491***	86.246***

Note:

*p<0.1; **p<0.05; ***p<0.01

Table 5: Panel Data Model Estimation and Placebo Tests.

The table reports estimated coefficients from the following regression:

$$Y_{it} = \beta_0 + \beta_1 \lambda_t + \beta_2 Exchange_i + \beta_3 D_{it} + \beta_4 X_{it} + \epsilon_{it},$$

estimated on the sample of liquid stocks traded on NASDAQ for the period around introduction of the M-ELO from February 12 to April 13, 2018. Y_{it} is the mini-flash crashes. $Exchange$ is a dummy variable which is equal to one if the stock is primarily listed on Nasdaq and zero otherwise, $D_{it} = Exchange_i \times d_t$ is the average treatment effect of the M-ELO introduction on Nasdaq. Columns 1–4 report panel data estimation results, Column 5 reports the estimation results for the case in which one third of the sample was discarded and the date of M-ELO implementation shifted to February 26. Column 6 reports the estimation results for the case in which stocks were randomly reassigned to treatment and control groups. Regressions are estimated with stock fixed effects and time effects. Standard errors are in parentheses.

	<i>Dependent variable:</i>					
	Average Number of Crashes					
	Panel Data Model Estimation				Placebo Tests	
	(1)	(2)	(3)	(4)	(5)	(6)
Post Period (d_t)			−0.002 (0.015)	−0.022 (0.020)		
Exchange					0.038 (0.023)	0.001 (0.003)
Exchange $\times d_t$	−0.054** (0.022)	0.006 (0.029)	−0.059*** (0.022)	0.007 (0.029)	0.021 (0.028)	0.002 (0.005)
Market Cap Rank	−0.082** (0.035)	0.043 (0.108)	−0.083** (0.035)	0.042 (0.108)	−0.022*** (0.006)	−0.005*** (0.001)
Turnover Rank	−0.011 (0.007)	−0.010 (0.009)	−0.014** (0.007)	−0.010 (0.009)	−0.012*** (0.005)	−0.001 (0.001)
Volatility Rank	0.105*** (0.005)	0.102*** (0.009)	0.100*** (0.005)	0.097*** (0.009)	0.090*** (0.004)	0.011*** (0.001)
Price Rank	0.003 (0.023)	0.037 (0.035)	0.006 (0.023)	0.042 (0.035)	0.0001 (0.005)	0.003*** (0.001)
Time dummies	<i>Yes</i>	<i>Yes</i>	<i>No</i>	<i>No</i>	<i>Yes</i>	<i>Yes</i>
Balanced Panel	⊙	✓	⊙	✓		
Observations	14,565	5,320	14,565	5,320	11,403	14,565
R ²	0.038	0.038	0.032	0.029	0.097	0.039
F Statistic	11.427***	5.229***	73.064***	26.105***	35.903***	12.556***

Note:

*p<0.1; **p<0.05; ***p<0.01

Figure 1: Example of the mini-flash crash.

Panel A plots the Procter&Gamble (P&G) share price on March 21, 2018. At 2 p.m. the price experienced a mini-flash crash which was identified by our methodology. The pre-market price of P&G was at around \$78.2 and then dropped to the region \$77.4 – \$77.7 where it stayed fairly stable until 2 p.m. and then experienced a massive spike to the levels of approximately \$78.3. Within the next five minutes the price dropped more than 1.4% down to \$77.16 and eventually returned back to the region of previous daily consolidation. Panel B zooms around the time of the crash. Each dot represents a trade. The duration of the crash is 26.2 seconds, cumulative return of the first spike is 0.98%, change in terms of standard deviations is $46.5 \cdot \sigma$ times, volume traded during the crash is \$633 thousand.

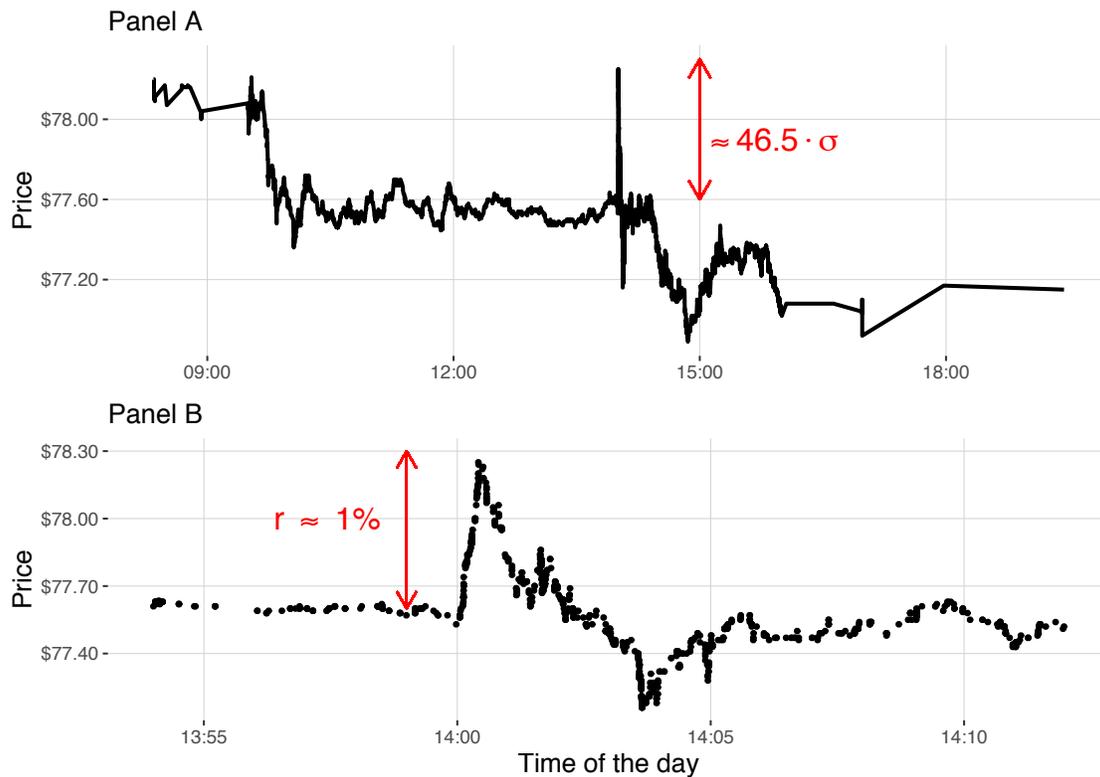


Figure 2: Trends in Number of Crashes.

Figure plots the average number of mini-flash crashes per a stock identified on Nasdaq and other national exchanges for the period from February 12, 2018 to April 13, 2018. The blue vertical line represents the day of the M-ELO introduction. Red and green dots represent the actual observations, while lines are estimations of trends. Shaded areas represent 95% confidence intervals.

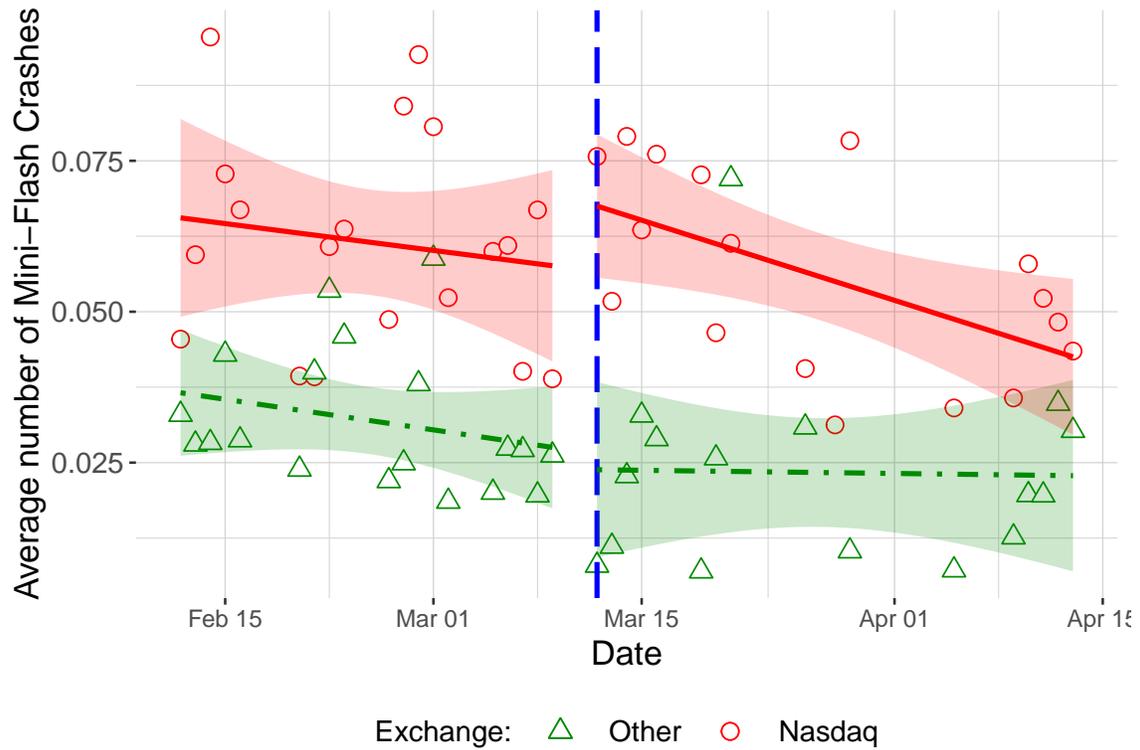
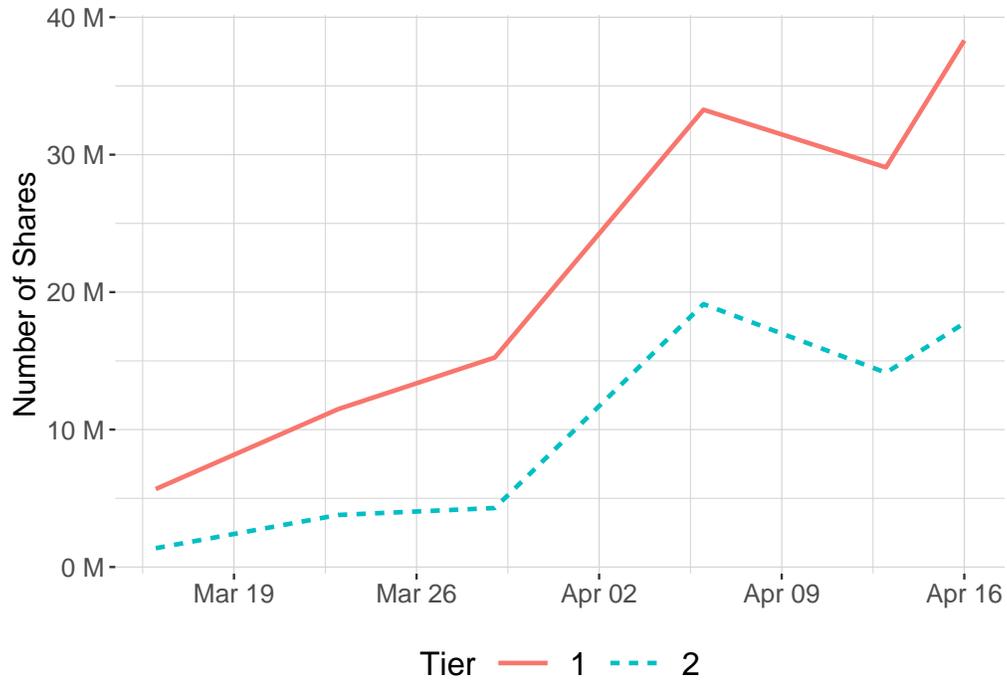


Figure 3: Volume of M-ELO orders.

Figure plots the number of shares traded in millions of M-ELO order type since its introduction on March 12, 2018. Tier 1 captures trades in stocks which belong to tier 1 and Tier 2 captures trades in stocks which belong to tier 2.



Appendix

A Data

Table 6: Example of the data tape.

The table reports an example of part of the data set. The data consists of order messages recorded by the exchange. Possible order message types: A (add an order without marker participant identifier (MPID)), C (execute an order with different price), D (delete an order), E (execute an order), F (add an order with MPID), P (trade message for hidden orders), U (replace an order), X (cancel an order). All messages are time-stamped by nanoseconds elapsed from the midnight. ORN stands for the order reference number, which makes possible to distinguish orders in a single security and to track their evolution (orn for message type P is always zero because the order was hidden). The side of the order might be either 'Buy' or 'Sell' (B/S) thus, making clear who was the initiator of the trade.

type	seconds	orn	side	shares	price	current.bid	current.ask
A	14400.01	13713	B	100	1.00	1.00	NA
D	14401.00	28705	B	2	1.00	1.00	NA
X	14432.36	287141	B	35	139.33	139.33	139.74
E	19922.60	515409	S	260	139.68	139.50	139.70
F	25200.25	2905093	B	100	0.01	139.50	139.63
P	26091.32	0	B	220	139.20	139.10	139.25
U	29423.20	4724289	S	100	138.95	138.91	138.95
C	34201.87	9851381	B	100	138.98	138.97	138.99

B Alternative Mini-Flash Crash Identification Technique

Here, we present the test which examines the presence of jumps in high-frequency price series which is based on the theory of [Aït-Sahalia et al. \[2009\]](#). The theoretical framework underlying the jump test is that the logarithmic price process X_t belongs to the class of Brownian semimartingales, which can be written as:

$$X_t = \int_0^t a_u du + \int_0^t \sigma_u dW_u + Z_t$$

where a is the drift term, σ denotes the spot volatility process, W is a standard Brownian motion and Z is a jump process defined by:

$$Z_t = \sum_{j=1}^{N_t} k_j$$

where k_j are nonzero random variables. The counting process can be either finite or infinite for finite or infinite activity jumps. Using the convergence properties of power variation and its dependence on the time scale on which it is measured, [Aït-Sahalia et al. \[2009\]](#) define a new variable which converges to 1 in the presence of jumps in the underlying return series, or to another deterministic and known number in the absence of jumps.

The test consists in comparing the multipower variation of equispaced returns computed at a fast time scale (h), $r_{t,i}$ ($i = 1, \dots, N$) and those computed at the slower time scale (kh), $y_{t,i}$ ($i = 1, \dots, N/k$). They find that the limit (for $N \rightarrow \infty$) of the realized power variation is invariant for different sampling scales and that their ratio is 1 in case of jumps and $k^{p/2} - 1$ if no jumps. Therefore the test detects the presence of jump using the ratio of realized power variation sampled from two scales. The null hypothesis is no jumps.

Assume there is N equispaced returns in period t . Let $r_{t,i}$ be a return (with $i = 1, \dots, N$) in period t . Also, there is N/k equispaced returns in period t . Let $y_{t,i}$ be a return (with $i = 1, \dots, N/k$) in period t . Then the test statistic is given by

$$AJ_{t,N} = \frac{S_t(p, k, h) - k^{p/2-1}}{\sqrt{V_{t,N}}}, \quad (2)$$

in which,

$$\begin{aligned} S_t(p, k, h) &= \frac{PV_{t,N}(p, kh)}{PV_{t,N}(p, h)} \\ PV_{t,N}(p, kh) &= \sum_{i=1}^{N/k} |y_{t,i}|^p \\ PV_{t,N}(p, h) &= \sum_{i=1}^N |r_{t,i}|^p \\ V_{t,N} &= \frac{N(p, k)A_{t,N(2p)}}{NA_{t,N(p)}} \\ N(p, k) &= \left(\frac{1}{\mu_p^2} \right) \left(k^{p-2}(1+k) \right) \mu_{2p} + k^{p-2}(k-1)\mu_p^2 - 2k^{p/2-1}\mu_{k,p} \\ A_{t,n(2p)} &= \frac{(1/N)^{(1-p/2)}}{\mu_p} \sum_{i=1}^N |r_{t,i}|^p \text{ for } |r_j| < \alpha(1/N)^w \\ \mu_{k,p} &= E \left(|U|^p |U + \sqrt{k-1}V|^p \right) \end{aligned}$$

where U and V are independent standard normal random variables; $h = 1/N$; p ; k , α , w : parameters. Below, we report a brief summary of results arising from this methodology. The results of model (1) estimation are in line with those reported in Table 3.

Table 7: Non parametric identification of Mini-Flash Crashes.

The table reports estimated coefficients from the following regression:

$$Y_{it} = \beta_0 + \beta_1 \lambda_t + \beta_2 Exchange_i + \beta_3 D_{it} + \beta_4 X_{it} + \epsilon_{it},$$

where Y_{it} is the average number of mini-flash crashes, identified by the test of Ait-Sahalia et. al. (2009). *Post Period* is a dummy variable indicating the post treatment period, *Exchange* is a dummy variable which is equal to one if the stock is primarily listed on Nasdaq and zero otherwise, $D_{it} = Exchange_i \times d_t$ is the average treatment effect of the M-ELO introduction on Nasdaq. Regressions are estimated with stock fixed effects and time effects. Standard errors are in parentheses.

	<i>Dependent variable:</i>	
	Avg. Number of Crashes	
	(1)	(2)
Post Period (d_t)		0.008 (0.017)
Exchange	0.049*** (0.017)	0.050*** (0.017)
Exchange $\times d_t$	-0.062** (0.025)	-0.068*** (0.025)
Market Cap Rank	-0.020*** (0.005)	-0.020*** (0.005)
Turnover Rank	-0.012*** (0.004)	-0.011*** (0.004)
Volatility Rank	0.082*** (0.004)	0.080*** (0.004)
Price Rank	0.003 (0.004)	0.003 (0.004)
Days' dummies	<i>Yes</i>	<i>No</i>
Observations	14,565	14,565
R ²	0.068	0.064
F Statistic	22.680***	141.625***

Note: *p<0.1; **p<0.05; ***p<0.01

Recent Issues

All CFS Working Papers are available at www.ifk-cfs.de.

No.	Authors	Title
635	Mathieu Aubry, Roman Kräussl, Gustavo Manso, and Christophe Spaenjers	<i>Machine Learning, Human Experts, and the Valuation of Real Assets</i>
634	Pascal Kieren and Martin Weber	<i>When saving is not enough – The wealth decumulation decision in retirement</i>
633	Volker Brühl	<i>LIBRA – a differentiated view on Facebook’s virtual currency project</i>
632	Yi Huang, Marco Pagano, and Ugo Panizza	<i>Local Crowding Out in China</i>
631	Annamaria Lusardi, Olivia S. Mitchell, and Noemi Oggero	<i>Debt Close to Retirement and Its Implications for Retirement Well-being</i>
630	Joelle H. Fong, Benedict SK. Koh, Olivia S. Mitchell, and Susann Rohwedder	<i>Financial Literacy and Suboptimal Financial Decisions at Older Ages</i>
629	Raimond Maurer, Olivia S. Mitchell, Ralph Rogalla, and Tatjana Schimetschek	<i>Optimal Social Security Claiming Behavior under Lump Sum Incentives: Theory and Evidence</i>
628	Refet S. Gürkaynak, Hatice Gökçe Karasoy-Can, and Sang Seok Lee	<i>Stock Market’s Assessment of Monetary Policy Transmission: The Cash Flow Effect</i>