CENTER FOR FINANCIAL STUDIES

CFS Working Paper Series

No. 691

Jorge Gonçalves, Roman Kräussl, and Vladimir Levin

Dark Trading and Financial Markets Stability

Center for Financial Studies | Goethe University | Theodor-W.-Adorno-Platz 3 | 60323 Frankfurt am Main | Germany | Phone +49 69 798 30050 | info-cfs@hof.uni-frankfurt.de | www.gfk-cfs.de

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.

Center for Financial Studies | Goethe University | Theodor-W.-Adorno-Platz 3 | 60323 Frankfurt am Main | Germany | Phone +49 69 798 30050 | info-cfs@hof.uni-frankfurt.de | www.gfk-cfs.de

Dark Trading and Financial Markets Stability☆,☆☆

Jorge Gonçalves[a,1], Roman Kräussl[b,c,d,1], Vladimir Levin[b,1,*]

[a] *University of Luxembourg, 6, avenue du Swing, Belvaux, L-4367, Luxembourg*
[b] *Department of Finance, University of Luxembourg, 6, rue Richard Coudenhove-Kalergi, L-1359, Luxembourg*
[c] *Hoover Institution, Stanford University, 434 Galvez Mall, Stanford, CA 94305-6003, USA*
[d] *Goizueta Business School, Emory University, 1300 Clifton Road, Atlanta, GA 30322, USA*

Abstract

This paper examines how the implementation of a new dark order – Midpoint Extended Life Order on NASDAQ – impacts financial markets stability in terms of occurrences of mini-flash crashes in individual securities. We use high-frequency order book data and apply panel regression analysis to estimate the effect of M-ELO trading on market stability and liquidity provision. The results suggest a predominance of a speed bump effect of M-ELO rather than a darkness effect. We find that the introduction of M-ELO increases market stability by reducing the average number of mini-flash crashes, but its impact on market quality is mixed.

Keywords: Market microstructure, financial market stability, mini-flash crash, dark trading, speed bump, investor protection
JEL: G10, G14

☆This work was supported by the Luxembourg National Research Fund.

☆☆We are grateful to NASDAQ Global Data Products for providing the data. We want to thank Yacine Aït-Sahalia, Tarun Chordia, Hans Degryse, Thierry Foucault, Anthony Haynes, Albert Menkveld, Per Mykland, Charles Nathanson, Roberto Steri, Jos van Bommel, Ulf von Lilienfeld-Toal, and Michael Weber for their valuable comments. We also benefited from the comments of all participants in the 4th International Workshop on "Financial Markets and Nonlinear Dynamics" in Paris, and the 2019 Summer School on Market Microstructure in Lugano. All errors are our own.

*Corresponding author.

Email addresses: `jorge.goncalves@uni.lu` (Jorge Gonçalves), `roman.kraussl@uni.lu` (Roman Kräussl), `vladimir.levin@uni.lu` (Vladimir Levin)

[1]Declaration of interest: none.

1. Introduction

For a couple of decades, market participants have been spending massive resources to obtain quick access to the richest data and invested in technologies to execute trades as fast as possible. According to Easley et al. (2012), for the period from 2009 to 2012, the share of high-frequency trading (HFT) firms have risen to more than 70% in the U.S. equity markets and approached 50% of the volume in futures markets. Since then, the percentage of HFT firms had increased further. Generally, HFTs use computer algorithms to look at patterns of prices, volumes, and past trading activity and react to any changes in those patterns at a matter of micro- or even nanoseconds. Some of them would not consider themselves as investing in fundamental information, but rather acquiring information about market dynamics and liquidity. Because fast actions of one algorithm may trigger responses of many others, small mispricing can rapidly self-reinforce itself and occasionally cause flash-crashes in securities' prices.

Golub et al. (2012) study the increase in the number of mini-flash crashes in individual securities between 2006 and 2011 and suggest that HFT causes those crashes. 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. On the other hand, Kirilenko et al. (2017) examine the structure of the E-mini S&P 500 stock index futures market on May 6, 2010, and observe that trading patterns of HFT did not change when prices fell during the Flash Crash.

Biais and Foucault (2014) and Biais et al. (2015) propose to create a segment of slow-friendly markets but to leave room for investment in the fast trading technology. In the spirit of this recommendation, exchanges started to introduce technology-based solutions to protect the interests of long-term investors. Those solutions were implemented in the form of latency delays and are commonly known as “speed bumps”. The Investors Exchange (IEX) applied the first such measure by introducing a 350-microsecond delay to all incoming and outgoing correspondences. Hu (2018) observes improvements in market functioning around the period when IEX became a national securities exchange. Moreover, he documents a positive impact of such speed bumps on market quality in terms of tighter spreads and improved liquidity. Several other exchanges followed the example of IEX and applied for the introduction of such a delay to the Securities and Exchange Commission (SEC).

Allowing for dark trading may also help in protecting the interests of long-term investors as dark orders hide trading intentions. When trading large amounts of stock using visible market or limit orders, one cannot prevent the price to be moved. To reduce price impact, a trader can submit the order to the dark pool and often even receive a better execution price. The downside of going to the dark pool is execution uncertainty since there is no guarantee that a trader will find a counterparty.

Major exchanges nowadays run their own dark pools where the execution price is referenced by the current mid-price (the average of the best bid and the best ask prices). It is difficult, however, to hide dark orders from HFT firms. They use their speed advantage to submit hidden orders inside the spread and quickly cancel them if they do not execute right away. If HFTs identified hidden orders, they could easily manipulate the best quotes to transact with those orders at comfortable prices.

In 2018, the NASDAQ exchange came up with a solution to improve dark orders and shield them from HFTs. Introduced on March 12, 2018, the Midpoint Extended Life Order (M-ELO) is targeted toward long-term investors. Anonymity and confidentiality of M-ELO are the critical tools to prevent potentially predatory counterparties from determining intentions and using that information to generate short-term profits at the expense of slow traders. This order becomes executable 500 milliseconds after submission and does not interact with other NASDAQ dark orders that have not met the 500 milliseconds holding period requirement. On May 11, 2020, NASDAQ revisited the design of the M-ELO order and decided to reduce the holding period to 10 milliseconds. The exchange motivated this change by optimization in M-ELO opportunities and execution. The decrease in the holding period was expected to open up M-ELO to use cases that were previously unavailable[2].

Despite increasing volumes traded in dark pools, there is still limited theoretical work discussing the effect of hidden orders on market quality, stability, and price discovery. Boulatov and George (2013) build a model where the strategies of informed traders can be adjusted in response to the visible and hidden liquidity on the market. They analyze venue and order choices of traders and find that hidden liquidity has a beneficial impact on market quality due to increased competition among informed market participants.

In contrast to Boulatov and George (2013), Zhu (2014) argues that adding a dark pool alongside the exchange decreases its liquidity. His model suggests that, since informed orders are much more correlated than uninformed ones, informed traders would rather choose the lit venue to avoid low execution probabilities in the dark pool. This relatively high presence of informed trading on the lit

[2]More information on the rule change may be found at https://www.nasdaq.com/articles/the-midpoint-extended-life-order-m-elo%3A-m-elo-holding-period-2020-02-13

exchange naturally increases price discovery which, in turn, leads to reduced liquidity. Buti et al. (2017) numerically solve a discrete model in which traders decide to submit their order to either an exchange or a dark pool. The authors obtain a set of equilibrium order submission probabilities and show that the introduction of a dark pool alongside the exchange widens bid-ask spread and reduces the depth available around the midquote. Those negative changes in market quality are partially mitigated when the initial liquidity of the limit order book increases.

Overall, theoretical works propose mixed results about the effects of dark trading on market quality and price discovery. It is still a challenge for regulators to decide on the degree of control of dark trading. Moreover, the aspect of market stability received much less attention. However, it remains a relevant topic in the current times of algorithmic trading proliferation since fragile markets may undermine investors' trust in the financial market system.

This paper adds to the existing literature on dark trading, speed bumps and also expands on its relation to market stability.

Our paper aims to identify the degree to which M-ELO is used in securities trading and its impact on the number of mini-flash crashes during a continuous trading period. We also investigate the association between M-ELO trading and market liquidity. The analysis employs high-frequency trades and quotes data from NASDAQ to identify mini-flash crashes in individual securities and to relate crash occurrences to the intensity of dark M-ELO trading through panel regression analysis. Recently implemented change in the order design allows also to disentangle the darkness and the speed bump effects of the M-ELO. Our results suggest a strong relationship between M-ELO volumes and measures of market quality and stability. This relationship is more pronounced through the speed bump effect of M-ELO rather than through the dark order effect. Overall, M-ELO is associated with fewer flash crashes, greater visible volumes in the limit order book, but widened spreads.

The remainder of the paper is organized as follows. Section 2 describes the data sources of intra-day trading and M-ELO volumes. It also presents measures of liquidity and summary statistics. Section 3 presents the methodology, empirical results, and describes various robustness checks. Section 4 concludes.

2. Data and Descriptive Statistics

2.1. Data sources

The Order Book Message data come from NASDAQ historical ITCH. This data set contains time-stamped in nanoseconds order submissions, executions, cancellations, and modifications on the NASDAQ equity market. The data, however, do not identify market participants and their activity in the dark. Submissions of hidden orders of any kind are not reported, while the executions are visible for all order types.

The data allow us to directly observe liquidity provision on each depth level of the limit order book at any time. The sample period covers three years of trading from January 2, 2018, to December 31, 2020. We consider a set of 1,342 firms traded on NASDAQ, where the trading activity is considerably high. We, therefore, in line with similar analyses of Andersen et al. (2001) and Brogaard et al. (2018), preserve the sufficiently large number of observations by focusing our analysis on large and medium firms.

The limitations of this data are straightforward and similar to those that previous research encountered (see, e.g., 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, 33% of trading activity for NASDAQ listed stocks, about 12.5% for NYSE stocks, and 16% for ARCA stocks. Despite the high fragmentation nature of financial markets, we share the reasoning of Brogaard et al. (2018) that liquidity transfers to other venues are unlikely due to the short period of interest and overall similar liquidity provision rules among exchanges. Thus, we argue that although the results obtained could not readily expand to other exchanges, still they should be taken into account in the matters of market design.

The M-ELO order became available on March 12, 2018, and quickly gained its share of trading volume which averages to be around 1.06% of NASDAQ's total matched volume for the period from 2018 to 2020. Weekly volumes of M-ELO trades come from NASDAQ's M-ELO Transparency Statistics[3]. Only weekly or lower frequencies of M-ELO trades are available. This is one of the biggest limitations of the present analysis. By shifting to weekly observations one might lose statistical power in identifying the effect of M-ELO trading. We resort to the fact, however, that the size of the sample is sufficient and the panel structure of the data allows us to obtain robust estimates.

Securities' characteristics to serve as covariates in our analysis were both computed from trades data and obtained from the SEC's MIDAS Market Structure Metrics database. As an instrument to the potentially endogenous M-ELO trading, we use the market-wide level of M-ELO trading in

[3] Available at https://www.nasdaqtrader.com

securities of the same market capitalization group.

2.2. Mini-flash crash identification

For the identification of mini-flash crashes, we only consider periods of continuous trading from 9:30 a.m. to 4:00 p.m. Our identification approach is similar to Brogaard et al. (2018) for extreme price movements (EPM) identification. In their methodology, the trading day is split into 10-second intervals between 9:35 a.m. and 3:55 p.m. The intervals are then ranked by the midquote return magnitude and those with returns exceeding the 99.99^{th} percentile are identified as EPMs.

However, our methodology differs from Brogaard et al. (2018) in several aspects. Firstly, we account for cross-sectional heterogeneity in trading activities for different firms. This is done by switching from a static time interval of 10 seconds to an interval with variable length. The length of the interval depends directly on the number of trades in the security for a day. To make sure the turnover is considerably high, we keep only those securities for which the number of trades for the period from 09:30 a.m. to 4:00 p.m. is not less than 11,700. This way, we require each time interval for mini-flash crash identification to contain at least 30 trades on average. Throughout the period from 2018 to 2020, we identified 1,342 firms, who satisfy the minimum number of daily trades.

Secondly, as flash crashes are known for their subsequent reversals a simple calculation of midquote returns has a flaw of missing those intervals within which the price suddenly jumps and quickly retraces back. To overcome this problem, for each time interval we compute the maximum possible midquote return. In such a way we obtain a series of intra-day extreme returns for a security.

Next, for each firm, we identify the intervals containing mini-flash crashes as those where the

Z-score for the midquote extreme returns at day t exceeds the value 7. These are the returns that satisfy $r_{t,i} > \mu_t + 7\sigma_t$, where $r_{t,i}$ is the i^{th} extreme return on day t, μ_t is the average extreme return on day t, and σ_t is the standard deviation of extreme returns on day t for this particular security.

Lastly, we separate systematic flash crashes from idiosyncratic ones. Systematic crashes are not limited to a particular security or exchange that is why it is unrealistic to expect any impact of M-ELO activity on a single stock on a single exchange. For 221,005 identified mini-flash crashes, for each stock, we check how many other stocks experienced a crash within a two-minute window around the initial crash. We mark flash crashes that have 25 or more instances in other stocks as systematic ones and remove them from the analysis.

The procedure gets a total of 30,942 idiosyncratic mini-flash crashes with 54.62% of them being negative. This identification technique is in line with previous works of Golub et al. (2012) and Johnson et al. (2013) that do not make a sharp distinction between crashes and spikes and require the price to move fast and severely. Similar to the approach of Bellia et al. (2020), we study only those mini-flash crashes which possess transitory dynamics. The average price reversal in the next 10 minutes following the end of the crash is 91.8%.

The first panel of Figure 1 shows the distribution of mini-flash crashes throughout the sample period. There is no much visual evidence of M-ELO introduction having a significant effect on the frequency of mini-flash crashes. We associate this fact with relatively low trading volumes of M-ELO orders. There is, however, a decrease in the number of crashes not long after the introduction of M-ELO that coincides with the spike in the M-ELO trading at this period. As we show in this work, there is a statistically significant association between the degree of M-ELO trading and the expected number of mini-flash crashes observed during a week.

The last two panels of Figure 1 show the time trends for the median values of the quoted spread and the depth available 30 basis points around the midquote relative to the daily trading volume. The dynamics of the depth suggests market quality is improving during times of active M-ELO trading in the first half of 2018. In 2019, the relative depth is at its highest levels, sometimes exceeding 2% of the daily trading volume, while the spread stays moderate around 3 basis points. In early 2020 we can clearly see the impact of the COVID-19 crisis. Spreads more than double while fewer shares are available close to the midquote.

Figure 1 around here

An example of a mini-flash crash, identified by our approach is presented in Figure 2. The crash occurred in the price of Procter & Gamble (PG) on March 21, 2018. Panel A spans the trading during the 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 about 1%. In the next five minutes, however, the price dropped more than 1.4% and eventually returned to the region of previous daily consolidation. The price became more volatile during the rest of the day.

Figure 2 around here

Panel B shows a zoomed representation of the crash event, where each dot represents a single trade. We note that the crash did not trigger the circuit breaker, even though the return associated with the initial spike had a Z-score of 7. The duration of this mini-flash crash in the price of PG was 26.2

seconds (the initial spike), with a cumulative return of 0.98%, and a trading volume that exceeded $633 thousand.

2.3. Midpoint Extended Life Order

Recent research on dark trading looks at the economics of liquidity provision in the dark. Academics start to distinguish two types of dark trading. The first, "one-sided" type, reflects dark trading at a single price which, in most cases, is the midquote. It is usually not regarded as a beneficial type of dark trading due to its low execution probability, absence of profitable market strategies of earning the spread, and limited abilities to hide trading intentions because of susceptibility to probing orders. The other, "two-sided" dark trading, allows dark buy and sell limit orders to exist simultaneously. This type of dark trading is believed to be benign to price discovery and market quality.

Foley and Putniņš (2016) and Comerton-Forde et al. (2018) analyze the effect of both types of dark trading by exploiting natural experiments in Canadian and Australian markets. They find that two-sided dark trading reduces quoted, effective, and realized spreads as well as market illiquidity measured by the price impact. On the other hand, they find no evidence that one-sided dark trading affects markets. The dark order studied in our paper can be mainly referred to as one-sided dark trading. However, it possesses some features of the two-sided type as well.

On March 12, 2018, NASDAQ launched a new order type: Midpoint Extended Life Order (M-ELO), which is designed to attract long-term investors to trade with each other at the midpoint of the National Best Bid and Offer (NBBO). M-ELO is a hidden order which interacts only with other M-ELO-type orders[4]. Because of this, M-ELO orders stay out of the way of the book clearing orders, with the aim to reduce information leakage and to provide a better execution price.

From the start, there was a 500 millisecond period called the “Holding Period” before an M-ELO order can be executed. This restriction protects market participants from the negative price impact as well as from adverse selection. If a bid or an ask price moves, M-ELO orders are automatically tagged to the new midquotes but the 500 milliseconds timer does not restart. If that was not the case, then one would expect to see fewer M-ELO orders executed during volatile markets, since prices will move a lot and sometimes exhibit mini-flash crashes. As the waiting locked-in times cannot be easily extended by the NBBO moves, the documented negative association between the number of matched M-ELOs and the number of mini-flash crashes may not be simply explained by built-in “protective” features of M-ELOs.

Starting from May 11, 2020, the NASDAQ exchange reduced the holding period of M-ELO orders from 500 to 10 milliseconds. Following the satisfactory M-ELO performance, the exchange decided to increase the opportunity set for its clients. With a 98% reduction of the waiting period, more trading strategies will be able to incorporate the benefits of the M-ELO. For the current research, this rule change is crucial as it helps to disentangle the dark trading effects and the speed bump effects of M-ELO orders.

Figure 3 shows that median M-ELO trading spikes to about 3% of the total matched by NASDAQ volume at the end of April 2018. In the second half of 2018, it starts to decline which can be associated with the change in the submission fees for this order type. M-ELO was fees-free until

[4] Specifications and more details of M-ELO order type can be found in the Appendix.

May 2018, with the possibility of a one-month extension, given that some trading activity-based milestones were reached. Since June 2018, M-ELO order submission incurs a fee for all stocks. This might be the reason for the gradual decline in relative M-ELO volumes already in the early autumn of 2018.

In 2019, relative M-ELO trading reached near maximum levels. On average M-ELO executions constituted 1.3% of the total amount matched by NASDAQ. The share of M-ELO executions reduced significantly during the first months of 2020. We observe an increase in relative M-ELO volumes right at the time of the design change. As the holding period became less restrictive, more market participants opted for this order. In our analysis, we control for this change by introducing a dummy variable that equals one if the current holding period is 10 milliseconds and zero otherwise. Also, we include time effects into the regression to rule out the end of the year effects as well as the COVID-19 related financial crisis.

Figure 3 around here

Another remarkable feature of M-ELO is that the sizes of executed orders are usually bigger than the sizes of visible limit orders, submitted to the book. Figure 4 plots density curves of sizes of visible limit orders and M-ELO orders. It can be seen that the distribution of sizes for M-ELO dominates the distribution of sizes of visible orders. The fact that M-ELO is associated with bigger order sizes can be a potential channel through which M-ELO activity may impact market stability. HFT would benefit from taking the opposite side relative to market participants who submit large orders. This contrarian trading by HFT may create unreasonable price pressure that can lead to crashes. It is possible that as more market participants opt for non-displayed M-ELO orders, fewer flash crashes occur in the market.

Figure 4 around here

2.4. Measures of liquidity and order imbalance

Liquidity is generally understood as the ability to quickly trade considerable volumes at a low cost. It is a multi-dimensional concept that includes trading costs, depth available to customers placing large orders, speed of execution, and protection against execution risks (Foucault et al. (2013)).

To measure market liquidity, we use the Order Book Message data from NASDAQ historical ITCH. It is the most granular type of data as it records every message sent to the exchange. High-frequency quote updates and trades allow for high-frequency estimation of the market liquidity metrics. Later, these metrics are aggregated to weekly frequencies by time-weighting.

Our first empirical measure of liquidity is the quoted half spread. It represents a scaled by the midpoint price difference between the lowest ask price (a_t) and the highest bid price (b_t) available at the moment:

$$QS_t = \frac{a_t - b_t}{2m_t} = \frac{a_t - b_t}{a_t + b_t}, \tag{1}$$

where $m_t = (a_t + b_t)/2$ is the midprice. The quoted spread for stock-day is time-weighted and based on the local limit order book.

The other measures are the 5-minute realized spread and price impact. They are calculated per trade and then averaged over the trading day. The 5-minute proportional realized spread for the t^{th}

transaction is defined as

$$RS_t = g_t \frac{p_t - m_{t+5min}}{m_t}, \tag{2}$$

where p_t is the trade price, g_t is the buy-sell indicator that equals $+1$ if the trade is a buy and -1 if the trade is a sell, and m_{t+5min} is a quote midpoint 5 minutes after the t^{th} trade.

A price impact measure is based on the extent to which a trade generates an adverse reaction in the market price. The midprice tends to rise when buy orders arrive, to an extent that is positively correlated with their size. Symmetrically, it tends to fall in the wake of sell orders. The 5-minute price impact of a t^{th} trade is defined as follows:

$$PI_t = g_t \frac{m_{t+5min} - m_t}{m_t}. \tag{3}$$

Our data set allows us to construct various depth measures and incorporate the limit orders beyond the best price levels. Similar to the work of Degryse et al. (2015), we measure the aggregate monetary value of shares offered within a fixed interval around the midpoint. We keep the original notation and refer to this measure as $Depth(X)$. Denote the price level $j = \{1, 2, \ldots, J\}$ on the pricing grid and the midpoint of the local limit order book as m, then

$$\text{Depth Ask}(X) = \sum_{j=1}^{J} p_j^a \cdot q_j^a \cdot \mathbb{I}\{p_j^a < m(1+X)\}, \tag{4}$$

$$\text{Depth Bid}(X) = \sum_{j=1}^{J} p_j^b \cdot q_j^b \cdot \mathbb{I}\{p_j^b > m(1-X)\}, \tag{5}$$

$$\text{Depth}(X) = \text{Depth Bid}(X) + \text{Depth Ask}(X), \tag{6}$$

where p_j^a (p_j^b) is the price of the limit sell (buy) order at price level j and q_j^a (q_j^b) is the number of shares available at this level. We use the indicator function $\mathbb{I}\{\}$ to determine if a limit order of a certain price is within the required interval around the midquote. The depth measure is expressed in U.S. dollars and calculated for $X = 30$ basis points.

Further, we calculate an imbalance measure similar to that of Belter (2007). This measure allows us to compare the liquidity supplied to different sides of the book beyond the best price levels. Having the interval of X basis points around the midquote and the price levels $j = \{1, 2, \ldots, J\}$, the depth imbalance is defined as follows:

$$DI(X) = \frac{\sum_{j=1}^{J} \left(\frac{q_j^a}{(p_j^a - m)} \cdot \mathbb{I}\{p_j^a < m(1+X)\} - \frac{q_j^b}{(m - p_j^b)} \cdot \mathbb{I}\{p_j^b > m(1-X)\} \right)}{\sum_{j=1}^{J} (q_j^a + q_j^b) \cdot \mathbb{I}\{1 - X < p_j/m < 1 + X\}}. \quad (7)$$

For each stock-day, we compute the average depth imbalance for $X = 30$ basis points. This measure is scaled by the total number of shares available in the given interval X.

2.5. Measure of algorithmic trading

It is hard to discriminate between orders placed by humans and orders placed by computer algorithms. A methodology to identify HFT depends heavily on the availability of data. If the data include information on HFT firms, it can be used directly to account for algorithmic trading (AT) activity. Most financial markets, however, do not provide information on whether an order comes from a human or an algorithm.

In case that HFT cannot be classified exactly, researchers use various proxies to quantify levels of HFT. Those proxies are constructed from trade and order submission data. We use the empirical measure developed in Hasbrouck and Saar (2013) as our major proxy for AT. This measure calculates the intensity of "strategic runs", which are series of linked messages. The linking results from HFT dynamically submitting and canceling orders to incorporate the latest information into prices.

Following their methodology, we connect a newly submitted limit order to a previously deleted order if the time between the two events does not exceed 100 milliseconds. The newly submitted order should have the same direction and size in shares as the previously deleted one. Only sufficiently long runs of 10 and more linked orders are kept. We scale the sum of durations of all runs, which are allowed to overlap, by the duration of the trading day. Our proxy for the AT activity on day t is defined as follows:

$$AT_t = \frac{1}{6.5 \cdot 3{,}600} \sum\nolimits_{j=1}^{N} T_{jt}, \tag{8}$$

where $6.5 \cdot 3{,}600$ is the total time in seconds from 9:30 a.m. to 4:00 p.m., N is the number of strategic runs on day t, and T_{jt} is the duration in seconds of run j on day t.

2.6. Summary statistics

Table 1 presents descriptive statistics for the sample stocks. There is considerable variation in average daily price and daily dollar volume traded. The average coefficient of variation of stock extreme returns (σ_r/μ_r) is centered around 0.84 with a relatively small standard deviation. This suggests the firms fall in pretty much the same volatility cohort. The NASDAQ's share shows what fraction of the consolidated volume in a particular security was matched by the NASDAQ exchange. The mean and the median share of NASDAQ across all listing venues are around 21%.

Table 1 around here

Table 1 also reports descriptive statistics on liquidity measures and measures of AT. The average daily quoted half spread throughout the sample period is about 9.2 basis points. Five-minute realized spread and price impact are centered around zero with the first (the third) quartile being about minus (plus) one basis point. There is a large variation in the dollar value of shares available 30 basis points around the midquote. Its mean of \$2.29 million is higher than the value of its 3$^{\text{rd}}$ quartile. The majority of depth available is sell-side volume, as the positive mean value of depth imbalance suggests.

The proxy of AT based on strategic runs suggests that somebody is engaging in dynamic order submission about 1.2% of the time during the period from 9:30 a.m. to 4:00 p.m. The messages-to-trade ratio clearly illustrates the fact that quoting activity nowadays is superior to trading activity. Its mean shows that there are about 36 times more quote update messages than actual trades.

Table 2 shows summary statistics for 30,942 identified mini-flash crashes. The Z-score of the return is the value $\frac{r_i - \mu_r}{\sigma_r}$, where r_i is the return during the mini-flash crash, μ_r and σ_r are the mean and standard deviation of maximum interval returns during the day. The threshold for the Z-score was chosen to be equal to 7, to identify extreme price movements. The table reports the average Z-score for the extreme returns around 8.5.

The average duration of a crash is approximately 31 seconds. In fact, 95% of identified mini-flash crashes do not last longer than 78 seconds. The mini-flash crash returns are distributed around -12.622 basis points, which is explained by the fact that only 54.62% of crashes are negative. The median absolute mini-flash crash return is about 56 basis points.

Table 2 around here

As Table 2 reports, mini-flash crashes are remarkable in subsequent price reversals. The 10-minute price reversal after the crash is on average 91.8%. The number of trades during the crash period is substantially higher than in normal times. With the mean number of trades of 185 and the mean crash duration, one obtains that the average number of trades per second to be 6 trades. This is a more than 7 times higher trading intensity compared to the average across the whole sample. The number of shares traded and the dollar volume during a crash tell a similar story. Also, the size of the trade increases slightly during crash times. The average size of the trade during normal times is around 127.3 shares, while it increases to 170.9 at the periods where we identify mini-flash crashes.

Table 2 also reports summary statistics on the crash aftermath volatility. This is the standard deviation of the extreme returns for the next 30 minutes following the end of the mini-flash crash. It is scaled by the standard deviation of the extreme returns throughout the day. The average volatility after the crash is approximately 22.7% higher than the volatility for that particular day.

Finally, descriptive statistics on the relative amount of M-ELO trading are provided. The average share of all M-ELO orders relative to all matched by NASDAQ orders is 1.18% for the stocks with identified mini-flash crashes. Also, the size of M-ELO orders is usually larger than the average size of visible orders on NASDAQ. These results suggest that M-ELO trading is more active in less liquid stocks when the market participants wish to trade large orders. This might seem unrealistic that such a small order type can impact any aspect of markets' behavior. However, in the next section, we show a steady coupling between the share of M-ELO trading and mini-flash crashes' intensity and liquidity provision quality.

3. Empirical Approach and Results

This section analyzes the effect of M-ELO trading on market stability which is approximated by the number of mini-flash crashes in individual securities. We also study the effect on various crash characteristics and liquidity measures. Our empirical approach involves relating market stability and liquidity characteristics to the M-ELO trading via stock-week panel regressions. For the panel regressions, we take two methods: (i) two-stage least squares (2SLS) instrumental variable regressions, and (ii) two-stage GMM estimations, which are efficient in the presence of heteroskedasticity of unknown type and apply heteroskedasticity and autocorrelation robust standard errors. As a robustness test, we apply difference-in-differences approach for the periods around the policy changes.

3.1. Mini-flash crashes

In this section, we identify the impact of M-ELO trading on the general number of mini-flash crashes and their characteristics. To account properly for both the cross-section variation and time variation, we employ the panel structure of the data and estimate the following panel regression with time and fixed effects:

$$y_{it} = \alpha_t + \beta_1 \cdot \text{M-ELO}_{it} + \beta_2 \cdot d_t \cdot \text{M-ELO}_{it} + \theta \cdot X_{it} + C_i + u_{it}, \tag{9}$$

where y_{it} is one of the following: (i) weekly number of mini-flash crashes in security i on week t, or (ii) one of the various crash characteristics like the maximum absolute return during the crash, the duration of the crash, subsequent price reversal, and others. The variable M-ELO_{it} is a fraction of M-ELO shares matched by NASDAQ to the total number of shares the exchange matched for security i on week t, d_t is the indicator that the holding period of M-ELO orders has been reduced from 500 milliseconds to 10 milliseconds. It equals one for all the days after May 11, 2020, and zero otherwise. X_{it} is a set of control variables that includes market capitalization and turnover ranks from the MIDAS database, messages-to-trades ratio, and the proxy for the AT activity. The set of controls also includes our proxy for the institutional trading activity that is based on the market participant identifier flag in the Order Add message.

Identifying the causal effect of dark M-ELO trading is generally problematic due to endogeneity. The possibility of reverse causality arises because M-ELO activity may affect market stability, but, at the same time, less stable markets may push participants to the dark. Econometrically, this means that endogenous regressors will make the estimates biased, inefficient, and inconsistent.

One potential solution to the endogeneity problem is the instrumental variable approach. A good instrument should be correlated with the potentially endogenous variable and should not be correlated with the model error. In the spirit of Hasbrouck and Saar (2013), Degryse et al. (2015), and Comerton-Forde and Putniņš (2015), we instrument M-ELO trading in stock i on week t with $\text{M-ELO}_{it}^{\text{other}}$ which is the average level of M-ELO trading across all stock in our sample in the same market capitalization rank excluding: (i) stock i itself, (ii) stocks from the same sector as stock i, (iii) stocks in the same index as stock i. If M-ELO activity has a significant market-wide component, then a market-wide average is likely to satisfy the first requirement. By excluding stocks from the same index and the same sector, we eliminate the potential cross-trading strategies candidates and maintain the economic intuition for using it as an instrument, in accordance with the second requirement. The formal tests of instruments' relevance and exogeneity can be found in Table A.1 of the Appendix.

Thus, we extend the linear regression model (9) by the following first-stage regression:

$$\text{M-ELO}_{it} = a_t + \pi_1 \text{M-ELO}_{it}^{\text{other}} + \gamma X_{it} + C_i' + v_{it}, \tag{10}$$

where a_t and C_i' are time and fixed effects, respectively, X_{it} is the set of control variables included to weaken the instrument exogeneity assumption. The set of control variables also includes such strong determinants of M-ELO activity as the dummy for the introduction of M-ELO and the holding period reduction dummy variable.

In the equation (9), both M-ELO_{it} and $d_t \cdot \text{M-ELO}_{it}$ are potentially endogenous. To get consistent estimates, we use $d_t \cdot \text{M-ELO}_{it}^{\text{other}}$ as an instrument for $d_t \cdot \text{M-ELO}_{it}$. The same economic argument used to support $\text{M-ELO}_{it}^{\text{other}}$ as an instrument for M-ELO_{it} will carry through to the interaction term. Therefore, we estimate an additional regression in first-stage:

$$d_t \cdot \text{M-ELO}_{it} = a_t + \pi_1 \cdot d_t \cdot \text{M-ELO}_{it}^{\text{other}} + \gamma X_{it} + C_i' + v_{it}. \tag{11}$$

In the second-stage regression, we can then estimate the causal effect of M-ELO trading on market stability. We regress y_{it} on $\widehat{\text{M-ELO}}_{it}$ and $d_t \cdot \widehat{\text{M-ELO}}_{it}$ to obtain the Two Stage Least Squares estimators $\hat{\beta}_{1,\,2SLS}$ and $\hat{\beta}_{2,\,2SLS}$.

An additional way to handle the endogeneity problem is the construction of a weekly panel of stocks and estimation of a dynamic model using the GMM system estimator developed by Blundell and Bond (1998). The property of GMM of not relying on any specific assumption of the distribution of the residuals makes it appropriate for our estimation. To mitigate the bias caused by endogenous regressors, the GMM estimation allows using lagged explanatory variables to eliminate correlations between explanatory variables and error terms. Under these conditions, the resulting estimator consistently estimates the impact of an exogenous change in M-ELO trading activity on the market stability of the stock.

The general form of our dynamic panel regression is as follows:

$$y_{it} = \alpha_t + \beta_1 y_{i\,t-1} + \beta_2 \text{M-ELO}_{it} + \beta_3 \cdot d_t \cdot \text{M-ELO}_{it} + \theta X_{it} + C_i + u_{it}. \tag{12}$$

To construct the set of moment conditions we assume sequential exogeneity. As GMM instruments for the lags of the dependent variable, we use its next three further lags. The GMM estimates are robust but typically weakened if the number of instruments is large. This is a common practice to either collapse the instruments to avoid the bias that arises as the number of instruments becomes high or to just use the most recent lags of the dependent variable as instruments.

The regression results for the models (9) and (12) are reported in Table 3. For the panel GMM, we estimate the model (12) including the first lag of the dependent variable (Column 2). The estimation results for both models suggest that the relative M-ELO trading effect is highly significant and is negatively associated with mini-flash crash occurrences during the week. Thus, in the linear panel specification (1), the loading on the M-ELO volume relative to the total NASDAQ matched volume indicates that a ten basis points increase in the relative volume of M-ELO orders is associated with the decrease in the average number of mini-flash crashes for that week by about

0.224 in that specific security for the period from March 2018 to May 2020.

In contrast, the reduction of the holding period from 500 milliseconds to 10 milliseconds resulted in a weakening of the initial market-stabilizing effect of M-ELO trading. After May 11, 2020, a ten basis points increase in relative M-ELO volumes is associated with on average 0.048 decrease in the weakly number of mini-flash crashes. We observe, that the 98% reduction in speed bump properties of M-ELO orders comes together with 79% reduction in the initial stability improvements associated with the M-ELO trading.

Table 3 around here

Table 3 shows that for the specification (2) of the panel GMM, the first lag of the number of mini-flash crashes turns out to be significant as well. Having all the parameters fixed, an additional mini-flash crash on the previous week is associated with a 0.34 decrease in the average number of crashes on a current week.

The estimation results from the panel GMM model suggest a smaller economic effect of M-ELO trading. With a ten basis points increase in relative M-ELO volume, during the first two years, the average number of crashes decreased by around 0.093. After the design change in May 2020, this effect diminishes by 92% and equals 0.0075 points of decrease in the average number of mini-flash crashes per security per week for a ten basis points increase in M-ELO activity.

We further investigate if M-ELO trading can explain the variation in different crash characteristics. Table 4 reports the estimation results of the panel regression where the dependent variables are the absolute value of the flash crash return, duration of the crash, a number of trades during the crash, 10-minute price recovery after the crash, and the relative volatility of extreme returns during the next 30 minutes after the crash.

Since we observe 30,942 mini-flash crashes throughout the sample period, and, more specifically, only 10,970 stock-week observations with a non-zero number of crashes, we decide to proceed with 2SLS linear panel estimation[5].

Table 4 around here

The effect of M-ELO volume relative to the total volume matched by NASDAQ is statistically significant for the absolute value of the mini-flash crash return and for the duration of the crash. An increase of ten basis points in M-ELO activity is associated with an average increase of 6.6 basis points in the absolute return during the crash and with an increase of the crash duration by 0.98 seconds on average. This result may suggest that for the securities with a higher degree of M-ELO trading flash crashes may stand out more in terms of how volatile they are compared to usual periods, but, at the same time, they happen less rapidly.

We also find a strong impact of M-ELO activity on the 10-minute price recovery and the following volatility 30 minutes after the crash. On average, the price reverts about 1.83% less for every additional ten basis points in relative M-ELO trading. The effect on the subsequent price volatility is positive and significant but is not sizable economically. Ten basis points higher M-ELO activity is associated with 0.88 basis points higher price volatility relative to its average level. The model finds no impact of M-ELO on the number of trades during the crash.

All the impact of M-ELO trades on crash characteristics is mitigated by approximately 77.9%

[5] It is known that when the number of instruments in the GMM setting is increasing the bias in estimates increases as well. That is why OLS regression is likely to have higher statistical power.

during the period when the holding time of M-ELO orders was reduced to 10 milliseconds. Clearly, the speed bump effect of M-ELO is dominating in its impact on the characteristics of the crashes.

Table 4 also reports the controls which are best in explaining various mini-flash crash characteristics. We observe that the security's turnover rank and the messages-to-trades ratio are statistically significant in almost all specifications. The effect of the messages-to-trades ratio is not sizable economically for any regression, while a higher turnover rank is associated with on average shorter mini-flash crashes, a greater number of trades during the crash, and a more pronounced price recovery. The security's market capitalization rank and the presence of institutional traders are statistically significant in the absolute crash return and the aftermath volatility regressions but are not noticeable economically.

On the other hand, our proxy for the algorithmic trading based on the intensity of "strategic runs" shows to be relevant for the price recovery regression. With a one percentage point higher algorithmic trading activity in a security, the price reversal after the crash tends to be 0.87 percentage points higher. This is in line with the previous results of Brogaard et al. (2014) that HFT firms on average trade in the opposite direction of the crash and supply liquidity to non-high frequency traders.

3.2. Liquidity provision measures

We further analyze the impact of relative M-ELO trading on various measures of liquidity. Table 5 reports the estimation results of model (9), where the dependent variable is one of the liquidity measures mentioned in Section 2.4. As spreads, depth, and depth imbalance are high-frequency liquidity measures, we do not expect to find a strong seasonality in those measures at weekly time frames. Therefore, a static panel regression is preferable for the analysis since the estimates have higher statistical power.

Table 5 around here

We find no statistically significant effect of M-ELO activity on the quoted half spread, the 5-minute price impact of trades, and on the absolute depth imbalance. At the same time, the model results suggest there is a strong effect of M-ELO trading on the relative depth available near the midprice. One percentage point increase in M-ELO activity in security i is associated with 1.4 percentage points increase in the share of available depth 30 basis points around the midprice relative to the average daily dollar volume in security i.

This effect was once again reduced after May 11, 2020, when the holding period of M-ELO orders was reduced from 500 milliseconds to 10 milliseconds. The positive effect on the available depth decreased by around 85.1%.

Table 5 also reports estimation results for control variables in each model specification. We observe that, in most cases, the market capitalization rank of the security is a significant determinant of its liquidity. Stocks with a higher rank in terms of market capitalization tend to have a lower quoted spread and a more balanced limit order book profile. A higher level of institutional investors' presence is associated with on average higher quoted spreads and at the same time deeper limit order books throughout a day. Also, more intense algorithmic trading is detected in the security that is associated with a lower relative depth in the book, but this effect is not large economically.

3.3. Robustness

In the following section, we detail additional robustness tests to support previous results. As the first robustness exercise, we estimate the model using alternative specifications of M-ELO trading. Previously, we used to relate M-ELO volumes to NASDAQ's total matched volumes. But M-ELO activity can also be compared to the overall dark volume handled by NASDAQ. This allows to distinguish M-ELO trading from other dark trading activity and to determine any additional or specific impact of M-ELO on market stability measures.

Also, we relate M-ELO volumes to the consolidated volume traded in each particular security. This specification of M-ELO activity takes into account the fact that M-ELO orders are available only to NASDAQ's participants, while NASDAQ may not have the biggest share in trading for some particular stock. Table 6 reports the results of the analysis for M-ELO volume related to NASDAQ's dark volume in columns 1 and 3, and total consolidated volume in columns 2 and 4. The estimation results for both static and dynamic panels suggest the effect of M-ELO stays highly significant and becomes more pronounced economically.

Table 6 around here

As an additional robustness test, we estimate the model separately for big and small stocks. We define a stock as a big one if its daily dollar trading volume is above all stocks' median trading volume throughout the sample period. Columns 1 and 2 in Table 7 report the estimation results for samples of big and small stocks separately. The estimates suggest that the results are driven by the most actively traded securities. There is only a marginally significant effect of M-ELO orders on the stability of small stocks, while the effect on the sample of big stocks is strongly significant at a

1% level. For both samples of small and big stocks, the economical effect of M-ELO is of the same order as for the full sample.

Table 7 around here

Additionally, in column 3 of Table 7, we report the estimation results of the linear panel model in Equation (9) after we remove the outliers in the stock trading activity. We drop stock-week observations where the average daily trading volume is located in the outside 1% of the empirical distribution. In total, about 15.3% of the observations were removed. The estimates indicate a strongly significant but less pronounced effect of the M-ELO activity on the market stability. For the period from March 2018 to May 2020 a higher degree of M-ELO trading was associated with fewer number of mini-flash crashes, while, after the holding period of M-ELO orders was decreased by 98%, this positive effect reduced by 80.2%.

Finally, as a robustness exercise, we apply a difference-in-difference strategy. As we do not have data on stocks where M-ELO is not implemented (e.g. outside the US), so we perform a Rajan and Zingales type of difference-in-difference (see Rajan and Zingales (1996)) by making two groups of stocks within our current sample based upon where we expect the benefit to be highest. We consider the introduction of M-ELO orders and the reduction in the holding period as separate events and estimate two difference-in-difference regressions with different control and treatment groups.

For the first event, of the introduction of M-ELO orders, we expect stocks with a previously high volume of dark trading to benefit the most from the M-ELO. Therefore, we mark stocks that before week 11 of the year 2018 had a relative dark trading share above the median in our sample as the treatment group and those that had a relative dark trading share below the median as the control group.

For the second event, when the holding period of M-ELO was reduced from 500 milliseconds to 10 milliseconds, we expect stocks that had a high level of M-ELO activity to be impacted by the rule change. Thus, the treatment group for the second event will consist of stocks that had a relative level of M-ELO trading above the median in our sample prior to the rule change.

Table 8 presents the results of estimation of the following regression:

$$y_{it} = \beta_0 + \beta_1 \text{Time}_t + \beta_2 \text{Intervention}_i + \beta_3 \text{Time}_t \times \text{Intervention}_i + u_{it}, \quad (13)$$

where y_{it} is an average weekly number of mini-flash crashes in security i at period t, Time_t is dummy variable for a treatment period that is equal to one if the rule change is active and zero otherwise, Intervention_i is a dummy variable for a treatment group that is equal to one if the security i is believed to be strongly affected by the rule change and zero otherwise. The interaction of Time and Intervention represents the causal effect of the rule change. The data for the analysis covers eight weeks before and after each rule change.

Table 8 around here

The estimation results are in line with our previous analysis. The difference-in-differences regressions suggest that there was a negative effect of M-ELO introduction on the average number of mini-flash crashes for the securities with a considerable level of dark trading. On the other hand, when the holding period of M-ELO decreased, the initial change to the average M-ELO frequency was gone. We would not, however, fully rely upon the above results as the treatment and control groups were created from the existing sample of all stocks traded on NASDAQ.

4. Conclusion

This paper provides novel evidence on market stability and liquidity provision due to the implementation of a non-displayed (dark) Midpoint Extended Life Order (M-ELO). M-ELO is a dark order that cannot interact with lit (visible) orders. It also possesses the speed bump effect due to the holding period prior to the execution. We use high-frequency order book message data from the NASDAQ exchange for the three years of M-ELO existence. The rule change applied on May 11, 2020, makes it possible to disentangle the dark and the speed bump impacts of M-ELO orders on market stability and liquidity.

For the period from January 2018 to December 2020, the degree of M-ELO activity is associated with a lower frequency of mini-flash crashes for NASDAQ traded securities. Results from panel regressions suggest the presence of significant effects of the M-ELO trading on crash returns, duration, volatility, and price recovery after the crash. Higher relative volumes traded via M-ELO are associated with less rapid crashes with a smaller price recovery, which brings them closer to natural information incorporation events. The effect of the M-ELO on the quality of the liquidity provision is mixed. We find no statistically significant effects of M-ELO trading on both spreads and depth imbalance. At the same time, the limit order book tends to be deeper for securities with a higher degree of M-ELO activity.

Our analysis shows that trading activity in M-ELO impacts market stability and liquidity mainly due to the speed bump effect. The reduction in the M-ELO's holding period by 98% decreases the influence of M-ELO on the market by around 79% on average. The robustness of the results to different specifications of the model strengthens the conclusion that only about 21% of the M-ELO market stabilizing effect comes from its dark properties and 79% from the speed bump properties.

As M-ELO volumes are relatively small, we are cautious about extrapolating the results of this analysis. The main goal of our research is to identify the effects of M-ELO on market stability during recent years. Our study delivers an important insight for market participants, policymakers, and researchers. The trade-off between execution speed and order transparency is capable of impacting the general stability of financial markets.

References

Andersen, T. G., Bollerslev, T., Diebold, F. X., Ebens, H., 2001. The distribution of realized stock return volatility. *Journal of Financial Economics* 61 (1), 43–76.

Bellia, M., Christensen, K., Kolokolov, A., Pelizzon, L., Renò, R., 2020. High-frequency trading during flash crashes: walk of fame or hall of shame? *Unpublished working paper*.

Belter, K., 2007. Supply and information content of order book depth: the case of displayed and hidden depth. *Unpublished working paper*.

Biais, B., Foucault, T., 2014. Hft and market quality. *Bankers, Markets & Investors* (128), 5–19.

Biais, B., Foucault, T., Moinas, S., 2015. Equilibrium fast trading. *Journal of Financial Economics* 116 (2), 292–313.

Blundell, R., Bond, S., 1998. Initial conditions and moment restrictions in dynamic panel data models. *Journal of Econometrics* 87 (1), 115–143.

Boulatov, A., George, T. J., 2013. Hidden and displayed liquidity in securities markets with informed liquidity providers. *Review of Financial Studies* 26 (8), 2095–2137.

Brogaard, J., Carrion, A., Moyaert, T., Riordan, R., Shkilko, A., Sokolov, K., 2018. High frequency trading and extreme price movements. *Journal of Financial Economics* 128 (2), 253–265.

Brogaard, J., Hendershott, T., Riordan, R., 2014. High-frequency trading and price discovery. *Review of Financial Studies* 27 (8), 2267–2306.

Buti, S., Rindi, B., Werner, I. M., 2017. Dark pool trading strategies, market quality and welfare. *Journal of Financial Economics* 124 (2), 244–265.

Carrion, A., 2013. Very fast money: High-frequency trading on the NASDAQ. *Journal of Financial Markets* 16 (4), 680–711.

Comerton-Forde, C., Malinova, K., Park, A., 2018. Regulating dark trading: Order flow segmentation and market quality. *Journal of Financial Economics* 130 (2), 347–366.

Comerton-Forde, C., Putninš, T. J., 2015. Dark trading and price discovery. *Journal of Financial Economics* 118 (1), 70–92.

Degryse, H., De Jong, F., Kervel, V. v., 2015. The impact of dark trading and visible fragmentation on market quality. *Review of Finance* 19 (4), 1587–1622.

Easley, D., López de Prado, M. M., O'Hara, M., 2012. Flow toxicity and liquidity in a high-frequency world. *Review of Financial Studies* 25 (5), 1457–1493.

Foley, S., Putniņš, T. J., 2016. Should we be afraid of the dark? Dark trading and market quality. *Journal of Financial Economics* 122 (3), 456–481.

Foucault, T., Pagano, M., Roell, A., Röell, A., 2013. Market liquidity: theory, evidence, and policy. Oxford University Press.

Golub, A., Keane, J., Poon, S.-H., 2012. High frequency trading and mini flash crashes. *Unpublished working paper.*

Hasbrouck, J., Saar, G., 2013. Low-latency trading. *Journal of Financial Markets* 16 (4), 646–679.

Hu, E., 2018. Intentional access delays, market quality, and price discovery: Evidence from IEX becoming an exchange. *Unpublished working paper.*

Johnson, N., Zhao, G., Hunsader, E., Qi, H., Johnson, N., Meng, J., Tivnan, B., 2013. Abrupt rise of new machine ecology beyond human response time. *Scientific Reports* 3, 2627.

Kirilenko, A., Kyle, A. S., Samadi, M., Tuzun, T., 2017. The Flash Crash: High-frequency trading in an electronic market. *Journal of Finance* 72 (3), 967–998.

Leal, S. J., Napoletano, M., Roventini, A., Fagiolo, G., 2016. Rock around the clock: An agent-based model of low- and high-frequency trading. *Journal of Evolutionary Economics* 26 (1), 49–76.

O'Hara, M., Yao, C., Ye, M., 2014. What's not there: Odd lots and market data. *Journal of Finance* 69 (5), 2199–2236.

Rajan, R., Zingales, L., 1996. Financial dependence and growth. *NBER Working paper series.*

Zhu, H., 2014. Do dark pools harm price discovery? *Review of Financial Studies* 27 (3), 747–789.

Tables and Figures

Table 1. Descriptive statistics of the sample firms.
The data set covers observations for 1,342 firms and exchange traded funds for the period from January 2, 2018 to December 31, 2020. The table shows the mean, standard deviation, and quartiles of all variables. The coefficient of variation of the extreme returns (σ_r/μ_r) shows the extent of return variability in relation to its mean, where extreme returns are maximum possible returns during crash identification intervals. NASDAQ's share denotes the share of consolidated traded volume handled by NASDAQ. QS is quoted half spread, RS_{5min} and PI_{5min} are realized spread and price impact in the following 5 minutes after the trade, respectively. Depth(30) is the U.S. dollar value of shares available 30 basis points around the midquote. DI(30) shows the imbalance of buy-sell orders 30 basis points around the midquote. The strategic runs variable shows the fraction of time, high-frequency traders (HFTs) engage in strategical order submission during the day. Msg/Trades represents the ratio of all order modification messages relative to executed trades. Msg/$100 shows how many order add messages are submitted for every $100 traded. The statistics are equally weighted based on a daily observations per firm.

	Mean	StDev	25th	50th	75th
General Characteristics					
Price	103.58	192.71	34.76	58.50	110.36
Trades '000	15.95	18.46	7.70	11.61	16.95
Shares, 'M	2.03	2.80	0.66	1.14	2.20
Volume $'M	161.26	430.13	39.03	70.76	131.80
σ_r/μ_r	0.84	0.17	0.73	0.82	0.93
NASDAQ's share	0.23	0.11	0.13	0.19	0.33
Liquidity Measures					
QS, bps	9.23	65.96	2.79	4.17	6.65
RS_{5min}, bps	0.01	11.41	-1.42	-0.20	1.01
PI_{5min}, bps	0.32	9.18	-0.97	0.16	1.40
Depth(30), $'M	2.29	4.85	0.42	0.85	1.88
DI(30)	3.50	126.14	0.18	1.47	3.57
Algorithmic Trading (AT) Measures					
Strategic Runs, %	1.21	4.69	0.11	0.29	0.75
Msg/Trades	35.56	23.52	20.64	29.95	43.64
Msg/$100	0.67	0.59	0.28	0.51	0.86

Table 2. Descriptive statistics of the identified mini-flash crashes.
The data set covers observations for 1,342 firms and exchange traded funds for the period from January 2, 2018 to December 31, 2020. The table shows the mean, standard deviation, and quartiles of various crash characteristics for 30,942 identified mini-flash crashes. The Z-score of the return is the value of $\frac{r_i - \mu_r}{\sigma_r}$, where r_i is the extreme return on the interval i, μ_r and σ_r are the mean and standard deviation of extreme returns on that day. The reversal shows what fraction of the initial jump did the price retrace 10 minutes after the crash. Number of trades, shares traded, and dollar volume traded are counted during the period of the crash. Aftermath volatility shows the size of the standard deviation of the extreme returns during the next 30 minutes after the crash ends relative to the standard deviation of the extreme returns throughout the day. M-ELO/Matched represents the relative amount of shares traded with M-ELO orders compared to the total amount of shares matched by the NASDAQ exchange. The statistics are equally weighted based on a daily observations per firm.

	Mean	StDev	25th	50th	75th
# of Crashes	0.874	1.807	0.00	0.00	1.00
Return, bps	-12.622	203.067	-59.55	-5.51	52.04
Abs. Return, bps	97.869	178.372	31.21	56.02	98.49
Return Z-score	8.486	1.560	7.49	8.06	8.93
Duration, s	30.929	23.782	13.00	25.43	42.50
Reversal	0.918	0.525	0.56	0.85	1.18
# of Trades	185.051	164.806	79.00	142.00	241.00
# of Shares, '000	31.621	62.176	7.13	16.04	36.01
Volume, $'M	2.153	4.066	0.46	0.99	2.23
Aftermath Volatility	1.227	0.434	0.98	1.17	1.39
M-ELO/Matched, %	1.181	1.842	0.22	0.63	1.47

Table 3. The effect of M-ELO trading on the average weekly number of mini-flash crashes.
The table reports the estimation results for the following regressions (with and without a lag of y_{it}):

$$y_{it} = \alpha_t + \beta_1 y_{i\,t-1} + \beta_2 \text{M-ELO}_{it} + \beta_3 \cdot d_t \cdot \text{M-ELO}_{it} + \theta X_{it} + C_i + u_{it},$$

which is estimated on a sample of 1,342 stocks traded on NASDAQ from January 2, 2018 to December 31, 2020. The specification in column (1) is a linear static panel instrumental variables model with time and fixed effects. The model specification in column (2) describes a dynamic panel estimated using GMM with the three most recent lags of the dependent variable as GMM instruments. The dependent variable is the average number of mini-flash crashes. The set of control variables includes market capitalization and turnover ranks from the MIDAS database, messages-to-trades ratio, and proxies for the AT activity of Hasbrouck and Saar (2013) and the institutional trading activity. This is an unbalanced panel with weekly observations. M-ELO is the share of the Midpoint Extended Life Order volume relative to the total volume handled by NASDAQ for the particular stock, d_t is a dummy variable that equals one at the time when NASDAQ decreased the M-ELO holding period by 98%, and zero otherwise. The value of M-ELO trading is instrumented by the market average M-ELO activity of the stocks within the same market capitalization rank and excluding: (i) stock i itself, (ii) stocks from the same sector as stock i, (iii) stocks in the same index as stock i. Heteroskedasticity corrected t-statistics are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable:	
	Average weekly number of crashes	
	Panel linear	*Panel GMM*
	(1)	(2)
$y_{i,t-1}$		-0.345^{***}
		(-26.745)
M-ELO	-224.079^{***}	-93.458^{***}
	(-3.262)	(-4.403)
$d_t \cdot$ M-ELO	176.382^{***}	85.989^{***}
	(3.139)	(4.174)
P value($\beta_2 = 0$ and $\beta_3 = 0$)	$9.093 \cdot 10^{-17}$	$< 2 \cdot 10^{-16}$
Controls	Yes	Yes
Observations	25,497	21,528
F-statistic	727.3^{***}	$1{,}011.5^{***}$

Table 4. The effect of M-ELO trading on crash characteristics.
The table reports the estimation results for the following linear panel regression:

$$y_{it} = \alpha_t + \beta_1 \text{M-ELO}_{it} + \beta_2 \cdot d_t \cdot \text{M-ELO}_{it} + \theta X_{it} + C_i + u_{it},$$

which is estimated on a sample of 1,342 stocks traded on NASDAQ from January 2, 2018 to December 31, 2020. The dependent variables are the absolute value of the crash return, crash duration in seconds, the number of trades executed during the crash, the price reversal 10 minutes after the crash and the relative volatility 30 minutes after the crash. M-ELO is the share of the Midpoint Extended Life Order volume relative to the total volume handled by NASDAQ for the particular stock, d_t is a dummy variable that equals one at the time when NASDAQ decreased the M-ELO holding period by 98%, and zero otherwise. The value of M-ELO trading is instrumented by the market average M-ELO activity of the stocks within the same market capitalization rank and excluding: (i) stock i itself, (ii) stocks from the same sector as stock i, (iii) stocks in the same index as stock i. The control variables are proxies for the AT activity of Hasbrouck and Saar (2013) and the institutional trading activity, market capitalization and turnover ranks from the MIDAS database, and the messages-to-trades ratio. This is an unbalanced panel with weekly observations. Heteroskedasticity corrected t-statistics are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable:				
	|Return|	Duration	# Trades	Reversal	σ_{aft}
	(1)	(2)	(3)	(4)	(5)
M-ELO	0.663***	977.804**	−943.399	−18.275**	0.088***
	(3.162)	(2.367)	(−0.401)	(−2.009)	(3.012)
$d_t \cdot$ M-ELO	−0.514***	−765.947**	−182.451	14.018*	−0.070***
	(−2.971)	(−2.296)	(−0.097)	(1.884)	(−2.913)
Institutional	0.117*	69.190	88.961	−0.188	0.020**
	(1.843)	(1.158)	(0.138)	(−0.161)	(2.264)
AT	−0.009***	−4.671	14.760	0.870***	$-8.252 \cdot 10^{-4}$***
	(−4.685)	(−0.894)	(0.518)	(6.536)	(−3.547)
Market Cap. Rank	−0.005***	−1.205	−0.745	−0.007	$-7.305 \cdot 10^{-4}$***
	(−4.467)	(−1.440)	(−0.132)	(−0.432)	(−5.271)
Turnover Rank	$-4.636 \cdot 10^{-4}$***	−5.646***	15.530***	0.064***	$8.075 \cdot 10^{-6}$
	(−3.925)	(−13.302)	(7.628)	(9.061)	(0.514)
Msg/Trades	$3.827 \cdot 10^{-5}$***	0.327***	−0.579***	−0.006***	$7.046 \cdot 10^{-7}$
	(3.241)	(9.454)	(−3.231)	(−7.140)	(0.404)
P value($\beta_1 = 0$ and $\beta_2 = 0$)	$7.782 \cdot 10^{-7}$	0.002	$1.699 \cdot 10^{-4}$	0.040	$6.508 \cdot 10^{-7}$
Observations	10,970	10,970	10,970	10,970	10,769
R^2	0.006	0.059	0.012	0.010	0.011
F-statistic	364.458***	$1.206 \cdot 10^{3}$***	152.441***	308.697***	573.865***

Table 5. The effect of M-ELO trading on market liquidity.
The table reports the estimation results for the following linear panel regression:

$$y_{it} = \alpha_t + \beta_1 \text{M-ELO}_{it} + \beta_2 \cdot d_t \cdot \text{M-ELO}_{it} + \theta X_{it} + C_i + u_{it},$$

which is estimated on a sample of 1,342 stocks traded on NASDAQ from January 2, 2018 to December 31, 2020. Dependent variables are quoted half-spread, 5-minute realized spread, dollar depth available 30 basis points around the midquote relative to the average daily dollar trading volume, and the absolute depth imbalance 30 basis points around the midquote. M-ELO is the share of the Midpoint Extended Life Order volume relative to the total volume handled by NASDAQ for the particular stock, d_t is a dummy variable that equals one at the time when NASDAQ decreased the M-ELO holding period by 98%, and zero otherwise. The value of M-ELO trading is instrumented by the market average M-ELO activity of the stocks within the same market capitalization rank and excluding: (i) stock i itself, (ii) stocks from the same sector as stock i, (iii) stocks in the same index as stock i. The control variables are proxies for the AT activity of Hasbrouck and Saar (2013) and the institutional trading activity, market capitalization and turnover ranks from the MIDAS database, and the messages-to-trades ratio. This is an unbalanced panel with weekly observations. Heteroskedasticity corrected t-statistics are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable:			
	QS	RS_{5min}	$\frac{\text{Depth(30)}}{\text{\$ Volumes}}$	$\lvert$DI(30)$\rvert$
	(1)	(2)	(3)	(4)
M-ELO	-121.522	-155.429	1.403^{***}	-722.069
	(-0.173)	(-1.565)	(2.656)	(-0.450)
$d_t \cdot$ M-ELO	129.268	121.760	-1.193^{***}	581.767
	(0.226)	(1.482)	(-2.757)	(0.443)
Institutional	861.277^{***}	55.550	0.120^{**}	-164.172
	(3.276)	(0.898)	(2.335)	(-0.334)
AT	2.742	2.825^{**}	-0.032^{***}	0.291
	(0.474)	(2.406)	(-3.840)	(0.033)
Market Cap. Rank	-5.656^{***}	0.086	-0.002^{**}	-7.302^{*}
	(-2.676)	(0.367)	(-2.384)	(-1.836)
Turnover Rank	0.212	-0.035	-0.002^{***}	0.222
	(0.665)	(-0.765)	(-6.953)	(0.450)
Msg/Trades	-0.014	-0.007	$3.709 \cdot 10^{-4***}$	-0.051
	(-0.385)	(-1.372)	(7.180)	(-0.726)
P value($\beta_1 = 0$ and $\beta_2 = 0$)	0.284	0.014	$1.829 \cdot 10^{-12}$	0.775
Observations	25,497	25,481	25,497	25,485
R^2	0.020	$3.484 \cdot 10^{-4}$	0.044	$6.520 \cdot 10^{-5}$
F-statistic	533.461^{***}	54.163^{***}	$1.310 \cdot 10^{3***}$	15.223^{**}

Table 6. Alternative M-ELO activity specifications.
The table reports the estimation results for the following regressions (with and without a lag of $y_{i,t}$):

$$y_{it} = \alpha_t + \beta_1 y_{i\,t-1} + \beta_2 \text{M-ELO}_{it} + \beta_3 \cdot d_t \cdot \text{M-ELO}_{it} + \theta X_{it} + C_i + u_{it},$$

which are estimated on a sample of 1,342 stocks traded on NASDAQ from January 2, 2018 to December 31, 2020. The specification in columns (1) and (2) is a linear static panel instrumental variables model with time and fixed effects. The specifications for columns (3) and (4) describe a dynamic panel estimated using GMM with three most recent lags of the dependent variable as GMM instruments for differenced equation. The dependent variable is the average number of mini-flash crashes. M-ELO (Dark) and M-ELO (Cons) represent a share of M-ELO volume relative to, respectively, the volume of dark trading on NASDAQ, and total consolidated volume across exchanges, d_t is a dummy variable that equals one at the time when NASDAQ decreased the M-ELO holding period by 98%, and zero otherwise. The set of control variables includes market capitalization and turnover ranks from the MIDAS database, messages-to-trades ratio, and proxies for the AT activity of Hasbrouck and Saar (2013) and the institutional trading activity. This is an unbalanced panel with weekly observations. The value of M-ELO trading is instrumented by the market average M-ELO activity of the stocks within the same market capitalization rank and excluding: (i) stock i itself, (ii) stocks from the same sector as stock i, (iii) stocks in the same index as stock i. Heteroskedasticity corrected t-statistics are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable:			
	Average weekly number of crashes			
	Panel linear		*Panel GMM*	
	(1)	(2)	(3)	(4)
$y_{i,t-1}$			-0.343^{***}	-0.342^{***}
			(-26.568)	(-26.441)
M-ELO (Dark)	-442.927^{***}		-180.401^{***}	
	(-2.665)		(-4.274)	
M-ELO (Cons)		$-1.295 \cdot 10^{3***}$		-652.454^{***}
		(-3.054)		(-4.415)
$d_t \cdot$ M-ELO (Dark)	370.549^{***}		170.462^{***}	
	(2.595)		(4.125)	
$d_t \cdot$ M-ELO (Cons)		$1.089 \cdot 10^{3***}$		605.685^{***}
		(2.939)		(4.207)
P value($\beta_2 = 0$ and $\beta_3 = 0$)	$1.829 \cdot 10^{-11}$	$4.221 \cdot 10^{-15}$	$< 2 \cdot 10^{-16}$	$< 2 \cdot 10^{-16}$
Controls	Yes	Yes	Yes	Yes
Observations	25,497	25,497	21,528	21,528
F-statistic	481.4^{***}	801.7^{***}	974.3^{***}	964.4^{***}

Table 7. The effect of M-ELO trading on the number of mini-flash crashes for big and small stocks and outlier-robust effects.

This table reports the estimation results for the following linear panel regression:

$$y_{it} = \alpha_t + \beta_1 \text{M-ELO}_{it} + \beta_2 \cdot d_t \cdot \text{M-ELO}_{it} + \theta X_{it} + C_i + u_{it},$$

which is estimated on a sample of 1,342 stocks traded on NASDAQ from January 2, 2018 to December 2, 2020. The dependent variable is the average number of mini-flash crashes. Column (1) reports results for the sub-sample of small stocks, and column (2) for the sub-sample of big stocks in terms of average daily trading volumes. Column (3) reports results for the full sample but after discarding the highest and the lowest 1% values of the average daily number of trades, shares and U.S. dollar volume. M-ELO is the share of the Midpoint Extended Life Order volume relative to the total volume handled by NASDAQ for the particular stock, d_t is a dummy variable that equals one at the time when NASDAQ decreased the M-ELO holding period by 98%, and zero otherwise. The value of M-ELO trading is instrumented by the market average M-ELO activity of the stocks within the same market capitalization rank and excluding: (i) stock i itself, (ii) stocks from the same sector as stock i, (iii) stocks in the same index as stock i. The set of control variables includes market capitalization and turnover ranks from the MIDAS database, messages-to-trades ratio, and proxies for the AT activity of Hasbrouck and Saar (2013) and the institutional trading activity. This is an unbalanced panel with weekly observations. Heteroskedasticity corrected t-statistics are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable:		
	Average weekly number of crashes		
	(1)	(2)	(3)
M-ELO	-117.831^{*}	-145.442^{***}	-114.053^{***}
	(-1.952)	(-3.181)	(-3.391)
$d_t \cdot$ M-ELO	96.414^{*}	109.031^{***}	91.494^{***}
	(1.839)	(2.960)	(3.238)
P value($\beta_1 = 0$ and $\beta_2 = 0$)	$1.135 \cdot 10^{-4}$	$8.782 \cdot 10^{-16}$	$7.702 \cdot 10^{-11}$
Controls	Yes	Yes	Yes
Sample	Small Stocks	Big Stocks	Full
Outliers Removed	No	No	Yes
Observations	6,874	18,623	21,501
R^2	0.004	0.018	0.010
F-statistic	154.240^{***}	991.815^{***}	531.058^{***}

Table 8. Difference-in-differences approach.
This table reports the estimation results for the following regression

$$y_{it} = \beta_0 + \beta_1 \text{Time}_t + \beta_2 \text{Intervention}_i + \beta_3 \text{Time}_t \times \text{Intervention}_i + u_{it},$$

for two rule changes: (i) M-ELO introduction in March 2018, and (ii) M-ELO holding period reduction in May 2020. Variable y_{it} is an average weekly number of mini-flash crashes in security i at period t, Time_t is dummy variable for a treatment period that is equal to one if the rule change is active and zero otherwise, Intervention_i is a dummy variable for a treatment group that is equal to one if the security i is believed to be strongly affected by the rule change and zero otherwise. For the first event the treatment group formation is based on the level of dark trading, while for the second event it is based on the level of M-ELO trading. The interaction of Time and Intervention represents the causal effect of the rule change. The data for the analysis cover 8 weeks before and after each rule change. Heteroskedasticity corrected t-statistics are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable:	
	Average weekly number of crashes	
	(1)	(2)
Time (2018)	−0.009	
	(−0.979)	
Intervention (2018)	0.181***	
	(12.103)	
Intervention × Time (2018)	−0.040**	
	(−1.970)	
Time (2020)		−0.114***
		(−8.090)
Intervention (2020)		−0.191***
		(−8.575)
Intervention × Time (2020)		0.069**
		(2.149)
Constant	0.140***	0.429***
	(21.330)	(39.511)
Observations	12,224	14,243
R^2	0.024	0.009
F-statistic	101.959***	40.928***

Figure 1. Mini-flash crashes, quoted spread, and book depth throughout the sample.

The first panel of the figure plots the total amount of identified mini-flash crashes across the sample of 1,342 liquid stocks traded on NASDAQ from January 2, 2018 to December 31, 2020. The second panel shows the median quoted spread (in basis points) across securities. The last panel presents the median aggregate monetary value (in thousands of dollars) of shares offered within 30 basis points around the midquote. The dashed line indicates the time when the M-ELO became available on March 12, 2018, the dotted line indicates the time of the holding period reduction from 500 ms to 10 ms on May 11, 2020. All observations are on weekly frequency.

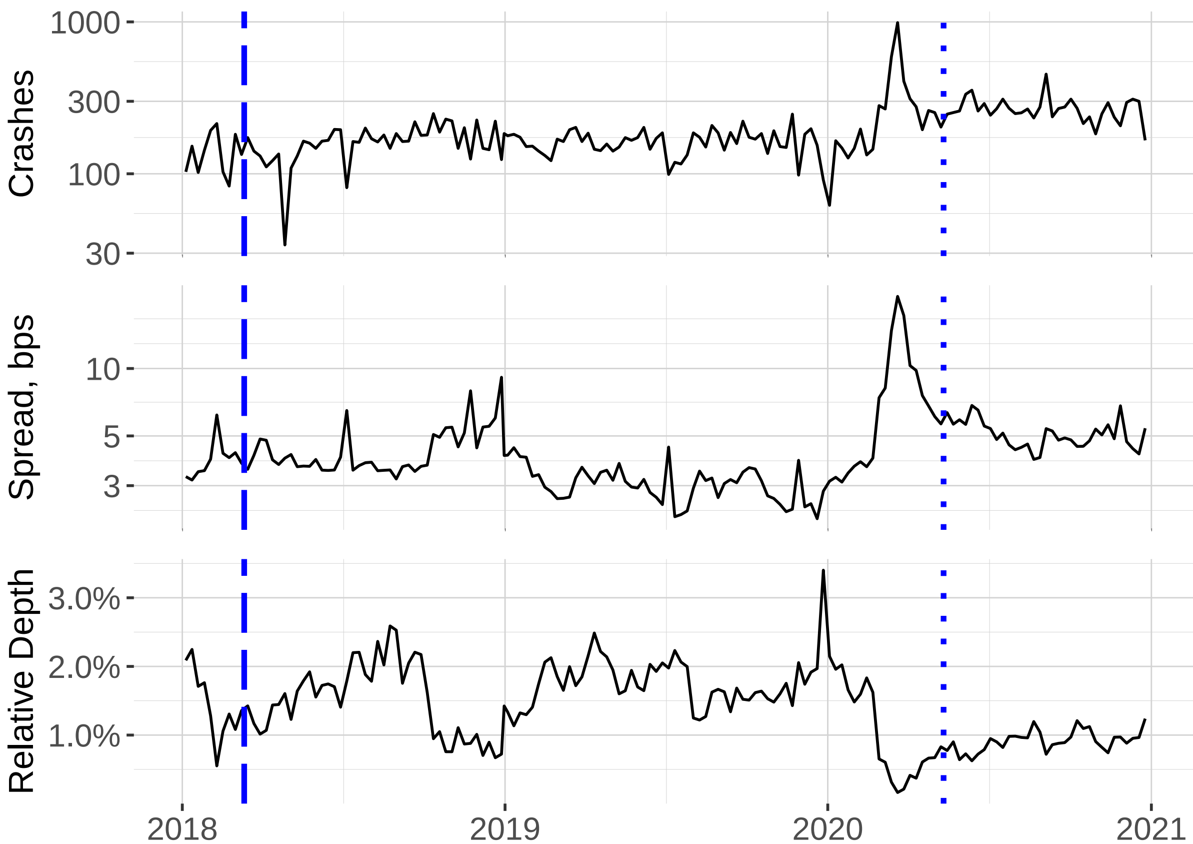

Figure 2. Example of the mini-flash crash.

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. The pre-market price of P&G was at around \$78.2, dropped to the region \$77.4 – \$77.7, where it stayed fairly stable until 2 p.m., and experienced then a massive spike to the levels of approximately \$78.3. Within the next five minutes, the price dropped more than 1.4% to \$77.16, and eventually returned to the region of its previous daily consolidation. Panel B zooms in around the time of the crash. Each dot represents a trade. The duration of the crash is 26.2 seconds, the cumulative return of the first spike is 0.98%, the volume traded during the crash is \$633 thousand.

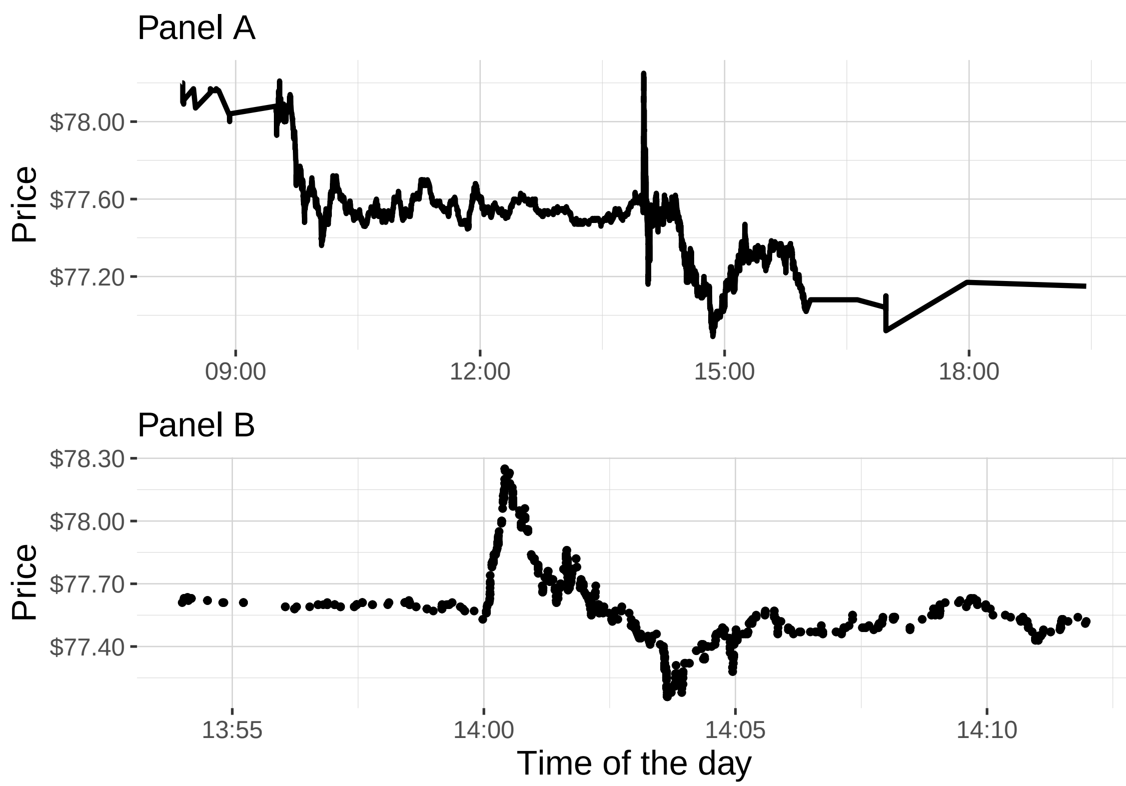

Figure 3. Share of M-ELO executions.

This figure plots time trends in M-ELO volumes relative to the total volumes matched by the NASDAQ exchange. The 10th, 50th, and 90th percentiles are depicted. M-ELO orders became available on March 12, 2018. The shaded area represents the period starting when NASDAQ decreased the M-ELO "Holding period" from 500 milliseconds to 10 milliseconds on May 11, 2020.

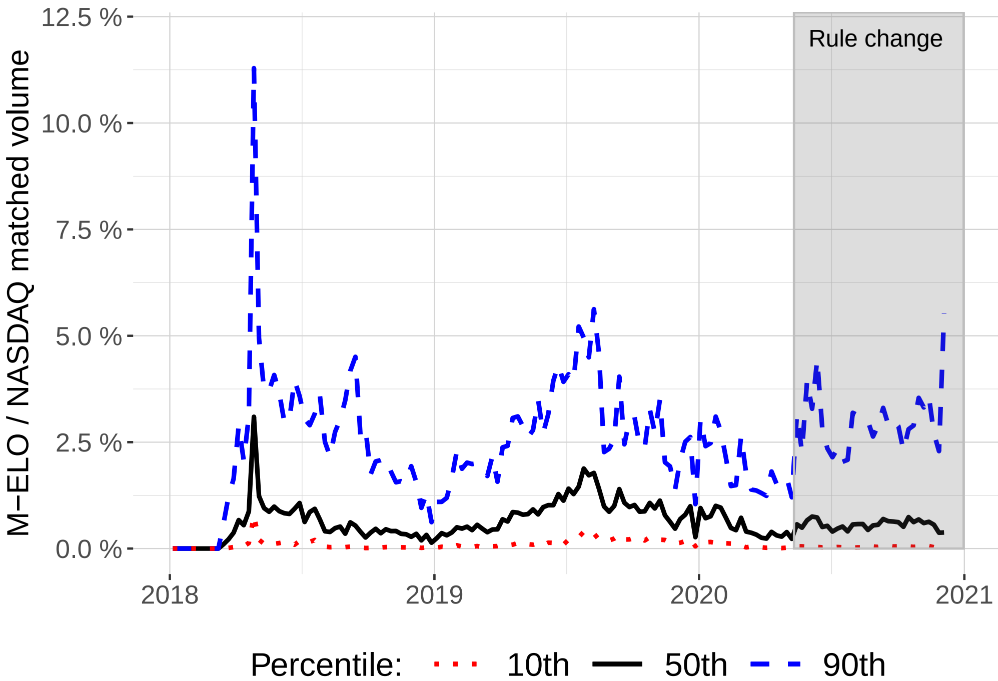

Figure 4. Densities of lit (visible) order sizes and M-ELO order sizes.

This figure plots densities of order sizes of two types of orders: (i) visible limit orders, and (ii) non-displayed, M-ELO orders. For better representation, only order sizes less than 400 shares are considered. All observations are stock-day averages.

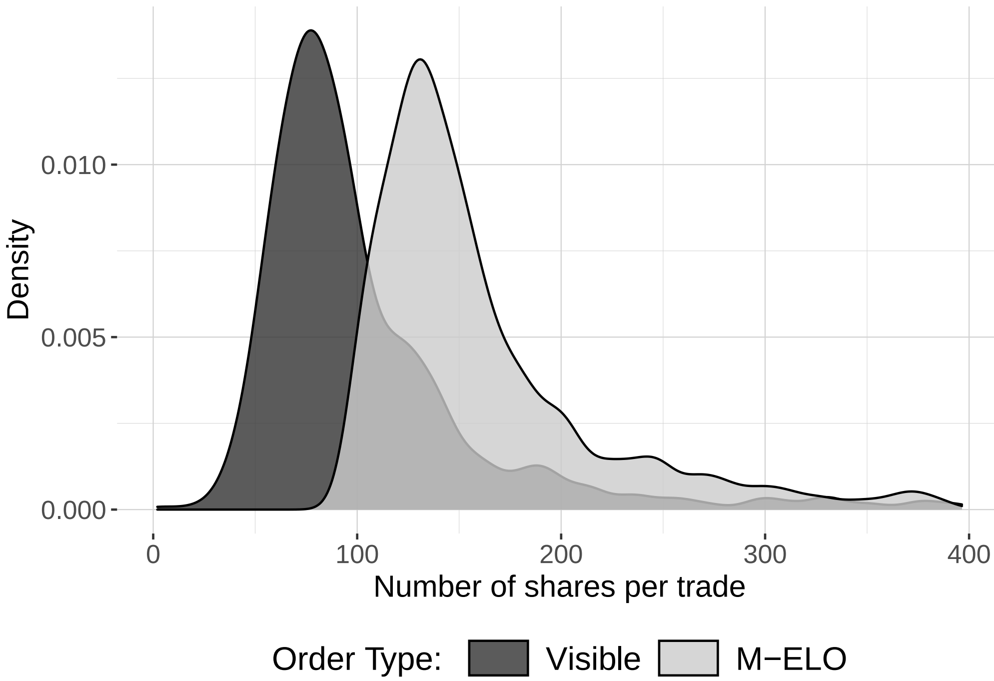

Appendix

M-ELO order specifications

- Non-displayed order type: M-ELO executions are reported to the Securities Information Processors and provided in NASDAQ's proprietary data feed in the same manner as all other transactions occurring on NASDAQ (i.e., without any new or special indication that a transaction is an M-ELO execution).
- The M-ELO timer for 500 millisecond waiting period starts upon entry if the order is marketable at the midpoint. If the order is not eligible to trade at the midpoint upon the entry, the M-ELO timer will start when the price of the order is at or better than the midpoint of the NBBO.
- Effective May 11, 2020, the holding period was reduced from 500 to 10 milliseconds.
- Any modification on a resting M-ELO order will result in a restart of the timer, except in the case of reducing the order quantity.
- The timer does not reset if the NBBO moves.
- An M-ELO order is ranked in time priority among other M-ELO orders at the time it becomes eligible to execute.
- Only round lots are accepted for the M-ELO submission.
- M-ELO orders may execute in a locked market but not in a crossed market[6].
- The M-ELO order type will never route out.

[6] In a locked market, a stock's bid price at one exchange and ask price at another exchange are identical, that is, the bid-ask spread is zero. In a crossed market, the bid price exceeds the ask price.

Tests for the instruments

Table A.1. Properties of the instrumental variables.
The table reports the results of instrument relevance and exogeneity tests. To explain the variation in the relative volume of M-ELO orders in the equation (9) the following instruments are used: Z_{it} is the average level of M-ELO trading across all stock in our sample in the same market capitalization rank excluding: (i) stock i itself, (ii) stocks from the same sector as stock i, (iii) stocks in the same index as stock i. The instruments' relevance ($\mathrm{Cov}(Z, \text{M-ELO}) \neq 0$ and $\mathrm{Cov}(d_t Z, d_t \cdot \text{M-ELO}) \neq 0$) is tested by estimating the first-stage regression, and obtaining F-statistics resulting from the test $H_0 : \pi_1 = 0$ against the the alternative $H_1 : \pi \neq 0$. The rule of thumb suggests that the F-statistics for significance of the instrument in the first-stage should exceed 10. The instrument exogeneity assumption is weakened by including the control variables X_{it} into the first-stage regression. The overidentifying test is implemented by, first, obtaining the residuals of the 2SLS model:

$$\hat{u}_{it}^{2SLS} = y_{it} - \widehat{\alpha}_t - \widehat{\beta}_1 \text{M-ELO}_{it} - \widehat{\beta}_2 \cdot d_t \cdot \text{M-ELO}_{it} - \widehat{\theta} X_{it},$$

and then regressing these residuals on the instruments and control variables. The resulting J-statistic of the test $H_0 : \eta_1 = \eta_2 = 0$ versus $H_1 : \eta_1 \neq 0$ or $\eta_2 \neq 0$ is distributed according to χ_q^2 where q is the number of instruments minus the number of endogenous regressors.

Panel A.1: Instrument Relevance	
Regression	$\text{M-ELO}_{it} = a_t + \pi_1 Z_{it} + \gamma X_{it} + C_i + v_{it},$
Hypothesis	$H_0 : \pi_1 = 0$
Statistics	$\mathrm{F}(1, 24577) = 1353.167$
p-value	$2.065 \cdot 10^{-288}$
Panel A.2: Instrument Relevance	
Regression	$d_t \cdot \text{M-ELO}_{it} = a_t + \pi_1 \cdot d_t \cdot Z_{it} + \gamma X_{it} + C_i + v_{it},$
Hypothesis	$H_0 : \pi_1 = 0$
Statistics	$\mathrm{F}(1, 24577) = 412.701$
p-value	$5.297 \cdot 10^{-91}$
Panel B: Instrument Exogeneity	
Regression	$\hat{u}_{it}^{2SLS} = a_t + \eta_1 Z_{it} + \eta_2 \cdot d_t \cdot Z_{it} + \gamma X_{it} + C_i + e_{it}$
Hypothesis	$H_0 : \eta_1 = \eta_2 = 0$
Statistics	$J = mF = 2 \cdot 0.0259 \sim \chi_1^2$
p-value	0.8201

CENTER FOR FINANCIAL STUDIES

CFS Working Paper Series

Recent Issues

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

No.	Authors	Title
690	Sandro Heiniger, Winfried Koeniger, and Michael Lechner	The Heterogeneous Response of Real Estate Asset Prices to a Global Shock
689	Alix Auzepy, Christina E. Bannier, and Fabio Martin	Walk the Talk: Shareholders' Soft Engagement at Annual General Meetings
688	Alix Auzepy, Christina E. Bannier, and Fabio Martin	*Are sustainability-linked loans designed to effectively incentivize corporate sustainability? A framework for review*
687	Lutz Kilian, Michael Plante, Alexander W. Richter	*Macroeconomic Responses to Uncertainty Shocks: The Perils of Recursive Orderings*
686	Lutz Kilian and Xiaoqing Zhou	*A Broader Perspective on the Inflationary Effects of Energy Price Shocks*
685	Lutz Kilian and Xiaoqing Zhou	*Heterogeneity in the Pass-Through from Oil to Gasoline Prices: A New Instrument for Estimating the Price Elasticity of Gasoline Demand*
684	Vanya Horneff, Raimond Maurer, and Olivia S. Mitchell	*Fixed and Variable Longevity Annuities in Defined Contribution Plans: Optimal Retirement Portfolios Taking Social Security into Account*
683	Volker Brühl	*The Green Asset Ratio (GAR) – a new KPI for Credit Institutions*