MINI FLASH CRASHES: REVIEW, TAXONOMY AND POLICY RESPONSES

Floris Laly, Mikael Petitjean

REPRINT | 2021 / 17

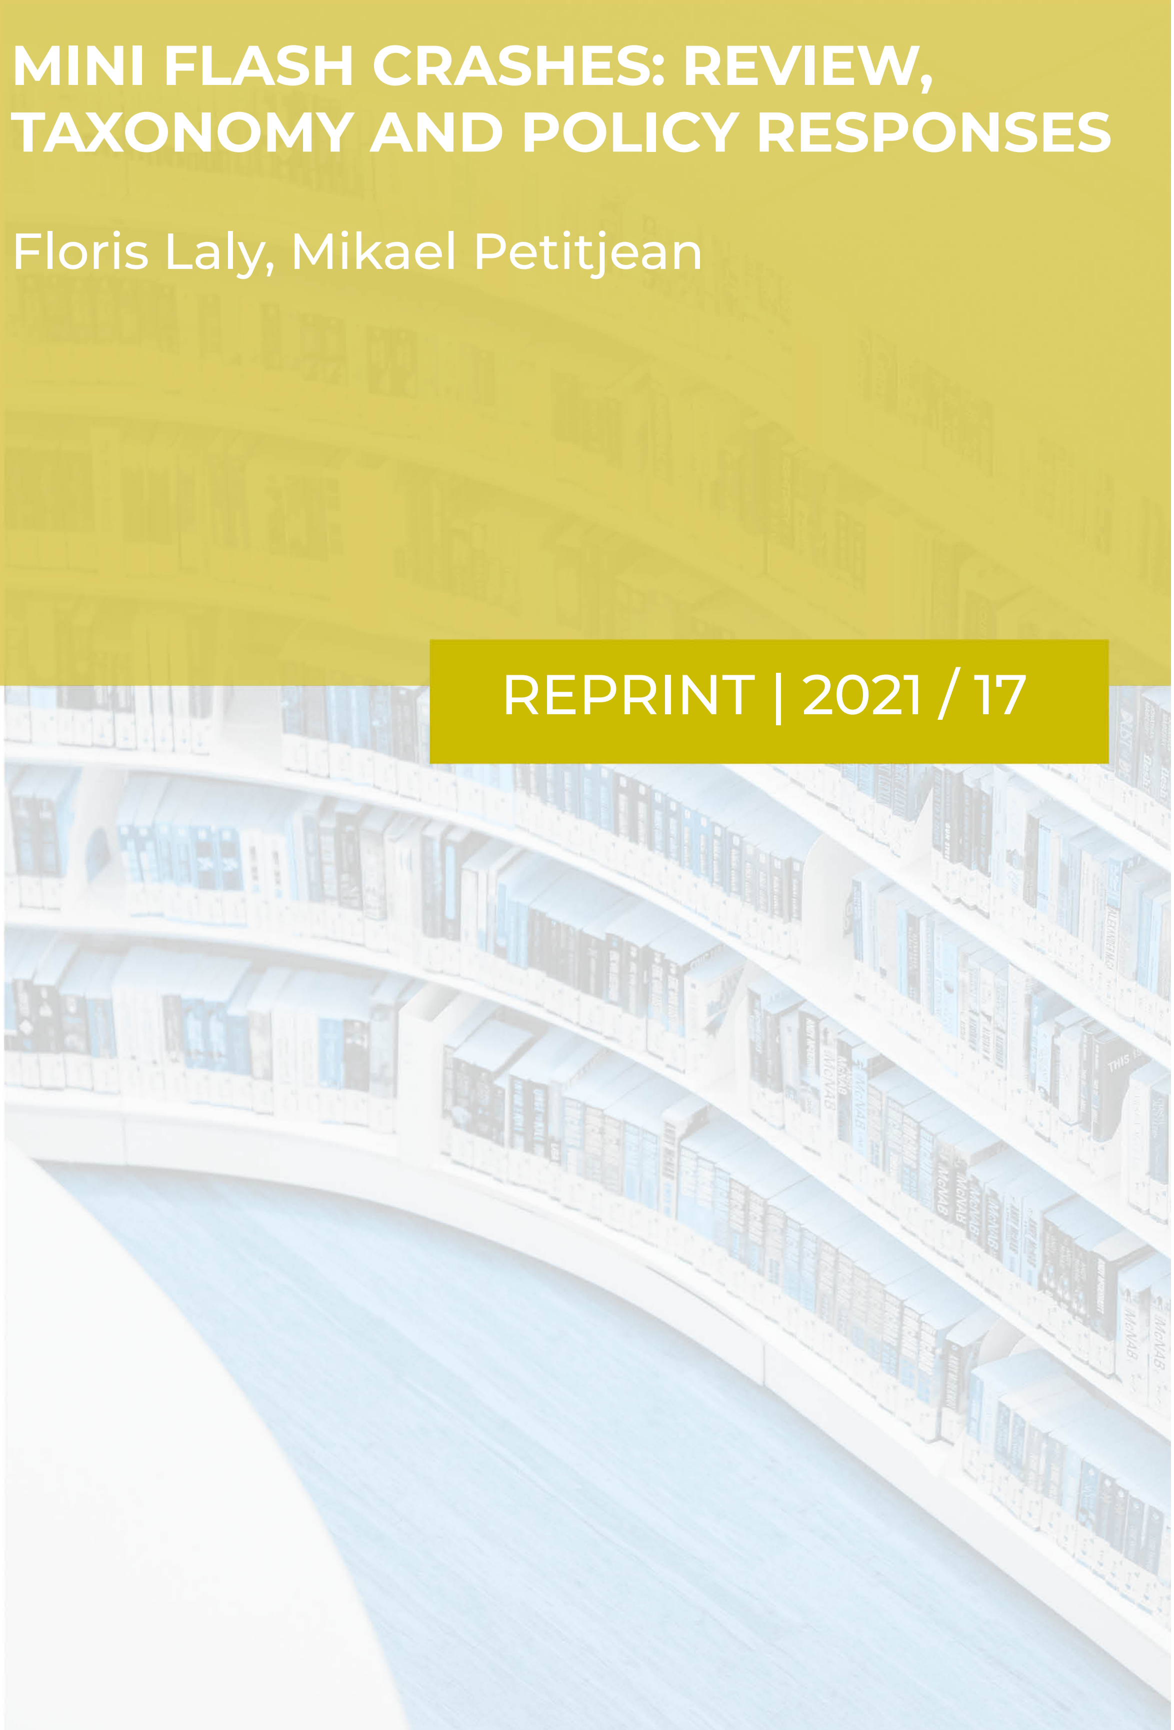

UCLouvain

LiDAM
LFIN | Louvain Institute of Data Analysis and Modeling in economics and statistics

LFIN

Voie du Roman Pays 34, L1.03.01 B-1348 Louvain-la-Neuve

Tel (32 10) 47 43 04

Email: lidam-library@uclouvain.be

https://uclouvain.be/en/research-institutes/lidam/lfin/publications.html

Mini Flash Crashes:
Review, Taxonomy and Policy Responses

Floris Laly* Mikael Petitjean† ‡

August 16, 2019

Abstract

We focus on extreme price movements known as mini flash crashes (MFCs). After reviewing the literature, we provide a taxonomy based on a sample of MFCs identified by Nanex on the U.S. financial markets over a three-year period. We detect significant differences between crashes and exchanges. In comparison to 'up crashes', we find that 'down crashes' exhibit lower absolute returns but have longer duration. We also show that the dynamics of MCFs varies across exchanges. For example, the MFCs on ARCA are on average both less severe and shorter in duration than those on the NASDAQ. We finally review all the key implications of MCFs in terms of public policy.

JEL classification: G18, G28

Keywords: mini flash crash, extreme price movement, high frequency trading, liquidity, policy responses, financial stability

*Louvain School of Management (LFIN-LIDAM, UCLouvain). Email: floris.laly@uclouvain.be.

†IÉSEG School of Management (LEM UMR 9221 CNRS) and Louvain School of Management (LFIN-LIDAM, UCLouvain). E-mail: m.petitjean@ieseg.fr.

‡We thank Charles Nokerman for helpful research assistance. We are also grateful to the discussants who helped us improve the paper during conferences and seminars. We apologize for not listing all the names. Any remaining error is solely the authors' responsibility.

1 Introduction

The speed and frequency of financial market crashes seem to have evolved side by side with the evolution of communication technologies, from slow motion crashes (1907, 1929, 1987), where a crash could unravel within a full day of trading, to flash crashes (2010), where a crash could unravel in less than an hour, to mini flash crashes (nowadays), where a crash can unravel in less than a second. The question is whether such extreme price movements occur more frequently than they should. According to Nanex (2010), MFCs frequently occur and impact every asset class without exception.[1]

Following the U.S. Flash Crash of May 6, 2010, flash events attracted the attention of both practitioners and the media who noticed sudden and extreme price movements here and there. For example, the stock of Progress Energy, a $2.5 billion utilities company with 11,000 employees listed on the New York Stock Exchange was impacted by a 90% price drop on September 27, 2010 within a one-second time interval. After the stock was halted for 5 minutes, the price quickly rallied to its initial level (Bowley, 2010).

Since then MFCs have become part of the trading landscape, impacting all types of assets, from stocks, to ETFs, futures contracts, bonds, currencies and even cryptocurrencies. Many observers now worry about the potential impact of such events on investors' confidence, not forgetting the obvious market stability and fairness issues it raises. Several scholars have also highligted the need for a better understanding of MFCs (Johnson et al. 2013; Aitken et al., 2015; Brogaard et al., 2018).

Despite the recurrence of such events, especially on the U.S. markets, the literature on the subject remains very limited and the mechanism that triggers these very short-term extreme price movements remains pretty mysterious so far. Some argue MFCs may be caused by high-

[1]Nanex LLC is a U.S. data analytics company recently awarded the Whistleblower Award by the SEC for its findings regarding the violation of Regulation NMS by NYSE and its parent company NYSE Euronext over an extended period of time.

frequency traders (HFTs) (Sornette and Von der Becke, 2011, Johnson et al., 2013) while others argue MFCs are either caused by human trades (Braun et al., 2017) or unrelated to HFTs (Brogaard et al., 2018). To sort things out, contributions in this area are deeply needed and we believe this paper can help academics, practitioners, regulators, and policymakers understand MFCs better.

The taxonomy that we propose is based on a sample of MFCs identified by Nanex on the U.S. financial markets over a three-year period. We detect significant differences between crashes and exchanges. In comparison to 'up crashes', we find that 'down crashes' exhibit lower absolute returns but have longer duration. We also show that the dynamics of MFCs varies across exchanges. For example, the MFCs on ARCA are both less severe and shorter in duration than those on the NASDAQ. We finally review all the key implications of MCFs in terms of public policy. Overall, very little is known on the exact role of high-frequency traders before, during, and after MFCs

The structure of the paper is as follows. Section 2 provides a definition of a mini flash crash. Section 3 presents a synthetic literature review on flash events. Section 4 describes the statistical properties of MFCs within our sample and provides a discussion of our empirical findings. Section 5 deals with the implications in terms of public policy. Finally, section 6 provides concluding comments.

2 Definition of a Mini Flash Crash

A mini flash crash is, as implied by the name given to it, a very short version of a flash crash, itself a short version of a crash. As such, an MFC is the fastest type of crash that presently exists. Moreover, and contrary to traditional crashes, flash crashes and mini flash crashes embody both crashes (down) and inverted crashes (up) so that they are sometimes called "flash rallies" and "mini flash rallies" when occurring on the way up (Table 1).

Table 1: Main types of crashes

Speed	Type	Maximum duration	Direction	Shape	Reverting Process
slowest	crash	years	down	\	no
↓	flash crash	minutes	down or up	V or ∧	yes
fastest	mini flash crash	seconds	down or up	V or ∧	yes

The table reports the main types of crashes that exist in modern markets and their specificities.

More precisely, an MFC can be defined as a sudden and very short-term extreme price movement which can either be negative (drop) or positive (spike), including a reverting process, which usually takes a V-shaped or inverted V-shaped form, that enables the price to come back to or close to its initial level. MFCs occur over extremely short time intervals and without any fundamental change in the underlying asset or without any specific news related to the underlying asset.

Figure 1 presents an MFC that impacted the stock of ConocoPhillips (NYSE: COP) on January 21, 2011, resulting in a 1.1% price drop in less than a second followed by the usual upward price correction, while Figure 2 presents an MFC that impacted the stock of IBM (NYSE: IBM) on January 25, 2011, resulting in a 2.2% price spike in less than a second followed by the usual downward price correction.

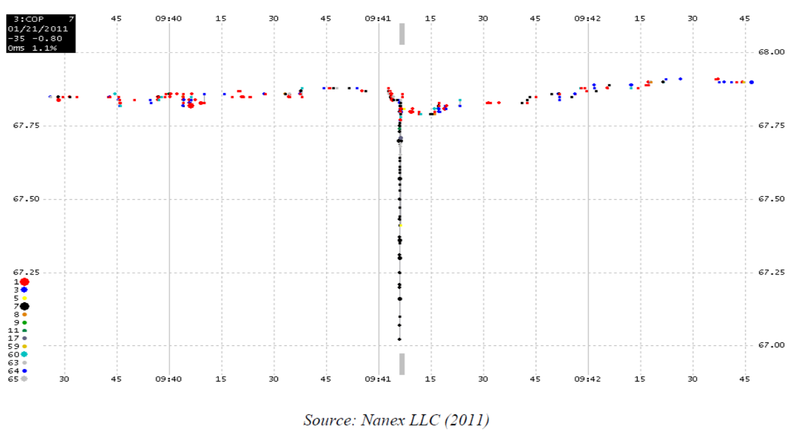

Figure 1: MFC on ConocoPhillips (Source: Nanex)

Nanex was the first, to the best of our knwoledge, to provide a definition of what an MFC is

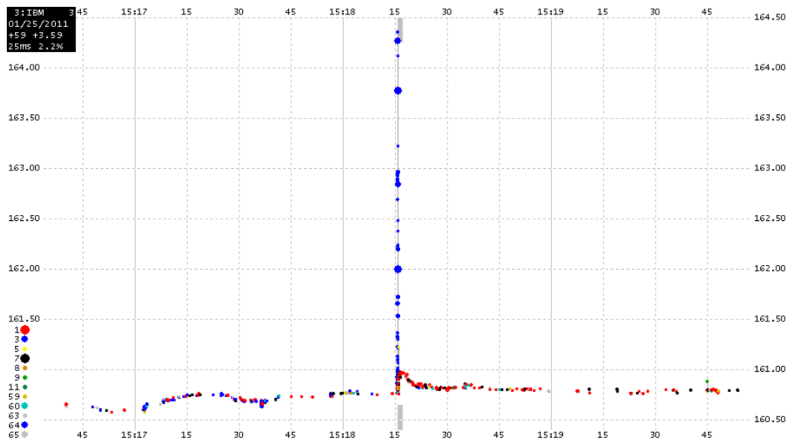

Source: Nanex LLC (2011)

Figure 2: MFC on IBM (Source: Nanex)

(Nanex, 2011). For a down (resp. up) move to be considered as an MFC, three conditions must be satisfied according to Nanex:

1) The absolute return has to exceed 0.8%.

2) The price change must occur within 1.5 seconds.

3) The price must tick down (resp. up) at least 10 times before ticking up (resp. down).

This definition is certainly not a golden rule, of course, so that the definition provided by Nanex may be too restrictive to incorporate all the very short-term extreme price movements. For example, according to Nanex (2011), the best example of an MFC occured on October 11, 2011 when the AMJ stock "plummeted from $34.90 to $32.61 (a 6.5% loss) and then recovered, all in just under 4 seconds". In this case, the second condition (within 1.5 seconds) is no longer met but the event can still be considered a very short-term extreme price movement. As a consequence, we do apply the three above-mentioned conditions in our study, but choose a time interval of 2.5 seconds instead of the restrictive 1.5-second time interval advocated by Nanex. As such, we are able to take into account more very short-term extreme price movements than would be the case if we followed a more restrictive definition. Still the time interval we choose falls within the 1.5 to 4-second rule of thumb recommended by Nanex.

We note that MFCs are sometimes called "black swan crashes" when the crash takes place on the way down and "black swan dashes" when the crash takes place on the way up (Johnson et al., 2012). For the sake of our study, we simply refer to them as "down crashes" and "up crashes" respectively, based on the terminology initially chosen by Golub et al. (2012). Depending on how suddenly they occur, MFCs are either refered to as extreme price movements (EPMs) when they occur within 10 seconds (Brogaard et al., 2018) or as ultrafast extreme events (UEEs) when they occur within 1.5 seconds (Johnson et al., 2013, Braun et al., 2017). We use the generic term "mini flash crash" to refer to any extreme price movement occurring within 2.5 seconds.

3 Literature Review on Flash Events

A crash can be defined as (1) a sudden price decline of more than 20% relative to the historical peak for developed markets and of more than 35% for emerging markets (Patel & Sarkar, 1998), that (2) may persist for days, weeks, months or even years. On the contrary, flash events are very short-term in nature (a few minutes for flash crashes; a few seconds only for mini flash crashes) even though they have the potential to be as systemic as traditional crashes as observed both during the flash crash of May 6, 2010 and during the flash crash of August 24, 2015. More importantly, flash events do not result from the publication of new information or data (Johnson et al. 2013; BIS, 2017) even though they may occur at a time of market stress (Kirilenko et al., 2017). We thus make a distinction between flash crashes and mini flash crashes. As pointed out by Johnson et al. (2013), flash crashes are fundamentally different from mini flash crashes since the former last several minutes and thus allow enough time for human involvement, while human involvement is more unlikely in the latter (although the trigger could be human). We first review the two most important flash crashes on the US equity market before studying the MCFs in more detail.

3.1 Two Major Flash Crashes on the US Equity Market

On May 6, 2010, the U.S. markets saw a billion dollars of market capitalization evaporate within minutes. The DJIA suddenly dropped by almost 1000 points before quickly bouncing back to its initial valuation. At the same time, other U.S. indices (S&P500, Nasdaq 100, Russell 2000) also collapsed and many stocks saw their valuation fall by 100% (e.g. Exelon Corp., Eagle Materials, Brown & Brown, Iowa Telecommunications, EnterPoint Energy, Boston Beer, Casey's General Stores, among others), while here and there exuberant spreads resulting from stub quotes appeared in the order books as was the case for Sotheby's (Bid: \$0.01 - Ask: \$100,000).

Kirilenko et al. (2017), in their empirical study of the Flash Crash, give a very precise description of the course of events. At 14:32, while the market was already in a volatile state (Greek crisis) - down 4% from the previous day's closing price - and thus infected by a lack of liquidity, an order to sell 75,000 futures contracts on the E-Mini S&P500 index (valued at \$4.1 billion) was sent by a mutual fund (whose identity - Waddell & Reed - was later revealed), by using an automated algorithm calibrated to execute the sell order at 9% of the volume calculated on the previous minute, in order to allow the fund to hedge its long positions on the U.S. equity market. Although the selling pressure resulting from this sell order - 4% of the total sell orders on that day (Menkveld & Yueshen, 2018) - was initially absorbed by market makers (HFTs and others) E-Mini S&P500 futures started to collapse (-5.1%), sellers being unable to find counterparties and contagion spreading almost instantly to other international indices such as the Canadian indices, until trading on the E-Mini futures was interrupted at 2:45 pm by the CME for five seconds (Menkveld & Yueshen, 2018). Once the market re-opened, prices bounced back, the E-Mini Futures S&P500 contracts canceling out almost all of their losses in 23 minutes (+6,4%). According to Kirilenko et al. (2017), HFTs did not cause the flash crash on that very day but contributed to it by exacerbating liquidity imbalances via a phenomenon well known to seasoned investors, the so-called 'hot potato' effect, which consists in passing the potato while waiting for it to cool down. According to Rzayev and Ibikunle (2017), order aggressiveness was a significant contributor to this flash crash and could even be used as a predictor of flash crashes in general.

A replica of the Flash Crash occurred on August 24, 2015 at the opening of the American trading session. The S&P500 index fell by 7.8% after a few minutes of trading only and generated the activation of circuit breakers on 1278 stocks. Indeed, a few minutes before the opening of the U.S. market, a circuit breaker was activated on the E-mini S&P500 futures after a sharp downward move beyond the 5% limit imposed on E-mini futures contracts. Negotiations on this instrument were interrupted between 9:25 am and 9:30 am, New York time. The S&P500 then opened at 9:30 am, down 5.2% from the closing price of the previous day. The decline continued, reaching 7.8% at 9.35 am. A few minutes later, the market bounced back, erasing much of its losses like during the flash crash of May 6, 2010. The SEC in its report published in December 2015 notes that: (1) stocks representing the largest market capitalizations were the most affected on that day, (2) a ban on short selling was activated on more than 2,000 stocks (representing 37% of the market capitalization of the S&P500 and 50% of the market capitalization of the Nasdaq 100), (3) volumes on the largest market capitalizations were five times greater than the usual average volumes, and (4) market depth was 70% lower than the usual average market depth at this time of the trading session (SEC, 2015).

3.2 Mini Flash Crashes

This brief literature review on MFCs is organized around four important themes in our view. Firstly, there is the study of the statistical properties of MFCs. Our paper relates most to this theme. Second, there is the debate on the danger of using Intermarket Sweep Orders. Third, there is the question of responsibility of machine and human traders in causing MFCs. Finally, there are the pros and cons of market fragmentation.

3.2.1 Statistical Features

Golub and Keane (2011) present statistical properties of MFCs using simple data mining practices on tick-by-tick transaction data. They find no difference between up and down crashes in terms of distribution during the trading day. MFCs happen when liquidity imbalances are the strongest and volatility is the highest, i.e. at the beginning and at the end of the trading day. The finance, insurance and real estate sector, which accounts for 31% of MFCs in their sample, as well as the manufacturing sector, which accounts for 30% of MFCs in their sample, appear to suffer the most. 39% of MFCs (up and down) occur on the NYSE, while 28% occur on both NASDAQ and ARCA.

Johnson et al. (2013) study ultrafast extreme events (UEEs) from January 3rd 2006 to February 3rd 2011 using the Nanex NxCore software package (data are timestamped to the millisecond). The authors find 18,520 MFCs with durations less than 1500 milliseconds. The authors note that ”since both crashes and spikes are typically more than 30 standard deviations larger than the average price movement either side of an event, they are unlikely to have arisen by chance since, in that case, their expected number would be essentially zero whereas we observe 18,520 [MFCs]”. Moreover, the authors argue that ”the fact that the occurrence of spikes and crashes is similar suggests MFCs are unlikely to originate from any regulatory rule that is designed to control market movements in one direction e.g. the uptick regulatory rule for crashes” and that ”their rapid subsecond speed and recovery suggests they are also unlikely to be driven by exogenous news arrival”.

3.2.2 Rule 611 and Intermarket Sweep Orders

Golub et al. (2012) conduct a second investigation on MFCs. They focus on regulation issues in the U.S., focusing more particularly on the so called Order Protection Rule, as known as Rule

611, a rule meant to prevent orders on one exchange from being executed at prices that are inferior to the NBBO that may prevail at another exchange.[2]

To satisfy the concerns of institutions looking for a quick and automatic execution in a specific market center, the SEC carved out an exemption to Rule 611. For institutions that have large marketable limit orders and need to sweep through multiple levels of the order book, they may use Intermarket Sweep Orders (ISOs) which enable them to execute those orders immediately on a chosen trading center, still assuming all liability for compliance with Rule 611 in the sense that all better-priced protected quotations displayed by other trading centers up to their displayed size must be executed as well (SEC, 2004 and 2015). However, there is no obligation to wait for better-priced quotations on other markets to be updated before executing the remaining orders on the chosen trading venue. In practice, the use of ISOs implies that Rule 611 only protects the top of the book of all exchanges and does not actually cover any additional depth-of-book prices that are outside the NBBO. Therefore, large trades can consume liquidity deep in the book of the domestic exchange *despite* more liquidity being still available in the books of other trading venues. The resulting 'local' MFC would primarily be a market structure issue related to the use of ISOs in fragmented markets by fast traders.

With these definitions in mind, Golub et al. (2012) attempt to classify the crashes between ISO-initiated MCFs and auto-routing-initiated MCFs. They estimate that 67.85% MFCs are ISO-initiated, 4.64% MFCs are auto-routing initiated and 27.51% MFCs are unclassified. Despite

[2]Rule 611 applies to all stocks on major U.S. indices and most of the OTC market. In other words, Rule 611 requires all types of trading venues, including registered exchanges, ECNs, dark pools, off-exchange market makers, as well as internalizers, to implement policies and procedures that are designed to prevent "trade-throughs", i.e. executing orders at prices that are less attractive than the NBBO. More precisely, a trade-through is defined as the purchase or sale of stock during regular trading hours (9:30 a.m. to 4:00 p.m. ET) at a price that is lower than a protected bid or higher than a protected ask. In the classical scenario, when there is no exception to Rule 611, marketable limit orders are executed sequentially, being automatically routed (or auto-routed) to the market center that posts the best execution price *which can change during the execution of the order set* (Chakravarty et al., 2012). This implies that if a large limit order is routed to a venue not displaying the NBBO, the venue must avoid the trade-through by either matching the NBBO or rerouting it to other exchanges displaying the NBBO. Then, if the order consumes all quantities available at the current NBBO, there is the obligation to wait for better-priced quotations on other markets to be updated before executing the remaining quantity on the trading venue where the order was initially routed to. The risk related to the use of auto-routed orders is still to chase liquidity that may be withdrawing rapidly, leading in turn to a potential MFC.

the fact that both types of MFCs differ in terms of mechanics, they are both the result of under-protection for depth in the book quotations.[3] They observe that additional liquidity may be available deeper in the order book of other trading venues but because Rule 611 applies only to the top of the book, trades can occur at inferior prices on the 'local' exchange, causing MFCs.

They also discover an association between fleeting liquidity and MFCs. In the case of MFCs, fleeting liquidity is identified when the quotes disseminated by the Securities Information Processor (SIP) are not reached while an MFC occurs.[4] It means the best displayed quotation is cancelled before the SIP can even disseminate the removal of the resting limit orders. According to Golub et al. (2012), 37.99% of the MFCs under scrutiny present signs of fleeting liquidity.

3.2.3 Humans or Robots

Braun et al. (2017) study the impact and recovery process of MFCs using a NYSE dataset of trades and quotes of all stocks of the S&P500 over the period 2007-2008. Using the commonly employed criterion advocated by Nanex (0,8% price change, within 1.5 seconds, with at least 10 tick movements), they find 5529 MFCs over the studied period, with the financial sector being hit the most (33.35 MFCs per company) due to the financial crisis of 2008. More importantly, the authors find that 60% of MFCs contain a large market order that already generates a 0.5% return, which suggests that most of the MFCs are not primarily caused by HFTs but by human traders, even though the analysis is limited by the use of a one-second timestamp precision (instead of a millisecond timestamp for example).

Brogaard et al. (2018) study the behavior of high-frequency traders (HFTs) around extreme price movements (EPMs), which they define as "ten-second returns in the 99.99th percentile according to magnitude" that consist of a series of sequential trades. In particular, they focus on a sample of more than 45,000 extreme price movements extracted from HFT data from NASDAQ

[3]This is less dramatic for non-ISO trades given the obligation to wait for better-priced quotations on other markets to be updated before executing the remaining orders on the trading venue where the order originated from.
[4]The SIP is the institution responsible for determining the NBBO and disseminating it to all its subscribers.

over the period 2008-2009. Using both this definition and the Lee and Mykland's (2012) jump detection methodology, the authors find that HFTs tend to provide liquidity during EPMs thus absorbing imbalances created by non-high frequency traders. They do not find evidence that HFTs are causing EPMs in spite of the fact HFT revenues are higher on days when EPMs occur. However, they note that the liquidity provided by HFTs is limited to EPMs in single stocks so that when several stocks experience simultaneous EPMs, HFTs tend to demand more liquidity than they supply. Moreover, they also admit that the dataset is limited to trades on Nasdaq, which only makes up 30%-40% of the overall trading activity in their sample, leaving a possibility that while providing liquidity on the NASDAQ during EPMs, HFTs may also be demanding liquidity on other exchanges.

Golub et al. (2012) also speculate on who (or what) might be the cause of MFCs. As mentioned before, most of the crashes occur due to an aggressive use of ISOs. In fact, retail and institutional investors cannot employ such mechanisms since only broker-dealers or traders with sponsored access can. Moreover, when looking at the magnitude and speed of MFCs, only one type of traders fits these conditions, i.e. the high-frequency trader. HFTs are the only ones being able to intervene at extremely mispriced levels and benefit from the reversal move during MFCs.

Bellia et al. (2018) come to similar conclusions. She uses a novel econometric methodology developed by Christensen, Oomen and Renó (2018), which enables them to differentiate directional price crashes from episodes of high volatility, study 65 flash crashes detected on 37 liquid French stocks of the NYSE-Euronext Paris market over the year 2013 and find (contrary to most of the existing empirical literature) that HFTs are responsible for initiating the crash in about 70% of the cases. Furthermore, they find that HFTs very often exacerbate the crash at its climax by selling more as the crash unfolds and do not contribute to the recovery phase since they keep filling up the order book with selling orders during this phase (while non-HFTs do provide liquidity thus contributing to the recovery).

3.2.4 Market fragmentation

Félez-Viñas (2018) investigates the impact of market fragmentation on liquidity during episodes of mini flash crashes on stocks from both the Italian FTSEMIB and the Spanish IBEX35. The author identifies MFCs following three requirements. First, the price change has to occur within 2 seconds. Second, the price change has to be of at least 0.8% within the 2-second time interval. Third, the recovery of the crash has to take place within 300 seconds (so as to avoid capturing crashes that are due to fundamental volatility) and the price of the stock has to recover at least 90% of the initial price. Félez-Viñas finds that market fragmentation improves liquidity, reduces the number of mini flash crashes and speeds up the price recovery of the stock when there is a crash. Still, she notes that small stocks appear to be more vulnerable to liquidity shocks so that they benefit the least from market fragmentation.

4 An Empirical Taxonomy

4.1 Data

In order to provide a taxonomy of mini flash crashes, we build a dataset based on the extensive list of MFCs reported by Nanex from May 2011 to September 2014. These flash events are not estimated volatility jumps, for which estimation error is always positive, but price jumps. Our initial dataset includes 332 MFCs as well as the following eleven variables:

- Symbol: ticker symbol of the stock, mutual fund, ETF, index or futures contract;
- Type of crash: up crash or down crash (base category = down crash);
- Date: day of occurrence of the MFC;
- Year: year when the MFC occurred (base category = 2011);

- Time of the day: (1) opening period (from 9:30 a.m. to 9:50 a.m. ET); (2) midday period (from 9:50 a.m. to 3:30 p.m. ET, base category); (3) closing period (from 3:30 p.m. to 4:00 p.m. ET); (4) extended hours ($\leq$9:30 a.m. ET, $\geq$4:00 p.m. ET);

- Crash duration: duration of the MFC in milliseconds (from peak to through for a down crash, or from through to peak for an up crash);

- Dominant exchange: exchange where most of the trades occurred during the crash. The exchanges are: NASDAQ (base category), NYSE, ARCA, NYMX, BATS, and 'OTHERS' (including CBOE, EDGE, ENID, NSX, and PBOT);

- Price at the beginning: last trade price at the beginning of the crash;

- Absolute return: absolute value of the simple arithmetic return of the crash;

- Sector: (1) finance, insurance and real estate (base category); (2) manufacturing; (3) services; (4) others (agriculture, forestry and fishing; mining; construction; transportation, communications and electric gas; wholesale trade; retail trade);

- Asset type: (1) stock (base category); (2) ETF; (3) futures

For example, the price of Nasdaq OMX Group Inc. suddenly crashed 2.72% on April 4, 2014, moving from $35.98 to $35 in 2.3 seconds before coming back to its initial price. The corresponding values encoded in our dataset for this MFC are shown in Table 2.

In this case, *crash duration* is estimated based on the chart provided by Nanex, the crash starting at 20:46:20:300 (peak) and ending at 20:46:22:600 (through), which represents 2300 milliseconds. *Price at the beginning* is directly provided by Nanex and we simply compute the *absolute return* based on the indicated price change provided by Nanex (from $35.98 to $35), which gives an absolute return of 2.72%. The other variables are found in the same way, i.e. either via the chart or via the information directly provided by Nanex. Finally, *sector* and *asset type* are determined via Google Finance and Yahoo Finance.

Table 2: Encoded set of information for a typical MFC

Variable	Encoded information
Symbol	NDAQ
Type of crash	Down
Date	04-04-14
Year	2014
Time of the day	Opening period
Crash duration (ms)	2300
Dominant exchange	Nasdaq
Price at the beginning	\$35.98
Absolute return	2.72%
Asset type	Stock
Sector	Finance, insurance and real estate

The table reports the encoded set of information for a typical mini flash crash. In this case, the encoded values correspond to an MFC that occurred on Nasdaq OMX Group Inc on April 4, 2014.

We eventually remove 140 flash events either because they do not lead to an absolute return greater than 0.8% within a 2.5 second-interval, or because they appear to be outliers.[5] Our final dataset thus includes 192 MFCs, or 57.83% of our initial dataset (Table 3).

Table 3: Number of MFCs per time interval

Interval (seconds)	Number of crashes	Ratio
0-2.5	192	57.83%
2.5-3.0	86	25.90%
3.0-4.5	41	12.35%
>4.5	13	03.92%
Total	332	100%

The table reports the number of mini flash crashes within our sample conditional on the selected time interval.

4.2 Empirical findings

First, we present some descriptive statistics for the following variables: *crash duration*, *price at the beginning* and *absolute return* (Table 5). As implied by our mini flash crash definition,

[5] All the MFCs with an absolute return greater than 100% are deleted and an outlier that is 1000 times higher than all the other values for *price at the beginning* is also removed.

the minimum absolute return is 0.8% and the maximum crash duration is 2500 milliseconds. We note that the mean absolute return is 10.41%, for a mean crash duration of 810 milliseconds and a mean price at the beginning of $117.77.

Table 4: Descriptive statistics

	N	Min	Max	Mean	Variance	St.Dev.
Absolute return	184	0.8%	99.99%	10.41%	371.744	19.28%
Crash duration (ms)	192	1,000	2,500	810	461,787.58	679.55
Price at the beginning ($)	191	3.2	9,194	117.77	748,371.76	865.08

The table reports descriptive statistics regarding the absolute return, crash duration and price at the beginning variables.

We then perform an ANOVA to evaluate the variation in the mean value of the variable of interest (e.g. variable A) when another variable (e.g. variable B) falls into different categories (or groups).

We first use *absolute return* as our variable of interest (i.e. variable A). Categories are defined according to three other variables (i.e. variables B): *time of the day*, *type of crash*, and *asset type*. As indicated in Table 4, the (unconditional) absolute return is 10.41% on average. In Table 5, we show the conditional mean values for the *absolute return* variable. For example, we find that down crashes are much more severe than up crashes, leading to an absolute return of 13.18% for down crashes to be compared to an absolute return of 5.57% for up crashes. We also observe a strong difference between crashes that occur in the opening period and crashes that occur during the midday period.

We perform a test of homogeneity of variances to determine the type of test we need for the difference in means. For *type of crash* (2 categories) and *time of the day* (4 categories), the hypothesis of variance equality is rejected so that we perform a test of means using the Welch test. We find that the mean value of *absolute return* significantly differs according to the type of crash and the time of the day, with p-values below 5% (indicated in italics in Table 5). We find a significant difference in the mean values for the following two groups “opening period” and “midday period”. As a result, we can affirm that these two groups, as well the “up” and “down”

Table 5: **Conditional mean values of absolute return**

Type of crash	N	Absolute return (mean)
Down	117	*13.18%*
Up	67	*5.57%*
Time of the day	N	Absolute return (mean)
Opening period	37	*17.84%*
Midday period	99	*5.92%*
Closing period	27	12.40%
Extended hours	21	15.94%
Asset type	N	Absolute return (mean)
Stock	156	11.26%
ETF	17	7.99%
Future	11	2.02%

The table reports the conditional mean values for the absolute return variable.

groups differ highly in terms of absolute return. For *asset type* (3 categories), the null of equality of mean cannot be rejected so that the mean value of *absolute return* does not significantly differ according to the asset type.

Second, we repeat the same analysis for *crash duration*. The three conditioning variables are identical: *time of the day* , *type of crash*, and *asset type*. As indicated in Table 4, *crash duration* is 810 milliseconds on average. Table 6 shows the conditional mean values of *crash duration*.

Table 6: **Conditional mean values of crash duration**

Type of crash	N	Crash duration (mean)
Down	117	716.41
Up	75	956.73
Time of the day	N	Crash duration (mean)
Opening period	41	956.05
Midday period	100	717.68
Closing period	27	857.37
Extended hours	24	894.17
Asset type	N	Crash duration (mean)
Stock	163	837.24
ETF	18	737.22
Future	11	530.46

The table reports the conditional mean values for the crash duration variable.

The null hypothesis of homogeneity of variances cannot be rejected, whatever the conditioning variables. When the 'up' and 'down' categories are used, we reject the null hypothesis of equality of means at 5%. Therefore, down crashes are significantly shorter than up crashes. However, we find no statistical difference in *crash duration* means when other categories (relative to the time of the day and the asset type) are used.

We also conduct a multiple linear regression with the objective of determining the ceteris paribus effects of the explanatory variables on the dependent variable. The dependent variable is *absolute return* and *crash duration* in turn.

Table 7: Initial multiple linear regression for the absolute return variable

Variable	Coef	Std.Error	T-Stat	Prob
C	0.153985	0.060134	2.560701	0.0114
Price at the beginning	-8.76E-06	1.04E-05	-0.838834	0.4028
Type of crash (up)	-0.046547	0.020961	-2.220663	0.0278
Opening period	0.079318	0.041834	1.896028	0.0597
Closing period	0.036292	0.038291	0.947798	0.3446
Extended hours	0.119845	0.073022	1.641216	0.1027
ETF	-0.025614	0.060390	-0.424136	0.6720
Future	0.040929	0.114176	0.358471	0.7205
Crash duration	-1.90E-05	1.75E-05	-1.086366	0.2789
NYSE	0.042702	0.061643	0.692735	0.4895
ARCA	-0.035639	0.023894	-1.491581	0.1377
BATS	-0.051656	0.025765	-2.004903	0.0466
NYMEX	-0.175330	0.102579	-1.709222	0.0893
Other exchange	-0.085371	0.056113	-1.521414	0.13001
Services	-0.021877	0.049855	-0.438814	0.6614
Manufacturing	-0.019508	0.036539	-0.533910	0.5941
Other sectors	-0.003355	0.043539	-0.077056	0.9387
Year2012	-0.059323	0.048553	-1.221813	0.2235
Year2013	0.004851	0.060569	0.080089	0.9363
Year2014	-0.090626	0.051055	-1.775080	0.0778
R-squared	0.161771	Mean dependant var		0.099726
Adjusted R-squared	0.064064	S.D. dependant var		0.184029
S.E. of regression	0.178037	Akaike info criterion		-0.510811
Sum squared resid	5.166619	Schwarz criterion		-0.160048
Log likelihood	66.73923	Hannan-Quinn criterion		-0.368630
F-statistic	1.655666	Durbin-Watson stat		1.666915
Prob (F-statistic)	0.048989	Wald F-statistic		1.070213
Prob (Wald F-statistic)	0.386109			
Included observations: 184				
White heteroskedasticity-consistent standard errors & covariance				

The table reports the results of the initial multiple linear regression for the absolute return variable.

Table 8: Final multiple linear regression for the absolute return variable

Variable	Coef	Std.Error	T-Stat	Prob
C	0.180312	0.033693	5.351700	0.0000
Type of crash (Up)	-0.072810	0.022480	-3.238878	0.0014
ARCA	-0.058596	0.022275	-2.630644	0.0093
BATS	-0.089686	0.024594	-3.646612	0.0006
NYMEX	-0.092273	0.026497	-3.482378	0.0006
Year2012	-0.064097	0.031468	-2.036879	0.0432
Year2014	-0.104932	0.027520	-3.812945	0.0002
R-squared	0.093093	Mean dependant var		0.104083
Adjusted R-squared	0.062350	S.D. dependant var		0.192807
S.E. of regression	0.186699	Akaike info criterion		-0.481335
Sum squared resid	6.169618	Schwarz criterion		-0.359028
Log likelihood	51.28281	Hannan-Quinn criterion		-0.431762
F-statistic	3.028142	Durbin-Watson stat		1.610600
Prob (F-statistic)	0.007645	Wald F-statistic		3.812078
Prob (Wald F-statistic)	0.001344			
Included observations: 184				
White heteroskedasticity-consistent standard errors & covariance				

The table reports the results of the final multiple linear regression after a cleaning step-by-step procedure for the absolute return variable.

In Table 7, the constant (i.e. 15.40%) measures the average absolute return for down crashes occurring during the midday period in 2011 on stocks included in the 'finance, insurance and real estate' sector and for which Nasdaq is the dominant exchange. We see that the two most significant independent variables are the "type of crash (up)" and "BATS" dummies, with p-values under 5%. So, when an up crash occurs (instead of a down crash), the absolute return is expected to decrease by 4.65 percentage points on average, all else equal. Now, if down crashes occur on BATS (instead of NASDAQ), the absolute return is reduced by 5.17 percentage points on average, all else equal. If we raise the level of significance at 10%, we can identify three additional explanatory dummies: 'opening period', 'Nymex', and 'Year2014'. For example, when down crashes occur during the opening period (instead of the midday period), the absolute return is estimated to increase by 7.93 percentage points. Although the model has a reasonable adjusted R^2 (given the nature of the dependent variable), the p-value of the F-test for global significance is only slightly below 5%, justifying a cleaning step-by-step procedure. We follow a classical 'general-to-specific' approach by excluding all the variables with a p-value over 5%, starting with the most insignificant variable and using White's heteroskedasticity-consistent standard errors and covariances at each step. The final regression is given in Table 8.

If we compare the final model to the initial one (which was potentially more affected by multicollinearity issues), we now have six independent dummies possessing a significant impact on the dependent variable at the 5% level. We can see that the occurrence of an up crash (instead of a down crash) reduces the absolute return by 7.28 percentage points on average, all else equal. Dummies characterizing the 'dominant exchange' play an important role in the final equation. All else equal, we note that NASDAQ is an exchange subject to more severe crashes on average than ARCA, BATS and NYMEX, since the three coefficients are negative. For example, the absolute return is reduced by 5.86 percentage points on average when the dominant exchange is ARCA instead of NASDAQ, ceteris paribus. The fact that the coefficient of the NYMEX dummy is significant and negative, points to the greater vulnerability of equities relative to futures contracts (given that most futures contracts are traded on the NYMEX in our sample). Interestingly, we do not find any statistical difference between the NASDAQ and the NYSE. Finally, we identify some cyclicality in the magnitude of crashes. Although there is no statistical difference between 2011 and 2013, we find that the absolute return of crashes on average is significantly lower in 2012 and 2014.

Regarding the multiple linear regression for *crash duration* (Table 9), the constant is equal to 801 milliseconds, corresponding to the average duration of down crashes occurring around midday in 2011 on stocks included in the 'finance, insurance and real estate' sector and for which the dominant exchange is NASDAQ. We also find two explanatory variables which have a significant ceteris paribus effect at the 5% level. This is the case for the *type of crash (up)* and *opening period* dummies. In comparison to down crashes, crashes on the upside lengthen the crash duration by 231 milliseconds on average, all else equal. In addition, when the crash occurs during the opening period instead of the midday period its duration increases by 325 milliseconds on average, all else being held constant. Again, given the relatively high p-value of the F-test for global significance, a ‘general-tospecific’ step-by-step procedure is justified.

In Table 10, we report the results for the final regression, after excluding all the variables with a p-value over 5% each in turn, starting with the most insignificant variable and applying the White correction on the standard errors. There are five significant variables left in the final equation at the 5% level. Again, we find the *type of crash (up)* and *opening period* dummies, albeit with lower estimated coefficients. There are also two dummies characterizing the dominant exchange. We see that the average duration of a crash is estimated to be 384 and 348 milliseconds (significantly) shorter on the NYSE and ARCA, respectively (and relatively to NASDAQ). Finally, the estimated crash duration for stocks belonging to the manufacturing sector is estimated to be 267 milliseconds (significantly) longer than for stocks included in the ’finance, insurance and real estate’ sector.

Overall, we detect significant differences between crashes and exchanges in terms of absolute return and duration. For example, we show that the MFCs on ARCA are both less severe and shorter in duration than those on the NASDAQ, and that ’down crashes’ exhibit lower absolute returns but have longer duration in comparison to ’up crashes’. There is a need for future research to explain these differences.

5 Implications for Public Policy

Flash events, whether flash crashes or mini flash crashes, have the power to undermine the integrity of securities markets. On the one hand, they tend to have a negative impact on investors’ confidence, and on the other hand, they may have a negative impact on investors’ returns. As a consequence, the increasing number of flash events raises questions in terms of public policy implication. How can such erratic price moves be dealt with from a regulatory perspective?

One possible response is to adapt circuit breakers and price limits to high-speed electronic

Table 9: Initial multiple linear regression for the crash duration variable

Variable	Coef	Std.Error	T-Stat	Prob
C	801.4888	156.0229	5.136994	0.0000
Price at the beginning	0.036263	0.040216	0.901716	0.3685
Type of crash (Up)	231.4333	110.0989	2.102049	0.0371
Opening period	324.7107	140.5273	2.310659	0.0221
Closing period	127.5788	141.9809	0.898564	0.3702
Extended hours	272.3490	163.3289	1.667488	0.0973
ETF	-88.32350	199.2111	-0.443366	0.6581
FUTURE	-724.1287	384.8497	-1.881588	0.0617
NYSE	-325.2592	175.5895	-1.852384	0.0658
ARCA	-283.1265	147.1116	-1.924569	0.0560
BATS	199.6656	135.2522	1.476246	0.1418
NYMEX	499.6648	406.9663	1.227779	0.2213
Other exchanges	87.01089	265.1317	0.328180	0.7432
Services	-52.74024	149.5805	-0.352588	0.7249
Manufacturing	213.5520	148.5443	1.437632	0.1525
Other sectors	-82.93187	155.7803	-0.532364	0.5952
Price change	-251.8798	224.7958	-1.120483	0.2642
Year2012	-103.5345	139.8840	-0.740145	0.4603
Year2013	-236.5515	131.7061	-1.796055	0.0743
Year2014	-319.2552	200.8967	-1.589151	0.1140
R-squared	0.151247	Mean dependant var		804.8907
Adjusted R-squared	0.052312	S.D. dependant var		665.7144
S.E. of regression	0.648.0680	Akaike info criterion		15.88871
Sum squared resid	68458722	Schwarz criterion		16.23948
Log likelihood	-1433.817	Hannan-Quinn criterion		16.03089
F-statistic	1.528756	Durbin-Watson stat		2.187358
Prob (F-statistic)	0.081704	Wald F-statistic		3.194681
Prob (Wald F-statistic)	0.000031			
Included observations: 184				
White heteroskedasticity-consistent standard errors & covariance				

The table reports the results of the initial multiple linear regression for the crash duration variable.

markets so as to better manage events of both exceptional volatility and speed. Introduced in the U.S. in 1988 after the 1987 stock market crash in order to protect markets from periods of extreme illiquidity (Brugler & Linton, 2014), circuit breakers have recently been coordinated between U.S. equity, options and futures exchanges[6] and tightened by the SEC following the U.S. Flash Crash of May, 2010. Whereas circuit breakers were previously triggered at the following three thresholds: 10% (Level 1), 20% (Level 2), and 30% (Level 3), based on the average closing

[6] As such, circuit breakers have been renamed "market-wide circuit breakers".

Table 10: Final multiple linear regression for the crash duration variable

Variable	Coef	Std.Error	T-Stat	Prob
C	659.7187	71.67559	9.204231	0.0000
Type of crash (Up)	214.4298	103.9816	2.062190	0.0406
Opening period	245.0054	122.2055	2.004864	0.0464
NYSE	-383.8690	161.0279	-2.383867	0.0181
ARCA	-348.3551	124.3143	-2.802213	0.0056
Manufacturing	267.2166	113.8485	2.347125	0.0200
R-squared	0.098939	Mean dependant var		810.2865
Adjusted R-squared	0.074717	S.D. dependant var		679.5495
S.E. of regression	653.6698	Akaike info criterion		15.83383
Sum squared resid	79474867	Schwarz criterion		15.93563
Log likelihood	-1514.048	Hannan-Quinn criterion		15.87506
F-statistic	4.084663	Durbin-Watson stat		2.105901
Prob (F-statistic)	0.001529	Wald F-statistic		4.751536
Prob (Wald F-statistic)	0.000414			
Included observations: 192				
White heteroskedasticity-consistent standard errors & covariance				

The table reports the results of the final multiple linear regression after a cleaning step-by-step procedure for the crash duration variable.

value of the DJIA for the month prior to the beginning of the quarter, they are now triggered at the following three (tighter) thresholds: 7% (Level 1), 13% (Level 2), and 20% (Level 3), based on the closing value of the S&P500 index for the prior day. In particular, in the case Level 1 or Level 2 circuit breakers are triggered *before* 3:25 p.m., ET, trading shall halt for 15 minutes. However, in the case Level 1 or Level 2 circuit breakers are triggered *at* or *after* 3:25 p.m., ET, trading shall not halt until the close, unless Level 3 circuit breaker is triggered. If Level 3 circuit breaker is triggered (whatever the time at which it is triggered during the trading day), trading shall halt and not resume for the rest of the day (SEC, 2012). The lack of homogeneous circuit breakers among trading venues around the world[7] may however be an issue at a time of high market fragmentation and high algorithm participation. Moreover, circuit breakers have the potential to slow down the price adjustment process and negatively affect resilience, which is an important aspect of the recovery process during flash events.

[7] Even though many exchanges around the world have introduced circuit breakers, some exchanges still do not use circuit breakers.

Adapting price limits, i.e. when trading is not permitted at any price above or below a predefined level for a predefined period of time (Brugler & Linton, 2014), is another way to better manage disorderly high-speed markets. In the U.S., price limits have been operational for several decades but the SEC approved, on a pilot basis, a new type of price limit called the "limit up-limit down mechanism" (LULD) in 2012 to address exceptional market volatility by preventing trades in listed equity securities (U.S. stock markets, stock option markets, single-stock futures markets) suffering large and sudden price moves during the following trading hours: from 9:30 a.m. to 4:00 p.m.[8], ET (SEC, 2012). This mechanism replaced former single-stock circuit breakers and starting in 2013, trades in individual equity securities have been prevented from occurring outside of a specified price band[9] set at a percentage level above and below the reference price[10] of the security over the last five minutes. In the case the security does not move back within the price bands within 15 seconds, there shall be a five-minute trading halt (SEC, 2012). The percentage parameters taken into account are the following: for stocks whose price is less than \$0.75 the percentage parameter is the lesser of \$0.15 or 75% for tier 1 and tier 2 securities.[11] For stocks whose price is comprised between \$0.75 and \$3 the percentage parameter is 20% for tier 1 and tier 2 securities. As for stocks whose price is greater than \$3 the percentage parameter is 5% for tier 1 securities and 10% for tier 2 securities. While this mechanism may be an adequate answer vis-à-vis flash crashes (this mechanism triggered more than 1200 5-minute trading halts during the U.S. flash crash of August 24, 2015 for example), it seems that this mechanism does not solve the issue of mini flash crashes, whose moves are shorter than 15 seconds. A mini flash crash of the same magnitude as the one endured by Progress Energy in 2010 (-90% on the downside; +1000% on the upside), remains a possibility

[8]These price bands are doubled during the opening and closing periods of the trading day.

[9]The price band is calculated in the following way: reference price ± (reference price x percentage parameter)

[10]The reference price is the arithmetic mean price of eligible reported transactions over the last five minutes. In case no transaction occured in the last five minutes, the reference price is the previous reference price.

[11]Tier 1 securities are all the securities belonging to the S&P500 index, the Russell 1000 index as well as some ETPs while tier 2 securities are all the other NMS securities. Rights and warrants are excluded from this mechanism.

under this mechanism as the stock price would come back within the specified price bands just before transactions in the stock be halted (the MFC duration being inferior to 15 seconds).

Another possible response is to introduce latency delays, also known as speed bumps. Such intentional slowing of order flow by exchanges would protect market makers from high-frequency arbitrage and therefore improve liquidity for uninformed investors via narrower spreads. Opponents nevertheless claim that the liquidity improvement is illusory because the ’improved´ quotes may fade before they are hit. The International Exchange (IEX) in the U.S. has increased the delay between trade decision and trade execution by 350 microseconds (or 0.35 millisecond) via the use of a speed bump on messages to and from its platform (thanks to a 38-mile coil of optical fiber placed near the trading engine). The objective of such a mechanism is to block HFTs specialized in latency arbitrage, whose strategy consists in exploiting latency gaps between the 13 U.S. trading platforms by calculating the National Best Bid and Offer (NBBO) a few microseconds before the exchanges themselves thus enabling them to profit from a single speed advantage. However, such an alternative has been challenged by some experts who believe the introduction of speed bumps, originally meant to make trading more fair for slow investors, may create a new kind of unfairness, while some practitioners (such as Citadel) even believe this could result in more market manipulation. In particular, Bell 2016 explains that due to the fact the SEC was forced to reinterpret a key provision of Reg NMS to approve IEX’s speed bump, declaring a delay inferior to one millisecond to be *de minimis*, i.e. insignificant to investors, many other exchanges (NYSE, Nasdaq, Chicago Stock Exchange, among others) have decided to develop new order types and add intentional delays so as to benefit and help other sophisticated traders benefit at the expense of the ordinary investor (Bell, 2016). Chen et al. (2017) are also critical: they find that disappearing offers on the TSX Alpha Exchange in Canada rose to 60 percent from 14 percent after the speed bump was introduced; and the speed bumps increase profits for liquidity providers on TSX Alpha but negatively impact aggregate liquidity. In Brolley and Cimon (2017)’s model on the contrary, latency delays lead to positive liquidity effects on both delayed and non-delayed exchanges provided the delays are sufficiently 'long´. In any case, the distribution of benefits among the different stakeholders, before and after the introduction of speed bumps, needs to be studied further.

Another potential policy response is to move from continuous-time trading to discrete-time trading, thanks to exchanges in which the trading day would be divided into extremely frequent equal-length time intervals (e.g.100 milliseconds) as suggested by Budish et al. (2015). The authors in their study note that correlations cease to exist at high-frequency time scales thus leading to arbitrage opportunities that can only be exploited by the fastest traders (i.e. HFTs) and thus exacerbating an unproductive arms race. To address this problem, Budish et al. suggest to develop a batch auction market where (1) time would be discrete and (2) orders would be processed in batch via a uniform-price auction. At the end of each batch interval, the market would clear where demand meets supply so that all transactions would occur at the same price and if demand does not meet supply then no trade would occur and all orders would remain outstanding for the next batch auction. Two orders sent during the same batch interval would have the same time priority while an order outstanding for a larger number of batch auctions would have time priority over a more recent order. As emphasized by Budish et al. the use of discrete-time trading would offer the following advantages: 1) it would reduce the value of "tiny speed advantage", 2) the use of batch auctions would eliminate the possibility for HFTs to snipe stale quotes in the order book and 3) it would be computationally more efficient than the continuous limit order book market design thus preventing processing delays during a surge of activity, which can potentially increase market stress. Even though we are unable to conclude at this stage that a batch auction market would prevent MFCs to occur, we do believe such a market could reduce the number of extreme price movements since the crashes with the minimum duration would de facto be eliminated. A question remains though. Would HFTs keep participating in such a market?

The use of ISOs has also been blamed for negatively contributing during flash events (Golub et al., 2012, McInnish et al., 2014). As discussed in Section 3.2, ISOs enable traders to access liquidity posted outside of the best prices by submitting the sufficient volume on all trading venues. They are faster than non-ISOs since there is no price check for a trade-through prior to the order execution as well as no rerouting delays, they are not impacted by changes in the market state during processing (ISOs are based on the state of the market at order submission) and they can be processed on all trading venues simultaneously (McInish et al., 2014). As stated by McInish et al. (2014), the problem of this type of orders is that they "can quickly deplete liquidity and lead to rapid price changes". In particular, McInish et al. (2014) show that the use of ISOs was substantially higher during the Flash Crash of May 6, 2010, ISOs representing over 65% of the sell volume for stocks that fell the most during the flash crash and 53% of the buy volume for these same stocks during the recovery phase. While this type of orders enables traders to access liquidity in a timely manner, it has also the power to 'walk the book', especially during periods of high stress and low market depth, thus depleting the whole posted liquidity until the full size of the order is filled or until the limit price is reached. If the limit price is set at \$0.01 or \$100,000, then these extreme values can be reached as was the case during the Flash Crash.[12] Arguably, Rule 611 of Regulation NMS needs to be analyzed further by academics and regulators and possibly reshaped so as to prevent flash depletion of order books during disorderly markets.

Last but not least, the Markets in Financial Instruments Directive II (MiFID II), which has been implemented since January 2018 in Europe, includes new provisions to deal with the recent evolution of algorithmic and high-frequency trading. First, firms using automated trading strategies[13] are required to make an important effort in terms of transparency. In particular, they will have to notify the regulatory authorities and the trading venue(s) where they engage

[12] ACN, ACOM, BGS, BRO, CASY, CNP, EXC, EXP, G, ITC, IWA, LEA, RDN, SAM, VRGY traded at \$0.01, whereas AAPL, AMLN, BID and AQIX traded at \$100,000 during the Flash Crash.

[13] These strategies are defined by ESMA as "trading in financial instruments where a computer algorithm automatically determines individual parameters of orders such as whether to initiate the order, the timing, price or quantity of the order or how to manage the order after its submission, with limited or no human intervention"

in algorithmic trading of their presence, and may be required to provide details of their trading strategies, including trading parameters, trading limits, key compliance and risk controls, as well as the details of the testing of their systems, which all imply the storage of a trading record. HFT market-makers in particular will be required to provide liquidity on a permanent basis during trading hours, except in exceptional circumstances, putting in place effective systems and risk controls to prevent any contribution to a disorderly market. As for broker-dealers providing HFT firms with a direct access to markets, they will have to monitor automated traders so as to identify infringements of MiFID II rules, disorderly conditions or potential market abuse and prevent algorithmic traders from exceeding pre-set trading and credit thresholds. Last but not least, the minimum tick size, i.e. the smallest possible increment when trading a security, will be harmonized in Europe so as to put an end to ever smaller tick sizes, which tends to benefit the fastest traders[14] (MiFID II, 2014). The implementation of MiFID II is under way but it remains to be seen whether these new rules will contribute to market stability significantly. In this respect, a comparison between the European Union and the US would be welcome.

6 Conclusion

By studying mini flash crashes (MFCs) on the U.S. markets over a three-year period from May 2011 to September 2014 (for stocks, ETFs, and futures), we provide a taxonomy of MFCs, i.e. describing, naming and organizing MFCs into groups that share similar qualities. We show that this taxonomy is useful to better understand both what MFCs are and how the mechanism behind them functions. In order to introduce a taxonomy of MFCs, we identify several dimensions that, we believe, do matter for classification.

First, our statistical analysis has demonstrated that "up crashes" need to be distinguished from "down crashes", both in terms of their expected absolute return and duration. Even though

[14] In order to attract liquidity, exchanges have been fighting one another by tightening their tick size. To stop this race to the bottom, MiFID II offers to apply minimum tick sizes common to all European equities.

up and down crashes share similar characteristics, down crashes are significantly shorter (by 240 milliseconds on average) and more severe (by 7.6 percentage points) than up crashes. This points to significant differences in terms of dynamics that need to be taken into account.

Second, MFCs differ according to the asset type, i.e. stocks, ETFs and futures contracts. Although we do not identify a different dynamic in terms of the absolute return and duration of an MFC when the traded asset is an ETF or a stock, we find that equities are potentially more subject to MFCs than futures contracts, since the absolute return during a MFC for futures contracts (traded on NYMX) is estimated to be significantly smaller relatively to stocks. More incidentally, we also find that low-priced stocks are more subject to MFCs than high-priced stocks.

Third, the time of the day also matters to understand the dynamics of MFCs. We identify four groups: the opening period (from 9:30 a.m. to 9:50 a.m. ET), the midday period (from 9:50 a.m. to 3:30 p.m. ET), the closing period (from 3:30 p.m. to 4:00 p.m. ET), as well as the extended hours (<9:30 a.m. ET and >4:00 p.m. ET). The different periods show great differences in terms of crash frequencies. We find that most MFCs in the equities market occur both at the open and at the close. On the futures market, we nevertheless find that MFCs seem to prevail during both the midday period and during the extended hours, justifying the need to account for the so called 'Form-T' trades. Our econometric analysis also shows that the duration of MFCs are significantly longer during the "opening period" relatively to the "midday period", all else equal.

Fourth, we show that the three most impacted sectors in terms of MFCs are the finance, insurance and real estate sector, the manufacturing sector, and the services sector, which tends to show that the three biggest sectors of the U.S. economy are all affected by MFCs. In the econometric analysis, we also show that the duration of MCFs in the financial sector is found to be significantly shorter than the duration of MCFs in the manufacturing sector, all else equal.

Finally, we find that NASDAQ is the exchange where most of the MFCs occur on the U.S.

equities market, while NYMEX is the exchange where most MFCs occur on the U.S. futures market. Indeed, we find that NASDAQ accounts for 72% of all the MFCs on U.S. equities in our sample, against 73% of the crashes on futures for the NYMEX exchange. Our econometric analysis also confirms that the identification of the exchange on which the asset is mostly traded, i.e. the so-called 'dominant exchange', matters when it comes to better estimating the absolute return and the duration of MFCs. For example, we estimate that the MFCs on ARCA are both less severe and shorter in duration, all else equal, than those on the NASDAQ.

There is a need in future research to cross-check these results by using an international database of MFCs and to include the information about the types of trades which are the most widespread during such events, in particular the ISOs. Most importantly, very little is known yet on the exact role of high-frequency traders before, during, and after MFCs. It is true that positive relationships between high-frequency trading and market liquidity have been pointed out in several papers. For example, Hendershott et al. (2011) as well as Hasbrouck and Saar (2013) show that the increased activity of high-frequency traders is correlated with a decrease in the bid-offer spread. Hasbrouck and Saar (2013) also demonstrate that high-frequency traders improve market depth. However, Hendershott et al. (2011) as well as Gresse (2017) show negative results when considering the relationship between high-frequency trading and market depth. Moreover, we are convinced that a detailed study of MFCs (based on order book data) would be very helpful in determining whether high-frequency trading improves market efficiency on average at the expense of bigger fat-tail risks.

References

Aitken, M., D. Cumming, and F. Zhan (2015). High frequency trading and end-of-day price dislocation. *Journal of Banking & Finance 59*, 330–349.

Bell, H. A. (2016). Who benefits from market speed bumps? the exchanges. *The New York*

Times.

Bellia, M., K. Christensen, A. Kolokolov, L. Pelizzon, and R. Renò (2018). High-frequency trading during flash crashes: Walk of fame or hall of shame? In *Paris December 2018 Finance Meeting EUROFIDAI-AFFI*.

BIS (2017). The sterling 'flash event' of 7 october 2016.

Bowley, G. (2010). The flash crash, in miniature. *New York Times Nov. 8*.

Braun, T., J. A. Fiegen, D. C. Wagner, S. M. Krause, and T. Guhr (2017). Impact and recovery process of mini flash crashes: An empirical study. *arXiv preprint arXiv:1707.05580*.

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

Brolley, M. and D. Cimon (2017). Order flow segmentation, liquidity and price discovery: The role of latency delays. *Available at SSRN 3005738*.

Brugler, J. and O. Linton (2014). Circuit breakers on the london stock exchange: do they improve subsequent market quality? Technical report.

Budish, E., P. Cramton, and J. Shim (2015). The high-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics 130*(4), 1547–1621.

Chakravarty, S., P. Jain, J. Upson, and R. Wood (2012, 6). Clean sweep: Informed trading through intermarket sweep orders. *Journal of Financial and Quantitative Analysis 47*, 415–435.

Chen, H., S. Foley, M. A. Goldstein, and T. Ruf (2017). The value of a millisecond: Harnessing information in fast, fragmented markets. *Available at SSRN 2860359*.

Christensen, K., R. C. Oomen, and R. Renò (2018). The drift burst hypothesis. *Available at SSRN 2842535*.

Félez-Vinas, E. (2018). Market fragmentation, mini flash crashes and liquidity. Working paper presented at the FMA European Conference, Kristiansand, Norway.

Golub, A. and J. Keane (2011). Mini flash crashes. *Working Paper*.

Golub, A., J. Keane, and S.-H. Poon (2012). High frequency trading and mini flash crashes. *Working Paper*.

Gresse, C. (2017). Effects of lit and dark market fragmentation on liquidity. *Journal of Financial Markets 35*, 1–20.

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

Hendershott, T., C. M. Jones, and A. J. Menkveld (2011). Does algorithmic trading improve liquidity? *The Journal of Finance 66*(1), 1–33.

Johnson, N., G. Zhao, E. Hunsader, J. Meng, A. Ravindar, S. Carran, and B. Tivnan (2012). Financial black swans driven by ultrafast machine ecology. *Working Paper*.

Johnson, N., G. Zhao, E. Hunsader, H. Qi, N. Johnson, J. Meng, and B. Tivnan (2013). Abrupt rise of new machine ecology beyond human response time. *Scientific reports (Nature.com) 3*(2627), 1–7.

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

Lee, S. S. and P. A. Mykland (2012). Jumps in equilibrium prices and market microstructure noise. *Journal of Econometrics 168*(2), 396–406.

McInish, T., J. Upson, and R. A. Wood (2014). The flash crash: trading aggressiveness, liquidity supply, and the impact of intermarket sweep orders. *Financial Review 49*(3), 481–509.

Menkveld, A. J. and B. Z. Yueshen (2018). The flash crash: A cautionary tale about highly fragmented markets. *Forthcoming in Management Science*.

MiFID, I. (2014). Directive 2014/65/ec of the european parliament and of the council.

Nanex (2010). Flash equity failures in 2006, 2007, 2008, 2009, 2010, and 2011. Technical report, Nanex, LLC. http://www.nanex.net/FlashCrashEquities/FlashCrashAnalysis$_{E}$*quities.html*.

Nanex (2011). Definition of a "mini-flash-crash". Technical report, Nanex, LLC. http://www.nanex.net/Research/MiniFlashCrashDef/MiniFlashCrashDef.html.

Rzayev, K. and G. Ibikunle (2017). Order aggressiveness and flash crashes. *Available at SSRN 3023230*.

Sarkar, A. and S. Patel (1998). Stock market crises in developed and emerging markets.

SEC (2012). Release no. 34-67091.

SEC (2015). Equity market volatility on august 24, 2015, technical report.

Sornette, D. and S. Becke (2011). Crashes and high frequency trading. Technical report, ETH working paper.