Financial Conduct Authority

Research Note

June 2018

How do participants behave during Flash events? Evidence from the UK equity market.

Matteo Aquilina, Brian Eyles, Jia Shao and Carla Ysusi

How do participants behave during Flash events? Evidence from the UK equity market.

Matteo Aquilina, Brian Eyles, Jia Shao and Carla Ysusi[1]

1. Introduction and main findings

There have been a number of flash events - quick and extreme price movements that recover in a few minutes - in financial markets around the world in recent years. Market participants and regulators have started to pay attention to these episodes and their potential impact on systemic risk and market resilience. Previous studies have analysed extreme individual price shocks. Easley et al (2011), Kirilenko et al (2017) and Menkveld and Yueshen (2016) focused on the May 2010 Flash Crash, Breedon et al (2018) on the 2015 Swiss franc cap removal and, Noss et al (2017) and Schroeder et al (2018) on the 2016 Sterling flash crash. As a contribution to the ongoing discussion, we use rich data from the UK secondary equity market to study large, but not extreme, price shocks, which we refer to as 'mini flash-crashes/rallies'.[2] Mini flash-crashes/rallies are significantly smaller than the 2010 Flash Crash or the 2016 Sterling flash crash but they happen more often. Using consolidated order-book data on FTSE350 stocks traded in the major UK trading venues, we identify extreme events for the 18 months between January 2014 and June 2015. We describe the liquidity around these stress periods and assess the role of different types of market participants.

On average, we find that hybrid firms appear to drive the extreme price movement by trading aggressively in the direction of the price change.[3] Hybrid firms are those firms that mainly provide agency trading services but that may use a similar technology to high frequency traders (HFTs) or provide direct electronic access (DEA) to HFTs; they are mostly large investment banks. HFTs lean against the wind initially, trading against the direction of the initial price movement. They subsequently tend to follow and exacerbate the price change, but are responsible for a substantially smaller amount of the aggressive trading. Both hybrid firms and HFTs keep providing liquidity (by submitting passive orders they add liquidity to the market, see footnote 3 for further explanation) and do not withdraw completely from the order book. However, their liquidity is consumed through trades more quickly than it is replenished with new orders. Furthermore, both participant types (but predominantly HFTs) move their existing orders in the order book to levels with less favourable prices, which amplifies the price movement. By the time the price starts to revert, the depth of the order book has dropped significantly. In most cases, HFTs take longer than hybrid firms to restore the level of liquidity they provide at best available prices (the top of the order book) to typical pre-event levels.

---

[1] Matteo Aquilina and Carla Ysusi work in the Economics Department of the FCA. Brian Eyles works in the Secondary Market Oversight Department of the FCA. This research was carried out while Jia Shao was working at the FCA.

[2] We do not analyse those sudden sharp movements in prices that triggered a circuit breaker on the London Stock Exchange (LSE); these episodes have been analysed in Bercich and Allan (2017).

[3] Participants are said to trade aggressively if they submit a new order (or modify an existing one) resulting in an execution against an order that is already present in the order book. The aggressive order removes liquidity from the order book. Instead, a passive order adds liquidity to the market; a passive order when entered will not result in a trade as it is priced below (above) the best offer (bid) price for a bid (offer). Here we say participants trade passively if their passive orders resting in the order book execute against an incoming aggressive order.

This is especially true for liquid stocks. Unlike hybrid firms that make a loss, HFTs on average make small profits from these events. This is probably due to their ability to better manage their risks and avoid adverse selection.[4]

2. Data and methodology

Defining a 'mini flash crash/rally'

We define a mini flash-crash/rally as a large price movement that reverts within a short time window and during which we see high levels of traded volume. The price shock could be upwards (a rally) or downwards (a crash). We identify recent episodes of such events using FCA order-book data. The data cover all order messages (new, modified and cancelled orders) and trades from all FTSE350 stocks on the London Stock Exchange (LSE), Bats Europe (including Bats and Chi-X) and Turquoise. We consider only the lit order book activity.[5] Our sample runs from January 2014 to June 2015.

More precisely, we identify mini flash crashes/rallies as price movements that:

1) exceed a pre-determined threshold (three times the average realised variation of the previous 20 trading days)

2) revert at least 50% within a 'short' time frame (less than 30 minutes) and

3) trigger high levels of trading volumes (at levels higher than the top 5% percentile of the distribution)

Some of the episodes we identify could be linked to the release of 'news'. Others could be caused by the execution of a single large order, sent either deliberately or by mistake, that the market can easily incorporate without 'contagion' to other venues or other participants' involvement. As we are interested in extreme events that could affect market resilience, we exclude any event that involves 'news' or that does not affect at least two trading venues. We also exclude events with a price change below 1%, as we view the magnitude of the change as too small and those where a 'circuit breaker' was triggered on the LSE. These last events have already been analysed in Bercich and Allan (2017). This leaves us with 40 events in total (Table 1) between January 2014 and June 2015.

	Number of events	Average % price change	Min duration of drop/spike	Max duration of drop/spike
FTSE100				
Crash	10	-2.8%	< 1 second	9 minutes
Rally	13	3.5%	< 1 second	10 minutes
FTSE250				
Crash	12	-2.7%	< 1 second	3 minutes
Rally	5	2.4%	< 1 second	14 seconds

Table 1: Basic features of the events in the sample

[4] Adverse selection is the risk that the market makers trade with someone better informed than them.

[5] In this paper we refer to lit order books as those with pre-trade transparency, where prices at which there are intentions to trade (the current bid and offer prices and the depth of trading interest) are publicly available. Dark order books are those where the transactions take place with no pre-trade transparency. Both lit and dark books are subject to post-trade transparency, ie there is a requirement to make public the time, price and volume of each trade after it has been executed.

Classification of market participants

We classify market participants in three categories: HFTs, hybrid firms and pure non-HFTs. HFTs are a subset of algorithmic trading participants that only carry out low-latency proprietary trading.[6] We classify HFTs by firm name based on supervisory knowledge and use the same classification that Bercich and Allan (2017) used to analyse the behavior of market participants when circuit breakers are triggered.[7] Hybrid firms are those firms that mainly provide agency trading services but may use a similar technology to HFTs or provide DEA to HFTs. They may undertake low-latency proprietary trading but it is not their main function.[8] This category mainly includes large investment banks. All remaining firms are classified as pure non-HFTs; firms that do not rely on low-latency strategies and do not provide DEA to HFTs. The pure non-HFT classification will give us an incomplete but cleaner proxy of trades done on behalf of institutional investors. Our sample includes 33 HFTs, 29 hybrid firms and 190 pure non-HFTs. From a previous study (Aquilina and Ysusi, 2016), we know that HFTs account for approximately 30% of the traded volume, hybrids for approximately 60% and pure non-HFTs for less than 10%.

3. Results

The 40 events we identified last for significantly different durations. For some, the time to the peak or trough is less than a second, while for some others it is several minutes. So we distinguish four groups of events based on duration of the price drop (or spike) and analyse these separately (see Table 2 for descriptions of these groups). As liquidity varies significantly between stocks even within groups, we scale the traded volumes and submitted order volumes before averaging the results.[9] Finally, as our groups contain both flashes and rallies, to aggregate them we invert the statistics of the crashes (ie multiply the crashes' statistics by -1). So, in what follows, we always present information as if we are discussing a rally.

	Observations	Duration of drop or spike
Group 1	10	< 1 seconds
Group 2	10	1 to 10 seconds
Group 3	13	10 seconds to 2 minutes
Group 4	7	2 to 11 minutes

Table 2: characteristics of the groups

[6] Lower latency implies higher speed.

[7] There are two different approaches to categorise firms as HFTs. The first relies on judgement and assesses whether the business model adopted by specific firms fits a given definition or not. The second relies on the data and categorises firms into HFTs on the basis of indicators such as the lifetime of orders, the order-to-trade ratio or others. As our data only identify members of trading venues at the firm level and not at the trading desk level, we get a cleaner proxy of HFT activity if we use the business model approach.

[8] With our data we are unable to separate the agency activity from the proprietary and DEA flow of these firms.

[9] The trading activity is divided by the stock's average traded volume by second. The order book activity (eg volume submitted or cancelled) is divided by the stock's average submitted volume by second. For Group 4 we use minutes instead of seconds because these events take longer to develop.

Mini-flash crashes and rallies are associated with greater trading, higher volatility and lower liquidity

We start by characterising what happens in the market as a whole before, during, and after these events.[10] Figure 1 shows the scaled trading volumes, quoted spreads, overall order submissions and intraday volatility around the time of the events. For Groups 1-3 this is shown in seconds, for Group 4 this is shown in minutes.[11]

During the event period, there is a significant spike in traded volume and volatility,[12] while quoted spreads sharply increase as do the number of submitted orders. Although we do not show it here, the pattern is consistent for all trading venues and is not driven by a single venue.

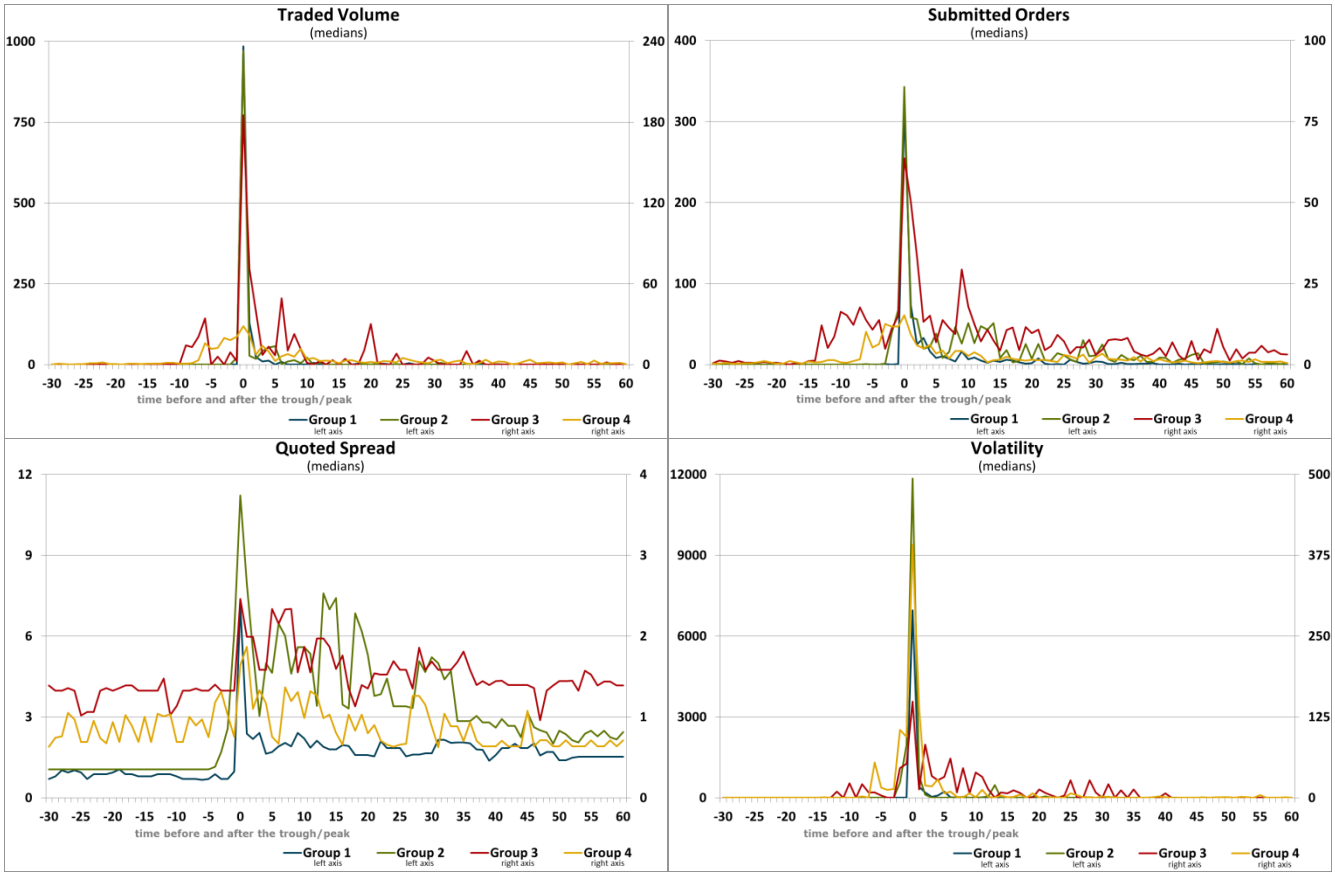

Figure1: features of the market before, during and after the events by group

The behaviour of market participants

In markets with limit order books, there are 2 distinct ways in which the best available bid/ask prices can change, namely:

[10] In all the analysis, t=0 is defined as trough of the price for a crash or peak of the price for a rally.
[11] The figures show the medians across stocks in the relevant group. Results are similar when using averages but with more noise.
[12] This is not surprising as we focus on events with large trading volumes and extreme price movements.

a) Aggressive trading. Where a participant takes the best available price to buy/sell a stock absorbing available liquidity, this may result in a change in the best available price.

b) Management of resting orders. Participants may modify or cancel their resting orders and this may change the best available price.

In the rest of this section we study both drivers as our data allow us to see how participants manage their resting orders as well as how they trade.

a) Hybrid firms give oxygen to mini-flash crashes and rallies by trading aggressively…

We begin our analysis by looking at the trading behaviour of market participants. Figure 2 compares the aggressive and passive trading of different types of participants in the 4 groups of events we identified. As explained further in footnote 3, aggressive price-taking orders remove liquidity from the order book while passive orders add to available market depth. Hybrid firms appear to be driving the extreme price changes in all groups by trading aggressively and building large positions. HFTs add some pressure by trading aggressively during the price spike but the overall magnitude is considerably smaller than hybrids.

Examining the trading of hybrid firms, we can see that they are the main liquidity providers. The majority of the traded volume executes against hybrids' orders resting on the book. But HFTs are also important providers of liquidity, especially at the beginning of the stress period and for those groups of events where the price drop or hike lasts for the shortest time (Groups 1 and 2).[13] Pure non-HFTs are not main players during these events: their participation is marginal in all cases.[14] Overall, hybrid firms carry out most of the trading in these events.

…while HFTs initially lean against the wind.

On average, HFTs as a group trade against the price movement more than they trade with it. But the best available prices can move even without trades, because participants modify the orders that are already resting on the order book. We now turn to examining this aspect.

[13] For the avoidance of doubt, it is not the same hybrid firms that are taking and providing liquidity simultaneously: the levels of self-trading are very low even during these stress periods. Different firms within the same category are on different sides of the trades.

[14] Pure non-HFTs' participation in the top-of-the-book is very rare so we do not show their behaviour in the next subsection when we discuss liquidity management.

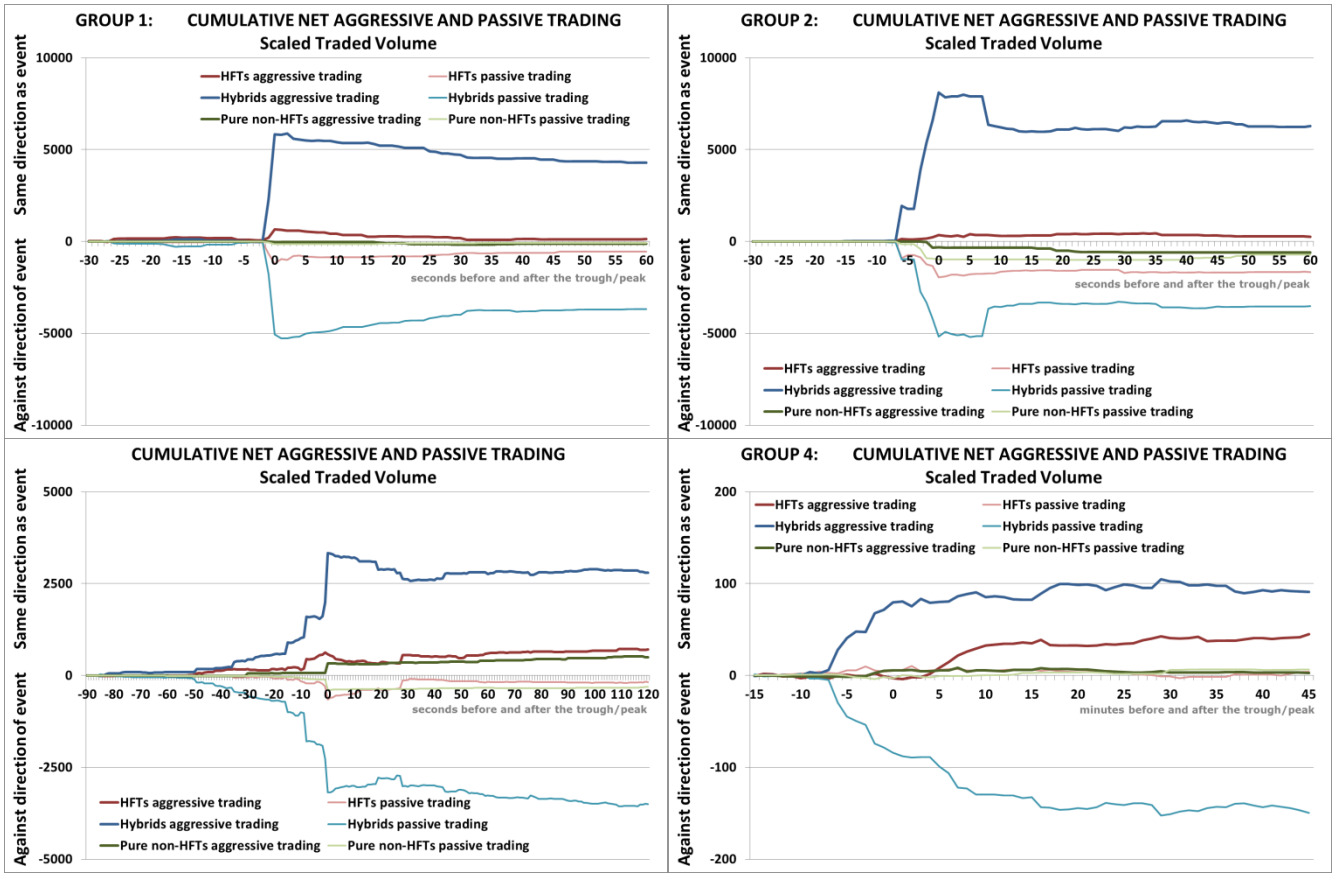

Figure2: the trading behaviour of market participants

b) The liquidity management of different market participants

There are 2 distinct ways in which changes in liquidity dynamics can intensify the pressure on prices:

i. participants not resubmitting the liquidity that is quickly consumed by others through trades

ii. participants cancelling or modifying their own orders that are resting on the order book

Both hybrids and HFTs keep submitting orders…

Focusing initially on the side of the order book in the same direction as the event (the bid-side for crashes and the ask-side for rallies), we study if there is an imbalance between liquidity being resubmitted and that being consumed by others that could intensify the price move. We examine the cumulative volumes that are being taken aggressively from HFTs (Figure 3 top plots) and hybrids (Figure 3 bottom plots) or, in other words, the cumulative volumes that each participant type is trading passively. We also examine the cumulative net volumes submitted by each participant type. The difference between the net submission of each group and the net liquidity taken by other market participants indicates if more liquidity is being taken than submitted in the order book. All the 4 panels in Figure 3 show that during the price spike, liquidity tends to be consumed more quickly than resubmitted. It is interesting to see however that the gap is on average larger for hybrids than for HFTs.

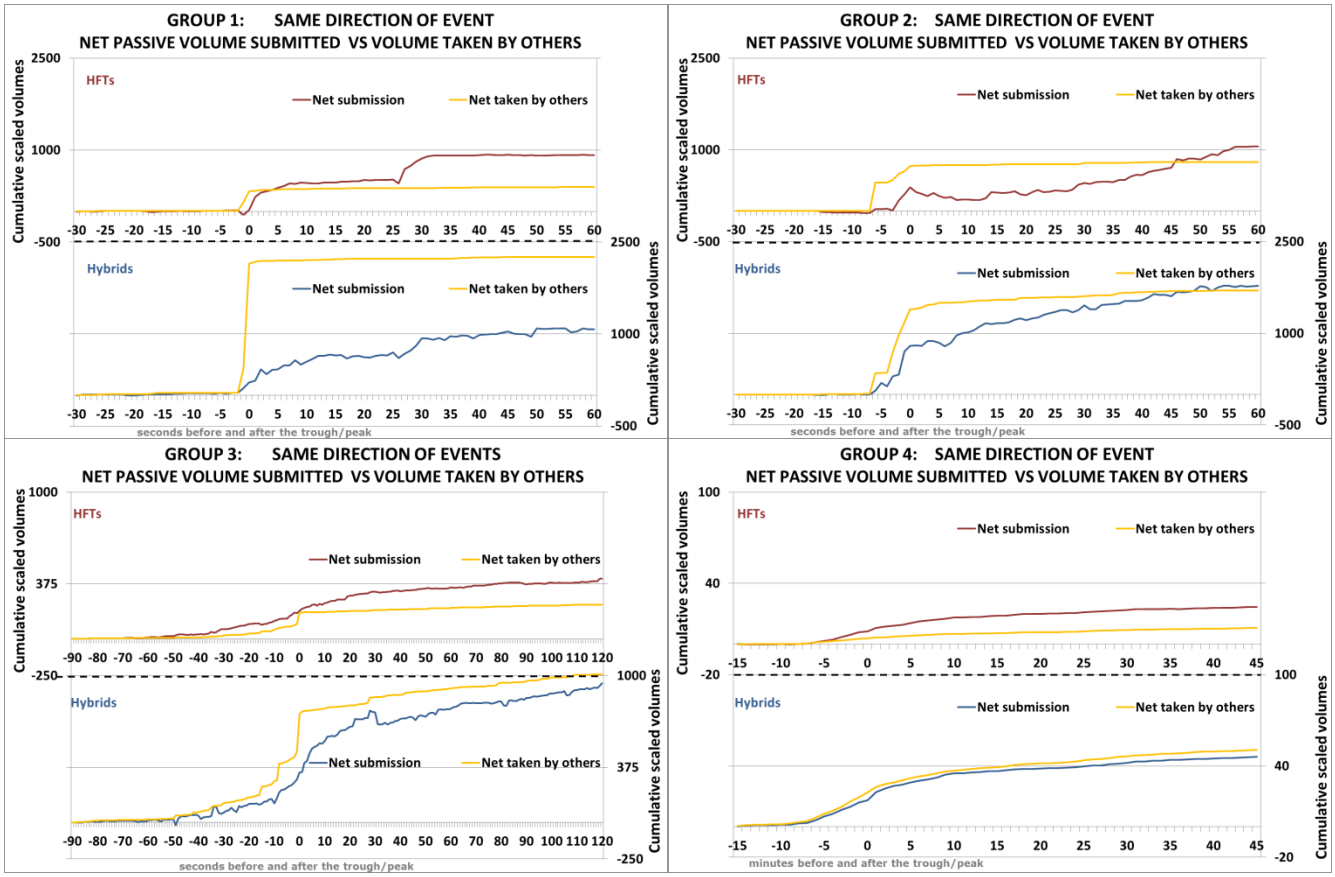

Figure 3: liquidity submission and consumption on the same side of the order book as the event

It is interesting to study the other side of the order book (the ask-side for a crash and the bid -side for a rally) to understand the price recovery. The submission of liquidity on this side of the order book starts later, ie during the recovery phase. And the overall volumes are much smaller than those we could see in Figure 3. So, the recovery phase takes much longer than the initial rally.

It is difficult to establish clear patterns for the 4 groups but it looks like during the recovery phase HFTs are submitting more liquidity than the amounts that are consumed by other market participants. Hybrids on the other hand do not resubmit enough liquidity to replenish the amounts that are being consumed by other traders, even if in most cases they submit higher volumes than HFTs.

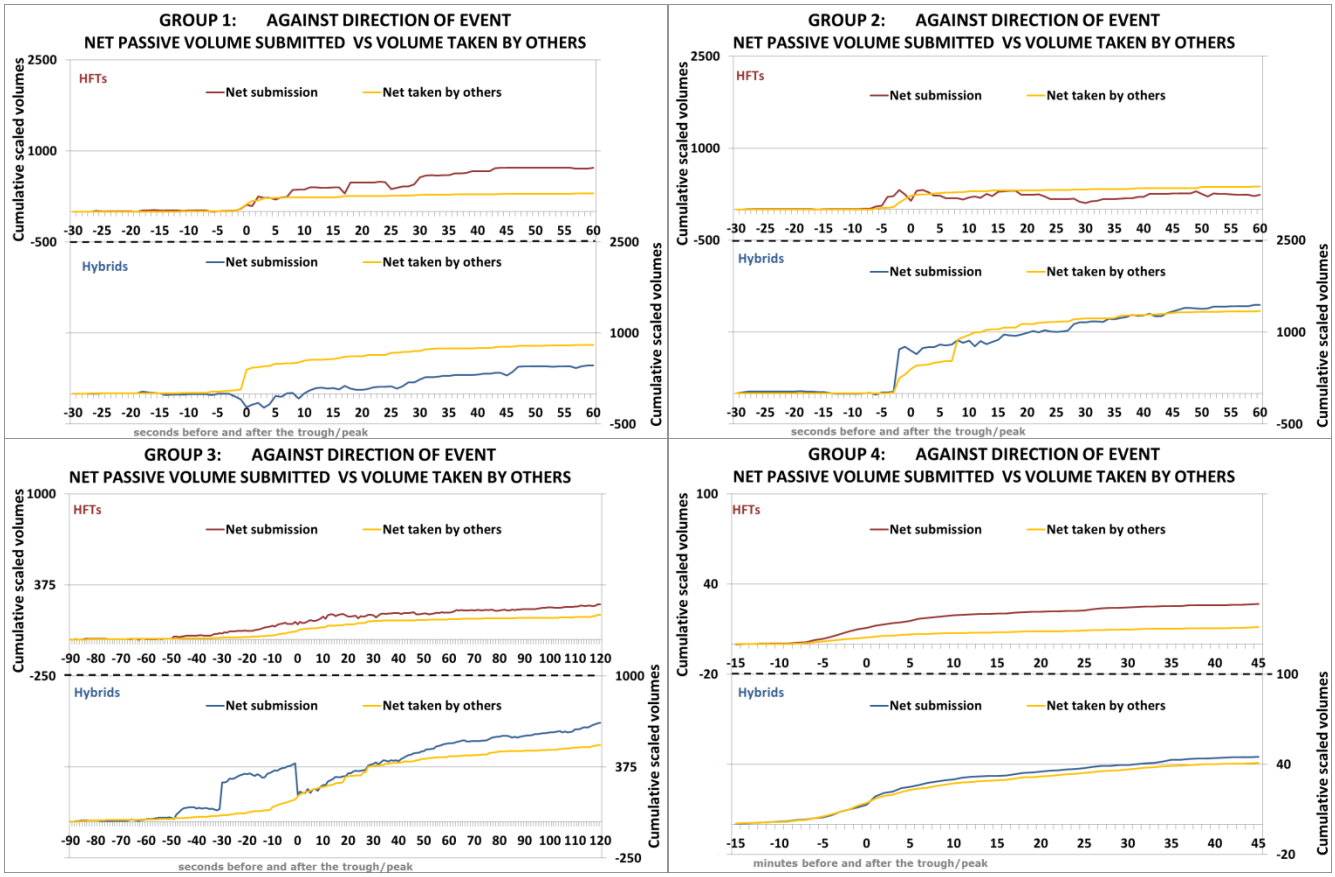

Figure 4: liquidity submission and consumption on the opposite side of the order book as the event

…but HFTs submit them far from the best bid and offer.

The previous analysis does not tell us where liquidity submission takes place in the order book. So far we have treated a new order the same if it was submitted very close to the prevailing price or much further away from it. However in terms of impacts on the market, it makes a difference where the order is placed in the order book.

For a more in-depth analysis of the management of liquidity in the order book, in Figure 5 we separate the order submissions into 3 groups: those that are at or very close to the best bid/offer (the top 3 levels of the order book), those that are a bit further away (levels 4 to 6) and, finally, all other submissions (all other levels above 6).[15] Positive values imply net submissions; negative values imply net cancellations.

It shows that, although HFTs' net liquidity submission is often positive, this liquidity is not submitted at the top-of-the-book or close to the best bid/offer. It also shows that HFTs tend to react differently depending on the duration of the stress period. For very fast shocks (less than a second long), HFTs simply cancel their liquidity at the top-of-the-book during the price fall/rally and only resubmit it until after the price pressure eased. For shocks lasting longer, HFTs resubmit liquidity far from the prevailing price. In some cases they even move some of their resting liquidity from the top levels of the book to lower levels. All these

[15] The scaled volumes are shown second by second for group 1 and 2, for group 3 the scaled volumes are aggregated every 2 seconds and for Group 4 the scaled volumes are aggregated every minute.

behaviours are likely to exacerbate the price movement as they affect the prices at which participants can trade.

Meanwhile, although not shown in this note, hybrids also cancel some of their liquidity at the top-of-the-book during extremely fast shocks. But in shocks lasting longer, they tend to resubmit liquidity at the top or mid-levels of the book (ie not too far from the best bid/offer). However, they don't resubmit enough quantities to substitute the liquidity taken from them. Overall, the depth of the order book is much thinner by the time the price starts to revert and even after the event has ended, because of the behaviour of both HFTs and hybrids.

In other results not shown here, in most cases, HFTs take much longer than hybrids to submit liquidity at levels comparable to those prevalent before these events.

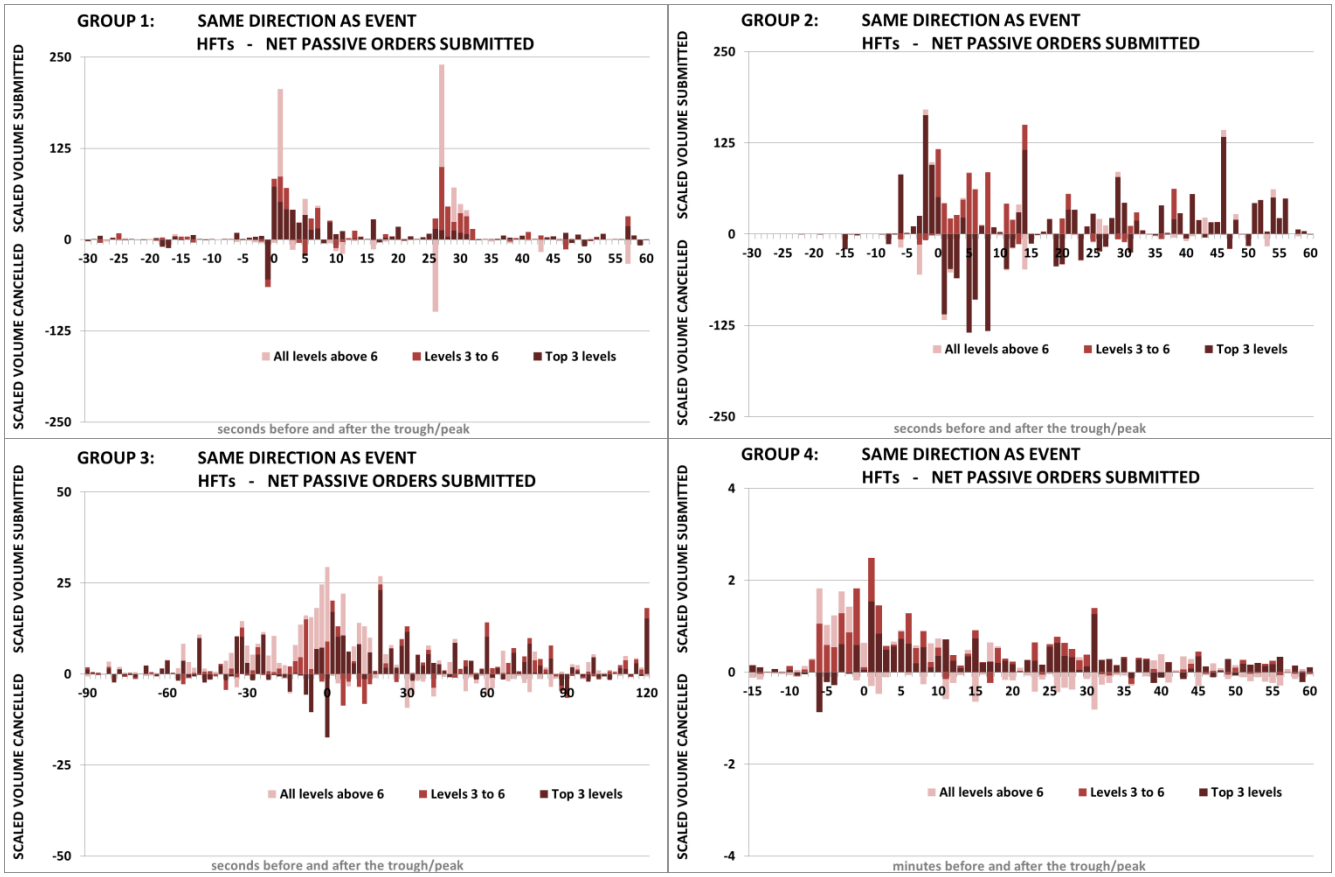

Figure 5: liquidity management by participants

On average HFTs make small profits, hybrids make small losses

In the previous section, we established that HFTs respond quickly to the risk by actively changing their quotes, widening their spreads, to avoid adverse selection. Probably due to their speed advantage, HFTs make on average a small profit during these stress periods (Table 3). Hybrids, on the other hand, make a small loss. Interestingly, pure non-HFTs make small profits probably because they do not face adverse selection as they are not liquidity providers and they can time better their executions. There are some episodes where HFTs make losses. But as shown in Figure 6, any HFT losses tend to be lower than those

made by hybrids. Figure 6 shows the distribution of the realised profits per event; the left plot includes extreme events and the right plot excludes them.

	Average realised profits per event (£)	Median realised profits per event (£)	Average trading share
HFTs	2,200	98	21%
Hybrids	-633	-21	74%
Pure non-HFTs	1,177	115	5%

Table 3: realised profits by participants' type

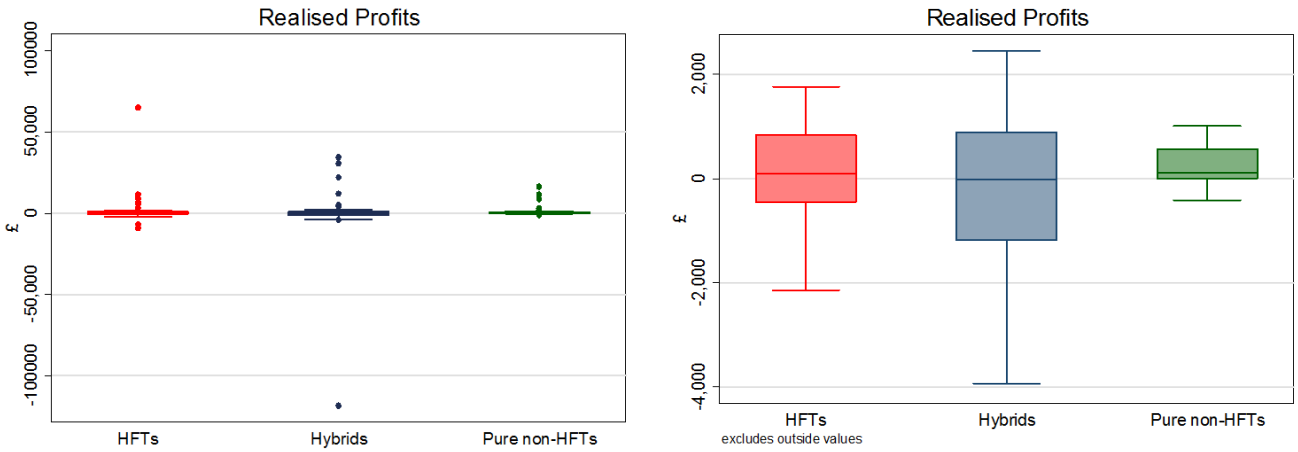

Figure 6: the distribution of realised profits

All participant types end up with long or short positions at the end of the stress periods. They may need to unwind these at some point, which will increase or decrease their realised profits/losses. It is impossible to quantify precisely these individual returns as we would need to know the price at which participants were able to buy or sell in the future. As an indicative measure, we calculate the unrealised profits of individual participants using best available prices at the end of the stress period. Table 4 shows the average and median unrealised profits. On average HFTs and hybrids experience unrealised losses but for HFTs this is driven by a single event. Excluding that event, HFTs make unrealised profits on average. These results need to be interpreted with care as these figures are estimated based on a strong assumption. However, together with evidence from realised profits, this proxy suggests that HFTs on average are better at managing their risks during stress periods.

	Average unrealised profits per event (£)	Median unrealised profits per event (£)
HFTs	-2,613	147
excluding one extreme value	606	
Hybrids	-1,587	-558
Pure non-HFTs	1,491	190

Table 4: unrealised profits by participants' type

A final remark

Our data show that the behaviour of hybrid firms contributes the most to the price spike/drop. However, we are unable to distinguish the types of order flow underlying hybrid transactions. Their trading during flash events could be driven by: 1) their own low-latency proprietary desks, 2) their agency business and the type of orders offered to their clients (eg stop-loss orders) or 3) other firms trading through them using DEA. Given what we know about HFTs through the result for the HFT category, it is less likely the main results are driven by HFTs trading through them using DEA.

4. Conclusions

In this note we analyse mini-flash crashes and rallies - large but not-extreme price movements. We identify and study 40 events in the UK equity markets, providing a characterisation of both the trading behavior and liquidity provision of different types of market participants.

On average, we find that hybrid firms (mainly large investment banks) compound the shock by trading aggressively in the direction of the price change. HFTs initially lean against the wind but subsequently tend to follow and ultimately also exacerbate the original price shock. But HFTs are responsible for a substantially smaller amount of aggressive trading.

Both hybrid firms and HFTs continue to provide liquidity during the event and do not withdraw completely from the order book. However, their liquidity is consumed more quickly through trades than it is replenished with new orders. Further, both participant types (but primarily HFTs) move their liquidity away from the top-of-the-book (best available prices) to further down on the side of the order book in the same direction as the event. This contributes to intensifying the initial price move as participants have to trade at worse prices. By the time price starts to revert, the depth of the order book has dropped significantly. In most cases, HFTs take longer than hybrid firms to restore the liquidity they provide at the top of the order book to a level similar to that present prior to the event, especially in liquid stocks. Unlike hybrid firms that make a loss, HFTs on average make small profits out of these events. This is probably due to their superior ability in speed which allows them to better manage their risks and avoid adverse selection.

References

Aquilina, M. and C. Ysusi (2016). Are High Frequency Traders anticipating the order flow? Cross-venue evidence from the UK market, FCA Occasional Paper 16.

Bercich, J. and M. Allan (2017). Catching a falling knife: an analysis of circuit breakers in UK equity markets. https://www.fca.org.uk/publications/research/analysis-circuit-breakers-uk-equity-markets

Breedon, F., L. Chen, A. Ranaldo and N. Vause (2018). Judgement Day: algorithmic trading around the Swiss franc cap removal, Bank of England Staff Working Paper 711.

Easley, D., M. Lopez de Prado and M. O'Hara (2011). The microstructure of the "flash crash": flow toxicity, liquidity crashes, and the probability of informed trading, 'The Journal of Portfolio Management', 37(2), 118-128.

Kirilenko, A., A. Kyle, M. Samadi and T. Tuzun (2017). The flash crash: the impact of high frequency trading on an electronic market, 'Journal of Finance', 72, 967-998.

Menkveld, A. and B. Yueshen (2018). The flash crash: A cautionary tale about high fragmented markets. 'Management Science' (forthcoming).

Noss, J., L. Pedace, O. Tobek, O. Linton and L. Crowley-Reidy (2017). The October 2016 sterling flash episode: when liquidity disappeared from one of the world's most liquid markets, Bank of England Staff Working Paper 687

Schroeder, F., M. Allan, A. Lepone, H. Leung and S. Satchell (2018). Flash Crash in an OTC Market, FCA Occasional Paper 37.