Scientific Annals of Economics and Business
67 (1), 2020, 139-155
DOI: 10.2478/saeb-2020-0006

You Need Three Butterflies to Cause a Hurricane

Gianluca Piero Maria Virgilio*

Abstract

The aim of this study is verifying the impact of high volatility, scarce liquidity and stop-loss orders on abnormal events like the May 6, 2010 Flash Crash. The paper assumes those three factors to be the main drivers, proposes a mathematical model based upon them and analyses audit trail data to verify whether those factors actually were at the origin of that event. It uses the concept of 'run', an uninterrupted sequence of trades all occurring in the same direction and compares volatility, liquidity and occurrence of stop-loss orders over the analysis period. The results found provide suggestive evidence that a combination of the three factors contributed to the crash. Each of them, taken individually, does not usually lead to extreme behaviours. Even two factors together may not disturb the orderly functioning of the markets but the combination of volatility, scarce liquidity and stop-loss orders may lead to a crisis.

Keywords: Flash Crash; abnormal market events; volatility; liquidity; stop-loss orders.

JEL classification: G12.

1. INTRODUCTION

May 6, 2010 (Flash Crash) shares with October 24, 1929 (Great Crash), October 19, 1987 (Black Monday) and a few others, the dubious honor of being an unforgettable day in the collective memory of the financial markets. On that day the market opened in a nervous mood. By 11 am Eastern Daylight Time (EDT) the Dow Jones was already down 60 points and at 2pm it reached 161 points (-1.5%) while the S&P500 was down 2.9%. At 2.24 pm the first stock was traded against a stub quote, more than 80% below previous day's closing. The investment firm Waddell & Reed (W&R) started a heavy sell programme of E-mini S&P 500 June 2010 futures contracts; in the meantime, over 200 securities had dropped 50% or more from their value just three quarters of an hour earlier. In the two minutes between 2.45pm and 2:47pm the DJIA, the S&P 500 and the NASDAQ 100 reached their daily low; at 2:45:27 the E-mini S&P 500 futures for June 2010, already down to 1062 from

* Universidad Católica Sedes Sapientiae, Nueva Cajamarca, Peru; e-mail: *consultgv@gmail.com*.

1165.75 at the beginning of the day, dropped 6 index points in just one second, down to 1056 (-9.4%) and launched a Stop Logic protection procedure. At that time US exchanges together were showing losses for one trillion dollars ($1,000,000,000,000). Despite the march to normalisation led by the Chicago Mercantile Exchange (CME), other markets kept behaving abnormally. The Dow Jones hit its daily low and so did the NASDAQ. Some large capitalisation stocks traded at ridiculously low prices: at 2:47pm Accenture changed hands at $0.01 from over $35 just ten minutes earlier (sources: CFTC-SEC Advisory Committee on Emerging Regulatory Issues, 2010b; Kirilenko *et al.*, 2017).

Several studies proposed explanations for the Flash Crash and actions to prevent it from happening again. This article contributes to that path of research by setting up a theoretical model and carrying out an in-depth analysis of audit trail data from the CME on the E-mini S&P 500 futures contract, one of the main characters of that day, with the purpose of identifying the causes of anomalous behaviours and to help understanding which factors contributed to the crash. Although there is a copious literature addressing the causes of the Flash Crash, no previous research used quantitative data to identify non-linear combination of the three factors listed above as the main cause of the crisis.

In the rest of this article Section 2 reviews the literature, Section 3 presents the theoretical model, Section 4 analyses the data, Section 5 carries out statistical analysis, Section 6 discusses the findings and Section 7 concludes.

2. REVIEW OF THE LITERATURE

It is widely accepted in the literature, although with some exceptions, that the Flash Crash originated at the CME, namely in the E-mini S&P 500 futures contract June 2010 market (Menkveld and Yueshen, 2019; CFTC-SEC Advisory Committee on Emerging Regulatory Issues, 2010b; Kirilenko *et al.*, 2017). Many authoritative sources (CFTC-SEC Advisory Committee on Emerging Regulatory Issues, 2010a, among others) identified a large sell program launched by W&R of 75,000 contracts on the CME, for a total worth of more than four billion dollars, as the triggering event of the Flash Crash. Yet, according to Menkveld and Yueshen (2019), that unusually large amount of futures contracts was not the cause of the crash, at least in a direct manner. The study finds that W&R only contributed 4% of the total net sells - a percentage that can hardly indicate it as the main culprit of the crisis, even more so as in the previous minutes (when the market was declining in an orderly manner) its trade intensity was even higher. As a last point, the paper notices that before the CME halt, only about half of the 75,000 contracts offered by W&R were actually sold, and other traders were the ones who sold aggressively. The remaining trades by W&R occurred when the CME was bouncing. One of the sharpest criticisms to the CFTC-SEC Report comes from the market analysis firm Nanex. By reporting findings from Nanex, Durden (2012) points out as "it was precisely HFT [High-Frequency Trading] quote churning that was the primary, if not sole, reason for the catastrophic chain of events". Nanex (2010b) analysis of the Flash Crash reports statements from some trading firms about detection of a data feed accuracy problem, which caused them to temporarily withdrawing from trading, leading to a further reduction in liquidity. Based on this evidence, Nanex (2010b) states that "the delay in NYSE quotes was at the root of this detection", contrary to CFTC-SEC's findings. Nanex (2010a) restates and arguments the above by showing evidence of the Consolidated Quotation System (CQS) not operating normally and within capacity by plotting CQS traffic at 2ms intervals versus the same plotted at one-second intervals. The former message rate trespasses several times the 250,000 messages/second proportional threshold whereas the latter graphs (by averaging the ratio over one second) misses to show those peaks. A coarser granularity of the analysis would thus completely miss to notice the CQS delay, therefore deflecting the focus of the Flash Crash investigation towards other causes. Zervoudakis *et al.* (2012) summarise the so-called Nanex Theory, a sharp refutation of the CFTC-SEC Report's findings. At 2:42:46 pm the NYSE began to experience delays in its quoting dissemination system and therefore the NYSE quotes transmitted to the CQS were no longer reflecting the actual quoted prices. Bid quotes crossed above ask for about 250 stocks, giving rise to a large amount of arbitrage opportunities. The High-Frequency (HF) traders that tried to exploit such opportunities by aggressively selling against limit order bids, found their market orders executed at lower prices than originally intended, as the market was sharply moving downward. Since quotes and execution data are reported to the CQS and to the Consolidated Trading System (CTS) respectively, and the latter's traffic is a fraction of the former's, trade data experienced little or no delay. Yet, HF traders detected the price decline, which led to a stream of sales, exacerbating the decline and causing a negative feedback loop. The stop-loss mechanism may have also had a role in exacerbating the process, as suggested by Foresight - Government Office for Science (2012). This is a comprehensive study that, underpinning its arguments on a set of nearly 60 papers, reaches the conclusion that "there is as yet insufficient evidence to what role HFT played either in the Flash Crash or other mini crashes that have occurred since HFT became established" (ibid. p.141). This uncertainty is shared by Kirilenko *et al.* (2017), who recommend further data analysis by making use of data from all venues, products and traders on May 6, 2010, in order to carry out an examination of Flash Crash hypotheses. By considering market fragility as not directly caused by HFT on the Flash Crash day, the authors implicitly suggest the former not being directly related to the latter. Instead, it could be a consequence of systemic risk but, according to Danielsson and Zer (2012), there is no clear consensus on what constitutes systemic risk, and translating the expression into the risk of collapse for the entire financial system does not help reaching a measurement of any practical usefulness. Indeed, lack of consensus on many important financial issues and even on some basic definitions thereof, is a recurring theme imperiously calling for settlement, if a solution to the financial stability problems is to be found. However systemic risk is defined, Cliff (2011) foresees it likely to grow rather than diminish in a future filled up with HFT activity, unless appropriate actions are taken. And although the probability of such events taking place is small, their potential consequences are so serious and so far-reaching in both space and time, that appropriate actions are seen as urgently needed. Yet, in her testimony rendered before the US Congress on the severe market disruption on May 6, 2010, Mary L. Schapiro, Chairwoman of the SEC, admitted that "the technologies used for market oversight and surveillance have not kept pace with the technology and trading patterns of the rapidly evolving and expanding securities markets" (Schapiro, 2010, p. 17). In order to cure this weakness many authors have advanced sensible proposals. Among them, Bullock (2011) highlights the emerging need for financial system simulation and Sornette and von der Becke (2011) reinforce this view by stating the “need to build policy making devices (a 'policy wind tunnel' [...] or an 'economic flight simulator')" (ibid. p.15).

3. THEORETICAL MODEL

Many papers analyse in depth the impact that volatility and liquidity have on sharp market movements and in some of them also the role of stop-loss orders is mentioned, although only on a qualitative basis. In particular, Cespa and Foucault (2014) present a theoretical model in which liquidity providers learn information about an asset from the price of another asset and therefore explain how shocks specific to liquidity supply in one asset class propagate to other asset classes. Then, signals returned from the latter influence the former, creating a feedback loop. According to the authors, this feedback loop provides an explanation for liquidity co-variations and crises. On the other side, Goldstein *et al.* (2013) propose a model in which traders with different motives operate and this difference may lead to opposite responses when facing the same market signal. Thus, adding more informed traders may add more contrasting signals, leading to learning complementarities, reducing price informativeness and eventually generating price jumps. The model proposed below makes use of, and links to, the models described in Cespa and Foucault (2014) and Goldstein *et al.* (2013), and also includes stop-loss orders.

Let P_1 and P_2 denote the price of risky asset 1 and risky asset 2, with supply functions X_1 and X_2, respectively. A cycle is made up of three dates. Liquidity X_n is provided at date 1 with price P_n (with $n \in \{1, 2\}$) when all aggressive traders observe the price and start analysing it. At date 2, traders post aggressive orders that execute immediately. If execution changes the price of top-of-book, at date 3 stop-loss orders get executed. For schematization purposes, within the trader community we will devise cautious traders (C-traders), who set stop loss at level $\pm\lambda\%$ of price, according to whether holding a short or long position. Price and liquidity of asset n is influenced by the market maker's analysis of θ_m (market signal), by its risk tolerance Υ_m, and by the price of asset n' (where n' = n mod 2 + 1) for a percentage of ρ. Moreover, price and liquidity depend on each other, as a recognition of market-wide consensus. Since all analyses and risk tolerances might be different, we shall take the most optimistic price as reference, since that is the price displayed at the top of the book. The system of equations relevant at date 1 is:

$$P_{nm} = P(\Upsilon_m, \theta_m, X_n, \rho P_{n'}) \qquad (1)$$

$$X_n = X(\Upsilon_m, \theta_m, P_n, N_{Cn\lambda}, \rho X_{n'}) \qquad (2)$$

where the suffix n identifies the asset to which price and liquidity refer, n' identifies the other asset, suffix m indicates market maker and $N_{Cn\lambda}$ is the number of C-traders who set a stop loss order for their trade on asset n, should the new price worsen $\lambda\%$, or more, of the originally traded price.

The only equation which enter the system at date 2 is:

$$P_{nt} = P(\Upsilon_t, \theta_t, X_n, \rho P_{n'}) \qquad (3)$$

where P_{nt} is the price resulting from the market taker's analysis of signal θ_t about asset n and Υ_t is the risk tolerance of market taker. The market clears when $P_{nm} \geq P_{nt}$ or $P_{nm} \leq P_{nt}$, according to whether the trade occurs at the bid or ask. All functions described above depend, via the signals θ_m and θ_t, on time series of the variable itself (price at time t also depends to some extent on price history at time t-1, t-2, …, and so does liquidity).

Moreover, the dependence of price on liquidity, and vice-versa, is not only a function of each other variable, but also of their first derivative, since the speed of price or liquidity change has an impact on the signals θ_m and θ_t. If price or liquidity show a variable speed, and therefore an acceleration, the dependence would also be on their second derivative. Should trading executed at date 2 consume all liquidity at the top of the book, and so doing move the asset price to P'_n, at date 3 all stop-loss orders hit by the price move will trigger, impacting on liquidity X_n and potentially on price P'_n, (changing to P''_n), should further liquidity levels be consumed by the stop-loss orders. The equation relevant to date 3 is:

$$P''_n = P_{SLn}(N_{Cn}, SL_n, P'_n, X_n), \text{ if } SL_n \text{ triggered at } P_n \text{ consumes enough liquidity } X_n \text{ to move price} \quad (4)$$
$$P''_n = P'_n, \text{ otherwise}$$

where SL_n is the number of stop-loss orders triggered at price P'_n by C-traders on asset n (N_{Cn}). In case N_{Cn} or SL_n are large enough, trade execution triggering stop-loss orders may, on its turn, move the price to be evaluated at date 1 of the following cycle and then pushing the system into a fiendish automatic loop in which prices take a definite direction without any need of external intervention, as shown in Figure no. 1.

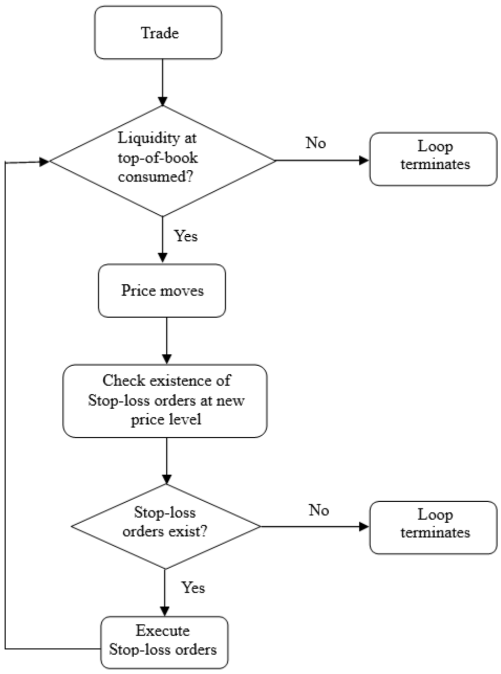

Figure no. 1 – Feedback loop

A small value of λ, correlated to a low value of Υ_t, indicating nervous C-traders, might further exacerbate the process. Moreover, the new price, P'_n, would also affect, on the following cycle, the price of related assets, $P_{n'}$, according to equation (1). The data analysis in the next section shows a case of this occurrence.

The purpose of this study is looking for evidence of causal relationship between the three factors described above and extreme market events. Data in themselves do not encompass causation evidence, and even causal statistics must be taken with extreme care as even Nobel Laureate C.W.J. Granger made it very clear that Granger-causation is not the same as causation in the common sense of the word (Granger, 1969). As it is well known, correlation cannot be taken as evidence of causal relationship: nevertheless, this study includes in section 5 statistical analysis on correlation between liquidity and abnormal price changes, and compares it with correlation on the combination of liquidity and stop-loss orders with price change. Although correlation does not imply causation, the findings could support suggestive evidence of directional relationship between these two factors combined together towards strong price movements.

4. DATA ANALYSIS

The literature recognises that stop loss orders may have had a role in the dramatic events of May 6, 2010. According to Foresight - Government Office for Science (2012) they have the potential to trigger negative feedback loops. In order to verify this hypothesis, detailed data would be necessary, showing which limit orders had a stop-loss order associated with them and at which price. Unfortunately, this level of details is not publicly available. The Chicago Mercantile Exchange provides market data messages needed to recreate the 10-level top of book for products traded on the CME Globex electronic trading platform: that includes all changes to the book including bids, asks, trades, bid volumes, ask volumes and trade volumes - ten levels deep, time-stamped to the millisecond. However, since a detailed analysis about the impact of stop-loss orders is not possible in a direct way, a different approach is required. The only way to work out the impact of stop-loss orders onto a volatility crisis is by using a proxy (a tool commonly used in the literature) and the proxy chosen in this research is the length of a 'run'. A 'run' is defined as an uninterrupted sequence of trades all in the same direction (that is, aggressive orders all against a bid quote or all against ask quote) and the 'length' of the run is the number of trades within a run. The runs are taken as an indication of the presence of stop-loss orders and all sections below dealing with runs can be grouped under the common purpose of showing stop-loss activity. Obviously there is no absolute guarantee that trades in a run are caused by a sequence of stop-loss orders and that is where a proxy helps. An example is given by a couple of runs occurred on the day before the Flash Crash, that can be assumed a ‘normal’ day. A run of length 8 started at 18:45:51 and 319 milliseconds GMT (that is 2:45:51.319pm EDT) on the ask book of the E-mini S&P 500 futures contracts traded at the CME. Another trade occurred within the same second but at millisecond 332, then one at 337, and then at 352, 358, 375, 378 and 394. The total time covered by the run was 75 milliseconds and this can hardly configure a sequence of stop-loss orders as they are normally launched with no latency in-between. In this case the latencies between successive trades were 13 milliseconds, then 5ms, then 15, 6, 17, 3, and 6. In another case on the same day and on the same book, another run of length 8 executed within one millisecond: the run started and terminated at 18:27:27.115. Seventy-five milliseconds versus less than one: it seems sensible to identify quite some exogenous activity in the former case and automatic stop-loss execution in the latter. In runs much longer than average, executed in a very short time, it can be sensibly assumed that a special automatic mechanism was in place, and the typical automatic mechanism that can increase the length of a run in a short time is the execution of stop-loss orders. It must be noticed that long runs do not necessarily cause large price movements or stress. The second longest run in the observation period on May 5 (lasting 72 milliseconds) displays a length of 684, trading a total of 1,248 contracts. The total price change was just one tick (0.25 index points).

The investigation presented here has been carried out with the purpose to identify abnormal values in periods of time, over different days, showing the same total number of market events and, once such anomalies have been identified, to understand the underlying causes. The date and time under primary observation are the six minutes and twenty-eight seconds (18:39:00.007 through 18:45:28.115 GMT) leading to the triggering of 5-second Stop Logic by the CME Globex platform, an event that started the recovery on the E-mini S&P 500 futures contract market and eventually brought the Flash Crash to an end. In order to evaluate such observations, data from May 6 will be compared with data including the same six minutes on May 3rd, 4th, 5th and 7th. The criterion chosen for deciding the length of the investigation period was the number of records produced by the CME Globex platform in that period. This way it is sure that the same number of market events will be taken into account, for all of the days observed. Since May 6 was a rather busy day, the same number of market events (580,864) occurred during six minutes on that day needed much more time to occur on the other days. An overview about the findings of this investigation is reported in Table no. 1.

Table no. 1 – Data about trade runs (bid book)

Date	**03-May**	**04-May**	**05-May**	**06-May**	**07-May**
Time window	173318.954-192748.770	181107.807-191047.481	183348.932-192318.863	183900.007-184528.115	182906.396-190434.858
Duration	01:54:30	00:59:40	00:49:30	00.06:28	00:35:28
Events	580,864	580,864	580,864	580,864	580,865
PANEL A					
Trade runs	11,399	9,293	8,843	12,824	6,656
Max length	273	260	239	324	156
Avg length	2.91876	3.40170	3.37114	2.79141	3.10592
>300	0	0	0	1	0
>200	3	2	4	4	0
>100	25	25	26	19	6
>50	87	90	103	63	39
>25	219	203	219	184	137
>10	503	543	438	575	380
PANEL B					
Delta price					
Max	0.25	0.25	0.25	4.5	0.50
Avg	0.00013	0.00073	0.00074	0.01482	0.00240
4.5 ip	0	0	0	1	0
4 ip	0	0	0	1	0
3.5 ip	0	0	0	1	0
3.25 ip	0	0	0	2	0
3 ip	0	0	0	1	0
2.75 ip	0	0	0	2	0
2.5 ip	0	0	0	3	0
2.25 ip	0	0	0	3	0
2 ip	0	0	0	3	0

Date	03-May	04-May	05-May	06-May	07-May
1.75 ip	0	0	0	6	0
1.5 ip	0	0	0	11	0
1.25 ip	0	0	0	15	0
1 ip	0	0	0	11	0
0.75 ip	0	0	0	49	0
0.5 ip	0	0	0	144	2
0.25 ip	6	27	26	1455	62

The data presented in this and all following tables have been computed by the author, based on Market Depth Data provided by the CME Group for the E-mini S&P 500 futures contracts expiring on June 2010. All raw data are time-stamped at millisecond granularity.

The header shows, for each day being investigated, the observation's time, its duration and the number of events. Then, Panel A displays the number of trade runs, the maximum and the average length of the runs, and the number of runs whose length was greater than 300, 200, 100, 50, 25, and 10 trades, respectively. Panel B details the difference between the initial price and the price at the end of the run. It shows the maximum price difference, the average, the number of runs in which the price difference was equal to 4.5 index points, 4, 3.5, and so on, at 0.25 index points interval, down to 0.25 index points, i.e. how many times at least one tick change occurred during a run (0.25 is the tick in the E-mini S&P 500 futures contract market.)

Some differences between the results in the columns of the table, representing the different days under investigation, are clear at first sight. Let us look at them more closely.

4.1 Traffic

The first thing to notice is that, in order to analyse an identical number of market events, very different periods in time had to be selected. Namely, the same number of events (580,864) which took place on May 6 during the six minutes and 28 seconds leading to the halt, needed one hour 54 minutes and a half on May 3 to execute, one hour less twenty seconds on May 4, forty-nine minutes and a half on the 5th, and 35 minutes and 28 seconds on the 7th (Table no. 1) This suggests a lot of trading activity that authoritative sources (CFTC-SEC Advisory Committee on Emerging Regulatory Issues, 2010a) state having been generated by High-Frequency traders: “HFTs began to quickly buy and then resell contracts to each other – generating a ‘hot-potato’ volume effect as the same positions were rapidly passed back and forth” (ibid. p.6).

4.2 Number of trade runs

Also the number of trade runs which occurred on the day of the Flash Crash was greater than the number of runs on the other days, the increase ranging from 12.5% (compared to May 3) to nearly 93% (compared to May 7), as shown in the third row of table 2, where again the percentage computes the increase on May 6 compared to the day on each other column.

Table no. 2 – Run rates

Date	03-May	04-May	05-May	06-May	07-May
Runs	11399	9293	8843	12824	6656
Comparison	12,50%	38,00%	45,02%	0,00%	92,67%
Duration	01:54:30	00:59:40	00:49:30	00:06:28	00:35:28
Runs/sec	1,7	2,6	3,0	33,1	3,1
Sec/run	0,603	0,385	0,336	0,030	0,320

This factor, together with the ones described in the next five sub-sections, are indicative of a higher-than average number of stop-loss orders being executed. It can be argued that this is an expected occurrence on a very volatile day, since wide price movements tend to trigger the stop-loss mechanism more often. Nevertheless, the wide range for this indicator suggests that the stop-loss mechanism is definitely a potential candidate to be one, although not the only one, of the main factors that contributed to the crisis. Moreover, the frequency of such events (Runs/sec in Table no. 2) is also worth discussing. On May 3rd, 4th, 5th, and 7th the event-per-second rate ranges between 1.7 (one event every 603ms, on the 3rd) and 3.1 (one event every 320ms, on the 7th), whereas the rate observed on May 6, that is 12,824 events in 6 minutes and 28 seconds, is equal to 33.1 events per second, or an event every 30 milliseconds. If in the non-Flash-Crash days the events, although at a very high rate, were at some extent still understandable by a well-trained human brain, the frequency on the 6th was far too high even for a human eye to grasp.

4.3 Maximum and average length of a run

On May 6 the maximum length of a run (shown in Table no. 1) was 324, versus an average of 250 max length for the days other than the 6th. Under the reasonable assumptions that a large number of trades occurring within a few milliseconds (and in many cases the whole run occurred within the same millisecond) cannot be but automatic, it looks very likely that after the initial investor-driven trades within a run, a rather large number of stop-loss orders were executed in sequence.

4.4 Number of runs longer than 10 trades

It can also be observed that the number of runs respectively longer than 100, 50, and 25 trades was not highest on May 6, as shown in Table no. 1. Instead it shows a peak in the number of runs longer than 300 trades and longer than 10 trades. This is relevant under the quite reasonable assumption that stop-loss order triggering definitely occurred at least in runs longer than 10 trades. This observation suggests that on the Flash Crash day the runs displayed extreme behaviours: a small number of very long runs and a larger than average number of short runs, showing a below-average number of medium-length runs. This observation again confirms the erratic behaviour of the market, at least during those critical six-and-a-half minutes.

4.5 Max price drop within a run

By far the most interesting finding in Table no. 1, panel B shows the maximum price difference between the beginning and the end of a run: on May 6 it was 4.5 index points (eighteen ticks, equivalent to USD 225) versus 0.25 index points (one tick or USD 12.50) on the 3^{rd}, 4^{th} and 5^{th}, and 0.5 index points on the 7^{th}. This provides a clear indication: stop-loss orders very likely contributed to the dramatic price drop on the day and at the time of the Flash Crash but that could not have happened if the long runs triggered by stop loss orders had not hit the vacuum, forcing the price to match sequentially downward to the next price level. As widely recognised by the literature, scarce liquidity was another major cause of the event. All large price movements occurring on May 6 during a run are summarised in Table no. 3.

Table no. 3 – Largest price movements on May 6, 2010

EVENT ID	INIT TIME	INIT PRICE	END TIME	END PRICE	TRADES	DELTA PRICE	MILLISECS
26129550	184518708	1071.25	184518708	1066.75	1	4.5	0
26129582	184518710	1070.75	184518710	1066.75	1	4	0
26142383	184526518	1070.25	184526525	1066.75	83	3.5	7
26119292	184513871	1074.5	184513881	1071.25	99	3.25	10
26145427	184527996	1066	184527998	1062.75	16	3.25	2
26116166	184512489	1074.75	184512491	1071.75	21	3	2
26117421	184512952	1075	184512960	1072.25	82	2.75	8
26129358	184518693	1071.75	184518699	1069	68	2.75	6
26117532	184512960	1075	184512960	1072.5	2	2.5	0
26143136	184526902	1065.5	184526902	1063	1	2.5	0
26145696	184528111	1059	184528114	1056.5	31	2.5	3
26115781	184512444	1075	184512459	1072.75	125	2.25	15
26127454	184517945	1075	184517961	1072.75	126	2.25	16
26143189	184526913	1065.25	184526913	1063	1	2.25	0
26114163	184511702	1076.75	184511709	1075	65	1.75	7
26114406	184511744	1076	184511748	1074.25	39	1.75	4
26116481	184512560	1073.5	184512560	1071.75	1	1.75	0
26143035	184526894	1064.75	184526897	1063	30	1.75	3
26143172	184526911	1064.75	184526911	1063	1	1.75	0
26145605	184528107	1062	184528109	1060.25	28	1.75	2
26106768	184508584	1077.5	184508589	1076	58	1.5	5
26116298	184512516	1074.75	184512516	1073.25	1	1.5	0
26116387	184512535	1073.25	184512535	1071.75	1	1.5	0
26117638	184512980	1073.75	184512980	1072.25	2	1.5	0
26119424	184513882	1074.5	184513882	1073	1	1.5	0
26129649	184518726	1068.25	184518726	1066.75	1	1.5	0
26135595	184522305	1068	184522307	1066.5	19	1.5	2
26143383	184526992	1064.25	184526992	1062.75	1	1.5	0
26143501	184527018	1064.5	184527018	1063	7	1.5	0
25897340	184321495	1103.75	184321515	1102.5	235	1.25	20
25944727	184343474	1100.25	184343487	1099	149	1.25	13
26089228	184458528	1081.75	184458532	1080.5	46	1.25	4
26117675	184512986	1073.75	184512986	1072.5	1	1.25	0
26129455	184518699	1069	184518705	1067.75	29	1.25	6
26129500	184518705	1067.75	184518707	1066.5	24	1.25	2
26133601	184521096	1070.75	184521098	1069.5	22	1.25	2
26134309	184521415	1068	184521417	1066.75	30	1.25	2
26134615	184521629	1068	184521630	1066.75	9	1.25	1

At 18:45:13.871 GMT a run of length 99 started and went on for 10 milliseconds, until 18:45:13.881. It caused the price of the E-mini S&P 500 futures contracts June 2010 falling down thirteen ticks (from 1074.50 down to 1071.25), while price change during a run was usually restricted to one or two ticks at the maximum on the other days. The total dollar amount involved in that run exceeded twenty-one millions. The run encompassing the third largest price movement, 3.5 index points or 14 ticks, lasted 7 milliseconds (an 83 trades-long run) and the fourth largest price drop (3.25 index points) lasted 10 milliseconds (99 trades-long run). Only two runs which experienced large price movements, for 2.25 index points, lasted more than 10ms and can therefore be considered, at least partly, generated by the so called 'hot-potato effect' (CFTC-SEC Advisory Committee on Emerging Regulatory Issues, 2010b, 2010a), namely the one started at 18:45:12.444 going on for 15ms and the one commenced at 18:45:17.954 lasting 16ms. On the other side, in four cases the run started and terminated within the same millisecond. In particular, the run which started at 18:45:12.960 GMT had a length 2 and a price change of 2.5 index points (10 ticks). This shows the interesting phenomenon of mixed HFT and non-HFT latency. At 18:45:12.952 GMT the prevailing bid price was 1075 index points. At that time a run started (seventh row in table 3) and in just 8 milliseconds it consumed all the liquidity available down to 1072.25. Just after that, at 18:45:12.960, a bid was quoted at 1075, which got gobbled immediately, causing a large quote drop on the bid book in no time. A sensible explanation is that an investor, probably a computer that, compared to a HF trader, had a higher latency, noticed the prevailing bid at 1075 at 18:45:12.952. So it launched its own bid at that price but, because of its latency, the quote only arrived 8 milliseconds later, at 18:45:12.960, after all the liquidity at 1075 (and down to 1072.25) had been taken away by the stop-loss mechanism. The result was that when the bid order at 1075 hit the book, the price had already gone down quite a lot and it found itself totally isolated from other bid limit orders which were, at that time, quoting at most 1072.25. The lonely quote was immediately taken up by a lucky HF trader which sold 2.75 index points above the prevailing bid price just exploiting the other's latency and its own rapidity. So, not only latency-prone market orders are at risk of nasty surprises but limit orders too could, and did, become stale in a matter of a few milliseconds - even before they reached the exchange server. In other words, in the Age of HFT, both limit and market orders may become obsolete between their conception and the time they are born.

4.6 Average delta price over all runs

The price difference (also called delta price in the text) between the beginning and the end of a run is indicative of the dramatic price drop caused by the stop-loss mechanism. But the peak (4.5 index points on May 6 against 0.25 or 0.5 on the other days) tells only half of the story. Table no. 4 shows the average price difference across the runs longer than 10 trades for each of the days under observation.

Table no. 4 – Average price difference normalized (Delta Price)

Date	03-May	04-May	05-May	06-May	07-May
AVG	0.00013	0.00073	0.00074	0.01482	0.00240
	1	5.51980	5.58585	112.59150	18.26763
		1	1.01197	20.39775	3.30947
				6.16344	1

The following three rows show the ratio between the averages normalised on May 3rd, 4th, and 7th, respectively. It means that on the third row the average for May 3 has been taken equal to 1 and the others as the multiple of that average. So, taking May 3rd as a basis, on May 4 and May 5 the average price difference was about 5 times higher, on May 7 it was 18 times higher and on the 6th the ratio was more than 112. The actual figures are not that important as the qualitative indication they provide: on the day and time of the Flash Crash the average price drop during a run was very much higher than on any other day. This is another indication of the role taken by stop-loss orders in a largely illiquid environment.

4.7 Number of runs showing price jumps

A similar information comes from the number of runs showing price drops. Whereas there is no point in comparing price drops higher than 0.5 index points, as only May 6 shows a number greater than zero in those cases, it is interesting to compare the number of runs experiencing at least one-tick price change across the days under scrutiny, as in Table no. 5.

Table no. 5 – Price change normalised (Delta Price)

Date	03-May	04-May	05-May	06-May	07-May
>0 ip	6	27	26	476	62
	1	4.5	4.33333	79.33333	10.33333
		1.03846	1	18.30769	2.38462
				7.67742	1

Taking May 3 as a reference, on the 4th and the 5th there were, respectively, four and a half and four and one third as many runs during which the price moved (down) by at least one tick, more than 10 times on May 7 and more than 79 times on the 6th. On the last row, the comparison between May 6 and the second most volatile day (May 7) shows a ratio as high as 7.677. So, the price change during a run clearly shows that on the Flash Crash day, price drops caused by stop-loss orders were not only larger in size per each run on average (as seen in the previous section) but they also occurred much more frequently, indicating very scarce liquidity.

4.8 Liquidity

Scarce liquidity has been considered one of the biggest issues of the Flash Crash. The Combination of price uncertainty and withdrawal of many market makers and other liquidity suppliers, led liquidity to virtually vanishing at the peak of the crisis. To better understand the extent to which this phenomenon materialised, Table no. 6 compares liquidity available at the ten top levels of the bid book at one of the most critical times on May 6 (heading: 06-May-b), compared to the beginning of the observation period on the same day (heading: 06-May-a) and to about the same time on the other days.

The aggregate top ten level liquidity varies a lot over the range investigated.

However, it seems reasonable to take the liquidity shown on the 3rd, the 4th and the 5th as somehow standard, because May 6 was affected by very negative news since the beginning of the trading day and because 18:39 was already a critical time for the E-mini S&P 500 futures contracts. Moreover, the 7th can hardly be considered a standard day as it is reasonable to assume the markets still rather shocked by the previous day's events and, understandably, liquidity suppliers much more cautious than usual. Nevertheless, as shown in Table no. 7, liquidity at the most critical time on May 6 decreased about 98% with respect to May 3 and May 4, nearly 97.5% with respect to May 5, more than 95% compared to just six minutes earlier, and still a frightening 88.34% with respect to the following day.

Table no. 6 – Comparison of liquidity

03-May				04-May			
time	**qty**	**price**	**value**	**time**	**qty**	**price**	**value**
184512999	1070	1199.75	64,186,625	184512091	518	1168.5	30,264,150
184512662	1381	1199.5	82,825,475.0	184509250	1439	1168.25	84,055,587.5
184513002	1854	1199.25	111,170,475	184511951	1135	1168	66,284,000
184511490	1743	1199	104,492,850	184508083	1261	1167.75	73,626,638
184512880	1582	1198.75	94,821,125	184506049	2644	1167.5	154,343,500
184510416	1976	1198.5	118,411,800	184511971	1545	1167.25	90,170,063
184510416	1522	1198.25	91,186,825	184513028	2961	1167	172,774,350
184512504	2289	1198	137,111,100	184506049	1345	1166.75	78,463,938
184510934	2435	1197.75	145,826,063	184511951	1365	1166.5	79,613,625
184509913	1591	1197.5	95,261,125	184510974	1090	1166.25	63,560,625
TOTAL	**17,443**		**1,045,293,463**	**TOTAL**	**15,303**		**893,156,475**
05-May				**06-May-a**			
time	**qty**	**price**	**value**	**time**	**qty**	**price**	**value**
184513074	667	1159.5	38,669,325	183900506	957	1121	53,639,850
184512746	956	1159.25	55,412,150.0	183900444	1432	1120.75	80,245,700
184512533	1719	1159	99,616,050	183900530	836	1120.5	46,836,900
184512643	1046	1158.75	60,602,625	183900540	602	1120.25	33,719,525
184512788	1521	1158.5	88,103,925	183900539	575	1120	32,200,000
184512555	1038	1158.25	60,113,175	183900407	503	1119.75	28,161,712.5
184511796	2021	1158	117,015,900	183900407	572	1119.5	32,017,700
184512920	1404	1157.75	81,274,050	183900358	414	1119.25	23,168,475
184512144	1058	1157.5	61,231,750	183900411	412	1119	23,051,400
184512935	929	1157.25	53,754,263	183900408	386	1118.75	21,591,875
TOTAL	**12,359**		**715,793,213**	**TOTAL**	**6,689**		**374,633,138**
06-May-b				**07-May**			
time	**qty**	**price**	**value**	**time**	**qty**	**price**	**value**
184513871	39	1074.5	2,095,275	184513147	84	1112	4,670,400
184513795	45	1074.25	2,417,062.5	184512502	242	1111.75	13,452,175.0
184513814	24	1074	1,288,800	184512922	269	1111.5	14,949,675
184513821	48	1073.75	2,577,000	184512926	307	1111.25	17,057,688
184513821	11	1073.5	590,425	184511207	390	1111	21,664,500
184513821	19	1073.25	1,019,587.5	184509505	247	1110.75	13,717,762.5
184513821	16	1073	858,4	184510267	305	1110.5	16,935,125
184513835	32	1072.75	1,716,400	184511611	308	1110.25	17,097,850
184513694	45	1072.5	2,413,125	184512389	258	1110	14,319,000
184513795	31	1072.25	1,661,987.5	184511808	249	1109.75	13,816,387.5
TOTAL	**310**		**16,638,063**	**TOTAL**	**2,659**		**147,680,563**

Table no. 7 – Loss of liquidity on May 6

Date	03-May	04-May	05-May	06-May-a	06-May-b	07-May
Liquidity	17443	15303	12359	6689	310	2659
	98.22%	97.97%	97.49%	95.37%		88.34%

Overall, during the critical time on May 6 the combination of high volatility, stop-loss orders and scarce liquidity had an enormous impact on the price drop. The runs were on average much longer than on any other day, there were more numerous and higher peaks, and even the short runs were larger in number. The beginning of a run was often the alarming start of a dramatic downward price movement, such movements were large and frequent, and down-ticks at millisecond level was a common occurrence. As far as the Flash Crash is concerned, it is worth noticing that all the price movements larger than 1.25 index point occurred in the last 30 seconds before the halt (see Table no. 3).

5. STATISTICAL ANALYSIS

In order to analyse correlation between liquidity and price movements, the following approach has been adopted. The analysis starts off with the pairs made up of price and liquidity at top level of the bid book. Since the data only provides information at irregular time intervals, namely when there is a change at the top-of-book, the first operation needed is normalising the data. In order to do so, it is necessary filling up the time gaps by adding information for the time intervals during which no change occurred, simply by increasing the time by one millisecond until the price changed. For all the newly included data, no change occurred in either price or liquidity, and therefore the same values were copied in. For example: at time 18:39:00.605 liquidity at top of the bid book was 966 contracts, with a bid price of 1121 USD. Both remained stable until time 18:39:00.611, when a new contract was inserted into the bid book, at a price of 1121.25 USD. The gap between millisecond 605 and 610 has been filled up with liquidity-price pairs of 966-1121. Then, for each price level, the average liquidity was computed, together with the resting time. The result is a list of 'Average liquidity' versus 'Resting time' pairs, meaning that for a certain price, over that period of time, average liquidity has been worked out. In the six minutes, 28 seconds and 114 milliseconds until the CME Stop Logic triggered there were 2,939 price changes, and for each one of them average liquidity and resting time has been computed. The correlation between these two factors is 0.06461, meaning a very weak correlation between liquidity and permanence of a specified price at top-of-book over time.

The second test takes price change, especially when coupled with long runs, as proxy of stop-loss order execution, as explained above. In this case the correlation between liquidity at top-of book and price change yields values very much higher, in the range 0.46 through 0.66 according to the time interval considered, as shown in Table no. 8.

Table no. 8 – Correlation between liquidity and stop-loss-driven price change

TIME INTERVAL	CORRELATION FACTOR
18:39:00-18:42:00	0.45961
18:39:00-18:43:00	0.65588
18:39:00-18:44:00	0.65431
18:39:00-18:45:00	0.50233

The two calculations are not directly comparable and it would be a gross mistake doing so. Nevertheless, where in the first test it could be concluded that liquidity is, at some approximation, uncorrelated to price change, in the second test stop-loss orders and liquidity show relatively high correlation, although other factors should also be taken into account.

It is worth repeating: correlation does not imply causation and therefore the analysis carried out above cannot be taken as a proof of causal relationship in any way. However, the huge difference in correlation factor between the two tests, goes exactly the same direction of a suggestive evidence found by the other analyses carried out in this study: an impact of liquidity and stop-loss orders toward strong, and sometimes even violent, price changes.

6. DISCUSSION

The result that can be drawn from the previous observations is that the Flash Crash on the E-mini S&P 500 futures was, at some considerable extent, caused by a combination of three apparently innocuous factors: volatility, stop-loss orders, and low liquidity. Similar to a butterfly flapping its flies in Beijing, each of these conditions can, and does, often appear in the market without necessarily resulting in a memorably negative day, and even two conditions at the same time can co-exist without leading to financial hurricanes. When all these three conditions are present, as on May 6, 2010, chances seem higher for a local crisis to materialise. Obviously, not all three factors carry the same weight and play the same function in a crisis. All three conditions identified as the main contributors to the Flash Crash, are, individually taken, rather common. Stop-loss orders frequently occur in day-to-day operations for risk management purposes; they are normal practice and no financial authority is worried about them. Nevertheless, they can be important crisis contributors, especially when many of them pile up, ready to trigger as the price movements take a definitive direction. It is also intuitive that in a stressed market most investors would be cautious enough to protect their trading with stop-loss orders, even against small swings (a low λ in the variable $N_{Cn\lambda}$ of equation 2). Volatility is also a rather common occurrence. Sharply falling, as well as rising, prices are frequent events; they are intrinsic to market practices - and fortunately so: frozen markets are not desirable from any participant's point of view. Investors look for price dynamics and lack of it would make trading activities unappealing. Scarce liquidity is a different kind of beast; it is a major threat on itself. Regulators and exchanges are engaging a full-time struggle to ensure more and more abundant liquidity. High-Frequency Trading has found several supporters on the basis that the practice tends to increase market liquidity. Nevertheless, on May 6 liquidity virtually disappeared (Foresight - Government Office for Science, 2012; Madhavan, 2012) and lack of it exacerbated the combined effect of the other two factors, where stop-loss orders were the trigger, as shown in the previous section. However, even scarce liquidity on itself is not the automatic cause of a major crisis: if prices are stable the macro-effect would be scarcely noticeable. There are securities, and even entire markets, occasionally or permanently affected by scarce liquidity but they do not necessarily experience daily crises. Therefore, from all the previous considerations it sounds sensible, in accordance to the Chaos Theory (Gleick, 2008), to state that the real cause of the Flash Crash was the non-linear combination of high volatility, stop-loss orders and scarce liquidity. Had the market had the capability to prevent apparently innocuous causes to turn into a violent outcome, to avoid prudential withdrawals from a downward market leading to a crash, the disaster of May 6, 2010 might not have happened. The butterfly effect seems to be at the root of unresolved problems markets are currently facing. Systems have apparently grown too complex and too rapidly for systems theory to be able to cope with.

7. CONCLUSION

The data analysis presented in this study showed a weak capability of the markets to cope with the effects of a few common factors combined together. The title of this paper recalls a famous example, proposed by the Theory of Chaos, according to which a butterfly flapping its wings in Beijing might cause hurricanes in the Caribbean. The underlying logic is that dynamic systems (and financial market are widely recognised to belong to that category) heavily depend on the amplifying effects described by the Theory of Chaos, where a small difference in the initial conditions (the butterfly) might cause huge differences in the outcome (the hurricane). A Bank of England paper (Anderson *et al.*, 2015) defines the concepts of market 'amplifiers' and 'stabilisers'. "Amplifiers are market dynamics that act to reinforce buying or selling pressure in response to an initial price move" (ibid. p.16), while stabilisers act in the opposite direction. Both dynamics depend on the market structure and the nature, preferences, goals, investment horizons, beliefs, feelings and strategies of the participants. The stability of a financial system is a feature encompassing mechanisms that swiftly bring it back to equilibrium, that is, in which stabilisers dominate amplifiers, whenever a disturbance arises. A disturbance can be a price misalignment, so in this respect HFT, with its arbitraging capability is definitely a factor increasing the stability of the system. HF traders, despite usually being non-registered market makers, are strong contributors to liquidity supply and the literature confirms that they are more often supplying liquidity by quoting limit orders rather than aggressively taking liquidity away from the market. Volatility is another hot topic in trading and once again, albeit with some very noticeable exceptions, HFT seems to be a mitigating factor in this respect, as fast traders are also quick to close positions as soon as a minuscule profit can be extracted from a price swing, in so doing driving prices back to the mean. All these effects are intrinsic to market behaviour, usually very well tolerated, and often sought after, by the system. The greatest insurance markets have to guarantee orderly functioning is provided by the large amount of participants - and the different strategies among them. Their number facilitates mean price reversion in case of misalignment, return to equilibrium when volatility is excessive and a nearly continuous supply of liquidity. A market may display abundant liquidity but perhaps at a different price from the one traded last: liquidity may be abundant somewhere but someone will have to accept a loss to make use of it. The opposite case is more troublesome: if liquidity is low even a moderate level of trading has the potential to make prices volatile. However, this is far from uncommon occurrence: no market can be guaranteed to be highly liquid at all times. But this is usually a temporary condition; sooner rather than later other participants, noticing the scarcity of liquidity, will judge it profitable to supply limit orders, restoring the normal functionality of the market. As long as there is plenty of players, dis-homogeneity is statistically more likely than homogeneity, different views are the norm and the amplification risk seems to be easy to keep under control. Yet, something different happened on May 6, 2010 and in other mini-flash crash crises. In the few minutes labelled as 'Flash Crash' volatility spiked high, liquidity virtually disappeared, and the number of participants dried up. At that point all the market stability mechanisms usually relied upon for restoring 'normality' (whatever it might mean) failed miserably with well-known results. Such an abnormal situation can only be cured by dramatic interventions - in the Flash Crash case, by the CME Globex Stop Logic mechanism. The Chaos Theory studies the behaviour of stable-to-unstable transition in dynamic systems and more use of it in the field of finance is desirable to better understand, and to cope with, complex market dynamics.

ORCID

Gianluca P. M. Virgilio http://orcid.org/0000-0002-1616-0557

References

Anderson, N., Webber, L., Noss, J., Beale, D., and Crowley-Reidy, L., 2015. The resilience of financial market liquidity. *Financial Stability paper, 34*(October). https://bit.ly/2WLgCAs.

Bullock, S., 2011. Prospects for large scale financial systems simulation. *The Future of Computer Trading in Financial Markets. Foresight Driver Review – DR 14*. https://bit.ly/2Un1Cax.

Cespa, G., and Foucault, T., 2014. Illiquidity Contagion and Liquidity Crashes. *The Review of Financial Studies, 27*(6), 1615–1660. http://dx.doi.org/10.1093/rfs/hhu016

CFTC-SEC Advisory Committee on Emerging Regulatory Issues, 2010a. Findings Regarding the Market Events of May 6, 2010. https://bit.ly/2xnYrX4.

CFTC-SEC Advisory Committee on Emerging Regulatory Issues, 2010b. Preliminary Findings Regarding the Market Events of May 6, 2010. https://bit.ly/33M68T0

Cliff, D., 2011. *Regulatory scrutiny of algorithmic trading systems: an assessment of the feasibility and potential economic impact*. Government Office for Science, EIA16. https://bit.ly/3dGv8Qi

Danielsson, J., and Zer, I., 2012. Systemic risk arising from computer based trading and connections to the empirical literature on systemic risk. *Government Office for Science, DR29*. https://bit.ly/2wE4I0N.

Durden, T., 2012. Was the SEC "Explanation" Of the Flash Crash Maliciously Fabricated or Completely Flawed Out Of Plain Incompetence? *ZeroHedge*. https://bit.ly/3dBLKsb.

Foresight - Government Office for Science, 2012. Future of computer trading in financial markets: final report. https://bit.ly/39iQsrG.

Gleick, J., 2008. *Chaos - Making a New Science*. New York: Penguin Books.

Goldstein, I., Li, Y., and Yang, L., 2013. Speculation and Hedging in Segmented Markets. *The Review of Financial Studies, 27*(3), 881-922. http://dx.doi.org/10.1093/rfs/hht059

Granger, C., 1969. Investigating Causal Relations by Econometric Models and Cross-spectral Methods. *Econometrica, 37*(3), 424-438. http://dx.doi.org/10.2307/1912791

Kirilenko, A., Kyle, A., Samadi, M., and Tuzun, T., 2017. The Flash Crash: The Impact of High Frequency Trading on an Electronic Market. *The Journal of Finance, 72*(3), 967-998. http://dx.doi.org/10.1111/jofi.12498

Madhavan, A., 2012. Exchange Traded Funds, Market Structure and the Flash Crash. *Financial Analysts Journal, 68*(4), 20-35. http://dx.doi.org/10.2469/faj.v68.n4.6

Menkveld, A., and Yueshen, B., 2019. The Flash Crash: A Cautionary Tale About Highly Fragmented Markets. *Management Science, 65*(10), 4470-4488. http://dx.doi.org/10.1287/mnsc.2018.3040

Nanex, 2010a. Analysis of the "Flash Crash" - Intro Empirical Analysis. *www.nanex.net*. from www.nanex.net

Nanex, 2010b. Analysis of the "Flash Crash" - Part 1 Empirical Analysis. from www.nanex.net

Schapiro, M., 2010. Testimony Concerning the Severe Market Disruption on May 6, 2010. *Before the Subcommittee on Capital Markets, Insurance and Government Sponsored Enterprises of the United States House of Representatives Committee on Financial Services*. https://bit.ly/3bu1xHE.

Sornette, D., and von der Becke, S., 2011. Crashes and High Frequency Trading. *Government Office for Science, DR7*. https://bit.ly/33LgJh8.

Zervoudakis, F., Lawrence, D., Gontikas, G., and Al Merey, M., 2012. Perspectives on High-Frequency Trading. *University College London White paper*. https://bit.ly/2wv6PUL.