NO. 16 | DECEMBER 2017

Defining the FX Flash Crash

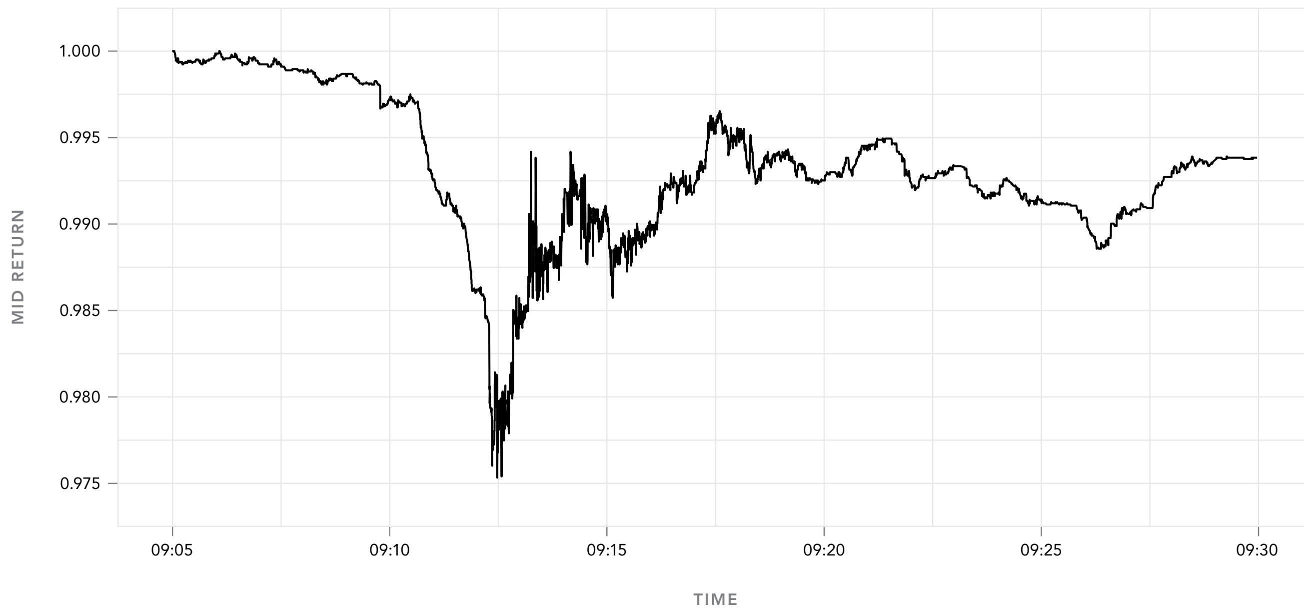

FIGURE 1 A "classic" flash crash. Midpoint return of AUD on 2015-08-24 around 09:10 New York time.

Introduction

FX flash crashes—sudden and dramatic price moves that are attributed at least partly to some kind of technical market failure—have been widely reported in the media and discussed at industry events. They have reportedly increased in frequency over the past several years, with possible culprits including banks pulling back from capital commitment; feedback effects among automated liquidity providers; increasing usage of algorithmic trading; and the juniorization of trading personnel.

Reports so far have tended to look at individual events in isolation, and discussion of the recent trend at industry events has been correspondingly anecdotal.

In this research note we look systematically at extreme volatility in the FX markets, and arrive at an operational definition of the *flash crash* phenomenon which identifies the widely publicized flash crashes, as well as a number of similar (if less extreme) events. We propose that systematically tracking the incidence of flash crashes using this definition would provide a useful quantitative index for regulators and industry participants interested in understanding how ongoing changes to technology, regulation, and industry practices affect market quality over time.

What is a flash crash?

The intuitive definition, reasonably summarized in the BIS' report on the October 2016 Sterling Flash Crash, is a large, fast, v-shaped price move and sudden widening of bid-offer spreads. Flash crashes are generally considered to be technical market failures that are distinct from ordinary volatility, presumably triggered when liquidity providers can't keep up with sudden demands of heavy but informationless directional volume. By contrast, market moving news

FIGURE 2A Frequency distribution of 5-minute high-low ranges.

FIGURE 2B Frequency distribution within the 99th percentile.

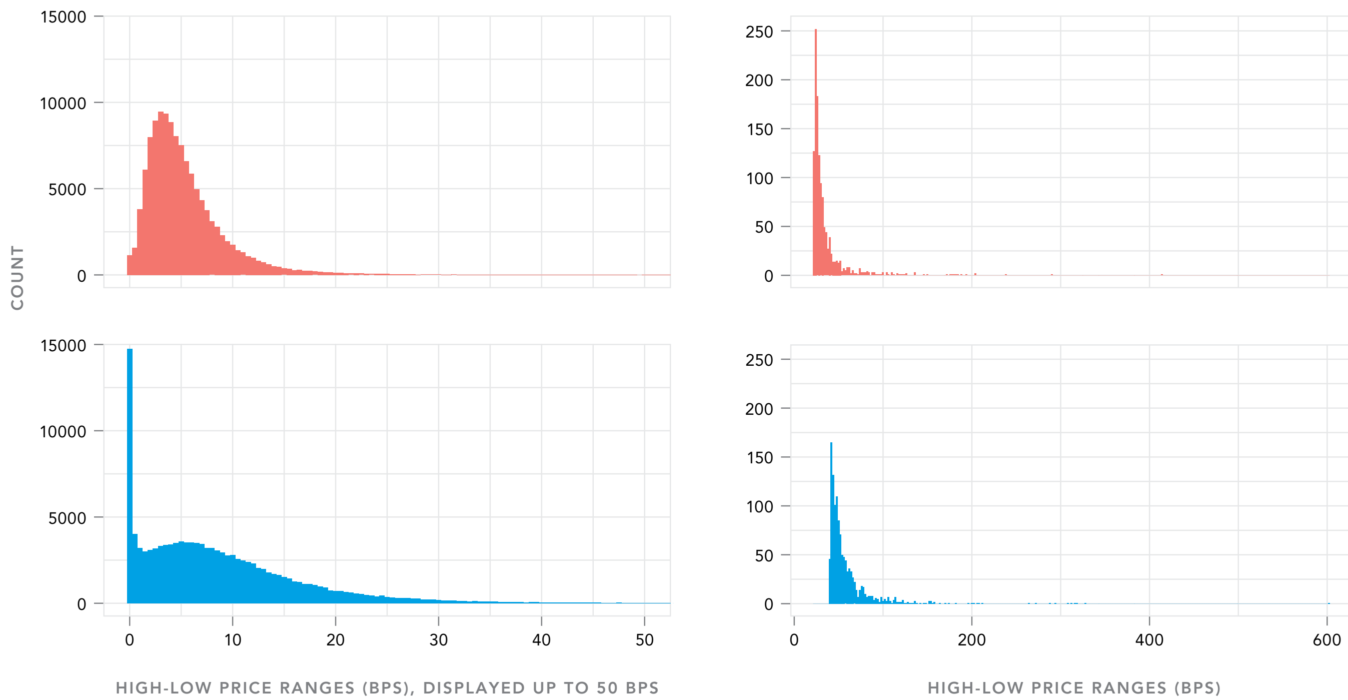

Frequency of 5-minute high-low price ranges, normalized by time-of-day variation in volatility, for EURUSD and USDZAR.

and events can cause sudden changes in price, and spikes in volatility and spreads, but are not considered flash crashes. In such events, price generally does not revert but converges to a new level reflecting the new information arriving to the market. We analyze each of these elements in turn.

Data and methodology

The core dataset for our analysis is the tick by tick quotes of some of the commonly traded currencies pairs. We mostly focus on the more liquid currencies—AUD, EUR, GBP, CAD, CHF, and JPY, and examine data across 2015 and 2016. From the tick data, we compute 5-minute bars of high, low, open, and close mid prices, number of quote updates, etc. Bar volatility of the close-to-close return is computed specific to the pair and hour of the day, as volatility varies over the course of the day in a characteristic pattern. Characteristic volatility is estimated over the past seven days. Average spreads within a bar are time-weighed. Finally, we filter out data in the 4PM-7PM New York time, to avoid the disproportionate number of data problems occurring during this time frame, for example resulting from market participants' daily operational cycle.

Speed and magnitude

The primary importance of speed and magnitude is suggested by the name itself—*flash* implies speed and *crash* implies magnitude. Correspondingly, we use the high-low price range within a 5-minute bar as our first metric.

Figure 2a above shows how often we see various sized moves in two pairs, one liquid, EURUSD, and one relatively illiquid, USDZAR. We can see from the figure that the bulk of moves within a 5 minute window are within a fairly narrow "ordinary" range: 80% are within 7 bps for EURUSD and 13 bps for USDZAR. Figure 2b zooms in on the distribution for the most extreme events—the 99th percentile. This figure shows that extreme events tens of times bigger than this "ordinary" size do occur with some frequency. However, not all extreme moves are necessarily flash crashes; in Figures 2-4 we look at a few examples. Note, the spike at zero for ZAR reflects the significant number of 5-minute periods in which there was no price change.

Figure 1 on the previous page shows what we might call a classic flash crash—rates crash by 230 pips within 3 minutes—on the order of 37 times the 5-minute volatility observed in the preceding week

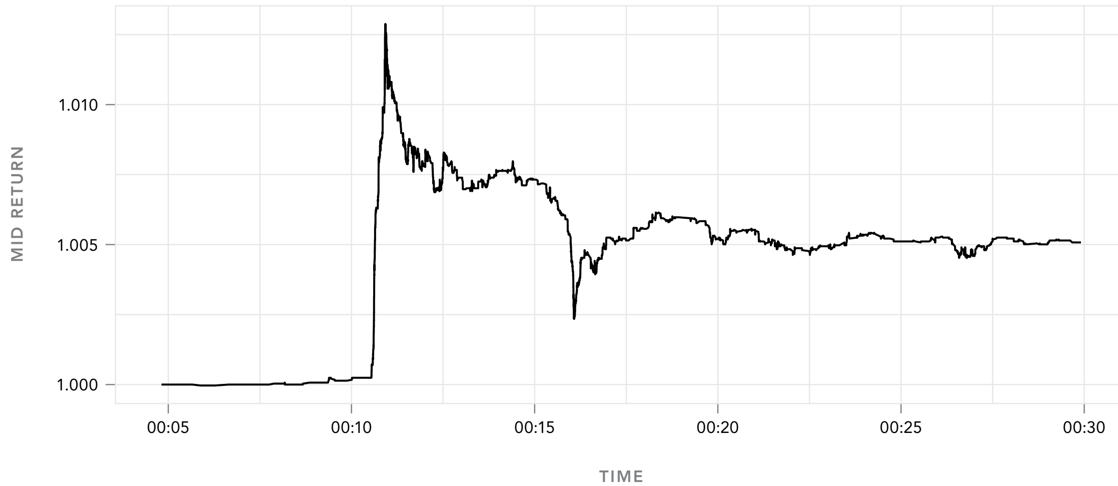

FIGURE 3

Ambiguous event with features of both news driven price move and technical failure. Midpoint return of GBP 2016-06-07 (~60% recovery).

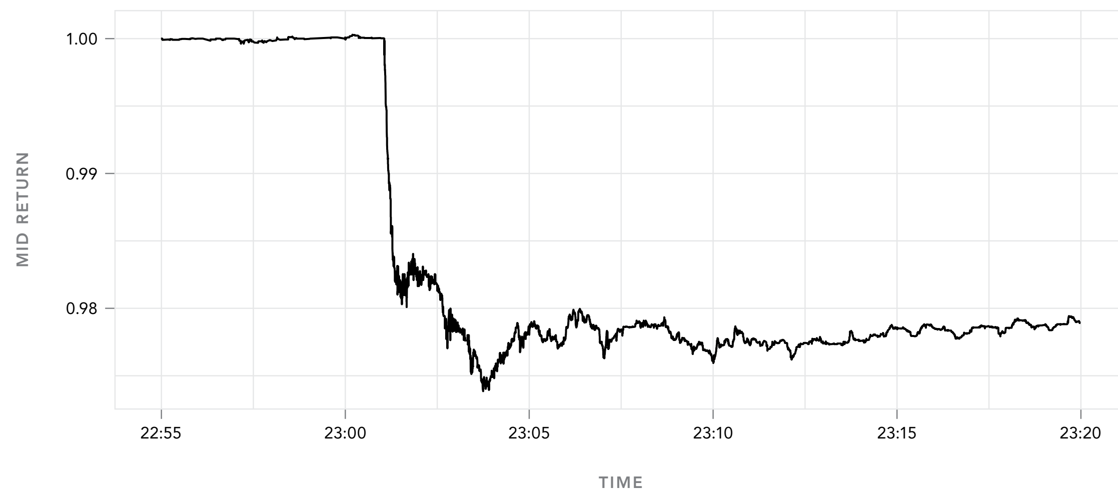

FIGURE 4

News driven market move. JPY 2016-04-27 (10% recovery).

during the same hour—and then retrace 90% of the move in the succeeding 5 minutes.

Figure 3 starts out looking similar to the event depicted in Figure 2—a 127 pip move within 1 minute—but after some recovery, the price settles down having retraced only about 60% of the move after 15 minutes. This example highlights a challenge in defining flash crashes. The price move was fast and large. Yet there was a very significant residual price move 20 minutes later. 40% of the original move was permanent at this time scale. The event looks like some rough price action around a major price move; that is, price formation and not a pure technical market failure.

Finally, Figure 4 starts out similarly again, with a 270 pip move within three minutes, but prices have only retraced about 19% of the move after 15 minutes. The context of this event was a decision by Bank of Japan to keep monetary policy steady, surprising market participants who were expecting further stimulus.

Reversion

With these examples in mind, we add a third component to our operational definition of flash crash—*reversion*, or what percentage of the initial high-low range has been retraced 15 minutes later. A V-shape, with prices quickly retracing the drop, as in Figure 1, strongly suggests that the initial movement was spurious. Conversely, if the price drops suddenly but then stays at a new equilibrium, as in Figure 4, it suggests price discovery, and arguably an efficiently functioning (albeit unpleasant) market—not a flash crash.

One might suppose that the difference between these examples can be attributed to the existence or non-existence of news. However, in examining many examples we found that events do not fall neatly into two categories. Many, like Figure 3 above, fall in the middle, and very few large price move events that persist were clearly explained by fundamental information.

In Figure 5 we look at reversion together with the

FIGURE 5

Extreme price move vs reversion. Magnitude, on the X axis, is plotted as the high-low range in the 5-minute bar as a multiple of the 5-minute volatility for that time of day for that pair. Reversion, on the Y axis, is plotted as a percentage of that magnitude that's retraced from the low price in that bar to price about 15 minutes later. The red '×' marks the AUDUSD event in Figure 1.

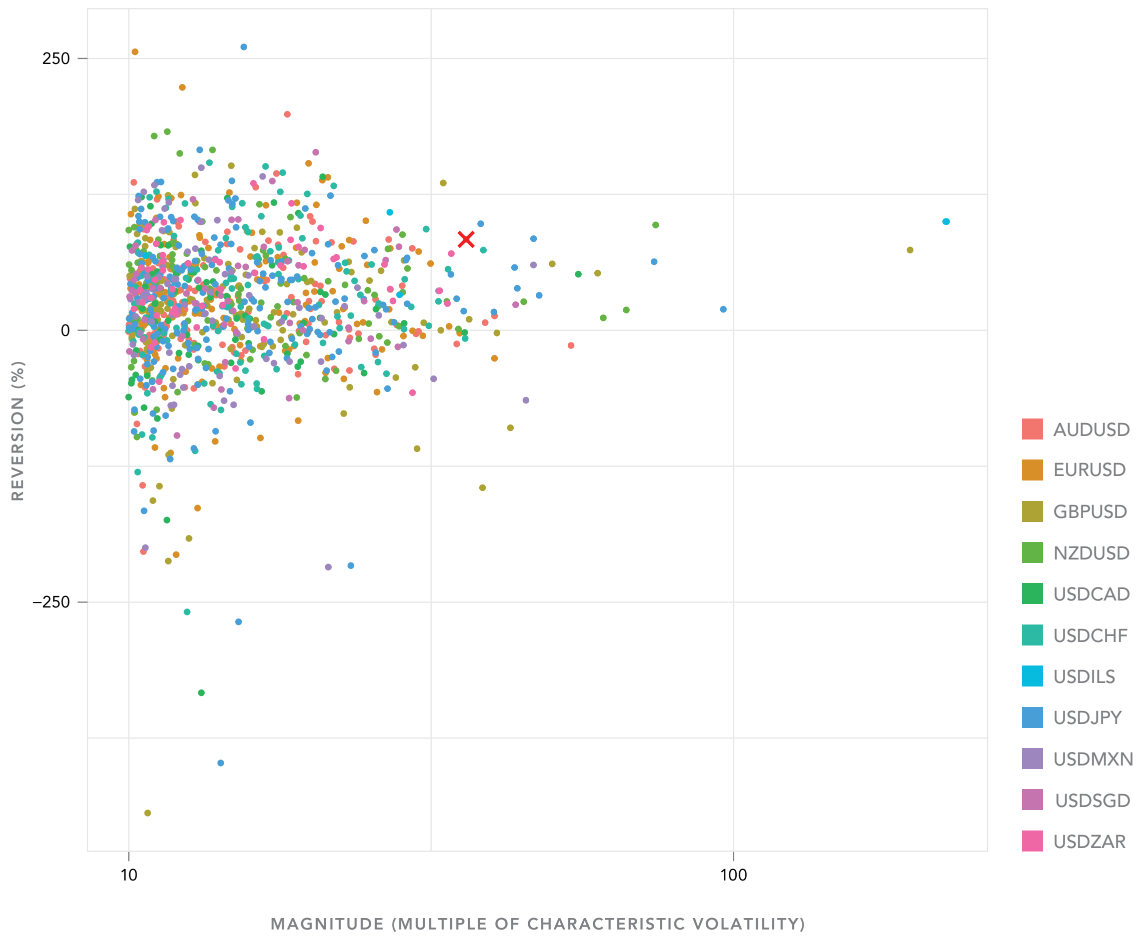

magnitude of the original price move for some of the most extreme events—those that have more than ten times the average 5-minute volatility. For the pairs shown in the plot, this is about 1100 events, or about 0.3% of the data.[1] We might conjecture that extreme price moves are more likely to be flash crashes and therefore show reversion, or that we will see a distinct cluster of large price move, high-reversion events that appear to be a different class of events than other extreme price moves. However, the figure shows no clear-cut boundary between flash crashes and ordinary extreme volatility. Rather, there is a continuum. This suggests that, at least with respect to magnitude and reversion, flash crashes are not a distinct class of event, but rather outliers on a continuum.

Spreads

If we think of a flash crash as a technical market failure, another defining characteristic should be wide spreads. For example, during the event illustrated in Figure 1, spreads spike from a few pips five minutes before the event to 40, even 70 pips during the event, before largely settling down about 10 minutes later, as shown in Figure 6. This kind of behavior is certainly consistent with our overall view of flash crashes as a failure of liquidity providers to keep up with demands of liquidity takers. However, the events of figures 3 and 4 also showed very wide spreads (not illustrated), though we asserted that because there wasn't dramatic reversion, they didn't really meet our criteria for a flash crash. Wide spreads can be caused by the uncertainty in price accompanying any major episode of volatility, and aren't sufficient to define flash crashes. But perhaps when we look at our criteria together, we will see a pattern?

Figure 7 brings magnitude, reversion, and spread

1 The 11 currency pairs in our dataset that predominantly trade against the USD are AUD, EUR, GBP, NZD, CAD, CHF, JPY MXN SGD, ZAR, and ILS. Our dataset includes about 342,000 5-minute bars.

FIGURE 6 Co-evolution of spread and price for the AUD flash-crash of Figure 1.

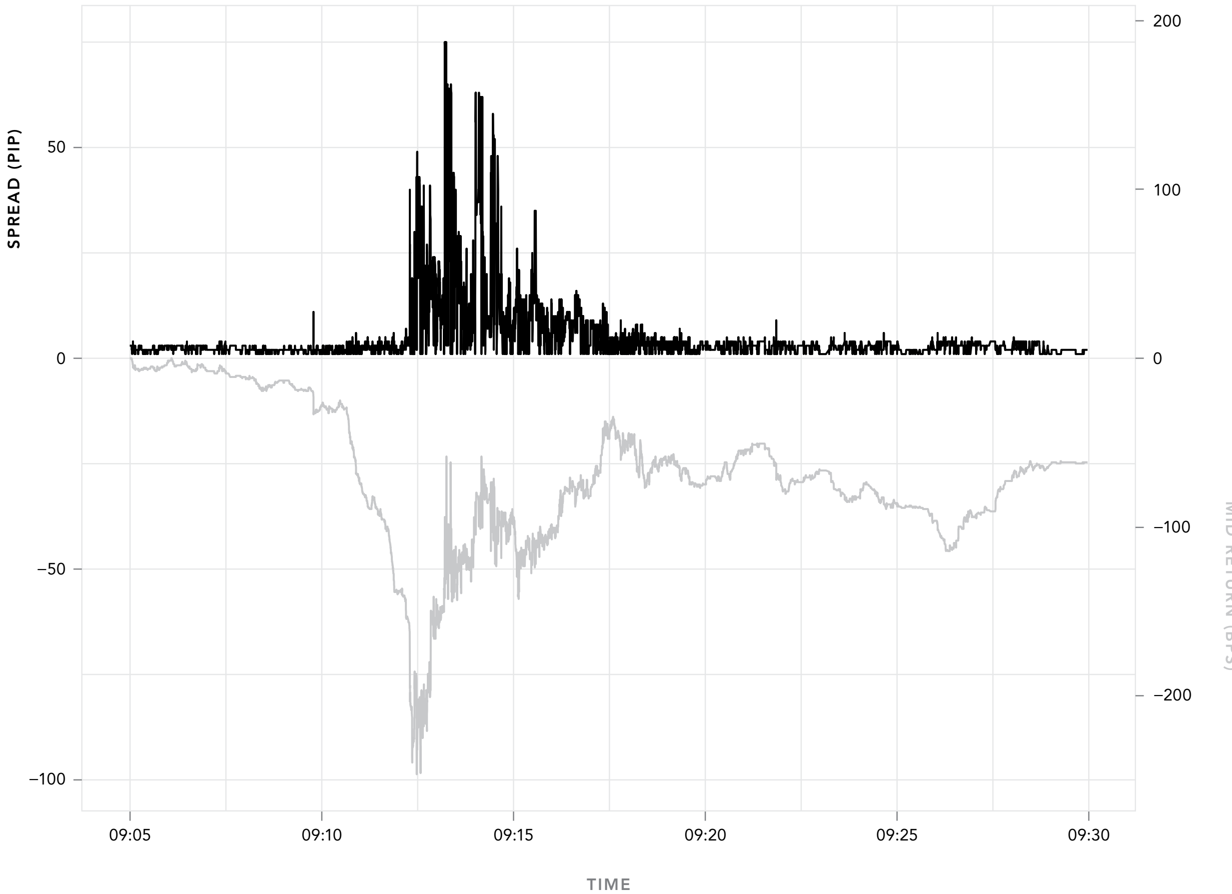

together in a single graph. Each point represents one extreme event with a magnitude of at least 10× the normal 5-minute volatility, in the same way as in Figure 5. As previously defined the X axis shows the magnitude of the price move, and the Y axis shows the amount of reversion after about 15 minutes. The size of the spread is indicated by the color, with ordinary spreads (less than 2× the normal spread for the time of day) in red and wide spreads (>2×) in blue.

Defining the flash crash

The dotted lines in Figure 7 divide the space into four quadrants. The blue points in the upper right quadrant represent our definition of flash crash: a large price move (>13× normal volatility), strong reversion (>70%), and widening spread (>2× normal spread; indicated here by blue points). Out of over 313K candidate events, this criterion flags 69 events, including the known flash crashes reported in the media, for example, the sterling flash crash on Oct 6, 2016, and the Swiss franc event on Jan 15, 2015. Our main objective, however, is to use this approach to identify similar though potentially smaller events that have not been widely reported in the popular press. This group of identified flash crash events in our 2-year sample for the 11 currency pairs is effectively the 0.02% percentile of 5-minute episodes in this data set. Because the distribution of events is fairly continuous, the precise thresholds are arbitrary, and small changes in the chosen values will cause relatively small changes in the set of events identified.

Summary

Flash crashes are extreme market events characterized by a large, fast price move followed by reversion, and a sudden and significant widening of the bid-offer spread. These criteria do not define a discrete cluster of events, but rather a continuous range of extreme events. We define flash crash using a set of quantitative thresholds for each of these dimensions.

FIGURE 7

Extreme price move vs reversion (same as in Figure 5). Blue indicates spread more than 2x the mean; Red indicates less than 2x. The vertical dotted line is at price move of 13x volatility. The horizontal dotted line is at 70% reversion.

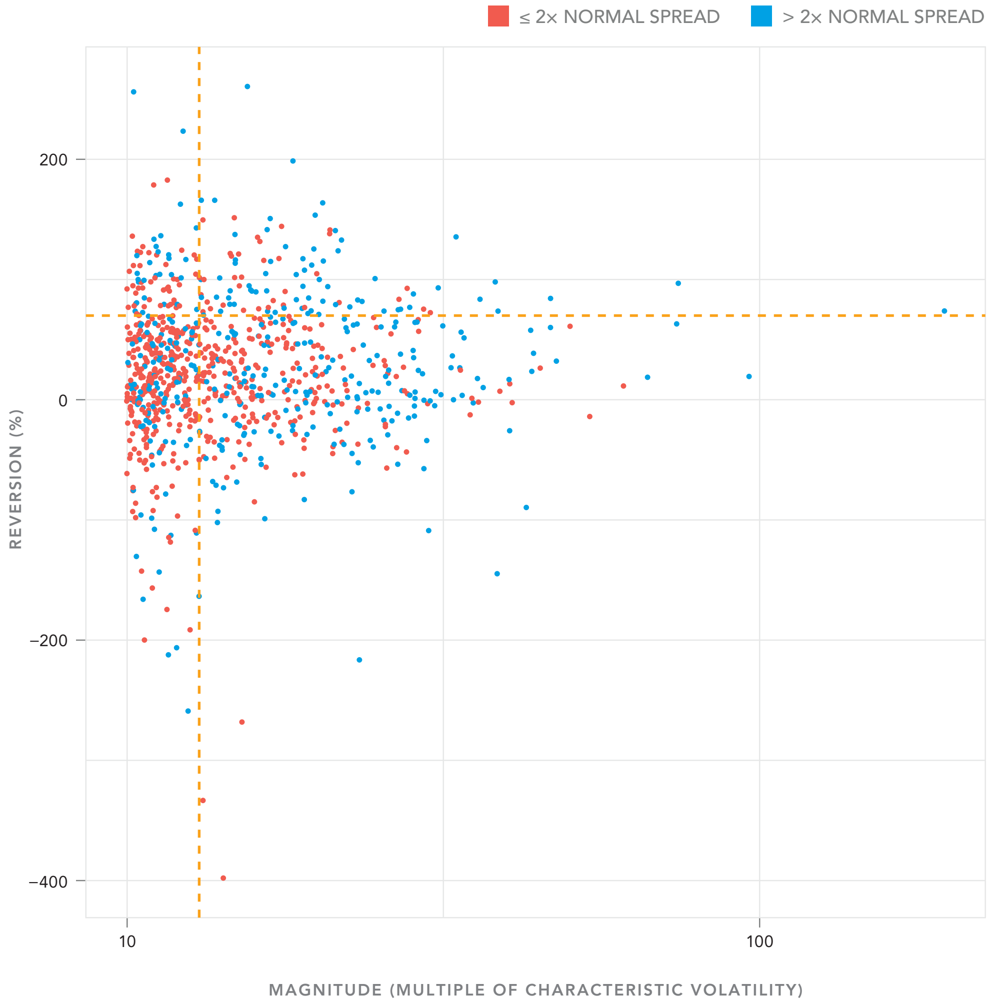

This definition identifies the flash crashes that have been widely reported in the media, as well as a set of similar (though less extreme) events.

Pragma's flash crash definition has several desirable features. It normalizes events based on typical volatility for the pair and time of day, so it allows comparison between and averaging across events. It also normalizes relative to recent history, making it more useful for tracking the incidence of extreme events against a background of volatility that rises and falls significantly. Finally, because it identifies dozens rather than a small handful of events per year, it escapes the realm of anecdote and provides a useful statistic.

Accordingly, we propose that systematically tracking the incidence of flash crashes using this definition would provide a useful quantitative index for regulators and industry participants interested in understanding how ongoing changes to technology, regulation, and industry practices affect market quality over time. Over the two-year period of the study, there was no clear trend—up or down—in the incidence of flash crash events.

PRAGMA IS AN INDEPENDENT PROVIDER OF MULTI-ASSET CLASS ALGORITHMIC TRADING TECHNOLOGY AND ANALYTICAL SERVICES.
FOR FUTURE RESEARCH UPDATES FOLLOW OUR BLOG AT BLOG.PRAGMATRADING.COM.

Copyright © 2017 Pragma Securities. All rights reserved. Do not reproduce or excerpt without permission.
Pragma Securities LLC. Member of FINRA and SIPC. C.A. #176