Modeling Flash Crash Behavior in a Stock Market using Multivariate Hawkes Processes

Feng Shi
Amazon
Seattle, WA, USA
and
University of North Carolina at Chapel Hill
Chapel Hill, NC, USA

John Paul Broussard
Estonia Business School
Tallinn, Estonia
and
Rutgers University
Camden, NJ, USA

G. Geoffrey Booth
The Citadel
Charleston, SC, USA
and
Michigan State University
East Lansing, MI, USA

June 26, 2020

Acknowledgement: We thank Xinyuan "Bernard" Niu for his assistance during the beginning stages of this research project, Gregory Forest for his ongoing encouragement, Elizabeth Booth for her careful editing and Nanex for supplying the intraday stock price and related data for the 30 Dow Jones Industrial Average stocks. We also thank Régis Gourdel and the other participants of the December 2019 Paris Financial Conference as well as the attendees of a February 2020 Stockholm Business School seminar. Substantial portions of this research were undertaken while Shi was at the University of North Carolina – Chapel Hill, prior to joining Amazon. The views contained herein are those of the authors and do not necessarily reflect the views of any of their named current or past employers or affiliations.

Modeling Flash Crash Behavior in a Stock Market using Multivariate Hawkes Processes

Abstract

This paper uses multivariate Hawkes processes to model the transactions behavior of the U.S. stock market as proxied by the 30 Dow Jones Industrial Average stocks before, during and after the May 6, 2010 flash crash, which lasted 36 minutes. The basis for our analysis is the excitation matrix, which describes the network of interactions among the stocks. Using high-frequency transactions data for individual stocks, we find, among other things, strong evidence of contagion that is self- and asymmetrically cross-induced. Our descriptive findings have implications for stock trading and corresponding risk management strategies, as well as stock market microstructure design.

Keywords: Stocks, Crash, High Frequency, Hawkes Processes, Networks, Contagion

Introduction

U.S. stock markets, similar to stock markets throughout the world, typically exhibit rapidly fluctuating share prices that at times are characterized by large and often unexpected changes. Taken together, a sequence of price declines in a broad group of stocks in a relatively short time period that results in significant monetary losses is referred to as a market crash. These price declines may be caused either by major negative news that directly or indirectly affects the entire economy, or by a serious breakdown in market quality because of the way in which the market is designed. Typically, not all of the prices of individual stocks fall at the same time or at the same rate, suggesting the presence of some sort of contagious behavior among the stocks.[1] If price declines are extremely rapid, the observed phenomenon is referred to as a flash crash.[2] As Gao and Mizrach (2016) point out, small crashes of very limited duration and breadth are common and are often referred to as mini-crashes.[3] In contrast, large flash crashes are much rarer regardless of the length of time it takes for the overall market to recover.

---

[1] Glasserman and Young (2016) argue that interconnectedness is a characteristic of modern financial systems both global and domestic. Although this phenomenon may provide transactional benefits, it also enhances the fragility of the system through, e.g., common risk exposure via the ownership of similar assets and liquidity shocks. These risks may be difficult to manage because of network opacity.

[2] Booth, Booth and Broussard (2014) point out that advances in information technologies that began in the 18th century have resulted in transactions in today's stock markets to occur at a rate faster than the human mind is able to comprehend. Relying on the interviews of professional traders, Lewis (2014) indicates that today's traders actively exploit current high-speed market technology. In this regard, also see Johnson et al. (2013).

The most pervasive flash crash in recent U.S. history occurred Thursday, May 6, 2010. Although the U.S. major markets had been highly volatile since their openings as a result of the ongoing disappointing economic news from Europe concerning its then-ongoing debt crisis, the Commodity Futures Trading Commission (CFTC) and the Securities Exchange Commission (SEC) in joint reports (CFTC-SEC, 2010a, 2010b) indicate that the crash began about 2:32 p.m. EST and lasted until 3:08 p.m. (or 14:32 to 15:08 in 24 hour time nomenclature). For the next 13 minutes or so after 14:32, stock prices continued to drop, with the largest declines occurring in the last few minutes. At 14:45, the market began to stabilize, although there were still extreme movements in the prices of some stocks. Between 15:00 and 15:08, the markets became noticeably less volatile as prices approached their pre-crash values and orderly trading resumed. At the nadir of the crash roughly $1 trillion in market capitalization was lost.

It is unclear as to the cause of this crash. Was it an unusually large transaction, algorithmic trading (high frequency and otherwise), spoofing (entering fake orders to manipulate prices) or some other type of illegal trading scheme, a flaw in the market's microstructure design, or some other reason?[4] Regardless of the cause, the prices of many stocks irregularly fell because of some sort of a shock, and it took over 30 minutes for their prices to recover. The presence of this type of price behavior over time most likely renders many, if not all, of the current volatility-based risk management techniques advocated or used by academics and practitioners to be of little practical use to active portfolio managers and traders who focus on intraday transactions.

---

[3] Mini crashes are negative ultrafast extreme events (UEE) and spikes are positive UEEs. In a study of Nanex NxCore stock market prices from 2006 to 2011, Johnson, et al. (2013) find 18,250 UEEs of both signs that endured less than 1,500 milliseconds.

[4] The CFTC-SEC (2010a, 2010b) attributes the start of the crash to a large fundamental trader placing an order for 75,000 ($4.1 billion) E-mini S&P 500 futures contracts to hedge against an existing equity position. Aldrich, Grundfest and Laughlin (2017), however, suggest that the flash crash was caused by continued presence of many large sell orders and the corresponding widespread withdrawal of liquidity, i.e., the decrease in the number of contracts quoted close to the best price. Nevertheless, Kirilenko et al. (2017) document that the trading pattern of high frequency traders did not change during the crash. Menkveld and Yueshen (2019) highlight the impacts of a fragmented marketplace on the flash crash from an arbitrage perspective using E-mini and SPY data.

To assist in developing market design initiatives, trading strategies, and risk management models that incorporate an intra-day, contagion perspective, the purpose of this paper is to model the behavior exhibited by individual stocks during the 2010 flash crash using Hawkes processes, with special attention given to the way stocks interact.[5] The Hawkes process is a stochastic model that describes the time of occurrences of events within a specific time interval. What distinguishes the Hawkes process from the Poisson process, which has also been used to model jump behavior in stock prices, is that it allows the occurrence of an event to increase the likelihood of triggering more or fewer events in the near future, i.e., the occurrence of an event is dependent on the occurrences of past events.[6] Because of this ability, Hawkes processes have been used to quantify the endogenous and exogenous price effects in various stock markets. However, just as any other statistical model, the Hawkes process mines the statistical association between the stocks but cannot reconstruct the true causal relationship between them. We thus use the Hawkes model as a new lens on the behavior of stocks during flash crash, but we are not looking for its root causes.

Hawkes (2018) provides a brief summary of his process and gives a useful and extensive bibliography pertaining to the model's development and finance applications, with a special mention of Bacry, Mastromatteo and Muzy (2015) for high frequency applications. Applications that are particularly relevant to our study are Filimonov and Sornette (2012) and Aït-Sahalia, Cacho-Diaz, and Laevin (2015). Filimonov and Sornette (2012) examine the behavior of the E-mini S&P 500 contract, which is traded on the Chicago exchange, from 1998 to 2010. They report that from 1998 to 2010 the portion of the price changes attributed to the endogeneity of this financial instrument increased dramatically and reached almost

[5] There is an extensive literature on models for stock prices as continuous or discrete stochastic processes, including random walk models and advanced models that incorporate statistical phenomena such as autocorrelation, jumps, volatility clustering, multifractality, and various combinations thereof. Louis Bachelier (1900) is often thought of the founder of the notion that stock prices follow a random walk. Fama (1965) was largely responsible for introducing this concept, which supports the theory of efficient markets (prices reflect all relevant information) to the academic and professional communities. Subsequently, Makiel (1973) popularized the notion and implications of market efficiency to the general public.

[6] Merton (1976) is the first to model jumps in financial asset prices using a Poisson process, and his modeling approach was subsequently adopted by many researchers, e.g., Akgiray and Booth (1988), Andersen, Benzoni and Lund (2002), and Cai and Kou (2011).

100% during the flash crash. Aït-Sahalia, Cacho-Diaz, and Laevin (2015) extend this approach by combining Hawkes and diffusion processes to model the joint time series behavior of the S&P 500 (U.S.), FTSE 100 (U.K.), Nikkei 225 (Japan), Hang Seng (Hong Kong), and IPC (Mexico) stock indexes. They use daily open and close data for various sub-periods (because of lack of data for some indexes) within the overall time span beginning January 2, 1980 and ending April 30, 2013. They present significant evidence of endogenous behavior within each market and similar relationships between various market pairs, with the latter phenomenon suggesting the presence of contagion among the markets examined.

We find that the influences between the 30 DJIA stocks increase on average during the flash crash and then revert to approximately their pre-crash level. The level of influence, however, varies greatly among stocks before, during and after the crash. Moreover, stocks that are strongly influenced by other stocks only are weakly influenced by their past behavior and vice versa. For example, ExxonMobil (XOM), as determined by a combined ordinal ranking, is the most influential stock that affects the 29 other stocks, but it drops to 30 with respect to the other stocks impacting it. In contrast, Travelers (TRV) corresponding ranks are 30 and 9. International Business Machines (IBM), however, ranks in the middle (14) for both its impact on others, as well as how it is impacted by others. These influence differences may suggest potential worthwhile dynamic diversification strategies and market microstructure policies.

The remainder of the paper is as follows: We first discuss our data and its source as well as provide some important descriptive statistics. We then describe and justify our use of Hawkes processes. Our focus is on the excitation matrix and its application to determining the interactions of the 30 DJIA stocks. Finally, we present our empirical results and place our findings in an economic context, giving special emphasis to risk measurement and management.

Data and Descriptive Information

To explore the viability of the Hawkes processes to model the pre-crash, crash, and post-crash stock price behavior, we focus on the 30 stocks that comprised the Dow Jones Industrial Average (DJIA) at the time of the crash.[7] These companies, which are listed in Table 1, are very large, publicly traded, U.S.-based, and most are listed on New York Stock (NYSE) with the remainder being listed on Nasdaq (NQNM). Their stocks can be traded on their listing exchange, as well as at 10 other smaller exchanges, which are often referred to as reporting exchanges or trading venues. Virtually all of the companies are household names and represent almost all of the major sectors in the U.S. economy, i.e., consumer staples, industrial materials, financials, telecommunications, energy, consumer discretionary, information technology, and heath care. Of the major sectors only transportation and utilities are not represented by a stock in the index. As a group the 30 DJIA companies accounted for approximately 22 percent of the market value of all traded U.S. stock around the time of the flash crash.[8]

Insert Table 1 about here.

Our data were obtained from Nanex, which provides real-time option and stock price data via its NxCore product. Data are archived by Nanex as transactions and quotes arrive from the various exchanges, and are time-stamped at millisecond time intervals. When the data were collected, Nanex's timestamp was the most granular available. Nanex's data are generated by activity from all U.S. exchanges where a given stock is traded, which is not necessarily where it is listed. The primary data extracted from the Nanex feed

---

[7] The DJIA is a price-weighted index that began on May 26, 1896 with 12 stocks and was calculated as a simple average of the 12 component stocks. The number of stocks was increased to 20 in 1916 and then to 30 (its current number) in 1928. To keep the average to be time-compatible it had to be adjusted to account for the stock additions and deletions from the index as well as adjusted for stock dividends, splits, spinoffs, mergers, acquisitions and so forth. The adjustment is handled by the Dow Divisor, which was initially 12 but at the time of the flash crash its value had declined to approximately 0.132. In addition, The DJIA is routinely edited by *The Wall Street Journal* to ensure that the stocks in the index fairly reflect the overall U.S. economy. Since May 6, 2010, Alcoa (AA), American Telephone & Telegraph (T), Bank of America (BAC), DuPont (DD), General Electric (GE), Hewlett-Packard (HPQ) and Kraft Foods (KFT) have been replaced. As of July 1, 2019 the replacement stocks are Apple (AAPL), Dow (DOW), Goldman Sachs (GS), Nike (NIKE), United Health Group (UHG), Visa (V), Walgreens Boots Alliance (WBA). In this study's data, General Electric is the only survivor of the original 12 companies.

[8] Even though there was a suggestion that the E-mini S&P 500 futures contract might have caused the flash crash (see fn. 4), we do not include this financial instrument in our empirical analysis. We make this exclusion because our interest is in the interaction of common stocks from the crash through the recovery period.

used in our analysis are transaction prices and their time stamps. Menkveld and Yueshen (2019) also use Nanex data in their analysis of fragmented market impacts on the flash crash.

For each of the DJIA 30 stocks listed in Table 1, we provide its open, close, and low prices on May 6. The time each stock reached its lowest price is also included. We also plot the transaction price standardized by its open price for each of the 30 stocks in Figure 1[9]. For comparison, we include the standardized price series on the day (May 6, Figure 1-middle) of the flash crash as well as the standardized prices from one day before (May 5, Figure 1-left) and one day after (May 7, Figure 1-right). There are noticeable breaks in the price series between the days. This is the result of the markets closing at the end of the trading day, thereby enabling the effects of news and overnight trading to be acted on by the market at its opening the next day.

Insert Figure 1 about here.

The price series behavior on May 6 is markedly different from the adjacent two days, with large abrupt drops occurring around 14:30. Before and after the flash crash and its recovery, prices tend to move up and down in small increments and do not seem to follow a trend. Statistically, this pattern has been often modeled using a continuous-time Markov process with Brownian motion after converting the price series to continuous returns by taking the first difference of the natural logarithm of prices. Economically, this type of pattern is typically attributed to normal transactions activity such as not wanting to buy or sell an unusually large position in a short period of time.

To clearly see the timing and the magnitude of the price drops of each of the 30 DJIA stocks on May 6th, in Figure 2 we show the lowest standardized price for each stock (y-axis) and the time each stock reaches its nadir (x-axis). It is clear that the time points for these prices are clustered between 14:45 and 14:48 with Walmart (WMT) the first stock (14:45:29.2) to reach its lowest price and Kraft Foods (KFT) to be the last (14:47:58.8). Although most stocks dropped about 10 percent, 3M (MMM) dropped 20 percent and Procter & Gamble (PG) dropped more than 35 percent. The observations in Figures 1 and 2 are

[9] Subtracting 1.0 from the result of this standardization procedure creates a measure of return based on the stock's price at the beginning of the standardization period.

consistent with the reports from CFTC-SEC (2010a, 2010b) and also reveal the distinguishing feature of a flash crash, i.e., large cumulative declines in a very short time period and a corresponding rapid recovery.

Insert Figure 2 about here.

Statistical Methods and Approaches

We use the multi-variate Hawkes process to model the DJIA 30 stocks. In particular, we model the activity level of each individual stock and the timing of its activities. For each stock, we study its price-changing events: a price-changing event is a transaction at a price different from that of its previous transaction. Without ambiguity, we will refer to such price-changing events simply as events. Note that those events are not returns – they are transactions that correspond to price changes. Because there are multiple stocks of interest, events from different stocks are treated as unique. Moreover, following a comment by Aït-Sahalia, Cacho-Diaz, and Laevin (2015), we do not incorporate some type of statistical transition mechanism in our model because we are modeling the dynamics of a single crash and its recovery, rather than a series of similar crashes and recoveries that need to be linked together.

<u>The Hawkes Process Model</u>

Mathematically, let t_i^s denote the time of the *i*-th event of stock s with the stocks indexed from 1 to 30, i.e., $s \in \{1,2,\ldots,30\}$. The Hawkes process assumes that at any time *t*, the probability P_s that an event of stock s will occur in the next *dt* time units is determined by the instantaneous rate $\lambda_s(t)$: $P_s \approx \lambda_s(t) \cdot dt$.[10] The rate of events $\lambda_s(t)$ is modeled as a function of the occurrences of previous events from all the stocks (including itself):

$$\lambda_s(t) = \mu_s + \sum_{s'=1}^{30} \sum_{i:t_i^{s'}<t} a_{ss'} g(t - t_i^{s'}), \tag{1}$$

where

- $(i: t_i^{s'} < t)$ corresponds to all the events of stock s' that occurred before time *t*.

[10] Additional details on the Hawkes process may be found in numerous sources but Liniger (2009), Embrechts, Liniger and Lin (2011), Rizoiu, et al. (2017) and Hawkes (2018) are particularly useful.

- $a_{ss'}$ captures the impact of stock s' on stock s and is estimated from the data. There are 30×30 such parameters $a_{ss'}$, $s = 1, \dots, 30, s' = 1, \dots, 30$. We can organize all of these parameters into a 30-by-30 matrix $A = [a_{ss'}]_{s=1,\dots,30; s'=1,\dots,30}$, typically referred to as the excitation matrix. In this work, we only consider excitatory interactions and hence all $a_{ss'} > 0$, and $a_{ss'}$ measures the expected number of events of stock s preceded by an event of stock s'.

- μ_s is the baseline rate of events of stock s that is independent of previous events. Since we focus only on the 30 DJIA stocks, it represents all exogenous effects to the market such as news of all types and effects emanating from all the stocks that are not included in the DJIA.

- $g(t - t_i^{s'})$ is the so-called memory kernel. It models how the effect from each previous event decays over time, because more recent events are generally regarded to have larger influences on the current event; hence, it is typically a decreasing but positive function. Here we adopt the logistic-normal density as in Linderman and Adams (2014). We pick this function because: 1) it is a probability density and integrates to 1, which endows $a_{ss'}$ with the units of the "expected number of events" (see the discussion on $a_{ss'}$ above) and allows a comparison of the strengths of interactions; 2) it has bounded support and naturally models the domain of $g(t - t_i)$, which is bounded below at 0 and above at a certain value[11]; and 3) it allows for an efficient Bayesian estimation framework. Compared to standard Hawkes model, Linderman and Adams (2014) show that their model version (which is the one used here) achieves better results in finding relationships between events.

The above model for $\lambda_s(t)$ applies to each stock s, and models for different stocks are coupled together as the events of stock s explicitly depend on the events of other stocks s'. We can rewrite Equation (1) compactly with matrices. Letting $\boldsymbol{\lambda}(t) = \left(\lambda_1(t), \dots, \lambda_{30}(t)\right)^T$, $\boldsymbol{\mu}(t) = \left(\mu_1(t), \dots, \mu_{30}(t)\right)^T$, $A = [a_{ss'}]$, and $\boldsymbol{g}(t) = \left(\sum_{i:t_i^1<t} g(t - t_i^1), \dots, \sum_{i:t_i^{30}<t} g(t - t_i^{30})\right)^T$, Equation (1) becomes:

$$\boldsymbol{\lambda}(t) = \boldsymbol{\mu}(t) + A \cdot \boldsymbol{g}(t). \tag{2}$$

From the equation above we can see that the excitation matrix A plays an important role in the Hawkes process. This matrix is the crux of our analysis because it captures the interactions between the

[11] Since we only consider impact from past events, $t - t_i > 0$, which means that the domain of $g(t - t_i)$ is bounded below by 0. On the other hand, typically events occurring earlier than a certain threshold are not considered, and hence the domain of $g(t - t_i)$ is bounded above, i.e., $t - t_i < \Delta t_{max}$.

stocks. Its columns depict the stock that triggers the effect and the rows denote the stock that is affected. Consequently, the principal diagonal represents the self-induced impact (self-influence) on each of the 30 stocks, and the other 870 cells in the matrix represent the impact of an individual stock on another individual stock (cross-influence). The number contained in each of the matrix's 900 cells is approximately the average number of events of the corresponding row stock triggered by one event of the corresponding column stock.

Model Estimation

The model is estimated using the Bayesian method by Linderman and Adams (2014) and their Python package "pyhawkes" (https://github.com/slinderman/pyhawkes), which has been proven to be efficient and reliable in other studies (see Linderman and Adams (2014) for details). They show that their model achieves better results in finding relationships between events than the standard Hawkes model does on both synthetic and real data. One advantage of their method is that it imposes a regularization (i.e., penalty) term on the excitation matrix through a prior distribution so that noisy interactions between the stocks are limited in the estimation process. Moreover, it directly models the influence structure between the stocks as a network by using random network models as the prior. In doing so, it separates influence structures from influence strengths. In brief, the method iteratively updates the network structure between the stocks and the influence strengths as follows: Given the network of influence between the stocks, it finds the strengths of the network links that are more likely than the current strengths; and given the strengths of the interactions, it finds the network structure that is more likely than the current structure. In the end, the excitation matrix output by this method is actually the weighted adjacency matrix of the influence network between the stocks.

Excitation Matrix Dynamics

As the behaviour of each stock may vary over time, especially during the flash crash, we do not fit the Hawkes model to all the data combined. Instead, to capture the dynamics of the stocks over time, we divide the data into overlapping time windows. The length of the rolling window is five minutes, and the window moves five seconds at each step. In total, there are 2,160 instances of the moving window between 13:00 and 16:00 plus one startup window immediately prior to 13:00. The Hawkes model is then fitted in every window to the DJIA 30 stocks. For the visualizations that follow, the parameter estimates from each window are plotted against the right boundary of the window. In other words, we sample time points every five seconds within the three-hour period in which the crash occurred and consider the 5-minute history before each time point. From the 5-minute history before each time point, we estimate the Hawkes model and use it to characterize the stocks at that time point. Accordingly, there will be one collection of parameters (e.g., baseline rate μ_s and excitation matrix A) estimated from each window. Since the parameters at each time point are estimated using information from the 5-minute time window prior to it, their calculated effects reflect a moving average and are not instantaneous. [12]

An example of the excitation matrix for five randomly picked stocks during the 13:00:45-13:00:50 window is shown in Figure 3, together with its network representation. We use Bank of America (BAC) and Travelers (TRV) to illustrate the interpretation of this matrix. BAC positively influences itself (0.44) and to a much lesser extent TRV (0.03). In contrast TRV's influence on itself is relatively small (0.10) and it has no effect on BAC (0.00). These relations graphically depicted in the network diagram with arrow heads showing the direction of influence and the thickness of the arrow shaft indicating the relative size of the influence. That TRV has no effect on BAC is shown by the lack of an arrow.

Insert Figure 3 about here

We explain the impact of each of the excitation matrices using the following summary measures. First, we consider the density of the influence network between the stocks, which is the number of links in

[12] A window spanning five minutes is used to ensure that there are adequate observations to estimate Equation (1). In doing so, we implicitly assume that there is no big shift in the dynamics in the window and, thus, choosing a short window is desirable. Permitting the window to move in five second intervals permits a smooth transition from one set of parameter estimations to the next. We also used a 10-minute window, and the empirical results are qualitatively the same but not reported.

the network[13]. Next, we consider the influence strengths. Following the extant convention, for a specific window we define the self-reflexivity of the market during that period as the average of the diagonal entries of the excitation matrix A. In similar manner, we label the average network link strength (i.e., the mean of the nonzero off-diagonal entries of the excitation matrix A) as cross-reflexivity (sometimes referred to as mutual-reflexivity), which can alternately be thought of as the average interaction strength of the market.

We also construct three measures for individual stocks: self-influence, out-influence and in-influence. The self-influence of a stock is the impact of the stock on itself and is the value indicated by the stock's position on the excitation matrix's principal diagonal, i.e., the intersection of the stock's row and column entries. In contrast, the out-influence of a stock is the impact of this stock on the 29 other DJIA stocks or the weighted out-degree of this stock in the influence network. Mathematically, this quantity is the sum of the corresponding column in the matrix A less the value of the diagonal entry. Correspondingly, the in-influence (weighted in-degree) of a stock is the impact of the 29 other stocks on it and is measured by the stock's row sum less the self-influence of the stock being measured. Thus, all the information contained in the excitation matrix is used by these three influence measures.

The excitation matrix is similar to the variance-covariance matrix that is often used to measure the risk associated with stock portfolios as it measures how the stocks are related to each other. However, the excitation matrix differs in three important ways. First, it does not require two time series to be synchronized, which high-frequency trading data are typically not. Second, the excitation matrix is asymmetric (or directed) so that a stock can have a specific impact on another stock, but the reverse need not be the case as it is entirely possible that the impacts may be asymmetric. Finally, it does not suffer from the Epps (1979) effect, which typically renders the variance-covariance matrix unreliable for high-frequency data.[14]

[13] Network density is typically defined as the number of links divided by the number of possible links. However, since our network size as well as the number of possible links are fixed, our network density is equivalent to the number of links.

[14] Epps (1979) points out that the covariance between two time series approaches zero as the observation frequency increases. Grossmass (2014), among others, suggests that this effect is a result of non-synchronous and asynchronous trading as well as microstructure noise. Nevertheless, the correlation

Economic Periods and Sample Sizes

All the figures and tables that follow are split into five economic periods: pre-crash, crash, nadir, recovery, and post-recovery. These periods are determined ex post and are defined in Table 2. The time period spanned by the crash, nadir and recovery periods corresponds to the SEC's (2010a, 2010b) crash period. Our nadir period is structured so that it contains the lowest price observation for each of the 30 stocks as depicted in Figure 2. Table 2 also contains information on the number of transactions in each of the five periods and these periods in aggregate. The table also provides the frequency of transactions, which is the average number of milliseconds between transactions, and the proportion of transactions associated with price increases (the complement to this proportion is associated with price decreases.)

Insert Table 2 about here.

As displayed in Table 2, the total number of transactions across the five economic periods combined is slightly over one million, which translates to, one transaction occurred every 10.1 milliseconds on average. The frequency of transactions monotonically increased from one every 22.1 milliseconds in the pre-crash period to one every 2.8 milliseconds in the nadir period and then monotonically decreased to one every 7.8 milliseconds in the post-recovery period.[15] The U-shape pattern is somewhat echoed by the proportion of trades associated with positive price increases. This proportion drops from 49.1 percent in the pre-crash period to 48.3 percent in crash period and rises to 50.3 percent in the nadir period and subsequently flattens for the remaining two periods.

timeframe depends on the trading (planning) horizon of the trader. Thus, flash crashes may not be of great interest to, say, the buy-and-hold investor whose planning horizon may be months or even years unless the flash crash is some sort of understandable and credible signal that indicates major market negative disruptions sometime in the future.

[15] To put these time intervals consider the following comparisons to the latency of humans and computers. It is generally accepted that the typical human reaction time (or latency) to visual signals is between 200 and 250 milliseconds. Saariluoma (1995) indicates that a chess grandmaster's latency to recognize that his king is endangered by an opponent's move is 650 milliseconds. Computer latency is the time it takes to deliver a message from one computer to the other, and is a function of distance and power of the computer system used. Kay (2009) suggests that at the time of his study that the computer technology at that time was around 70 percent the speed of light or about 130 miles per millisecond. Thus, minimum one-way latencies for information from stock trading centers in New York to San Francisco and Chicago are approximately 19.1 and 5.5 milliseconds, respectively.

Recall that our Hawkes model specification requires that a price change signals the existence of an event despite the price change being negative or positive. Because our event measure is a binary variable and transaction profits can be made on price changes regardless of direction, we would expect that each type of price change would account for 50 percent of the total. Z-tests for the proportion of positive transactions for each of the five economic periods and the five periods combined are provided in Table 2. In this case all the Z-test values for the five periods are highly significant with the exception of the nadir period, which is barely insignificant at the traditional 0.05 critical p-value. Taken together the five periods are significantly negative as are the three periods (crash, nadir and recovery), which coincide with the official crash period. Although, the beginning of the flash crash is reasonably straightforward, its end is not. Nevertheless, adding the post-recovery period to the official crash time interval, the percent of positive price changes remains statistically significant with a p-value of 0.0024.

Thus, it seems that in all of the above cases we should reject the null hypothesis that the incidence of the two price changes are each 50 percent. This conclusion, however, is an example of the "Large Sample Fallacy (LSF)." In the case of the Z-test and similar statistical tests, the LSF is the result of dividing the statistic's denominator by the square root of the sample size so that as the sample becomes larger the Z-test value becomes larger and the p-value of the test becomes smaller.

As pointed out by Lin, Lucas and Shmueli (2013), among others, several approaches have been suggested to mitigate this problem, but the two most popular appear to be: (1) decreasing the acceptable p-value, and (2) focusing on whether the actual finding is meaningful in the context of the phenomenon under investigation. Of course, the two approaches are not mutually exclusive but de facto they often are. The first approach minimizes Type 1 error, which makes the null hypothesis more difficult to reject, and should require the p-value signaling statistical significance to be specified ahead of time.[16] The second involves

[16] This is clearly an ad hoc approach, but then so is the use of the standard use of 0.10. 0.05, and 0.01 as often used p-values for small samples, a practice that dates back to the 1930s. Often many p-values were not recorded numerically but instead were simply signified by asterisks, i.e., *, **, and ***, for the above standard p-values. Many in the past have used this technique. Although the above three p-values are ad hoc, they do have the advantage of not conveying unwarranted precision.

determining whether the difference between the null hypotheses and what is observed is substantive or, in our case, economically meaningful. As argued and demonstrated by Ziliak and McCloskey (2008), failing to address the latter may result in dire consequences. We adopt the second approach and advance two arguments supporting the position positive price and negative price changes should be considered the same since the purpose of the price change variables is only to count how many trades occurred.

First, a popular view of a crash is that some major negative information is noted by the market participants, and as stock prices begin to fall, some sort of market contagion takes effect and prices drop in unison. The reverse holds true during a recovery. Our results do not support this view, which is most likely caused by media hype and the fact that the public does not have access to nor is aware of high frequency data. For example, as displayed in Table 2, 48.26 percent of the price-changing events are positive during the crash period, and 50.41 percent are positive during the recovery period. For the official crash period, which includes the previous two periods plus the nadir period, the positive events account for 49.62 percent of the total. Numerically, these are very close to the 50 percent neutral value and indicate that market participants engage in price discovery, so that they tend to find and exploit profit opportunities regardless of the direction of the market's general movement to learn true value of the market after the shock. Frank et al. (2019) suggest that traders who engage in this process together are an example of Adam Smith's often mentioned "invisible hand".

Second, market regulators throughout the world are concerned about the effectiveness of the price discovery process. A majority of them, included the United States, have adopted some sort of circuit breaker, which temporarily stops trading for a short period of time, despite the fact that theoretical and empirical research is mixed concerning its usefulness (see, e.g. Ackert (2012) and Sifat and Mohamad (2018).) Circuit breakers can be market-wide or focused on individual securities. Typically they are concerned with falling prices. Thus, the main argument favoring this approach is that it provides a cooling off period, which gives market traders an opportunity to better evaluate the reason for the price drop in an effort to make better trading decisions or to adjust parameters in their algorithms. The contrary view is that the delay only postpones trading and may exacerbate the price decline when traders try to change their trading strategies in an attempt to game the system as the circuit breaker trigger price approaches.

On June 19, 2010, approximately one month after the May 6 flash crash, U.S. regulators established a single stock circuit breaker to guard against falling prices. On April 5, 2011, the Financial Industry Regulatory Authority, along with several security exchanges, suggested replacing the single stock circuit breaker with a "limit up-limit down" (LULD) circuit breaker. The main reason for the replacement was that this type of mechanism would not only handle downside volatility but also upside volatility. On May 31, 2012, the Securities Exchange Commission approved the proposal for a trial run, and on April 11, 2019, it gave the LULD permanent status.

Conceptually, the LULD is a straightforward reference price constructed by calculating a simple average of the transaction prices during the previous five minutes. Market opening is an exception since there are no data to average; in this case the opening price is used as the reference price. Then the upper and lower price bands are determined by multiplying the reference price by one plus or minus the preset percentage parameter, respectively. This calculation is done every 30 seconds and is updated if the new reference price is at least one percentage point away from the current posted reference price. (See www.luldplan.com for details.) For our purposes, the main point is that the Securities Exchange Commission band measures price volatility, and that large price changes of equal size in either direction are equally unwanted as both increase price volatility.

Empirical Findings

We first present the empirical results for the 30 DJIA as a group then focus on the 30 stock's impact upon themselves and upon each other. For convenience we refer to the 30 stocks taken together as the "market" and by themselves as simply "stock" or by company name or stock symbol. All results are based on the excitation matrices calculated using the five minute rolling window.

DJIA Market

We first examine the density of the influence network (i.e., number of edges in the network) between the stocks, which is shown in Figure 4-top. Recalling that the Hawkes process only uses information about the stocks' trading events, it is notable that the two time series – average market price and network density – almost collapse on each other.[17] Specifically, the two series drop down simultaneously when the crash starts, reach their bottoms at the same time, and recover concurrently. While the network density follows closely with the market price, the influence strength shows a different pattern. The cross-reflexivity of the market (i.e., average of the cross-influences between the 30 stocks) is shown in Figure 4-bottom. The abrupt increase of cross-reflexivity around 14:32 reflects the beginning of the sudden decline in stock prices. The two plots together suggest that, when the crash starts, network density decreases, but the interaction strength for the remaining links increases dramatically, indicating that although market activity increases it is concentrated between fewer stocks. The average cross-reflexivity reaches its highest point several minutes after the average price reaches its lowest value, possibly because traders are unable to determine exactly when the market reached its lowest point and the activity level of the market is still high. After reaching their extreme points, price and cross-reflexivity both tend to return to their approximate pre-crash levels, although price is not quite as high nor is cross-reflexivity quite as low.

Insert Figure 4 about here.

Recalling that our model includes three types of effects - the exogenous effect (the average baseline rate (μ_s)), the self-reflexivity, and the cross-reflexivity - we further examine the relative strengths of these three effects over time. Figure 5 shows the proportion of each type of effect to their sum. The exogenous effect, which includes stocks not included in the DJIA and exogenous information to the market, is small (blue curve), taking a proportion of less than one percent most of the time. The self-reflexivity (orange curve) takes a larger proportion than the cross-reflexivity (green curve), but the former starts to decay and the latter to grow when the crash starts (around 14:32), and the two become closer during the crash. They roughly return to pre-crash levels after the crash. The overall pattern of the cross-reflexivity is similar to

[17] The average price in Figure 4 is the simple mean the 30 stocks in the DJIA. It is not the official DJIA, which is also an average. See fn. 7 for details on the difference between these two averages.

that of the exogenous effect, which suggests that the latter effect may be dominated by stocks that are not part of the DJIA and are not explicitly modeled in our analysis.

Insert Figure 5 about here.

DJIA Stocks

To determine the major influencers in the market, for each company we calculate the average out-influence or its weighted out-degree, which is the impact of a particular stock on the other 29 stocks, in the pre-crash through post-recovery periods (see Table 2). As we show in Table 3, the pattern of out-influence varies among the stocks and it varies among the five economic periods. In addition to the average out-influence for each stock in each economic period, we provide the stock's average ranking of out-influence. As indicated in Table 3, the three strongest out-influence stocks over time are Bank of America (BAC), ExxonMobil (XOM), and JPMorgan Chase (JPM), and the three weakest are 3M (MMM), Dupont (DD) and Travelers (TRV).

Insert Table 3 about here

As a complement to Table 3, Table 4 shows the average in-influences, i.e., the impact on a stock by the other 29 stocks, of the stocks before, during, and after the crash. Similar to average out-influences, the pattern of average in-influences varies among stocks, but the variation is less even. The three weakest in-influence stocks over time are ExxonMobil (XOM), Microsoft (MSFT) and JPMorgan Chase (JPM), two of which are among the top three in out-influence. The three strongest in-influence stocks are 3M (MMM), Alcoa (AA), and American Express (AXP); and 3M (MMM) and Alcoa (AA) rank among the bottom four with respect to out-influence.

Insert Table 4 about here.

In a format similar to Tables 3 and 4, Table 5 reports the self-influence data for each of the 30 DJIA stocks. As shown in this table the three strongest stocks with respect to self-influence are Bank of America (BAC), Microsoft (MFST) and ExxonMobil (XOM) and the three weakest stocks are Chevron (CVX), United Technologies (UTX) and DuPont (DD). Two observations merit special mention. First, ExxonMobil is one of the strongest stocks with respect to out-influence (rank: 01) and self-influence (rank: 03) but is one of the weakest stocks with respect to in-influence (rank: 30). Second, although Chevron is not in either the strong or weak category with respect to out-influence (rank: 10) or in-influence (rank: 10) but it is clearly in the weakest self-influence category (rank: 28). Taken together these observations suggest that ExxonMobil exhibits more market power, but this market power may not be related to industry sector since both ExxonMobil and Chevron are in the energy sector and they are the only two DJIA stocks that are in this category (see Table 1).[18]

Insert Table 5 about here.

Out-influences and in-influences are much larger than self-influences, although the averages of all three influences changed size as the market moved through the pre-crash, crash, nadir, post-recovery and post-recovery periods. As shown in Table 6, self-influence and out-influence are positively correlated in all periods. In contrast, in-influence is always negatively correlated with self-influence and out-influence measures. In absolute terms the correlations of the three pairs of influences in the crash period are smaller than those experienced in the pre-crash period. Beginning in the recovery period, these correlations tend to move toward their pre-crash levels.

Insert Table 6 about here.

To further investigate the behavior of the three different types of influence, we calculate the means of the out-influence, in-influence, and self-influence for the 30 Dow Jones stocks when the environment changes from pre-crash to crash, crash to nadir, and so forth. We conjecture that the sequential means may be dependent in some way to the previous mean, e.g., the mean in the nadir period is dependent on the mean in the crash period. Thus, we use a paired t-test to examine the way in which the means evolve over time. The results of these tests are given in Table 7. We cannot reject the null hypothesis that the mean out-influence and the mean in-influence do not change between the sample periods. This is not the case for

[18] Numerous studies have suggested that there may be an industry effect. See, e.g., King (1966), Cavaglia, Brightman and Aked (2000), and Fan, Furger and Xiu (2016).

self-influence. All of the paired t-tests are statistically significant. The mean self-influence decreases during the crash and increases as the market recovers.

Insert Table 7 about here.

As previously mentioned, the stock market can be thought of as a dynamic network with its nodes being the individual stocks and the movements being the price changes of these stocks. A useful descriptor of large-scale structures of a network is modularity.[19] Modularity quantifies the degree to which a network can be divided into communities, or in our case clusters of stocks that are related to one another[20]. Networks with high modularity have dense connections between nodes within same communities but sparse connections between nodes contained in different communities. Mathematically, modularity is a function of the excitation matrix *A*, and a partition *C* of nodes (where $C_i = k$ indicates that node *i* belongs to community or cluster *k):*

$$modularity\ = \frac{1}{M} \sum_{ij} \left(A_{ij} - \frac{d_i^{out} d_j^{in}}{M} \right) \delta\left(C_i, C_j\right), \quad (3)$$

where d_i^{out} is the weighted out-degree (out-influence) of node *i*, d_j^{in} is the weighted in-degree (in-influence) of node *j*, *M* is a normalizing constant, and $\delta(C_i, C_j)$ is a delta function such that $\delta\left(C_i, C_j\right) = 1$ if $C_i = C_j$ and $\delta\left(C_i, C_j\right) = 0$ otherwise. Thus, modularity is different for different partitions on a network, but typically, in applications, modularity refers to the maximum modularity according to the best partition. In our case, we also refer to the optimal modularity in the discussions below and we use the Python package leidenalg (https://github.com/vtraag/leidenalg) by Traag, Walton, and van Eck (2019) to find the best partitions and the corresponding modularity scores. For completely random networks modularity will be close to zero, and the larger the modularity, the more fragmented a network will be.

[19] For a description of modularity and its applications to a variety of networks, including club membership, scholarly citations, and fictional characters in *Les Misérables* by Victor Hugo, see Newman and Girvan (2004) and Leicht and Newman (2008).

[20] Community is a term used in Network Science to describe a cluster of nodes. We use it interchangeably with cluster in this paper. Further, a division of a network into communities or clusters is called a partition.

Similar to some of our earlier analyses, the modularity of our stock network is plotted in Figure 6 along with the stock price series from the pre-crash period to the post-recovery period. As shown in Figure 6, the modularity increases during the crash, then peaks during the nadir period (i.e., the network is the most fragmented) and decreases in the recovery period. Compared to these three middle periods, the pre-crash and post-recovery periods are both more homogeneous and less fragmented.

Insert Figure 6 about here.

Previously, we indicated that several earlier studies (see fn. 14) reported an industry effect suggesting that the prices of stocks in the same industry may tend to move together. We suggested that this effect may not be the case for energy stocks. We explore this issue in more detail by comparing the network communities to industry sectors using the normalized mutual information (NMI) statistic. The NMI quantifies the similarity that exists between two partitions on the same set of objects (i.e., stocks): If the two partitions are the same the value of NMI is one, and if they are independent the value is zero (see Vinh, Epps, and Bailey (2010) for a formal definition). Accordingly, for each excitation matrix we cluster the stocks using the modularity method (see discussion on modularity above and Traag, Waltman, and van Eck (2019)), and compute the NMI statistic between the identified communities and the industry sectors. The calculated NMI values for the communities and the industries (see Table 1) are plotted in Figure 7.

Insert Figure 7 about here.

A review of Figure 7 indicates that the NMI does not have any significant trend and its value remains relatively small, i.e., about 0.3, throughout the time period under examination.[21] This finding in conjunction with the modularity results depicted in Figure 6 suggests that the network becomes more clustered during the crash, but these clusters do not materially overlap with industry sectors. In other words, the impact of the crash is more likely to spread among sectors than within sectors.

[21] Newman and Girvan (2004) report that in their studies typical value range from 0.30 to 0.70, with values above 0.70 being quite rare.

Discussion and Concluding Remarks

Our empirical findings indicate that the DJIA 30 stocks exhibit the characteristics of self- and cross-reflexivity as well as out-influence and in-influence, suggesting that past price movements in the prices of these stocks not only influence the future prices of the stocks themselves, but also of other stocks that make up the index. The out- and in-influence interactions between the stocks vary before, during and after the 2010 flash crash. Nevertheless, with respect to rank correlation, the behavior of stocks is strongly negative for in-influence vs. out-influence or self-influence. Taken as a whole, the self-influence of the 30 Dow Jones stocks declines from the pre-crash period to the bottom (nadir) and then increases in the recovery period.

With respect to trading profits, algorithms, possibly based on artificial intelligence techniques, should be able to be (or possibly have been) constructed to exploit the patterns that we uncover. Nevertheless, as an increasing number of traders do this, trading profitability from this source will undoubtedly diminish. Thus, portfolio risk management may be a more useful application.

Modern risk management of a stock portfolio has focused on various measures of expected return and volatility. The initial basis for this approach was initiated by Markowitz (1952) using the statistical concepts of mean and variance. Relying on this development, Sharpe (1964, 1994) developed a risk-return performance ratio (the Sharpe Ratio), which implicitly assumes that the relevant return distributions are Gaussian.[22] Unfortunately, this assumption is often not true. Recognizing that return distributions are typically asymmetric, Sortino (2001) designed a performance measure that focuses on downside risk (the Sortino Ratio) that is concerned only with the left tail of the return distribution.[23] Applying either the

[22] The Sharpe ratio is the expected return in excess of the risk free rate divided by the standard deviation of return. Its theoretical basis is the Mean-Variance Capital Asset Pricing Model (CAPM), which was developed and refined by Markowitz (1952, 1959), Sharpe (1964), Lintner (1965), and Mossin (1966).

[23] The work of Roy (1952), Markowitz (1959), Hogan and Warren (1974), Bawa and Lindenberg (1977), Harlow and Rao (1989), and others resulted in the Lower Partial-Moment Capital Asset Pricing Model (LPMCAPM). Satchel (2001) points out that this asset pricing model is the foundation of the Sortino Ratio. The numerator of this ratio is the expected return less an investor determined target rate and its denominator is the square root of the second lower partial moment, which is also defined by the target rate.

Sharpe Ratio or the Sortino Ratio to high frequency data, however, is problematic because of the Epps effect (see fn. 14).

Fortunately, Aït-Sahalia and Hurd (2016) develop a capital asset pricing model where the stocks being considered are described by mutually exciting Hawkes processes. They show in a dynamic context that the optimal portfolio composition changes in responses to changes in the jump intensities of individual stocks since these jumps predict future jumps. An interesting and important result of their model when applied to risky assets and a risk-free asset is that because of the excitation relationship among stocks, a jump in one causes the investor to sell all the stocks and invest the proceeds in the risk-free asset. This behavior suggests that in this type of environment that recognizes the possibility of contagion, a very few of or even single stock could trigger a crash, flash or otherwise.

In sum, because of the high level of technology that is being used in the major stock and similar markets worldwide, and the likely continuing of high-frequency trading, additional work should be done exploring the implications of the Aït-Sahalia and Hurd (2016) capital asset pricing model using real transactions data with a focus on measuring portfolio performance and trading strategies as well as the implications of market fragmentation as suggested by the Staff of the Division of Trading and Markets of the U.S. Securities and Exchange Commission (2013) and the references contained therein.

References

Ackert, Lucy F., 2012, *The Impact of Circuit Breakers on Market Outcomes*, Economic Impact Assessment EIA9, Foresight Project – The Future of Computing in Financial Markets, U.K. Government Office for Science, August 31.

Aït-Sahalia, Yacine, Julio Cacho-Diaz and Roger J.A. Laeven, 2015, Modeling financial contagion using mutually exciting jump processes, *Journal of Financial Economics* 117:3, 585 – 606.

Aït-Sahalia, Yacine, and Thomas Robert Hurd, 2016, Portfolio choice in markets with contagion, *Journal of Financial Econometrics* 14:1, 1 – 28.

Akgiray, Vedat, and G. Geoffrey Booth, 1988, Mixed diffusion-jump process modeling of exchange rate movements, *Review of Economics and Statistics* 70:4, 631 – 637.

Aldrich, Eric M., Joseph A. Grundfest and Gregory Laughlin, 2017, The flash crash: A new deconstruction. *https://dx.doi.org/10.2139/ssrn.2721922.*

Andersen, Torbin G., Luca Benzoni and Jesper Lund, 2002, An empirical investigation of continuous-time equity return models, *Journal of Finance* 57:3, 1239 – 1284.

Bachelier, Louis J.-B. A., 1900, *Theory of Speculative Prices*, Dissertation, University of Paris. English version: Mark Davis and Alison Etheridge (trans.), 2007, *Louis Bachelier's Theory of Speculation: The Origins of Modern Finance* (Princeton University Press, Princeton, N.J.).

Bacry, Emmanuel, Iacopo Mastromatteo and Jean-Francis Muzy, 2015, Hawkes processes in finance, *Market Microstructure and Liquidity*, 1:1, 1550005.

Bawa, Vijay S., and Eric B. Lindenberg, 1977, Capital market equilibrium in a mean-lower partial moment framework, *Journal of Financial Economics* 5:2, 189 – 200.

Booth, Elizabeth B., G. Geoffrey Booth and John Broussard, 2014, Communication technology and exchanging financial assets: A historical perspective, *Business and Economic Research* 4:2, 308 – 322.

Cai, Ning, and S.G. Kou, 2011, Option pricing under a mixed-exponential jump diffusion model, *Management Science* 57:11, 2067 – 2081.

Cavaglia, Sefano, Christopher Brightman, and Michael Aked, 2000, The increasing importance of industry factors, *Financial Analysts Journal* Sept/Oct, 41 – 54.

Commodity Futures Trading Commission-Securities Exchange Commission (CFTC-SEC), 2010a, *Preliminary Findings Regarding the Market Events of May 6*, Staff report. Available at https://www.sec.gov/sec-cftc-prelimreport.pdf.

Commodity Futures Trading Commission-Securities Exchange Commission (CFTC-SEC), 2010b, *Findings Regarding the Market Events of May 6*, Staff report. Available at https://www.sec.gov/news/studies/2010/marketevents-report.pdf.

Embrechts, Paul, Thomas Liniger and Lu Lin, 2011, Multivariate Hawkes processes: An application to financial data, *Journal of Applied Probability* 48:A, 367 – 378.

Epps, Thomas W., 1979, Comovements of stock prices in the very short run, *Journal of the American Statistical Association* 74:366, 291 – 298.

Fama, Eugene F., 1965, Behavior of stock market prices, *Journal of Business* 38:1, 34 – 105.

Fan, Jianqing, Alex Furger and Dacheng Xiu, 2016, Incorporating global industrial classification into portfolio allocation: A simple factor-based large covariance matrix estimator with high frequency data, *Journal of Business & Economic Statistics* 34:4, 489 – 503.

Filimonov, Valdimir, and Didier Sornette, 2012, Quantifying reflexivity in financial markets: Toward a prediction of flash crashes, *Physical Review E* 85:5, 056108-2 – 056108-9.

Frank, Kenneth A., Yun-Jia Lo, G. Geoffrey Booth, and Juha-Pekka Kallunki, 2019, The market for socially embedded trading, *Rationality and Society* 31:2, 152 – 181.

Gao, Cheng, and Bruce Mizrach, 2016, Market quality breakdowns in equities, *Journal of Financial Markets* 28, 1 – 23.

Glasserman, Paul, and H. Peyton Young, 2016, Contagion in financial networks, *Journal of Economic Literature* 54:3, 779 – 831.

Grossmass, Lindan, 2014, Obtaining and predicting bounds of realized correlations, *Swiss Journal of Economics and Statistics* 150:3, 191 – 226.

Hawkes, Alan G., 2018, Hawkes processes and their applications to finance: A review, *Quantitative Finance* 18:2, 193 – 196.

Harlow, W.V., and Ramesh K.S. Rao, 1989, Asset pricing in a generalized lower partial moment framework, *Journal of Financial and Quantitative Analysis* 24:3, 285 – 311.

Hogan, William W., and James M. Warren, 1974, Toward the development of an equilibrium capital market based on semivariance, *Journal of Financial and Quantitative Analysis* 9:1, 1 – 11.

Johnson, Neil, Guannan Zhao, Eric Hunsader, Hong Qi, Nicholas Johnson, Jing Meng, and Brian Tivnan, 2013, Abrupt rise of new machine ecology beyond human response time, *Scientific Reports*, Article 2627, 1 – 7.

Kay, Rony, 2009, *Pragmatic Network Latency Engineering*, working paper (cPacket Networks, Inc., Mountain View, Calif.)

King, Benjamin F., 1966, Markets and industry behaviors in stock price behavior, *Journal of Business* 39:1, 139 – 190.

Kirilenko, Andrei, Albert S. Kyle, Mehrdad Samadi and Tugkan Tuzun, 2017, High-frequency trading in an electronic market, *The Journal of Finance* 72:3, 967 – 998.

Lin, Mingfeng, Henry C. Lucas, Jr. and Galit Shmueli, 2013, Too big to fail: Large samples and the p-value problem, *Information Systems Research* 24:4, 906 – 917.

Leicht, E.A., and M.E.J. Newman, 2008, Community structure in directed networks, *Physical Review Letters* 100:11, 118703-1 – 118703-4.

Lewis, Michael, 2014, *Flash Boys: A Wall Street Revolt* (W.W. Norton & Company: New York, N.Y.).

Linderman, Scott, and Ryan Adams, 2014, Discovering Latent Network Structure in Point Process Data, *31st International Conference on Machine Learning* (JMLR: W&CP: Beijing, China), vol. 32, pp. 1413–21.

Liniger, Thomas, 2009, Multivariate Hawkes Processes, Doctoral Thesis, Department of Mathematics, ETH Zürich, Switzerland.

Lintner, John, 1965, The value of risk assets and the selection of risky investments in stock portfolios and capital budgeting, *Review of Economics and Statistics* 47:1, 13 – 37.

Madhavan, Ananth, 2012, Exchange-traded funds, market structure, and the flash crash, *Financial Analysts Journal* 68:4, 20 – 35.

Malkiel, Burton G., 1973, *A Random Walk Down Wall Street* (W.W. Norton & Company: New York, N.Y.).

Markowitz, Harry M., 1952, Portfolio selection, *The Journal of Finance* 7:1, 77 – 91.

Markowitz, Harry M., 1959, *Portfolio Selection* (John Wiley & Sons, New York, N.Y).

Menkveld, Albert J., and Bart Zhou Yueshen,2019, The Flash Crash: A cautionary tale about highly fragmented markets, *Management Science* 65:10, 4470 – 4488.

Merton, Robert C., 1976, Option pricing when underlying stock returns are discontinuous, *Journal of Financial Economics* 3:1-2, 125 – 144.

Mossin, Jan, 1966, Equilibrium in a capital asset market, *Econometrica* 34:4, 768 – 783.

Newman, M.E.J., and M. Girvan, 2004, Finding and evaluating community structure in networks, *Physical Review E* 69:2, 026113-1 – 026113-15.

Rizoiu, Marian-Andrei, Young Lee, Swapnil Mishra and Lexing Xie, 2017, Hawkes processes for events in social media, *Frontiers in Multimedia Research*, Shih-Fu Chang, ed. (Association for Computing Machinery and Morgan & Claypool: New York, N.Y.), pp. 191 – 218.

Roy, A.D., 1952, Safety first and the holding of assets, *Econometrica* 20:3, 431 – 449.

Saariluoma, Pertti, 1995, *Chess Player's Thinking: A Cognitive Psychological Approach* (Routledge, New York, NY).

Satchell, Stephan E., 2001, Lower-partial moment capital asset pricing models: A re-examination, *Managing Downside Risk in Financial Markets: Theory, Practice and Implementation*, Frank S. Sortino and Stephan E. Satchell, eds. (Butterworth-Hienemann Finance, Oxford, U.K.) pp. 156 – 168.

Sharpe, William F., 1964, Capital asset prices: A theory of market equilibrium under conditions of risk, *Journal of Finance* 19:3, 425 – 442.

Sharpe, William F., 1994, The Sharpe ratio, *Journal of Portfolio Management* 21:1, 49 – 58.

Sifat, Imtiaz Mohammad, and Azhar Mohamad, 2018, Circuit Breakers as market stability levers: A survey of research, praxis, and challenges, *International Journal of Finance & Economics*, 1 – 40.

Sortino, Frank A., 2001, From alpha to omega, *Managing Downside Risk in Financial Markets: Theory Practice and Implementation*, Frank A. Sortino and Stephan E. Satchell, eds. (Butterworth-Hienemann Finance, Oxford, U.K.), pp. 3 – 25.

Staff of the Division of Trading and Markets of the U.S. Securities and Exchange Commission, 2013, *Equity Market Structure Literature Review* (U.S. Securities and Exchange Commission. Washington, DC.

Traag, V.A., L. Waltman, N.J. van Eck, 2019, From Louvain to Leiden: Guaranteeing well-connected communities, *Scientific Reports* 9, Article number: 5233 (12 pp.).

Vinh, Nguyen Xuan, Julien Epps, and James Bailey, 2010, Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance, *Journal of Machine Learning Research* 11, 2837-2854.

Ziliak, Stephen T., and Deidre N. McCloskey, 2008, *The Cult of Statistical Significance: How the Standard Error Cost Us Jobs, Justice, and Lives* (The University of Michigan Press, Ann Arbor, Michigan).

Table 1. The 30 Dow Jones Industrial Average (DJIA) Stocks on May 6, 2010 with Selected Industrial Sector and Price Information.

Company	Sector	Symbol	Open	Close	Low	Low time
3M	I	MMM	86.06	84.24	67.98	14:46:05.7
Alcoa	IM	AA	12.34	11.92	11.25	14:47:35.1
American Express	F	AXP	44.41	42.39	40.16	14:45:52.4
AT&T	T	T	25.68	25.11	24.04	14:46:03.7
Bank of America	F	BAC	17.48	16.26	15.50	14:46:36.2
Boeing	I	BA	70.66	67.90	62.00	14:45:41.8
Caterpillar	I	CAT	65.85	63.49	58.00	14:45:33.1
Chevron	E	CVX	79.42	77.41	71.50	14:47:03.4
Cisco Systems	IT	CSCO	26.41	25.48	23.23	14:45:32.6
Coca-Cola	CS	KO	53.67	52.23	51.21	14:47:23.0
DuPont	IM	DD	37.70	36.69	33.66	14:46:29.3
ExxonMobil	E	XOM	65.79	63.72	58.46	14:46:52.0
General Electric	I	GE	18.00	17.33	15.00	14:46:11.0
Hewlett-Packard	IT	HPQ	50.53	48.32	41.94	14:46:13.3
Home Depot	CD	HD	34.92	33.92	32.07	14:45:56.7
IBM	IT	IBM	126.29	123.86	116.00	14:46:32.9
Intel	IT	INTC	22.15	21.51	19.90	14:47:30.1
Johnson & Johnson	HC	JNJ	65.04	63.39	60.03	14:46:09.7
JP Morgan Chase	F	JPM	42.63	40.78	39.29	14:45:45.5
Kraft Foods	CS	KFT	29.63	29.20	27.49	14:47:58.8
McDonalds	CD	MCD	70.45	69.30	67.49	14:47:52.7
Merck & Company	HC	MRK	35.43	34.21	30.70	14:46:10.7
Microsoft	IT	MSFT	29.60	28.97	27.91	14:46:39.0
Pfizer	HC	PFE	17.16	16.72	15.85	14:46:06.2
Procter & Gamble	CS	PG	61.91	60.71	39.37	14:47:15.3
Travelers	F	TRV	50.58	49.76	48.53	14:45:46.0
United Technologies	I	UTX	73.04	71.14	65.17	14:46:38.0
Verizon	T	VZ	28.61	28.00	26.49	14:45:47.9
Walmart	CS	WMT	54.35	53.21	51.53	14:45:29.2
Walt Disney	CD	DIS	35.15	33.94	31.00	14:45:44.8

Note: For each company, columns two through seven provide (1) the stock's sector code, (2) the stock's ticker symbol, (3) the price at market opening, (4) the price at closing, (5) the lowest price of the day, and (6) the time (hour-minute-second) that the lowest price was recorded. International Classification Benchmark (ICB) sectors (codes) are Consumer Staples (CS), Industrial Materials (IM), Industrials (I), Financials (F), Telecommunications (T), Energy (E), Consumer Discretionary (CD), Information Technology (IT), and Health Care (HC).

Table 2. Statistical and Economic Crash and Recovery Periods and Selected Trade information on May 6, 2010

	Pre-Crash	Crash	Nadir	Recovery	Post-Recovery	Total
Clock Time						
Start	13:00:00.0	14:32:00.0	14:45.29.2	14:47:58.9	15:08:00.1	13:00.00.0
Stop	14:31:59.9	14:45:29.1	14:47:58.8	15:08:00.0	16:00.00.0	16:00:00.0
Trades						
Number	250,101	150,819	53,884	210,933	400,604	1,066,131
Frequency	22.063	5.365	2.787	5.694	7.788	10.130
Price Incr.						
% of Total	49.07	48.26	50.35	50.41	50.24	49.73
Z-test	-9.28	-13.5	1.59	3.72	3.00	15.7
p-value	$<10^{-20}$	$< 10^{-41}$	.0559	.0001	.0013	$<10^{-8}$

Note: Time periods are expressed in 24-hour clock time and are recorded in hours, minutes and seconds, which are separated by colons. The SEC (2010a, 2010b) considers the crash period to begin at 14:32:00.0 and end at 15:08:00.0. Frequency is expressed as average number of milliseconds between trades. Price Increase is the percent of total trades that are characterized by a price increase. Its complement is the percent of price decreases. The Z-test entry tests the null hypothesis that the portion of the total number of trades are associated with price increases is 0.50. The price increase percent, Z-test and p-value for the duration of the official crash, i.e., crash, nadir and recovery, are 49.62, -9.82 and 0.0004, respectively. Adding the post-recovery period to the official crash period results in a price increase percent of 49.92, a Z-test of 2.82 with a p-value of .0024.

Table 3. Average Out-Influences and Corresponding Ranks for Each of the 30 Dow Jones Industrial Average (DJIA) Stocks on May 6, 2010

Company	Pre-Crash	Crash	Nadir	Recovery	Post-Recovery	R
3M	1.18 (26)	1.11 (27)	1.00 (27)	1.05 (27)	1.23 (26)	***28***
Alcoa	1.21 (25)	1.13 (26)	1.04 (26)	1.20 (23)	1.08 (29)	26
American Express	1.42 (13)	1.54 (11)	1.44 (17)	1.39 (16)	1.55 (11)	13
AT&T	1.38 (14)	1.38 (15)	1.56 (10)	1.61 (11)	1.32 (23)	14
Bank of America	3.26 (01)	2.63 (03)	2.45 (03)	2.87 (02)	2.59 (02)	**02**
Boeing	1.33 (17)	1.37 (16)	0.91 (29)	1.19 (24)	1.44 (17)	22
Caterpillar	1.57 (08)	1.81 (05)	1.37 (20)	1.38 (18)	1.48 (15)	12
Chevron	1.64 (07)	1.75 (06)	1.40 (18)	1.31 (20)	1.57 (09)	10
Cisco Systems	1.52 (10)	1.61 (07)	2.22 (04)	1.76 (06)	1.49 (14)	07
Coca-Cola	1.26 (21)	1.39 (14)	1.29 (22)	1.30 (21)	1.63 (06)	19
DuPont	1.14 (29)	1.11 (27)	1.07 (25)	1.00 (28)	1.09 (28)	***29***
ExxonMobil	2.09 (02)	2.67 (01)	2.96 (01)	3.08 (01)	2.70 (01)	**01**
General Electric	1.89 (04)	2.65 (02)	2.04 (06)	2.32(04)	2.45 (03)	04
Hewlett-Packard	1.82 (05)	1.56 (09)	1.59 (09)	1.63 (10)	1.59 (07)	06
Home Depot	1.23 (23)	1.23 (24)	1.45 (15)	1.22 (22)	1.44 (17)	21
IBM	1.47 (12)	1.48 (12)	1.25 (23)	1.35 (19)	1.59 (07)	14
Intel	1.71 (06)	1.58 (08)	2.16 (05)	1.67 (09)	1.56 (10)	05
Johnson & Johnson	1.29 (20)	1.31 (21)	1.61 (08)	1.48 (14)	1.53 (12)	16
JPMorgan Chase	2.04 (03)	2.18 (04)	2.53 (02)	2.36 (03)	2.25 (04)	**03**
Kraft Foods	1.15 (28)	1.08 (29)	1.13 (24)	1.39 (16)	1.36 (21)	25
McDonalds	1.36 (16)	1.27 (22)	1.34 (21)	1.15 (25)	1.23 (26)	23
Merck & Company	1.31 (18)	1.35 (19)	1.53 (14)	1.51 (13)	1.47 (16)	18
Microsoft	1.55 (09)	1.56 (09)	1.84 (07)	1.75 (07)	1.53 (12)	08
Pfizer	1.52 (10)	1.48 (12)	1.54 (12)	1.57 (12)	1.42 (19)	11
Procter & Gamble	1.38 (14)	1.37 (16)	1.44 (17)	1.82 (05)	1.78 (05)	09
Travelers	0.94 (30)	1.00 (30)	0.79 (30)	0.79 (30)	0.99 (30)	***30***
United Technologies	1.22 (24)	1.23 (24)	0.97 (28)	0.95 (29)	1.25 (25)	27
Verizon	1.31 (18)	1.27 (22)	1.39 (19)	1.69 (08)	1.38 (20)	17
Walmart	1.25 (22)	1.36 (18)	1.55 (11)	1.43 (15)	1.33 (22)	20
Walt Disney	1.17 (27)	1.34 (20)	1.53 (14)	1.15 (25)	1.27 (24)	23

Note: In each of the first five statistical columns, average out-influence is listed first and average rank of the influences follows in parentheses. The sixth statistical column provides the rank (R) of the five combined ranks. The three highest combined ranks are signaled by boldface and the three lowest are designated by boldface italics.

Table 4. Average In-Influences and Corresponding Ranks for Each of the 30 Dow Jones Industrial Average (DJIA) Stocks on May 6, 2010

Company	Pre-Crash	Crash	Nadir	Recovery	Post-Recovery	R
3M	1.77 (02)	1.95 (01)	1.95 (04)	1.95 (01)	1.82 (02)	**02**
Alcoa	1.78 (01)	1.88 (02)	2.06 (01)	1.87 (03)	1.97 (01)	**01**
American Express	1.69 (03)	1.80 (03)	1.98 (03)	1.90 (02)	1.78 (04)	**03**
AT&T	1.63 (06)	1.71 (05)	1.68 (07)	1.75 (07)	1.82 (02)	05
Bank of America	1.10 (30)	1.41 (23)	1.53 (14)	1.47 (16)	1.53 (14)	18
Boeing	1.62 (08)	1.74 (04)	2.05 (02)	1.82 (05)	1.74 (06)	04
Caterpillar	1.64 (07)	1.70 (06)	1.86 (06)	1.80 (06)	1.78 (04)	06
Chevron	1.53 (13)	1.58 (10)	1.68 (07)	1.71 (09)	1.65 (08)	10
Cisco Systems	1.42 (21)	1.43 (21)	1.34 (25)	1.42 (20)	1.56 (12)	20
Coca-Cola	1.61 (09)	1.61 (08)	1.89 (05)	1.60 (10)	1.59 (10)	08
DuPont	1.66 (04)	1.69 (07)	1.62 (10)	1.83 (04)	1.74 (06)	07
ExxonMobil	1.37 (23)	1.34 (28)	1.18 (30)	1.20 (30)	1.24 (30)	***30***
General Electric	1.35 (24)	1.14 (30)	1.40 (21)	1.42 (20)	1.39 (26)	26
Hewlett-Packard	1.34 (26)	1.56 (12)	1.44 (17)	1.44 (19)	1.46 (22)	17
Home Depot	1.45 (15)	1.57 (11)	1.41 (19)	1.51 (15)	1.55 (13)	14
IBM	1.61 (09)	1.53 (14)	1.39 (22)	1.58 (14)	1.53 (14)	14
Intel	1.31 (28)	1.42 (22)	1.27 (27)	1.38 (25)	1.46 (22)	27
Johnson & Johnson	1.55 (11)	1.44 (19)	1.43 (18)	1.45 (18)	1.47 (21)	16
JPMorgan Chase	1.32 (27)	1.28 (29)	1.37 (23)	1.27 (27)	1.30 (28)	***28***
Kraft Foods	1.45 (15)	1.38 (26)	1.54 (13)	1.59 (11)	1.51 (16)	13
McDonalds	1.44 (19)	1.55 (13)	1.58 (11)	1.59 (11)	1.51 (16)	12
Merck & Company	1.46 (14)	1.48 (16)	1.21 (29)	1.37 (26)	1.41 (25)	24
Microsoft	1.24 (29)	1.35 (27)	1.24 (28)	1.25 (28)	1.34 (27)	***29***
Pfizer	1.35 (24)	1.48 (16)	1.45 (16)	1.40 (23)	1.48 (20)	20
Procter & Gamble	1.45 (15)	1.46 (18)	1.41 (19)	1.21 (29)	1.30 (28)	24
Travelers	1.65 (05)	1.51 (15)	1.68 (07)	1.74 (08)	1.60 (09)	09
United Technologies	1.55 (11)	1.60 (09)	1.58 (11)	1.59 (11)	1.59 (10)	11
Verizon	1.43 (20)	1.44 (19)	1.31 (26)	1.39 (24)	1.49 (19)	23
Walmart	1.45 (15)	1.40 (24)	1.35 (24)	1.42 (20)	1.51 (16)	20
Walt Disney	1.40 (22)	1.40 (24)	1.49 (15)	1.47 (16)	1.46 (22)	19

Note: In each of the first five statistical columns, average in-influence is listed first and average rank of the influences follows in parentheses. The sixth statistical column provides the rank (R) of the five combined ranks. The three highest combined ranks are signaled by boldface and the three lowest are designated by boldface italics.

Table 5. Average Self-Influences and Corresponding Ranks for Each of the 30 Dow Jones Industrial Average (DJIA) Stocks on May 6, 2010

Company	Pre-Crash	Crash	Nadir	Recovery	Post-Recovery	R
3M	0.15 (26)	0.10 (29)	0.11 (17)	0.15 (20)	0.16 (13)	25
Alcoa	0.15 (26)	0.14 (20)	0.09 (27)	0.17 (11)	0.13 (23)	27
American Express	0.18 (14)	0.17 (11)	0.12 (13)	0.15 (20)	0.15 (19)	18
AT&T	0.18 (14)	0.17 (11)	0.11 (17)	0.16 (14)	0.16 (13)	16
Bank of America	0.31 (01)	0.24 (02)	0.16 (04)	0.19 (06)	0.19 (04)	**01**
Boeing	0.18 (14)	0.15 (18)	0.09 (27)	0.14 (25)	0.16 (13)	23
Caterpillar	0.18 (14)	0.13 (23)	0.11 (17)	0.12 (28)	0.13 (23)	25
Chevron	0.16 (24)	0.14 (20)	0.09 (27)	0.14 (25)	0.13 (23)	***28***
Cisco Systems	0.22 (06)	0.21 (06)	0.12 (13)	0.20 (03)	0.17 (09)	05
Coca-Cola	0.17 (20)	0.15 (18)	0.11 (17)	0.15 (20)	0.16 (13)	21
DuPont	0.12 (29)	0.12 (27)	0.11 (17)	0.12 (28)	0.12 (28)	***30***
ExxonMobil	0.20 (10)	0.16 (16)	0.18 (02)	0.22 (02)	0.21 (01)	**03**
General Electric	0.22 (06)	0.27 (01)	0.14 (09)	0.16 (14)	0.20 (02)	04
Hewlett-Packard	0.22 (06)	0.13 (23)	0.11 (17)	0.18 (10)	0.18 (07)	14
Home Depot	0.21 (09)	0.13 (23)	0.13 (11)	0.16 (14)	0.13 (23)	19
IBM	0.17 (20)	0.17 (11)	0.20 (01)	0.17 (11)	0.16 (13)	11
Intel	0.23 (05)	0.17 (11)	0.13 (11)	0.20 (03)	0.17 (09)	06
Johnson & Johnson	0.16 (24)	0.16 (16)	0.14 (09)	0.20 (03)	0.17 (09)	13
JPMorgan Chase	0.19 (12)	0.22 (04)	0.11 (17)	0.19 (06)	0.18 (07)	08
Kraft Foods	0.26 (02)	0.24 (02)	0.15 (07)	0.15 (20)	0.15 (19)	10
McDonalds	0.19 (12)	0.13 (23)	0.11 (17)	0.17 (11)	0.15 (19)	20
Merck & Company	0.18 (14)	0.17 (11)	0.16 (04)	0.19 (06)	0.19 (04)	06
Microsoft	0.24 (03)	0.22 (04)	0.12 (13)	0.19 (06)	0.19 (04)	**02**
Pfizer	0.24 (03)	0.18 (09)	0.08 (30)	0.16 (14)	0.17 (09)	15
Procter & Gamble	0.17 (20)	0.14 (20)	0.11 (17)	0.23 (01)	0.20 (02)	12
Travelers	0.12 (29)	0.18 (09)	0.15 (07)	0.12 (28)	0.12 (28)	24
United Technologies	0.13 (28)	0.10 (29)	0.11 (17)	0.14 (25)	0.12 (28)	***29***
Verizon	0.20 (10)	0.19 (07)	0.17 (03)	0.16 (14)	0.16 (13)	09
Walmart	0.17 (20)	0.12 (28)	0.16 (04)	0.15 (20)	0.14 (22)	22
Walt Disney	0.18 (14)	0.19 (07)	0.12 (13)	0.16 (14)	0.13 (23)	17

Note: In each of the first five statistical columns, average self-influence is listed first and average rank of the influences follows in parentheses. The sixth statistical column provides the rank (R) of the five combined ranks. The three highest combined ranks are signaled by boldface and the three lowest are designated by boldface italics.

Table 6. Correlations among Out-Influence, In-Influence and Self-Influence Ranks for the 30 Dow Jones Industrial Average (DJIA) Stocks on May 6, 2010 by Economic Period

Influence	Pre-Crash	Crash	Nadir	Recovery	Post-Recovery	Combined
In vs. Out	-0.63 (<.0001)	-0.37 (.0206)	-0.67 (<.0001)	-0.75 (<.0001)	-0.53 (.0025)	-0.64 (.0001)
In vs. Self	-0.77 (<.0001)	-0.60 (.0002)	-0.59 (.0003)	-0.74 (<.0001)	-0.66 (.0001)	-0.79 (<.0001)
Out vs. Self	0.58 (.0003)	0.35 (.0282)	0.29 (.0627)	0.72 (<.0001)	0.72 (<.0001)	0.72 (<.0001)

Note: Out-influence, in-influence and self-influence values for each stock are contained in Tables 3, 4 and 5, respectively. R in each of these three tables denotes the rank for the combined ranks (last column). P-values the one-sided t-test testing the null hypothesis that the correlation coefficient is zero are in parentheses.

Table 7. Out-Influence, In-Influence and Self-Influence Changes for the 30 Dow Jones Industrial Average (DJIA) Stocks on May 6, 2010 by Economic Period

Influence	Pre-Crash to Crash	Crash to Nadir	Nadir to Recovery	Recovery to Post-Recovery
In-Influence Mean Difference	0.04 (.0508)	0.02 (.5150)	0.00 (.9754)	0.01 (.6597)
Out-Influence Mean Difference	0.04 (.3557)	0.02 (.7123)	0.00 (.9871)	0.01 (.8581)
Self-Influence Mean Difference	-0.02 (.0012)	-0.04 (< .0001)	0.04 (<.0001)	-0.01 (.0401)

Note: In-influence, out-influence and self-influence values for each stock are contained in Tables 3, 4 and 5, respectively. P-values for the two-sided paired t-test with the null hypothesis that the difference between two dependent means is zero are in parentheses.

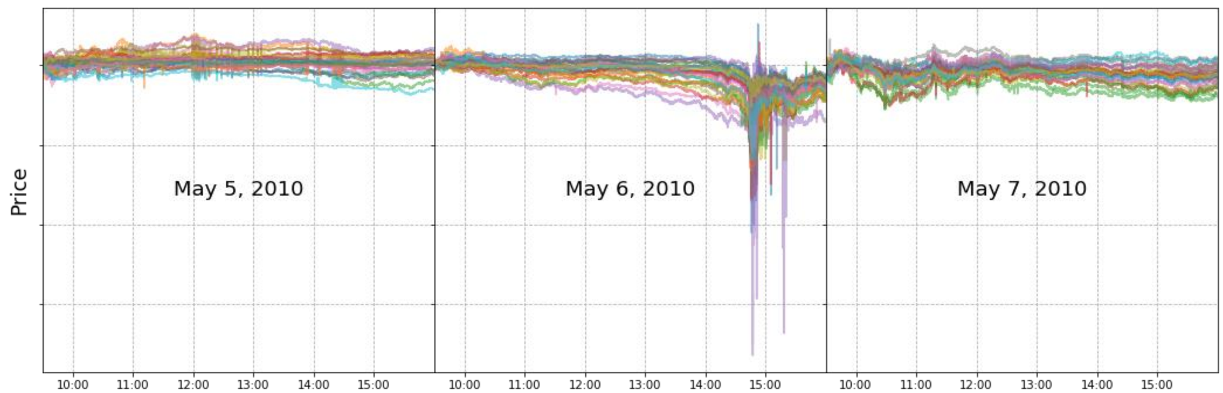

Figure 1. Prices for the DJIA 30 stocks on 5/5/2010 (Left), 5/6/2010 (Middle), and 5/7/2010 (Right) from 9:30 to 16:00 (x-axis) each day. The price series of each stock is standardized by its opening price in order to fit all the series in the same plot.

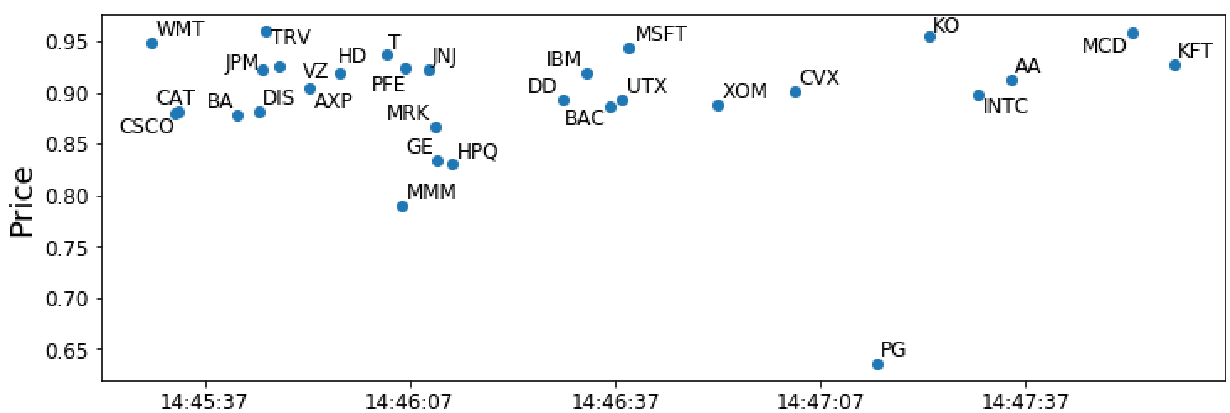

Figure 2. Time (x-axis) each stock reached its lowest price on May 6, 2010. Stock prices (y-axis) are standardized by their corresponding opening prices. Stock names corresponding to the stock symbols in the figure are given in Table 1.

	AA	BAC	GE	MSFT	TRV
AA	.14	.07	.06	.05	.03
BAC	.02	.44	.05	.05	0
GE	.04	.29	.17	.03	.05
MSFT	0	.14	.08	.09	0
TRV	0	.03	0	.13	.10

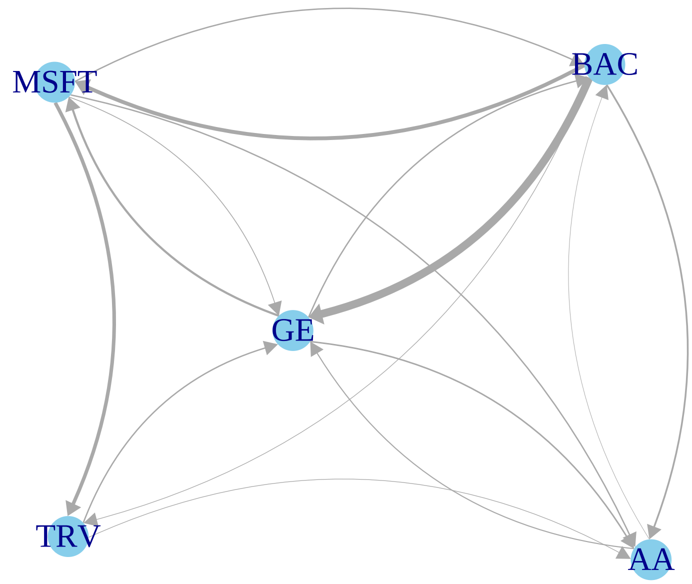

Figure 3. The excitation matrix for five randomly picked stocks during the 13:00:45-13:00:50 window (left) and the corresponding influence network (right). Directional arrows indicate source and recipient of the influence. Stock names corresponding to the stock symbols in the figure are given in Table 1.

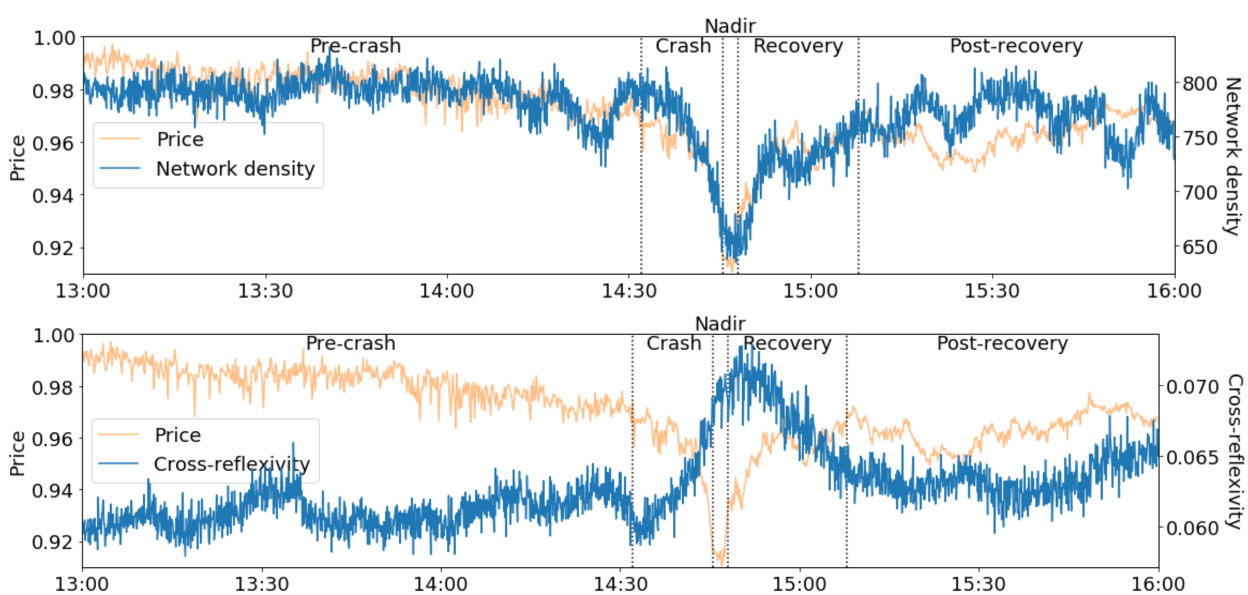

Figure 4. Network density (i.e., number of links in the influence network) (Top) and cross-reflexivity (i.e., average of the cross-influences between the 30 stocks) (Bottom) of the market and the average standardized price across all the DJIA 30 stocks from 13:00 to 16:00 (market close).

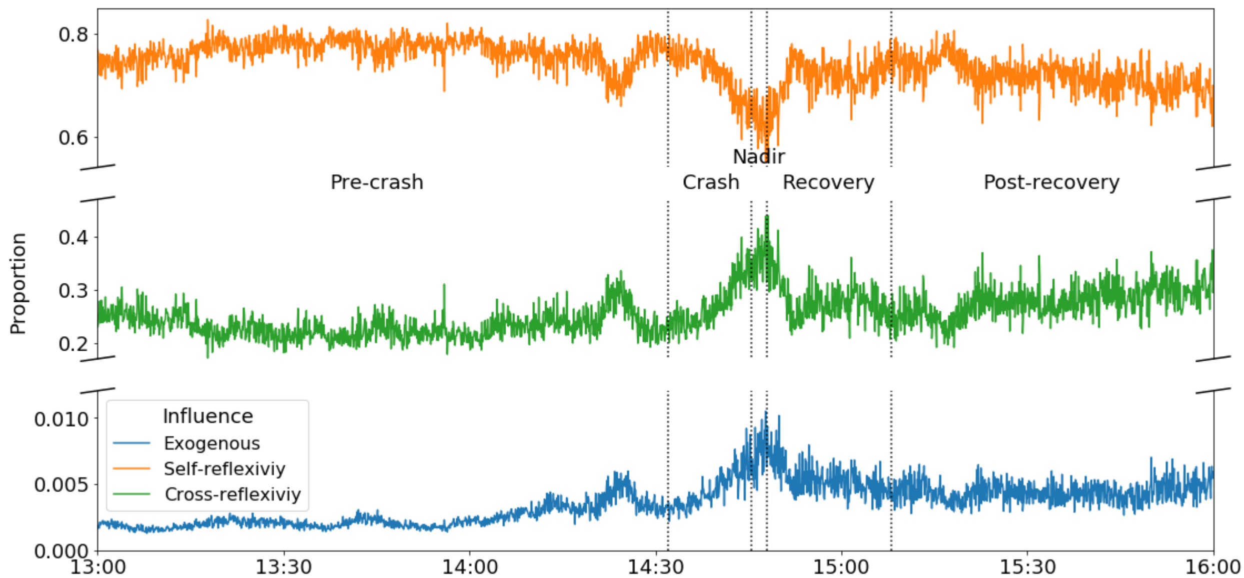

Figure 5. Proportions of exogenous effect, self-reflexivity, and cross-reflexivity from 13:00 to 16:00 (market close**).** The sum of the three influences at each time point equals one.

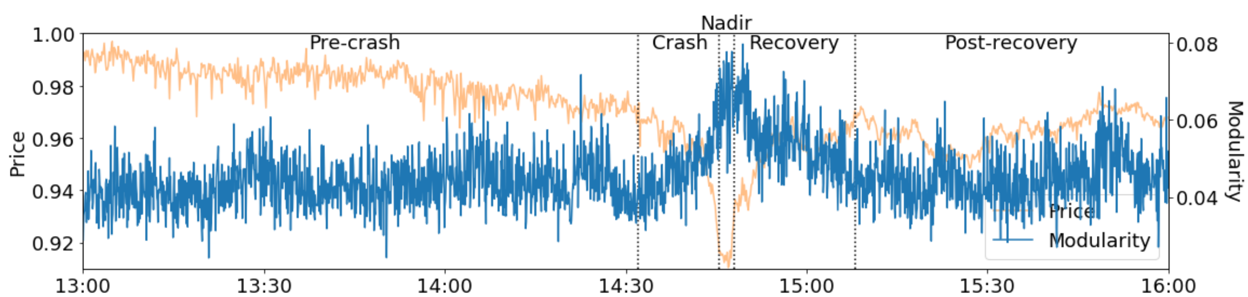

Figure 6. Network modularity from 13:00 to 16:00 (market close).

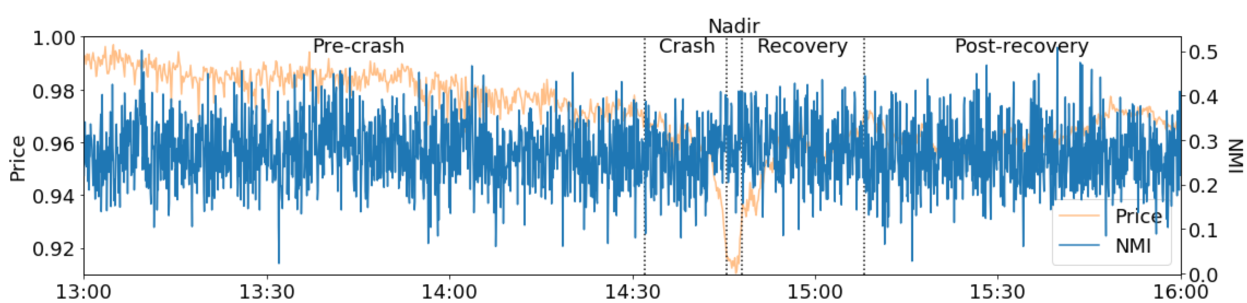

Figure 7. Normalized Mutual Information (NMI) between network communities and industry sectors from 13:00 to 16:00 (market close).