



The complex nature of financial market microstructure: the case of a stock market crash

Feng Shi¹ · John Paul Broussard^{2,3}  · G. Geoffrey Booth^{4,5}

Received: 16 April 2021 / Accepted: 11 December 2021

© The Author(s), under exclusive licence to Springer-Verlag GmbH Germany, part of Springer Nature 2022

Abstract

This paper uses multivariate Hawkes processes to model the transactions behavior of the US stock market as measured by the 30 Dow Jones Industrial Average individual stocks before, during and after the 36-min May 6, 2010, Flash Crash. The basis for our analysis is the excitation matrix, which describes a complex network of interactions among the stocks. Using high-frequency transactions data, we find strong evidence of self- and asymmetrically cross-induced contagion and the presence of fragmented trading venues. Our findings have implications for stock trading and corresponding risk management strategies as well as stock market microstructure design.

Keywords Stock market · Contagion · Networks · Hawkes processes · Granger causality · Adaptive learning · Behavioral finance

JEL Classification G12 · G14 · G23 · G41 · D47 · D80 · P1

1 Introduction

Stock markets are a key component of an economy based on the principles of capitalism. Their primary function is to provide a straightforward and convenient mechanism to transfer ownership of an asset claim in a formal trading venue that is governed by rules of trade. A trade is consummated when a price is agreed upon. Prices, however,

✉ John Paul Broussard
jpb@ou.edu

¹ University of North Carolina, Chapel Hill, NC, USA

² University of Oklahoma, Norman, OK, USA

³ Rutgers University, Camden, NJ, USA

⁴ The Citadel, Charleston, SC, USA

⁵ Michigan State University, East Lansing, MI, USA

often exhibit rapid fluctuations that at times are characterized by large and often unexpected changes. Taken together, a sequence of price changes involving a broad group of stocks in a relatively short time period that results in a meaningful cumulative decrease in their paper values is referred to as a market crash. As Glasserman and Young (2016) point out, interconnectedness is a defining characteristic of global and domestic modern financial systems. Although this phenomenon may provide transactional benefits, it also enhances the fragility of the system through common risk exposure via the ownership of similar assets, liquidity shocks, and macroeconomic shocks in general, making these risks difficult to manage because of network complexity and market crashes.

The purpose of this paper is (1) to model the price movements exhibited by individual stocks during the 2010 Flash Crash using a Hawkes process excitation matrix and (2) to interpret its entries in the context of complex networks and crowd behavior from the Granger causality and Adaptive Market Hypothesis perspectives. The analyses provide insights into the development of market design initiatives, trading strategies, and risk management methods that incorporate an intra-day, endogenous perspective.

Stock market crashes are not uncommon. In addition to the Flash Crash, recent crashes include Black Monday (October 19, 1987), the Dotcom Bubble Burst (March 10, 2000), the 2008 Financial Crisis (September 29, 2008), and the 2020 Pandemic (February 19, 2020). We explore the nature of the Flash Crash because its short length (36 min from the start through recovery) severely limits the number and resultant impacts of external factors, such as market intervention by government or stock exchange regulators that might affect the behavior of the market during its drawdown and subsequent recovery.

We focus on the 30 stocks that make up the Dow Jones Industrial Average (DJIA). The DJIA companies are large, blue chip firms chosen to represent the bulk of US economic activity. The Hawkes excitation matrix measures how the influence these 30 stocks have on each other and themselves and how these influences may evolve over time as traders learn about and react to the ever-changing environment. We use individual stock sub-second transactions data, which are the most granular data available at the time of the crash. This granularity permits the analysis of sequential transactions that often occur within milliseconds of each other, making it possible to construct a complex system containing multiple feedback loops that depict the impacts of a stock not only on itself but also on the other 29 stocks in the network. Thus, we provide a new lens on the behavior of stocks during the Flash Crash, which will help to understand the nature of other market crashes and other disruptions.

The Hawkes (2018) process is a stochastic model that describes the frequency of events within a specific time interval, allowing the occurrence of one or more of these events to increase the likelihood of triggering more or fewer events in the near future. Because of this characteristic, Hawkes processes have been used to quantify the endogenous and exogenous price effects in various asset trading markets. Our main focus is on the endogenous effects measured by the excitation matrix, which is a square matrix where its columns represent the stocks that are the influencers and the rows indicate the stocks being influenced. A stock's self-influence is measured by the matrix's principal diagonal. To provide a dynamic perspective, we estimate the excitation matrix using a rolling sample. This approach allows us to examine

how the excitation matrix evolves before, during and after the crash. Similar to any other statistical model, however, the Hawkes process mines the statistical association between the stocks but it cannot reconstruct the true causal relationship between them in terms of specific types of human cognitive biases.

Nevertheless, Etesami et al. (2016) demonstrate that the specification of the linear multivariate Hawkes model permits us to interpret the feedback relationships in the Granger-cause sense (Granger 1969, 2004). We analyze these relationships from the perspective of the Adaptive Market Hypothesis (Lo 2004), which uses the ideas and concepts associated with behavioral finance to help understand why stock markets behave as they do. The Adaptive Market Hypothesis recognizes that traders, being human, are not perfectly rational, although at times they might act as they are, and are subject to social and psychological biases. These biases may affect their assessments of current and future stock prices and these assessments may change as trading progresses.

We find that the influences between the 30 DJIA stocks increase on average during the Flash Crash and then revert close to their pre-crash level. The level of influence, however, varies greatly among stocks before, during and after the crash. Moreover, stocks that are strongly (weakly) influenced by other stocks are only weakly (strongly) influenced by their past behavior. Furthermore, although the industry sector of the 30 stocks does not seem to have an effect on the influences between them, the specific trading venue in which the transaction occurs does have a noticeable impact, suggesting that at times the ability to provide liquidity may vary among venues, which supports the notion of fragmented markets. These influence differences suggest the possibility that there may be worthwhile dynamic diversification strategies and market microstructure policies that have not yet been fully identified or exploited.

The remainder of the paper is divided into five sections. In Sect. 2, we present background information related to the 2010 Flash Crash and discuss possible price discovery network implications. We devote Sect. 3 to describing the role of complexity in the stock market, including the specification of the multivariate Hawkes model and how trader behavior can be thought to be embedded in a model's excitation matrix.¹ In Sect. 4, we discuss our data and their source, as well as provide several important descriptive statistics and caveats where appropriate. We also illustrate the interpretation of the excitation matrix and its corresponding network diagram for a subsample of DJIA stocks. We present our empirical results in Sect. 5 in three parts: (1) the 30 DJIA as a group using a simple average, (2) the 30 individual stocks and (3) the 30 stocks divided into industry sector and trading venue communities. In the last section, Sect. 6, we discuss our findings in a broader economic context, giving emphasis to their potential usage in portfolio risk measurement and management, stock index construction and the impact of various types of composition changes, and implications for fragmented markets.

¹ Traders can represent themselves, their firms, or individual institutional and private investors. We treat all relations the same and, therefore, use the terms "trader(s)", "investor(s)" and "stockholder(s)" interchangeably.

2 Crashes and price discovery implications

Crashes and their recoveries may be long or short and may or may not be followed by an economic recession or even a depression. Moreover, other than using arbitrary rules of thumb, it is often difficult to determine precisely when crashes begin and end. Numerous studies have been conducted to pin down the cause(s) of these crashes and many possibilities have been suggested. The cause(s) may be large or small, exogenous or endogenous, but, whatever the size or origination, the cause(s) must be able to promote the strong need for many stockholders, with large or small positions, to attempt to divest themselves of their holdings, at least temporarily. We focus on the timelines, possible reasons for flash crashes that have been suggested in the related literature, as well as implications for modeling crash behavior.

2.1 2010 Flash crash

The Flash Crash began on Thursday, May 6, 2010. US stock markets had been highly volatile since their openings as a result of, some believe, the disappointing economic news from Europe concerning its then-ongoing Greek debt crisis. As a result, O'Hara (2015) points out, buy and sell orders were becoming increasingly unbalanced. The Commodity Futures Trading Commission (CFTC) and the Securities Exchange Commission (SEC) in joint reports (CFTC-SEC 2010a; b) 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-h clock time. For the next 13 min or so after 14:32, stock prices continued to drop, with the largest declines occurring in the last few minutes. At the nadir of the crash, roughly \$1 trillion in the paper value of market capitalization was lost. By 14:45 markets began to stabilize, although there were still extreme movements in the prices of some stocks. Between 15:00 and 15:08 (the generally agreed upon end of the recovery), markets became noticeably less volatile as prices approached their pre-crash values and orderly trading resumed, although it took four additional trading days to reach May 10th's high. No external stimulus, government or otherwise, was involved and the decrease was not quite large enough to enact the market-wide circuit breaker that was initiated as a result of the 1987 Black Monday Crash. In the following days, stock market trading and prices appear to have proceeded in a "normal" manner and an economic depression or even a recession did not follow the day's seemingly unexpected and unsettling intraday event.

The cause(s) of the Flash Crash is unclear. The (CFTC-SEC 2010a; b) 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 et al. (2017), however, suggest that the flash crash was the result of the 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. Similarly, Easley et al. (2012) posits that this decrease may have been the result of the inability of cross-arbitraders to make a profit on their transactions and, hence, withdrew from the market. They also indicate that the Volume-Weighted Probability of Informed Trading Index (VPIN), a statistical measure of market toxicity, increased about 60 min

prior to the start of the crash, thereby signaling a possible short-term liquidity crisis. Menkveld and Yueshen (2019) highlight the impacts of a fragmented marketplace on the flash crash from a cross-arbitrage perspective using E-mini contracts that trade on the Chicago Mercantile Exchange and S&P 500 ETF Trust (SPY) positions that are tradable on eight different exchanges. They suggest that liquidity differences among the various SPY trading venues may have contributed to the crash, although Kirilenko et al. (2017) document that the trading pattern of high-frequency traders did not change during the crash.

2.2 Price discovery hypotheses and stock markets crashes

The function of a stock market is to bring various traders who want to buy or sell stock together so that they can jointly determine the value of a specific stock. Bid prices and ask prices are made by buyers and sellers, respectively. When they are equivalent in price, a transaction occurs. A market for a stock is considered liquid if there are many posted bid and asked prices that are very close to each other. An illiquid market is one in which there are only a very small number of bid prices and ask prices quoted and these quotes are far apart. In this type of market, a transaction will not be made unless there are major price concessions by one or both of the two parties. This buyer–seller interaction, whether or not there is a transaction, is typically referred to as the price discovery process. There are three main hypotheses concerning the way that investors engage in determining what they believe to be the true price of a stock: (1) the Efficient Market Hypothesis, (2) the Fractal Market Hypothesis, and (3) the Adaptive Market Hypothesis.

The Efficient Market Hypothesis (EMH) is the foundation of much of current finance theory and its relevant applications. It maintains that stock market participants make transaction decisions based on rational expectations created by rational individuals using all relevant and available information. The EMH is validated, at least on average, by the actual prices and how fast the information is incorporated in the price. For example, for the weak version of this hypothesis (the information set contains only market generated information), validation occurs when daily stock prices are characterized by a random walk, which in economic parlance means that the stock market is “efficient.”² Market efficiency implies that each transaction price is an equilibrium price because it considers all necessary information. Thus, Fama (1998) maintains that instances in which the random walk hypothesis does not appear to explain the behavior of stock returns, such as stock market crashes, should be considered anomalies, i.e., situations that occur by chance and should be discarded. After all, he (1998, p. 291) opines “...a model should be judged ... on how it explains the big picture.”

² Bachelier et al. (2007) is often thought of the founder of the notion that stock prices follow a random walk. Fama (1965, 1970), however, is largely responsible for introducing, developing and championing this concept, which supports the theory of efficient markets, to the academic and professional communities. Subsequently, Makiel (1973) popularized the notion and implications of market efficiency to the general public.

The Fractal Market Hypothesis (FMH) developed by Peters (1991, 1994) relies on Mandelbrot's (1963) concepts that are embedded in fractal geometry.³ In the stock market context, the geometric shape is usually thought of as a triangle with its left (right) side depicting a positive (negative) price change and its base measured in units of time, which, in a general sense is referred to as scale (Mandelbrot 1963, 2001; Mandelbrot and Hudson 2004). Triangles with different scales are combined to create patterns that resemble stock movements over time and are characterized by the presence of continuous discontinuities. The FMH posits that a stock market consists of many types of investors who have different investment horizons (scales) with correspondently different information needs, although all investors may be aware of all available information.⁴ For instance, fundamental investors may consider their time frame to be the business cycle while short-term traders may focus only on the day in which they are trading. Any significant change in trading horizons by either group may result in an overall shift of demand and supply of various stocks and this shift is reflected in stock price turbulence, which, among other factors, may reflect a lack of liquidity. As a result, the FMH is consistent with the notion that market crashes are not anomalies. Nevertheless, it is silent on the way that traders make the transition from one investment horizon to another.⁵

The Adaptive Market Hypothesis (AMH), which is largely attributed to Lo (2004), rests firmly on the notion that traders are human, have various cognitive biases, and do not operate in isolation but instead engage in various types of crowd behavior. In doing so, they not only gather objective information but also learn from others and other traders learn from them ad infinitum. Lo (2004) mentions that these biases include, but are not limited to, overconfidence, overreaction, loss aversion, herding (a.k.a. flocking), mental accounting, and regret.⁶ Bouchard (2010) emphasizes that herding is an important, persistent bias and typically is a result of traders imitating each other, a behavior that reflects the concern of traders believing that other traders may be smarter or have better information. Burzryn et al. (2014) maintain that the learning channel between traders is more effective when those who are less informed observe

³ "Fractal" is derived from the Latin word "fractured", which in Mandelbrot's lexicon refers to discontinuities at all scales. From Mandelbrot's (1963, 1982, 2001) perspective, fractals are complex patterns that are identical, infinite, and self-similar regardless of scale and are found in nature (e.g., shore lines and ice crystals) and the behavior of stock prices. Shenger (1994) does not agree with the accuracy of Mandelbrot's claim but indicates that fractals may be approximate estimations of the observed patterns.

⁴ The Homogeneous Market Hypotheses (HMH) developed by Müller et al. (1990) and Müller et al. (1995) is a special case of the FMH. As Weron and Weron (2000) explain, the HMH involves future expectations compared to what has been experienced. If future expectations differ from those that have been experienced, the change in homogeneity has the potential to disturb the market as investors attempt to change their portfolios accordingly.

⁵ Sohn and Sornette (2017) integrate the economic notions of the FMH and EMH. They posit that all time scales are not necessarily efficient. Thus, if a substantial portion of the market changes slowly its investment time scale from one that is efficient to one that is not and in doing so requires a different type of information, a bubble may be the result.

⁶ See Lo (1997, 2002, 2017), Lo and MacKinlay (1999) and Lo and Repin (2002) for detailed discussions on AMH and its dynamics.

the behavior of traders who are thought to be more sophisticated (successful).^{7,8} Moreover, Park and Sabourian (2011) show that traders herd when information is widely dispersed and, consequently, believe that extreme outcomes are likely. In this case, herding is a rational response in the “bounded” sense that tends to result in more volatile stock prices and lower liquidity and, thereby, is self-reinforcing.⁹ In summary, traders’ beliefs concerning stock prices may contain biases and these beliefs may evolve over time as trading progresses.

From a philosophical point of view, hypotheses are formal conjectures that are made to explain what is observed. Before a conjecture is accepted as potential truth, it should be based on logical reasoning and subject to rigorous testing. As more data become available and new tests are devised, existing hypotheses that previously could not be rejected are found to be acceptable under certain conditions or rejected in their entirety. Thus, a hypothesis, unless it is shown to be true by irrefutable evidence available now or in the future, such as a mathematical proof, can never be considered true in an absolute sense (Popper 1959/2002).

How do the EMH, FMH, and AMH fare with respect to Popper’s (1959/2002) philosophical view? Despite his defense of the EMH, Fama (1998) characterizes stock market crashes as anomalies. Kuhn (1970), however, believes evidence that is not consistent with the current paradigm should cause the emergence of new paradigms. Christensen and Raynor (2003, p. 27) agree with this sentiment and believe that instead we should do “...*anomaly-seeking* (italics theirs) research, not anomaly-avoiding research.” In contrast, the FMH tackles the issue of crashes by stating that they are not anomalies but are a result of investors moving from one fractal scale to a fractal with a smaller scale and eventually returning to their original scale. At the present time, there does not appear to be an in-depth analysis of the ways in which investors obtain the information relevant to changing fractal scales nor how they accomplish this change over time. The AMH differs from both the EMH and the FMH because it involves sequential investor learning not only from outside sources but also the behavior of other investors through feedback loops and cognitive biases.¹⁰

⁷ In the context of a housing market bubble, Bayer et al. (2021) show that the less informed investors typically underperform the more informed investors. They often refer to the less informed investors as being “infected”, meaning that they cannot miss out on the possibility of making a profit and, accordingly, follow the behavior of the crowd of more informed investors.

⁸ An unusually clear example of herding occurred in the spring of 2021. Investing in stocks became a game, according to Zweig (2021). Following 12 months of general price increases in U.S. stocks, individuals, mostly millennials, who openly admitted that they had no knowledge of stocks or stock markets began to trade using Robinhood and similar on-line stock buying apps and provided related comments on TikTok and WallStreetBets. One of these millennial traders garnered about half a million followers on TicToc. Some have referred to this behavior as gambling-style entertainment.

⁹ Simon (1979) points out that individuals, who need to make a decision are sometimes confronted with information that they do not understand, seek out alternatives to assist them in understanding the consequences and resolving uncertainties prior to making the decision. Since the requisite information may not be obtainable and action must be taken, the goal for the decision is not to yield the best possible result but, instead, a satisfactory one. Thus, the decision-maker makes a bounded rational decision.

¹⁰ Frankfurter and McGoun (2001) maintain that the AMH is not just a minor adjustment to the EMH; it is an alternate hypothesis that provides a reasonable rationale for many, if not all, the anomalies mentioned by Fama (1998), including stock market crashes.

AMH's perspective on human behavior is compatible with the econophysics view that stock markets are complex systems. For instance, Sornette (2003/2017, p. 280) maintains that a stock market is a complex system and similar to other complex systems "...has an endogenous origin and exogenous shocks only serve as triggering or tipping point factors." The endogenous origin refers to the cooperative behavior of market's participants, and the exogenous shocks to events that occur external to the market. Arthur (1999) opines that the notion of complex systems applied to a stock market contradicts the EMH because complexity in this context involves cooperative behavior of the market's participants.

3 Stock market complexity, modeling, and human behavior

In this section, we discuss market complexity in terms of a network, its relationship to the concepts of endogenous and exogenous, and how we statistically measure and interpret these results from the perspective of the AMH.

3.1 Complex networks

Much of the literature on complex systems focuses on physical or biological topics.¹¹ Ladyman et al. (2012), however, maintain that important complex features are also found in social systems, including stock markets. They indicate that these markets are highly complex because there are numerous potential buyers and sellers who often randomly interact with one another, thereby providing ongoing feedback to each other and themselves. Buyers and sellers need not be human; transactions can be undertaken by algorithms developed by humans or by artificial intelligence. It is the feedback process that creates the endogeneity present in the stock market. This process, however, is dynamic and traders may face Knightian uncertainty, i.e., the probability of stock price moves or the factors that might cause these changes are unknown. As a result, these traders learn about and adapt to the changing market environment, which is sometimes made more difficult in a changing technological environment.¹²

According to Kuhlmann (2014), crashes can occur endogenously or exogenously. Endogeneity results in a market where traders search for an equilibrium in a stock's price but never attain it. As a result, extreme events like crashes can occur when some of the traders in the system adapt to or adopt the behaviors of other traders. In this regard, Ladyman et al. (2012) maintain that, although traders may interact with each other in a way that is disordered, their actions as a group are organized because the traders not only communicate information to each other, but they do so by using a defined

¹¹ Examples of early research in applying the concept of complexity to economics and finance in particular are given in Anderson et al. (1988) and Arthur et al. (1997). For a discussion of the philosophical role of complexity in science and its impact see Ladyman et al. (2012).

¹² Arthur (1999) suggests that instead of the term "complex" system a more descriptive name might be "adaptive nonlinear network" and attributes this term to J. H. Holland. Ladyman et al. (2012) indicate that the network in the case of the stock market can be considered linear since the interactions among traders can be represented by a standard matrix. Easley and Kleinberg (2012) offer a *tour de force* of networks and their concepts applied to markets.

set of formal and informal rules. Frank et al. (2019) provide evidence that in upstairs markets traders interact with each other and sometimes appear to engage in reciprocal agreements, especially if the traders belong to the same investment house.

Any specific exogenous event can be relatively large or small. Yet, regardless of the event type, it acts as a trigger that initiates a crash with a significant number of stockholders trying to sell all or a significant number of their stock holdings, thereby driving stock prices down. Fleeing from the market is an example of what psychologists have dubbed as “fight or flight” behavior, which are two physiological responses by humans facing real or perceived danger (Canon 1929).¹³ In this context, the act of fleeing (selling) is, according to Goldstein and Kopin (2007), ingrained behavior triggered by a person’s sympathetic nervous system. This behavior does not depend on whether the event or events occur; it can also exist if there is a fear (or an anxiety) that they will occur. From this perspective, market recoveries occur in the same manner. Previous stockholders seek to regain their earlier exposure to equities or new stockholders enter the market because they fear missing the possibility of significant profits in the future.¹⁴

Some types of events can be either endogenous or exogenous, and their designation depends on the context in which they occur. For example, the lack of liquidity is often proffered as a potential cause of a market crash. Liquidity could be exogenously decreased by a halt in trading as a result of a circuit breaker being activated. It could also be endogenously lessened by market-makers reducing their market presence because of dwindling trading profits. The first reason is the result of market microstructure rules and the second by market toxicity, i.e., an imbalance between buyers and sellers. With respect to toxicity, Liu et al. (2021) provide empirical evidence that there is a strong negative relationship between herding and liquidity.

How does this behavior manifest itself in a stock market context? We address this issue in the next section, which focuses on the ways in which individual stocks are connected in complex networks and how these connections exhibit their strength. According to Kuhlmann (2014, p. 1124), these networks seem to be able to “...identify the common underlying structural mechanism.”

3.2 Statistical methods and approaches

We use the multivariate Hawkes process to model the complex microstructure behavior of the DJIA 30 stocks. Hawkes (2018) provides a brief summary of his process and gives a useful and extensive bibliography pertaining to the model’s development and

¹³ Bracha et al. (2004) update Canon’s (1929) concept to “freeze, flight, and fight.” The adding of “freeze” and the reordering of “fight” and “flight” seem even more applicable to a crash situation. “Freeze” suggest that the stockholder delays the decision to transact before acting. It also suggests that “fighting” may indicate that a stockholder must fight the urge to sell and stay in the market.

¹⁴ Another twist to this idea is the fear of missing out (FOMO), i.e., individuals doing something different than their peers. In a rational general equilibrium setting, Demarzo et al. (2008) show that investors are greatly influenced by the wealth of their social cohort. In rising markets investors want to keep up with their peers, and in a declining market losing together is more palatable than losing alone. In addition, a larger loss may signal that the investor has more wealth than his contemporaries and a smaller loss may suggest that the investor is smarter.

finance applications, with a special mention of Bacry et al. (2015) for high-frequency applications.¹⁵ Previous applied research on this area that is particularly relevant to our study are Filimonov and Sornette (2012) and Aït-Sahalia et al. (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 100% during the Flash Crash. Aït-Sahalia et al. (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.¹⁶ 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.

In contrast to the above studies, 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 immediately previous transaction. Without ambiguity, we refer to such price-changing events simply as events. These events act as a proxy for a change in the traders' view of the paper value of the stock in question as a result of their on-going learning experiences. O'Hara (2015) maintains that positive price changes signal a mixture of good news and negative changes are consistent with bad news, with both types of price changes containing some noise. Passive trades are transactions not associated with a price change. Because there are multiple stocks considered, price-changing events from each stock are treated as unique.

3.2.1 The Hawkes model¹⁷

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., $s \in \{1, 2, \dots, 30\}$. 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: $\lambda_s(t): P_s \approx \lambda_s(t) \cdot dt$. The rate of events $\lambda_s(t)$ is modeled as a function of the occurrences of previous events from all the stocks (including itself):

¹⁵ In addition to finance related applications, the Hawkes process has been applied to model and analyze network systems in many areas such as bioinformatics, seismology, criminology, terrorism, and social networks in general (Zadah and Shards 2015).

¹⁶ A similar approach has been used in other seemingly discontinuous stock return data. Instead of a Hawkes process a Poisson process is used. A Poisson process is similar in concept to a Hawkes process but the occurrence of a jump is not influenced by past observations and, therefore, does not address the concept of learning and adapting. Merton (1976) is the first to model financial asset price behavior using a mixture of diffusion and Poisson processes. His modeling approach was subsequently adopted by many researchers, e.g., Akigray and Booth (1988), Andersen et al. (2002), and Cai and Kou (2011).

¹⁷ Additional details on the Hawkes process may be found in numerous sources but Liniger (2009), Embrechts et al. (2011), Rizoïu, et al. (2017) and Hawkes (2018) are particularly helpful.

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

where $(i : t_i^{s'} < t)$ corresponds to all the events of stock s' that occurred before time t ; $a_{ss'}$ captures the impact of stock s' on stock s ; μ_s is the baseline rate of events of stock s that is independent of previous events; and $g(t - t_i^{s'})$ is the memory kernel that models how the effect from each previous event decays over time. More detailed descriptions of the variables are given in Table 1.

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 Eq. (1) compactly with matrices. Let $\lambda(t)$, $\mu(t)$, A , and $g(t)$

Table 1 Hawkes model symbols and definitions

Symbol	Definition
t_i^s	Time of the i th event of stock s
$\lambda_s(t)$	Rate of events associated with stock s at time t
$\lambda(t)$	Vector of rates of events associated with all stocks: $\lambda(t) = (\lambda_1(t), \dots, \lambda_{30}(t))^T$
μ_s	Baseline rate of events of stock s independent of previous events. Since our focus is on the 30 DJIA stocks, it represents all exogenous effects to the market such as news and effects emanating from all the stocks not included in the DJIA
$\mu(t)$	Vector of baseline rates of events: $\mu(t) = (\mu_1(t), \dots, \mu_{30}(t))^T$
$a_{ss'}$	Influence of stock s' on stock s . It is the expected number of events of stock s preceded by an event of stock s' . There are 30×30 such parameters: $a_{ss'}, s = 1, \dots, 30, s' = 1, \dots, 30$. All $a_{ss'} > 0$
A	Excitation matrix for all stocks: $A = [a_{ss'}]_{s=1, \dots, 30; s'=1, \dots, 30}$
$g(t - t_i^{s'})$	Memory kernel. The kernel 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. We adopt the logistic-normal density as in Linderman and Adams (2014). We pick this function because it is a probability density function and it integrates to one, which endows $a_{ss'}$ with the units of the “expected number of events” and allows a comparison of the strengths of interactions. Moreover, it has bounded support and naturally models the domain of $g(t - t_i^{s'})$. Since we only consider impact from past events, $t - t_i > 0$, the domain of $(t - t_i)$ is bound below by zero. 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., $t - t_i < \Delta t_{max}$
$g(t)$	Vector of memory kernels for all stocks: $g(t) = \left(\sum_{i:t_i^1 < t} g(t - t_i^1), \dots, \sum_{i:t_i^{30} < t} g(t - t_i^{30}) \right)^T$

represent the combined matrices for λ_s , μ_s , $a_{ss'}$, and $\sum_{i:t_i^{s'} < t} g(t - t_i^{s'})$, respectively.

Eq. (1) then becomes:

$$\lambda(t) = \mu(t) + A \cdot g(t). \quad (2)$$

Equation (2) highlights the important role that excitation matrix A plays in the Hawkes process. It is the crux of our analysis because it captures the interactions between all of the stocks. Etasami et al. (2016) point out that the excitation matrix is equivalent to a minimal generative model and this type of model, according to Quinn et al. (2011), measures directed information that can be interpreted as Granger causality (Granger 1969, 2004). Thus, A 's columns depict the stocks that trigger the effects, and its rows denote the stocks that are 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.

3.2.2 Model estimation

The model is estimated using the Bayesian method as adopted by Linderman and Adams (2014) and their Python package "pyhawkes" (<https://github.com/slinderman/pyhawkes>), which has been found efficient and reliable in discovering latent network structure in both synthetic and real data. Below we briefly describe the intuition behind their Bayesian inference method. (See Linderman and Adams (2014) for details.)

Let $p(\{t_i^s\} | \mu, A, g)$ be the likelihood function of observing the sequences of events $\{t_i^s\}_{s=1}^{30}$ given the parameters μ , A , g , which is well-defined as a result of the assumptions of the Hawkes model. A traditional approach to estimating the parameters is to compute their maximum likelihood estimate by maximizing this likelihood function. Because of the complex structure of the likelihood function, however, there is no closed-form solution nor is there a straightforward approach to optimize it. Nevertheless, Linderman and Adams (2014) develop a novel and efficient Bayesian inference algorithm by taking advantage of the superposition property of the Hawkes model, which simply states that the total response equals the sum of the individual responses.

Before introducing the Bayesian inference method, we note one extra step in Linderman and Adams (2014) and decompose the excitation matrix A into two parts: $A = A' \times W$, where A' is a binary matrix modeling the structure of the network ($A'_{ij} = 1$ if there is an edge between nodes i and j and $A'_{ij} = 0$ otherwise), and W is a non-negative weight matrix that models the strength of the edges between nodes. The advantage of this separation is that recent advances in random graph models can be used to describe the network structure and impose separate beliefs about the strength and the structure of the network as Bayesian priors.

Given the likelihood function $p(\{t_i^s\} | \mu, A', W, g)$, the Bayesian inference method proceeds as follows, which is an example of Gibbs sampling (a procedure

that uses a Markov Chain—Monte Carlo algorithm) with modifications (Linderman and Adams 2015). In particular, assuming $\Pi(\boldsymbol{\mu})$ is a prior distribution of $\boldsymbol{\mu}$, by Bayes' rule, the posterior distribution of $\boldsymbol{\mu}$ is: $p(\boldsymbol{\mu}|\{t_i^s\}, A', W, \mathbf{g}) = p(\{t_i^s\}|\boldsymbol{\mu}, A', W, \mathbf{g})\Pi(\boldsymbol{\mu})/p(\{t_i^s\}, A', W, \mathbf{g})$. We sample $\boldsymbol{\mu}$ from this posterior as its estimate. Similarly, we derive the posterior distributions of other parameters A' , W , \mathbf{g} , respectively, and sample from their posteriors. This process is then repeated multiple times and the average of the samples is used as the estimate for each parameter.

Linderman and Adams (2014) show that their method achieves better results in discovering network structures than standard estimation methods on both synthetic and real data. One contributing factor to their success is the direct modeling of network structure. Another advantage of their Bayesian method is the capacity to impose prior distributions on the parameters, which to some extent regulates noise in the interactions between the stocks. In addition, the Bayesian approach is the equivalent of updating its parameters, including the components of the excitation matrix, from samples of past values. Conceptually, this is similar to the approach exhibited by human behavior underlying the AMH. Both involving learning from the past and adapting their behavior to attempt to optimize their behavior.

4 Data and descriptive information

In this section, we describe the sample stocks and their sources as well as demonstrate with actual data how to use and interpret the excitation matrix. We also divide the Flash Crash timeline into five economic periods, provide some important statistics on price changes, and discuss the issues using statistics based on large samples.

4.1 Data source and composition

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. These companies, which are shown in Table 2, are very large, publicly traded, US-based, and most were originally listed on the New York Stock (NYSE) with the remainder on Nasdaq exchange (NQN).¹⁸ Their stocks can be traded on their original listing exchange as well as at eight other exchanges as long as they have a dual listing with them. These venues

¹⁸ 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, making "Industrial" a misnomer. Since May 6, 2010, Alcoa (AA), American Telephone & Telegraph (T), Bank of America (BAC), DuPont (DD), ExxonMobil (XOM), General Electric (GE), Hewlett-Packard (HPQ), Kraft Foods (KFT), Pfizer (HC) and United Technologies (UTX) have been replaced on or before August 31, 2020. As of this date the replacement stocks are Apple (AAPL), Amgen (AMGN), Dow (DOW), Goldman Sachs (GS), Honeywell International (HON), Nike (NIKE), Salesforce.com (CRM), United Health Group (UHG), Visa (V), and Walgreens Boots Alliance (WBA).

Table 2 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

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)

include either NYSE or NQNM and seven smaller exchanges. Trades can also be made off-exchange but these transactions must be reported to a separate unit overseen by the industry. The behavior of stocks not in the DJIA and all other financial instruments related to any stock or group of stocks is considered exogenous and their impacts are treated accordingly.

Virtually all of the 30 companies are household names and represent almost all of the major sectors in the US economy, i.e., consumer staples, industrial materials, financials, telecommunications, energy, consumer discretionary, information technology, and health 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% of the market value of all traded US stock around the time of the flash crash.

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 US exchanges where a given stock is traded, which is not necessarily where it was originally listed. The primary data extracted from the Nanex feed used in our analysis are transaction prices and their time stamps.

For each of the DJIA 30 stocks listed in Table 2, 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 Fig. 1. For comparison, we include the standardized price series on the day (May 6, Fig. 1-middle) of the flash crash as well as the standardized prices from one day before (May 5, Fig. 1-left) and one day after (May 7, Fig. 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.

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 see clearly the timing and the magnitude of the price drops of each of the 30 DJIA stocks on May 6th, we show the lowest standardized price for each stock (y -axis) and the time each stock reaches its nadir (x -axis) in Fig. 2. The time points for these prices are clustered between 14:45 and 14:48 with Walmart (WMT) being the first stock (14:45:29.2) to reach its lowest price and Kraft Foods (KFT) to be the last

Footnote 18 continued

In this study, General Electric is the sole surviving member of the original 12 companies. Santilli (2020) provides a graphic history of the DJIA stocks with respect to the dates of their entry to and exit from the index.

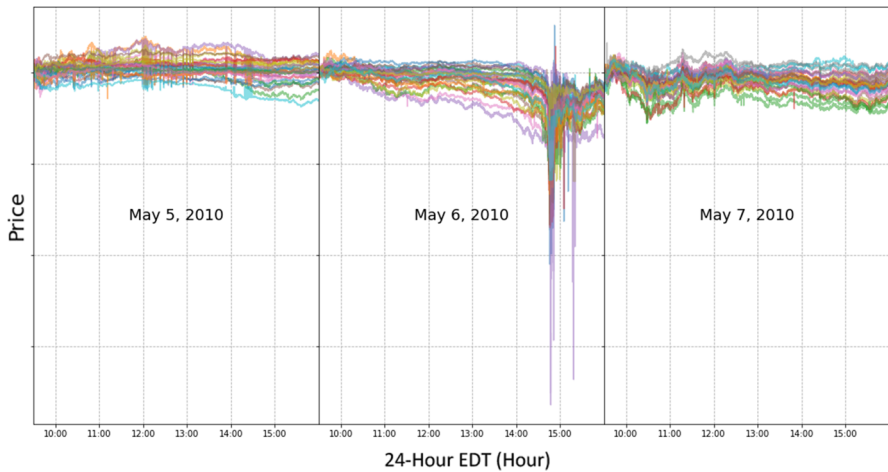


Fig. 1 Prices for the DJIA 30 stocks on 5/5/2010 (Left), 5/6/2010 (Middle), and 5/7/2010 (Right) from 9:30 to 16:00 (x -axis) each day. The price series of each stock is standardized by its opening price in order to fit all the series in the same plot. Subtracting 1.0 from the result of this standardization procedure creates a measure of return based on the stock's price at the beginning of the standardization period

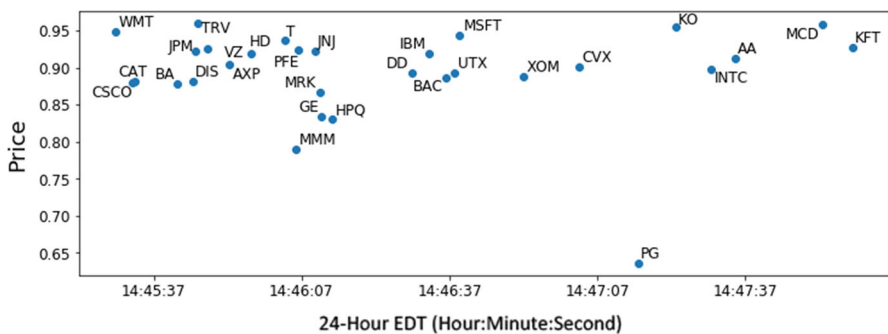


Fig. 2 Time (x -axis) each stock reached its lowest price on May 6, 2010. Stock prices (y -axis) are standardized by their corresponding opening prices. Stock names corresponding to the stock symbols in the figure are given in Table 2

(14:47:58.8). Although most stocks dropped about 10%, 3 M (MMM) dropped 20% and Procter & Gamble (PG) dropped more than 35%. The observations in Figs. 1 and 2 are consistent with the reports from the CFTC-SEC (2010a; b) 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.

4.2 Modeling excitation matrix dynamics

Because the behaviors of stocks may vary over time, especially during the flash crash, we do not fit the Hawkes model to all the data combined as this would mask the




Fig. 3 The excitation matrix for five randomly picked stocks during the 13:00:00–13:05:00 window (left) and the corresponding influence network (right). Directional arrows indicate source and recipient of the influence, and their thickness represents the strength of influence. Stock names corresponding to the stock symbols in the figure are given in Table 2. As indicated by the principal diagonal of the matrix, all stocks exhibit self-influence with BAC having the highest value and MSFT having the lowest. Arrows in the diagram that indicate self-influence are not provided for visual simplicity

dynamics of the stocks over time. Instead, 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. Thus, each five-minute moving window, on average, lengthens the rolling sample by 1.67% to accommodate new observations and shortens the sample by the same percent by deleting the oldest observations. The Hawkes model is then fitted in every window to the DJIA 30 stock data.

For the visualizations that follow, the parameter estimates from each window are plotted against the right boundary of the window. From the 5-min 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-min time window prior to it, their calculated effects reflect a moving average and are not instantaneous and better reflect the investor learning process.¹⁹

An example of the excitation matrix for five randomly picked stocks during the 13:00:00–13:05:00 window is shown in Fig. 3, together with its network representation. Column names denote the stock that influences, and row names indicate the stocks being influenced. 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

¹⁹ A window spanning five minutes is used to ensure that there are adequate observations to estimate Eq. (1). In doing so, we implicitly assume that there is no big shift in the dynamics in the window and, thus, choosing 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 use a 10-min window, and the empirical results are qualitatively the same but are not reported.

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. Self-influence is not shown in the network diagram because it would only add unnecessary detail. Nevertheless, the excitation matrix entries show that BAC exhibits the most self-influence and is substantially higher than the other three stocks in our example.

To extend this example consider two distinct scenarios: (1) none of the five stocks influenced themselves or each other, and (2) each of the five stocks only influenced themselves. In the first case, the excitation matrix would only contain zero entries. In the second situation, the principal diagonal cells of the matrix would have non-zero values, but all the other entries would be zero. Including noise to the example makes the scenarios a bit more difficult to picture because non-zero values would be added to the excitation matrix's entries. Nevertheless, because excitation matrices are calculated using a rolling sample approach in conjunction with a Bayesian algorithm, the impact of noise on the matrices should be mitigated.

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 the network.²⁰ 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 a 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 its diagonal entry. Correspondingly, the in-influence 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 the 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 the impacts may be asymmetric. Finally, it does not suffer from the Epps

²⁰ 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.

(1979) effect, which typically renders the variance–covariance matrix unreliable for high-frequency data.²¹

4.3 Economic periods and sample size implications

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 3. The time period spanned by the crash, nadir and recovery periods corresponds to the SEC's (2010a; b) crash period. Our nadir period is structured so that it contains the lowest price observation for each of the 30 stocks as depicted in Fig. 2. Table 3 also contains information about the number of transactions in each of the five periods and these periods in aggregate. The table also provides the frequency of transactions (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.)

As displayed in Table 3 price-changing trades, often referred to as active trades, never exceed 50% of total trades. In the pre-crash stage, they account for only 8.49% of the total number of trades. The percentage monotonically increases to 44.62 in the nadir period and decreases to 19.31% in the post-recovery stage. The total number of active trades across the five economic periods combined is slightly over one million, which translates to one transaction occurring every 10.1 ms, on average. The frequency of trades monotonically increases from one every 22.1 ms in the pre-crash period to one every 2.8 ms in the nadir period and then monotonically decreases to one every 7.8 ms in the post-recovery period.²² This U-shape pattern is roughly echoed by the proportion of trades that are associated with positive price increases. This proportion drops from 49.1% in the pre-crash period to 48.3% in crash period and rises to 50.4% in the nadir period and subsequently becomes relatively flat for the remaining two periods.

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

²¹ Epps (1979) points out that the covariance between two individual 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 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.

²² To put these time intervals into perspective, 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 ms. Saariluoma (1995) indicates that a chess grandmaster's latency to recognize that his king is endangered by an opponent's move is 650 ms. 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% 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 ms, respectively.

Table 3 Economic crash and recovery periods and selected trade information on May 6, 2010

	Pre-Crash	Crash	Nadir	Recovery	Post-Recovery	Total
<i>Clock time</i>						
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
Total trades	2,960,429	499,528	120,270	591,977	1,348,492	5,520,693
<i>Price-changing trades (active)</i>						
Number	250,101	150,819	53,884	210,933	400,604	1,066,141
<i>Percent of total</i>						
Trades	8.49	30.19	44.62	35.63	29.71	19.31
Frequency	22.07	5.365	2.780	5.694	7.788	10.13
<i>Price increase</i>						
% of total	49.07	48.26	50.35	50.41	50.24	49.73
Z-score	-9.302	-13.52	1.653	3.766	3.038	-5.576
<i>p value</i>						
Two-tail	< 10 ⁻²⁰	< 10 ⁻⁴¹	.1040	.0002	.0023	< 10 ⁻⁸
One-tail	< 10 ⁻²¹	< 10 ⁻⁴²	.0521	.0001	.0012	< 10 ⁻⁸

The time periods are expressed in 24-h 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. Total trades are the sum of active and passive trades. Passive trades are trades that are not associated with a price change. Price-changing trades are active trades. Frequency is the average number of milliseconds between active trades. Price Increase is the percent of active trades characterized by an increase in price; its complement is the percent of price decreases. The Z-score's two-tailed *p* value tests the null hypothesis that the portion of the total number of trades associated with price increases and decreases are both 50% of the active trades. The one-tailed *p* value tests the null hypothesis that the portion of price increases (decreases) is 50% or less (greater). The price increase percent, Z-score, and two-tailed *p* value for the duration of the official crash, i.e., crash, nadir, and recovery, are, respectively, 49.62, -4.900, and < 10⁻⁷. Adding the post-recovery period to the official crash period results in a 49.92% of price increase with a Z-test of -1.446 with a two-tailed *p* value of .1482.

usually account for 50% of the total. Z-scores for the trade proportion for each of the five economic periods and the five periods combined are provided in Table 3. We use a two-tailed Z-score to test the null hypothesis that the positive and negative price-change trades each account for 50% of the active trades. A one-tailed Z-score is used to test the hypothesis that a positive (negative) price change is less (more) than or equal to 50%.

Turning first to the two-tail Z-scores, we find that values for all periods combined, and four of the five sub-periods are statistically significant using any standard critical p value, indicating that the null hypothesis that the negative and positive price changes are equally balanced is rejected. The exception is the nadir period, which is insignificant using the traditional 0.10 critical value. The one-tail Z-scores help to pinpoint the source of the imbalances. Negative price changes in the pre-crash and crash periods statistically dominate the positive price changes and the reverse is the case for the recovery and post-crash periods. Over the entire test period the number of negative price changes dominate.

Thus, from the above statistics, it would seem that in all of the above cases we should reject the null hypothesis that the behavior of the two types of price changes are the same even if the actual numbers are slightly different from the hypothesized values. 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 a result, the size of the universe from which the sample was extracted is not considered so that in relative standards a large sample may be only a small fraction of the total population.

As pointed out by Lin et al. (2013), among others, several approaches have been suggested to mitigate this problem, which is becoming increasingly important in the era of “Big Data.” 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 requires the p value signaling statistical significance to be specified ahead of time.²³ The second involves 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 that positive 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 that contained new information.

²³ 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 researchers 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.

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 may be exacerbated 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% of the price-changing events are positive during the crash period, and 50.41% 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% of the total. Numerically, these are very close to the 50% 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 the 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 (1776/2020) well known and often mentioned "invisible hand" metaphor at work, which in our case is moving price toward a moving equilibrium value.

Second, market regulators throughout the world are concerned about the effectiveness of the price discovery process. A majority of them, including those in the USA, 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. This delay gives traders an opportunity to come to grips with their cognitive biases and, as a result, to better evaluate the reasons 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.²⁴

On June 19, 2010, approximately one month after the May 6 flash crash, US regulators established a single stock circuit breaker to guard against falling prices. On April 5, 2011, the Financial Industry Regulatory Authority (FINRA), 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.²⁵

²⁴ Whatever type of regulation is imposed, it must be able to promote fair competition among all parties involved. This is especially important when advance computer technology permits stock market access to both small and large as well as new investors to trade on intraday basis. Doing so requires an understanding of the endogenous nature of the market and the various cognitive biases of its participants.

²⁵ Conceptually, the LULD is a straightforward reference price constructed by calculating a simple average of the transaction prices during the previous five minutes. The 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 s and is updated if the new reference price is

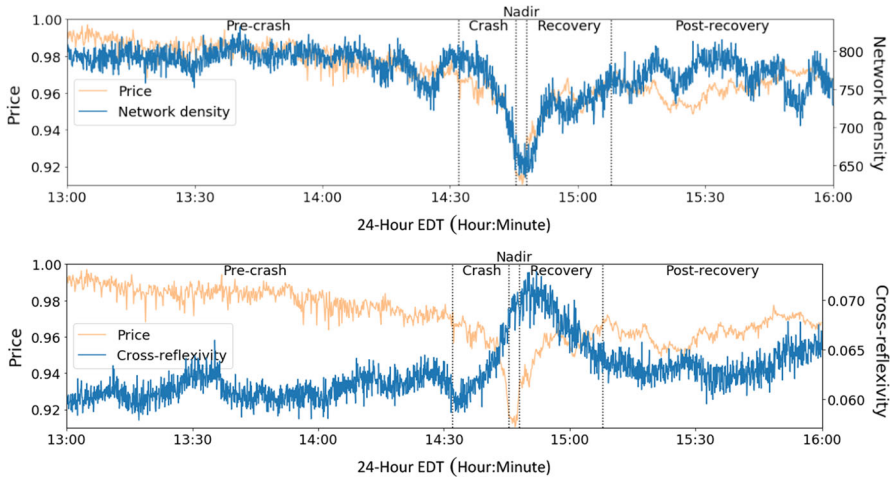


Fig. 4 Network density (blue), i.e., number of links in the influence network (Top) and cross-reflexivity (blue), i.e., average of the cross-influences between the 30 stocks (Bottom), and the average standardized price of all DJIA 30 stocks (orange) from 13:00 to 16:00 (market close) (color figure online)

5 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. We also report the possible presence of subgroups based on industry sector and trading venue. For convenience we refer to the 30 stocks taken together as the “market,” individually as simply “stock,” by company name or stock symbol, and in subgroups as communities. All results are based on excitation matrices calculated using the five-minute rolling window.

5.1 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 Fig. 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.²⁶ 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 Fig. 4-bottom. The abrupt increase of cross-reflexivity around 14:32 reflects the beginning of the sudden decline in stock prices.

Footnote 25 continued
 at least one percentage point away from the current posted reference price. (See www.luldplan.com for details.).

²⁶ The average price in Fig. 4 is the simple mean of the normalized prices of the 30 stocks in the DJIA. It is not the official DJIA, which is also an average. See fn. 18 for details on the difference between these two averages.

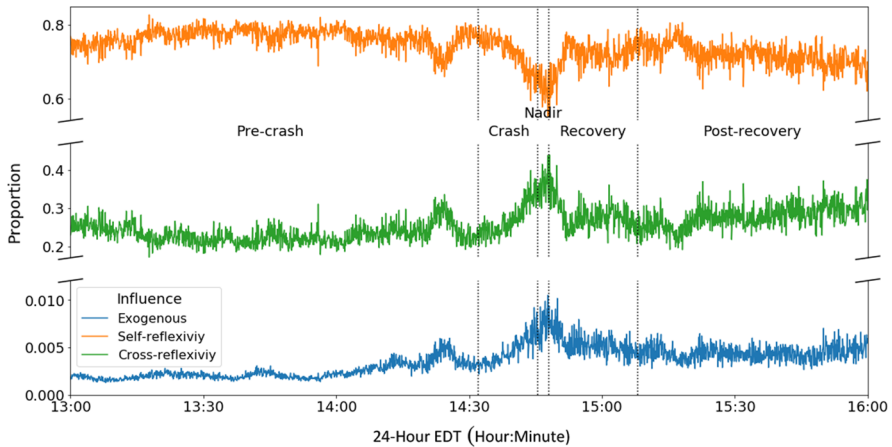


Fig. 5 Proportions of exogenous effect (blue), self-reflexivity (orange), and cross-reflexivity (green) from 13:00 to 16:00 (market close). The sum of the three influences at each time point equals one (color figure online)

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 while 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.

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 1% 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 that of the exogenous effect, which suggests that the latter effect may be dominated by stocks that are not part of the DJIA.

5.2 DJIA stocks

To determine the major influencers in the market, for each company we calculate the average out-influence, which is the impact of a particular stock on the other 29

stocks, in the pre-crash through post-recovery periods (see Table 3 for specific dates). As we show in Table 4, 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 4, the three strongest out-influence stocks over time are Bank of America (BAC), ExxonMobil (XOM), and JPMorgan Chase (JPM), and the three weakest are 3 M (MMM), Dupont (DD) and Travelers (TRV).

In a format similar to Tables 4 and 5, Table 6 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), 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 2).²⁷

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, recovery, and post-recovery periods. As shown in Table 7, 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.

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 8. 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 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.

²⁷ Matthews et al. (2021) report that in late 2020 the top executives of Chevron and Exxon Mobil discussed the possibility of the merger of their companies as a result of the financial stress to their firms brought about by the Covid 19 pandemic that began in early 2020. If such a merger occurs, it would reunite the new companies that were part of the Standard Oil monopoly prior to its break-up in 1911 by U.S. Supreme Court as a result of its finding that the monopoly violated the Sherman Antitrust Act of 1890.

Table 4 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

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 5 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
3 M	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

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 6 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

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 7 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 versus out	- 0.63 ($< .0001$)	- 0.37 (.0206)	- 0.67 ($< .0001$)	- 0.75 ($< .0001$)	- 0.53 (.0025)	- 0.64 (.0001)
In versus self	- 0.77 ($< .0001$)	- 0.60 (.0002)	- 0.59 (.0003)	- 0.74 ($< .0001$)	- 0.66 (.0001)	- 0.79 ($< .0001$)
Out versus self	0.58 (.0003)	0.35 (.0282)	0.29 (.0627)	0.72 ($< .0001$)	0.72 ($< .0001$)	0.72 ($< .0001$)

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 8 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)

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

5.3 DJIA communities

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 number of price-changing trades associated with these stocks. A useful descriptor of large-scale structures of a network is modularity.²⁸ Modularity quantifies the degree to which a network can be partitioned into communities, or, in our case, groups or clusters of stocks that may be related to one another in some fashion. Networks with high modularity have dense connections between nodes within same communities but sparse

²⁸ 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).




Fig. 6 Network modularity (blue) and average standardized price for all DJIA 30 stocks (orange) from 13:00 to 16:00 (market close) (color figure online)

connections between nodes contained in different communities.²⁹ Thus, modularity is different for different network partitions, but typically, in applications, modularity refers to the maximum modularity according to the best partition and we do as well. To find the best partitions and the corresponding modularity scores, we use the Python package `leidenalg` (<https://github.com/vtraag/leidenalg>) by Traag et al. (2019). For completely random networks, modularity will be close to zero, and the larger the modularity, the more fragmented is the network.

Similar to some of our earlier analyses, the modularity of our stock network is plotted in Fig. 6 along with the stock price series from the pre-crash period to the post-recovery period. As shown in Fig. 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 as indicated by their very low modularity value.

Numerous previous studies, e.g., King (1966), Cavaglia et al. (2000), and Fan et al. (2016), suggest the presence of an industry effect such that the prices of stocks in the same industry move together because they face the same production and demand issues. Other studies, e.g., O'Hara and Ye (2011), Menkveldt and Yueshen (2019), and Tivan et al. (2020), argue that stock markets may be fragmented, i.e., there are many markets that serve the same general purpose, but, at times, they may not be well connected to each other and, thus, have the potential to create current or latent liquidity problems.

We explore these two topics in more detail by examining the existence of industry sector and stock trading venue communities using modularity and the normalized

²⁹ 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):

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

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 $\delta(C_i, C_j) = 1$ if $C_i = C_j$ and $\delta(C_i, C_j) = 0$ otherwise.




Fig. 7 Normalized mutual information (NMI) statistics between identified stock communities and by industry sectors (blue) and the average standardized price for all DJIA 30 stocks (orange) from 13:00 to 16:00 (market close) (color figure online)

mutual information (NMI) statistic.³⁰ In our case, the NMI measures the similarity that exists between any two partitions of the same set of objects (i.e., the 30 DJIA stocks). If the NMI value is one, the partitions completely overlap each other, indicating that there is no difference between them. However, if the NMI value is zero, then the partitions do not overlap, signaling that there is no similarity between them. Newman and Girvan (2004) report that in their studies typical values range from 0.30 to 0.70, with values above 0.70 being quite rare. Accordingly, for each excitation matrix we cluster the stocks using the modularity method and compute the NMI statistic between the identified stock communities and the industry sectors. We make the same calculation for trading venues and stock trade reporting units.

The calculated NMI values between the nine industry sectors (see Table 2 for their names and symbols) and the corresponding identified communities are plotted in Fig. 7. A review of Fig. 7 indicates that the NMI does not have any significant trend and its value remains relatively small, i.e., about 0.30, throughout the time period under examination. This finding in conjunction with the modularity results depicted in Fig. 6 suggests that the network becomes more clustered during the crash, but these clusters do not appear to materially overlap with industry sectors. In other words, the impact of the crash is more likely to spread among the different industry sectors than within the industry sectors.

Another possible candidate for understanding the network modularity pattern is the trading venue. As mentioned on Sect. 4.1, the DJIA stocks are traded on 10 different trading venues, nine of which are stock exchanges. Each of the exchanges is responsible for accounting for all of the information corresponding to each trade. To be able to be traded on any of the exchanges, the stock must be listed on the exchange in question.

³⁰ Normalized mutual information (NMI) measures the similarity between any two partitions. Mathematically, for any two partitions U and V of the same objects, NMI is defined as: $NMI(U, V) = \frac{1}{C(U, V)} \sum_{i=1}^{|U|} \sum_{j=1}^{|V|} \frac{|U_i \cap V_j|}{N} \log \frac{N|U_i \cap V_j|}{|U_i||V_j|}$, where U_i, V_j are clusters U and V , respectively, and $|U_i|, |V_j|$ are their respective cluster sizes. N is the number objects in total, and $C(U, V)$ is a normalization factor to force the NMI value into $[0, 1]$. The intuition behind this definition is to measure the overlap between clusters in U and clusters in V . If none of the clusters in U overlap with any cluster in V , then $NMI = 0$, suggesting that the two partitions are very different. If, however, the two partitions are the same, then $NMI = 1$. (See Vinh et al. (2010) for additional details.) Because each stock is associated with only one industry, we partition the 30×30 excitation matrix for our calculations. In contrast, the stocks may be traded on 10 different venues, which equates to 300 stock-venue categories and a 300×300 excitation matrix.

Since the DJIA companies are large and well known, it is common for these companies to be listed on many exchanges. Although each listing must be purchased, most chief financial officers believe that additional costs are worth the potential additional liquidity. The tenth venue is the Nasdaq Trade Reporting Facility (NTRF), which is a partnership between Nasdaq and Financial Industry Regulatory Authority (FINRA). It is responsible for the collection of trade data for all of the internal trades, including crossing networks and dark pools operated by broker-dealer firms as well as trading desks that execute internal trades for large investment houses. In the 1980s, these ten trading venues were linked together under the auspices of the National Market System (NMS). To improve competition among these venues, the NMS creates a consolidated tape, thereby permitting all market participants the opportunity to view the transactions and quotes from each venue at the same time. This arrangement led O'Hara and Ye (2011) to suggest that the USA does not have multiple markets but rather has a single market that supplies its participants with multiple access points with each access point being one of the ten trading venues.

The NMI between the identified stock communities and the ten trading venues are plotted in Fig. 8. Unlike the NMI plot for industries, it contains patterns that are significant when compared to the lower benchmark of 0.30. In particular, 30 min before the crash the NMI increases and becomes significant around 14:15 but it then decreases and returns to insignificance a few minutes before the crash. At the beginning of the crash (14:32) it again significantly increases throughout the entire crash period. It then begins to decrease until 14:55 but it rises again until the end of the recovery period. Except for a very few seconds all the NMI values are significant. Throughout the post-recovery period, the NMI remains relatively steady and significant. Although not a perfect match, NMI pattern for the trade venue community segmentation is clearly reminiscent of the modularity pattern for the market as a whole (Fig. 6).

To probe more deeply into NMI fragmentation pattern, we examine the market shares of each of the 10 venues. The names of these trading (and reporting) venues are shown in Table 9, along with market share percentages for each of the five market periods individually and their total. Nasdaq led the venues' market share with approximately 46%. The New York Stock Exchange (NYSE) and BATS Global Markets (BATS) are ranked second and third. Together the top three venues account for slightly over 79% of the price-changing trades. On the low end of the market shares




Fig. 8 Normalized mutual information (NMI) statistics between identified stock communities defined and trade reporting market venues (blue) and the average standardized price for all DJIA 30 stocks (orange) from 13:00 to 16:00 (market close) (color figure online)

Table 9 Percentage trade volume of trade reporting venue for the 30 Dow Jones Industrial Average (DJIA) Stocks on May 6, 2010 by Economic Period

Trade reporting venue	Pre-crash	Crash	Nadir	Recovery	Post-recovery	Total
BATS Global Markets (BATS)	11.73	13.70	11.68	12.35	12.94	12.58
Boston Stock Exchange (BOST)	7.30	7.00	9.00	6.71	6.67	6.99
Cboe Global Markets (CBOE)	0.03	0.07	0.00	0.02	0.03	0.03
Chicago Stock Exchange (CHIC)	0.07	0.18	0.20	0.07	0.04	0.08
Cincinnati Stock Exchange (CINC)	0.55	0.10	1.38	1.54	0.09	0.10
International Securities Exchange (ISE)	1.09	2.09	3.75	2.42	1.73	1.87
Nasdaq Exchange (NQEX)	53.40	46.53	42.25	42.85	43.31	45.99
Nasdaq Trade Reporting Facility (NTRF)	3.44	1.60	0.68	1.12	1.77	1.96
New York Stock Exchange (NYSE)	14.85	18.71	21.18	20.60	22.54	19.74
Pacific Stock Exchange (PACF)	7.55	9.10	9.88	12.34	10.03	9.77
All Venues	100	100	100	100	100	100

The sum of each column (All Venues) may not equal 100 because of rounding

are the Cincinnati Stock Exchange (CINC), the Chicago Stock Exchange (CHIC), and CBOE Global Markets (CBOE), which together amount to 0.21%. The venue composition of the three market share categories varies little over the five economic periods. Although the differences in market share are relatively small, the Cincinnati Stock Exchange (CINC) increased its market share in the nadir and recovery periods and switched positions with the Nasdaq Trade Reporting Facility (NTRF), which dropped market share in the same two periods.

There are, however, notable differences in venue trading activity between adjacent market periods. For example, the top three venues (NQEX, NYSE, and BATS) accounted for nearly 80% of the trades in the pre-crash period, dropped to 79% during the crash, and continued to fall to 75% at the market's nadir for a market share drop of approximately five percentage points. Together their market share began to increase during the recovery period and continued to increase to approximately 79%. These large market share venues did not, however, operate in tandem. From the pre-crash period to the nadir periods NQEX market share dropped by 11.15 percentage points. The corresponding changes for NYSE and BATS are an increase of 6.33 percentage points and a decrease of 0.05 percentage points, respectively. As a result, almost five percentage points were picked up by the seven other trading venues, although up and down patterns are also exhibited by these smaller trading venues. These changes in

market share strongly suggest that traders are willing to move from one venue to another in an effort to find the best price for the size of their transactions.

Gomber et al. (2016) conclude their extensive survey on the nature of market fragmentation versus consolidation by suggesting that the economic welfare costs and benefits of stock markets mostly depend on the way they handle issues of price discovery and adequate liquidity. Although there are economies of scale that strongly favor consolidation, they maintain that because of the differences in trader behavior such as trading motives, order sizes as well as the need for quickness, it may be difficult to design a single market that can satisfy the needs of all traders. Recent research by Nicole et al. (2020), however, posit that the presence of heterogeneous agents may not be a reason for the existence of multiple markets. Instead, their work on market fragmentation using agent-based modeling leads them to suggest that traders' preference for multiple markets may be the result of their adaptive behaviors.

The industry and trading venue community results differ greatly. We suggest that this difference is the length of the crash as well as high-frequency trading. When the crash and its recovery are very short as is the case of the Flash Crash, a large majority of the trading is endogenous as a result of traders seeking the best deal regardless of venue in order to gain profit or minimize loss. In contrast, industry information does not change and, hence, is exogenous. For longer crash and recovery periods (months or even years), we expect there would be an industry effect since industry information would most likely be part of the price discovery process employed by the traders.

6 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.

In addition, we find that the US trading venues are usefully viewed as a tightly connected network that traders tend to use to their advantage. As O'Hara and Ye (2011) suggest, the venues compete with each other to attract volume and, hopefully, increase profits by offering lower transaction costs and faster execution speeds, although these advantages may also be associated with greater volatility in the short run, which may cause some traders, e.g., market-makers, to temporarily exit the market, possibly resulting in the loss of liquidity. Traders have the ability to access all of the trading venues and can quickly search for the best deal for them and their clients. For extremely large buy or sell orders, fragmentation allows the traders to split their orders among trading venues to try to lessen the impact of the transaction.

Our findings have implications for the performance of stocks in terms of risk and return, market microstructure design, and stock index composition. Turning first to

stock portfolio performance, modern risk management 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 the mean and variance of returns. Relying on this development, Sharpe (1964, 1994) creates a risk-return performance ratio (the Sharpe Ratio), which implicitly assumes that the relevant return distributions are Gaussian.³¹ This assumption, however, is often not true. Recognizing that return distributions are typically asymmetric and thick-tailed, 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.³² Applying either the Sharpe Ratio or the Sortino Ratio to high-frequency data, however, is problematic because of the Epps effect (see fn. 21). Ait-Sahalia and Hurd (2016), however, have developed 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. An 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 an environment that recognizes the possibility of contagion, a very few or even a single stock could trigger a crash, flash, or otherwise.

Second, as we mention in Sect. 4.2, the Securities and Exchange Commission (SEC) adopted the limit up-limit down (LULD) circuit breaker rule to mitigate the possibility of a market crash. The idea was to try to constrain the volatility of a specific stock in the hope that this would reduce the possibility of a market crash induced by this stock. Our results, however, show that the contagion (or influence) between stocks is not symmetric, i.e., in-influence and out-influence are not identical and are negatively related. This finding suggests that a more refined microstructure arrangement might have different LU and LD limits with the magnitude of these limits determined by the magnitude of the two respective influences. If having different limits is not administratively practical, an alternative to consider is to have symmetric limits with the values being determined by the most out- and in-influencing stocks.

Third, the composition of stock indexes is routinely changed to reflect structural changes in the economy. This is typically done by adding some stocks of strong companies in growing industries and eliminating stocks in companies that have become less strong in declining industries. Our results, however, indicate that industries may not be the important factor in the construction of a high-frequency stock index. Instead, what is important is the influence effect that stocks have on each other and themselves. Moreover, although current daily indexes are designed to handle the effects of stock price differences when their composition is being modified, there does not appear to be

³¹ 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).

³² 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 (LPMCAMP). 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.

a rigorous quantitative assessment of the possible cross-influence differences between the stocks being added and those being deleted. Not accounting for these differences in some way (even qualitative) may render the index behavior before and after the change as not comparable with respect to its dynamical behavior.

In sum, because of the high level of technology that is being used in the major stock and similar markets worldwide and the likelihood of continued high-frequency and algorithmic 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. Work should also be done on the implications of stock cross-influences and on the rules of trade and their impact on prices and liquidity restrictions, as well as the design of stock indexes that are not only used to monitor the overall performance of the market, but also as an input to various asset pricing models. Additionally, efforts should be directed toward the ways that market fragmentation affects the overall self-, in- and out-influences of individual stocks. Finally, the Hawkes model quantitatively describes the behavior of the stock market during the 2010 Flash Crash. Nevertheless, more intricate versions of the process may be needed to better understand a market crash that may last weeks, months, or even years. One possibility is to add higher-order nonlinear terms or jump processes to the model, e.g., stochastic cusp catastrophe models.

We leave these and other similar interesting topics to future research. Nevertheless, we note that Kuhlmann (2014) believes that social complexity in economics/finance is not often explicitly included in its models or theory because its impact through institutional arrangements and interaction are not sufficiently recognized. We concur with his view and urge that financial research, especially that involving market microstructure research, should be guided not only by statistical models and measures, but also by the notion of the complex nature of human behavior and how this behavior is reflected in the research scheme.

Acknowledgements We thank 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 other participants of the December 2019 Paris Financial Management Conference as well as the attendees of a February 2020 Stockholm Business School seminar for useful comments on an earlier version of this paper. 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.

References

- Ackert LF (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
- Akgiray V, Geoffrey Booth G (1988) Mixed diffusion-jump process modeling of exchange rate movements. *Rev Econ Stat* 70(4):631–637
- Aldrich EM, Grundfest JA, Laughlin G (2017) The flash crash: a new deconstruction. <https://doi.org/10.2139/ssrn.2721922>
- Andersen TG, Benzoni L, Lund J (2002) An empirical investigation of continuous-time equity return models. *J Finance* 57(3):1239–1284
- Anderson PW, Arroxxw K, Pines D (eds) (1988) *The economy as an evolving system*. Addison Wesley, Reading

- Arthur WB (1999) Complexity and the economy. *Science* 284:107–109
- Arthur WB, Durlauf SN, Lane DA (eds) (1997) *The economy as an evolving system II*. Addison Wesley, Reading
- Aït-Sahalia Y, Hurd TR (2016) Portfolio choice in markets with contagion. *J Financ Economet* 14(1):1–28
- Aït-Sahalia Y, Cacho-Diaz J, Laeven RJA (2015) Modeling financial contagion using mutually exciting jump processes. *J Financ Econ* 117(3):585–606
- Bachelier LJ-BA (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
- Bacry E, Mastromatteo I, Muzy J-F (2015) Hawkes processes in finance. *Mark Microstruct Liq* 1(1):1550005
- Bawa VS, Lindenberg EB (1977) Capital market equilibrium in a mean-lower partial moment framework. *J Financ Econ* 5(2):189–200
- Bayer P, Mangun K, Roberts JW (2021) Speculative fever: Investor contagion and the housing bubble. *Am Econ Rev* 111(2):609–651
- Bouchaud J-P (2010) The endogenous dynamics of markets: price impact and feedback loops, pp 1–15. [arXiv:1009.2928v1](https://arxiv.org/abs/1009.2928v1)
- Bracha HS, Ralston TC, Matsukawa JM, Williams AE, Bracha AS (2004) Does “fight or flight” need updating. *Psychomatics* 5:448–449
- Burzyn L, Federer F, Ferman B, Yuchtman N (2014) Understanding the mechanisms underlying peer effects: evidence from a field study on financial decisions. *Econometrica* 82(4):1273–1301
- Cai N, Kou SG (2011) Option pricing under a mixed-exponential jump diffusion model. *Manag Sci* 57(11):2067–2081
- Cannon WB (1929) *Bodily changes in pain, hunger, fear and rage: an account of recent research into the functional of emotional excitement*, 2nd edn. Appleton-Century-Crofts, New York
- Cavaglia S, Brightman C, Aked M (2000) The increasing importance of industry factors. *Financ Anal J* Sept/Oct:41–54
- Christensen CM, Raynor ME (2003) *The innovator's solution: creating and sustaining successful growth*. Harvard Business School Publishing, Boston
- Commodity Futures Trading Commission-Securities Exchange Commission (CFTC-SEC) (2010a) Preliminary Findings Regarding the Market Events of May 6. Staff report. <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. <https://www.sec.gov/news/studies/2010/marketevents-report.pdf>
- DeMarzo PM, Kaniel R, Kremer I (2008) Relative wealth concerns and financial bubbles. *Rev Financ Stud* 21(1):19–50
- Easley D, Kleinberg J (2010) *Networks, crowds and markets*. Cambridge University Press, Cambridge
- Easley D, Marcos M, de Prado L, O'Hara M (2012) Flow toxicity and liquidity in a high frequency world. *Rev Financ Stud* 25(5):1457–1493
- Embrechts P, Liniger T, Lin L (2011) Multivariate Hawkes processes: an application to financial data. *J Appl Probab* 48:A:367–378
- Epps TW (1979) Comovements of stock prices in the very short run. *J Am Stat Assoc* 74(366):291–298
- Etesami J, Kiyavash N, Zhang K, Singhal K (2016) Learning network of multivariate Hawkes Processes: a time series approach. In: *Proceedings of the thirty-second conference on uncertainty in artificial intelligence* (AUAI Press, Arlington, VA), pp 162–171
- Fama EF (1965) Behavior of stock market prices. *J Bus* 38(1):34–105
- Fama EF (1970) Efficient capital markets: a review of theory and empirical work. *J Finance* 25(2):383–417
- Fama EF (1998) Market efficiency, long-term returns, and behavioral finance. *J Financ Econ* 49:283–306
- Fan J, Furger A, Xiu D (2016) Incorporating global industrial classification into portfolio allocation: a simple factor-based large covariance matrix estimator with high frequency data. *J Bus Econ Stat* 34(4):489–503
- Filimonov V, Sornette D (2012) Quantifying reflexivity in financial markets: toward a prediction of flash crashes. *Phys Rev E* 85:5:056108-2-056108–9
- Frank KA, Yun-Jia Lo G, Booth G, Kallunki J-P (2019) The market for socially embedded trading. *Ration Soc* 31(2):152–181
- Frankfurter GM, McGowan E (2001) Anomalies in finance: what they are and what are they good for? *Int Rev Financ Anal* 10(2):407–422

- Glasserman P, Peyton Young H (2016) Contagion in financial networks. *J Econ Lit* 54(3):779–831
- Goldstein DS, Kopin IJ (2007) Evolution of the concepts of stress. *Stress Int J Biol Stress* 10(2):109–120
- Gomber P, Sagade S, Theissen E, Weber MC, Westheide C (2017) Competition between equity markets: a review of the consolidation versus fragmentation debate. *J Econ Surv* 31(3):792–814
- Granger CWJ (1969) Investigating causal relations by econometric methods and cross-spectral methods. *Econometrica* 37(3):424–438
- Granger CWJ (2004) Time series analysis, cointegration, and applications. *Am Econ Rev* 94(3):421–425
- Grossmass L (2014) Obtaining and predicting bounds of realized correlations. *Swiss J Econ Stat* 150(3):191–226
- Harlow WV, Rao RKS (1989) Asset pricing in a generalized lower partial moment framework. *J Financ Quant Anal* 24(3):285–311
- Hawkes AG (2018) Hawkes processes and their applications to finance: a review. *Quant Finance* 18(2):193–196
- Hogan WW, Warren JM (1974) Toward the development of an equilibrium capital market based on semi-variance. *J Financ Quant Anal* 9(1):1–11
- Kay R (2009) Pragmatic network latency engineering, working paper. cPacket Networks, Inc., Mountain View, CA
- King BF (1966) Markets and industry behaviors in stock price behavior. *J Bus* 39(1):139–190
- Kirilenko A, Kyle AS, Samadi M, Tuzun T (2017) High-frequency trading in an electronic market. *J Financ* 72(3):967–998
- Kuhlmann M (2014) Explaining financial markets in terms of complex systems. *Philos Sci* 81(5):1117–1130
- Kuhn TS (1970) *The structure of scientific revolutions*, 2nd edn. University of Chicago Press, Chicago
- Ladyman J, Lambert J, Wiesner K (2012) What is a complex system? *Eur J Philos Sci* 3(1):33–67
- Leicht EA, Newman MEJ (2008) Community structure in directed networks. *Phys Rev Lett* 100(11):118703–118703–4
- Lin M, Lucas HC, Shmueli G (2013) Too big to fail: Large samples and the p-value problem. *Inf Syst Res* 24(4):906–917
- Linderman SW, Adams RP (2014) Discovering latent network structure in point process data. In: 31st international conference on machine learning (JMLR: W&CP: Beijing, China), vol 32, pp 1413–1421
- Linderman SW, Adams RP (2015) Scalable Bayesian inference for excitatory point process networks, pp 1–16. [arXiv:1507.03228v1](https://arxiv.org/abs/1507.03228v1)
- Liniger T (2009) Multivariate Hawkes processes, doctoral thesis. Department of Mathematics, ETH Zürich, Switzerland
- Lintner J (1965) The value of risk assets and the selection of risky investments in stock portfolios and capital budgeting. *Rev Econ Stat* 47(1):13–37
- Liu Q, Tse Y, Zheng K (2021) The impact of trading biases on market liquidity under different volatility levels: Evidence from the Chinese commodity futures markets. *Financ Rev*
- Lo AW (2002) Bubble, rubble, finance in trouble? *J Psychol Financ Mark* 3:76–86
- Lo AW (2004) The adaptive market hypothesis: market efficiency from an evolutionary perspective. *J Portf Manag* 30(2):15–29
- Lo AW (2017) *Adaptive markets: financial evolution at the speed of thought*. Princeton University Press, Princeton
- Lo AW, Craig MacKinlay A (1999) *A non-random walk down wall street*. Princeton University Press, Princeton
- Lo AW, Repin DV (2002) The psychophysiology of real-time financial risk processing. *J Cogn Neurosci* 14:323–339
- Lo AW (ed) (1997) *Market efficiency: stock market behavior in theory and practice*, vols 1 and 2. Edward Elgar Publishing Company, Cheltenham
- Madhavan A (2012) Exchange-traded funds, market structure, and the flash crash. *Financ Anal J* 68(4):20–35
- Malkiel BG (1973) *A random walk down wall street*. W.W. Norton & Company, New York
- Mandelbrot BB (1963) The variation of certain speculative prices. *J Bus* 9:53–66
- Mandelbrot BB (1982) *The fractal geometry of nature*. W.H. Freeman and Co., San Francisco
- Mandelbrot BB (2001) Scaling in financial prices: I. Tails and Dependence. *Quant Finance* 1:113–123
- Mandelbrot BB, Hudson RL (2004) *The (mis)behavior of markets: a fractal view of risk, Ruin and Reward*. Basic Books, New York
- Markowitz HM (1952) Portfolio selection. *J Finance* 7(1):77–91
- Markowitz HM (1959) *Portfolio selection*. Wiley, New York

- Matthews CM, Glazer E, Lombardo C (2021) CEO's explored chevron-Exxon Merger in '20. Wall Street J February 1: A1 and A2
- Menkveld AJ, Yueshen BZ (2019) The Flash Crash: a cautionary tale about highly fragmented markets. *Manag Sci* 65(10):4470–4488
- Merton RC (1976) Option pricing when underlying stock returns are discontinuous. *J Financ Econ* 3(1–2):125–144
- Mossin J (1966) Equilibrium in a capital asset market. *Econometrica* 34(4):768–783
- Müller UA, Dacorogna MM, Olsen RB, Pictet OV, Schwartz M, Morgengegg C (1990) Statistical study of foreign exchange, empirical evidence of price change scaling law, and intraday analysis. *J Bank Finance* 14(6):1189–1208
- Müller UA, Dacorogna MM, Davé RD, Pictet OV, Schwartz M, Olsen RB, Ward JR (1995) Fractals and intrinsic time—a challenge to econometricians, research report- UAM.1993-08-16, June 28, Olsen & Associates, Zurich, Switzerland
- Newman MEJ, Girvan M (2004) Finding and evaluating community structure in networks. *Phys Rev E* 69(2):026113-1-026113–15
- Nicole R, Aloric A, Sollich P (2020) Fragmentation in trader preferences among *multiple* markets: market coexistence versus single market dominance, pp 1–25. [arXiv:2012.04103v1](https://arxiv.org/abs/2012.04103v1)
- O'Hara M (2015) High frequency market microstructure. *J Financ Econ* 116(2):257–270
- O'Hara M, Ye M (2011) Is market fragmentation harming market quality. *J Financ Econ* 100(3):459–474
- Park A, Sabourian H (2011) Herding and contrarian behavior in financial markets. *Econometrica* 79(4):973–1026
- Peters EE (1991) Chaos and order in capital markets: a new view of cycles, prices and markets. Wiley, New York
- Peters EE (1994) Fractal market analysis: applying chaos theory to investment and economics. Wiley, New York
- Popper KR (1959) Logic of Scientific Discovery, Routledge Classics edition, 2002 (Routledge, London, U.K. and New York, N.Y.). First published: 1935, Logik der Forschung (Verlag von Julius Springer, Vienna, Austria)
- Quinn CJ, Kiyavash N, Coleman TP (2011) Equivalence between minimal generative model and directed information graphs. In: IEEE international symposium on information theory proceedings, information theory society, St. Petersburg, Russia, pp 293–297
- Rizoiu M-A, Lee Y, Mishra S, Xie L (2017) Hawkes processes for events in social media. In: Chang S-F (ed) Frontiers in multimedia research. Association for Computing Machinery and Morgan & Claypool, New York, pp 191–218
- Roy AD (1952) Safety first and the holding of assets. *Econometrica* 20(3):431–449
- Saariluoma P (1995) Chess player's thinking: a cognitive psychological approach. Routledge, New York
- Santilli P (2020) The Dow's evolution from the 1800s to the 21st century. *Wall Street J* 29–30:B2
- Satchell SE (2001) Lower-partial moment capital asset pricing models: a re-examination. In: Sortino FS, Satchell SE (eds) Managing downside risk in financial markets: theory, practice and implementation. Butterworth-Hienemann Finance, Oxford, pp 156–168
- Sharpe WF (1964) Capital asset prices: a theory of market equilibrium under conditions of risk. *J Finance* 19(3):425–442
- Sharpe WF (1994) The Sharpe ratio. *J Portf Manag* 21(1):49–58
- Shenker OR (1994) Fractal geometry is not the geometry of nature. *Stud Hist Philos Sci A* 25 6:967–981
- Shiller RJ (2014) Speculative asset prices. *Am Econ Rev* 104(6):1486–1517
- Sifat IM, Mohamad A (2018) Circuit breakers as market stability levers: a survey of research, praxis, and challenges. *Int J Finance Econ* 56:1–40
- Simon HA (1979) Rational decision making in business organizations. *Am Econ Rev* 69(4):493–513
- Smit A (1776) The Wealth of Nations, Vintage Classics 2020 edition (Vintage Classics, London, U.K.). Originally published as: An Inquiry into the Nature and Causes [sic] of the Wealth of Nations (W. Strahan and T. Cadell, London, U.K.).
- Sohn H, Sornette D (2017) Bubbles as violations of efficient time-scales, working paper no. 65, Institute for New Economic Thinking, New York, NY
- Sornette D (2003) Why stock markets crash: critical events in complex financial systems. Princeton Science Library edition, 2017, with new forward by the author. Princeton University Press, Princeton
- Sornette D (2009) Dragon-kings, black swans and the prediction of crises. *Int J Terraspace Sci Eng* 2(1):1–18

- Sortino FA (2001) From alpha to omega. In: Sortino FA, Stephan E (eds) *Managing downside risk in financial markets: theory practice and implementation*. Butterworth-Hienemann Finance, Oxford, pp 3–25
- Taleb NN (2007) *The black swan: the impact of the highly improbable*. Random House, New York
- Tivan BF, Dewhurst DR, Van Oort CM, Ring JH, Gray TJ, Tivan BF, Koehler MTK, McMahon MT, Slater DM, Veneman JG, Danforth CM (2020) Fragmentation and inefficiencies in the US equity markets: Evidence from the Dow. *PLoS ONE* 15:e0226968. <https://doi.org/10.1371/journal.pone.0226968>
- Traag VA, Waltman L, van Eck NJ (2019) From Louvain to Leiden: Guaranteeing well-connected communities. *Sci Rep* 9, Article number: 5233
- Vinh NX, Epps J, Bailey J (2010) Information theoretic measures for clusterings comparison: variants, properties, normalization and correction for chance. *J Mach Learn Res* 11:2837–2854
- Weron A, Weron R (2000) Fractal market hypothesis law and two power laws. *Chaos Solitons Fractals* 11:289–296
- Zadeh AH, Sharda R (2015) Hawkes point processes for social media analytics. In: Iyer LS, Power DJ (eds) *Reshaping society through analytics, collaboration, and decision support*, *annals of information systems* 18. Springer, Switzerland, pp 51–66
- Ziliak ST, McCloskey DN (2008) *The cult of statistical significance: how the standard error cost us jobs, justice, and lives*. The University of Michigan Press, Ann Arbor
- Zweig J (2021) The know-nothings running with the bulls. *Wall Street J* 27–28:B1 & B7

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.