High-frequency Trading, Algorithmic Finance, and the Flash Crash

Reflections on Eventalization

Borch, Christian

Document Version
Accepted author manuscript

Published in:
Economy and Society

DOI:
[10.1080/03085147.2016.1263034](https://doi.org/10.1080/03085147.2016.1263034)

Publication date:
2016

Citation for published version (APA):
Borch, C. (2016). High-frequency Trading, Algorithmic Finance, and the Flash Crash: Reflections on Eventalization. *Economy and Society*, 45(3-4), 350-378. <https://doi.org/10.1080/03085147.2016.1263034>

[Link to publication in CBS Research Portal](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us (research.lib@cbs.dk) providing details, and we will remove access to the work immediately and investigate your claim.

Download date: 26. Apr. 2023

High-frequency Trading, Algorithmic Finance, and the Flash Crash: Reflections on Eventalization

Christian Borch

Journal article (Post print version)

Cite: High-frequency Trading, Algorithmic Finance, and the Flash Crash: Reflections on Eventalization. / Borch, Christian. In: *Economy and Society*, Vol. 45, No. 3-4, 2016, p. 350-378.

This is an Accepted Manuscript of an article published by Taylor & Francis in *Economy and Society* on 18 January 2016, available online:

<http://dx.doi.org/10.1080/03085147.2016.1263034>

Uploaded to Research@CBS: April 2017

High-frequency trading, algorithmic finance and the Flash Crash: Reflections on eventalization

Christian Borch

Christian Borch, Department of Management, Politics and Philosophy,
Copenhagen Business School, Denmark. Email: cbo.mpp@cbs.dk

ORCID 0000-0001-8217-5880

Abstract

The Flash Crash of 6 May 2010 has an interesting status in discussions of high-frequency trading, i.e. fully automated, superfast computerized trading: it is invoked both as an important illustration of how this field of algorithmic trading operates and, more often, as an example of how fully automated trading algorithms are prone to run amok in unanticipated frenzy. In this article, I discuss how and why the Flash Crash is being invoked as a significant event in debates about high-frequency trading and algo-financial markets. I analyze the mediatization of the event, as well as the variety of eventalizations of the Flash Crash, i.e. the different ways in which the Flash Crash is being mobilized as an illustrative event. I further critically discuss the impact often associated

with the Flash Crash – and on that basis, inquire into why the event nonetheless attracts so much attention. I suggest that a key reason why the Flash Crash is widely discussed is that eventalizations of 6 May 2010 evoke familiar tropes about the fear of technology and the fear of herding. Finally, and given their emphasis on herding, I argue that the Flash Crash eventalizations may contribute to economic sociology discussions about resonance in quantitative finance.

Keywords: algorithmic trading; Émile Durkheim; events; financial markets; Flash Crash; herding; high-frequency trading; resonance; Gabriel Tarde

Introduction

On Thursday 6 May 2010, the US financial markets experienced one of the biggest intraday declines in the history of the Dow Jones Industrial Average – a drop of 998.5 points (or more than 9 per cent), equivalent to around one trillion dollars. The greater part of the decline took place over just four-and-a-half minutes, but the markets, after having been temporarily suspended, recovered almost as quickly. The trading day began with a dip due to the sovereign debt crisis in Europe, but the subsequent rapid market decline was triggered – at least according to the official report subsequently published jointly by the two US regulatory authorities: the Commodity Futures Trading Commission (CFTC) and the Securities and Exchange

Commission (SEC) – when a large fundamental trader executed a sizeable sell order that was initially partly absorbed by high-frequency traders, i.e. high-speed algorithmic traders, who then ‘aggressively sold’ their position in a manner that rapidly drove down prices (CFTC-SEC, 2010, p. 3). What ensued was a situation of extreme market volatility, in which some stocks were traded at prices far from their normal value: at ‘irrational prices as low as one penny or as high as \$100,000’ (2010, p. 5). A large proportion of these trades were later cancelled by the exchanges, due to being ‘clearly erroneous’ (2010, p. 6).

The event was soon dubbed ‘the Flash Crash’ due to the rapid decline and rebound of prices. What caused the Flash Crash? What characterized its dynamics? And what might be its broader implications for financial markets? These questions have been the subject of much subsequent inquiry, controversy and discussion by the public, media and scholars, and a range of explanations and perspectives have been offered. For example, some observers suggest that the Flash Crash laid bare deeper structural problems with the market microstructure (Schapiro, 2010), whereas others draw parallels with other famous cases from the domain of algorithmic finance, where orders are executed by fully automated computer algorithms without any direct human intervention (e.g. Kenett *et al.*, 2013; Kirilenko and Lo, 2013). Cases cited include Knight Capital, a major player in the US financial markets, which in August 2012 lost around USD 460 million in just 45 minutes due to an error in their trading algorithm (e.g.

Seyfert, 2016); and a so-called ‘hack crash’ in 2013, in which a fake Twitter update led to rapid market activity (Karppi and Crawford, 2016). Although the Flash Crash occurred back in 2010, it continues to attract attention (e.g. Lange, Lenglet and Seyfert, 2016; Seyfert, 2016; Thompson, 2016). For example, in 2015 the event gained renewed media attention due to the arrest in April that year of the British trader Navinder Singh Sarao, who was charged with fraud and market manipulation. Among other things, he was accused of playing a critical role in the 2010 Flash Crash – allegedly bringing down US markets from his parents’ home on the outskirts of London. In November 2016, then, Sarao pleaded guilty to having spoofed markets – but whether his market-manipulating behaviour could indeed produce an event such as the Flash Crash remains contested.

The Flash Crash is no doubt interesting due to its inescapably spectacular nature. In fact, it is not uncommon for observers to refer to the Flash Crash primarily to create dramatic suspense. In such accounts, the Flash Crash is often flagged at the outset, only to be subsequently disregarded. However, as I will argue in this article, the Flash Crash is also interesting because it is often invoked by commentators (journalistic and academic alike) as a window into the workings and potential effects of present-day algorithmic finance in general, and high-frequency trading in particular. High-frequency trading is a subset of algorithmic finance in which orders are executed by fully automated algorithms in fractions of a second. However, while the Flash Crash serves as a window into this form

of trading and its implications, observers tend to see rather different things through that window. One reason for this is that there is considerable confusion as to what exactly constitutes the Flash Crash. There are, in other words, conflicting accounts about its ‘eventness’.

This is not specific to the Flash Crash. As Robin Wagner-Pacifici argues, events are often ‘intrinsically restless’ (2010, p. 1356). In other words, a lot of effort may be required to put such events to rest. This relates to a more fundamental problem examined by Wagner-Pacifici: ‘what makes an event an “event”?’ (2010, p. 1358). She argues compellingly for answering this question on the basis of a political semiosis that attends to three key features (performative, demonstrative and representational) that collectively produce an event. The focus of the present article is slightly different. It revolves around how and why certain incidents lend themselves to eventalization. The issue at stake might be illustrated by comparing the event under consideration, the Flash Crash, with the central event in Wagner-Pacifici’s analysis: the terrorist attacks of September 11 2001.

Wagner-Pacifici details the initial confusion about the September 11 events, which were first described as ‘accidents’, then as ‘incidents’, and then as ‘terrorist attacks’ (2010, p. 1352). The central point here is that, despite the initial confusion, an event narration eventually took shape that, while multidimensional and contested, was founded on the real experience of physical destruction and the deaths of nearly 3,000 people. In other words, the eventalization of the September 11 attacks had a real experien-

tial base (although some of the matter-of-factness relating to the eventalization of the attacks would soon be questioned, e.g. by conspiracy theorists). The Flash Crash, I will argue, is a rather different type of event.

First, while a massive market drop was certainly registered on 6 May 2010,¹ the spatial and temporal delimitations of the event are contested, giving rise to rather different ‘eventalizations’ of the Flash Crash. Second, and relatedly, while the effects of the September 11 attacks can be reasonably (though not comprehensively) quantified in terms of the number of victims and the extent of the physical destruction, the effects of the Flash Crash are rather less clear. Indeed, I will argue, much of the interest in this latter event seems to be founded not in its *actual* but rather in its *potential* effects. As I will show, the real economic effects of the Flash Crash appear rather limited and much less grave than those of previous ‘analog’ crashes, such as the 1987 crash. While prices declined rapidly on 6 May 2010, they also swiftly recovered. And so, according to some eventalizations of the Flash Crash, the most damaging aspects of the event may relate less to actual losses and more to consequences we have yet to fully experience: returning to the earlier metaphor, its real significance might be that it provides a window into how markets dominated by automated, superfast computer algorithms, operating (largely) beyond human control, might be prone to running amok in devastating ways.

In this article, I discuss a series of dominant eventalizations of the Flash Crash, i.e. important ways in which the crash has been delineated in

space and time, how it has been perceived, the main processes attributed to it, etc. This will include some of the journalistic coverage of the event, as well as a range of academic commentary from different fields. I also subject the event's alleged impact to critical scrutiny. I evaluate the experimental impact (the Flash Crash's effect on market participants), the real economic impact (the actual economic effects) and the potential systemic impact (the systemic risk involved in algorithmic trading, as illustrated in the crash), and subsequently argue that each impact is contestable. This raises the question of why the Flash Crash continues to attract so much attention. However, it also raises a more fundamental question about why certain occurrences lend themselves to eventalization. Why has the Flash Crash invited so much study as an important event? The answer, I suggest, lies in how it echoes wider, well-established tropes. More specifically, the Flash Crash is recognizable: it seemingly instantiates older concerns regarding technology and crowd/herding behaviour. These types of tropes are mobilized in the eventalizations of the Flash Crash and contribute to constituting it as something easily recognized as significant.

In other words, my assertion is that the Flash Crash was eventalizable not simply because it appeared dramatic, but equally because it conjured up a broader event register. So, similar to Wagner-Pacifici's observation that eventalization often involves mimetic features, which 'take their forms from templates developed in the past and brought to bear on [the] emergent event' (2010, p. 1362), I suggest that the tropes about the fear of technolo-

gy and of herding constitute the central template for the interest devoted to the Flash Crash. This attempt to situate the discourse on the Flash Crash in a broader cultural context aligns the article with work on how present-day derivate capitalism is embedded in broader socio-cultural logics (e.g. Arvidsson, 2016; Lee and Martin, 2016).

As is hopefully clear from this, the aim of the article is not to provide novel insights into what actually happened on 6 May 2010, nor to shed new light on what constitutes an event (for which reason I leave aside work on this matter in, e.g. the Whitehead and Deleuze traditions). Rather, in line with Janet Roitman's analysis of how the subprime mortgage crisis of 2007–8 prompted a series of crisis narratives that treated 'crisis' itself as a blind spot and rarely grasped the ways in which these narratives 'allow[ed] certain questions to be asked while others [were] foreclosed' (2014, p. 94), so I strive in this article to demonstrate how, in the various eventalizations of the Flash Crash, the event's importance operates as a blind spot that is rarely critically examined.

However, I also wish to discuss the event's importance in the light of debates in economic sociology. So, after having examined how various observers eventalize the Flash Crash, i.e. after treating discussions of the event as a matter of what Niklas Luhmann (2012) called 'second-order observation', I will push the analysis forward by asking what economic sociology might learn from the event if we assume that there is some substance to the eventalizations and how they portray contemporary algorithm-

mic finance. I shall discuss this with particular emphasis on discussions of dissonance, resonance, cognitive interdependence and market sociality as they pertain to quantitative finance (Beunza and Stark, 2012).

The majority of the article's analysis is based on written sources, especially journalistic accounts, government reports, and academic papers. However, the discussion of the impact of the Flash Crash also takes into account other sources, such as interviews I have conducted with US market participants involved in various forms of algorithmic finance, including high-frequency trading. The majority of these interviews (c. 30 in total) were conducted in April and May 2016 with CEOs, CTOs, traders, exchange officials, etc. from especially, but not exclusively, New York, Chicago and San Francisco.

The article comprises five sections. In the first, I analyze the mediation of the Flash Crash – both how the name of the event was attached to it almost immediately, and how the event operates as a turning point in the American journalist Michael Lewis's critical examination of high-frequency trading. The next section discusses how the Flash Crash has been eventalized by scholarly commentators, i.e. how the trading activity on 6 May 2010 has been discursively portrayed as a significant event. Two main types of event signification can be differentiated, which testifies to what Michel Foucault referred to as 'the tactical polyvalence of discourse' (1990, p. 100). Thus, while some commentators emphasize the allegedly exceptional status of the Flash Crash event, others argue that it signifies a

common occurrence in present-day financial markets. The third section critically discusses the types of impact usually attributed to the Flash Crash. The fourth section analyzes the two central tropes that are mobilized in eventalizations of the Flash Crash. The fifth section then relates the eventalizations of the Flash Crash to discussions of dissonance and resonance in quantitative finance. A brief conclusion rounds off the article.

The mediatization of the Flash Crash

As mentioned, Wagner-Pacifici argues that events are ‘intrinsically restless’ (2010, p. 1356). This restlessness entails that controversy is integral to the way in which events are shaped. There are different ways of delineating events (including their beginnings and endings, the main actors and processes involved, etc.), and event contestation is likely to revolve around such delineation and framing. Importantly, the framing of events can have performative effects. As Wagner-Pacifici notes, ‘the business of event framing is part and parcel of the continuing effect flow of events’ (2010, p. 1354). Events can obviously be framed by a variety of sources, including media coverage, government reports and academic commentary. As was the case with September 11, as analyzed by Wagner-Pacifici, the media have played a significant role in constituting the Flash Crash as an event. Mirroring the immense speed of the trading reportedly involved in the Flash Crash, the mediatization of the collapse and recovery of prices on 6 May 2010 was almost instantaneous. The most dramatic market drop

took place between 2:41 pm and 2:45:27 pm, at which point trading was suspended for five seconds. When trading resumed, ‘broad market indices recovered’, but individual securities still saw considerable price fluctuations. At approximately 3:00 pm, ‘prices of most individual securities significantly recovered and trading resumed in a more orderly fashion’ (CFTC-SEC, 2010, p. 9).

The media coverage was immediate. At 2:58 pm, CNBC published an online article describing how the ‘selloff is feeding on itself, bringing down almost all stocks’ (Melloy, 2010). The *Wall Street Journal* website was similarly quick to report on the unfolding events on 6 May. The first posting in which it not just reported, but reflected on what had happened was published as early as 3:22 pm, with follow-up commentaries at 3:46 pm and 4:00 pm. These initial postings (all of which were published on the so-called MarketBeat blog, which comments on market developments throughout the day) were characterized by confusion. No explanations were offered; there were only questions being posed, including whether ‘computerized trading [was] to blame’; if the ‘market swoon was caused by technical factors’, such as a ‘technical glitch’; or whether there was ‘a fat-finger trade’ behind it all (Phillips, 2010a; 2010b; 2010c).² Then at 4:15 pm, Michael Corkery, also writing for the *Wall Street Journal*, published a post entitled ‘The Stock Market’s Flash Crash: How to Destroy \$1 Billion in 60 Minutes’ (2010), apparently the first *WSJ* journalist to use the term ‘Flash Crash’. Within an hour, the term was gaining traction. At 5:04 pm,

the MarketBeat blog published a post that described how the notion of the ‘Flash Crash’ was already achieving hegemony:

The ‘Flash Crash.’ Here at MarketBeat HQ that seems to be the moniker we’re hearing for what happened during Thursday afternoon. [...] The usage of ‘flash’ echoes some of [*sic*] controversy over the component of electronic trading known as ‘flash trading’. (Phillips, 2010d; see also the 5:47 pm post by Grocer, 2010)

It therefore took just two hours for the term ‘Flash Crash’ to crystallize as the central point of reference. The term’s catchiness is arguably part of the reason why the event has attracted so much attention, but beyond that, its narrative quality consists in equally describing the swiftness of the market drop and recovery, *and* the immediate suspicion that high-frequency traders were profoundly implicated in the event.

Journalistic accounts of the Flash Crash abound in the steady stream of media articles about the event, as well as in the coverage of the subsequent CFTC-SEC report, the Sarao case, etc. The journalistic framing of the event as revealing the inner workings of financial markets also plays an important role in Michael Lewis’s bestselling book *Flash Boys* (2014), the publication of which was followed by considerable controversy, because it alleged that the markets dominated by high-frequency trading are rigged. However, that debate is not of central importance here. Instead, I am inter-

ested in how the Flash Crash event is portrayed at a critical juncture in Lewis's book (the title of which obviously alludes to the Flash Crash and flash – i.e. high-speed – trading).

The main protagonist of *Flash Boys* is Brad Katsuyama, a trader who, while working for the Royal Bank of Canada, began to notice that the prices he received upon completing a trade were slightly worse than those he saw on his screen before submitting the order. Something was happening between sending his orders and them being filled, something that was working consistently to his, and hence to his clients', disadvantage. Much of *Flash Boys* narrates Katsuyama's attempt to understand what was happening in this split second. He eventually discovered that he was being front-run by high-frequency traders who were working at much faster speeds than him, and were thus able to profit from his trades. He also discovered that not many investors, not even those he had considered sophisticated, were aware of how they too were falling prey to aggressive high-frequency trading algorithms. As a result, he convinced the Royal Bank of Canada to assemble a team tasked with developing an alternative trading system, Thor, designed to neutralize the advantage of suspected predator algorithms.³ Building the system was one thing, but attracting clients willing to use it was another. *Flash Boys* recounts Katsuyama's considerable efforts as he set up meeting after meeting with potential investors, but all to little avail. 'Then came the so-called flash crash', writes Lewis (though he misstates the timing):

At 2:45 on May 6, 2010, for no obvious reason, the market fell six hundred points in a few minutes. A few minutes later, like a drunk trying to pretend he hadn't just knocked over the fishbowl and killed the pet goldfish, it bounced right back to where it was before. (2014, p. 80)

According to *Flash Boys*, the Flash Crash was a turning point for Katsuyama and his attempts to counter the influence of predatory high-frequency traders. Many people began to ask 'the same question he was asking himself [about the Flash Crash]: Isn't there a much deeper question of how this one snowball caused a deadly avalanche?' (2014, pp. 81–2). His ideas now, suddenly, had considerable resonance in the investment community, and so '[a]fter the flash crash, Brad no longer bothered to call investors to set up meetings. His phone rang off the hook' (2014, p. 82).

What are the important factors in this account of the Flash Crash? Most notably, Lewis's account portrays it as something exceptional. It happened 'for no obvious reason' and appeared to be triggered by a minor event that eventually, but with extreme speed, spiralled both downwards and upwards. Indeed, in order to stress the exceptional, arguably even traumatic nature, of the event, Lewis draws a parallel between the Flash Crash and the crash of October 1929, effectively suggesting, therefore, that later generations might look back upon the Flash Crash as one of our era's defining moments (2014, p. 81).

A normal or exceptional occurrence?

Lewis's rendering of the event's exceptionality and singularity finds support in the academic literature. One illustration of this is found in an article by Donald MacKenzie entitled 'Mechanizing the Merc: The Chicago Mercantile Exchange and the Rise of High-Frequency Trading' (MacKenzie, 2015). This article offers a fascinating and thorough analysis of the historical backdrop to high-frequency trading, but it also proposes an interesting eventalization of the Flash Crash. Three features in particular stand out. The first relates to the delimitation of the event. The opening sentence of MacKenzie's article reads: 'At 2:40 p.m. on 6 May 2010, the US financial markets went into spasm' (2015, p. 646). The subsequent pages narrate the events of the day and how the Flash Crash, which began in Chicago, eventually spilled over, such that 'the sell pressure began to swamp the stock markets' in New York (2015, p. 648). For MacKenzie, then, the event is not reducible to what happened in Chicago; the ramifications on the New York markets equally form part of the Flash Crash event. Second, this spill-over logic makes MacKenzie characterize the Flash Crash as the 'first generalized crisis' of the 'new world of automated high-frequency trading' (2015, p. 648), i.e. the first illustration of how tightly interconnected markets are vulnerable to this type of trading. Third, since the Flash Crash is seen to embody a particular type of crisis that is intimately tied to high-frequency trading, the event plays an important role in MacKenzie's text as

a window into the new world of automated trading and its consequences. In other words, MacKenzie portrays the Flash Crash as an event that allows us to see attributes or dynamics that are crucial but usually escape attention.

While MacKenzie's analysis presents the Flash Crash as a window into present-day financial markets, he also treats it as an exceptional event, a rare occurrence. Other scholarly accounts question the exceptional status accorded to the Flash Crash, instead conceiving of the event as a window into the now normal workings of algorithmic finance, as exemplified by high-frequency trading. One example of this is a paper jointly written by Susanne von der Becke and the physicist and co-director of the Financial Crisis Observatory Didier Sornette, entitled 'Crashes and High Frequency Trading: An evaluation of risks posed by high-speed algorithmic trading' (a shorter version of the paper appeared in the UK Foresight report 'The Future of Computer Trading in Financial Markets'). In the paper, Sornette and von der Becke set out to find answers to the following question: 'Can high-frequency trading lead to crashes?' (2011, p. 3). This is answered in the affirmative, with the Flash Crash serving as a key example throughout the analysis. In fact, the Flash Crash is first mentioned in the paper's opening paragraph:

The May 6 flash crash in 2010, which saw the Dow Jones lose about 1 trillion USD of market value and individual stocks trading at fractions

or multiples within minutes, put increased focus on HFT. Even though high frequency traders were subsequently mostly cleared from having caused the crash, doubts remain as to whether this new form of trading bears potentially destabilizing risks for the market. (2011, p. 4)

Two things are interesting about this quote. One is that, whereas MacKenzie presents an expansive account of the Flash Crash, linking the initial price movements in Chicago to subsequent developments in New York, Sornette and von der Becke offer a narrow demarcation. In their analysis, the Flash Crash event is reduced to the price drop, with no mention of the fact that prices quickly recovered, and therefore that the overall 1 trillion USD losses were rapidly regained. This narrow eventalization of the Flash Crash adds weight to the second interesting feature of the quote: that the event is linked directly to the ‘potentially destabilizing risks’ that high-frequency trading is said to introduce to financial markets (for similar accounts, see Lange, Lenglet and Seyfert, 2016; Snider, 2014, p. 754). This is indeed how Sornette and von der Becke invoke the Flash Crash throughout their paper: it serves the strategic function of exemplifying what might go wrong with algo-financial markets, i.e. markets dominated by fully automated computerized algorithms, and how they are allegedly prone to crashes.

One of the central messages conveyed by Sornette and von der Becke (and repeated by, e.g. Leal *et al.*, 2014) is that the Flash Crash may appear

to be a singular and exceptional event, but should rather be seen as illustrative of a destabilizing, crash-prone dynamic that is inherent in present-day high-frequency-trading-dominated financial markets. In other words, its exceptionality is only apparent. A similar conclusion is reached by other commentators who, in effect, argue that the Flash Crash constitutes an important event because it captures something that is increasingly typical and representative of current financial markets. One example of such a reconfiguration of the Flash Crash – from an exceptional toward an increasingly typical or representative eventalization – is apparent in the work of Nanex. This firm, a financial data provider, has not only been an outspoken critic of the official CFTC-SEC report on the 2010 Flash Crash, but has also devoted a substantial part of its research to demonstrating the repeated recurrence of flash crashes in algo-financial markets. The Nanex homepage cites more than 100 flash crashes in 2013, and more than 40 between January and June 2014. According to the Nanex definition, a (down) flash crash occurs when the following applies to a stock price (the definition implies that flash events can also refer to rapid *positive* price changes):

- (i) it has to tick *down* at least 10 times before ticking *up*,
- (ii) price changes have to occur within 1.5 seconds,
- (iii) price change has to exceed -0.8%.

(quoted from Golub, Keane and Poon, 2012, p. 5, italics in the original)

Several academic papers have adopted this definition in their analysis of algo-financial markets' alleged flash crash-prone nature. One example is the article 'Financial black swans driven by ultrafast machine ecology', by the physicist Neil Johnson *et al.* (one co-author of which is the founder of Nanex, Eric Hunsader). In it, the authors examine what they refer to as 'ultrafast black swan events' or 'black swan crashes' (Johnson *et al.*, 2012, pp. 1, 3; see also Johnson *et al.*, 2013). They refer to the Flash Crash as a critical, destabilizing event that illustrates the need for further understanding of such financial black swans. Thus, they argue:

Society's drive toward ever faster socio-technical systems means that there is an urgent need to understand the threat from 'black swan' extreme events that might emerge. On 6 May 2010, it just took five minutes for a spontaneous mix of human and machine interactions in the global trading cyberspace to generate an unprecedented system-wide Flash Crash. (2012, p. 1)⁴

Specifically, Neil Johnson *et al.* identify a total of 18,520 ultrafast black swan events in a dataset comprising stock-price movements from 3 January 2006 to 3 February 2011. They take the notion of black swans from the work of Nassim Taleb, according to whom a black swan event has three attributes:

First, it is an *outlier*, as it lies outside the realm of regular expectation, because nothing in the past can convincingly point to its possibility. Second, it carries an extreme impact [...]. Third, in spite of its outlier status, human nature makes us concoct explanations for its occurrence *after* the fact, making it explainable and predictable. (2010, p. xxii, italics in the original)

For present purposes, what is especially interesting about this definition is the first attribute, and how Johnson *et al.* relate to it. On the one hand, the flash crashes examined by Johnson *et al.* are indeed seen as black swan events, i.e. as outliers or exceptions. Referencing the work of sixteenth-century philosopher Francis Bacon, Johnson *et al.* note that ‘the scientific appeal of extreme events [such as ultrafast black swan events] is that it is in such moments that a complex system offers glimpses into the true nature of the underlying fundamental forces that drive it’ (2012, p. 3). On the other hand, however, the sheer number of black swan events analyzed undermines the notion of them being exceptional or outliers. Johnson *et al.* seem aware of this, noting that ‘[w]e uncovered 18,520 such black swan events, which surprisingly is more than one per trading day on average’ (2012, p. 3). This, however, is a vast underestimation. If we say that there are about 261 trading days per year (holidays excluded), then the period under scrutiny includes around 1,325 trading days in total, meaning that there are, on average, close to 14 black swan events per trading day. Of course, to get a

full picture of the relative frequency of these events, one would have to compare them to the actual number of orders executed – a number that is much higher in present-day algo-financial markets than in pre-algorithmic markets. Still, if about 14 black swan events can be identified per trading day on average, it hardly makes sense to speak of *exceptional* events.

The same sense of the normality rather than exceptionality of flash crashes is conveyed in an analysis by Anton Golub, John Keane and Ser-Huang Poon (2012). They too subscribe to Nanex’s definition of a flash crash, but particularly focus on the four months between 2006 and 2011 that had the highest frequency of flash crashes (namely, September, October and November 2008, as well as May 2010). The resulting dataset contains a total of 5,140 flash crashes, which amounts to more than 63 flash crashes per day on average. The flash crashes analyzed by Golub, Keane and Poon are ‘scaled down versions of the May 6th 2010 Flash Crash’, for which reason they refer to them as ‘Mini Flash Crashes’ (2012, p. 2). Although these mini crashes may not possess Taleb’s second attribute of black swan events (having an *extreme impact*), the authors, in the introduction to the paper, devote a full paragraph to the 2010 Flash Crash. The strategic use of the Flash Crash is obvious: as in the paper by Johnson *et al.* (2012), it serves as a cautionary reminder of how algo-financial markets currently operate, and of the kinds of crashes to which they are supposedly prone. Thus, while the many scaled-down flash crashes analyzed by Golub, Keane and Poon may not have been as severe as the 6 May 2010 Flash Crash,

they are characterized by the same underlying logics and thus *could* have produced far more devastating effects than they in fact did – this is the argumentative framework promoted by Golub, Keane and Poon.⁵

I have tried to demonstrate that the Flash Crash is indeed a restless event. While there is widespread agreement about the event's importance, its demarcation in time and space is more contested. Similarly, there is disagreement about whether the event represents something normal or exceptional in algo-financial markets. Interestingly, in its emphasis on the importance of the Flash Crash, much of the scholarly commentary explicitly asserts or takes for granted that the event indeed signifies a major change in financial markets, in the sense that it represents a type of crash-like tendency now inherent in algo-financial markets, and/or in the sense that it constituted a significant crisis. The latter interpretation is visible both in MacKenzie's characterization of the Flash Crash as the 'first generalized crisis' in the world of high-frequency trading (2015, p. 648) and when, e.g. Johnson *et al.* liken the event to a black swan event that has 'extreme impact' (Taleb, 2010, p. xxii). Similarly emphasizing the alleged impact of the Flash Crash, Andrei Kirilenko, former Chief Economist of the CFTC, has argued that 'the algorithmic nature of HFT can cause significant impact. The pre-eminent example of this, of course, is the Flash Crash of May 6th of 2011 [*sic*]' (Kirilenko, Sowers and Meng, 2013, p. 59). However, as I shall argue in the following section, there is little evidence to support this notion of significant impact.

What was the impact of the Flash Crash event?

The eventualizations of the Flash Crash analyzed so far point to at least three ways in which the event might be of considerable importance to present-day financial markets: (a) as an event that significantly changed how market participants perceive markets; (b) as an event that generated a massive loss of value, and hence had or could have had considerable economic effects; and/or (c) as an event that is symptomatic of a novel set of systemic risks associated with algorithmic finance. I shall critically discuss each of these in turn.

The first potential impact, an experiential one, is alluded to in Lewis's account of the Flash Crash, when he sheds light on how Katsuyama and other market participants experienced the Flash Crash, how it seriously affected them, and the questions it raised in the trading community. Lewis's image of the event as being dramatic, even a turning point, is echoed in other journalistic accounts. For example, in their coverage of the Flash Crash the day after the event, the *Financial Times* quoted a senior trader as saying 'Was this the end of the world? No, but it sure felt like it for a few minutes' (Bullock, Mackenzie and Guerrera, 2010, p. 8). In Marije Meerman's 2013 documentary *The Wall Street Code*, David Lauer, a former high-frequency trader at Citadel (one of the main US high-frequency-trading players), describes the Flash Crash in even more dramatic parlance. According to Lauer, the Flash Crash was for him 'a defining event'. He

relates how, during the escalating market drop, he and his fellow high-frequency traders followed the ‘drifting’ of orders in the order book, and how he felt that ‘something terrible has just happened. Something indescribably horrible just happened’. The horrible thing was that ‘the market was gone’; ‘for moments, for seconds, there was no market’, no sets of buyers or sellers to be matched. ‘Even 9/11 didn’t have that impact’, he states.

It is difficult to say whether the Flash Crash generated similar responses in other market participants, as this has never been systematically examined. But to the extent that it had long-lasting experiential effects, these are likely to have been distributed differentially across different market actors (high-frequency traders, retail investors, brokers, pension funds, etc.), depending on whether (and how) they were engaged in trading on the afternoon of 6 May 2010. Interestingly, however, my interviews with people involved in algorithmic finance, including high-frequency traders, do not confirm the dramatic portrayal of the event in *Flash Boys* or *The Wall Street Code*. Indeed, none of my interviewees referred to the Flash Crash in a manner remotely parallel to Lauer’s account. Illustratively, one experienced New York-based high-frequency trader took a rather relaxed, unaffected approach to the event, and placed particular emphasis on how the downward part of the Flash Crash (the market drop) was followed by what he referred to as a ‘Flash Recovery’, most likely generated by the high-speed trading algorithms. Contrary to Lauer’s account, this suggests that

the most spectacular aspect of the Flash Crash was not the absence of the market, but rather its immediate auto-recuperation.

Similarly, when asked about the central risks characterizing contemporary algorithmic markets, a former high-frequency trader from one of the leading US high-frequency trading firms, now CEO of a firm that specializes in a combination of high-frequency trading and advanced machine learning, asserted that flash-crash events do not amount to any substantial risk. There should be much greater concern, he argued, with the consolidation taking place among algorithmic firms where a few major players trade very significant parts of the overall volume (the biggest firm being in charge of around 20% of the order flow in some markets), with all the market fragility and risk this entails.

Market participants are obviously scared of losing large amounts of money. And so, if the Flash Crash did in fact generate the sense of the world ending, then this was most likely tied to the fear of markets collapsing and profits crumbling. Interestingly, however, there has been hardly any systematic examination of the actual economic effects of the Flash Crash – and, it turns out, for good reasons, since it is difficult to ascertain its precise economic impact. However, there are a small number of indicators that provide a partial picture. Some were described in a 2010 talk entitled ‘Strengthening Our Equity Market Structure’, delivered by the then SEC chair, Mary L. Schapiro, a few weeks prior to the publication of the joint CFTC-SEC report on the Flash Crash. The Flash Crash was central to

the talk and was described by Schapiro as a clear ‘market failure’ that required the most detailed examination. As such, it was viewed not as an ‘aberration’, not as something exceptional, but as an event that may reveal some fundamental problems in the current market structure and its dynamics (2010, p. 3).

Discussing the impact of the Flash Crash, Schapiro made three observations. The first concerned stop-loss orders, which are ‘designed to help limit losses by selling a stock when it drops below a specified price’ (2010, p. 3). During the Flash Crash, the reliance on and use of this safety tool ‘backfired’ (2010, p. 3):

A staggering total of more than \$2 billion in individual investor stop loss orders is estimated to have been triggered during the half hour between 2:30 and 3 p.m. on May 6. As a hypothetical illustration, if each of those orders were executed at a very conservative estimate of 10 percent less than the closing price, then those individual investors suffered losses of more than \$200 million compared to the closing price on that day. (2010, pp. 3–4)

USD 200 million is obviously a lot of money. Yet, compared to the 1987 crash, for example, where 19 out of 23 global markets declined by more than 20% (Roll, 1988), the Flash Crash appears relatively slight. The 1987 crash also had much longer-lasting effects on markets. These differences

are nicely captured in a recent discussion of the Flash Crash by Imad Moosa, which has a rather more sympathetic view of high-frequency trading than most accounts (and essentially concurs with the abovementioned notion of a ‘Flash Recovery’):

On 6 May 2010, the market dropped as much as 9.2 per cent, but only stayed down for a matter of minutes before recovering most of its losses. On 19 October 1987, the market dropped 23 per cent and stayed down for weeks and months. If anything, it was high-frequency traders who saved the day [on 6 May 2010] by staying in when long-term traders left the market. (Moosa, 2015, p. 84)

While the effects of the 1987 crash certainly lasted longer, Schapiro noted a second economic impact of the Flash Crash, namely that both individual investors and mutual funds ‘pulled back from participating in the equity markets’ for months afterward (2010, p. 3). I have not been able to find a quantified estimate of what that might have meant for the financial markets. However – and this is the third central impact listed by Schapiro – the most important consequence of the Flash Crash might also lie elsewhere, in something less quantifiable. According to a survey mentioned by Schapiro, ‘Less than 50 percent of the buy-side respondents [...] expressed confidence in the current market structure’ (2010, p. 4). It goes without saying that lack of confidence in markets could potentially have a devastating ef-

fect, but it is unclear from Schapiro's talk precisely how this might materialize. Some insight into this is available in an analysis by Thomas Boulton, Marcus Braga-Alves and Manoj Kulchana (2014), who find signs that the Flash Crash led to increased investor uncertainty over a period of at least two weeks, with consequences – including widening bid-ask spreads and market withdrawal – that meant that it became more expensive to trade in markets in the immediate aftermath. However, as Boulton, Braga-Alves and Kulchana note, it is difficult to single out which consequences were due specifically to the Flash Crash and which were due to the sovereign debt crisis in Europe in which the Flash Crash was at least partly embedded (2014, p. 154). In fact, it turns out that the evidence about the economic impact of the Flash Crash is inconclusive. In a survey conducted by BlackRock, the world's largest asset manager, 380 retail financial advisors were asked about the impact of the Flash Crash. The survey demonstrated 'that the majority of advisors were minimally affected by the market disruption', and BlackRock therefore concluded that 'the final impact on investors was relatively limited due to widespread trade cancellations' (2010, p. 3). So while the evidence regarding the actual economic impact is inconclusive, it suggests that the Flash Crash had little economic effect – not least because markets quickly rebounded.

This leads me to the third potential impact that is implied in eventualizations of the Flash Crash: even if the Flash Crash did not have a massive actual economic effect, the event is nonetheless significant because it is

symptomatic of new types of systemic risk associated with algorithmic finance – and these point far beyond the immediate economic aftermath of 6 May 2010. In the words of Andersen and Bondarenko, the Flash Crash is important as an example of ‘highly erratic market dynamics’ which might propagate ‘into systemic disruptions to the financial systems’ (2015, p. 2; see also Ivanov, 2011; Knight, 2012). This observation is echoed by scholars such as Johnson and Sornette, as described in the previous section. They essentially analyze the market as a complex system or ecology in which Flash Crash-type events constitute an imminent risk. For Johnson *et al.*, the Flash Crash is significant as an illustration (a) of the transition into a fully automated, machinic ‘population of adaptive trading agents’; and (b) of the ‘systemic risk’ looming in such an ecology (Johnson *et al.*, 2013, pp. 3, 5). A similar image is conveyed by Sornette and von der Becke, who describe financial markets as ‘truly “complex systems” in a technical sense. As such, they are intrinsically characterized by periods of extremity and by abrupt state-transition and spend much time in a largely unpredictable state’ (2011, p. 15). It is due to such unpredictable states, they argue, that high-frequency trading can produce ‘[p]ro-cyclicality mechanisms, also known as positive feedbacks’, something which allegedly ‘leads to unsustainable regimes, ending in crashes and crises’ (2011, p. 16). In other words, if markets are already unstable, as they were on the morning of 6 May 2010 (in the light of the Greek debt crisis), then the instability can be further exacerbated by high-frequency trading algorithms herding one an-

other and thereby transforming instability into widespread crashes (a view also expressed by Kirilenko *et al.*, 2014).

Looking critically at such accounts, it is no coincidence that they tend to work with a narrow eventalization of the Flash Crash in which, as per Sornette and von der Becke (2011, p. 4), only the market *drop* of 6 May 2010 is taken into consideration. The notion of a potential flash recovery, itself generated by algorithmic trading, has no place in this logic of algorithmic ‘panic herding’ (2011, p. 11). This is not the only problematic feature of such accounts, which often originate within the field of econophysics. They also suffer from an overly naturalistic analysis of financial markets that draws explicitly on epidemiology and population ecology as key points of reference. For example, Sornette and von der Becke argue that the alleged ‘endogenous self-excitation nature’ and ‘viral epidemic’ of high-frequency trading algorithms warrant depiction via market ‘pandemics’ (2011, pp. 7, 11–14; for a critical discussion of the adoption of an epidemiological lexicon in the financial realm, see Peckham, 2013).⁶ In that light, it is interesting to see that Sornette and von der Becke, if only briefly, point to ‘human oversight’ of automated trading algorithms as ‘a stabilising factor for the HFTs and a mechanism to mitigate the risk of herding’ (2011, p. 12). While, unsurprisingly, none of my interviewees shared the impression that algorithmic trading in itself poses an imminent threat to financial markets, some did express concern that specific innovations in algorithmic finance might eliminate human oversight, with potentially dire

implications for the control of the algorithms. For example, an experienced high-frequency trading system builder, now head of an IT department that develops electronic trading systems at a major bank in New York, voiced what he saw as widely held concerns with algorithmic trading based on machine learning:

many people are reluctant to use [machine learning] because many models that do machine learning are not intuitive. For people, it's kind of hard to understand those models, and a lot of what traders do is still based on some intuition. When the models become counterintuitive, it's very hard to say when the models are going off and when the models are doing right. You don't have a human control of the system's behaviour because you 100% trust the model and not many people are willing to do that – some do, but not everybody.

In that sense, some market participants do appear to second, e.g. Sornette and von der Becke's observation that, without proper human monitoring – and, one might add, *understanding* – the interaction of fully automated trading algorithms might spiral into, for instance, Flash Crash-type events. In other words, if algorithms move too far beyond intuition, they might have a huge negative impact on financial markets.

Crash tropes

The fact that the types of impact often associated with the Flash Crash are at least to some extent contestable does not seem to have lessened interest in the event. The Flash Crash continues to attract attention as a significant event, both in the media and in scholarly circles. I would suggest that this is partly because the eventalizations of the Flash Crash evoke particular well-known tropes. More precisely, the discussions of the Flash Crash often centre on familiar notions that contribute to making what happened on 6 May 2010 recognizable as an *event*. In other words, a dual logic appears to be at play here: (a) even if the Flash Crash is, as per Wagner-Pacifici's formulation, a restless event, in that its eventness (its main dynamics, temporal and spatial boundaries, key actors, central impact, etc.) is differently articulated in different eventalizations, certain common tropes seem to crystalize in these eventalizations; at the same time, it might be argued (b) that the Flash Crash lends itself to eventalization because the key dynamics attributed to it echo wider, well-established tropes.

I will discuss two such tropes in particular – tropes that are usually considered distinct from one another, but which tend to coalesce in eventalizations of the Flash Crash. The first concerns the kind of systemic risks thought to be generated by automated (especially high-frequency) trading. While the immediate media reports on the event conjectured that the market drop was triggered by some 'fat finger' error, i.e. by some (principally avoidable) human mistake (e.g. *Financial Times*, 2010), the focus swiftly changed to the allegedly inbuilt propensity of algorithmic trading to run

amok in an unanticipated frenzy. This is precisely the image mobilized by, e.g. Sornette and von der Becke in their claim regarding the supposed ‘endogenous self-excitation nature’ of high-frequency trading (2011, p. 11). Such references tap into broader concerns with technology, specifically how it increasingly replaces humans or fundamentally alters their societal status, as well as the fear of technology running wild and being opaque to humans (including regulators). While this trope about technological risk is as old as modern society itself, it acquires a particular framing in the eventualizations of the Flash Crash. On the one hand, several accounts characterize the Flash Crash as symptomatic of what might happen once trading becomes automated and humans are more or less left behind. This is the imaged conjured up in MacKenzie’s notion of a ‘generalized crisis’ and in the econophysics-inspired literature on the crash-prone nature of high-frequency trading. It is equally conjured up in notions of the dangerous ‘nontransparency of computerized trading’ (Golumbia, 2013, pp. 286–7), and in related discussions of how the complexity and speed of high-frequency trading have outmanoeuvred regulators and overwhelmed their capacity to effectively monitor modern markets (Lewis, 2014; Malmgren and Stys, 2011; Snider, 2014).

On the other hand, some accounts, especially that of Lewis, emphasize that while the Flash Crash was indicative of algorithmic markets that are out of control – as evinced by the seemingly obscure market swings of 6 May 2010 that defied reasonable explanation (they happened ‘for no obvi-

ous reason', in a manner comparable to the irrational behaviour of a drunk) – algorithmic finance is not simply a matter of inter-machine wars. Rather, according to Lewis, algorithmic finance is at once embedded in mundane human profit-seeking behaviour and recasts the relations between humans and machines. Indeed, running through Lewis's entire narrative in *Flash Boys* is the assertion that high-frequency trading divides markets into an elite of machinic traders and an underbelly of non-machinic ones – or, in Lewis's parlance, between algorithmic predators and human prey.⁷

This connection between the alleged self-excitation nature of trading algorithms and the ways in which such algorithms' increasing dominance of financial markets reshuffles the relations between humans and machines aligns the fear of technology with a different trope, one that is also evoked in eventalizations of the Flash Crash. This latter trope concerns another modernist worry: the alleged destabilization and de-individualization that are the consequences of crowd and herd behaviour. The articulation of this concern was particularly strong in the formative years of the sociological discipline, around the turn of the nineteenth century (Borch, 2012). One manifestation of this appeared in Émile Durkheim's reflections on social 'currents': 'in a public gathering the great waves of enthusiasm, indignation and pity [...] come to each one of us from the outside and can sweep us along in spite of ourselves' (1982, pp. 52–3). Durkheim would later further elaborate on this topic in his discussion of the exultant efferves-

cence that characterizes the religious practice of the *corrobbori*, as outlined in his sociology of religion. According to Durkheim:

When [people] are once come together, a sort of electricity is formed by their collecting which quickly transports them to an extraordinary degree of exaltation. Every sentiment expressed finds a place without resistance in all the minds, which are very open to outside impressions; each echoes the others, and is re-echoed by the others. The initial impulse thus proceeds, growing as it goes, as an avalanche grows in its advance. (1947, pp. 215–16)

The force of Durkheim's observation lies in the fact that, while the latter quote described a non-modern incident, it also captured an experience that was precisely, and especially, modern. This modern anchoring was clearly articulated in other sociological accounts, which conveyed a widespread late-nineteenth and early- to mid-twentieth-century sense of modern society as being intimately linked to collective forms of being-carried-away (see e.g. Fromm, 1941; Sidis, 1898; Simmel, 1950; Tarde, 1962). What people experience in the collective formations of modern society, it was held, is a loss of self: strict boundaries between people disappear as a result of the social avalanche of crowd and herding behaviour. Significantly, this experience was not restricted to the collective formations on the street or in particular religious settings – it was soon associated with financial markets,

too. Examples and discussions of panicking and being-carried-away in speculative frenzy have a long history, with the 1637 Tulip mania and the 1720 Mississippi bubble being oft-mentioned (e.g. Mackay, 2002). However, what is special about the early twentieth century is that the financial markets were increasingly described in a crowd-psychology lexicon. For example, in the 1920s and '30s, advocates of so-called contrarian investment philosophy argued that financial markets are widely characterized by contagious crowd dynamics, in which individual investors are easily swept away by the market crowd's avalanching impulses (see Borch, 2007; Hansen, 2015; Stäheli, 2006; 2013; see also Arnoldi and Borch, 2007). It is this notion of a sudden social avalanche – a destabilizing and de-individualizing crowd or herding process – that is also evoked in some eventalizations of the Flash Crash. This not only applies to Lewis's explicit likening of the Flash Crash to 'a deadly avalanche' (2014, p. 82), but also to e.g. Sornette and von der Becke's analysis of the crash-prone nature of algorithmic herding (2011, pp. 10–12). They argue that Flash Crash-like events essentially display dynamics in which, to use Durkheim's words, each algorithm 'echoes the others, and is re-echoed by the others'.

I have suggested that two tropes in particular are being mobilized in many eventalizations of the Flash Crash: one concerning the fear of technology, the other concerning the fear of crowds/herding and the social avalanches they are said to produce. Although these tropes have distinct genealogies, they coalesce in discussions of 6 May 2010, not least around hu-

man-machine relations. So, while concerns regarding crowds and their social avalanches have usually been framed as a primarily, if not exclusively, inter-human affair (by Durkheim, Simmel, Tarde, etc.), the mobilization of crowd tropes to understand the Flash Crash essentially applies this human-oriented trope to a non-human, algorithmic domain. What was previously reserved for human beings is now attributed to trading algorithms that operate independently, without direct human intervention. This entails that, contrary to the notion that machines and humans exist in a mutually antagonistic relationship (the former challenging the latter), the eventalization of the Flash Crash as a matter of herding behaviour essentially suggests that some form of sociality is conferred upon algorithms, namely the sociality of interacting actors captured in an avalanching maelstrom. In other words, one implication of these eventalizations of the Flash Crash is that individual algorithms, too, can be swept away.

However, there is one important respect in which the algorithmic herding of 6 May 2010 is not quite equivalent to the being-carried-away described by modernist observers such as Durkheim, Simmel and Tarde. At least according to the official CFTC-SEC (2010, p. 5) report, ‘a significant number [of automated trading algorithms] withdrew completely from the markets’ during the downward-spiralling part of the Flash Crash, as they were programmed to shut down in cases of extreme volatility. As such, while some algorithms were caught in an avalanching crowd process, others were pre-programmed or instructed (by the humans monitoring their

activities) to stand aside, to seek refuge from the homogenizing movement, as it were. This is an attribute rarely ascribed to humans in sociological accounts of crowd dynamics (see e.g. Tarde, 1968). Here, individuals are typically described as being entirely swept away, unable to evade the maelstrom. This might suggest that herding and social avalanching assume different shapes and intensities in human and non-human domains. However, an arguably more important observation is that, in the overall eventualizations of the Flash Crash, this standing aside is often ignored or downplayed (e.g. Johnson *et al.*, 2012; Johnson *et al.*, 2013; McInish, Upson and Wood, 2014).

In the context of discussing relations between humans and machines, it is relevant to return briefly to Sarao. The Sarao case seemingly confirms the observation that the kinds of algorithmic herding identified in present-day algorithmic finance are not simply a matter of algorithms being carried away by one another; humans may also play a highly active role in governing inter-algorithmic dynamics. And yet, it is remarkable how even the commentary on Sarao's supposed involvement in the Flash Crash actually supports the notion that humans are relegated to the margins of contemporary algo-financial markets. Illustratively, commenting on the initial charges against Sarao, *The Economist* notes that it seems rather unlikely that a single trader could manipulate markets to such an extent that they would temporarily collapse (*The Economist*, 2015). According to this type of ac-

count, the profound technological complexity and interconnectedness of markets render significant interventions of isolated humans less likely.

Implications for economic sociology: Dissonance and resonance reconsidered

So far, I have treated eventualizations of the Flash Crash from the perspective of second-order observation, meaning that I have observed how different observers observe the Flash Crash and how they constitute it as an event. Luhmann (e.g. 2012) separates such second-order observations from first-order observations, the latter of which are observations about ‘something’, e.g. the causal logics allegedly inherent to modern markets. Luhmann’s sociological project (at least in its final phase, see Borch, 2011) is characterized by a strong preference for second-order observations, since, according to Luhmann, such observations bracket ontological assumptions about the world. More precisely, by observing how other observers observe, the sociologist need not make any ontological claims about how the social reality is constituted, but can instead bring into relief the ontological assumptions of other observers. In spite of Luhmann’s reservations about first-order observations, they can be of analytical value (and nor, indeed, are they entirely absent from his work). In that spirit, I now wish to discuss what the eventualizations of the Flash Crash might entail for economic sociology if they are not treated (in a distanced fashion) as second-order observations, but are taken seriously as first-order observations about how pre-

sent-day algo-financial markets actually operate. The discussion will revolve around Daniel Beunza and David Stark's (2012) analysis of dissonance and resonance in quantitative finance, which is conceived as a correction to debates on, e.g. herding behaviour and black swan events, both of which play a central role in eventualizations of the Flash Crash.

Beunza and Stark's analysis focuses on how quantitative merger arbitrageurs deploy a set of socio-technical devices in order to develop their own estimates of particular variables, as well as to 'check their own estimates against those of their rivals', i.e. other merger arbitrageurs (2012, p. 384). This taking-into-account of essentially social cues amounts, in Beunza and Stark's terms, to a modelling endeavour turned reflexive. Importantly, quantitative traders deploy reflexive modelling as a means through which to incorporate *dissonance* into their trading: if the models indicate that rival traders estimate important variables differently, this is an invitation not so much to follow them blindly, but to scrutinize matters more carefully. However, while dissonance thus plays a productive role for quantitative traders, Beunza and Stark stress that there is a risk that the incorporation of other market participants' assessments can generate calamitous 'cognitive interdependence', with traders making unprofitable decisions based on inaccurate social cues (2012, p. 410). For example, the trading community might collectively ignore important features of a potential merger or be (incorrectly) overconfident about other features, meaning that individual traders might make investment decisions without being ex-

posed to any real dissonance. In such cases, dissonance is replaced by *resonance*. So, '[r]eflexive modelling amplifies individual errors when a sufficiently large number of arbitrage funds have a similar model. Whereas reflexive modelling improves trading on the basis of *dissonance*, it can lead to financial disaster in the presence of *resonance*' (2012, p. 410; italics in the original).

Beunza and Stark present their notion of reflexive modelling as superior to notions of herding and black swans. Referencing Taleb, among others, they argue that black swan accounts have the merit of pointing to how crises might occur if market participants deploy models that significantly 'underestimate uncertainty', e.g. by assuming that stock returns follow normal distributions, rather than (as would be more realistic) fat-tailed ones (2012, p. 386). That said, however:

the black swan is ultimately an under-socialized explanation of the risks created by models. The black swan presents financial actors as hopelessly unreflexive about the limitations of their models. Confronted by uncertainty about the model, we would expect market actors to rely on the social cues around them. (2012, p. 386)

Beunza and Stark acknowledge that such cues are included in herding models. However, these models 'do not account for the existence of technology in the decision-making process' (2012, p. 387). This allegedly

means that, while herding and imitation might have offered ‘a realistic portrayal of financial actors before the 1980s’, ‘the introduction of computers, equations and models into financial markets’ has rendered herding behaviour irrelevant, as the ‘handling and manipulating a body of codified knowledge [...] cannot simply be put to the side for the sake of copying someone else’s decision’ (2012, p. 387). It is in the light of such critiques that the notion of reflexive modelling is considered superior, because it better accounts for the incorporation of social cues and for the socio-technical reality of quantitative finance.

Beunza and Stark’s article is based on ethnographic fieldwork conducted between 1999 and 2003. At that time, as the authors report, human traders considered quantitative tools to be important and helpful – but quantitative finance had not yet developed into the type of fully automated trading that dominates current financial markets. To pointedly emphasize (and slightly exaggerate) this difference, we might say that in the early 2000s, quantitative tools supported human traders and their decision-making, whereas today human traders support fully automated algorithms. As the eventualizations of the Flash Crash suggest, this shift has significant implications for the ways in which quantitative finance generates and relates to dissonance and resonance. There are four dimensions to this.

First, present-day algorithmic finance, including high-frequency trading, is far more technology-driven than the kind of quantitative finance analyzed by Beunza and Stark. While Beunza and Stark pay considerable

attention to the merger arbitrageurs' use of (rather low-tech) spreadsheets and phone calls, high-frequency trading is based on sophisticated forms of high-speed data-processing, server co-location, microwave transmission, etc. (Borch, Hansen and Lange, 2015; Borch and Lange, 2016; MacKenzie, 2014). In this extremely technological and deeply interwoven market reality, particular forms of 'cognitive interdependence' seem to take shape, not so much among human actors (as per Beunza and Stark's analysis), but rather among algorithms. This is precisely what the Flash Crash laid bare, according to a host of the eventalizations analyzed above: this event demonstrated that human oversight and human cognitive interdependence become both less feasible and less important in algo-financial markets, which instead become dominated by tightly entangled algorithms operating at extreme speeds, beyond human perception.

Second, eventalizations of the Flash Crash do not seem to suggest that the kind of mature quantitative finance epitomized by high-frequency trading is based on reflexive modelling. Contrary to Beunza and Stark's example of merger arbitrageurs, high-frequency traders do not seek to systematically incorporate the observations of rivals *in order to check their own strategies*. While they certainly take into account how others act in the market and how rivals seek to tweak the order book (in order, e.g. to limit market footprint; for sociological discussions of such behaviours, see Borch and Lange, 2016; Coombs, 2016; Lange, 2015; MacKenzie, 2014; forthcoming), they do not do so in order to critically question their own

practice. Rather, they do it to better grasp the direction in which the market is moving and to profit from imitative and counter-imitative strategies (see again Lange, 2015; 2016).

This leads me, third, to the question of herding. Sornette and von der Becke emphasize in particular that high-frequency trading algorithms can engage in herding. This runs counter to Beunza and Stark's dismissal of herding as a relevant phenomenon today. The contrast at stake here is due in part to a significant difference between their conceptualizations of herding. Beunza and Stark discuss the concept by referencing financial economics literature (e.g. Scharfstein and Stein, 1990) and neoinstitutionalist accounts (DiMaggio and Powell, 1983) in which herding is conceived as a form of instrumental action: in situations of information uncertainty, it makes sense to imitate others, as they might have better information than oneself. By contrast, in Sornette and von der Becke, the notion of herding has more in common with traditional sociological tropes of people being carried away in collective ecstasy. This conception of herding is ignored in Beunza and Stark's instrumental rendering. And yet, eventalizations of the Flash Crash suggest that it might be precisely such forms of non-instrumental herding that are best suited to accounting for what happened on 6 May 2010. The high-speed trading algorithms' mimicking of each other is not necessarily due to instrumental reasons, but could be a consequence of their resonance-oriented design, i.e. due to the inter-observational ways in which they are programmed. Such non-instrumental

forms of herding can exist even if imitative strategies are otherwise instrumentally designed in the algorithms of individual high-frequency traders (Lange, 2015). Therefore, Beunza and Stark's finding that quantitative merger arbitrageurs 'were emphatically not mimicking their rivals' (2012, p. 402) does not apply to the Flash Crash. On 6 May 2010, algorithms were emphatically mimicking their rivals, downwards and upwards.

Fourth, the eventalizations of the Flash Crash may challenge Beunza and Stark's point that 'arbitrageurs can collectively be wrong' (2012, p. 391). This point is central to their analysis of how collective ignorance or overconfidence might lead to disastrous resonance. However, perhaps being right or wrong is a false dichotomy. The eventalizations of the Flash Crash, at least, forego such a distinction: neither the market drop nor the subsequent market recovery were seen as a matter of algorithms being 'right' or 'wrong'. Neither part of the process can be adequately understood in those terms. What the Flash Crash eventalizations point to – in particular, those inspired by the work of Nanex – is simply that herding can go both ways: resonance can spiral downward, generating massive instantaneous market drops, as well as upward, leading to instantaneous price rises or market recoveries.

I have argued, on the basis of eventalizations of the Flash Crash, that Beunza and Stark's interesting analysis of reflexive modelling in quantitative finance needs qualification when applied to fully automated algorithmic trading. However, there is one important respect in which their analy-

sis is apt when it comes to understanding algo-financial markets. One of the central achievements of their analysis is that it emphasizes that quantitative finance might generate particular forms of sociality that are not easily understood by extant approaches in economic sociology. Specifically, they argue, the sociality of quantitative finance bridges social cues with socio-technical aspects. In the words of Beunza and Stark:

Reflexive modelling thus brings quantitative finance full circle: whereas the introduction of models and information technology in the capital markets brought in anonymity and a semblance of objectivity in the data, reflexive modelling makes it clear that traders are modelling not just the economic but also the social. Although anonymous and impersonal, quantitative finance brings back the interdependence among the actors – and, for that reason, its social aspect. But this form of sociability around models does not easily fit existing frameworks in economic sociology – it is disembedded yet entangled; anonymous yet collective; impersonal yet, nevertheless, emphatically social. (2012, p. 412)

Eventualizations of the Flash Crash suggest that something similar could be said about fully automated trading algorithms, despite these not being based on reflexive modelling. These algorithms, too, are ‘disembedded yet entangled; anonymous yet collective; impersonal yet, nevertheless, emphatically social’. This social dimension of an otherwise emphatically non-

human domain – a realm of fully automated algorithms operating without human intervention – might appear controversial or at least counterintuitive. However, as argued earlier (and further analyzed by Lange, 2015; MacKenzie, forthcoming), the tropes evoked in eventalizations of the Flash Crash indicate that it may indeed make sense to conceive of a distinct social realm of interacting algorithms.

Conclusion

This article has critically examined how the Flash Crash is often invoked as an event that offers a window into the inner workings and potential problems of high-frequency trading and present-day financial markets. I have discussed the restlessness of the event, i.e. how the delimitation of the event is contested, as well as the controversy about whether it signifies something normal or exceptional in algo-financial markets. While this points to a tactical polyvalence of the event, which is interesting in its own right, the more crucial observation has been that some of the impact often associated with the Flash Crash is contestable. This applies in particular to the event's economic effects, but also to its imprint on market participants' perceptions of markets: market participants certainly do not uniformly see the Flash Crash as a defining event or a turning point. Similarly, while concerns regarding advanced algorithms might find some support among present-day market participants, eventalizations of the Flash Crash that emphasize the crash-prone nature of algorithmic finance, and especially

high-frequency trading, tend to strategically stress only the negative aspects of the event (the market drop) while neatly disregarding the subsequent recovery. In such accounts, substantial arguments enter into a melee with more dramatic renderings and usages of the Flash Crash.

On the basis of the critical analysis of the various eventalizations of the Flash Crash, as well as the contested nature of the different types of impact accorded to it, I then pondered why the Flash Crash nonetheless continues to attract so much attention. I proposed that to understand this, it is important to see how Flash Crash lends itself to eventalization. Specifically, I suggested that discussions of the Flash Crash mobilize two familiar modern tropes, each of which contributes to making the Flash Crash recognizable as an event. Thus, concerns regarding technology and herding occupy a central place in eventalizations of the Flash Crash. More precisely, these tropes coalesce in eventalizations of the Flash Crash. Technological concern in effect becomes a concern regarding machinic algorithms that herd in ways that are otherwise primarily characteristic of the human domain. This particular combination of tropes is arguably the central unique feature of the Flash Crash and its eventalizations.

Finally, I suggested that it might make sense to consider the eventalizations of the Flash Crash not so much in the light of how they delineate this particular event, but rather in the light of what they might contribute to economic sociology discussions about quantitative finance. In particular, I suggested that if eventalizations of the Flash Crash are treated as offering

substantial insights into how algo-financial markets actually operate – or if there is at least some accuracy to what they suggest – then this might prompt a reconsideration of how we conceive of dissonance and resonance, herding and cognitive interdependence within sociological accounts of quantitative finance. In addition, the Flash Crash eventualizations might stimulate economic sociologists to seriously examine how the interplay of fully automated algorithms might constitute new, distinct forms of market sociality, the understanding of which may well call for new sociological categories.

Funding

This work was supported by the Danish Council for Independent Research's Sapere Aude Grant for the project 'Crowd Dynamics in Financial Markets'.

Acknowledgements

Previous versions of this paper were presented at seminars and workshops at the University of Konstanz; the Department of Sociology, New School for Social Research; and the Institute for Public Knowledge, New York University. I am grateful to the participants and audiences for their valuable feedback. I owe special thanks for helpful comments to Nathan Coombs, Kristian Bondo Hansen, Ann-Christina Lange, Donald MacKen-

zie, Emilio Marti, Robin Wagner-Pacifici and the anonymous *Economy and Society* referees. I would also like to thank Thomas Lauronen and Benjamin Schwarz for research assistance, as well as my interviewees for taking the time to discuss their work with me. The usual disclaimers apply.

Disclosure statement

No potential conflict of interest was reported by the author.

Notes

1 The drop was registered in the so-called ‘order book’, which lists a trading venue’s buy and sell orders in a way that visualizes the market at any given moment. I stress this experiential anchoring of the market because high-frequency trading, which reportedly was partly to blame for the Flash Crash, takes place at timescales beyond human perception (see Borch, Hansen and Lange, 2015). So, on one level, high-frequency trading, in its actual operations, renders financial markets non-experiential. On another level, however, high-frequency trading orders leave visual traces in the order book, and thus lend themselves to human experience.

2 Reflecting this initial confusion, market participants engaged in active internet searches in order to understand what was happening. Bloomberg reported that ‘Yahoo! Inc., Google Inc. and at least one brokerage [Charles Schwab Corp] sustained slowdowns on Web pages that provide financial

information as U.S. stocks tumbled and users swarmed the Internet for market updates’ on 6 May 2010 (Womack, 2010).

3 Katsuyama later developed the New York-based IEX exchange (recently approved by the SEC), which similarly aims to neutralize the speed advantage of high-frequency traders.

4 David Golumbia (2013, p. 295) similarly characterizes the Flash Crash as a black swan event.

5 In one important respect, the analysis by Golub, Keane and Poon differs significantly from that of both Sornette and von der Becke, and Johnson *et al.* Whereas the latter focus on how the interplay of algorithms is responsible for flash crashes, Golub, Keane and Poon (2012) argue that most flash crashes occur as an effect of particular types of market regulation, specifically the use of so-called Intermarket Sweep Orders (cf. Johnson *et al.*, 2013, p. 6). For a similar finding, see McInish, Upson and Wood (2014).

6 Observing this particular feature of the econophysics literature on the Flash Crash may be seen as a supplement to the characterization of econophysics offered by Jovanovic and Schinckus (2013). In their illuminating genealogy of this new field, they especially stress ‘econophysics’ major

distinguishing feature, which is the use of stable Lévy processes' (2013, p. 444), but place less emphasis on the particular types of language that econophysicists deploy alongside their statistical approaches.

7 This connection between the systemic risks of high-frequency trading and its embeddedness in strategies in which the market is tweaked to benefit a few (in particular, to the detriment of human traders) is echoed in a recent *MIT Technology Review* article: 'High-frequency traders are able to make pennies off of individual trades but execute them millions of times a day, while regular investors are left in the dust. And it could be a destabilizing force, where software gone haywire erases huge chunks of a company's value in a matter of minutes. That has happened enough that it has a name: a flash crash' (Reilly, 2016).

References

- Andersen, T. G., & Bondarenko, O.** (2015). Assessing Measures of Order Flow Toxicity and Early Warning Signals for market Turbulence. *Review of Finance*, 19, 1–54.
- Arnoldi, J., & Borch, C.** (2007). Market Crowds between Imitation and Control. *Theory, Culture & Society*, 24(7–8), 164–80.
- Arvidsson, A.** (2016). Facebook and Finance: On the Social Logic of the Derivative. *Theory, Culture & Society*, 33(6), 3–23.

- Beunza, D., & Stark, D.** (2012). From dissonance to resonance: cognitive interdependence in quantitative finance. *Economy and Society*, 41(2), 383–417.
- BlackRock** (2010). Understanding the “Flash Crash”: What Happened, Why ETFs Were Affected, and How to Reduce the Risk of Another. Retrieved from <https://www.blackrock.com/corporate/ench/literature/whitepaper/understanding-the-flash-crash-nov-2010.pdf>.
- Borch, C.** (2007). Crowds and economic life: bringing an old figure back in. *Economy and Society*, 36(4), 549–73.
- Borch, C.** (2011). *Niklas Luhmann (Key Sociologists)*. London and New York: Routledge.
- Borch, C.** (2012). *The Politics of Crowds: An Alternative History of Sociology*. Cambridge: Cambridge University Press.
- Borch, C., Hansen, K. B., & Lange, A.-C.** (2015). Markets, bodies, and rhythms: A rhythmanalysis of financial markets from open-outcry trading to high-frequency trading. *Environment and Planning D: Society and Space*, 33(6), 1080–97.
- Borch, C., & Lange, A.-C.** (2016). High-frequency trader subectivity: Emotional attachment and discipline in an era of algorithms. *Socio-Economic Review*, doi: 10.1093/ser/mww013.

- Boulton, T. J., Braga-Alves, M. V., & Kulchania, M.** (2014). The Flash Crash: An examination of shareholder wealth and market quality. *Journal of Financial Intermediation*, 23(1), 140–56.
- Bullock, N, Mackenzie, M., & Guerrera, F.** (2010, May 7). Trading goes wild on Wall St. *Financial Times*, p. 8.
- Coombs, N.** (2016). What is an algorithm? Financial regulation in the era of high-frequency trading. *Economy and Society*, 45(2), 278–302.
- Corkery, M.** (2010, May 6). The Stock Market’s Flash Crash: How to Destroy \$1 Billion in 60 Minutes. *Wall Street Journal*. Retrieved from <http://blogs.wsj.com/deals/2010/05/06/the-stock-markets-flash-crash-how-to-destroy-1-billion-in-60-minutes/>.
- CFTC-SEC** (2010). *Findings Regarding the Market Events of May 6, 2010: Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues*. Washington, DC: The CFTC and SEC.
- DiMaggio, P.J., & Powell, W.W.** (1983). The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields. *American Sociological Review*, 48(2), 147–60.
- Durkheim, E.** (1947). *The Elementary Forms of the Religious Life: A Study in Religious Sociology*, trans. J. W. Swain. Glencoe, Illinois: The Free Press.

- Durkheim, E.** (1982). *The Rules of Sociological Method. And selected texts on sociology and its method*, trans. W. D. Halls. London: Macmillan.
- The Economist** (2015, April 25). Flash boy, p. 67.
- Foucault, M.** (1990). *The History of Sexuality, Vol. 1: An Introduction*, trans. Robert Hurley. New York: Vintage Books.
- Fromm, E.** (1941). *Escape from Freedom*. New York: Rinehart & Company.
- Golub, A., Keane, J., & Poon, S.-H.** (2012). High Frequency Trading and Mini Flash Crashes. Retrieved from <http://ssrn.com/abstract=2182097>.
- Golumbia, D.** (2013). High-frequency trading: Networks of wealth and the concentration of power. *Social Semiotics*, 23(2), 278–99.
- Grocer, S.** (2010, May 6). Six Mega Drops of the Flash Crash. *Wall Street Journal*. Retrieved from <http://blogs.wsj.com/deals/2010/05/06/four-mega-drops-of-the-flash-crash-sam-adams-goes-flat/>.
- Hansen, K. B.** (2015). Contrarian investment philosophy in the American stock market: On investment advice and the crowd conundrum. *Economy and Society*, 44(4), 616–38.
- Ivanov, S. I.** (2011). The Effects of Crisis on the Cointegration between the S&P 100 and the S&P 500 Indexes. *International Journal of Finance*, 23(2), 6783–97.

- Jovanovic, F. & Schinckus, C.** (2013). The Emergence of Econophysics: A New Approach in Modern Financial Theory. *History of Political Economy*, 45(3), 443–74.
- Johnson, N. et al.** (2012). Financial black swans driven by ultrafast machine ecology. Retrieved from <http://arxiv.org/abs/1202.1448>.
- Johnson, N. et al.** (2013). Abrupt rise of new machine ecology beyond human response time. *Scientific Reports*, 3, Article number: 2627.
- Karppi, T., & Crawford, K.** (2016). Social Media, Financial Algorithms and the Hack Crash. *Theory, Culture & Society*, 33(1), 73–92.
- Kenett, D. Y. et al.** (2013). How High Frequency Trading Affects a Market Index. *Scientific Reports*, 3, Article number: 2110.
- Kirilenko, A. & Lo, A. W.** (2013). Moore’s Law versus Murphy’s Law: Algorithmic Trading and its Discontents. *Journal of Economic Perspectives*, 27(2), 51–72.
- Kirilenko, A., Sowers, R., & Meng, X.** (2013). A multiscale model of high-frequency trading. *Algorithmic Finance*, 2, 59–98.
- Kirilenko, A. et al.** (2014). The Flash Crash: The Impact of High Frequency Trading on an Electronic Market. Retrieved from http://www.cftc.gov/idc/groups/public/@economicsanalysis/documents/file/oce_flashcrash0314.pdf.
- Knight, W.** (2012). Watch High-Speed Trading Bots Go Berserk. *MIT Technology Review*. Retrieved from

<https://www.technologyreview.com/s/428756/watch-high-speed-trading-bots-go-berserk/>.

Lange, A.-C. (2015). *Crowding of Adaptive Strategies: Swarm Theory and High-Frequency Trading*. Copenhagen: unpublished manuscript.

Lange, A.-C. (2016). Organizational ignorance: An ethnographic study of high-frequency trading. *Economy and Society*, 45(2), 230–50.

Lange, A.-C., Lenglet, M., & Seyfert, R. (2016). Cultures of high-frequency trading: Mapping the landscape of algorithmic developments in contemporary financial markets. *Economy and Society*, 45(2), 149–65.

Leal, S. J. et al. (2014). Rock around the Clock: An Agent-Based Model of Low- and High-Frequency Trading. Retrieved from <http://arxiv.org/abs/1402.2046>.

Lee, B. & Martin R. (eds) (2016). *Derivatives and the Wealth of Societies*. Chicago: University of Chicago Press.

Lewis, M. (2014). *Flash Boys: Cracking the Money Code*. London: Allen Lane.

Luhmann, N. (2012). *Theory of Society*, trans. R. Barrett. Stanford, California: Stanford University Press.

Mackay, C. (2001). *Extraordinary Popular Delusions and the Madness of Crowds*. New York: MetroBooks.

- MacKenzie, D.** (2014). *A Sociology of Algorithms: High-Frequency Trading and the Shaping of Markets*. Edinburgh: unpublished manuscript.
- MacKenzie, D.** (2015). Mechanizing the Merc: The Chicago Mercantile Exchange and the Rise of High-Frequency Trading. *Technology and Culture*, 56(3), 646–75.
- MacKenzie, D.** (forthcoming). How Algorithms Interact: Goffman’s “Interaction Order” in Automated Trading. *Theory, Culture & Society*.
- Malmgren, H., & Stys, M.** (2011). Computerized Global Trading 24/6. *The International Economy*, Spring, 30–32, 62.
- McInish, T., Upson, J., & Wood, R. A.** (2014). The Flash Crash: Trading Aggressiveness, Liquidity Supply, and the Impact of Intermarket Sweep Orders. *Financial Review*, 49(3), 481–509.
- Melloy, J.** (2010, May 6). Traders See Black Hole Effect as Sell-off Feeds on Itself. *CNBC*. Retrieved from <http://www.cnbc.com/id/36997706>.
- Moosa, I.** (2015). The regulation of high-frequency trading: A pragmatic view. *Journal of Banking Regulation*, 16(1), 72–88.
- Peckham, R.** (2013). Economies of contagion: Financial crisis and pandemic. *Economy and Society*, 42(2), 226–48.
- Phillips, M.** (2010a, May 6). Procter & Gamble: Something Strange Was Going On. *Wall Street Journal*. Retrieved from

<http://blogs.wsj.com/marketbeat/2010/05/06/procter-gamble-something-strange-was-going-on/>.

Phillips, M. (2010b, May 6). Stock Market Fear Gauge up 60%, and Staying There. *Wall Street Journal*. Retrieved from <http://blogs.wsj.com/marketbeat/2010/05/06/stock-market-fear-gauge-up-60-and-staying-there/>.

Phillips, M. (2010c, May 6). Stock Tumble: What Have You Heard Explain the Selloff? *Wall Street Journal*. Retrieved from <http://blogs.wsj.com.escweb.lib.cbs.dk/marketbeat/2010/05/06/stocktumblewhathaveyouheardexplaintheselloff/>.

Phillips, M. (2010d, May 6). Behold, The Flash Crash! *Wall Street Journal*. Retrieved from <http://blogs.wsj.com/marketbeat/2010/05/06/behold-the-flash-crash/>.

Reilly, M. (2016). High-Frequency Trading is Nearing the Ultimate Speed Limit. *MIT Technology Review*. Retrieved from <https://www.technologyreview.com/s/602135/high-frequency-trading-is-nearing-the-ultimate-speed-limit/>.

Roitman, J. (2014). *Anti-Crisis*. Durham and London: Duke University Press.

Roll, R. (1988). The International Crash of October 1987. *Financial Analysts Journal*, September–October, 19–35.

- Schapiro, M. L.** (2010). Strengthening Our Equity Market Structure.
Retrieved from
<https://www.sec.gov/news/speech/2010/spch090710mls.htm>.
- Scharfstein, D. S., & Stein, J. C.** (1990). Herd Behavior and Investment.
American Economic Review, 80, 465–79.
- Seyfert, R.** (2016). Bug, predations or manipulations? Incompatible
epistemic regimes of high-frequency trading. *Economy and Society*,
45(2), 251–77.
- Sidis, B.** (1898). *The Psychology of Suggestion: A Research into the
Subconscious Nature of Man and Society*. New York and London:
D. Appleton and Company.
- Simmel, G.** (1950). On the Significance of Numbers for Social Life. In K.
H. Wolff (ed.), *The Sociology of Georg Simmel* (pp. 87–104). New
York: The Free Press.
- Snider, L.** (2014). Interrogating the Algorithm: Debt, Derivatives and the
Social Reconstruction of Stock Market Trading. *Critical Sociology*,
40(5), 747–61.
- Sornette, D., & von der Becke, S.** (2011). *Crashes and High Frequency
Trading: An Evaluation of Risks Posed by High-Speed Algorithmic
Trading*. Zurich: Swiss Finance Institute, Research Paper Series,
Nos. 11–63.

- Stäheli, U.** (2006). Market Crowds. In Schnapp, J. & Tiew, M. (eds.), *Crowds* (pp. 271–87). Stanford, California: Stanford University Press.
- Stäheli, U.** (2013). *Spectacular Speculation: Thrills, the Economy, and Popular Discourse*. Stanford, California: Stanford University Press.
- Taleb, N. N.** (2010). *The Black Swan: The Impact of the Highly Improbable*, revised ed. London: Penguin Books.
- Tarde, G.** (1962). *The Laws of Imitation*, trans. E. C. Parsons. Gloucester, Mass.: Peter Smith.
- Tarde, G.** (1968). *Penal Philosophy*, trans. R. Howell. Montclair, New Jersey: Patterson Smith.
- Thompson, G.** (2016). Time, trading and algorithms in financial sector security. *New Political Economy*, <http://dx.doi.org/10.1080/13563467.2016.1183116>.
- Wagner-Pacifici, R.** (2010). Theorizing the Restlessness of Events. *American Journal of Sociology*, 115(5), 1351–86.
- Womack, B.** (2010, May 7). Yahoo, Google, Schwab Sites Slowed Amid Market Dive. *Bloomberg*. Retrieved from <http://www.chinapost.com.tw/business/company-focus/2010/05/08/255680/Yahoo-Google.htm>.

Christian Borch is Professor of Sociology at the Department of Management, Politics and Philosophy, Copenhagen Business School, Denmark. His research interests include crowd theory, financial markets and urban sociology. His current work focuses on algorithmic finance. Christian is the author/editor of several books, including *Foucault, Crime and Power: Problematisations of Crime in the Twentieth Century* (Routledge, 2015); *Urban Commons: Rethinking the City* (ed. with Martin Kornberger, Routledge, 2015); *Architectural Atmospheres: On the Experience and Politics of Architecture* (Birkhäuser, 2014); and *The Politics of Crowds: An Alternative History of Sociology* (Cambridge University Press, 2012), which was awarded the 2014 Theory Book Prize by the American Sociological Association. He is currently working on a book entitled *Avalanche Modernity: Crowds, Cities, Markets*.