Contagious Markets: On Crowd Psychology and High-Frequency Trading

Ann-Christina Lange and Christian Borch
Department of Management, Politics and Philosophy,
Copenhagen Business School, Denmark
ala.mpp@cbs.dk; cbo.mpp@cbs.dk

25 September 2014

Abstract: This paper presents a sociological study of how contemporary high-frequency trading (HFT) practices and strategies are organized in response to alleged crowd dynamics in financial markets and how they nevertheless, and somewhat paradoxically given this ambition, seem to reinforce such crowd dynamics. Drawing on qualitative data, including interviews with and ethnographic observations of HFT traders, the paper lends support to recent agent-based models (ABM) that characterize HFT as prone to exhibit contagion and negative crowd effects. Specifically, the paper supplements ABM accounts, which are often based on assumptions derived from theory, by providing empirically grounded insights into the daily practices of HFT traders and the designs of their black boxes. Furthermore, the paper argues that the references to contagion and herding that can be found in much financial economics on HFT evokes tropes from late-nineteenth-century crowd psychology that appear more apt than realized in financial economics. Thus, not only are HFT black boxes indeed designed in ways that place contagion and adaptive feedback loops centrally. HFT traders interact with their algorithms in ways that seek to produce psychological detachment from market crowds.

Keywords: agent-based modelling; black box; crowd psychology; high-frequency trading; sociology

Introduction[1]

Recent technological developments within the financial sector have sparked controversy on the merits and dangers of automated computer-based high-speed trading technologies, so-called high-frequency trading (HFT). In HFT sophisticated technological tools and computer algorithms are used to analyze

[1] Research for this paper was supported by a 'Crowd Dynamics in Financial Markets' Sapere Aude Grant from the Danish Council for Independent Research.

multiple markets and execute orders based on real-time market conditions. Proprietary trading strategies are carried out in order to move in and out of positions in fractions of a second (down to nanoseconds, which is one billionth of a second (10^{-9})). They do this many times a day to earn a single tick (and possible exchange rebates).

Neil Johnson et al. have argued that HFT marks a transition from a human–machine phase to a new 'all-machine phase' (Johnson et al. 2012: 5). What we see is a transition from the trading floor, to the trader in front of the screen (Knorr Cetina 2009) to pure machine-based types of trading. And it is a type of trading that constitutes a substantial part of the trading volume. According to TABB Group (a financial vendor and research firm), HFT provided 53% of the trading volume in mature markets (UK+US) in 2011 (see Cui and Lauricella 2011).[2] The significance and proliferation of HFT have caused heated debate in public media and among regulators and financial institutions about its consequences for financial stability and future regulation. The publication of Michael Lewis' book *The Flash Boys* (2014) especially contributed to the rise of critical voices against issues related to HFT such as front-running and quote stuffing. However, years before the release of Lewis' book HFT has been investigated by regulators (such as the CFTC and SEC), although rarely in ways that have attracted as much public attention as is presently the case.

Partly accounting for the heated debates over HFT is its relation to the Flash Crash of May 6, 2010 where more than one trillion dollars evaporated within a few minutes. The Dow Jones Industrial Average plummeted by more than 600 points (approx. 5% of its total value) in a matter of minutes. According to the UK Foresight report published by the UK Government Office for Science, what happened during May 6, 2010 was that 'a fundamental trader's algorithm started selling based on previous trade volume, creating a positive feedback between its own selling and the trading activity of other market participants' (Sornette and von der Becke 2011: 6). More specifically, a mutual fund attempted to sell a large portion of E-mini S&P 500 contracts (CFTC and SEC Staff Report 2010: 15). This triggered a negative feedback loop where HFT programs attempted to sell at lower and lower prices to minimize short-term losses. This drove down the price of the E-mini 3%, which in turn spilled over to the equities markets. The negative trends continued until computer systems paused trading, effecting a rapid recovery when trading resumed.

Some observers have likened the flash crash to an unwanted domino effect where preprogrammed algorithms triggered other preprogrammed algorithms.

[2] There has been a decrease in HFT in equities since 2009.

Sornette and von der Becke describe the incidence as a 'self-excitation component' or 'viral epidemic' (2011: 11), an argument which is supported by the Nanex Report on the May 6 Flash Crash (http://www.nanex.net/FlashCrash Final/FlashCrashSummary_II.html). Even more disturbing, HFT is considered by some to contribute to similar 'mini' crashes on a continuous basis, thereby representing an inherently destabilizing factor in financial markets (Golub, Keane and Poon 2012). Furthermore, Ben-David, Franzoni and Moussawi (2012) argue that HFT strategies may induce contagion due to the correlated nature of securities and accelerated time scales and that such mini crashes continuously occur in single stocks.[3]

What is interesting about such accounts about the unintended consequences, negative externalities, feedback loops, technological dangers, and regulatory challenges allegedly involved with the use of HFT strategies is how they resonate with an old vocabulary typically associated with late-nineteenth-century crowd psychology. Thus, the HFT reconfiguration of financial markets is said in much financial economics to exhibit a number of traits usually associated with crowd behavior such as, in particular, herding and contagious processes that play out in non-intentional ways, i.e. are not driven by deliberate, intentional individuals and their decisions, but instead assume self-organizing, emergent, and seemingly irrational properties.

Identifying a connection between crowd psychology and HFT may appear surprising, given that HFT is generally conceived as fully automated trading, hence leaving hardly any room for individuals and their psychological motives and dispositions. Yet, as we shall return to below, HFT does not necessarily rule out psychological aspects. Specifically, we will argue, on the level of how HFT traders interact with their algorithms, psychological elements are manifest, in particular with respect to ensuring emotional detachment from markets. This also means that establishing a connection between current financial economics on HFT, on the one hand, and crowd psychology, on the other, is not simply a matter of identifying a curious parallel between two modes of thought. By bringing in the vocabulary of crowd psychology we claim to be able to account in fuller detail for the present realities of HFT, for example, with respect to how individual HFT traders relate to market crowds.

In this paper we establish connections between crowd psychology and HFT by shedding light on actual HFT strategies as they are deployed in practice. Consequently, in contrast to much existing work on HFT which, as is the case of the above-cited references, tend to account for HFT on the basis of an

[3] See http://www.nanex.net/FlashCrashEquities/FlashCrashAnalysis_Equities.html for a graphic illustration.

extrapolation from the May 6, 2010 flash crash (something that is likely to address only a few aspects of the wider HFT reality) or employ e.g. agent-based modelling (ABM) to forecast HFT developments (an approach which is, of course, only as good as its assumptions, which tend to be theory-driven rather than empirically derived), we suggest a different take where emphasis is on how HFT firms and traders conceive of their strategies on a day-to-day basis. More specifically, this paper is a sociologically informed study of HFT where our primary data consist of ethnographic observations and interviews conducted inside a Wall Street HFT firm by one of the authors in spring 2014, as well as of 50 interviews with a broad range of actors involved with HFT, including programmers, software developers, broker-dealers, exchange officials, investment bankers, and regulators (conducted in Copenhagen, London, and New York since October 2013).

The ethnographic work focused on the daily practices and conversations amongst HFT traders, including how traders and programmers trade at their desk while monitoring preprogrammed algorithms, but the ethnographic work also followed their activities around designing and building black-box market automata. The focus was on HFT in Treasury bond futures and index futures. Our data offer insights into the ways in which traders reflect upon their own and other markets participants' trading behavior.[4] We trace if they refer explicitly or implicitly to crowd thinking in the ways they reflexively constitute an account of their own practice and trading strategies (Garfinkel 1967).

What this 'inside look' offers is, we claim, a more complete understanding of how crowd dynamics may be playing out within HFT. So, on the one hand, we are able to lend support to e.g. Sornette and von der Becke's ABM account of how the negative feedback loops of adaptive strategies may be a risk inherent to HFT – only, we arrive at this conclusion from the *inside*, namely on the basis of how HFT traders design their black boxes rather than on the basis of abstract modelling. On the other hand, we can address interactions between traders and algorithms not accounted for by ABM, and can demonstrate that in such interactions classical (pre-HFT) challenges about emotional detachment from market crowds continue to exist.

This way of studying financial markets on the basis of interviews and ethnographic observations follows in the footsteps of a range of recent sociological studies of financial markets, including HFT (although the latter is still a little-researched field within sociology, a key expection being MacKenzie, 2014). One aspect that has been central to such sociological studies – in contrast

[4] The interviews were semi-structured and open, maintaining a sensibility for the situations in which they took take place (Emerson, Fretz and Shaw 2001).

to, say, financial economics – is how particular assumptions about markets may have far-ranging practical implications for financial markets when deployed by market participants (e.g. MacKenzie, 2006; MacKenzie and Millo, 2003). In other words, when market participants such as traders start using particular vocabularies or modes of thought, or when they begin using specific models or formulas to guide their work, then these forms of thinking or modelling assume a real character that may have real effects on the markets – even if the scientific status of these models or modes of thought might be highly contested. What we propose in the present paper is that something similar applies to crowd psychology: while this theoretical tradition has been widely contested by sociologists and social psychologists (for an overview, see Borch 2012), some of its basic ideas are being performed and thus have real effects on HFT.

The paper is structured as follows. In the first part we illustrate how recent financial economics attributes to HFT inherent risks relating to the crowding and negative contagion of adaptive strategies. We further show how this way of analyzing HFT evokes particular tropes about non-intentional contagion and imitation as found in late-nineteenth-century crowd psychology. In the second part of the paper, we turn to our empirical data and demonstrate how actual black box strategies are designed in ways that focus on adaptive feedback loops. In the third part, we discuss how HFT traders interact with their black box systems in ways that seek to produce psychological detachment from market crowds. In the conclusion we discuss some regulatory implications of our analysis.

Non-Intentional Contagion and Imitation: Crowd Psychological Inspirations

In a seminal paper on the possible relations between 'Crashes and High Frequency Trading', prepared as part of the UK Foresight report on 'The Future of Computer Trading in Financial Markets', Sornette and von der Becke (2011) provide a detailed analysis of HFT and its alleged structural propensity to produce negative herding in and across markets. Key to Sornette and von der Becke is the argument that HFT stimulates 'the crowding of adaptive strategies that are pro-cyclical' (2011: 3). According to the authors, '[a]s HFT use short-term information as well as adaptive algorithms, there is potential for herding as the strategies can crowd to the same signal, synchronize and lead to transient large instabilities' (2011: 12). Indeed, commenting explicitly on the flash crash of May 6, 2010 Sornette and von der Becke suggest that HFT renders negative spill-over effects more likely:

> it could be possible that deeper markets are more prone to pandemics as their impact and connection to other markets is larger. […] One reason is that the large number of participants can herd and therefore form large destabilizing crowds, whose size may be comparable to the global market size. (2011: 7, see also 16)

In order to better conceive how such processes may unfold Sornette and von der Becke propose ABM as a means of understanding possible dynamic interaction effects between markets. Briefly put, ABM 'can be used to group various market participants, assign behavioural preferences (for example, short-term systematic vs. long-term fundamental trading) and simulate their behaviour over time', taking into account feedback mechanisms centered on 'adaptation and imitation' (2011: 14).

For present purposes, two things are particularly important about Sornette and von der Becke's approach. First, their general diagnosis of HFT is one that attributes special importance to contagious dynamics that can spiral out of control (cf. the references to pandemics); as well as to the notion that this is driven by patterns of imitation and adaptation, and that the entire process is unforeseen, self-organizing, and emergent, meaning that no individual actors can sensibly be held accountable for what are truly systemic properties and risks. This is a diagnosis that resonates with findings from other parts of financial economics, including analyses that do not focus explicitly on HFT, but contribute for example to discussions on predatory trading and contagion (Brunnermeier and Pedersen 2005). A key similarity between these strands of research is how their analytical starting point is individual traders pursuing their individual (sub-optimal) strategies, and how this behavior then leads to negative contagious processes.

Second, Sornette and von der Becke's approach presents a global view of the market, as it were, in that they are interested in the overall market effects of HFT. This is different from other analyses of HFT and contagion. For example, Easley, López de Prado and O'Hara define and examine 'contagion as the natural consequence of market makers revising their orders in one market in response to changing liquidity conditions in *related markets*' (2013: 146, italics added). The example offered in their analysis is the contagion dynamics of order flow between heating oil and gasoline. In Easley, López de Prado and O'Hara's analysis, contagion is not so much about non-intended, self-organizing, and emergent processes. It is rather a notion deployed to understand how interrelated products are indeed related. Specifically, they argue, such contagion dynamics of order flow occur as a result of 'market makers hedging their risk of adverse selection in one instrument by taking liquidity in another, with the hope that over time they will be able to unwind their position in both instruments at a profit' (2013: 146). For some, describing this as contagion might appear slightly misleading, as it essentially refers to a (mere) structural correlation. In order, therefore, to be able to distinguish between this form of contagion and the notion of contagion analyzed by Sornette and von der Becke, we suggest that the former be characterized as a *weak* form of contagion and the latter as a *strong* form of contagion.

In our ethnographically informed analysis below, we will demonstrate how HFT seems to be implicated in both types of contagion, i.e. both weak and strong forms. For now, however, we wish to attend to an aspect of contemporary financial economics which has escaped systematic analysis, namely that approaches on HFT contagion such as the ones cited above share important affinities to tropes from crowd psychology. Such tropes have for some time infiltrated the field of behavioral finance, with Nobel Laureate Robert J. Shiller being the key exponent for taking seriously crowd and mass psychology when analyzing financial markets (for a seminal illustration of this argument, see especially Shiller 1984). But as indicated, these tropes are also, and increasingly so, entering debates about HFT.

While the reference to behavioral finance is illustrative of the endeavor to deploy crowd psychology and notions of contagion in the field of finance, it easily obscures the fact that, actually, there is a long (pre-behavioral finance) history of associating financial markets with contagion. This is a point also made by Peta Mitchell (2012: 131–2) who draws attention to how John Stuart Mill theorized contagious markets back in the 1840s, and how such ideas later resurfaced in, for example, William F. Taussig's *Principles of Economics* (Taussig, 1911). Notions of contagion were linked explicitly to crowd psychology especially in the early twentieth century, i.e. at a point when crowd psychology had been established as a separate sub-discipline.

The field of crowd psychology emerged in France in the late-nineteenth century as a response to rapid industrialization and urbanization as well as, in particular, the long series of uprisings that had characterized the country ever since the French Revolution, reaching, according to many contemporaneous observers, a culmination with the Paris Commune in 1871 (see Borch, 2012). The general experience that crowd psychology seemed to address was the transformation that allegedly takes place when people become part of a crowd: left on their own, individuals usually behave in ways that comply with society's norms; however, as soon as they become part of a crowd they engage in all sorts of actions they would never resort to otherwise (such as violence, destruction, and unlawful behavior more generally). The central point for now is not whether this was a fitting diagnosis of crowd behavior (indeed, many subsequent observers would question the allegedly inherent criminal dispositions of crowds). The more interesting thing to note is how crowd psychologists explained crowd behavior, and how this explanatory framework soon travelled to the field of financial markets.

Beginning with the former aspect, crowd psychologists such as Gustave Le Bon and Gabriel Tarde argued that what takes place in crowds is a kind of mental transformation. The crowd is not a physical entity, but rather the name given to a particular (dis)organizing principle where the individual crowd members suspend their singular thoughts, reflections, and deliberations and instead become subsumed under a kind of 'mental unity', as Le Bon put it (2002: 2). In such a situation, crowd psychologists argued, contagion becomes central to how sociality unfolds: what happens one place in the crowd transmits instantaneously to the rest of the crowd. The central explanation offered by crowd psychologists to account for this state of non-friction was that of hypnotic suggestion. That is, crowd psychologists likened the behavior that takes place in crowds to the relationship between a hypnotist and a hypnotized: the latter is not aware that he or she imitates the former; demands are followed non-intentionally without any friction.

This state of hypnotism provides a sense of the image of the crowd individual depicted by many early crowd psychologists. They essentially described the members of a crowd as de-individualized, non-intentional 'automatons', who slavishly submit to whatever the leader of the crowd demands (Le Bon 2002: 8). Furthermore, given his or her de-individualized state of being, the crowd member was said to be reduced to a person who reacts impulsively and in a highly affective manner, which explains why crowd psychologists would typically characterize crowd behavior as inherently irrational.

Putting this together, crowd psychology in effect sought to account for a particular form of (abnormal) behavior, characterized by intense imitative, affectively charged contagion, which basically rendered crowd members into automatons, momentarily deprived of any sense of conscious, rational, deliberate thought. Another way of saying this is that crowd psychology essentially pitted the individual against the crowd: the entire explanatory edifice of crowd psychology was based on a stark opposition between individual and crowd. To repeat, the ideas was that, once you become part of a crowd, your conscious faculties are suspended, and you blindly imitate whatever the leader says or does.

Our point of bringing attention to this tradition of crowd psychology is not to revive it as a kind of explanatory horizon in a present-day context, in the sense of arguing that scholars and regulators need to hark back to this tradition in order to explain how HFT really works (stating, for example, that HFT black boxes are inherently irrational). Instead, our twofold point is, first, to argue that, when much current financial economics associates HFT with contagion, it does so in order to emphasize 'a departure from the normal, the expected, or the rational' (Kolb 2011: 5). And this emphasis on the abnormal, unexpected, and irrational or non-rational is exactly what motivated early-twentieth century attempts to adopt crowd psychology to the field of finance. Thus, second, crowd psychology brought to the fore a particular tension or challenge that early-twentieth-century observes of financial markets sought to utilize, namely the tension between individual and crowd. Indeed, what would become known as contrarian speculation theory, emerging in USA in the 1920s, was from the very outset an attempt to utilize crowd psychology as an integrated means of a broader investment strategy (for discussions of this tradition, see Borch 2007; Stäheli 2006; 2013). Basically arguing that financial markets display the same forms of irrationality that crowd psychologists would associate with crowds, contrarian speculators asserted that in order to beat the market, the clever investor would need to speculate contrary to the irrational market crowd. Interestingly, this entailed among many other things that the investor should keep some distance to the market in order not to be subsumed under its hypnotic spell. Put differently, the clever investor would have to discipline him- or herself so as not to be seized by the affective contagion of market imitations. Interestingly, as we shall return to later on, this problem of market detachment, as it were, remains of crucial importance also to HFT, so in that sense some important continuity can be detected between present-day HFT and early-twentieth-century financial adoptions of crowd psychology.

The more general point we wish to make is that, as is hopefully clear from this brief overview, when contagious market dynamics are referred to in much contemporary financial economics, this – unwittingly or not – essentially recaps older attempts to understand, and respond to, alleged crowd dynamics in financial markets. What resonates between early-twentieth-century contrarian speculation accounts and present-day analyses of financial markets is the notion of a market prone to run amok in non-intentional, irrational ways where rational individual calculation and decision-making are replaced by blind imitation.

Informational Feedback Loops and Adaptive Strategies

We will now leave aside for a moment the more historical-conceptual relations between financial economics on HFT, on the one side, and crowd psychology, on the other, and move to our empirical analysis of present-day HFT strategies.

The aim of this analysis is to show how such HFT strategies are designed in ways that place contagion centrally. More specifically, the present section demonstrates that HFT strategies can be delineated according to the distinction we established in the previous section between weak and strong forms of contagion.

To begin with the former, we may recall how for Easley, López de Prado and O'Hara, contagion refers to a process where 'market makers [are] hedging their risk of adverse selection in one instrument by taking liquidity in another' (2013: 146). Here, contagion is employed in a classical diagnostic sense, namely to diagnose a particular type of process within HFT. Our interviews suggest, however, that HFT traders describe similar types of behavior in *strategic* rather than diagnostic terms. As one HFT trader put it, 'we profit from correlation and hedge ourselves. We exploit securities that move in sync due to them being tightly hedged'. Similarly, a programmer from a research firm specializing in HFT stated that 'what [HFT traders] do is to empirically measure the correlation between securities. Virtually every pair of securities in the market has a positive correlation'. The correlation that some traders profit from was, according to these traders, created by large institutional investors' tendency to over-hedge their positions in order to reduce their risk. This was, from an HFT-trader perspective, seen as an irrational or emotional response to risk calculation creating a price impact.[5] Put differently, it was said to be an inefficiency or noise that HFT could exploit. In the words of a CEO of a small HFT firm in New Jersey:

> people are in the business of propagating that price impact to other securities […] So what we are doing, basically, is transferring the price impact of one security to a large set of other securities. That's where liquidity comes from, we're sourcing liquidity from other securities and transferring them to a specific future contract and then we're taking the price impact from that future and spreading it to other securities.

What this suggests is that weak contagion (in the form of a mere structural correlation) is not simply a feature that can be identified diagnostically in HFT, it is rather a fundamental *condition* for HFT strategies: it is something HFT acts upon and exploits. Yet this is not the whole story, for HFT is also more directly implicated in contagion, in the sense of (potentially) fostering it. Indeed, actively

[5] A price impact occurs in financial markets when an investor issues a large buy or sell order and other traders start to buy or sell on the assumption that that investor has some fundamental knowledge about the specific asset. This, in the view of the HFT traders we talked to, skews the 'fair' price of that specific asset.

exploiting weak contagion also further reinforces it. This may materialize in various ways. One is described by a trader who designs his algorithms to exploit contagion effects:

> What you do [in one HFT strategy] is making markets. So you are offering and bidding competitively on one exchange. That way when someone pays the spread, when someone buys the offer or sells the bid, they are first to know because they got filled. If they are part of that sell or buy, they find out immediately and that gives them the time-jump to go on to the next exchange and if they sold they can buy on that exchange and make profit on the difference.

So, here, HFT traders act upon an assumption of contagious order flow and at the same time participate in the resulting price move. They do so by constantly issuing and cancelling orders to be in front of the price move that they aim to profit from. Another trader, also acting CEO of a major HFT firm in Chicago, described a similar strategy:

> The fact that I am participating on the market gives me time to speed-jump because the information were a fill and that preempts market data significantly […] and when you receive that fill, that's what triggers your hedge orders essentially, to these other exchanges.

In other words, taking active part in contagion dynamics here becomes a strategy that itself produces further contagion.

Propensities to stronger contagion (relating to adaptive strategies) may be identified in the designs of HFT black boxes which seek to utilize controlled feedback loops to exploit contagion by detecting order-flow information and at the same time leave the traders' strategies unaffected by directional market moves. Taking a look inside one trader's black box we find that it is composed by various blocks of software connected to the trading venues' matching engines, different types of algorithms, and a dozen different strategies. The different blocks of software and their connections are illustrated in Figure 1.

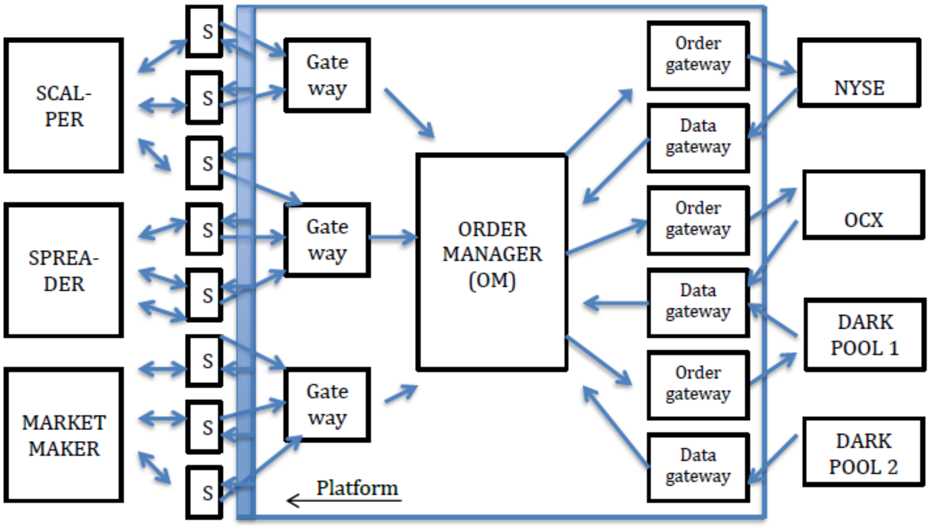

Figure 1: HFT black box

The trader designing this black box described an order manager (OM) as sitting on top of the system.[6] The OM is responsible for loading various strategies (S) – quotes go in and positions come out. The gateways translate quotes to the order books, which show bids and offers. The platform can dynamically load and unload strategies during the day.[7] Each of these strategies is loaded to the platform as configuration files (xml text files). A number of numeric parameters are defined in the xml files. Here, the trader enters the limit (the volume to be traded) as well as and the unit time (the specific time-scale that determines how long the quotes stay in the order books before they are automatically removed).

[6] In other HFT firms the order manager was referred to as the 'strategy engine' or the 'price engine'. This is an example of a broader point, namely that HFT is far from a fully standardized field. Similar terms can have different meanings and similar functions can be described using different terminologies across the field. This non-standardization makes it difficult to extrapolate the observations presented in this article to the wider field of HFT and algorithmic trading. However, in order to ensure some generalizability the traders' statements were in all cases contextualized through complementary interviews conducted with professionals working in the field of HFT (experts, analysts, other HFT traders, and HFT managers). Doing so the notion transpired that HFT profit lies less in a preprogrammed strategy or highly complex mathematical models, and more in optimizing order execution with the use of sophisticated technology. That is, traders build their own black box system to optimize execution time.

[7] The platform refers to an interface (one of the software blocks) between the OM and the trader sitting in front of the screen. The trader would load his codes in an xml file for the platform to load it to the strategies.

The algorithms composed by the system are divided into three basic types. The first type is called a spreader. It buys one instrument and sells another with as little internal latency as possible. For instance, the algorithm buys government bonds traded at the New York Stock Exchange (NYSE) and futures traded at the Security Futures Exchange in Chicago (OCX). There is a 13 millisecond delay in the transmission of data from New York City to Chicago. This delay creates arbitrage opportunities of exploiting the price discrepancies between U.S. Treasury bonds traded at the NYSE and futures traded at the OCX. When the price of a government bond on the NYSE and its corresponding futures contract at OCX are out of sync, the algorithm would buy the less expensive one and sell it on the more expensive market.

The second type of algorithm is a scalper. This type of algorithm earns minimum incremental profits in a single instrument by buying and selling that same instrument many times a day across different trading venues. As one trader explained, 'we take advantage of the noisy motions on instruments where you'll have price fluctuations that are not linked to any meaningful information, and in that case you know you can profit from that noise'. This involves reading the depth of the order book and taking advantage of the probability that 'there's a large resting size at a certain level'. The third type of algorithm is a market maker which seeks to quote bids and offers in the same instrument and makes the market buys and sells according to certain basic rules to control the risk in the same way that a scalper seeks to take advantage of noise in a single instrument.

The OM covers a summarized abstract version of all the variable inputs needed to take advantage of real-time price movements (determined by structural correlation, noise, and adverse selection). Different types of contagion may be said to be implicated by the black box design. On one level, the traders in general described that their algorithms would 'mimic' or try to understand crowd behavior in order to profit from it. The trader designing the above-mentioned black box explained that 'all of the algorithms I have described to you to this point have been counter-trend. They profit not from trending behavior but from reversion behavior. […] Purely trending markets will be a loss-scenario for these kinds of algorithms'. The central point is that crowding and contagion are taken as the condition for specific HFT strategies, but not necessarily in the weak form previously described (i.e. structural correlations between related instruments). The kind of counter-herding strategy described here would respond to principally any kind of market herding.

On another level, the black box is designed in a fashion where adaptive feedback loops (linked to contagion by Sornette and von der Becke) are placed centrally. According to one HFT trader:

> The strategy will determine from the position change how to react, sends out another quote, that is a feedback loop […] and then there are also feedback loops between the order manager and the exchange. You can see the larger feedback loops encompass smaller feedback loops between the order manager and the exchange. So basically the whole thing is controlled feedback loops.

In other words, the black box is composed of feedback loops where one order is conditioned by the previous ones and by the ways in which the algorithms adapted to that information. A key facet of this feedback modelling is speed, or perhaps rather relative speed. Thus, one trader explained that an important aspect of his price discovery is to measure the latency in the order book (the speed by which orders from buyers and sellers are matched), which determines the probability of HFT market participation. The lower this is, the more probable it is that other HFT agents participate. In that sense, the black box demonstrates a probabilistic system that connects and seeks to determine the behavior of one algorithm to the behavior of others.

Similarly, sophisticated tools are built to detect market moves initiated by 'real money' (i.e. institutional investors and banks) in order to act upon or counteract expected price moves. A programmer explained his activity as 'seeing if there are other people obscuring the signal, i.e. the edge that you are trying to capture, and part of that is doing constant market recognos'. Another trader offered a specific example of this kind of market recognos the purpose of which is to detect the rhythms in buying and selling interests that the rest of the market is not aware of or does not know about:

> The shop that I started trading at, first thing they did – you know, I came from an automation background – was that they introduced me to markets and they immediately said, 'we know that banks are using iceberg orders', you know, hidden size, and they wanted to be able to detect the hidden size, because they are market makers and hidden size changes the typology of the market in ways that they can't readily identify. So the first thing I did when I entered this business was to build an iceberg detector. And that is very much that kind of recognos where you're looking for patterns that indicate other high-frequency or micro-structure activity and base decisions on that.

Here, the key strategic point of HFT is to move faster than the crowd effect spreads, i.e. to beat the speed by which contagion unfolds. In fact, the trader building the iceberg detector emphasized market crowds directly in reflecting on the effects of making this device:

> these algorithms are trying very hard to mimic or understand crowd mentality and profit from it. What it is that drives the market or market motions is identical to the crowd mentality that came out of the medieval period and even before. You know, nothing that we are doing now is new.

So, not only is the algorithm's next move determined by its previous results, but also by the behavior of other algorithms within and outside the black box system. One trader explained that his system is set up to accommodate for negative feedback loops between unrelated algorithms: 'one algorithm will quote, the other will respond the quoting, the other one will change its quote and they'll flip it back and forth'.

This feedback relation between interactive algorithms does not only apply to structurally correlated instruments but also to the interactional algorithmic responses between the exchanges' matching algorithms, institutional investors' and broker-dealers' executing algorithms and other HFT algorithms. What this suggests, in more general terms, is that the black box optimizes microstructural mechanisms that spread through informed order flow. This may take place across related instruments (as in the case of weak contagion), but the design of the black box in fact also lends support to Sornette and von der Becke's more abstract ABM approach to HFT, in that it shows how adaptive feedback loops (strong contagion) are an integral part of black boxes. With numerous market participants, each operating their own feedback-modelled black boxes, our inside view into HFT suggests that contagious crowding may indeed be reinforced on a more systemic level by the inter-adaptive nature of HFT black boxes.

Human–Algo Interaction: Ensuring Detachment from Market Crowds

The previous section described a generic, yet specific model of a HFT black box. A futures trader who has managed traders and desks across different asset classes (futures and equities) defined the four generic components of a trading system as data-in (price feeds or other data such as news, etc.), an order manager, gateways, and human monitoring. In this section we look further into the latter component.

The human aspect of HFT deserves analysis for several reasons. One is that for understandable reasons the majority of the existing research on HFT has focused on the technological aspects of algorithmic trading (including the technological arms race associated with HFT), but at the expense of studying what the HFT configuration might entail for the people operating the systems. This is not to suggest that the fetishization, as it were, of the technological aspects of HFT should be replaced by the notion that in the end it is always human beings who program the algorithms and that to understand HFT these individuals should therefore occupy a central analytical status. Far from it. What we wish to point out is rather that our empirical research suggests that the interactions between human traders and their HFT algorithms are both complex and interesting. Indeed, and this is a second reason why this interaction is worthwhile exploring, it appears that HFT traders are preoccupied with potential market contagion in ways that strikingly resemble early-twentieth-century contrarian concerns about being subsumed under irrational market crowds. Even more to the point, similar to contrarian speculation theory, developed in an era of open-outcry trading, HFT traders interact with their algorithms in ways that place the relation between the individual trader and the market crowd center stage: what should certainly be avoided is that one's emotions determine one's trading behavior. For this reason, particular forms of emotional market detachment are continuously exercised by HFT traders.

A starting point for this discussion is the widespread notion that HFT and other forms of algorithmic trading derive part of their legitimation from their alleged superiority over human individuals when it comes to executing orders efficiently. While humans are portrayed as fickle and easily captured by emotions and crowd following, algorithms are seen as a means to ensure rationality in markets. Indeed, it is a common assumption among HFT traders that machines can avoid the errors humans make, especially those that are the result of human susceptibility to emotions and irrationality. Not surprisingly, then, algorithms are described as a bulwark against crowd behavior. In the words of one HFT trader:

> human beings are naturally trend followers. If the market is selling off, then they start panicking and everybody else does. But that's not what machines do. What machines do is they'll say, 'okay, historically, when everybody else is selling, it's more profitable to be a buyer'. Because that's what the data says. So machines are not emotional. They don't really care. They can be preprogrammed so that if they hit a stop loss then they can stop trading. But that's not an emotional decision. Very often traders may have a stop loss too, but they might ignore it because they are emotional and they don't want to go home losing money. So they'll keep trading and just get whipped out.

Similarly, another HFT trader states that 'machines never disobey their rules and their rules are more intelligent because they are not arbitrary, they are not based on emotions […], they are based on what the data says you should do'. So, a central motivation for doing HFT is that it can exclude or minimize incorrect decisions or human biases that stem from irrational and contagiously spreading emotional responses. Such human faults might not directly lead to a loss, but they do compromise the consistency of the trading approach and thereby lead to taking greater or unintended risks. A trader managing a HFT firm said that:

> I like the quantitative, very computer-based approach. Get rid of the emotion, you know. If you just lost […] much money yesterday you might be pushing more to make it up. A computer program never pushes to make up for an error from yesterday. If it lost money yesterday, it's not taking more risk to make more money. It's going to take exactly [the risk it is programmed to].

Such views are strong and often heard among HFT traders. And yet, as we shall demonstrate, in spite of this persistent celebration of the virtues of the non-emotionality of HFT algorithms, significant efforts are made by HFT traders to keep ensuring that emotions do not suddenly interfere with the algorithms. This materializes in various ways.

Physically, we have seen HFT trading rooms, which are designed so as to avoid that unintended imitative dynamics play out across traders working in the same room. In some places separate walls are erected between the traders to keep them psychically apart, just as safety screen are used together with secretive displays. The general notion here is that communication between HFT traders can be disastrous, as one trader's emotions may spread contagiously to the others, potentially undermining the rational operations of the algorithms if emotionally affected traders start interfering in their algorithms. Indeed, in addition to this physical and organizational layer, HFT traders spend much energy avoiding that they interfere emotionally in their algorithms. Part of this relates to how to deal with profits and losses. One trader put it the following way:

> I try not to get too happy on a winning day. I try to temper myself in both directions. When it's a losing day it's very much part of the strategy, it's very much part of the strategy. So I try as much as possible not to let myself experience the emotional swings.

The importance of controlling oneself and not becoming emotionally affected by the market and its apparent pulls evokes a key idea in much twentieth-century contrarian investment theory. According to the contrarians, affect control is central to the clever investor: only by disciplining oneself and by maintaining one's individuality (and hence rational abilities) can one avoid being subsumed under the market crowd – and thereby avoid being misled to make irrational investments (for a discussion of this self-disciplining affect control, see Stäheli 2006: 282–7).[8]

In HFT this shows not only in how traders seek not to let recent losses and profits interfere in their strategies, but also more generally in how they strive not to interfere too much in the algorithms. For example, the futures market is open 23 hours. In the one-hour close traders seek to tame themselves not to 'overreact'. One trader explained how he struggles to keep calm in the one-hour period and often he would need to leave the screen not to 'overadjust' the algorithms by, say, changing the value or risk parameters. Indeed, he said, the desire to adjust the algorithms, while knowing that this may reflect a short-term emotional impulse, means that 'that hour can be complex emotionally'.

This, of course, does not entail that no interference with algorithms is needed. It is generally acknowledged that HFT algorithms need to be regularly supervised, adjusted, and further innovated, in part because specific strategies are rarely efficient for more than six months. For present purposes the most interesting aspect of this relates to how HFT traders can often end up having very intense relations with their algorithms. One trader noted that 'I don't go more than 45 minutes without checking my strategies, that's absolute sure'. He further explained that this was due to his own anxieties – it was in order 'to shake less'. This personal nurturing of the algorithms further materializes in a general reluctance to employ night-traders to monitor the algorithms. Not only are the skills required to do this highly specialized, many HFT traders express a discomfort with having other people interfering in their algorithms.

To sum up this section, we have demonstrated two things. First, the turn to HFT does not make human beings irrelevant, but it does reconfigure the relations between individual HFT traders and the market since this relation is now mediated through algorithms. Second, in spite of this reconfiguration, and in spite of the rationality and non-emotionality ascribed to algorithms, emotional crowd contagion remains a concern to be constantly addressed in practice. So, even within HFT the relation between individual and crowd is of key importance in the sense that many efforts are made to avoid partly that market emotions spread contagiously among traders and partly that, should one be captured by

[8] Somewhat relatedly, many HFT algorithms are programmed to suspend trading around the announcement of key numbers. This too is a means to avoid being captured by erratic market contagions.

such emotions (either due to contagion or because of losses/profits), traders do not interfere too much in their algorithms. Consequently, the alleged need for market detachment – a notion derived from crowd psychology in early twentieth-century speculation theory – is very alive also in HFT.

Concluding Remarks

This paper has examined HFT and how it is associated, by financial economists and HFT traders alike, with contagious dynamics. We have argued that particular notions from late-nineteenth-century crowd psychology are being evoked within financial economics on HFT, such as the alleged supraindividual contagious, self-organizing, emergent properties that are often being attributed to HFT, and which are said to be prone to have irrational effects on financial markets. We have further argued that the link between current financial economics on HFT and crowd psychology is not only a curious one – a mere coincidental articulation of similar ideas. Rather the reference to crowd psychology is informative in that it invites a discussion of how crowd dynamics play out in a HFT reality. Following that lead, we have drawn on interviews and ethnographic observations to study sociologically how HFT looks from the inside, i.e. from the point of view of the people who trade actively on the basis of HFT algorithms. This analysis has demonstrated that contagious crowd dynamics are indeed central to HFT strategies in several ways. Thus, on one level, HFT strategies seek to exploit contagion in the form of structural correlations – what we have referred to as weak contagion; at the same time, HFT black boxes are designed in a fashion where adaptive feedback loops play a crucial role, thereby rendering strong contagion across markets more likely. As we have pointed out, this finding lends support to ABM approaches to HFT. On another level, contagious crowd dynamics appear as something to be avoided by all means. This is especially the case for how human HFT traders interact with their algorithms: it is a widespread concern that traders interfere too much, and too emotionally, in their allegedly rational algorithms – and that irrational market contagion may therefore creep into the HFT algorithms. As a result efforts are made to ensure detachment from contagious, irrational market crowds. Summing up, our analysis suggests that HFT at once takes for granted market contagion, is implicated in reinforcing it, and seeks to avoid it.

What this analysis offers for discussions of the regulation of financial markets is an inside view on the relation between what we have defined as weak and strong contagion. The ways in which HFT traders perform these kinds of contagion may well have an impact on the stability of the financial markets. As we have demonstrated, HFT practices are organized in ways that seek to eliminate the risk of contagion dynamics (i.e. the widespread concern among HFT traders about the risk of emotional contagion affecting the algorithms), yet contagion informs the basic assumption upon which HFT trading strategies rely. This tension might pose specific kinds of risks in financial markets. Comparing the design of a black-box system to the way in which the human trader sitting in front of the screen on a daily basis interacts with the system sheds light on a certain kind of *cultural contagion* that spreads not via the interaction between human traders preprogramming the same assumptions into their algorithmic systems, so that they might all react in similar ways. Rather, the opposite may actually be the case: the traders relate and refine their strategies on the basis of previous trading volume (informed by controlled feedback loops inside the black box). One might imagine that such interaction between the traders and the feedback loops they design to imitate other traders' trading behavior might pose new risks not linked to the intentional acts of human traders but to more collective dynamics. On a related note, the exploitation of contagion (even in a weak form) might lead to strong contagion. This conclusion, when looking at how the assumption of weak contagion might inform the design of HFT algorithms, contradicts Easley, López de Prado and O'Hara's argument about seeing weak contagion as 'a natural consequence' of market-making behavior (2013: 146). Weak contagion can very likely be seen as a basic condition for the behavior of adaptive feedback loops.

Finally, as we mentioned in the introduction, the vocabularies that are being deployed in financial markets (and elsewhere) are not inconsequential. How traders describe their own reality provides insights into the assumptions upon which they program their algorithms, detect meaningful information, and design more or less controlled feedback loops. This suggests a need for regulators of financial markets to take seriously conceptions of collective dynamics when addressing risk in relation to HFT. What our analysis suggests is that crowd psychology, or particular tropes from it, is deeply rooted among traders – it constitutes the way traders make sense of and account for their own practice. Importantly, it is likely that such interpretative accounts (enacted by traders on a daily basis) affect the ways they define meaningful information and extract signals defining the risk limits enacted by the algorithms especially when programmed into algorithms entering interactive feedback loops. Put differently, the vocabulary of crowd contagion may well have performative (self-reinforcing) effects on how HFT is organized and, thereby, on its broader market effects.

About the Authors

Christian Borch is Professor of Political Sociology at the Department of Management, Politics and Philosophy, Copenhagen Business School, Denmark. He is the principal investigator of a sociological research project on 'Crowd Dynamics in Financial Markets', funded by the Danish Council for Independent Research. He is the author of several books, including *The Politics of Crowds* (Cambridge University Press, 2012), which was awarded the 2014 Theory Book Prize by the American Sociological Association.

Dr. Ann-Christina Lange is Assistant Professor at the Department of Management, Politics and Philosophy, Copenhagen Business School, Denmark. She is part of the 'Crowd Dynamics in Financial Markets' research project, where her work focuses on high-frequency trading. She obtained her PhD from the Department of Sociology, Goldsmiths, University of London, in 2013.

References

Ben-David, I., Franzoni, F. and Moussawi, R. (2012) ETFs, Arbitrage, and Contagion. Available at: http://www.ilsole24ore.com/pdf2010/SoleOnLine5/_Oggetti_Correlati/Documenti/Finanza%20e%20Mercati/2012/02/etf-testo-ricerca.pdf

Borch, C. (2007) 'Crowds and economic life: bringing an old figure back in', Economy and Society 36(4): 549–73.

Borch, C. (2012) *The Politics of Crowds: An Alternative History of Sociology*. Cambridge: Cambridge University Press.

Brunnermeier, M. K., and Pedersen, L. H. (2005) 'Predatory Trading', *Journal of Finance* 40(4): 1825–63.

CFTC and SEC Staff Report (2010) 'Findings regarding the market events of May 6', Avaiable at: http://www.sec.gov/news/studies/2010/marketevents-report.pdf

Cui C. and Lauricella T. (2011) 'Mini 'Crashes' Hit Commodity Trade', The Wall Street Journal. Retrieved 7 May 2011

Easley, D., López de Prado, M., and O'Hara, M. (2013) 'Liquidity and Toxicity Contagion', in D. Easley, M. López de Prado and M. O'Hara (eds.), *High-Frequency Trading: New Realities for Tarders, Markets and Regulators* (pp. 143–57). London: Risk Books.

Garfinkel, H. (1967) *Studies in ethnomethodology* Oxford: Blackwell Publishing Ltd.

Golub, A, Keane, J. and Poon, S. H. (2012) High Frequency Trading and Mini Flash Crashes. Available at: http://ssrn.com/abstract=2182097 or http://dx.doi.org/10.2139/ssrn.2182097

Johnson N., et al. (2012) Financial black swans driven by ultrafast machine ecology. Available at: http://arxiv.org/ftp/arxiv/papers/1202/1202.1448.pdf

Knorr Cetina, K. (2009) The Synthetic Situation: Interactionism for a Global World, Symbolic Interaction, 32(1): 61–87.

Kolb, R. W. (2011) 'What is *Financial Contagion*?', in R. W. Kolb (ed.), *Financial Contagion: The Viral Threat to the Wealth of Nations* (pp. 3-10). Hoboken: John Wiley.

Le Bon, G. (2002) *The Crowd: A Study of the Popular Mind*. Mineola, New York: Dover.

Lewis, M. (2014a) *Flash Boys: A Wall Street Revolt*. New York: Allen Lane.

MacKenzie, D. (2006) *An Engine, not a Camera: How Financial Models Shape Markets*. Cambridge, MA: MIT Press.

MacKenzie, D. (2014) *A Sociology of Algorithms: High-Frequency Trading and the Shaping of Markets*. Edinburgh: Unpublished paper.

MacKenzie, D., and Millo, Y. (2003) 'Constructing a Market, Performing Theory: The Historical Sociology of a Financial Derivatives Exchange', *American Journal of Sociology* 109(1): 107–45.

Mitchell, P. (2012) *Contagious Metaphor*. London: Bloomsbury.

Shiller, R. J. (1984) 'Stock Prices and Social Dynamics', *Brookings Papers on Economic Activity* 2: 457–510.

Sornette, D., and von der Becke, S. (2011) 'Crashes and High Frequency Trading: An Evaluation of Risks Posed by High-Speed Algorithmic Trading', *Zurich: Swiss Finance Institute, Reseach Paper Series* No 11 – 63.

Stäheli, U. (2006) 'Market Crowds', in J. T. Schnapp and M. Tiews (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.

Taussig, F. W. (1911) *Pinciples of Economics, vol. 1*. New York: Macmillan.