Government Office for Science

Foresight

Feedback effects and changes in the diversity of trading strategies

The Future of Computer Trading in Financial Markets - Foresight Driver Review – DR 2

Contents

1 Introduction 4
2 Diversity, and the link to Positive Feedback Loops 5
3 Algorithms and Feedback Loops 7
4 Algos and Market-Wide Effects 11
5 Algos and Financial Innovation 13
6 Computer Trading and Information 14
7 Computer Trading and Common Knowledge 15
First 16
Second 16
8 Common Knowledge and Coordination in Distributed Environments 17
9 Network Effects in Decentralised Markets 18
10 A Sociological Approach to the Risk of Positive Feedback in Automated Trading 20
References 21

Feedback effects and changes in the diversity of trading strategies [1]

Jean-Pierre Zigrand

London School of Economics

31/8/2011

Review Author: Dr. Jean-Pierre Zigrand, LSE

Additional Authors: Prof. Hyun Song Shin, Princeton, and Dr. Daniel Beunza, LSE

1 We would like to thank Philip Bond, Kevin Houstoun, Eric Hunsader and Pete Kyle for helpful discussions.

1 Introduction

A driver for future risk and catastrophes lies in the fact that the seemingly large bio-diversity of traders may be illusory and that in a stress situation many algorithms quickly and unwittingly coordinate, act in unison and feed on each other in a feedback loop, thereby leading to a disproportionate value destruction.

The first of the two main themes concerns *uniformity* (or its antonym *diversity*) and *synchronisation*. Even though various strategies and algorithms can appear to be genuinely diverse, there is the risk that a market event - even one quite unrelated to the strategies underlying the algorithms - can under a given set of circumstances make the actions underlying the strategies nearly instantaneously uniform and synchronised. This paper explores a number of possible situations that could lead to such unintended coordination. Furthermore, the trigger leading to lock-stepping can be very nonlinear in the driving variables, which means that to an outside observer who does not see the inner workings of the markets, the ganging-up appears very unpredictable.

The second, and often complementary, theme concerns *feedback effects* or *feedback loops*, mainly positive. Positive feedback effects act as a mutually reinforcing propagation mechanism. This mechanism is strengthened precisely the more players act upon it at the same time, i.e. the more they synchronise, which is the first theme. In finance, positive feedback effects often arise from dual roles played by endogenous variables. For instance, prices both *reflect* fundamentals and actions, and *drive* actions, too. In standard models, the latter comes from various constraints, such as leverage caps, VaR constraints, solvency constraints, policy responses (such as inflation targeting by central banks) and so forth. But with computer driven trades, the link may be more mechanical, systematic and immediate, and hence more direct and more powerful. The feedback loop emerges when both aspects of endogenous variables (as signal and as imperative to action) are in place.

In our mind there is little doubt that positive feedback loops are underlying forces behind many, if not most, financial crises. The fact that neither in the *Findings Regarding the Market Events of May 6, 2010* (The Staffs of the CFTC and SEC (2010)) nor in the *Recommendations regarding regulatory responses to the market events of May 6, 2010* (Joint CFTC-SEC Advisory Committee on Emerging Regulatory Issues (2010)) there appears to be any explicit mention of the feedback loop dynamics inherent in such events is a cause of some worry. Such nonlinearities,[1] and the inherent unpredictability and speed of unravelling that go with such feedback nonlinearities, underlie many surprising financial crises and deserve to be singled out as a major issue. In principle positive feedback loops could lead either to welfare improving or to welfare destroying outcomes, although economic logic more often than not leads to deleterious effects.[2] We believe that a very simple – though necessarily superficial – criterion that market regulators ought to keep a watchful eye on is the minimisation of net positive

[1] Feedback loops give rise to strong nonlinearities since the system feeds on itself as long as the driving variables stay in a given range. Some violent feedback loops gradually soften and then disappear as the variable that is fed back is pushed out of range while others excite themselves in a never ending spiral to a critical point. We shall come back to this theme repeatedly in the sequel.

[2] Options markets for instance recognise this through the implied volatility skew that exhibits higher implied volatilities for out-of-the money puts – those options that pay out if markets tank – than for out-of-the money calls – those options that pay out if markets rally.

feedback loops, those naturally present in markets as well as those caused by the very regulations aimed to prevent crises.

In this paper, we outline some of the risk drivers from such deleterious loops in computer-based environments. The seeds of such crises episodes are to a non-negligible extent already sown and exist today as potentials, even if the conditions or realisations that would bring them to the forefront have not materialised yet. In that sense many of the risk drivers of the next 5 to 10 years are the same as the ones present now, with the caveat that given the perceived trends in technology, the unravelling may become more forceful and more rapid. For other risk drivers, the technological conditions that would make them dangerous are yet to be fulfilled through potential future developments.

2 Diversity, and the link to Positive Feedback Loops

Computer algorithms come in a few broad flavours, such as based on momentum or on mean-reversion for HFT for instance, or based on a variety of market-impact reducing techniques for buy-side execution algorithms, such as volume-weighted average price VWAP, time-weighted average price TWAP or arrival price, AP. There is also a library of infamous algorithms and practices that have been sometimes accused in the press of being socially harmful, such as liquidity detection, order anticipation, momentum ignition, quote stuffing, tape painting, pushing the elephant or towing the iceberg etc. algorithms (see for instance Durbin (2010) or Donefer (2010) for introductory surveys). Some low-latency algorithms operate across alternative trading venues, attempting to lock in tiny mispricings. This is obviously not an exhaustive list, and other drivers within this Foresight exercise are concerned with painting a fuller picture of the algorithmic bestiary.

More diversity of trading participants is usually welcomed. For instance, depending on the strengths of income and substitution effects, there is nothing pathological about an upward sloping individual demand curve. But unusual and strange market phenomena can occur with upward sloping *aggregate* demand curves.[3] Along one line of reasoning going back to Hildenbrand 1983, provided that agents' characteristics are dispersed in a sufficiently heterogenous and complementary fashion, then the Law of Demand does hold in the aggregate (under some assumptions, including the absence of market frictions) even if individual agents can have downward sloping demands. Diversity produces regularity in this case through a law-of-large-numbers averaging reasoning.

---

[3] For instance, the portfolio insurance–led market decline in 1987 effectively exhibits an upward sloping demand. Delta hedging of negative convexity exposures requires the delta hedger (in this case the portfolio insurer) to sell after a price fall, i.e. the delta hedger's demand is upward sloping since less is demanded at lower prices. A price fall leads to more selling, which in turn leads to a lower price, which in turn leads to more selling etc, until prices have fallen low enough into a region far away from at-the-money where convexity has become negligible:

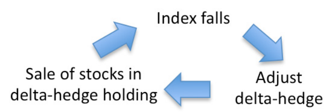

More instances can be found in Footnote 2.

This does not mean that there cannot be positive feedback loops in the presence of heterogeneity (given or emergent), to the contrary. For instance, the "dying seminar" unravelling model of Schelling (1978) or the simple riot model of Granovetter (1978) illustrate this point (also refer for instance to the introductory textbook by Page (2011) for further details and examples). During an episode of social upheaval, a group of people are pondering whether to riot, and in doing so they base their decision to join or not on the number of other people who do. If all of their thresholds are identical, then they will not riot. For instance, assume that each of a thousand people have a uniform threshold of 20. Any person who witnesses at least 20 people rioting joins the riot. Then there will be no riot if agents act sequentially.[4] On the other had, in a more diverse group of a thousand people, assume that the first 10 people riot anyhow, the next 10 people riot provided at least 10 other people riot, the next 10 riot provided at least 20 other people riot, etc. Then all will riot. It is precisely the variation of thresholds (together with the assumptions that the first ten people riot anyhow and that decisions are taken sequentially) that is creating the tipping point through positive feedback effects. When tipping points and cascades occur at endogenous thresholds - for instance influenced in part by contemporaneous prices - as they do in finance, a more subtle analysis is required. Examples include the herding literature in finance along the lines of Avery and Zemski (1998), the applications of global games or even the analysis of financial networks whereby cascades can occur when the failure of one entity in a network implies the failure of a linked entity.

More subtly, one of our main points would be that many of the greatest risk drivers lie in the fact that diversity is not a fixed hard-wired characteristic of modern markets. Even if diversity is of the stabilising (negative feedback) sort, the influence of the diversity of fundamental characteristics can very rapidly disappear, pushing agents to act in unison regardless of their fundamental diversity. In other words, if market phenomena induce agents to effectively coordinate, any effects that fundamental diversity might otherwise have had disappear, and agents are effectively homogenous. To go back to the downward sloping aggregate demand function in the presence of diversity, if investors are interacting even indirectly, demand functions can quickly morph into upward sloping ones and lead to positive, mutually reinforcing, feedback effects and therefore to situations of extreme market stress. In the paper by Danielsson and Zigrand (2008), for instance, this occurs through marked-to-market value-at-risk (VaR) constraints imposed upon a heterogenous population of traders. In the event of a negative shock to the economy, the VaR binds for the less risk-averse agents and effectively constrains them to act uniformly, selling the same risky securities. This in turn depresses prices further and raises risk warnings trough the dual nature of prices as both indicators of value and imperatives for action. Market demand functions are effectively *upward sloping*,[5] leading to renewed selling the lower the prices are. In turn, yet more agents become subjected to uniform regulatory constraints, and so forth.[6]

---

[4] The outcome that all riot is a Nash equilibrium of a simultaneous move game.

[5] From Zigrand (2006) we know that absent intertemporal effects, downward sloping aggregate demands negate manipulation opportunities in financial markets. As we shall see later, the phenomenon of quote stuffing creates intertemporal effects through the slowing down of quotes on one of the trading venues and therefore leads to manipulation opportunities somewhat reminiscent of Jarrow (1992).

[6] In contrast to the riot model, however, the positive feedback effects would have occurred in a uniform society as well, and could have been possibly more violent still. But it matters greatly how you exactly transform a diverse society into a uniform one: if you make each agent in a diverse society as unhinged as the most unhinged in that diverse society does not lead to the same conclusions as if you

What is more, this mechanism of endogenous risk[7] does not require explicitly interdependent preferences or choices as in the tipping models of Granovetter or Schelling (1978). This is relevant in the context of computer trading, for it would be unrealistic to assume that one algorithm directly decides on its actions as a function of the past observed actions of another algorithm. In fact, it works even if all interactions are market-based and price-intermediated. In that sense, the interdependence is not exogenous and direct, but arises endogenously as an indirect mechanism via market-wide risk appetite. Because this interdependence is endogenous and not part of the hard-wired setup, its dynamics are subtle and hard for agents to estimate or take into consideration. Banks and hedge-funds appear to have been largely unaware of these interactions before the last few crises, and showed great astonishment after the fact (see LTCM in 1998, The Quants in 2007 etc.). To make matters even worse, formal models of endogenous risk such as Danielsson et al. (2010) and the fact-finding mission by Arup following the wobbly Millenium Bridge event found that this interdependence is very *non-linear*. This means that for much of the time the interdependence (and the resulting feedback loops) are undetectable, but once the driving variables reach certain critical thresholds (a critically low capitalisation level of banks in Danielsson et al. or a critically high number of pedestrians on the Millenium Bridge (estimated to be 156)), idiosyncrasies no longer cancel out and (inter)actions become very rapidly mutually reinforcing.[8] Nonlinearities such as these are one of the reasons why critical episodes and flash-crashes are so hard to predict since market participants do not possess the “true model” and need to estimate and second-guess it in real-time (estimation is also not helped by the fact that most econometric tools normally used are based on *linear* regressions).

3 Algorithms and Feedback Loops

We believe that deleterious non-linear effects of the endogenous risk type could plausibly arise in algorithmic trading environments. Regardless of the perceived benefits of heterogeneity in normal times, such as providing markets with an attractive biodiversity and double coincidence of wants, under situations of stress, algorithms may gang up and very quickly become one, all on the same side of the market and heading for the exit.

There are two main reasons why that may be. Either the algorithms have actions hard-coded into their programmes that directly lead to positive feedbacks. We call this the *direct* mechanism. Alternatively, the algorithms do not have such a behaviour coded directly into the programmes, but some higher level interventions by the controlling or the supervising entity effectively can decide to overrule the algorithm, and thereby create the feedbacks.[9] This we call the *indirect* mechanism. For instance, a high-level operator (either the owner of the algorithms or the market regulator) may decide in times of stress to overrule the algorithms and

---

had transformed each heterogenous agent into the equivalent of the most risk-averse and responsible one.

[7] Loosely speaking, *endogenous risk* is the risk and volatility increment (or multiplier) to the fundamental risk and volatility that is due entirely to amplifying feedback loops generated purely by some institutional features of the structure of the system itself.

[8] The stabilising averaging benefits of the law of large numbers fail in that case due to the interdependence.

[9] In HFT systems, this would be done through what is sometimes referred to as the *Strategy Server* in the *Thinker* component.

liquidate all positions, say, which then leads to further losses and higher perceived risk, which may lead to yet more sales by this or other operators, etc.[10]

The most basic situation would be one whereby stop loss orders on a single stock are implemented through algorithms and where there is a sequence of decreasing stop loss thresholds inherent in the algorithms, perhaps belonging to one firm or to a sequence of firms. An initial market sell order, or a HFT company trying to sniff out stop losses, hits the bids and pulls the price through the highest stop loss threshold. The algorithm reacts by selling, which in turn leads to the touching of the second highest stop loss barrier, which in turn forces some algorithm to sell and breach the next threshold etc, closing the loop. If the sequence of stop losses leads to a sale that is insufficient to lower the price below the next highest barrier, the further presence of momentum strategies may keep the loop operating by helping push prices through subsequent barriers.

Whether the official version of the events of May 6th 2010 turns out to be accurate and complete or not, it does illustrate a risk driver in the sense that it outlines a possible scenario whereby some algorithms may *directly* create feedback effects due to a lack of common sense in the coding. In a nutshell, a simplified scenario would be as follows. An execution algorithm by a buy-side firm is directed to sell a large number of securities. The algorithm basically gauges the market impact it may have by looking at the market volume and is instructed to sell more if volume is higher. The original large sale finds buyers, most likely HFTs given their speed advantage. If it turns out that there are no real-money investors stepping in, the HFT algorithms may pass the securities around like a hot potato, generating more volume. This volume then teases the algorithm to sell even more, closing the feedback loop until far out-of-the-money limit orders are hit and the order books are emptied. The destabilising feedback loop in this instance has been brought about through the interaction of two distinct algorithms. Diagrammatically,

Another direct feedback can be created from a market making algorithm of the following type: "I cancel and resubmit my limit orders with a larger bid-ask spread if realised volatility in the market picks up." If volatility then does pick up for whatever reason, then the market-making algorithm raises the bid-ask spread. As buy and sell market orders come in, prices bounce off

[10] The opposite scenario is possible also. Perhaps human operators may realise, at least after some delay due to the speed of trading, the potential to a feedback loop in a given situation and engage in evasive action that a computer algorithm would not be able to see. For instance, Leland states that portfolio insurers intended to sell almost triple the amount they were actually able to execute on October 1987. It would appear that by the afternoon of October 20th 1987 many asset holders turned their sales programmes off.

the wider spread which, depending on how realised volatility is measured, may lead to higher perceived volatility by the algorithm, which in turn forces it to widen its spreads, closing the feedback loop.

Closely related is the potential feedback loop described by Angel (1994) and Zovko and Farmer (2002), which one might call the *shallowness feedback loop*. The data on limit-order books studied in the latter paper seem to suggest that the realised volatility of the best bid leads (in tick-time) the average dispersion of bids and of asks centered around the best bid and the best ask respectively. The positive feedback loop this entails works as follows. Assume an initial increase in volatility, perhaps due to news. The distribution of bids and asks in the limit order book adjusts and becomes more dispersed. Everything else constant, as market orders or crossed limit orders arrive, they “walk the book” and move trade prices further away from the mid-point. The increased realised volatility in turn feeds back into yet more dispersed quotes, and the loop is closed. Diagrammatically,

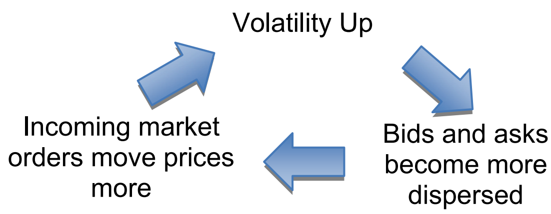

Purely electronic markets exhibit another instance of the fallacy of composition whereby the behaviour of individually prudent and conservative algorithms can mechanically and very quickly generate destabilising and overall imprudent dynamics. The loops in the last two instances are strongest if the interaction is between the market making algorithms that spread out more in volatile markets and liquidity absorbing algorithms that exhibit momentum type behaviour, such as VWAP, Arrival Time, or that themselves already exhibit positive feedback loops, such as Value-at-Risk or certain delta-hedging strategies, and that will be further accentuated through the shallowness feedback loop.

The news-tagging features of newsfeeds also may lead to reinforcing feedbacks that we label *news feedbacks*. The *news listener* component of HFT systems scans headlines for tags and acts upon them immediately by broadcasting the tag to all other components of the HFT system. That there are HFT algorithms that react to news in a mechanical and uncritical way has been colourfully illustrated in the September 8, 2008 event where a six years old story of a United Airlines reorganisation made it into the newsfeeds, including Bloomberg, as breaking news and led to a massive sell-off of its stock. Now suppose that an algorithm detects and executes a profitable sell trade to satisfaction, moving the market down as a result. In general, such a market move by itself may make it into the newsfeed, which through the news listener may generate further sales orders, both from separate HFTS and even from the original HFT that was the root of the newsfeed in the first place. Even though the pricing algorithm of the original HFT believed that the price after its original sale is equal to its *microprice* (the HFT’s perception of fair value), after the newsfeed came in it may have revised its microprice either because the HFT does not know that the newsfeed referred to its own original sale, or because the pricer now takes its effect on other traders into account through the dissemination of the newsfeed. Diagrammatically,

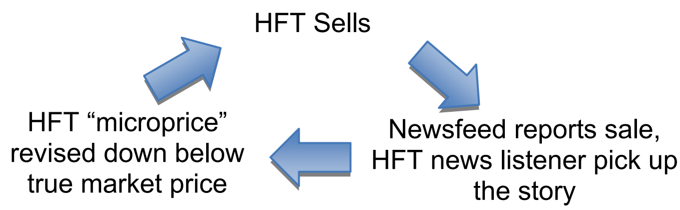

When reading Eric Hunsader's data analysis of the events from May 6th (for instance Hunsader (2010)), the following possible strengthening of feedback effects comes to mind. Hunsader (2010) shows that during a number of occasions both before and after the Flash Crash, the NYSE CQS quote feeds have suffered from delays varying from a few seconds to minutes, perhaps due to deliberate quote stuffing. If the delays are real but small, as they were during the Flash Crash, the delays are hard to detect. The trade reporting system CTS can be expected to be less affected by such delays, although CTS was similarly affected during the Flash Crash. In a fragmented market suffering from overall selling pressure of a basket of high-capitalisation stocks (e.g. originating from the sales of E-Mini futures), assume the NYSE CQS quotes lag by a bit. Since the market is falling, the NYSE bids appear to be the most attractive to sellers, and all sales are routed to NYSE, regardless of the fact that actual bids were lower.[11] Prices fall by more than they would have in the absence of delays since the pent-up limit buy orders on the remaining TVs are not called upon and all sales hit the bids on the NYSE only. This large and sudden price drop can lead algorithmic momentum robots that trade bellwether stocks to short those stocks, leading to further downward price pressure across all markets. Furthermore, through the indirect mechanism HFTs may simply sell their inventories and temporarily exit market-making altogether when price aberrations and odd delays lead to uncertainty and fear. Since the CQS quotes remain delayed, all renewed selling pressure goes to the NYSE, and the strengthened *pricing feedback loop* delays creep in, the increased flurry of activity arising from the previous feedback loop can cause further misalignments in bid/ask time stamps, closing the *delay feedback loop* that is amplifying the pricing feedback loop. Schematically:

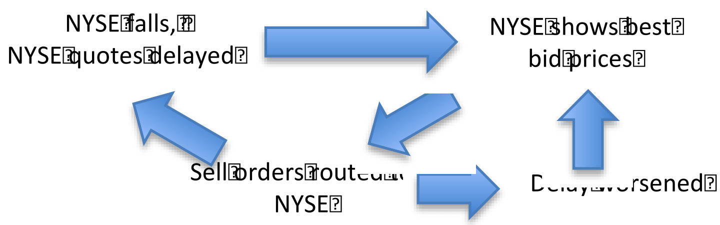

This list is not intended to be comprehensive but merely illustrative. There are infinitely many possible versions of feedback effects that could occur, most of which are still unknown.

As mentioned in the previous section on diversity, the reinforcing feedback effects created by this one hypothetical algorithm are worsened further if the reactions spread to become system-wide, say because these propagating dynamics force other players or distinct algorithms to act in the same fashion and mimick the reinforcing actions of the first category of algorithms. Take for instance the second scenario of the market making algorithm. The larger bid-ask spreads may indicate less liquid markets, or the higher volatility may indicate a riskier environment, both would lead some investors who are required to invest only in safe and liquid

[11] Thoughtful HFTs are not likely to be such sellers since most HFTs, for a fee, get their quotes directly from NYSE through OpenBook as opposed to through CQS.

securities (or who are subjected to VaR-type constraints) to divest themselves from the securities, leading to transaction prices blowing through the limit order book, further raising volatility and bid-ask spreads, etc.

One might say that some of the aforementioned feedback loops also can occur in non-computer based markets. That is correct, although the force of the loops is determined by how mechanical the rules are, and the speed determines how quickly humans can jump in and either turn the algorithm off or trade against it. A human trader would be unlikely to sell 100 Accenture shares at one cent as happened at 2:47:53 p.m. on May 6th 2010 when prices traded around $40 seconds before, a ballpark level that is justifiable in terms of fundamentals.

Finally, with the trend towards computer-based trading follows the trend for longer opening hours for trading venues. Many contracts today can be traded electronically throughout the day and night, with perhaps one or two brief periods of closure. Some contracts that retain open-outcry during waking hours can be traded electronically out-of-hours. We expect that over the next 5 to 10 years most contracts will trade around the clock. While the feedback mechanisms described above can apply anytime, it is clear that they propagate faster if markets are more shallow, which they tend to be after hours. Since the trend towards electronic trading enables both new sorts and more violent feedback loops as well as nearly 24h continuous trading, the probability of such a feedback loop hitting a thinly traded patch out-of hours it itself reinforced naturally.

4 Algos and Market-Wide Effects

Following Mifid and RegNMS, price discovery and allocation efficiency rely on arbitrageurs who are linking the diverse trading venues and integrating prices through buying low and selling high across venues. This is on top of the arbitrage trades linking derivatives and ETFs to the underlying replicating securities. Linking venues and securities is deemed to be a popular market-making and HFT strategy, although it is hard to find precise figures as to what fraction of HFT profits is due to such “cross-sectional” market making, as opposed to time-series or statistical trading. Be that as it may, the integrity of market signals relies on HFTs equalising, for a fee, the marginal rates of substitution across multiple venues. In that sense, HFTs are now performing the job of the Walrasian auctioneer. They are de facto the invisible hand (for a formal model establishing this fact, see Zigrand (2004)). Because that service is socially useful, thoughts ought to be given to assure its flawless and continued operation.

There is a risk driver linked to this observation. If the current decentralised market structure delegates the socially useful job of the Walrasian auctioneer to the HFTs, then the economy relies on HFTs. Social welfare requires that this role be filled in a permanent fashion. There is a small risk that HFTs, under situations of market stress, margin calls or collateral pressure, will not be able or willing to perform this job. Prices on the various venues will not be synchronised, trades and transactions will happen at socially inefficient prices, marking will need to be done to multiple and illiquid marks and doubts will creep in as to the true value of securities.[12] All of those second-round effects will have as a consequence that market stress is exacerbated; liquidity and pricing efficiencies deteriorate further, leading to a socially harmful positive feedback loop.

[12] With multiple prices, a depth-weighted average would be a proxy for the true, unobserved, price, see Zigrand (2006) for a proof.

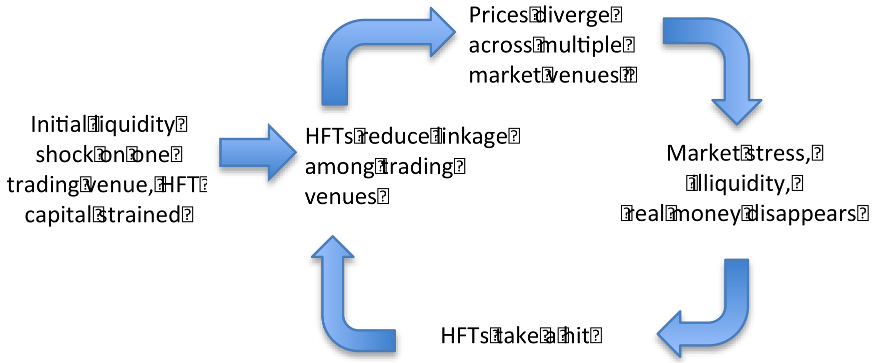

For instance, The Staffs of the CFTC and SEC (2010) report such findings on ETFs, whereby the extreme volatility of the individual component securities spilled over into the ETF markets and led market makers to pause their market making activities. In return, the illiquid and stub ETF prices for aggregates provide false systematic factor signals, feeding back into the pricing of individual securities etc.[13] Schematically, with SS a shorthand for single stocks,

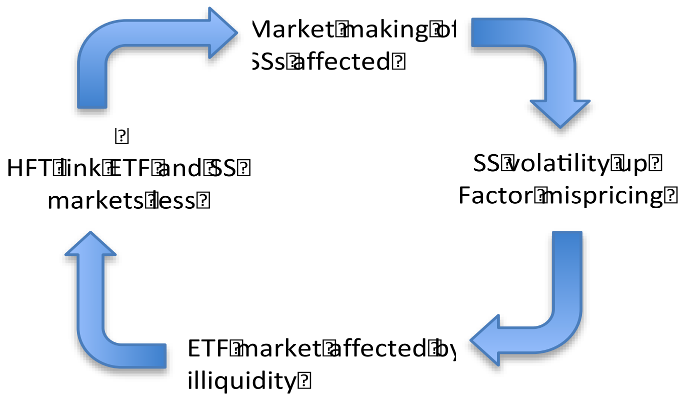

A closely related risk driver deals directly with marking-to-market requirements. If the future trading venues remain competitive / fragmented, and if future computer-based financial systems automatically and mechanically mark portfolios to market, this has the potential risk to make the feedback loop worse. This is because prices determined in an off-main-exchange market may have an impact on the larger market if traders look to the off-main-exchange market for cues for actions. Since the off-main-exchange market in this scenario is assumed to be very thinly traded, the effect of even small trades might be felt very powerfully in other markets if those other markets need to mark their prices to that thin mark. This is like the dog chasing its tail. A small ripple in the off-exchange market will push *all* traders to react in their respective markets, which they would not have done had the markets been integrated because the outlier demand or supply originally submitted to the thin venue would have been absorbed by the aggregate limit order book on an integrated exchange. Furthermore, strategic traders

[13] These feedback effects were amplified by the fact that measured market depth for the top 100 ETFs is substantially lower than for the top 100 individual stocks, see The Staffs of the CFTC and SEC (2010).

would be tempted to manipulate executed marks to maximise the MTM value of their portfolios or to inflict maximum damage to their competitors.

It is true that if we have a system of exchanges and MTFs or ECNs (hence-forth the MTFs and ECNs are called “Alternative Trading Venues,” ATVs, while any one venue, be it an exchange or an ATV, is labelled “trading venue,” or TV) in mind that trade the same securities, then an aberrant price would create an arbitrage opportunity that in normal circumstances would be very quickly undone. Cross-venue arbitraging is akin to a *negative feedback loop* since the further the mispricing, the stronger the arbitrage forces that attempt to profit from this mispricing and therefore the stronger the realignment forces undoing the arbitrage. In that sense the scenario we depicted may be more realistic in crises episodes, or in normal times for less liquid and less homogenous securities, such as derivatives and structured products, where there is less consensus as to price. For instance, trading firms are particularly worried by centralised clearing of illiquid derivatives or structured products where mispricing is very frequent given the small number of transactions and the large uncertainty about the right model to use. If one bank completes a deal with a buy-side client and this deal is cleared and the terms publicly posted post-trade, and if no other deal completes around that time, all banks and clients holding this (or a closely related) security will need to mark their own positions to that last mark, no matter how far off the price is. In a fully automated environment, this happens automatically with little human intervention that would have been able to qualify the mark on logical or intuitive grounds. But too low a mark will lead to MTM capital losses and will be flagged up by the risk systems of the market participants, who will be required to react mechanically to the new risks by lowering overall risk, in the short term through the reduction of risky positions held. Fire sales may ensue, in turn leading to lower prices, higher capital losses and more risk, and the positive feedback loop is closed once more. The endogenous risk loop can be represented schematically:

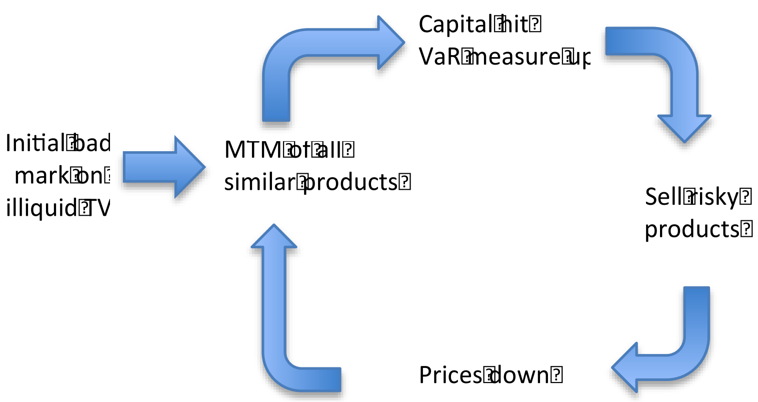

5 Algos and Financial Innovation

The ease of creating faster algorithms may lead to more dynamic trading, such as the dynamic self-financing delta-hedging that allows a trader to construct a non-linear final payoff. If the final payoff so constructed is convex, then the strategy implementing this payoff will lead to more buying after markets rise and more selling after markets fall. The feedback effect generated can be quite large, as evidenced in the 1987 crash, depending on the dollar gamma of the options.

A new modern version of this risk, though dampened due to the much lower gammas and vegas involved, appears in the hugely successful synthetic ETF markets.

6 Computer Trading and Information

In this section we review some of the risks to the degree of informational efficiency in computerised markets. In one view of the world, algorithms and computer trading operate as an invisible hand, as a veil, bringing *impersonal efficiency* (see the driver by Beunza (2011)) into the allocation of resources to long-term investors. In another view of the world, algorithms impound their own trading patterns into price dynamics, occasionally leading (inter alia through positive feedback effects) to prices and allocations far away from the contract curve, at least in the short run. In other words, prices may be equal to fundamental values plus a mean-reverting noise term driven by system-induced endogenous risk.

Most computer algorithms seem at the moment to largely consist of relatively short (and therefore simple) computer code. It seems that the inputs into the algorithm largely consist of the past behaviour and data series of this very asset itself, and perhaps of some related securities, mechanical news tags as well as engineering state variables such as the state of latency in the network, the temperature in the data center, sunspots that may perturb global communication network and the like. In other words, algorithms follow simple Markov rules. If much of the trading volume is generated by such univariate and low-dimensional rules, then fundamental research about the deep systematic macro pricing factors will perhaps find it harder to impact the market.[14] By the very fact that computer algorithms are able to crunch large amounts of data and that many algorithms probably data-mine shared time-series, they are more prone still to engage in similar strategies and coordinate. Of course, human traders may follow simple Markov rules as well, but they may not follow them quite as strictly and quite as quickly. As mentioned before, it would seem implausible to ascribe the sale of 30000 shares of Accenture for one penny a share at 2:47:53 p.m. on May 6th 2010 to a human trader since the human trader would have a fundamental value in mind to compare the price to. Some algorithms know the price of everything and the value of nothing. Human traders rely on the soft information and on “colour” that brokers tended to relay to traders. Algorithms on the other hand need to overcome their lack of intuition and overly rely on public signals. If such a trade is one that leads to self-reinforcement, such as momentum or carry-trade types of opportunities, the feedback loop leads to confirmation and to further reinforcement. Prices move further and further away from fundamental information-based prices. Since fundamental human investors require time to locate capital and analyse data, it takes time for them to push prices back towards value, and mean reversion towards fundamental values is slow as a result.

In a nutshell, market dynamics dominated by very short horizon robots run the risk of being “freer” in the sense of being more decoupled from a fundamental anchor, such as the fundamental valuation of payoffs and earnings. Markets that are more decoupled from fundamentals can depart more readily and be pushed further by self-reinforcing feedbacks. Of course one might say that there have always been technical traders, momentum followers and so forth among human traders, but the arbitrage instinct and the instinct of a bargain compared to a fundamental value that some of the human traders possess tend to act as a countervailing negative feedback force over the medium term. This is one risk driver. It is possible, of course, that trading robots with longer term horizons will be developed and fed with solid valuation tools based on historical, accounting or economic modelling signals. Such algorithms would be working towards prices that are glued more closely to fundamental values as defined by such signals, at least at certain time scales.

---

[14] A Bloomberg article (Foroohar (2010)) quotes a HFT hedge fund manager as saying “you don’t need to look at anything other than the price of Microsoft to determine if it is rich or cheap.”

7 Computer Trading and Common Knowledge

An event *E* is common knowledge if all know it has obtained, and all know that all know, and all know that all know that all now etc, ad infinitum. Suppose now that all market participants know that a certain event *E* has not occurred, but that it is not commonly known that the event has not occurred. Then no participant can simply disregard event *E* in his deliberations if he is engaged in a strategic situation interacting with other participants. The reason is that the payoff of any participant depends on the actions of other participants, as well as on the true (unknown) state of the world, and the actions of other participants depend on their own beliefs. So when I am informed that an event *E* has not obtained, this event cannot be ignored since it may contain information about your beliefs, not least on your beliefs about my own beliefs.

As Shin (1996) writes (our emphasis added),

> When unrealized states exercise an influence on the equilibrium allocation, there are unavoidable welfare consequences. Since the optimality or otherwise of the final allocation hinges on what the fundamentals are (in terms of preferences and endowments), a well-functioning trading system is one which ensures that post-trade allocations are determined in the appropriate way in relation to the fundamentals. However, when unrealized states exert an influence on the final allocation, *this link from fundamentals to the final outcome is subject to interference*. The mark of a *well-functioning trading system is one which minimizes such interference, and which ensures that the final allocation is as close as possible to that justified by the fundamentals*. In contrast, a poorly performing trading system is unable to insulate the equilibrium outcome from the influence of unrealized states.

Shin compares a very stylised version of a dealership market with a decentralised market à la Shapley and Shubik (1977) (but not allowing for limit orders). One can view the dealership market as a market where demand schedules are posted, firm and attributable to a given market maker, whereas the decentralised market is an order driven anonymous market where orders are not firm. He finds that

The dealership market has strictly higher trading volume, and delivers the Pareto efficient allocation in most states. In contrast, the decentralized market (employing the Shapley- Shubik rules) suffers from low trading volume, and the post-trade allocation is bounded away from the efficient allocation everywhere on the state space. (…) The apparent fragility of the Shapley-Shubik market to departures from common knowledge can be traced to the large effect of unrealized states on the equilibrium allocation.

In this sense, a “transparent market” with displayed and firm limit orders minimises interference since the liquidity provider posts the limit order first before liquidity users choose their actions and hit the limits, so that the liquidity provider does not need to worry what limits he would have posted had he observed different signals. This restores the direct link between the true fundamentals of the economy and the outcomes in terms of allocations and prices. One general risk driver therefore is the risk that anonymous computer based trading environments sever the links between market allocations and efficiency.

More restrictively, let us apply these ideas to algorithms. While crashes initiated by algorithms, such as the Flash crash, may revert quickly as the market realises that the issue is not a fundamental one, there is the possibility of further complication due to two channels.

First, the market may not realise that what is creating havoc is only an uninformed algorithm that is inducing other algorithms to contribute to the damage (see the seminal paper by Gennotte and Leland (1990)) rather than smart money trading. As Donefer (2010) rephrased it more recently: " A decade ago, regulators might see this activity [a selling spiral] as coming from a single source, a single fat-fingered trade, but now we just see many small orders across multiple markets. (…) It is not easily identified as a single algo; it is amorphous and widespread. (…) Is it "smart money" - does someone know something? "

In decentralised computer networks, the concept of what is and what is not common knowledge warrants deeper study. A panic stops once it is common knowledge that the driver is an algorithm of no fundamental importance and carrying no fundamental information, in which case traders are not second-guessing and running for cover just in case, and the market dislocation was temporary. So in centralised open outcry markets, the centrality of the pit meant that it was more likely that the cause of a dislocation became common knowledge, because of the mathematical fact that an event E is common knowledge if some self-evident event that entails the event E obtains. In the limit, we have the powerful theorem that says that if agents have asymmetric information but if their actions are common knowledge, then information is effectively symmetric.

In modern computer based markets, common knowledge may therefore be harder to obtain (also see next section for further evidence). For instance, limit orders are cancelled at an astoundingly fast rate, partly because of the advent of new information, but partly also for reasons unrelated to new information but rather as a result of strategic trading strategies. Furthermore, in fragmented MTF markets, traders and trading firms interact less with each other and the anonymity no longer allows the same extent of knowledge creation as in a specialist market where it is commonly known that the limit orders reflect the specialist's beliefs, or as in a pit-based market where orders and to some extent their originators are both publicly observed. With higher order uncertainty due to the impersonality of computer-based trading, markets may be more prone to crashes and large swings, as shown in the herding model of Avery and Zemski (1998). These observations are compounded by the fact that the speed of computer trading itself may make it less likely that events become common knowledge and market events may lack clarity.

Second, even if it turns out rather quickly that the dislocation was unwarranted, it might in the mean time have led to further indirectly related self-fulfilling feedback effects that cannot be reversed. For instance, one can imagine that as a result of a temporary crash some barriers are breached which in turn lead to further consequences, say to margin calls which cannot be honoured and therefore lead to a credit event.[15] Or the internal risk systems might flag up the

[15] A classical example is described in Clark (2010) who quotes comments by Andreas Preuss, CEO of Eurex, given to the Futures Industry Association Expo 2007 in Chicago as saying that "It is only a matter of time before such a scenario becomes reality, and it may already have happened, albeit on a smaller scale. In 2003, a U.S. fund was sunk in 16 seconds when an algorithmic application was accidentally switched on. The scary part wasn't how fast the company went down, it was that it took 47 minutes for the company to realize it, and call its clearing bank. That's a situation that we as an industry need to work against. The increasing speed with which trading is being conducted requires the ever-increasing speed of detection capabilities and reaction capabilities. To this end Eurex clearing

dual squeeze of capital losses and of higher risk, and therefore lead to forced fire sales of a variety of risky assets, which in turn contribute to further falls in the market and to heightened uncertainty, the so-called “endogenous risk.” In other words, even if events become common knowledge eventually or even quickly, in fast markets much irreversible damage can be created before this occurs. We say that markets are strongly *path-dependent* in such situations.

We have witnessed the erosion of open-outcry trading. On CME for instance, some contracts moved from open-outcry to a hybrid model with open-outcry during waking hours and electronic markets outside of those hours. Over the next 5 to 10 years, one may see further substitution of floor based trading in favour of computer based trading, rendering the effects detailed above more salient still.

8 Common Knowledge and Coordination in Distributed Environments

Interestingly, and confusingly, there also is the opposite possibility that a lack of common knowledge *removes* - not adds - some instability from the system. Consider a coordination game, defined to be any situation in which an individual player has an incentive to act in unison with the other players. Many coordination games have multiple equilibria. Such situations arise naturally in finance, for instance individual depositors deciding whether or not to run on a bank, or currency investors deciding whether or not to sell a pegged currency. If everyone else runs on the bank, then it is in my interest to run also, and if no one runs, then neither should I. In order to allow for perfect coordination, the underlying structure of the economy must be common knowledge (see Rubinstein (1989) for instance), for not only must we all know the state of the world, each one must also know that everyone else knows and so on. If there are grains of doubt about whether some of the depositors know that the others know that others still know etc, then perfect coordination is impossible. A *global game*, in a nutshell, introduces a bit of fuzziness about fundamentals into the picture which removes the common knowledge assumption, see Carlsson and Van Damme (1993) and Morris and Shin (2002). Under some conditions, it can be shown that there is a unique equilibrium in this slightly fuzzy economy. This equilibrium of the global game is of the threshold type: if the publicly observed signal is larger than some threshold, then we all run, if it is below then we do not, and this is individually optimal. In coordination games with common knowledge, the outcomes are unstable in as far as they are simply the outcome of the whim of people, and not grounded in any sense to observed fundamental values. So changing whims or sunspots lead to runs, while in global games, people cannot coordinate perfectly and therefore they run only if the fundamentals of the bank are bad enough to warrant a run.

As argued above, in impersonal computer based distributed trading environments, generating common knowledge seems to be impossible (see for instance Halpern and Moses (1990)), since there are no situations where all players have eye contact with each other and witness events as a group, finding themselves all together in the same room at the same time. This used to be possible to a higher degree when traders did gather physically in the trading pits, so in coordination game situations, violent perfectly coordinated jumps from no-run to run outcomes that are independent of any fundamental news are possible in pits but not in decentralised computer based environments.

---

members have been provided an “emergency stop button” so they don’t have to call the exchange if an algorithmic box goes berserk.”

9 Network Effects in Decentralised Markets

Given that algorithmic and high-frequency traders link alternative trading venues that form an arbitrage network through common or complementary listed securities, the question arises as to what the network structure of TVs linked through such low-latency traders looks like. Co-location and other costs linked to low-latency links make it implausible that all algorithmic traders collocate on all TVs.

Little is known in general about current arbitrage networks in finance (see Rahi and Zigrand (2010a) for a formal model on endogenous network formation), and next to nothing is known about how such networks will develop over the 5 to 10 years that form the foresight horizon of the current exercise. We do not know exactly which algorithmic traders are linking which exchanges empirically, nor have we got a solid understanding of the outcomes of such network formation games. Our knowledge of network effects and of the stability properties of various network architectures is incomplete as well. For instance, is a fully connected market – where each TV is linked to all other TVs through HFTs – more or less resistant to a large shock than a star network, say, whereby all TVs are linked to some primary exchange but no two TVs are linked directly, or indeed to a wheel whereby each trading venue is linked to exactly two other trading venues in a circular graph?

Some preliminary research papers (for instance Allen and Gale (2000) and Cabrales et al. (2010)), although pursuing very different avenues, model balance sheet links between financial intermediaries. The setting is therefore quite distinct from the one we have in mind here, but let me summarize in a sentence or two the gist of their findings. If the network is very dense, then it can withhold many small and medium shocks, but it will fail when the big one strikes. If a network is very sparse, some institutions will fail if smaller shocks occur, and since small shocks occur frequently, a very sparse network does not provide enough diversification. A medium-dense network may be most able to resist a mixture of small, medium and large shocks. The small or medium shocks that would frequently bring down one stand-alone institution will be unable to bring it down when it has links due the diversification additional diversification/mutualisation. The big one will still bring down large parts of the system, but since the network is not very dense, some of the components may resist the onslaught and keep on providing the financial services that society requires.

Some of those insights may inform our debate about fragmentation of the trading environment as well as of the clearing environment. Having one node only makes the node a systemic node. For instance, in light of the recent outages at the main UK exchange, the LSE, trading was able to continue on MTFs. Unfortunately, a preliminary reading of the last few outages on the LSE appears to indicate that only the most pressing trades are executed on MTFs when the main exchange is out of action. Using the formal definitions found in of Rahi and Zigrand (2010b), it would appear, counter-intuitively perhaps, that TVs in Europe form a network of complements as opposed to a network of substitutes, at least in periods of stress. A network of complements is shown to give rise to contagion, whereby a negative liquidity shock on one TV leads to further negative liquidity shocks on all complementary TVs (and to positive liquidity shocks on substitutable TVs).

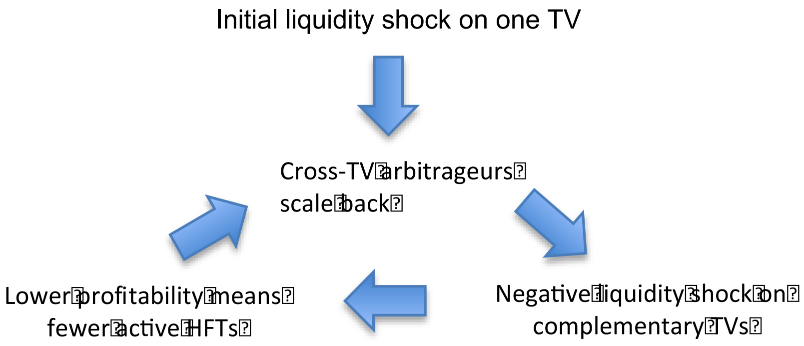

The themes of decentralised markets on one hand, and of the desire for coordination in markets that are lacking common knowledge (reviewed in a previous section) on the other hand, can now be joined. The discussion here is based on Zigrand (2005) and Zigrand (2003). A market participant whose trading strategies involve multiple securities that are executed electronically through algorithms faces the following situation. The desired trades for multiple securities are not independent one of the other. For instance, a well diversified portfolio may require fixed proportions invested in different securities, an arbitrage trade trying to benefit from a failure of the law of one price of one given stock across multiple TVs requires a purchase of a number of shares on the cheap TV and the sale of the same number of shares on another TV, and a delta-hedging algorithm may involve the sale of one option on one TV and the purchase of delta stocks on another TV. Since market or limit orders for one traded security cannot be made dependent on prices or quantities of any other security, and furthermore since the various securities that need to be traded may not be listed on the same TV anyhow, perfect coordination between the various single-stock algorithms is impossible, no matter how desirable coordination may be. In a spread trade for instance where two legs need to be implemented, even if one leg is successfully filled, the trade may still fail if the second leg is not – the so-called *leg-in risk.* Coordination here is across strategies belonging to the same investing entity, and coordination would be stabilising and welfare improving in this instance.

What prevents perfect coordination is the lack of common knowledge that arises from the fact that once the algorithm interacts with the limit order book for one security, it cannot change its behaviour instantaneously as a function of the contemporaneous information that occurs on any other limit order book, whether on the same TV or on another one. Of course, the investing entity can quickly cancel one order for one security on one TV if information comes in from some other order book, but since this requires some time, perfect coordination is impossible. It can even be the case that each algorithm "knows" that a certain event has occurred, and even knows that each algorithm knows that each other algorithm knows that a certain event has obtained, and yet the induction must stop at a finite level of depth. It follows that optimal strategies must from the outset build in robust mechanisms into any single algorithm to deal with the higher order uncertainty arising from market developments affecting any other single security algorithm.

The market outcomes of this impossibility to coordinate can be summarised as follows (Zigrand (2005)). First, markets suffer from a degree of allocational inefficiency since the marginal rates of substitution that are priced into different securities need not be equalised due to the limiting effect of the lack of common knowledge on trading quantities. This is the welfare cost that has already been described in more detail in section **Error! Reference source not found.**. Second, markets remain informationally inefficient since information does not get impounded simultaneously across all securities markets but rather is diffused with some delay through revised trading algorithms. Third, any one market price tends to overreact to news. The reason is that any one algorithm that belongs to the same investor tries to second-guess the information that any other algorithm of this same investor (and recall that these algorithms attempt as best as possible to coordinate) will have seen and acted upon. The optimal decision rule for any one security will overly rely on the information revealed by that one security (be it price or volume).

The proliferation of trading venues and clearing houses may well continue over the next 5 to 10 years, although the dynamics may be affected by regulatory changes that have the potential to undo some of the liberalisations that gave rise to the ATVs in the first place, and they also may be affected by mergers and acquisitions of exchanges and clearing houses.

10 A Sociological Approach to the Risk of Positive Feedback in Automated Trading

Although initially developed in the context of systems theory and cybernetics, the concept of positive feedback has long been deemed central to the sociology of markets. From a sociological standpoint, the rise of automated trading creates new possibilities for positive feedback, along with the corresponding dangers of market crashes.

Positive feedback was theorized in sociology in 1968 as Robert Merton's (Merton (1968)) theory of "self-fulfilling prophecies," that is, predictions that directly or indirectly cause themselves to be true. In the classic example, a run on a bank, when a large number of customers come to withdraw funds, their number and prominence trigger a rumour that lead others to follow suit, to the point that the bank lacks enough funds to meet all demands.

More recently, sociological research has built on Merton by emphasizing the role of tools and devices. In the past three decades, the sociological study of markets has evolved from a focus on the role of norms and beliefs (such as the belief that a bank is insolvent), to an interest on social networks and, more recently, to attention on the material devices that are employed by economic actors in their decisions. By "material devices," sociologists mean the decision-making tools that actors employ to lighten the cognitive load of processing information. Some examples of these tools are mathematical equations such as Black-Scholes, visualizations such as a spread plot and software such as an Excel spreadsheet or a commercial program to value credit default swaps. The importance of these material devices lies in that they allow actors with limited cognitive capacity to engage in calculative decision-making. By offering decision-makers the possibility of augmenting their cognitive ability, calculative devices help bring about the calculating decision maker hypothesized by orthodox economics. At the same time, the theory offers a twist: the outcome of these calculations will not simply be an ideal, impersonal "optimum," but it will be shaped by the ideas and interests that are inscribed in the specific device that is used for the purpose of the calculus.

This emphasis on materiality has provided a new lens to understand positive feedback. The notion of self-fulfilling prophecies, according to the influential work of French sociologist Michel Callon (Callon (2007)), is unsatisfactory in many ways. In the standard Mertonian setup, self-fulfilling prophecies entail an over-abstracted, almost tautological portrait of how crises happen. If a sufficiently large number of depositors fear a crisis, the run on the bank will surely happen. But as Callon asks, how do these beliefs arise in the first place? One answer might be that these beliefs are a shared convention. But this poses the additional question of how the convention arises in the first place, or comes to dominate competing conventions. The answer, Callon suggests, points to the material basis of belief formation. For instance, a line forming outside a retail bank branch can be enough to prompt fears of a bank run.

However, in a context of financial models and automated trading, the sightings of the naked eye or the judgments of floor traders have little bearing on the ultimate outcome. In a contemporary setting, the central issue at hand becomes the ways in which quantitative trading strategies, trading algorithms and other technological components of market engagement interact to create positive feedback Ð whether it is in the form of self-fulfilling prophecies or the reverse, “self-destructing prophecies.”

The recent stream of literature in the social studies of finance, and specifically the work of Donald MacKenzie, provides some important pointers on how this might happen. According to Donald MacKenzie, the crash of 1987 was an instance in which an automated trading strategy - portfolio insurance - became ineffective as it diffused among investors. As MacKenzie writes, “when portfolio insurance was small-scale, the assumption that the stock and futures markets were external ‘things’ in which prices would not be affected by the insurer’s purchases or sales was plausible enough (…) but by June 1987, the portfolios ‘insured’ by LOR and its licensees were sufficiently large that Leland was pointing out that “if the market goes down 3 per cent É we could double the volume of trading in the New York Stock Exchange” (MacKenzie (2004), pp. 312-3). Thus, in other words, the introduction of an automated trading strategy, portfolio insurance, contributed crucially to the crash because its widespread diffusion undermined the assumptions on which the model was initially premised.

A similar problem can be found with the rise of credit derivatives in the 2000s. As MacKenzie (2011) argues, the success of CDOs priced out the gatekeepers of mortgage quality in the market, thus opening the possibility for the mispricing of subprime mortgages. Again, the premise of the model behind a sophisticated (and initially correct) trading strategy was called into question by the very success of the strategy.

A similar dynamic might have been at play in the Flash Crash. One of the risks highlight by the recent Flash Crash is the possibility of positive feedback loops induced by automated technology. Specifically, the report by the Commission pointed that one reason why the initial order by the large fund had an outsized impact on the market was that the execution algorithms of this fund lock in with the trading algorithms of high frequency trading firms. The fund’s Sell Algorithm mistook the high volume of trading created by the high frequency traders as a sign of a high level of liquidity, thereby increasing the pace of selling without the corresponding buy-side depth. This was one factor that led to the sharp fall in prices.

The situation can be read in a similar light as the crises of 1987 and 2008. The practice of algorithmic execution is based upon the notion that the algorithmic slicing of a large order will conceal the social cues provided by the order itself, thereby reducing its impact. This takes place through the use of algorithms such as HVWAP, or Historical Volume Weighted Average Price, which adjust the pattern of purchases on the basis of a few selected variables. However, the Flash Crash suggests that the Sell Algorithm did not take into account the fact that its counterparties were other algorithms – high frequency traders. The widespread diffusion of algorithms undermined the basic idea behind the slicing of orders. The algorithm achieved the opposite effect it was intended to, increased the magnitude of the price impact and unleashed another phase of the crisis.

References

Allen, F. and Gale, D. (2000). Financial contagion. *Journal of Political Economy*, 108(1):1–33.

Angel, J. (1994). Limit versus market orders. Georgetown Working Paper.
Avery, C. and Zemski, P. (1998). Multidimensional uncertainty and herd behavior in financial markets. *American Economic Review*, 88(4):724–748.
Beunza, D. (2011). Impersonal efficiency and the dangers of a fully automated securities exchange. Technical report, BIS, Foresight Exercise on the Future of Computer Trading.
Cabrales, A., Gottardi, P., and Redondo, F. V. (2010). Risk sharing and contagon in networks. Technical report, EUI.
Callon, M. (2007). What does it mean to say that economics is performative? In Donald Mackenzie, F. M. and Siu, L., editors, *Do Economists Make Markets? On the Performativity of Economics*, pages 311–358. Princeton University Press.
Carlsson, H. and Van Damme, E. (1993). Global games and equilibrium selection. *Econometrica*, 61:989–1018.
Clark, C. (2010). Controlling risk in a lightning-speed trading environment. *Chicago Fed Letter*, 272.
Danielsson, J., Shin, H. S., and Zigrand, J.-P. (2010). Balance sheet capacity and endogenous risk. Mimeo, www.riskresearch.org.
Danielsson, J. and Zigrand, J.-P. (2008). Equilibrium asset pricing with systemic risk. *Economic Theory*, 35:293–319.
Donefer, B. S. (2010). Algos gone wild: Risk in the world of automated trading strategies. *Journal of Trading*, Spring:1–4.
Durbin, M. (2010). *All About High-Frequency Trading*. McGraw Hill.
Foroohar, K. (2010). Trading Êpennies Êinto Ê$7 Êbillion Êdrives Êhigh-frequency's cowboys. *Bloomberg*.
Gennotte, G. and Leland, H. (1990). Market liquidity, hedging, and crashes. *American Economic Review*, 00:999–1021.
Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology*, 83(6):1420–1443.
Halpern, J. and Moses, Y. (1990). Knowledge and common knowledge in a distributed environment. *Communications of the ACM*, 37(3):549–587.
Hunsader, E. (2010). Analysis of the flash crash, date of event: 20100506, complete text. *Nanex*, http://www.nanex.net/20100506/ FlashCrashAnalysis_CompleteText.html.
Jarrow, R. (1992). Market manipulation, bubbles, corners, and short squeezes. *Journal of Financial and Quantitative Analysis*, 27(3):311–336.
Joint CFTC-SEC Advisory Committee on Emerging Regulatory Issues (2010). Recommendations regarding regulatory responses to the market events of may 6, 2010. Technical report, CFTC, SEC.
MacKenzie, D. (2004). The big, bad wolf and the rational market: portfolio insurance, the 1987 crash and the performativity of economics. *Economy and Society*, 33(3):303–334.
MacKenzie, D. (2011). The credit crisis as a problem in the sociology of knowledge. *Forthcoming in the American Journal of Sociology*.
Merton, R. K. (1968). *Social Theory and Social Structure*. Free Press.
Morris, S. and Shin, H. S. (2002). Global games: Theory and applications. In M. Dewatripont, L. H. and Turnovsky, S., editors, *Advances in Economics and Econometrics, the Eighth World Congress*. Cambridge University Press.
Page, S. (2011). *Diversity and Complexity (Primers in Complex Systems)*. Princeton University Press.
Rahi, R. and Zigrand, J.-P. (2010a). Arbitrage networks. Mimeo, London School of Economics.
Rahi, R. and Zigrand, J.-P. (2010b). Endogenous liquidity and contagion in a world with multiple trading venues. Mimeo, London School of Economics.
Rubinstein, A. (1989). The electronic mail game: Strategic behaviour under almost common knowledge. *American Economic Review*, 79:385–391.
Schelling, T. (1978). *Micromotives and Macrobehavior*. Norton.

Shapley, L. and Shubik, M. (1977). Trade using one commodity as a means of payment. *Journal of Political Economy*, 85:937–968.
Shin, H. S. (1996). Comparing the robustness of trading systems to higher order uncertainty. *Review of Economic Studies*, 63:39–59.
The Staffs of the CFTC and SEC (2010). Findings regarding the market events of may 6, 2010, published september 30, 2010. Technical report, CFTC, SEC.
Zigrand, J.-P. (2003). Common knowledge, coordination and rational limits to arbitrage. Mimeo, FMG/LSE.
Zigrand, J.-P. (2004). A general equilibrium analysis of strategic arbitrage. *Journal of Mathematical Economics*, 40(8):923–952.
Zigrand, J.-P. (2005). Rational asset pricing implications from realistic trading frictions. *Journal of Business*, 78(3):871–892.
Zigrand, J.-P. (2006). Endogenous market integration, manipulation and limits to arbitrage. *Journal of Mathematical Economics*, 42(3):301–314.
Zovko, I. and Farmer, D. (2002). The power of patience: a behavioural regularity in limit-order placement. *Quantitative Finance*, pages 1–7.

Foresight
1 Victoria Street
London SW1H 0ET
www.foresight.gov.uk
URN: 11/1221