Government Office for Science

Foresight

Crashes and high frequency trading

The Future of Computer Trading in Financial Markets - Foresight Driver Review – DR 7

Contents

Crashes and High Frequency Trading

An evaluation of risks posed by high-speed algorithmic trading

D. Sornette and S. von der Becke

August 2011

This review has been commissioned as part of the UK Government's Foresight Project, The Future of Computer Trading in Financial Markets. The views expressed do not represent the policy of any Government or organisation**.**

Executive summary

We present a partial review of the potential for bubbles and crashes associated with high frequency trading (HFT). Our analysis intends to complement still inconclusive academic literature on this topic by drawing upon both conceptual frameworks and indicative evidence observed in the markets. A generic classification in terms of Barenblatt's theory of similarity is proposed that suggests, given the available empirical evidence, that HFT has profound consequences for the organization and time dynamics of market prices. Provided one accepts the evidence that financial stock returns exhibit multifractal properties, it is likely that HFT time scales and the associated structures and dynamics do significantly affect the overall organization of markets. A significant scenario of Barenblatt's classification is called "non-renormalizable", which corresponds to HFT functioning essentially as an accelerator to previous market dynamics such as bubbles and crashes. New features can also be expected to occur, truly innovative properties that were not present before. This scenario is particularly important to investigate for risk management purposes. This report thus suggests a largely positive answer to the question: "Can high frequency trading lead to crashes?" We believe it has in the past, and it can be expected to do so more and more in the future. Flash crashes are not fundamentally a new phenomenon, in that they do exhibit strong similarities with previous crashes, albeit with different specifics and of course time scales. As a consequence of the increasing inter-dependences between various financial instruments and asset classes, one can expect in the future more flash crashes involving additional markets and instruments. The technological race is not expected to provide a stabilization effect, overall. This is mainly due to the crowding of adaptive strategies that are pro-cyclical, and no level of technology can change this basic fact, which is widely documented for instance in numerical simulations of agent-based models of financial markets. New "crash algorithms" will likely be developed to trade during periods of market stresses in order to profit from these periods. Finally, we argue that flash crashes could be partly mitigated if the central question of the economic gains (and losses) provided by HFT was considered seriously. We question in particular the argument that HFT provides liquidity and suggest that the welfare gains derived from HFT are minimal and perhaps even largely negative on a long-term investment horizon. This question at least warrants serious considerations especially on an empirical basis. As a consequence, regulations and tax incentives constitute the standard tools of policy makers at their disposal within an economic context to maximize global welfare (in contrast with private welfare of certain players who promote HFT for their private gains). We believe that a complex systems approach to future research can provide important and necessary insights for both academics and policy makers.

Introduction

The way stocks are traded on stock exchanges has evolved enormously over time as a result of technological advancement and arguably changes in regulation[1]. Nowadays the majority of volume is traded electronically, based on systematic computer algorithms. The ultra-high-speed version of algorithmic trading, high frequency trading, is estimated to account for over 77% of transactions in the UK market according to Tabb Group[2], but there are lower estimates of about 25% for futures in 2010[3]. The May 6 flash crash in 2010, which saw the Dow Jones lose about 1 trillion USD of market value and individual stocks trading at fractions or multiples within minutes, put increased focus on HFT. Even though high frequency traders were subsequently mostly cleared from having caused the crash, doubts remain as to whether this new form of trading bears potentially destabilizing risks for the market.

Being a fairly new phenomenon, academic research on this subject is still limited in numbers and to some extent inconclusive with respect to potential risks posed by HFT. Advocates often point to HFT's role as liquidity providers, hence reducing transaction costs, lowering spreads and volatility. Several studies lend support to this view. Brogaard (2010) analyses the impact HFT has on US equities market and finds that high frequency traders add to price discovery, provide best bid offer quotes for most of the day and do not seem to increase volatility but may even reduce it. Studies on algorithmic trading (not necessarily at high frequency) also find it improves liquidity and price discovery (Hendershott and Riordan (2009); Hendershott, Jones and Menkveld, (2010)). Menkveld (2011) demonstrates that the success of a new equities market critically benefited from HFTs role as “new market makers”. Gsell (2008) simulates the impact of algorithmic trading on markets and finds indications that, for large trade volumes, algorithmic trading has a negative impact on market prices and that lower latency decreases market volatility. In contrast, Zhang (2010) finds a positive correlation between HFT and stock price volatility, which seems to be stronger in times of high market uncertainty. He finds that HFT hinders the incorporation of fundamental information into asset prices, causing stock prices to overreact to fundamental news. Smith (2010) finds that HFT have an increasingly large impact on the microstructure of equity trading and finds trades are showing significantly higher degrees of self-similarity.

An important point is that the focus on liquidity provision by HFT may be misguided. First, liquidity is not equal to volume. HFT arguably increases the volume of transactions. But volume is roughly the product of order sizes by their number per unit time (a kind of velocity). In the same way that the momentum of a body is the product of its mass by its velocity, volume can be large with just small order sizes contributing with a very large velocity or rate of transactions. Second, the hypothesis that HFT is a positive development is often based on the underlying assumption that more liquidity is necessarily good for investors and companies.

In the following section of this report, we will develop the above points and question the value of liquidity provided by HFT. We will develop the hypothesis that the utility of the real economy is likely to be an increasing function of liquidity, concave and asymptotically saturating to a plateau, and perhaps even decreasing when liquidity is many times the need of the real economy. It is conceivable that liquidity reaches a point beyond which the real economy does

---

[1] Reg NMS in the US and MiFID in Europe have been argued to have changed the incentive structure of stock exchanges.
[2] http://www.bloomberg.com/news/2011-01-24/high-frequency-trading-is-77-of-u-k-market-tabb-group-says.html
[3] http://www.futuresmag.com/Issues/2011/May-2011/Pages/Highfrequency-trading-Good-bad-or-just-different.aspx

not benefit anymore and where additional liquidity increases the risk of herding, of strong correlations, possibly leading to systemic instabilities and ultimately to crashes and their aftermath. Consequently, the third section of this article will examine risks posed by high-speed algorithmic trading. Preliminary evidence suggests the system of price formation to be indeed "non-renormalizable" which leads us to believe that HFT is likely to accelerate future market crashes. The systemic danger lays in the possibility of cross-excitations to other markets causing additional herding of low latency and/or fundamental traders. We believe that a complex systems approach to future research is crucial in the attempt to capture this inter-dependent and out-of-equilibrium nature of financial markets. Particularly relevant in assessing HFT risk is to understand their behaviour in conjunction with other market participants across various asset classes. Agent-based-modelling (ABM hereafter) offers a key method to improve our understanding of the systems' dynamics. Section 4 of this report will discuss the complex systems approach (and often associated agent-based modelling) and identify possible avenues for future research. Finally in section 5, we develop considerations for policy makers and regulators drawing upon our current analysis of the system.

2 Rethinking the value of liquidity

As shown above, several academic papers suggest that HFT increases liquidity in the market and advocates of HFT point to this liquidity provision as a key contribution to the well-functioning of financial markets. What is consequently overlooked is firstly that high frequency traders can also be significant liquidity takers and, secondly, that there are indications that larger liquidity increases herding effects and crashes thereby potentially reducing the value the real economy derives from liquidity above a certain threshold.

Market makers vs. market breakers

HFT can play an important role as market makers, for example, generating trading volume on new electronic exchanges (King and Rime, 2010; Menkveld, 2011). Trade volume, however, is not liquidity but all too often mistaken for it. Liquidity means "there is a bid/offer on the other side when I need it, for the amount I need it (market depth) at a reasonable level (market breadth). Volume is not the same as liquidity, since volume is approximately like the product of liquidity x velocity, and a large volume does not necessarily imply a large liquidity. This is illustrated by the May 6 flash crash when 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.

The same event also demonstrated that HF Traders can turn into significant liquidity takers[4]; while they are liquidity providers when it suits them (they have no obligation to make quotes[5]). This is also described as "flow toxicity", when market makers provide liquidity at their own loss or when informed traders take liquidity from uninformed traders[6]. In fact it seems HFT provides liquidity in good times when it is perhaps least needed and takes liquidity away when it is most needed, thereby contributing rather than mitigating instability (a point that will be discussed in the following section).

---

[4] "…it appears that the 17 HFT firms traded with the price trend on May 6 and, on both an absolute and net basis, removed significant buy liquidity from the public quoting markets during the downturn." (CFTC-SEC report on May 6, 2010, p. 48).
[5] We need to keep in mind, however, that even brokers with this obligation turn to sub quotes during the flash crash.
[6] See D. Easley et al. "Measuring flow toxicity in a high frequency world" (2010) for a development of a volume based metric to measure flow toxicity.

A recent report showed that the frantic development of HFT has slowed down in developed markets, and there is a transfer of activity to emerging markets such as Russia, Brazil and Mexico where exchanges are beginning to revamp their systems to attract such players. Low market volumes and stiff competition have led to a sharp fall in “high-frequency” (Grant, 2011). This illustrates the fact that, as HFT market participants flock into a given market, the opportunities shrink, dispelling the possibility for further growth.

It is also conceivable that HFT liquidity is provided at the expense of other market participants. Short term traders may be specifically prone to herd to the same information, driving the price further away from its fundamentals (Froot et al., 1992). The more momentum traders there are in a market and the higher the diversion from fundamentals, the fewer fundamental traders survive, further strengthening momentum traders. Various equilibria are possible between short and long-term investors[7]. The question is what is the right mix of investment strategies and horizons that best serves the well-functioning of financial markets and ultimately social welfare?

Liquidity vs. volatility

If the main argument for the benefits of HFT is liquidity, we need to ask whether liquidity is always “a good thing”, whether additional liquidity above a certain threshold becomes only marginally more useful or if too much liquidity can even be a “bad thing”. Finding answers to these questions is crucial for balancing potential benefits with potential risks associated with HFT.

Higher liquidity and higher trading volumes are generally associated with lower transaction costs, narrowing bid/ask spreads and thereby reducing volatility. Several studies have shown that HFT or algorithmic trading have improved market efficiency in this way. The real economy benefits from lower volatility as for example a company’s stock price is an indication of market confidence in the management of the company. In contrast, higher volatility is perceived as higher riskiness and may translate in higher funding costs, lower consumer, supplier and investor confidence. Investors will expect higher returns for higher volatility.

There is, however, evidence that there can be indeed too much so-called “liquidity” (actually trading volume). While volatility appears to be reduced at the level of individual stocks’ bid/ask prices, it may have amplified tail risk and increased volatility at the macro level. Dichev et al. (2011) analysed the effect of higher trading volumes and stock volatility and find that higher trading volumes can be destabilizing and produce “its own volatility above and beyond that based on fundamentals”. Interestingly, there appears to be an inflection point at which an increase in trading volume increases volatility to the extent that only a small circle of investors benefit and that “benefits to investors dominate at low to medium levels of trading”. In similar vein, Haldane recently pointed to the danger of normalising deviance at the micro level and concluded suggesting “thinner technological slices may make for fatter market tails. Flash Crashes, like car crashes, may be more severe the greater the velocity” (Haldane, 2011).

The flash crash of May 6, 2010 started in one of the most liquid markets, the E-Mini S&P 500 futures contracts. What is the role of liquidity or its scarcity in the occurrence of crashes? On the one hand, one can argue that deep markets should absorb new players more easily. On the other hand, it could be possible that deeper markets are more prone to pandemics as their

[7] See A.G. Haldane “Patience and Finance” (2010) for a good discussion of this point.

impact and connection to other markets is larger. It turns out that crashes occur in both types of markets. The fact that the flash crash started in one of the most liquid markets provides additional support for the hypothesis that flash crashes are not incompatible with large and deep markets. 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. In this perspective, it is important to distinguish this type of flash crash in the presence of large volume (herding effect) and the more localized ones occurring in single stocks with low volume (liquidity effect). Indeed, there is also evidence of crashes in not-so-liquid stocks. It seems single stocks are equally prone to crash but have less pandemic like consequences. The linking of assets through derivative instruments amplifying correlations is an important factor which will be discussed in more detail later in this document.

This leads to the question of whether there is a relationship between "incidents", large volatility events, market disruptions, flash crashes on the one hand and the volume of transactions on the other hand. We conjecture that there exist regimes in which larger so-called liquidity (actually volume) leads to larger extreme risks. We would like to see plotted some measures of extreme risks as a function of volume of transactions. This would give an important insight and is suggested as a line of future actions.

Summarizing, we formulate the hypothesis that the utility of the real economy could be an increasing function of liquidity with decreasing marginal gains, that is, with a concave shape and asymptotically saturating to a horizontal plateau for large liquidity. There could even be a decline in utility through possibly negative consequences for very large liquidity (or volume), such as the risks associated with e.g. increased volatility and crash risks. We suggest that testing this hypothesis has large priority in order to pose the problem in its core fundamentals.

Finally, the utility derived from liquidity provided by HFT could be argued to be lower than from other market participants. Why? Because HFT does not absorb risks. If a fundamental trader sells a position as on May 6, it can only be absorbed by counterparties wanting to be long. HFT books are flat by the end of the day. They carry no inventory; there is no transfer of risk apart for some milliseconds. HFT are opportunistic, they arbitrage what is referred to as "inefficiencies", but may often result from differences in time scales and technology. It remains to be seen if liquidity is a real robust externality of the behaviour of HFT.

The "law of the few"

Another aspect in assessing the value of liquidity could be the concentration of liquidity providers. If from around 12,000 market participants, 30% of liquidity is provided by 15 participants, the liquidity is less reliable than if it was more evenly distributed[8]. It is probable that Zipf's law applies here[9]. At least with a rough approximation, the number of participants contributing a liquidity of L or larger is roughly proportional to 1/L. This means in general that a few largest participants contribute a major fraction. In addition, there could be "dragon-kings" that correspond to even more concentration (Sornette, 2009). An important question is how does the concentration evolve as a function of time, by following secular trends versus short-term adaptive transients, or both? Obtaining data on trading volumes of individual players

---

[8] See table "Summary statistics of E-mini traders" from the CFTC-SEC report on market events of May 6, 2010.
[9] The evidence for concentration is fuzzy, and it is not clear that investment banks and the biggest funds are the major players. A possible clue is provided by the advertisement for HFT programmers by headhunting agencies, which are almost all for "a top investment bank" (http://www.wilmott.com/messageview.cfm?catid=5&threadid=83435 or http://www.wilmott.com/messageview.cfm?catid=5&threadid=85840).

would be very meaningful. Some sources suggest that HFTs make up 2% of approximately 20,000 trading firms in the U.S. and account for about 60-70% of equity trading volume. These are suggested to include a small number of investment banks, less than 100 hedge funds and hundreds of specialist prop shops[10]. This suggests that Zipf's law applies (Saichev et al., 2009), as we can presume that the few investment banks are the largest players and the many prop shops the smallest ones with specific hedge funds in the middle. So far we cannot prove it.[11] In the presence of the scarcity of data, we propose to use Zipf's law as a normative model, rather than an empirical question at this stage (which will be falsified as data becomes available). Indeed, Zipf's law is so well verified in many contexts that it can be used as a guide in the absence of sufficient data (Saichev et al., 2009).

3 Understanding the risks of HFT

3.1 Similarity classification of possible HFT regimes

In order to gain a better appreciation of the nature of the possible impact that HFT has on price dynamics, we can frame the question as follows: How do the different properties of price dynamics change as the rate of trading increases or, equivalently when the time τ between trades shrinks to smaller and smaller values. In other words, is calendar time playing a role? It is well documented that financial time series exhibit self-similarity properties, to a first approximation (Mandelbrot and Hudson, 2006; Mantegna and Stanley, 2007; Calvet and Fisher, 2008). This means that time units are arbitrary and the properties are related at all time scales. In details, of course, this simple "fractal" or scale invariance view is naïve and needs updating, as for instance with the development of the most sophisticated multifractal models (see Filimonov and Sornette (2011) and references therein). In addition, with the development of HFT, the issue of how microstructure impacts large scale properties (at the time scale of minutes, hours, days, weeks, months and years) is of paramount importance and a priori not trivial.

To make the discussion more precise, let us consider some statistical property denoted $P(\{r\}, t, \tau, \theta)$ of the price dynamics $\{r\}$ at time scale t, given a minimum time scale for trading τ and the presence of other unspecified control parameters θ. The time scale τ results from physical and regulatory limits that control the smallest time scale at which transactions can occur (we neglect for simplicity that in reality there will be a distribution of time scale; here τ is the typical micro-time scale). It can be called the ultraviolet (UV) cut-off (using the physical analogy that a short period is the same as high frequency, and the UV spectrum is in the high frequency range above visible light). As τ approaches zero (as with the development of HFT that reduces the waiting time between trades to milliseconds or lower), there are three possible impacts this could have on price dynamics.

As $\tau \Rightarrow 0$, $P(\{r\},t, \tau, \theta)$ converges to a function $P^*(\{r\}, t, \theta)$ that is independent of τ. This means that the UV cut-off is not important: the characteristic of price dynamics and other important properties are the same at the minute, hourly, daily time scales if the frequency of trades is per second or per millisecond. This is called "complete similarity of the first kind" in the classification of Barenblatt (1996). Concretely, suppose that, at some time T, by innovation and

---

[10] Rob Iati from TABB Group in "The real story of trading espionage": http://advancedtrading.com/algorithms/showArticle.jhtml?articleID=218401501

[11] In the UK, HFT make up 77% of transactions and 35% of turnover according to TABB group. They estimate that there are between 35 and 40 independent HFT firms http://www.advancedtrading.com/articles/229100205.

via the introduction of new technologies, the UV cut-off is reduced from τ1 to τ2 that is much smaller than τ1, say by a factor of 10 or 100 or 1000. Then, the characteristics of the price dynamics at time scale t > τ1 are unchanged. This implies that HFT should not have any impact if "complete similarity of the first kind" holds. Only the properties at time scales between τ2 and τ1 may be novel. Concretely, suppose that t1=1 second and t2=1 millisecond. The above claim means that the properties of market price dynamics at scales above 1 second are unchanged by the introduction of HFT below the 1 second time scale. In the language of Barenblatt's (1996) classifications of similarity, such "complete similarity of the first kind" holds in this instance, which implies that HFT would have essentially no effect on the price dynamics. The microscopic time scale τ is irrelevant to the organization of the financial markets.

The second possible regime is called "complete similarity of the second kind", which holds if, as τ⇒0, the properties embodied in P({r}, t, τ, θ) converge to the product f(t/τ) P*({r}, t, θ). Note that the presence of the UV cut-off τ only appears via some function of the dimensionless ratio t/τ. This means that, in absence of physical or regulatory constraints, the dynamics at a given time scale should consequently be similar to that of any other time scales and HFT should not have any significant impact, apart from a rescaling of the properties that can be absorbed into a suitable regulatory or risk management framework. In other words, as τ⇒0, time accelerates and the movie of the price dynamics is just running faster. The implication is that one should observe more often bubbles and crashes, in proportion to an acceleration of the rate of trading captured by the function f(t/τ). This scenario suggests that HFT adds to market instability *per unit of calendar time* and increases significantly the probability of crashes and crises. Indeed, consider the limit where HFT accounts for all the volume of transactions. Then the whole market is moving at the HFT rate. Now, imagine a movie in which you slow down frame by frame. Then, HFT slows down and become low frequency trading, such as daily trading. If the correspondence is 1 second of HFT corresponds to 1 day of low frequency trading in 1962, say, then one crash per year in 1962 would corresponds to one crash every 4 minutes in HFT time!

In the third scenario, P({r}, t, τ, θ) does not converge as τ⇒0. This means that the presence of a minimum time scale τ has profound consequences for the organization and time dynamics of market prices. This regime is called "non-renormalizable" in mathematics and physics and is for instance the case of hydrodynamic turbulence, for which the limit τ⇒0 leads to multifractal properties. Given the significant evidence of the multifractal properties in financial stock prices (Calvet and Fisher, 2008), it is thus plausible that the HFT time scale and the associated structures and dynamics do cascade and affect profoundly the overall organization of markets. In this "non-renormalizable" regime, HFT can be expected to have significant impact on market price dynamics, essentially functioning as an accelerator to previous market dynamics such as bubbles and crashes. New features can also be expected to occur, truly innovative properties that were not present before. This scenario is particularly important to investigate for risk management purpose.

3.2 Destabilizing effects of HFT

The present section aims to present several observations that could provide clues as to what kind of regime we are faced with when analysing the risks of HFT. Are the underlying algorithms in fact stabilizing or the reverse?

Indications of herding during the May 6, 2010 flash crash

On May 6, 2010, HFT created a “hot potato effect”[12]. Paired with other sellers, this translated in a negative spiralling effect. This can give a hint to what can happen if HFTs start trading mainly among each other. Why did they keep buying? Minority agent-based games in the crowding phase provide examples of such behaviour, in which buys and sells alternate as a result of the anti-correlation between recent individual actions and aggregate behaviour (Challet et al., 2005; Coolen, 2005).

According to the CFTC-SEC report, it was the algorithm of a mutual fund and not HFT that created a negative feedback loop on “volume” during the May flash crash. As it sold at 9% of previous volume, this was to a large degree absorbed by HFT, who later needed to sell their net long position, thereby increasing volume and selling pressure. As mentioned previously, this is an example of when “volume” is mistaken for liquidity. However, in our view, the event as described in the report was indeed a problem created through an algorithm of a fundamental trader, but it was then amplified by the strategic behaviour of the HFT.

Nanex analysed the trade flows on May 6, 2010 and found that it was aggressive high frequency selling that “would often clear out the entire 10 levels of depth before the offer price could adjust downward. As time passed, the aggressiveness only increased, with these violent selling events occurring more often, until finally the e-Mini circuit breaker kicked in and paused trading for 5 seconds, ending the market slide”[13]. Nanex also reports that the algorithm from the mutual fund was mostly active after stocks had already fallen. This should not be misconstrued as removing the influence of the mutual fund, since its algorithm was targeting volume in its strategy to pass its larger order.

The second liquidity crisis on May 6, 2010, in stocks, seems to have been impacted more by fundamental and discretionary traders than by HFT. Trading of HFTs after 2.45pm was back to previous levels, nevertheless many stocks saw drawdowns of 60%. How does this fit in with the findings of the E-mini crash? This is an illustration of the interplay between an overall climate of uncertainty with the on-going Greek crisis developing since April 2010, and panic herding with pro-cyclical mutual excitations between HFT and the rest of the investor population. Preliminary unpublished calibrations of self-excited Hawkes processes performed in our group suggest that the self-excitation component (or viral epidemic) of trades was indeed abnormally large during May 6, 2010 compared to other trading days. This supports the hypothesis that HFT may have a destabilizing effect through its endogenous self-excitation nature within the (small) pool of participants.

On May 6, 2010, the fundamentals in the stock market were shaky (Greece, flight to quality, higher volatility) so we cannot claim that a purely technical issue caused the E-mini and later individual stocks to tank. Perhaps, there needs to be a general instability and increased volatility in the market, which is then pushed over the edge through HFT. This is in line with the understanding of the authors of this report concerning the causes of financial crashes: a proximate cause triggers the crash that is rooted fundamentally on the existence of an intrinsic

[12] Between 2:45:13 and 2:45:27, HFTs traded over 27,000 contracts, which accounted for about 49 per cent of the total trading volume, while buying only about 200 additional contracts net. (CFTC-SEC report on May 6, 2010, p. 15).
[13] Nanex Report on May 6th Flash Crash http://www.nanex.net/FlashCrash/FlashCrashAnalysis.html

instability that has matured progressively, preparing the field for major disruptions that are triggered by local or proximate causes (Sornette, 2003).

General arguments for the existence of herding

The Kauffman Report (Bradley and Litan, 2010) points out the unprecedented level of correlation in equity markets and identifies ETFs as a major destabilizing factor, more so than HFT[14]. JP Morgan concludes that the high level of correlation can be explained by "macro-driven environment, record use of index derivatives such as futures and to a lesser extent ETFs, and high-frequency trading". They point out that "the share of futures and ETFs steadily grew over the past five years, and is now ~140% of cash equity volume (i.e., futures and ETFs are roughly ~60% of all equity volumes – perhaps not a coincidence that realized stock correlation is ~60%). The growth in index volumes coincided with a rise in correlation over the past ten years. More importantly, the growth of index volumes is directly driving excess market correlation (levels of correlation above the levels implied by macro volatility)"[15]. The reason why J.P. Morgan gives relatively less weight to the impact of ETFs on correlation than Kauffman could be explained by the fact that J.P. Morgan differentiates between broad based ETFs and specialist ETFs.

When different instruments (equity ETFs/Futures) are connected to the same underlying (stocks), it seems logical that there would be mutual excitation and increased inter-dependencies (correlations). Our own (still unpublished) research at the Financial Crisis Observatory (www.er.ethz.ch/fco/) supports the hypothesis that bubbles and crashes are in general mutually excited, with interesting dynamics and fluxes in different sectors. This occurs not only at times of great crashes but also for events developing over 6 months characterized with loss amplitudes of 10 percent or so. This supports the notion that HFT in broader market indices may pose greater systemic risk.

As 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. Froot et al. (1992) find that "if speculators have short horizons, they may herd on the same information, trying to learn what other informed traders also know. There can be multiple herding equilibria, and herding speculators may even choose to study information that is completely unrelated to fundamentals."

Finally, we note that HFT generally has stop losses built in the algorithms plus human oversight that can withdraw trading in critical market circumstances altogether[16]. This can be seen as a stabilising factor for the HFTs and a mechanism to mitigate the risk of herding.

---

[14] http://ftalphaville.ft.com/blog/2010/11/08/397431/kauffman-etfs-are-the-problem-not-hft and http://www.kauffman.org/uploadedFiles/etf_study_11-8-10.pdf

[15] J.P. Morgan „Why we have a correlation bubble" (5 October 2010).

[16] 6 of the 12 HFTs scaled back their trading during some point after the broad indices hit their lows at about 2:45 p.m. Two HFTs largely stopped trading at about 2:47 p.m. and remained inactive through the rest of the day. Four other HFTs appear to have each significantly curtailed trading for a short period of time, ranging from as little as one minute (from 2:46 p.m. to 2:47 p.m.) to as long as 21 minutes (from 2:57 p.m. to 3:18 p.m.). (CFTC-SEC report on May 6th, 2010, p.45).

Observations of mini-crashes

Mini-flash-crashes in single stocks seem to happen rather frequently[17]. In the case of the mini crash of "Progress Energy" in September 2010, the company reported a "big ramp-up in trades, hundreds of trades a second" for an otherwise "sleepy stock", a fact that points to HFT[18]. Also an interesting account of events is the day when Apple saw a mini flash crash – followed by other tech stocks and even a sharp move in USD/Yen[19]. The involvement of HFT is not evident. However, other mini flash crashes, for example April 28[20] or 27 September 2010[21] also seem to have been accompanied by increased frequency in quotes.

Most of HFT trading is currently taking place in stocks, futures and options. Even though the focus so far has been mostly on equity markets, there is no reason other markets will not become equally vulnerable to mini-crashes. It seems that we already see this happening. For example, on March 9, 2011 the price of Cocoa plummeted 12.5% in less than a minute. Raw sugar dropped 6% in 1 second on February 3 this year and trading in cotton was halted several times[22]. Mini-crashes (or flash rises) have also been observes in the currencies markets although, as often, the link to HFT is not clear[23].

Nanex identifies "thousands" of mini-crashes over the last few years. We are not sure though how useful this is because in order "To qualify as a down-draft candidate, the stock had to tick down at least 10 times before ticking up -- all within 1.5 seconds and the price change had to exceed 0.8%". We are probably looking for more severe impacts. Moreover, the frequency of those occurrences has not increased over time. This makes it hard to link to HFT[24].

Our experience with other complex systems, including earthquakes and epileptic seizures, suggests that one should not discard the precursory information of such smaller events that can announce the "great ones" (Helmstetter and Sornette, 2003; Ouillon and Sornette, 2004; Osorio et al., 2010; Sornette and Osorio, 2010). They may also be symptomatic of more structural changes, as argued in a recent work of our group concerned with nuclear risks (Sornette et al., 2011).

Evidence of operational risk

Another important risk dimension is associated with so-called operational risks, resulting from infrastructure disruptions, computer bugs, hacking and others, collectively known as "cyber-risks". Such cyber-risks are now at the top of concerns of the 500 largest companies in the USA, according to a study by Swiss Re in 2006. Maillart and Sornette (2010) report the first quantitative analysis of cyber-risks, quantified by the number of identity thefts per event. They find a power law distribution with tail exponent equal to 0.7, implying that the variance and the mean of the losses do not exist mathematically, and that larger and larger risks are expected to surface in the future. This corresponds to the regime of "wild" risks, according to the

[17] http://ftalphaville.ft.com/blog/2010/09/28/354876/the-apparent-rise-of-the-mini-flash-crash or http://wallstreetpit.com/61727-apple-aapl-mini-flash-crash-renews-market-structure-worries or http://www.observer.com/2010/wall-street/mini-flash-crash-nucor-shares
[18] http://www.nytimes.com/2010/11/09/business/09flash.html
[19] http://www.ft.com/cms/s/0/ddf14a8a-506f-11e0-9e89-00144feab49a.html#axzz1Ifuk6rCi
[20] http://www.nanex.net/FlashCrash/FlashCrashAnalysis_042810_MiniFlashCrash.html
[21] http://www.nytimes.com/2010/11/09/business/09flash.html?pagewanted=1&_r=1
[22] http://www.ft.com/cms/s/0/dbfb15d6-4a83-11e0-82ab-00144feab49a.html#axzz1Ifuk6rCi
[23] http://ftalphaville.ft.com/blog/2010/10/25/381381/flash-a-ha-hell-strike-everyone-of-our-security-markets/ or http://www.ft.com/cms/s/0/ddf14a8a-506f-11e0-9e89-00144feab49a.html#axzz1Ifuk6rCi
[24] http://www.nanex.net/FlashCrashEquities/FlashCrashAnalysis_Equities.html

terminology of Mandelbrot, for which conventional methods are invalid. One should thus keep in mind the possible collision and interplay between the endogenous dynamics of investors that can lead to crashes and such operational risks. In the case of “Infinium Capital Management”, HFT is blamed for a significant market disruption in the Oil Market and a $1 million loss within a few seconds of trading. Some highlights from a Reuters report: “The algorithm was turned on at 2:26:28 p.m. (Eastern) on Feb. 3, less than four minutes before NYMEX closed floor trading and settled oil prices. It immediately started uncontrollably buying oil futures… Infinium placed 2,000 to 3,000 orders per second before its flooded order router "choked" and was "dead in the water" a few seconds later… Infinium's burst of buying and selling represented about 4 percent of average daily trading volume in the contract, and caused a brief 1.3 percent jump in oil prices, from $76.60 to $77.60, before settling at $76.98, Reuters’ data show. Trading volume spiked nearly eight-fold in less than a minute”[25] Infinium blamed the mistake on a “broken algo”, a flaw with the computer that did not properly record the order. This could be an example of “operational risk” or an example of deliberate market manipulation such as “flash orders” or “banging the close”, something HFT is frequently accused of.

This leads us naturally to the issue of “market manipulation”. A market can also be manipulated through rumours, insider trading, etc. However, the algorithms seem to offer very specific new ways of manipulation. Remember an ex Goldman Sachs employee being arrested by the FBI because of fears that the stolen algorithm could be used to manipulate markets. Why do we worry so much about 32MB of this specific code but not about all the codes that are still with and in constant development at Goldman Sachs?

Synthesis and recommendations

Some defenders of HFT point to the fact that we already had a flash crash in May 1962, before high speed trading was invented. There are indeed some parallels[26]. Even though this shows flash crashes can occur without HFT, the question is whether HFT has increased the likelihood of flash crashes occurring (either frequently in single stocks with minor market impact or less frequently in large, connected markets (Futures/ETFs) with larger impact). The above presented observations lead us to conclude that this is indeed the case.

If nothing else, HFT can be understood as accelerating time, so to speak. Indeed, take in consideration that HFT accounts for most of the volume. Take the limit where it is all the volume. Then the whole market is moving at the HFT rate. Now, imagine a movie in which you slow down frame by frame. Then, HFT slows down and becomes low frequency trading, such as daily trading. If the correspondence is 1 second of HFT corresponds to 1 day of low frequency trading in 1962, say, then one crash per year in 1962 would corresponds to one crash every 4 minutes in HFT time ! This reasoning is of course naïve and misses a lot of ingredients, but it nevertheless captures what is a key aspect of the problem. By definition and intrinsically by its time-acceleration nature when it dominates the trading volume, HFT will give many more crashes per unit calendar time (not per unit transaction time or HFT time).

What is needed is a better understanding of the relationship between trade volume and systemic risk. Is there a “right amount of liquidity”? What is the right mix of trading strategies that maximises social welfare? We also need a better understanding of the interplay between different trading methods (e.g. high frequency, low frequency, fundamental, technical) and

[25] http://af.reuters.com/article/energyOilNews/idAFN2511929020100825
[26] http://online.wsj.com/article/SB10001424052748703957604575272791511469272.html

investment instruments (e.g. stocks, ETFs, futures) and markets (e.g. equities, FX, commodities). In our view, current research does not make the distinction between these factors explicit enough. We believe that in the quest for answers to those questions, much is to be gained by a complex systems approach to future research as we will argue in the following section.

4 A complex systems approach to future research

4.1 Brief summary of main insights about crises from agent-based models

Key to understand financial market risk of various players is to gain a better understanding of the interactions of various market participants and the resulting patterns emerging in the markets. We propose that greater transparency allowing the identification and observation of market participants should be used to improve monitoring and early warning systems. These warning systems should be based on models that capture the complex evolutionary nature and the non-linearity of financial markets (Hommes and Wagener, 2009; Evstigneev, Hens and Schenk-Hoppe, 2009). Traditional macro-economic models such as the DSGE (dynamic stochastic general equilibrium) models, based on the assumption of rational agents that are drawn to equilibrium, have proven insufficient in forecasting or identifying large systemic financial risks. Agent-based models (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. Harras and Sornette (2011) used ABM to study the emergence of bubbles and consequently crashes in financial markets. Their findings demonstrate how feedback mechanisms lead to the development of transient collective herding regimes resulting in unsustainable high prices that are then corrected by a crash. They state that “Paradoxically, it is the attempt for investors to adapt to the current market regime which leads to a dramatic amplification of the price volatility. A positive feedback loop is created by the two dominating mechanisms (adaptation and imitation) which, by reinforcing each other, result in bubbles and crashes”. These findings applied to HFT demonstrate how “learning algorithms” can indeed be expected to accelerate the formation of bubbles.

Chiarella, Dieci and He (2009) emphasise the importance of agent heterogeneity and the formation of expectations through economically intuitive rules of thumb. They find that bounded rational heterogeneous agent (BRHA) models are “able to accommodate market features that seem not easily reconcilable for the standard financial market paradigm, such as fat-tail behaviour, volatility clustering, large excursions from the fundamental, and bubbles”.

Minority games (Challet et al., 2005; Coolen, 2005), and more generally first-entry games, can be taken to be simplified models of investors trying to be first movers either to buy in or to sell out before their competitors. These ABM show clearly that agents who try to optimize their utility function tend to crowd in with similar strategies that have been recent winners. This leads in general to enhanced financial volatility and sometimes to big swings. It seems indeed that adaptive and learning algorithms interacting by buying and selling on the same market tend to develop collective dynamical modes that are prone to large moves.

One of the most pressing subjects is to come up with a realistic agent-based model where crisis and complexity arise from simple rules and interactions in a universal way, robust against specific assumptions. Indeed, current economic theories are inadequate and one of the most promising alternative of today is agent-based modelling. From there, we need to build policy making devices (a "policy wind tunnel", as Nigel Gilbert would call it, or an "economic flight simulator"). In the limit of large system sizes, universality classes may appear, and we need to identify and classify them for economic systems. But the high dimensionality of agent-based models makes them hard to calibrate and validate. There is an extreme sensitivity to control parameters, a single factor can matter ("butterfly effect") and we need to pay careful attention to the calibration-overfitting-validation problem (Sornette et al., 2007; 2008; Satinover and Sornette, 2011).

4.2 Financial markets as truly "complex adaptive systems"

We argue that markets and economies in general are truly "complex systems" in a technical sense. As such, they are intrinsically characterized by periods of extremity and by abrupt state-transition and spend much time in a largely unpredictable state. The world of financial engineering has seen extraordinary growth in recent years and many extremely sophisticated methods have been developed to assess and distribute risk, largely by the development of new instruments derived from packaging and then re-dividing large number of underlying simpler instruments. In addition, technology innovations, such as HFT, have led to new markets and new opportunities. While many of these developments are quite sophisticated, they generally fail to incorporate the crucial insight that financial markets and economies have become very complex systems. Financial instruments are designed making a convenient and desirable assumption of independence. Complex systems typically contain many instances of hidden interdependences, tight couplings and other subtle (and inconvenient!) features (Satinover and Sornette, 2010).

i. Complex systems are usually open and dynamic - the underlying components of the system are in flux. Nonetheless, complex systems usually demonstrate stability of patterning, which lends itself to a mistaken presumption of equilibrium, as in classical economics and control theory.

ii. Most frequently the stability of patterning may be considered a "meta-stability" - it includes multiple quasi-stable states with dynamic, abrupt and difficult to predict transitions among states.

iii. Complex systems often have a memory. The future state depends not only on one or more preceding states but upon the dynamic sequence of preceding states - i.e., they demonstrate path dependency. This feature lends these systems the power of self-organization.

iv. Complex systems often consist of nested hierarchies of smaller-scale complex systems. This is most evident in the neurobiology of the brain where, e.g., cortical brain tissue forms a self-organizing complex systems "sheet," but is itself at a lower level composed of cortical processing units which "compute" an output passed up to the higher level. The cortical units are themselves composed of individual units (neurons) whose computational capacity arises from the complex systems nature of their internal components. In the other direction, economies and markets may be thought of as composed of many individual brains, or agents (Satinover, 2002).

v. Complex systems often yield outputs that are emergent: the interactions among agents/individual units may be deterministic, but the global behaviour of the system as a whole conforms to rules that are only rarely deducible from knowledge of the interactions and topology of the system. Financial services systems are intractable, which means that it is impossible fully to specify them.

vi. In complex systems, the relationship between input and output is typically non-linear so that a small perturbation may yield a very large overall disturbance while a large perturbation may be absorbed with little or no effect. Complex systems are typically exquisitely sensitive and at the same time resilient, in ways that are difficult to predict.

vii. Complex systems are also characterised by having both negative (damping) and positive (amplifying) feedback loops. The output of the system alters the nature of the (next) input.

4.3 Financial bubbles and crashes: implications of and for HFT

Let us come back to the simple picture according to which HFT corresponds approximately to an accelerated movie of the financial time series prior to its existence, as described above. Given that algorithms involved in HFT adapt and learn (and therefore imitate) somewhat similarly to human investors, we can draw on our previous research on financial market crashes to suggest that financial instabilities can be expected to flourish in a world dominated by HFT (Jiang et al., 2010; Johansen and Sornette, 2010; Sornette, 2003; Sornette and Johansen, 2001; Sornette and Zhou, 2006). Pro-cyclicality mechanisms, also known as positive feedbacks, are numerous in a world in which automated hedging strategies are implemented. This leads to unsustainable regimes, ending in crashes and crises.

Specifically, our previous works support the proposition that (i) the presence of a bubble can be diagnosed quantitatively before its demise and (ii) the end of the bubble has a degree of predictability. We hypothesize that the same holds true, probably to an even larger degree, for instabilities occurring at the intraday HFT time scales. This opens the road for systematic studies and a large research program. Of course, these two claims are highly contentious and collide against a consensus both in the academic literature (Rosser, 2008) and among professionals. For instance, in his recent review of the financial economic literature on bubbles, Gurkaynak (2008) reports that “for each paper that finds evidence of bubbles, there is another one that fits the data equally well without allowing for a bubble. We are still unable to distinguish bubbles from time-varying or regime-switching fundamentals, while many small sample econometrics problems of bubble tests remain unresolved” (page 1). Similarly, the following statement by former Federal Reserve chairman Alan Greenspan (2002), at a summer conference in August 2002 organized by the Fed to try to understand the cause of the ITC bubble and its subsequent crash in 2000 and 2001, summarizes well the state of the art from the point of view of practitioners: “We, at the Federal Reserve recognized that, despite our suspicions, it was very difficult to definitively identify a bubble until after the fact, that is, when its bursting confirmed its existence. Moreover, it was far from obvious that bubbles, even if identified early, could be pre-empted short of the Central Bank inducing a substantial contraction in economic activity, the very outcome we would be seeking to avoid.”

To break this stalemate, Didier Sornette, together with Anders Johansen (from 1995 to 2002), with Wei-Xing Zhou (since 2002 (now Professor at ECUST in Shanghai)) and with the FCO group at ETH Zurich (since 2008, www.er.ethz.ch/fco) have developed a series of models and techniques at the boundaries between financial economics, behavioural finance and statistical physics. Our purpose here is not to summarize the corresponding papers, which explore many different options, including rational expectation bubble models with noise traders, agent-based models of herding traders with Bayesian updates of their beliefs, models with mixtures of nonlinear trend followers and nonlinear value investors, and so on (Sornette (2003) and references therein for the period 2002 and the two recent reviews in Kaizoji and Sornette (2010) and Sornette and Woodard (2010) and references therein). In a nutshell, bubbles are identified as “super-exponential” price processes, punctuated by bursts of negative feedback spirals of crash expectations. These works have been translated into an operational methodology to calibrate price time series and diagnose bubbles as they develop. Many cases are reported in Chapter 9 of the book (Sornette, 2003) and more recently successful applications have been presented with ex-ante public announcements posted on the scientific international database arXiv.org and then published in the referred literature, which include the diagnostic and identification of the peak time of the bubble for the UK real-estate bubble in mid-2004 (Zhou and Sornette, 2003), the U.S. real-estate bubble in mid-2006 (Zhou and Sornette, 2006), and the oil price peak in July 2008 (Sornette et al., 2009).

Kindleberger (2000) and Sornette (2003) have identified the following generic scenario developing in five acts, which is common to all historical bubbles: displacement, take-off, exuberance, critical stage and crash. Applied to HFT, the development of unsustainable mispricing follows similar pro-cyclical mechanisms, in particular amplified by the tendency for algorithms and technical trading that dominates at short time scales to crowd.

In fine, let us conclude this subsection by mentioning that preliminary calibrations of intraday high-frequency price time series with the bubble-diagnostic model developed in our group (see for instance Jiang et al. (2009) and Filimonov and Sornette (2011)) do support the evidence for the presence of bubble-like behaviours at arbitrary short time scales, that are often followed by strong corrections and swings.

5 Considerations for regulators and policy makers

Potential consequence of increased market crashes

If the system is indeed “non-renormalizable” so that HFT leads to singular or non-convergent limiting behaviours, it is an almost certainty that HFT will lead to a higher frequency of crashes. Anecdotal evidence seems to confirm this statement as we pointed out earlier with respect to different mini-crashes that have been observed mostly but not only in equities related markets.

One may rightly ask why we should seek to prevent these crashes from occurring, especially as markets so far have demonstrated an equally fast recovery. Perhaps these volatility bursts are the price to pay for higher liquidity in the market (during normal market situations)? In this respect, there is an interesting paradox in the fact that HFT is justified by innovations that are thought to provide large liquidity and lower cost to investment and access to capital. But these innovations also create the risks of liquidity freezes. In a sense, “more on average” is associated with “much less” or zero at certain times. This phenomenon is well-illustrated by Louzoun et al. (2003), who use a simple auto-catalytic model of innovation and growth, with positive feedbacks and varying interaction range. Louzoun et al. (2003) show that, in such models, the total measure of welfare (wealth) is maximum when the dynamics is the most turbulent and risky, with huge spikes and collapses punctuating a very intermittent dynamics (as shown in the figure below). This is suggestive that more liquidity may similarly be associated with high turbulence, volatility and crashes.

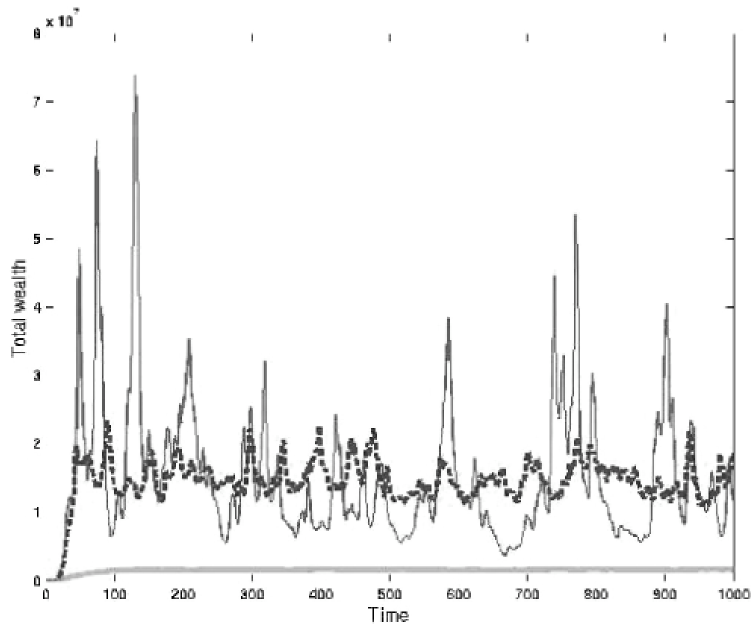

Figure: The time evolution of the total wealth for R = 0 (bottom smooth solid line), R=∞ (spiky solid line), and intermediate R (dashed line), where R is the interaction range between agents. Note that the intermediate competition case R = 40 (dashed line) ensures the maximal average. Note also that the totally globalized case (R = system size = 200) is better even at the worst moments than the completely localized market case (R = 0). However, the dramatic crashes involve of course massive human suffering (loss). This is avoided for R = 40. Reproduced from Louzoun et al. (2003).

This higher turbulence in the form of an increasing number of crashes could raise overall market risk due to the possibility of excitation to other markets and possible herding, specifically when the market is already fragile (e.g. May 6, 2010; the Greece crisis). It is conceivable that, as a consequence, we could see much larger "system failures", specifically if paired with general pessimism and a mistrust in the financial system. Those failures might result in losses in the real economy that take much longer to recover than the recovery after flash crashes so far.

Regulating complex systems

Two images arise in our mind when we think about the role of a governing agency with respect to financial markets: (1) a conductor in front of an orchestra, who coordinates very talented musicians – who can all play beautifully on their own, but need a coordinator to play a symphony in harmony or (2) a biologist in charge with the complex ecology of a forest.

The conductor does not need to know how to play each individual instrument but he needs to understand them and know when each instrument/musician should play its part. There are two problems with this image. Firstly, the conductor decides which piece to perform and the musicians are willing to follow him because their interests are aligned. This, in financial markets is not a given, as the individual players are generally only concerned with the sound of their own music (profits) and whether or not this goes in harmony with other instruments to create a symphony (growth in the real economy) has so far been of secondary concern (if at all). The second problem with this image of government agencies as conductors is that it is a very powerful position. This view undermines the general free market ideology underlying capitalism. Alternatively, the image of the agency as a biologist seems more appropriate. She plays an active role but lets nature and ecology complexity do its part. She understands the role of all the different species in the ecosystem and when necessary (for example, if excessive growth of certain animals or plants threatens the balance of the system), steps in to help the system maintain or return to a balance. The difference to a conductor is an understanding for a constantly emerging and evolving system and the consciousness to never fully be in control. This view of regulation emphasises the need to add to current academic research on financial markets and HFT with out-of-equilibrium dynamical system theory.

It is always difficult to forbid something per se via regulations, as it will emerge in a different form somewhere else because professionals are rational in their quest to work at the limit of legality. This can even be transformed as a theorem: optimization of utility in the presence of many constraints is bound to occur somewhere on the boundary of the authorized simplex of possible states, hence at the limit of legality.

That banks serve their own interests on the one hand and play a key role in lubricating the economy, thus serving as public good entities, on the other hand has been widely recognized in recent debates. Many discussions, with different emphasis across the Atlantic, focus on what kind of regulations should therefore be imposed to align the private interests of banks with the public interests. The recent Dodd-Frank act (2010) can be seen as a rather timid step towards a working solution, if not just because many of the changes implied by its implementation are not expected to be fully enacted until 2015 (five years is really like eternity for financial markets!).

Consider in contrast that the fifty years following WWII have constituted arguably the most stable economic period in the history of the United States and of Europe. Most scholars attribute a key role for this stability to the Glass-Steagall Act of 1932, which successfully prevented the occurrence of “super-spreader” instabilities, by separating by law investment banking, commercial banking, retail banking and insurance. This disaggregation provided completely separated waterproof compartments to prevent any Titanic like event of crisis spreading. Only with deregulation that started taking place in the 1980s culminating in the repelling of the Glass-Steagall act by the Gramm–Leach–Bliley Act of 1999, banking mutated into a new highly interconnected form that recovered basically its pre-1929 role within the ecosystem. Much of the risks that we currently face both in Europe and in the US originate from too much leverage and uncontrolled indebtedness spreading across all networks that build on the incorrect belief that transfers of debts to bigger and bigger entities will solve the problem.

We cannot afford and do not need to wait another decade or more until new super high tech models are developed. Faster solutions are possible by revisiting policies that worked in the past and by relearning and expanding some of the old wisdom in economics, specifically related to the role of banks. These theories should be anchored on rigorous analyses of empirical evidence and enhanced by fertilization with various branches of the natural sciences, network analysis, and out-of-equilibrium agent-based models (Sornette and von der Becke, 2011).

References

Aldridge, I. (2010) "High Frequency Trading: a practical guide to logarithmic trading strategies and trading systems", *John Wiley & Sons, Inc*.

Barenblatt, G. I. (1996) Scaling, self-similarity, and intermediate asymptotics, *Cambridge University Press.*

Black, F. (1976) "Studies of Stock Price Volatility Changes", Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economical Statistics Section, 177–181.

Bradley, H. and R.E. Litan (2010) "Choking the recovery (why new growth companies aren't going public and unrecognized risks of future market disruptions)", Ewing Marion Kauffman Foundation, Published Nov. 8, 2010; Revised Version with Corrections and Clarifications Nov. 12, 2010. (http://www.kauffman.org/uploadedFiles/etf_study_11-8-10.pdf)

Berry, M. (1991) "Asymptotics, Singularities and the reduction of Theories", *Logic, Methodology and Philosophy of Science IX, 597-607.*

Brogaard, J. A. (2010) "High Frequency Trading and its Impact on Market Quality", http://www.gsb.stanford.edu/facseminars/events/finance/documents/fin_02_11_brogaard.pdf

Calvet, L.E. and A.J. Fisher (2008) Multifractal Volatility: Theory, Forecasting, and Pricing (Academic Press Advanced Finance), Academic Press.

Cantono, S. and S. Solomon (2010) "When the collective acts on its components: economic crisis autocatalytic percolation", *New Journal of Physics 12.*

Chaboud, A., B. Chiquoine, E. Hjalmarsson, C. Vega, "Rise of the machines: algorithmic trading in the foreign exchange market", International Finance Discussion Paper (2009) http://ssrn.com/abstract=1501135

Challet, D., M. Marsili and Yi-Cheng Zhang (2005) "Minority Games: Interacting Agents in Financial Markets" (Oxford Finance), *Oxford University Press, USA.*

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

Chiarella, C. Dieci R., and He X.-Z. (2009) "Heterogeneity, Market Mechanisms, and Asset Price Dynamics" in *Handbook of Financial Markets: Dynamics and Evolution*. Chapter 5, 277-344.

Coolen, A.C.C. (2005) "The Mathematical Theory of Minority Games: Statistical Mechanics of Interacting Agents" (Oxford Finance), *Oxford University Press, USA*.

Dichev, I. D., Huang, K and Zhou, D. (2011) "The dark side of trading", Law & Economics Research Paper Series, Research Paper No. 11-95.

Easley, D. M. M. Lopez de Prado, M. O'Hara (2010) "Microstructure of the flash crash", http://ssrn.com/abstract=1695041

Easley, D., M. M. Lopez de Prado, M. O'Hara (2010) "Measuring Flow Toxicity in a High Frequency World", http://ssrn.com/abstract=1695596

Evstigneev, I.V., Hens, T.; Schenk-Hoppe R. K.,(2009) "Evolutionary Finance" in *Handbook of Financial Markets: Dynamics and Evolution, Chapter 9, 507-566.*

Filimonov, V.A. and D. Sornette (2011) "Self-Excited Multifractal Dynamics", *Europhysics Letters 94, 46003 (http://arxiv.org/abs/1008.1430).*

Filimonov, V.A. and D. Sornette (2011) A Stable and Robust Calibration Scheme of the Log-Periodic Power Law Model, submitted to Computational Statistics (http://arxiv.org/abs/1108.0099)

Froot, K.A., D.S. Scharfstein, and J.C. Stein, (1992) "Herd on the street: Informational inefficiencies in a market with short-term speculation". *Journal of finance*, 47(4):1461–1484.

Goldstein, M.A., J.E. Evans, and J.M. Mahoney, "Circuit breakers, volatility, and the U.S. equity markets", 385-405 (http://www.imes.boj.or.jp/cbrc/cbrc-16.pdf)

Grant, J. (2011) The Financial Times, April 12 2011 (http://www.ft.com/cms/s/0/e364be1a-6529-11e0-b150-00144feab49a,_i_email=y.html)

Greenspan, A. (August 30 2002) Economic volatility, symposium sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming. www.federalreserve.gov/boarddocs/speeches/2002/20020830/default.htm

Greenwood, J. and P. Krusell (2007) "Growth accounting with investment-specific technological progress: A discussion of two approaches", *Journal of Monetary Economics 54, 1300–1310.*

Gsell, Markus (2008), "Assessing the Impact of Algorithmic Trading on Markets: A Simulation Approach" ECIS 2008 Proceedings. Paper 225. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1134834

Gurkaynak, R.S. (2008) "Econometric tests of asset price bubbles: Taking stock", *Journal of Economic Surveys 22~(1), 166-186.*

Haldane, A. G. (2010) "Patience and Finance", Speech at Oxford China Business Forum, Beijing. http://www.bankofengland.co.uk/publications/speeches/2010/speech445.pdf

Haldane, A. G. (2011) "Race to zero, Speech at International Economic Association Sixteenth World Congress, Beijing, China. http://www.bankofengland.co.uk/publications/speeches/2011/speech509.pdf

Harras G., Sornette D., (2011) “How to grow a bubble: A model of myopic adapting agents,” in press in J. Economic Behavior and Organization (in press) (http://arXiv.org/abs/0806.2989)

Helmstetter, A. and D. Sornette (2003) “Predictability in the ETAS Model of Interacting Triggered Seismicity”, *J. Geophys. Res., 108, 2482, 10.1029/2003JB002485*.

Hendershott, T. and R. Riordan (2009) “Algorithmic Trading and Information”, Working Paper #09-08, NET Institute

Herndershott, T., C. M. Jones, A. J. Menkveld (2011) “Does Algorithmic Trading Improve Liquidity?” *The Journal of Finance, LXVI, NO.* 1.

Hommes, C., Wagner F. (2009) “Complex Evolutionary Systems in Behavioral Finance” in *Handbook of Financial Markets: Dynamics and Evolution*. Chapter 4, 217-276.

Jiang, Z.-Q., W.-X. Zhou, D. Sornette, R. Woodard, K. Bastiaensen and P. Cauwels (2010) “Bubble Diagnosis and Prediction of the 2005-2007 and 2008-2009 Chinese stock market bubbles”, *Journal of Economic Behavior and Organization 74, 149-162.*

Johansen, A. and D. Sornette (2010) “Shocks, Crashes and Bubbles in Financial Markets”, Brussels Economic review (Cahiers economiques de Bruxelles) 53 (2), 201-253 (summer 2010) (http://arXiv.org/abs/cond-mat/0210509)

Kaizoji, T. and Sornette, D. (2010) “Market bubbles and crashes”, *Encyclopedia of Quantitative Finance (Wiley, West Sussex)*, http://www.wiley.com//legacy/wileychi/eqf/ (long version of the preprint at http://arXiv.org/abs/0812.2449).

Kindleberger, C.P. (2000) “Manias, panics, and crashes: A history of financial crises”, 4th Edition. *Wiley, New York.*

King, M. R. and Rime, D. (2010) “The $4 trillion question: what explains FX growth since the 2007 survey?”, BIS Quarterly Review, December 2010

Louzoun, Y., S. Solomon, J. Goldenberg and D. Mazursky (2003) “World-Size Global Markets Lead to Economic Instability”, *Artificial Life 9, 357-370.*

Maillart, T. and D. Sornette (2010) “Heavy-tailed distribution of cyber-risks”, *Eur. Phys. J. B 75, 357-364.*

Mandelbrot, B.B. and R.L. Hudson (2006) The Misbehavior of Markets: A Fractal View of Financial Turbulence, Basic Books; annotated edition.

Mantegna, R.N. and H. E. Stanley (2007) Introduction to Econophysics: Correlations and Complexity in Finance, Cambridge University Press.

Matheson, T. (2010) “Taxing Financial Transactions: Issues and Evidence”, IMF Working Paper

Menkveld, A, J. (2011) “High Frequency Trading and The New-Market Makers”(http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1722924)

Moffatt, H.K. (2000) “Euler’s disk and its finite-time singularity”, *Nature 404, 833-834.*

Osorio, I., M.G. Frei, D. Sornette, J. Milton and Y.-C. Lai (2010) “Epileptic Seizures, quakes of the brain?”, *Physical Review E 82 (2), 021919.*

Ouillon, G. and D. Sornette (2004) “Search for Direct Stress Correlation Signatures of the Critical Earthquake Model”, *Geophys. J. Int., 157, 1233–1246.*

Rosser, J.B. (2008) “Econophysics and economic complexity”, *Advances in Complex Systems 11, 745-760.*

Saichev, A., Y. Malevergne and D. Sornette (2009) “Theory of Zipf's law and beyond”, Lecture Notes in Economics and Mathematical Systems 632, Springer (November), ISBN: 978-3-642-02945-5.

Satinover, J. B. (2002) “The Quantum Brain*”. New York: Wiley.*

Satinover, J. and D. Sornette (2010)” Taming Manias: On the Origins, Inevitability, Prediction and Regulation of Bubbles and Crashes”, chapter of the book “Governance and Control of Financial Systems: A Resilience Engineering Perspective,” Gunilla Sundström and Erik Hollnagel, eds., published by Ashgate Publishing Group in their Resilience Engineering Perspectives series. (http://ssrn.com/abstract=1646428)

Satinover, J. and D. Sornette (2011) Cycles, determinism and persistence in agent-based games and financial time-series, in press in Quantitative Finance (2011) (http://arXiv.org/abs/0805.0428)

Securities and Exchange Commission, 17 CFR Part 242, Concept Release on Equity Market Structure; Proposed Rule http://www.sec.gov/rules/concept/2010/34-61358fr.pdf

Smith, R. (2010) “Is high-frequency trading inducing changes in market microstructure and dynamics?” http://arxiv.org/abs/1006.5490v2

Sornette, D. (2003) “Why Stock Markets Crash (Critical Events in Complex Financial Systems)” *Princeton University Press.*

Sornette, D. (2009) “Dragon-Kings, Black Swans and the Prediction of Crises”, *International Journal of Terraspace Science and Engineering 2(1), 1-18 (2009)* (http://arXiv.org/abs/0907.4290) and (http://ssrn.com/abstract=1470006)

Sornette, D.,A. B. Davis, K. Ide, K. R. Vixie, V. Pisarenko, and J. R. Kamm (2007) Algorithm for Model Validation: Theory and Applications, Proc. Nat. Acad. Sci. USA 104 (16), 6562-6567.

Sornette, D., A. B. Davis, J. R. Kamm, and K. Ide (2008) A General Strategy for Physics-Based Model Validation Illustrated with Earthquake Phenomenology, Atmospheric Radiative Transfer, and Computational Fluid Dynamics, Book series: Lecture Notes in Computational Science and Engineering, vol 62, Book Series: Computational Methods in Transport: Verification and Validation

F. Graziani and D. Swesty (Eds.), Springer-Verlag, New York (NY), 2008, pp. 19-73. (http://arxiv.org/abs/0710.0317).

Sornette, D. and A. Johansen (2001) "Significance of log-periodic precursors to financial crashes", *Quantitative Finance 1 (4), 452-471.*

Sornette, D., T. Maillart and B. Sovacool (2011) "Reassessing nuclear power accidents after Fukushima", Nature (submitted 26 April 2011), available upon request to the first author.

Sornette, D. and I. Osorio (2010) Prediction, chapter in "Epilepsy: The Intersection of Neurosciences, Biology, Mathematics, Physics and Engineering," Editors: Osorio I., Zaveri H.P., Frei M.G., Arthurs S., CRC Press, Taylor & Francis Group (2010) (http://arxiv.org/abs/1007.2420).

Sornette, D. and S. von der Becke (2011) "Complexity clouds finance-risk models", *Nature 471, 166.*

Sornette, D. and R. Woodard (2010) Financial Bubbles, Real Estate bubbles, Derivative Bubbles, and the Financial and Economic Crisis, in Proceedings of APFA7 (Applications of Physics in Financial Analysis), M. Takayasu, T. Watanabe and H. Takayasu, eds, Springer (http://arxiv.org/abs/0905.0220)

Sornette, D., Woodard, R. and Zhou, W.-X. (2009) "The 2006-2008 oil bubble: Evidence of speculation and prediction", *Physica A 388, 1571-1576.*

Sornette, D. and W.-X. Zhou (2006) "Predictability of Large Future Changes in Major Financial Indices", *International Journal of Forecasting 22, 153-168.*

Thurner, S. (2010) Systemic financial risk: agent based models to understand the leverage cycle on national scales and its consequences, OECD project "Future Global Shocks".

Vives, X. (1995) "Short-term investment and the informational efficiency of the market", *Oxford Journals, Review of Financial Studies, Vol. 8, Issue 1, pp 125-160*

Zhang, X.F. (2010) "The Effect of High-Frequency Trading on Stock Volatility and Price Discovery", http://ssrn.com/abstract=1691679

Zhou, W.-X. and Sornette, D. (2003) "2000-2003 real estate bubble in the UK but not in the USA", *Physica A 329, 249-263.*

Zhou, W.-X. and Sornette, D. (2006) "Is there a real-estate bubble in the US?" *Physica A 361, 297-308.*

Foresight
1 Victoria Street
London SW1H 0ET
www.foresight.gov.uk
URN 11/1226