Government Office for Science

Foresight

Brave new world: quantifying the new instabilities and risks arising in subsecond algorithmic trading

Driver Review DR27

Foresight, Government Office for Science

Contents

Brave new world: quantifying the new instabilities and risks arising in subsecond algorithmic trading

Neil Johnson and Guannan Zhao[1]

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.

---

1 NFJ gratefully acknowledges the support of the UK Foresight program in preparing this document. He is also grateful to Pak Ming Hui, Brian Tivnan, Eric Hunsader of Nanex and Paul Jefferies for a long-term research collaboration on financial complexity, which provides the basis for the discussions in this paper.

1. Introduction

Looking to the next decade, there can be little doubt that computers will become increasingly involved in the functioning of financial markets at all levels – from new smart-phone apps aimed at individual investors, through to the mechanics of accounting and information transfer in exchanges and financial institutions. The most dramatic impact is likely to arise at the ultrafast, subsecond timescale: Operating beyond human response times, even relatively modest computer trading platforms can already digest information that they have been fed, take buy or sell decisions based on the internal algorithms that they have been given, and then execute these trades, all before a human has had a chance to draw breath. It is this new world that we explore in the present report.

This report shows that the behaviour of market prices within this subsecond world in which computers can trade freely in real-time – but humans cannot -- is fundamentally new, and that its understanding will require new sets of tools, new theoretical results, and new 'rules of thumb' for both practitioners and regulators. One might counter-argue that since financial markets have always tended to use the latest technologies, they have always been 'fast' compared to many other aspects of human life, and hence the subsecond world will just be a faster version of the everyday market phenomena which we already know. However this report shows that this statement is false. One might suggest that standard mathematical results concerning risk calculations can simply be re-applied at this shorter timescale. Again this is false. One might declare that financial instabilities on ultrafast timescales are not new -- after all, the Flash Crash of 6 May 2010 indeed happened very quickly, being over within a few minutes. However, such fast crashes are still in principle slow enough for humans to be directly involved with the trading in real time. Given a big enough 'stop' button, a human operator could in principle step in and stop such rapid buying or selling - even if it took place on the scale of a few seconds. By stark contrast, the subsecond regime on which we focus here, lies beyond the limits of human response times. Looking at detailed and reliable millisecond data recorded by our collaborators at the US company Nanex (www.nanex.net), we instead show that new breeds of extreme behaviour can – and have already started to -- emerge on the subsecond scale, at timescales where no human can physically react, let alone think strategically. We examine the properties of this new breed of behaviours, and use it to develop a broader theoretical picture of what is likely to emerge over the next decade on the subsecond timescale, and how such behaviours might be described theoretically using relatively simple mathematical analyses. In addition to providing a quantitative interpretation of the subsecond price behaviour to date, our proposed theoretical framework suggests that this new subsecond machine regime can be usefully seen as an ecology of competitive trading machines, fighting it out on the millisecond scale, and hence is entirely consistent with the ecological perspective of Farmer and Skouras (2011). Our findings are also remarkably consistent with the detailed and careful studies of Cliff and co-workers (De Luca et al. (2011). Our model framework is also fairly consistent with the idea of financial instability proposed by May and Haldane (2011), though not in substance and on a completely different timescale.

The main implication of this report's findings, in addition to the specific technical deliverable of a framework for understanding and even estimating future behaviours as a function of the algorithmic diversity etc., is that the behaviour at these subsecond timescales is not simply a faster version of what is happening on timescales above one second. This in itself is surprising since the approximate self-similar nature of financial market price movements on larger timescales is now well-established: To a reasonable approximation, the patterns observed over months are similar to those over weeks, which are similar to those over days etc. We find that this is not the case as one moves through the subsecond time barrier beyond which only machines can operate. The self-similarity stops abruptly, with a fundamental system-wide transition arising near the limits of human response times (approximately 600-800 milliseconds, Liukkonen (2009) and Saariluoma (1995)). Indeed, instead of simply postulating the 'rise of the robots', we are able to actually observe the signatures of this fundamental transition in the data.

The accompanying model that we present is simple in form – indeed it is a deliberately oversimplified representation of what is effectively the world's largest technosocial system. Yet its ability to reproduce a fundamental transition akin to the one we uncover in subsecond price behaviour, suggests that it is capturing some essential ingredients of the complexity arising in populations of machine-like trading objects, as they operate on very short timescales without direct human intervention. We show that the emergent properties of this model system (e.g. price volatility) are relatively simple to formulate mathematically, and yet offer concrete quantitative predictions of how instabilities will likely develop as a function of the physical variables in the system – from the intrinsic memory of the machines through to their diversity in terms of trading algorithms. It is this quantitative picture of an ecology of machines which we believe will prove useful to both regulators and participants over the next decade, and which can provide a solid platform for a new generation of financial derivative and risk models. We outline how this can be done in the penultimate section of this report.

2. Technological advances drive market behaviour into a new regime

The potential benefits to a financial entity of having an advantage over a competitor are so large, and worth so much money, that competition within the financial markets alone should drive technological developments quickly toward the microsecond and even nanosecond operating timescale over the next decade (Haldane (2011), Perez (2011)). Such competition-driven speed-up did arise in the past – however, what is remarkable now is that these technologies are set to push hard up against the physical limitations of the laws of nature in terms of the ultimate speed limit, which is the speed of light, and the laws of quantum physics in terms of the physical switching of logic gates or transistors. For example, a new dedicated transatlantic cable is being built just to shave 5 milliseconds off transatlantic communication times between US and UK traders (Popular Mechanics (2012)) while a new purpose-built chip iX-eCute is being launched which prepares trades in 740 nanoseconds (Wall Street Journal (2011)).

But perhaps most interesting as a sign of things to come, are the new hybrid ventures beginning to spring to life, involving collaborations between traders, engineers and basic scientists, with the aim of pushing both the financial and physical boundaries for trading. This includes the re-emergence of traders and fund managers who were themselves previously doing cutting-edge research in electromagnetic theory for signal propagation. One U.S. example is a new venture (see http://www.mckay-brothers.com/about-us/) co-founded by Dr. Bob Meade, a Harvard PhD in theoretical physics who previously did internationally leading research at MIT into the speed of electromagnetic radiation in particular types of smart media, producing several patents before switching to a career in finance. After heading the derivative research group at JPMorgan and then quant-trading Fleet Bank's Robertson Stephens investment bank, he ran a High Frequency Trading group at Ronin Capital before co-founding McKay Brothers. The McKay Brothers initiative aims to link the stock market trading in Chicago and New York using electromagnetic technology in the microwave spectrum, operating at speeds faster than fiber optic transmission can deliver. Their main competitor, Spread Networks, operates fiber optic links and is reported to have spent 300 million dollars developing a low latency connection, which will soon be outstripped by McKay Brothers' microwave routing. This type of initiative, which manages to simultaneously be creative financially and scientifically, is likely to become the norm over the next decade rather than the exception, with other such hybrid examples arising across the globe. In the end, the pressure to succeed will drive speeds down toward their physical temporal limits in the same way that Moore's Law drove processor sizes down to their physical spatial limits. Since it ultimately only requires the presence of machines, not people, to profit from these reductions in latency, this trend toward increasingly fast and increasingly automated systems will likely continue unbounded (Haldane (2011)).

These technological developments raise important questions about the future added value of existing financial centers such as London. From a purely technological point of view, it would make perfect sense to site hubs of microwave information flow (as required in the McKay Brothers proposal) at isolated sites with little electromagnetic interference, and hence *away* from major cities -- particularly if the information is being transferred through microwaves in the open air. Dense co-location hubs built around optic fiber communications, such as those recently built in Essex so that they are close to London's East End financial area, would also become redundant because of optic fiber slowness compared to a raw microwave link. Regardless of the eventual winning technology for establishing fast communication links between computers, it is clear from the above discussion that as communications become faster, so too will the competition between companies to develop faster trading machines. As this competition to build faster machines hots up, the relevant timescale for significant volumes of trade will move beyond the millisecond scale toward the microsecond and even nanosecond scale.

Looking to the future, however, one might also argue that while the number and speed of subsecond scale trading machines is set to increase rapidly, the diversity and complexity of their trading algorithms may not necessarily match this rapid increase. On ultrafast subsecond timescales, the information concerning recent price movements needs to be assimilated quickly by the machine, then the trading algorithm run, then the trading decision implemented. Hence there may end up being a practical limit to the amount of complexity that a trading algorithm can have. Given the short timescale requirement for a trade to be made, one could argue that the number of lines of code, as well as the number of data look-ups, matrix manipulations, iterative loops etc., will be restricted in order that the code can run quickly, efficiently and yet remain manageable and explainable – in other words, a complex code that takes too long to run, debug, and explain to the company's risk manager, may simply be impractical. This would act as a natural restriction in diversity which – when combined with the fact that similar 'hot' ideas can proliferate throughout trading circles at any one time as a result of a common pool of employment, similar background training and trade magazines, as well as attendance at the same conferences -- suggests that the same type of trading algorithm (or even the exact same algorithm) could inadvertently become part of the trading repertoire of a significant number of otherwise unconnected trading institutions. Given the natural secrecy of financial institutions, such inadvertent posession of similar strategies would go unnoticed and uncorrected. Should conditions then become favourable for use of a given algorithm, a significant fraction of all market participants would switch to it around the same time. The resulting 'crowd' effect, by which the trading algorithms of many trading institutions suddenly become similar or even identical, will (as we show later in this report) produce large movements in the market on the subsecond timescale. Moreover this will likely continue to happen in a way that is not seen at longer timescales where humans become actively involved in real-time trading, and hence where the natural diversity of human decision-making and 'free will' will tend to greatly expand the space of possible strategies, thereby diluting any such crowding around a given deterministic (algorithmic) trading strategy.

This report provides evidence to suggest that such algorithmic crowding is already happening on the subsecond scale, and our accompanying model provides a theoretical description of the likely knock-on effects on price volatility. Specifically, we develop a quantitative expression for how market fluctuations are likely to vary according to future computer algorithm diversity. We also present quantitative results for the market volatility that is likely to emerge under such crowded conditions, i.e. where multiple algorithms with essentially the same composition are all in use at the same time. We look at how the resulting market behaviour takes the system well away from the typical regime of a near-random walk (see Fig. 1) characterized by a near-Gaussian, albeit fat-tailed, distribution (see Figs. 2 and 3) and hence nearly perfectly hedgable derivative contracts following a Black-Scholes type risk analysis (Bouchaud and Potters (2003), Johnson et al. (2003)). Instead it launches the market toward a new regime in which a new class of risk calculation must be developed. We suggest an alternative measure of the risk for this new subsecond, machine-driven regime in the presence of finite latency. Our empirical evidence and support comes in the form of subsecond extreme events in the stock time-series across stock and exchanges between 2006 and 2011, with the total number of such events undergoing a huge increase through the period of global market instability in the latter part of 2008.

3. Typical financial market dynamics

In order to differentiate these new market price dynamics which arise in the subsecond regime, we need to briefly review the 'usual' dynamics of financial markets on longer timescales. There are of course many hundreds, or even thousands, of econometric reports concerning the properties of stock price movements (see, for example, Campbell et al. (1996) and references within, as well as Bouchaud and Potters (2003) and Johnson et al. (2003)). Indeed, the goal of characterizing the movements of financial markets on the scale of years, months, weeks, days - and most recently, hours, minutes and seconds -- has been the long-term focus of academics and practitioners for many decades. It is a basic truth of science that in order to observe, and hence ultimately understand the nature of, objects of a certain size, a microscope is required with a resolution which is at least one order of magnitude greater -- in order to be able to distinguish detail from the blur. This principle applies not only to objects which are small in terms of spatial size, but also for events that only last a small amount of time. A sports photographer looking to capture a picture of a football flying into a goal, requires a shutter speed many times faster than the time-of-flight of the football. The finite resolution on which we examine a system, necessarily restricts the range of phenomena that we might see, and automatically rules out the part of the spectrum of behaviours for which the analysis method is too slow. The same applies to financial market behaviour: It is only as data become recorded in an accurate way on increasingly small timecales – from weeks to days, to hours, minutes and seconds – that increasing insight can be gained. Indeed, it is the ability of the company *Nanex* to capture and reliably store time-stamped stock price data on the millisecond scale, that has made the present study possible (see www.nanex.net).

Although the passage of time from years to months, to weeks and days, differentiates between many human activities (e.g. annual vacation rest as opposed to nightly sleep), it turns out that financial market trading patterns – at least to a reasonable approximation – show a remarkable degree of self-similarity in terms of their price behaviour (e.g. Bouchaud and Potters (2003), Mantegna et al. (1995), Johnson et al. (2003), Gabaix et al. (2003)). More recent analysis has shown that the price series is actually even more subtle than this simple self-similar view, and instead exhibits multifractal characteristics -- however, in order to explain the concept, we stick with the simpler fractal version here. If one looks at the price chart of a typical liquid stock, stripped of its time units on the horizontal axis and price units on the vertical axis, it is difficult to tell by eye whether the chart referred to price-changes by month, by week, or by day. Recent work by Preis and Stanley has shown that this approximate self-similarity in price charts can also exist on shorter timescales, down to the typical second scale (Preis et al. (2011)), although we note that more recent work by Filimonov and Sornette casts doubt on these authors' analysis (see http://arxiv.org/abs/1112.3868 for details of this debate). The broad feature whereby scale does not seem to matter, is called a statistical fractal and means that the pattern of price movements is approximately scale independent in the same way that the coastline of Britain is scale independent, as well as a whole set of other phenomena from the natural world (Bouchaud and Potters (2003)). In reality, no perfect fractal exists due to finite size effects, e.g. eventually the coastline of Britain is bounded by the size of the island itself -- but the point is that this self-similarity holds approximately over a wide range of scales, from months to weeks to days etc. The resulting distribution of price-changes then tends to be fat-tailed. This means it is not Gaussian as one would expect for a random walk as in Figs. 1 and 2, however nor is it a perfect power-law, as in the specific case of the Lorentzian in Fig. 2. Instead, it tends to lie between the two, as shown explicitly for the Shanghai stock index in Fig. 3. Other stock markets produce data with distributions that are remarkably similar (Bouchaud and Potters (2003)).

Figure 1. Random walk model of financial price movements, with price on the vertical axis and time on the horizontal axis. Although this coin-toss incarnation or price movements is far simpler than many versions used in practice, it illustrates the basic principle of markets being described by ongoing stochastic changes in the price as opposed to some more microscopic, yet physically realistic model, comprising a population of trading agents. The price changes here all occur for a given pre-determined time-interval (i.e. one timestep may be 1 hour, or 1 day etc.).

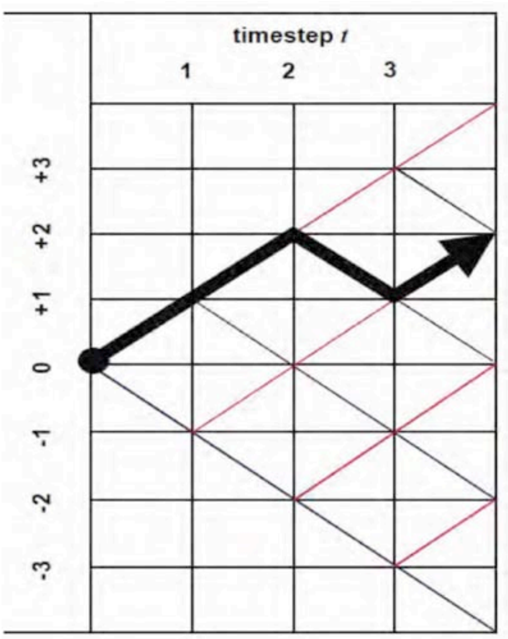

Figure 2. Gaussian vs . Lorentzian distribution for price-changes in a financial market, i.e. returns. The Lorentzian, as shown, behaves like a power-law in its tail (i.e. as *x* becomes large) since it varies as the inverse square of *x* and hence has an exponent of value -2. It is therefore referred to as a 'fat-tailed' distribution, whereas the Gaussian is not since it decays exponentially and hence far quicker as *x* increases.

$$p(x) = \frac{C}{x^2 + C^2\pi^2} \quad \text{(Lorentzian)} \qquad\qquad p(x) = \frac{1}{\sqrt{2\pi\sigma_x^{\,2}}}\, e^{-\frac{x^2}{2\sigma_x^{\,2}}} \quad \text{(Gaussian)}$$

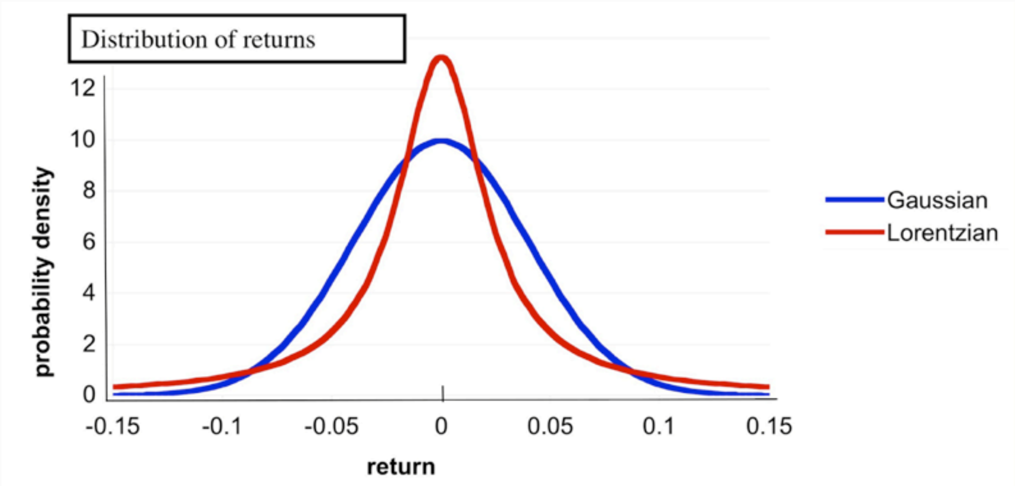

In addition to traditional studies of market prices, sophisticated quantitative descriptions have emerged recently from the new field of 'econophysics' (see www.unifr.ch/econophysics) although such efforts have met with some resistance from encumbant finance researchers. Economists studying financial market fluctuations might rightly point to the detailed GARCH-type models and generalizations that have already been developed outside the econophysics field (Campbell et al. (1996)). They might also complain of the new invaders' frequent lack of detailed referencing to past economics papers. Financial mathematicians might themselves claim huge strides in the development of complex stochastic models aimed at describing the pricing of risk associated with the burgeoning derivatives markets. Exotic options have for some time been priced using elegant mathematics, which may involve complicated jump processes and also memory in the time-series itself.

In their defense, econophysicists -- who are literally physicists trained in the theoretical tools of statistical mechanics in physical systems -- might in turn claim that they do indeed reference existing finance papers when they are relevant, but that they are actually trying to focus on aspects of market complexity which are not addressed by economists or by financial mathematicians (Bouchaud and Potters (2003)). They have a point: the domain of the physicist is one in which real-world data takes center-stage as representing the best 'measurement' of the system, and hence the best indicator of what is actually going on inside the system -- and hence the goal of any market study should be to analyze the properties of the data, identify any generic common patterns, and then build a model which is consistent with these observed features. This approach, which has after all worked well for physics since the time of Sir Isaac Newton, involves a continual iteration between more refined measurements, model development, output prediction and model adjustment. In the econophysics domain, the resulting ‘model’ tends to comprise a population of physical pieces which would include both humans and machines in principle – as opposed to some statistical ‘model’ based on a uni- or multivariate description in which the parameters have no clear meaning in terms of the microscopic workings of the system. Going further, the econophysicists would argue that a sensible approach to understanding markets is to observe the data, deduce the characteristics of the data that are explicable and those that are surprising and/or inexplicable, and then develop a model with minimal parameters and details but maximal insight, through a process of model modification and iterative comparison with the data. A paper plane hence becomes a good model to explain flight, while a child’s plastic model (which has seats, and dolls as passengers) cannot fly and hence is not. Since the future behaviour of financial markets in the presence of computer trading, is essentially unknown and yet data is now available, we tend to adopt the econophysicist philosophy in this review, i.e. we seek to analyze the subsecond data in order to deduce some stylized facts. Then we develop a minimal model which has a plausible micro-level interpretation, in order to reproduce these stylized facts, and hence infer a reasonable scenario for what might actually be going on in the market on this timescale.

Figure 3. Distribution of price returns z for Shanghai market data, for timescale $\Delta t = 1$ seconds (i.e. second-by-second price-changes). Also shown is a power-law (so-called Levy) distribution for comparison purposes. The agreement is very good over the main central portion, with deviations for large z. We show two attempts to fit a Gaussian: The wider Gaussian is chosen to have the same standard deviation as the empirical data, however the peak in the data is much narrower and higher than this Gaussian and the tails are fatter. The narrower Gaussian is chosen to fit the central portion. However the standard deviation is now too small. The data has tails that are much fatter and furthermore have a non-Gaussian functional dependence.

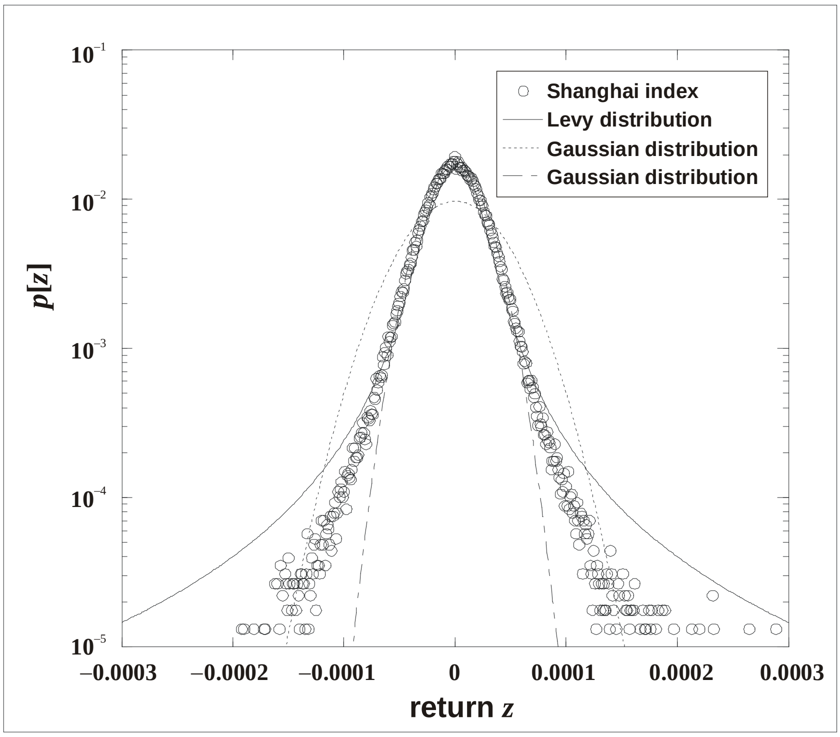

Most studies of stock market movements look at a time-series of price changes, and then chart the distribution or correlations for a fixed time increment (e.g. price-change from day to day). By contrast the study of extreme behaviours and events – which we turn to in the next section -- has been more traditionally the domain of the insurance and risk field. More recently, there has been a move to correct this: e.g. works by Sornette (2009) and co-workers as well as multiple studies in the area of Extreme Value Theory (Bouchaud and Potters (2003)). However, very little of this work looks for a mechanistic explanation of what is going on prior to, and during, such extreme events. This might be somewhat acceptable if one takes the stance that such events are so rare that they are not worth worrying about – however this is not the case since their effects can be long-lasting. Indeed, the extreme behaviours that we show in the rest of this report as arising in the subsecond regime beyond human response, are rare on the scale of milliseconds -- however, there are so many millisecond intervals during a day that we end up finding approximately ten such extreme behaviours per day in the data.

4. Large subsecond changes with variable duration

The problem with many studies of financial price changes – such as those outlined above and shown in Fig. 3 -- is that they adopt a fixed, pre-determined time increment over which to determine price changes (e.g. 1 day). A fixed time increment is unable to capture the wide variety of shapes and durations of extreme behaviours exhibited by financial markets. This is particularly true of the subsecond extreme events discussed in the rest of the article, and shown explicitly in Figs. 4(a) and 4(b). Indeed, extreme behaviours are often referred to as extreme 'events' on the assumption that they have a well-defined change Δx (e.g. price-change) in some macroscopically measurable quantity x (e.g. stock price) occurring at a particular point in space (e.g. Dow Jones) and time t (e.g. at 10am), over a specific time-interval Δt (e.g. 1 hour). If this idealization is indeed the case, then histograms can be obtained using historical data and approximate point probabilities deduced. However, as emphasized by Sornette (2009) extreme behaviors in principle invoke an entirely different layer of difficulty, because (1) they do not have a well-defined duration Δt, and hence may be missed when evaluating histograms of changes for a particular fixed, pre-defined time increment Δt (e.g. 1 minute, 1hour or 1day); and (2) even if their duration Δt and maximum size are well defined, they can take on an effectively infinite number of possible temporal profiles during that period, i.e. x has its own characteristic time-dependence during Δt. Hence for a given maximum drop size and duration Δt, there are a priori myriad possible temporal forms of x versus t. Such extreme behavior represents a fascinating departure from 'typical' behavior, and helps highlight the failings of mean-field theories upon which most of our quantitative descriptions of financial and natural systems are currently built.

A quick consideration of what makes a system complex, provides insight into the properties of a system which enable it to exhibit extreme behavior of the type observed on the subsecond scale in the markets in Fig. 4. Consider the outcome from tossing *N* coins. Assigning 1 as heads and -1 as tails, the famous Central Limit Theorem (CLT) guarantees that the net outcome value approaches a normal (i.e. Gaussian) distribution as $N \to \infty$. Such a normal distribution has an infinitesimally small probability of showing any extreme behavior (i.e. 99.73% of outcomes lie within three standard deviations from the mean), hence there is negligible likelihood of approximately *N* heads appearing as $N \to \infty$.

By contrast, real-world extremes, such as market crashes larger than three standard deviations, are far more common (Bouchaud and Potters (2003)). This relative abundance of extreme behavior in the real world (e.g. stock crashes) as opposed to a coin-toss world, suggests that real-world systems represent the effective opposite of a collection of independent stochastic processes. Indeed, current thinking within the scientific community suggests that for extreme behaviors to arise with appreciable frequency, the system needs to exhibit collective behavior – for example, crowd effects in a system comprising a population of *N* interacting objects which may adapt to past outcomes using the feedback of information, while continually competing to win. It is this idea that we will develop into a fuller model of the crowding of computer algorithms later in this review.

Figure 4. Traded price during black swan events. (a) Spike. Stock symbol is SMCI. Date is 10/01/2010. Number of sequential up ticks is 31. Price change is +2.75. Duration is 25ms (i.e. 0.025 seconds). Percentage price change upwards is 26% (i.e. spike magnitude is 26%). Dots in price chart are sized accordingly to size of trade. (b) Crash. Stock symbol is ABK. Date is 11/04/2009. Number of sequential down ticks is 20. Price change is -0.22. Duration is 25ms (i.e. 0.025 seconds). Percentage price change downwards is 14% (i.e.crash magnitude is 14%). (c) Cumu;ative number of crashes (red) and spikes (blue) compared to overall stock market index (Standard & Poor's 500)in black, showing daily close data from 3 Jan 2006 until 3 Feb 2011.

!

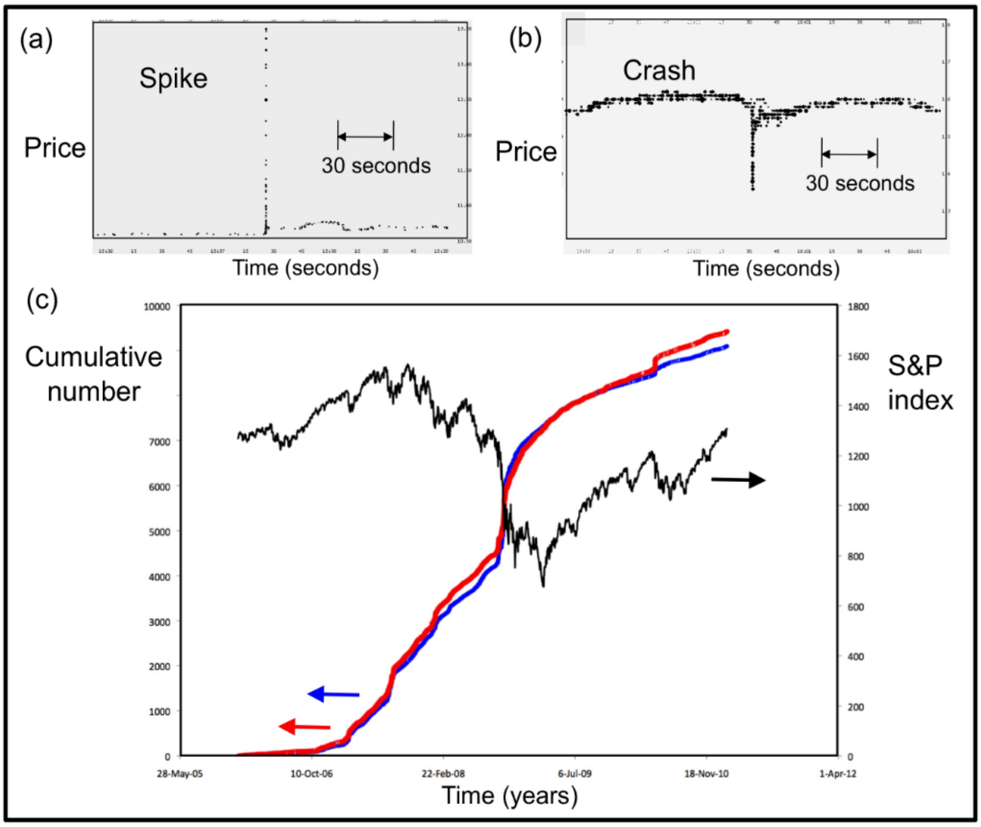

!

5. The new world of subsecond black swans

We now discuss the emergence of a fundamentally different regime of financial market behaviour on the subsecond scale. Our findings are surprising since one might have imagined that a move to subsecond timescales would simply reveal faster versions of the same phenomena that one observed at larger timescales. In particular, given that markets are known to have an approximate self-similar structure, in that the movements on the scale of months look like an expanded version of the movements on the scale of weeks, and the movements on the scale of weeks look like an expanded version of the movements on the scale of days, and so on, one might think we would simply get more of what we already now – it might just come and go more quickly, and possibly be accompanied by an increase or decrease in intensity depending on details of the system. However the subsecond timescale is different for a fundamental reason: At every timescale above a second or greater, a human trader – if sufficiently attentive – can in principle intervene in an automated trading system, no matter how complicated, by hitting a 'stop' button on the trading machine or even cutting the power. However, this does not hold for the subsecond timscale: Instead it takes a chess grandmaster approximately 650 milliseconds just to realize that she is in trouble (i.e. her king is in checkmate), without any physical action (Saariluoma (1995), Liukkonen (2009)). In many other areas of human activity, the quickest that someone can notice such a cue and physically react, is approximately 1000 milliseconds (1 second). The relevance of this subsecond timescale in financial markets would be relatively minor if it were not for the fact that this regime is already populated by computers which *can* operate this fast, even though the human participants cannot.

For reasons given in the previous section, it does not make sense to analyze extreme behaviours in this new ultrafast machine regime in terms of price-changes for a given fixed time increment. Instead we will analyze the size and duration of the extreme events themselves (see Fig. 4). Specifically, our collaborators at *Nanex* undertook a search for ultrafast extreme events in a high-throughput millisecond-resolution stream of prices for multiple stocks across multiple exchanges between 2006-2011. This data includes all financial and non-financial company stock, and looks across all major exchanges such as Nasdaq Exchange (NQEX), New York Stock Exchange (NYSE), American Stock Exchange (AMEX), Boston Stock/Options Exchange (BOST), Chicago Stock Exchange (CHIC), through to the Winnipeg Commodity Exchange (WCE) and the London Stock Exchange (FTSE). We refer to the documentation at www.nanex.net for details of specific extreme events within individual markets and stock. As shown by Fig. 5, there is a rapid explosion in the total number of subsecond events within a given duration range, as we move to smaller durations.

Figure 5. Number of black swan events with duration within a given 100ms time-window, as a function of the time-window label. For example, the first (*i*=1) entry on the horizontal axis shows approximately 3000 black swans with durations which lie between 100(*i*-1)=0ms and 100*i*=100ms. Likewise, the second (*i*=2) entry on the horizontal axis shows approximately 2000 black swans with durations which lie between 100(*i*-1)=100ms and 100*i*=200ms, etc. The number of black swans within each bin decreases rapidly as the bin index *i* (and hence the black swan duration) increases. The number of black swans with small durations (e.g. 200ms) is therefore much larger than the number with larger durations (e.g. 1200ms). The duration of the black swan crash or spike, is simply the length of time (i.e. clock time) during which the price ticked down or up respectively.

Since the clock time between ticks varies, the duration of a crash or spike with the same number of ticks can vary considerably, as shown.

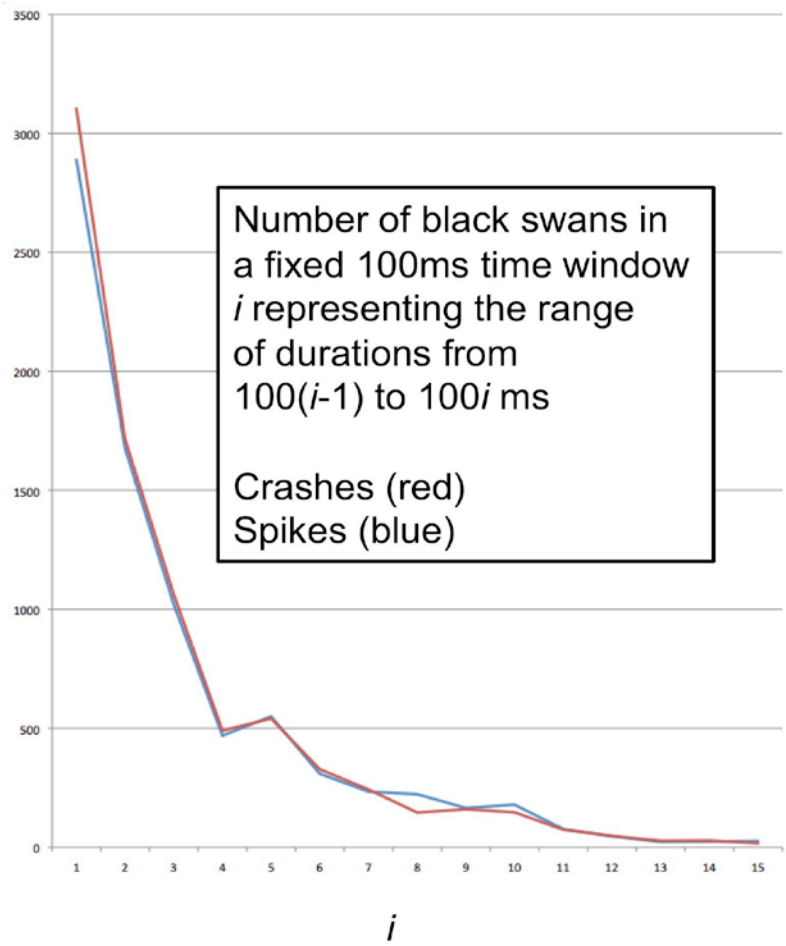

For convenience, we use the popular term ‘black swan’ (Taleb (2010)) for each extreme event. We might also usefully refer to them as ‘fractures’ given their visual similarity to microscopic fractures in a material. For a large price drop to qualify as an extreme event (i.e. black swan crash) the stock price had to tick down at least ten times before ticking up and the price change had to exceed 0.8%. For a large price rise to qualify as an extreme event (i.e. black swan spike) the stock had to tick up at least ten times before ticking down and the price change had to exceed 0.8%. The duration of the black swan crash or spike, is simply the length of time (i.e. clock time) during which the price ticked down or up respectively. Since the clock time between ticks can vary considerably according to issues such as liquidity, the duration of a crash or spike can also vary considerably, even if the number of down or up ticks is fixed.

In order to explore timescales which go beyond typical human reaction times, we focus on black swans with durations less than 1500 milliseconds. We uncovered 18,520 such black swan events, which surprisingly is approximately ten per trading day on average. Figure 4 illustrates a spike (Fig. 4(a)) and crash (Fig. 4(b)) from our dataset, both with duration 25 milliseconds (0.025s), while Fig. 4(c) suggests a systemic coupling between these sub-second black swan events in individual stock (blue and red curves) and long-term market-wide instability on the scale of weeks, months and even years (black curve) – in particular, in relation to the global financial crisis starting in 2008. Each black swan feature in Figs. 4(a) and 4(b) is huge compared to the size of the fluctuations either immediately before or after it, while the quick recovery from the initial drop or rise probably results from an automatically triggered exchange response or predatory computer trades. The coupling in Fig. 4(c) across such vastly different timescales is made even more intriguing by the fact that the ten stock with highest incidences of ultrafast black swans are all financial institutions -- and yet it is financial institutions that have been most strongly connected with the late 2000’s global financial collapse (e.g. Lehmann Brothers filing for Chapter 11 bankruptcy protection on 15 September 2008). This suggests an analogy to engineering systems where it is well known that a prevalence of micro-fractures can accompany, and even precede, large changes in a mechanical structure (e.g. tiny cracks in a piece of plane fuselage which then eventually breaks off). As shown in Fig. 5, our dataset shows a far greater tendency for these financial fractures to occur, within a given duration time-window, as we move to smaller timescales, e.g. 100-200ms has approximately ten times more than 900-1000ms. The fact that the instantaneous rate of occurrence of spikes and crashes is similar (i.e. blue and red curves are almost identical in Fig. 4(c)) suggests that these ultrafast black swans are not simply the product of some pathological regulatory rule for crashes. An immediate implication of these observation in Fig. 4(c) for regulators is that extreme behaviors on very short (i.e. <1s) and long timescales (e.g. 1 year, or $\sim 10^{-7}$s) cannot a priori be separated: In particular, a large change in the behaviour on the monthly scale as in Fig. 4 can be accompanied by an explosion in the number of subsecond instabilities (black swans) on the subsecond scale (see Fig. 4(c)). This coupling between long and short timescales means that rules targeted solely at controlling ultrafast (e.g. subsecond) fluctuations can induce unexpected feedback effects at the scale of months or years – likewise, rules targeted solely at calming markets on the scale of months or years, can induce unexpected feedback on the intraday or even subsecond scale.

As shown in Fig. 6, there is a general tendency for the number of these black swans to increase as the overall market volatility increases, however one is not simply a mirror of the other. Likewise, Fig. 7 confirms that the size and frequency of the black swans are not trivially inter-related.

Figure 6: Number of separate spikes (blue) and crashes (red) in individual months, as a function of the volatility in that month. As can be seen, there is some support for a coupling between the number of ultrafast black swans (i.e. spikes and crashes) and the overall volatility of the market on much longer timescales (i.e. months).

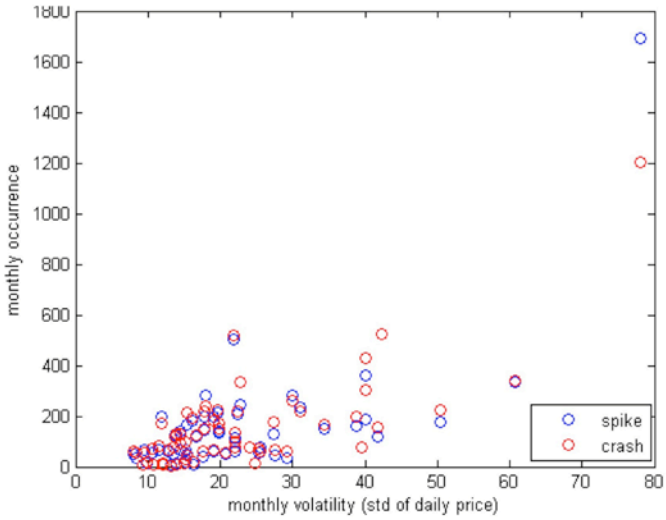

Figure 7. Fractional size vs. the duration for all subsecond black swans including both crashes and spikes. For the sake of clarity in the plot, we truncate events with size larger than 100%, i.e. we have truncated the plot at 100 on the vertical scale. As can be seen, there is no well-defined relationship between spike/crash size and duration. The fact that size and duration are not trivially linked helps confirm the surprising nature of our findings.

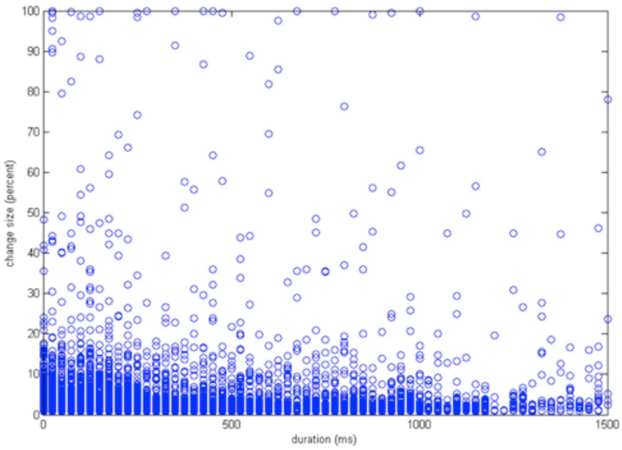

Figure 8. Distribution of subsecond black swan sizes for fixed, consecutive, non-overlapping time-windows for the duration (e.g. durations between 400-500ms, 500-600ms etc.). The colours (with values indicated) show the results of the Kolmogorov-Smirnov two-sample test to check the similarity of the different distributions within 15 different time-windows. The fact that there is little similarity between distributions on the longer timescale (> 1 second) and ones a few hundred milliseconds below, is consistent with the claim that black swan events of duration below about 800-900ms are fundamentally different from those above, i.e. there is a phase transition. We also carried out the power law test on these same consecutive, non-overlapping time-windows for the duration (e.g. power-law test on durations between 400-500ms, 500-600ms etc.). Only the time-windows containing long duration black swans pass the *p*-value test, which is consistent with our claim of a fundamental phase transition just below the human response time.

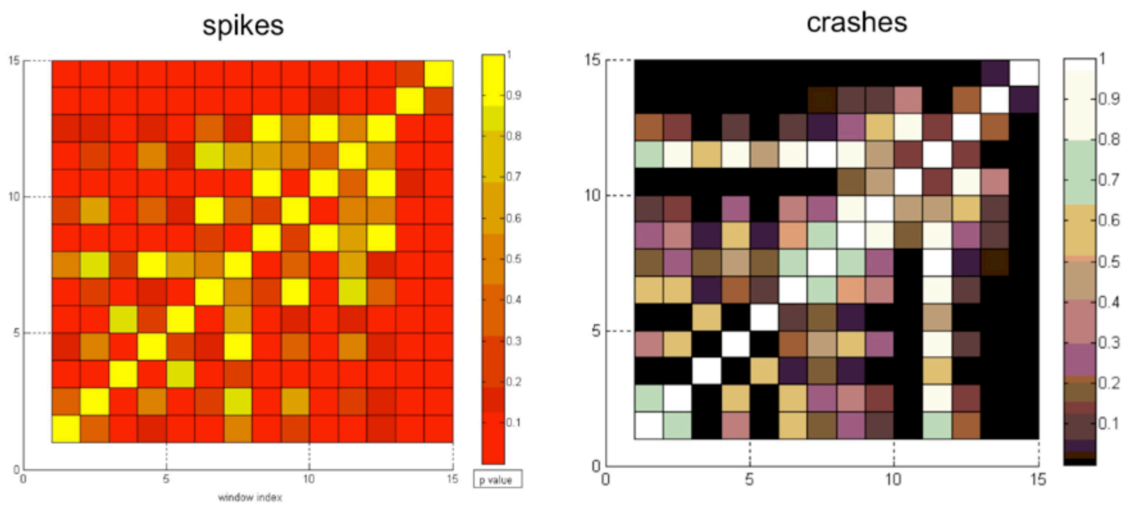

Figure 8 analyzes how the distribution of the sizes of subsecond black swan events changes as a function of the timescale. The low values for the similarity between distributions above and below 1 second in Fig. 8, combined with the rapid increase observed in Fig. 5 for the number of black swans below 1 second, suggests that a phase transition might arise at subsecond timescales. Figures 9-15 provide even stronger evidence to support the existence of such a transition, by showing a number of properties that undergo a visibly abrupt change as we move through the timescales at which humans become too slow to intervene and act.

Figure 9. Left panels show the average size of black swans within a given window of duration as a function of the upper value of the duration for this window. The windows of duration are fixed, consecutive, non-overlapping with each one having a size of approximately 150 ms (there are 10 windows in total). Right panels show the standard deviation of the size of the black swans within a window. Again, the subsecond regime appears fundamentally different from the regime of >1 second. The mean size is seen to increase dramatically as the duration moves above 1 second, and yet the mean number is also decreasing dramatically as shown earlier in Fig. 5. Although there are fewer black swans at larger durations, the mean size of them is larger simply because they have more time to develop and grow. The fact the size increases so abruptly is another indication that the underlying distribution has changed. The emergence of a power-law-like distribution above 1 second will generate a large mean for durations >1 second, as expected for such a fat-tailed distribution -- in stark contrast, non-power-law distributions below 1 second will generate a smaller average, exactly as observed.

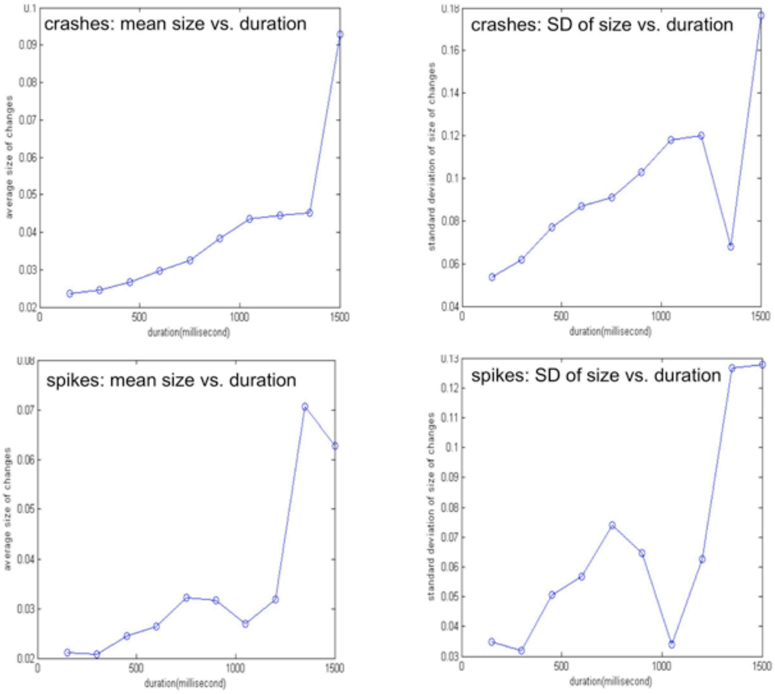

To remove any suspicion that our results could be seen more trivially, we also analyzed our dataset using more conventional techniques available in typical pre-packaged statistical software products. We found that the distribution of durations failed to match any of the standard distributions, including Normal, Lognormal, Weibull, Exponentional, Logistic, Smallest extreme value, and other multi-parameter variations of these distributions.

Figure 10. Average and standard deviation in the number of transactions making up the individual black swans which lie within a given duration window (e.g. top left panel shows average number of transactions per crash). This again supports the claim that there is a fundamental difference between the black swans in the subsecond and >1 second regimes. In particular, in the lower panels, the standard deviation appears to diverge for black swans with durations just below 1 second. Such divergent behaviour is typical of a phase transition in a physical system, where the scale of the fluctuations diverges at the transition point.

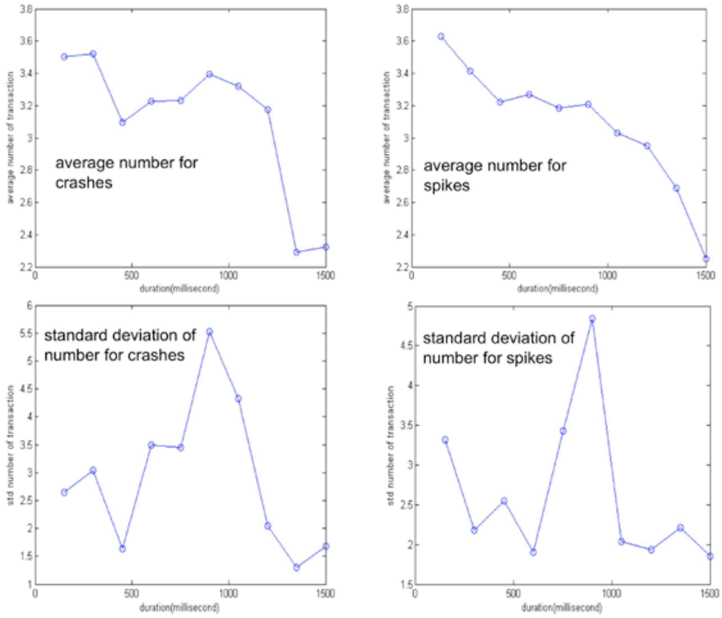

Figure 11. Lognormal distribution (straight dashed line) fit to complementary distribution function for black swan durations. The large deviation from the straight line, where the straight line represents a lognormal distribution, provides evidence supporting our claim of a transition point for durations around 1000 milliseconds (i.e. 1 second).

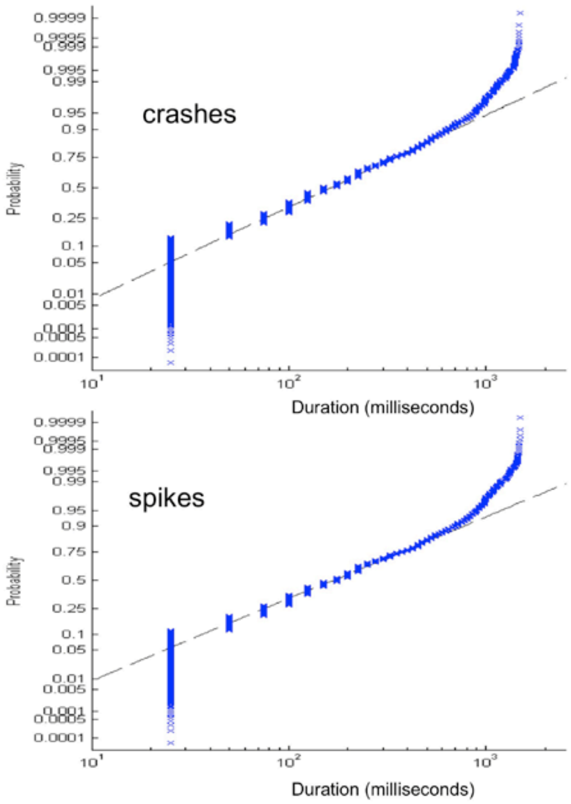

Figure 12. Extent to which the cumulative distribution for all crashes follows a power-law (top), and the subset with durations less than 1 second (lower left panel) and greater than 1 second (lower right panel). For crashes with durations more than 1 second, there is strong evidence for a power-law (*p*-value is 0.974). The appearance of a power-law for timescales larger than 1 second is consistent with the appearance of power-laws for the distribution of financial price-changes in the many studies in the literature for which increments of time larger than 1 second are chosen. By stark contrast, for crashes with durations less than 1 second, there is no evidence for a power-law – which offers support for the notion that subsecond black swans represent a new class of extreme event behavior.

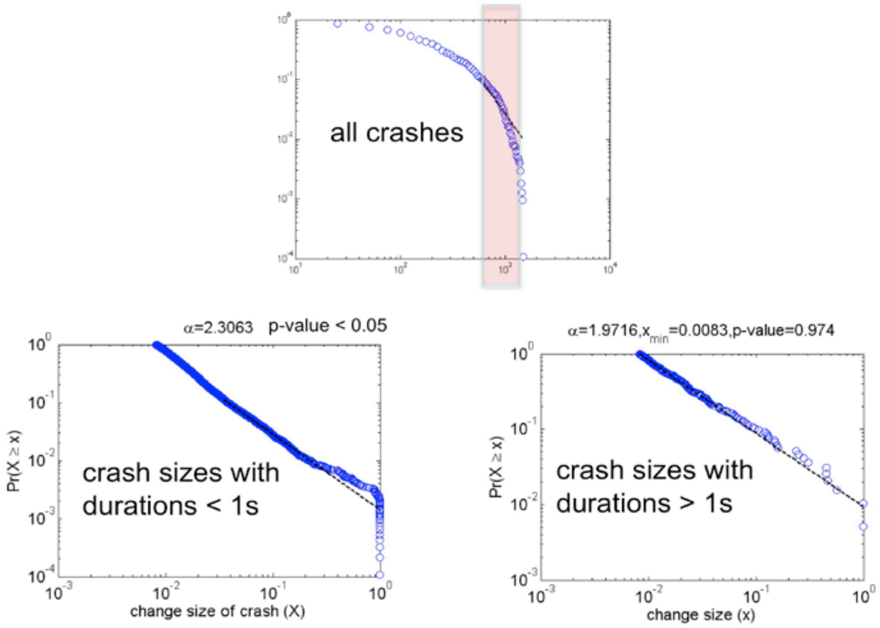

Figure 13. Extent to which the cumulative distribution for all spikes follows a power-law (top), and the subset with durations less than 1 second (lower left panel) and greater than 1 second (lower right panel). For spikes with durations more than 1 second, as for crashes in Figure 12, there is strong evidence for a power-law (*p*-value is 0.912). The appearance of a power-law for timescales larger than 1 second is consistent with the appearance of power-laws for the distribution of financial price-changes in the many studies in the literature for which increments of time larger than 1 second are chosen. By stark contrast, for spikes with durations less than 1 second, there is no evidence for a power-law – which again offers support for our claim that subsecond black swans represent a new class of extreme event behavior.

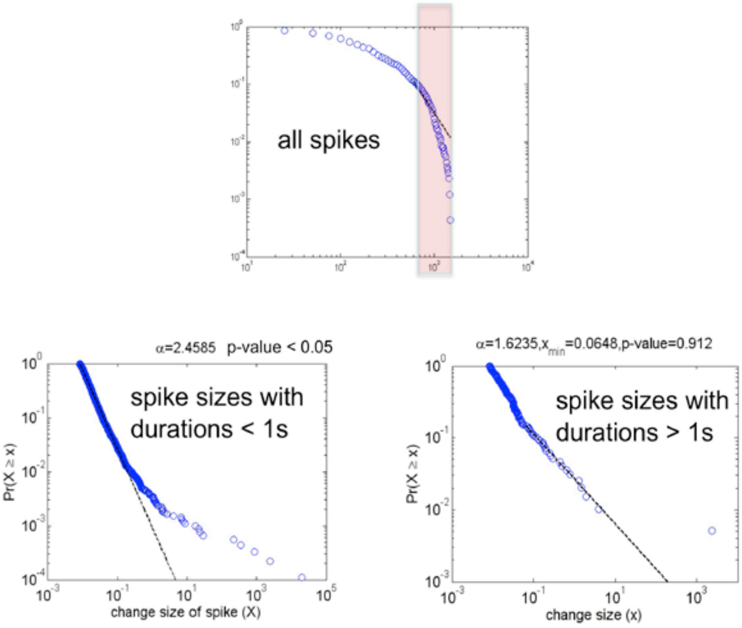

Figure 14. Reverse cumulative distribution function (CDF) for the fraction of crashes (left) and spikes (right) with durations within a given millisecond range, having a change size which is at least as big as the fraction shown on the horizontal axis. There is a gradual shifting of these curves with the change of duration range – which is again consistent with our claim that there is a different distribution for subsecond black swans as compared to black swans that have duration larger than 1 second.

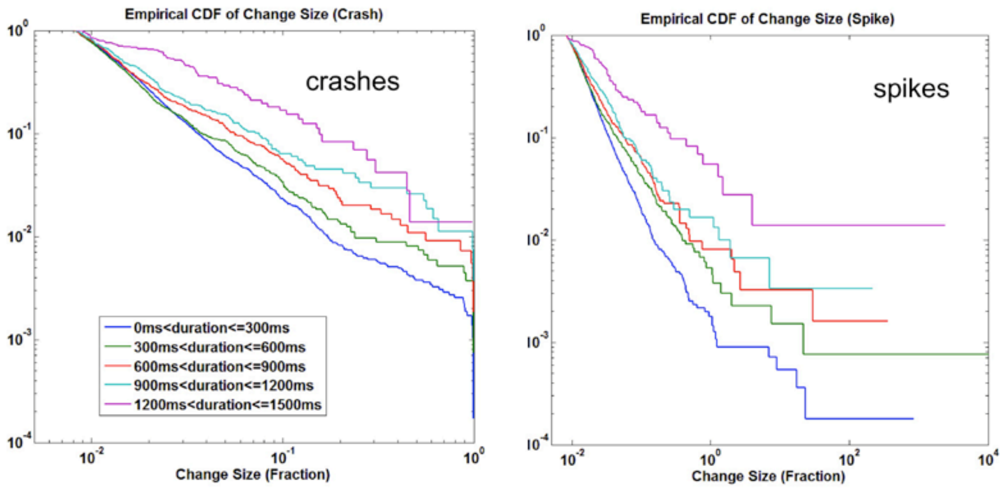

Figure 15. Empirical transition in size distribution for black swans with duration above duration threshold τ, as function of τ. Plots show results of the best-fit power-law exponent (black) and goodness-of-fit (blue) to the distributions for size of crashes and spikes separately. These plots show a visible transition for both the spikes and crashes.

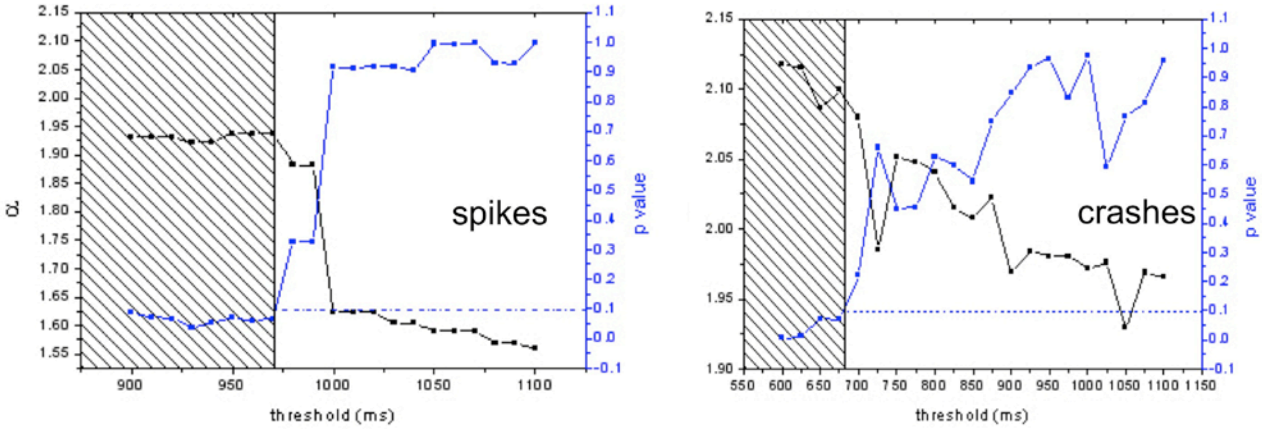

6. A new model for the ecology of subsecond markets

The results of the data analysis in the previous section provide evidence of an unexpected, yet fundamental phase transition arising below the one second timescale, to a regime where large jumps in the price are frequent. Given that a market is a collection of autonomous human and machine (i.e. computer algorithm) agents watching the latest prices before their next move, and yet human beings have limitations on how fast they can notice a particular situation and act on it, it is reasonable to try to relate this observed transition to the decreasing ability of human beings to influence price movements at smaller timescales.

Many models have been proposed to mimic real-world complex systems, however the famous 'El Farol' problem – upon which our model is based -- has long been considered archetypal (Arthur (1999), Johnson (2003)). Here we analyze a binary version of this model, which mimics how heterogenous agents (human or machine) use strategies and information about the recent past, to determine trading actions. While a suitable starting-point for populations of humans and/or computers, the machine-like binary nature of the model makes it ideally suited to discuss the subsecond regime dominated by computer trading algorithms – and indeed, the transition to this regime from the second-scale where machines and humans co-exist.

Our model comprises a population of *N* agents repeatedly competing for some limited resource (e.g. seating *L*) in a potentially crowded place (e.g. bar). They make decisions as to whether to attend on a given night based on the limited number of strategies *s* that they each have at their disposal, together with some limited information μ about the *m* most recent global outcomes. If the bar was undercrowded two nights ago, and overcrowded yesterday, this implies that there was an under-demand two days ago (i.e. excess demand $D<0$ which we denote as 0), and an over-demand yesterday (i.e. excess demand $D>0$ which we denote as 1). We make the reasonable approximation that the price is proportional to the excess demand, and that the excess demand D is proportional to the number of buyers minus sellers. Hence at each timestep, there is either an upward ($D>0$) or downward ($D<0$) pressure on the price. Translating this to our example in Figure 16, this means the price went down two days ago (i.e. outcome is 0 since there were more sellers than buyers) and up yesterday (i.e. outcome is 1 since there were more buyers than sellers). This down-up pattern is shown in Figure 16 as $\mu=\{0,1\}$. The agents then make a decision as to whether to buy or sell today, with each agent using his own highest-scoring strategy chosen from his own set of *s* strategies. As shown in Figure 16, the result is that the next day's outcome is a price fall (i.e. the outome is 0 since there are more sellers than buyers). This produces an updated bit-string for the price-change history $\mu=\{1,0\}$ for the next timestep. To ensure an unbiased market we set the bar seating capacity $L=(N-1)/2$, thereby turning the El Farol problem into a Minority Game in which the agents effectively compete to be in a minority group – below we discuss the financial relevance of the minority mechanism. This Minority Game has been studied extensively in the literature (Johnson et al. (2003)). In order to more closely mimic the real market in our model, we add an additional 'Grand Canonical' modification which means that any agent with poorly performing strategies at a given timestep (i.e. score over the previous *T* timesteps is below some threshold value *r*), does not play at that timestep. It is this feature that generates fluctuations in the 'volume' *V* of agents actively trading at a given timestep. This so-called Grand Canonical Minority Game, has been shown to reproduce the well-known stylized facts of financial markets (Johnson et al. (2003)).

Figure 16. The iterative decision-making process forming the core of our model. In this example, the length of the price history bit-string is *m*=2. Agents may be humans or machines, and are heterogeneous since they each have their own *s* strategies pulled from the space of available strategies (i.e. the strategy space). The information input at the beginning of each timestep is the current price history bit-string μ which lists the signs of the price-changes over the previous *m*=2 timesteps, e.g. down-up is encoded in binary form as $\mu = \{0,1\}$ as shown, where 1 means up and 0 means down. At each timestep each agent adopts his/her own current best (i.e. highest-scoring) strategy from his/her own set of *s* strategies. The strategy space is shown as a table, with the top row being the possible price-history bit-strings for *m*=2, and the entries in each subsequent row comprise actions +1 (i.e. buy) and -1 (sell). Hence each row represents a single strategy, i.e. it gives a well-defined response for each possible μ and hence each possible situation. At a given timestep, each agent chooses his best performing strategy and hence follows the action shown in the entry under column μ. All the agents follow this same procedure: They choose their own best strategy and receive its recommendation, i.e. buy (+1) or sell (-1). In the example timestep shown, more agents sell (-1) than buy (+1). The new price-change, given by an excess demand, is downward, i.e. it gets added to the price history μ as a 0. Therefore the new *m*=2 price history bitstring for the next timestep is $\mu = \{1,0\}$ as shown. This process then iterates in the same manner for all timesteps.

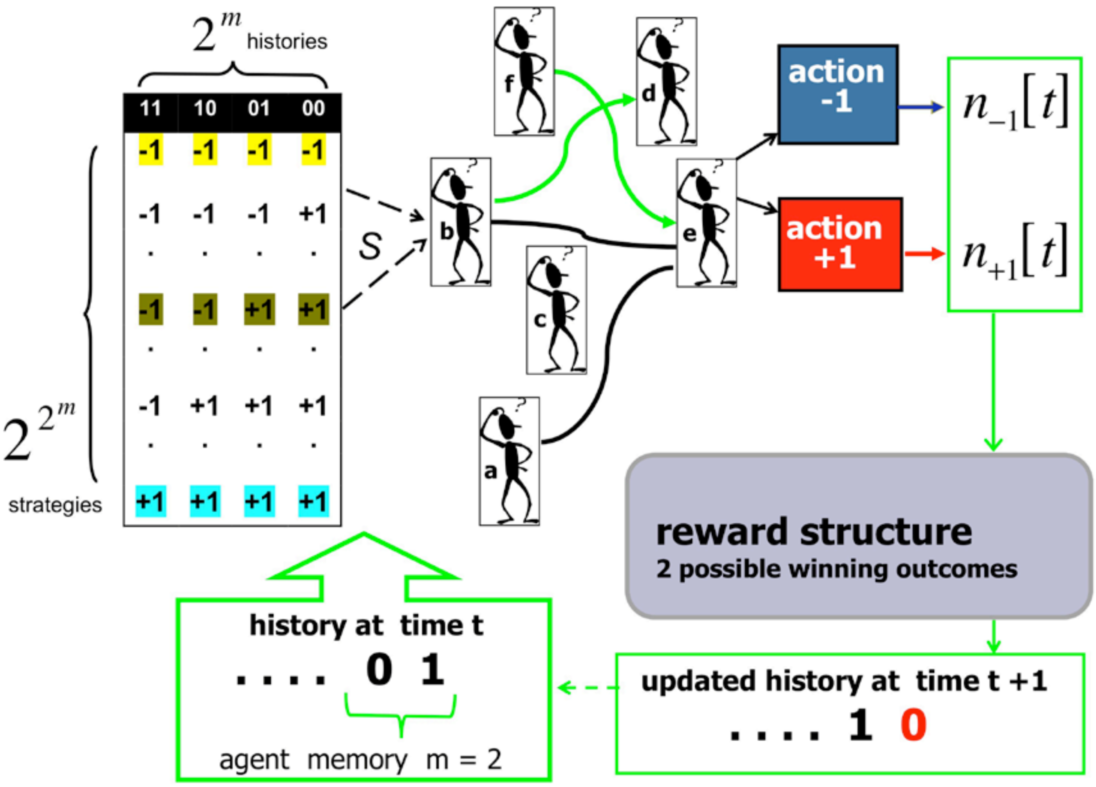

Tables 1 and 2 show explicitly why the goal of choosing the minority group (i.e. being a buyer when there is an excess of sellers, or vice versa) is a sensible goal for financial market agents – be they humans or machines -- with short-term, high-frequency trading goals. For example, an agent could be an automated trading platform with *s* being the number of algorithms that this platform manages. We define a notional wealth W_i of an agent i at time t as follows:

$$W_i[t] = \phi_i[t]x[t] + C_i[t]$$

where ϕ_i is the number of assets that it holds at time t, C_i is the amount of cash it holds at time t, and $x[t]$ is the asset price at time t. Since an exchange of cash for assets does not affect the agents' overall wealth at that moment, $W_i[t]$ is a notional wealth. The real measure of wealth is C_i, which is the amount of capital that the agent has available to spend. An agent has to do a 'round trip' (i.e. buy (sell) an asset then sell (buy) it back) to discover whether a real profit has been made.

Table 1. Trading with the minority mechanism. For the buy action, the agent (human or machine) is a buyer when the majority are sellers. Because of the negative market impact of an excess of sellers, the price $x[t]$ at which the trade is finally executed is below the advertised price $x[t]=10$ (i.e. it is executed at $x[t]=9$ in our example and hence only costs the trader 9 units of cash). For the sell action, the trader is a seller when the majority are buyers, hence the price $x[t]$ at which the trade is finally executed is above the advertised price $x[t]=9$ (i.e. it is executed at $x[t]=10$ in our example). Hence the agent ends up with 101 units of cash, having started with only 100. The agent has therefore made a profit of 1 unit of cash after the round-trip, simply by trading in the minority group, hence providing support for our model's mechanism of agents trying to trade in the minority.

t	**Action** $a[t]$	$C_i[t]$	$\phi_i[t]$	$x[t]$	$W_i[t]$
1	submit buy order	100	0	10	100
2	buy…, submit sell order	91	1	9	100
3	sell	**101**	0	10	101

Tables 1 and 2 show two examples of such a round trip, in which the agent (human or machine) trades with the minority decision and the majority decision respectively: Trading with the minority decision creates wealth for the agent on performing the necessary round-trip, whereas trading with the majority decision loses wealth. Of course, if the agent were a longer-term investor and had held the asset for a length of time between buying it and selling it back, his wealth would also depend on the rise and fall of the asset price over the holding period – however we wish to reflect the conditions of current and future markets which have, and will increasingly have, a large proportion of high frequency traders and/or automated trading platforms holding no long-term positions hence using the Minority Game reward mechanism (in which traders tend to buy/sell on one timestep and sell/buy back on the next) is a reasonable assumption.

Table 2. Trading with a majority mechanism. For the buy action, the agent (human or machine) is a buyer when the majority are also buyers. Because of the positive market impact of an excess of buyers, the price $x[t]$ at which the trade is finally executed is above the advertised price $x[t]=10$ (i.e. it is executed at $x[t]=11$ in our example and hence costs the trader 11 units of cash). For the sell action, the trader is a seller when the majority are also sellers, hence the price $x[t]$ at which the trade is finally executed is below the advertised price $x[t]=11$ (i.e. it is executed at $x[t]=10$ in our example). Hence the agent ends up with 99 units of cash, having started with 100. The agent has therefore lost 1 unit of cash after the round-trip, simply by trading in the majority group. Being in the majority, as opposed to the minority, is therefore undesirable.

t	action $a[t]$	$C_i[t]$	$\phi_i[t]$	$x[t]$	$W_i[t]$
1	submit buy order	100	0	10	100
2	buy…, submit sell order	89	1	11	100
3	sell	99	0	10	99

Figure 17. 'Strategy space' for an example case where two prior recent outcomes form the global information, i.e. the bit-string price history μ which represents the global information for each agent, has length $m = 2$ as in Fig. 16. The table on the left represents the $2^{2^m} = 16$ different strategies, given that $m = 2$. Here + represents action +1 (i.e. buy) and – represents -1 (i.e. sell). The greyed strategies are either totally uncorrelated or anticorrelated to each other (i.e. their respective dot-products are either zero or maximally negative). This set of greyed strategies form a skeleton representation of the full strategy space, and are referred to as the Reduced Strategy Space (RSS). There are $2.2^m = 2^{m+1} = 8$ strategies in the RSS. The figure on the right is a $2^m = 4$ dimensional hypercube which demonstrates the Hamming distance between strategies. The minimum number of edges linking strategies is the Hamming distance; for example, the dotted line shows a Hamming distance of 4 between strategies $----$ (i.e. strategy mandates to sell irrespective of history μ) and $++++$ (i.e. strategy mandates to buy irrespective of history μ).

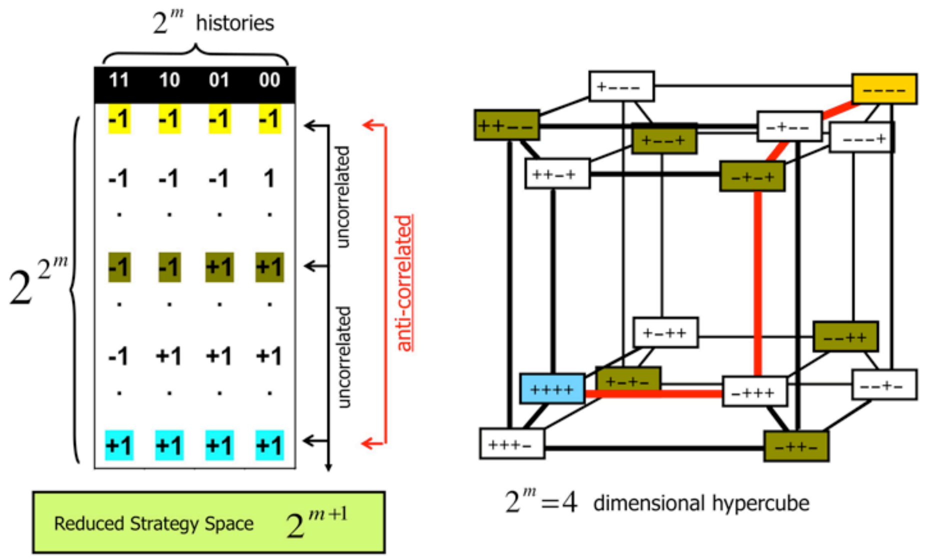

To formalize Figure 16, we assume a simple linear price formation process:

$$x[t+1] - x[t] = \frac{D\left[(t+1)^{-}\right]}{\lambda} \qquad (1)$$

where $(t+1)$ represents the time at which the new price $x[t+1]$ is announced and the buy/sell orders are executed, while $D\left[(t+1)^{-}\right]$ represents the excess demand in the market just prior to this time $(t+1)$. The market-maker uses the interval of time between $(t+1)^{-}$ and $(t+1)$ to deduce the new price. The scale parameter λ represents the market depth, i.e. how sensitive a market is to an order imbalance. We set $\lambda = 1$ for convenience. We assign a 0 to a downward price movement (i.e. $D < 0$) and a 1 to an upward one. In the unlikely event of the price-change at timestep *t* being zero, we flip a coin to determine the outcome (i.e. $D > 0$). Since the global information available to the agents is given by the m most recent price-change outcomes (0 for a price decrease and 1 for a price increase), the number of possible global information states is finite and equal to $P = 2^m$. For $m = 2$, there are only $P = 2^{m=2} = 4$ possible patterns in price, which are given by up-up (i.e. 11), up-down (i.e. 10), down-up (i.e. 01) and down-down (i.e. 00). These $P = 2^m$ different possible states of the global information variable form a space, with one state for each unique price history (e.g. 00, 01, 10 and 11 for $m = 2$). We will denote the state in this 'history space' at time t as the *decimal* equivalent of this string of m zeros and ones: $\mu[t] \in \{0 \ldots P-1\}$. For example, the history bit-string 00 corresponds to $\mu = 0$, 01 corresponds to $\mu = 1$, 10 corresponds to $\mu = 2$, and 11 corresponds to $\mu = 3$ (i.e. $\mu = 2^{m=2} - 1 = 3$). We denote the mandated action of strategy R, given global information (i.e. previous price pattern) $\mu[t]$, to be $a_R^{\mu[t]}$. The Rth element $a_R^{\mu[t]}$ corresponds to the action for strategy R given global information state $\mu[t]$, with action +1 meaning 'buy' and -1 meaning 'sell'. Each strategy therefore maps the present available global information $\mu[t]$ to an action $a_R^{\mu[t]} \in \{-1, +1\}$. The *s* strategies per agent are assigned in a random fashion before the simulation begins. The space of strategies can be broken down as shown in Figs. 16 and 17 for the example $m = 2$, using $\{-, +\}$ to denote the two possible actions $\{-1, +1\}$ for each $\mu[t]$. Each agent (human or machine) chooses the strategy, from his set of *s* strategies, which is the most successful judging from the past history of the market. He does this by using the tally of the success rate $S_R[t]$ for each of his strategies, with the score $S_R[t]$ of each strategy *R* being increased or decreased at each timestep by 1 according to whether it would have predicted a winning action or not. Hence different agents holding the same strategy *R* will agree on its relative merit. The model's dynamics can be described by trajectories on a DeBruijn graph as shown in Fig. 18, with each transition incurring a particular increment (plus or minus 1) to the score vector**.** This use of plus or minus ones as score increments, means that there are $P = 2^m$ orthogonal increment vectors $\underline{a}^{\mu}$ for the score vector $\underline{S}[t]$, one for each node μ. To mimic the general decrease in relevance attached to the more distant past by traders in a real market, we assume for simplicity that the agents only evaluate their strategy scores using the past *T* timesteps, and that they do so by attaching equal weight to each timestep. As mentioned earlier, our model then adds the 'Grand Canonical' generalization to this original Minority Game model, by allowing agents to only participate in a decision at a particular timestep *t* if their strategies have performed sufficiently well in the recent past (i.e. the score is above some threshold value *r* over the past *T* timesteps). The number of agents actively trading hence fluctuates over time, reminiscent of the real trading volume $V[t]$.

The resulting dynamics of our model are driven by the interplay between the deterministic dynamics of 'decided' traders, and the stochasticity of 'undecided' traders. Decided traders are the ones who, at a given timestep t, have a unique best strategy and hence a unique predicted action (i.e. buy or sell) for any given $\mu[t]$ – or equivalently, have two tied strategies with the same predicted action. Hence they do not need a coin-toss in order to decide which strategy recommendation to follow and hence which action to take (buy or sell). In other words, these agents would always take the same action when faced with a particular state of the system, hence the label 'decided'. This group is itself dynamic, i.e. a decided trader at one timestep can become undecided in the next timestep and vice versa. Undecided traders are the ones who, at a given timestep t, have tied strategies with different predicted actions, and hence invoke a coin-toss in order to act. These agents would only have a 50% chance of taking the same action if faced again with a particular state of the system. This group is also dynamic in that an undecided trader at one timestep may become a decided trader at the next timestep. It is this interplay of determinism and stochasticity – and in particular the existence of pockets of determinism associated with 'decided' groups – that underlie the ability to produce analytic analysis of the extreme events (see for example, Johnson et al. (2003)). In short, the net effect is to produce stochastically-perturbed deterministic dynamics, i.e. an underlying deterministic signal with a particular kind of added noise.

Figure 18. 'History space' for an example case where prior recent outcomes form the global information, i.e. the bit-string price history μ which represents the global information for each agent, has length 1, 2 and 3 respectively.

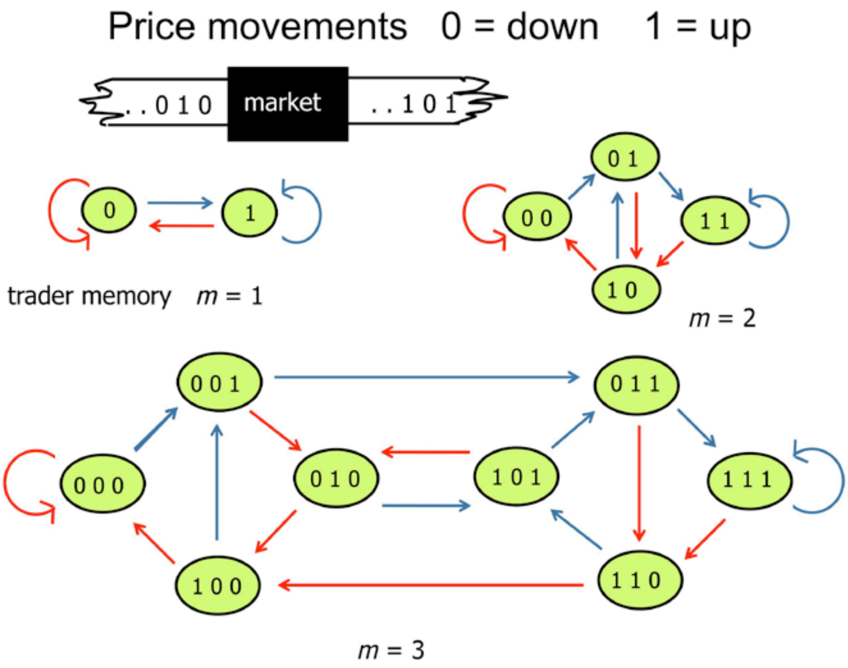

7. Phase transition within model consistent with behaviour observed for subsecond black swans

The task of understanding the details of the market dynamics in the subsecond regime, and hence providing a comprehensive and definitive quantitative explanation of the underlying causal mechanisms which explain the observed phase transition, is an open problem which will undoubtedly require many years of careful study, trawling through the entire range of data at the level of prices and individual orders. This is obviously beyond our scope. Instead we will present a plausible explanation of the general behaviour, building on the generalized El Farol framework described above. We will not consider detailed issues concerning the asynchronicity of order placement, preferring instead to keep a steady clock running with decisions and subsequent trades taking place at the ticks of this clock. We also do not consider the detailed complexity of possible algorithmic trading strategies. Despite these shortcomings, we are able to identify an interesting phase transition that emerges from our model, just as in the empirical data – moreover, we find that the features of the resulting price series change abruptly in a similar way to that observed in the empirical data. Although this in itself does not offer a proof that the model is correct, it does provide a concrete framework in which to discuss the variation of the various system parameters, each of which has a simple physical interpretation for market regulators. As a result, we are able to focus in on what we believe to be one *crucial, yet little understood, issue facing the high frequency regime of computer trading: the knock-on effects of crowding in strategy space*.

With these caveats in mind, we first consider our model (Fig. 16) in the regime in which the total number of different strategies in the market (which is 2^{m+1} in our model) is typically larger than the total number of agents N (i.e. $\eta > 1$ where $\eta = 2^{m+1}/N$). This is the right-hand regime in Figs. 19 and 20. We associate this regime with a market in which both humans and machines are dictating prices, and hence timescales above the transition (>1s), for the following reasons: The presence of humans actively trading -- and hence their individual 'free will', together with the myriad ways in which they can manually override algorithms -- means that the effective number (i.e. diversity) of strategies should be extremely large (i.e. $\eta > 1$). Moreover $\eta > 1$ implies m is large, hence there are more pieces of information available which suggests longer timescales (there will be more millisecond price movements in the past 1000ms than in the past 500ms). Since by definition $N/2^{m+1} < 1$ in this $\eta > 1$ regime, the average number of agents per strategy is less than 1, hence any crowding effects due to agents coincidentally using the same strategy will be small. This lack of crowding leads our model to predict that any large price movements arising for $\eta > 1$ will be rare and take place over a longer duration (see Fig. 20, right-hand panel) – exactly as observed in our data for timescales above 1000ms. Indeed, we have already shown that our model's price output in this $\eta > 1$ regime can reproduce the stylized facts associated with financial markets over timescales longer than 1 second, including a power-law distribution (Johnson et al. (2003)).

Our model then undergoes a transition around $\eta \approx 1$ to a regime characterized by significant strategy crowding and hence large fluctuations. (See Fig. 19, left-hand regime). The price output for $\eta < 1$ (see Fig. 20, left-hand panel) shows frequent abrupt changes due to agents moving as unintentional groups into particular strategies. Our model therefore predicts a rapidly increasing number of ultrafast black swan events as we move to smaller η and hence smaller subsecond timescales – mimicking what we observed in the actual black swan data as discussed earlier. Our association of the $\eta < 1$ regime with an all-machine phase is consistent with the idea that trading algorithms in the sub-second regime are likely to be designed to be executable extremely quickly and hence be relatively straightforward, without calling on much memory concerning past information: In this case m will be small, so the total number of strategies will be small and therefore $2^{m+1} < N$ which means $\eta < 1$. Our model also predicts that the size distribution for the black swans in this ultrafast regime ($\eta < 1$) should not have a power law since changes of all sizes do not appear – this is again consistent with the results presented earlier.

Although our model obviously ignores many potentially important details about the real market -- including the fact that actual markets do not run in a perfectly synchronous way – the model's simplicity allows us to derive precise mathematical formulae for the scale of the price fluctuations in each phase if we make the additional assumption that the number of agents playing each timestep is similar to N. For $\eta < 1$ (see Figs. 19 and 20) we can prove analytically that the standard deviation σ of the price fluctuations has a lower bound given by $\sigma = 3^{-\frac{1}{2}} 2^{-\frac{(m+1)}{2}} N\left(1-2^{-2(m+1)}\right)^{\frac{1}{2}}$ for $s = 2$, and an upper bound which is a factor of $\sqrt{2}$ bigger given by $\sigma = 3^{-\frac{1}{2}} 2^{-\frac{m}{2}} N\left(1-2^{-2(m+1)}\right)^{\frac{1}{2}}$. Full derivations of these results are available from the author on request. For $\eta > 1$, σ is given approximately by $\sigma = N^{\frac{1}{2}}\left(1-2^{-(m+1)}N\right)^{\frac{1}{2}}$ for general s. Our model's prediction that σ is proportional to N for $\eta < 1$ as compared to $N^{\frac{1}{2}}$ for $\eta > 1$, provides an analytic explanation for the empirical finding that there are many more black swans at shorter durations, while the transition in σ around $\eta \approx 1$ explains the abrupt change in the character of their distribution. Since σ plays a fundamental role in traditional finance as a measure of risk, these explicit formulae and their parameter dependencies could be used to help quantify the effect of changes in regulations on conventional risk measures.

One might attempt to offer an alternative hypothesis, claiming that the extreme subsecond behaviour that we observed in the real data are generated by external news arrival, as opposed to being truly endogenous and hence generated from within the system by the actions of the agents themselves. However, given that typical daily news is neither 'good' nor 'bad' for the entire market, and given that the black swan movements happen so quickly on the subsecond scale, it seems highly unlikely that these effects are being generated by news arrival. Instead, it is as if – prior to each subsecond black swan -- the market has built up some kind of internal pressure like a wound-up spring, which is then quickly released. Indeed, as we will show, this is exactly the causal effect that arises endogenously within our model.

Figure 19. Schematic showing the phase transition which arises in our model, for constant *N*, where $\eta = 2^{m+1}/N$. The term 'memory' is used to denote the length of past information bit-string *m* that the agents use to make their trading decision. Small *m* implies $\eta < 1$ and hence many agents per strategy, leading to large crowding which produces frequent, large and abrupt price-changes, i.e. high number of short-duration ($<< 1$ second) black swans, as observed empirically. Large *m* implies $\eta > 1$ and hence very few, if any, agents per strategy, hence small crowding. Therefore large changes are rarer and last longer for $\eta > 1$, i.e. low number of longer-duration black swans, as observed empirically.

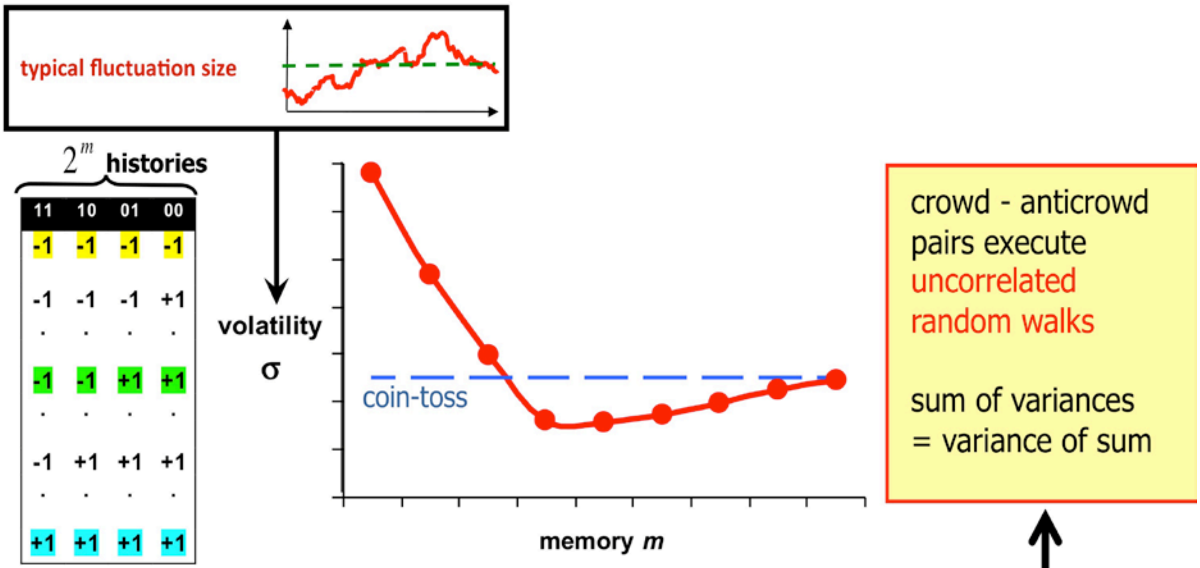

Memory m	$2^{m+1} << N.s$	$2^{m+1} \sim N.s$	$2^{m+1} >> N.s$
Crowd size	large	medium	~ 1
Anticrowd size	small	medium	~ 0
Net crowd – anticrowd pair size	large >> 1	small	small 1
# crowd - anticrowd pairs	$\sim 2^m$ $<< N$	$\sim 2^m$ $< N$	$< 2^m$ $\sim N$

We now provide an intuitive explanation of the expressions that we presented for the market volatility in the new subsecond machine phase, and hence the behaviour in Fig. 19 for the predicted market volatility. As shown in Fig. 19, the low-m phase is characterized by a decrease in σ as m increases: this is the 'crowded' phase we discussed earlier where the number of strategies 2^{m+1} (in the RSS) is small compared to the number of agents N. The high-m phase is characterized by a slow increase in σ toward some limiting value as m increases. This phase is called the 'dilute' phase since the number of strategies 2^{m+1} (in the RSS) is now large compared to the number of agents *N.* Our simple qualitative picture relates back to the correlation between strategies. In the 'crowded' phase, i.e. at small m, there will at any one time be a large number of machines who are using the same (e.g. the perceived best) strategy and so will flood into the market as large groups or *crowds*, producing large swings in demand and hence a high volatility as shown. If instead the length of information m being used by the agents is large, then the crowd of agents using the same strategy will be smaller simply because many may not hold the best strategy – the chances of a given agent holding the instantaneous best strategy decrease as m increases. There will also be groups of agents who are forced to use the anti-correlated (e.g. the perceived worst) strategy: these can be thought of as *anti-crowds* since they cancel out the market action of the crowds at every timestep *t* *regardless* of the particular history bit-string at that timestep. This cancellation effect causes a reduction in the size of the market volatility. In the dilute phase of very large memory m, it is very unlikely that any two agents will hold the same strategies and so the market can be modelled as a group of independent coin-tossing agents.

Figure 20. Theoretical transition. Model output for the two regimes of strategy distribution among agents ($\eta = 2^{m+1}/N$) together with timescales (top). The regime $\eta < 1$ represents a 'crowded' strategy space phase, while $\eta > 1$ is the 'dilute' phase. $\eta < 1$ implies many agents per strategy, hence large crowding which produces frequent, large and abrupt price-changes, i.e. high number of short-duration ($<< 1$ second) black swans, as observed empirically. $\eta > 1$ implies very few, if any, agents per strategy, hence small crowding. Therefore large changes for $\eta > 1$ are rarer and last longer, i.e. low number of longer-duration black swans, as observed empirically for $\eta > 1$.

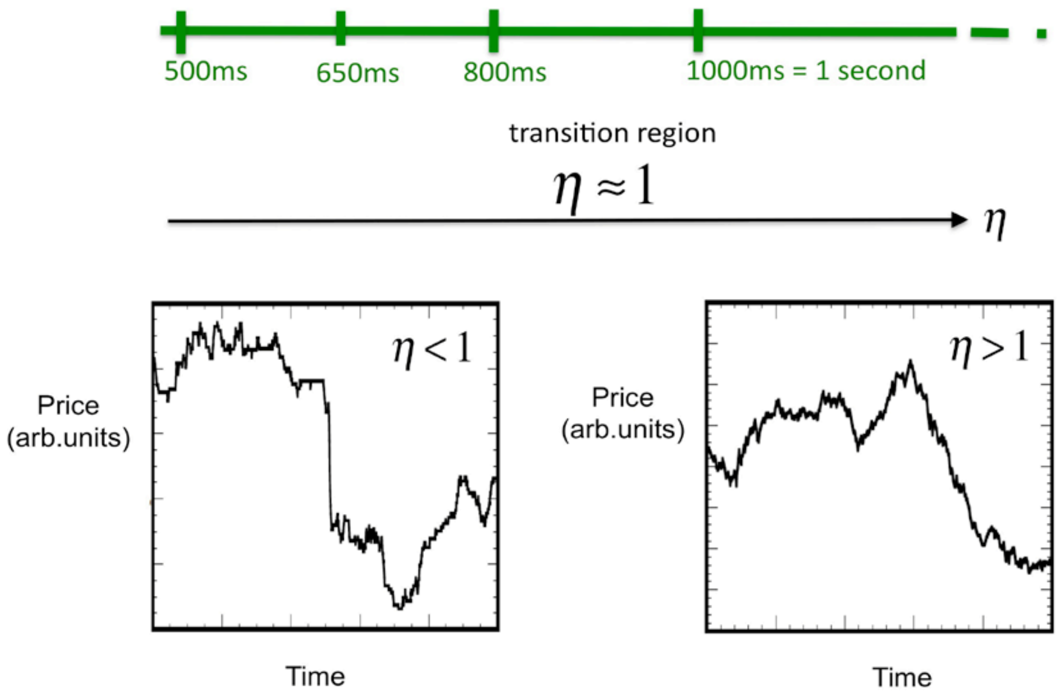

We can now make this new understanding of subsecond computer trading volatility slightly more quantitative, after first establishing a few basic results of statistics. Consider a random walk along the *y*-axis, with step-size = d and number of steps = *N*. The probability of moving in a positive (negative) direction at each step is p, (q) where $p+q=1$. The mean displacement $y_{N=1}$ for $N=1$ is given by: $\langle y_{N=1}\rangle = p\,d + q(-d) = (p-q)d$ hence $\langle y_{N=1}\rangle = 0$ if $p=q=\frac{1}{2}$. To calculate the variance $\sigma^2_{N=1}$ for $N=1$, we start with $\langle {y_{N=1}}^2\rangle = p\,d^2 + q(-d)^2 = (p+q)d^2 = d^2$

hence

$$\sigma^2_{N=1} \equiv \langle {y_{N=1}}^2\rangle - \langle y_{N=1}\rangle^2 = d^2 - (p-q)^2 d^2 = d^2\left[1-(2p-1)^2\right] = 4\,p\,q\,d^2$$

Let us consider uncorrelated steps: It is a well-known result that the variance (or average) of the sum is equal to the sum of the variances (or averages). The mean displacement y_N for $N \geq 1$ is therefore given by $\langle y_N\rangle = N\langle y_{N=1}\rangle = N(p-q)d$ and hence $\langle y_N\rangle = 0$ if $p=q=\frac{1}{2}$. The variance σ^2_N for $N\geq 1$ is therefore given by $\sigma^2_N = N\sigma^2_{N=1} = 4N\,p\,q\,d^2 \equiv \sigma^2$. Hence $\sigma^2_N \equiv \sigma^2 = Nd^2$ if $p=q=\frac{1}{2}$. Note that $\sigma^2_N \equiv \sigma^2 = N$ if $p=q=\frac{1}{2}$ and $d=1$. So, turning to the model, we first consider the oversimplified case of N independent agents each deciding on an investment decision by tossing a coin. Each agent therefore provides a random-walk process in terms of increasing or decreasing the demand by 1. Assume for the moment that these coin-tosses are uncorrelated. Using a standard result of undergraduate statistics, the total variance σ^2 for this random-walk in excess demand, is given by the *sum* of the individual variances produced by each of the N agents. If the agent decides 1, then he contributes 1 to the excess demand. If, by contrast, the agent decides -1, then he contributes -1 to the excess demand. In both cases the random-walk 'step-size' is $d=1$. This coin-tossing agent chooses 1 with probability 1/2, and -1 with probability 1/2. The variance contributed to σ^2 by each agent is therefore given by 1 since $d=1$. Summing over all N agents, the total variance in the excess demand σ^2 is given by *N*. Hence the standard deviation (i.e. volatility) of demand is given by $\sigma=\sqrt{N}$ which, for $N=101$, gives $\sigma=10.0$ which is the dashed 'coin-toss' line in Fig. 19.

In reality, on any given turn of the game, there will be a number of agents using the same, or similar, strategies. Consider the subset of agents n_R using a particular strategy R. Although there is no information available to a given agent about other individual agents, nor is any direct communication allowed between agents, this subset of agents n_R using a particular strategy R will all make the same investment decision at each timestep *irrespective* of the particular history bit-string for that timestep. Hence they will act as a *crowd*. Since the corresponding random-walk 'step-size' that this crowd contributes is $d \equiv n_R$, this crowd should contribute a variance $4pqd^2 = 4\frac{1}{2}\frac{1}{2}{n_R}^2 = {n_R}^2$ to the total variance. However because of the initial strategy allocation, there may also be a subset of agents $n_{\overline{R}}$ who are using the anticorrelated strategy to R, i.e. $\overline{R}$. This second group, the *anti-crowd*, makes the *opposite* investment decision to the crowd at each timestep *irrespective* of the particular history bit-string for that timestep. Over the timescale during which these two opposing strategies R and $\overline{R}$ are being played, the fluctuations are determined only by the net crowd-size $n_R^{\mathrm{eff}} = n_R - n_{\overline{R}}$ which constitutes the net step-size of the crowd-anticrowd pair. Hence the net contribution by this crowd-anticrowd pair to the random-walk variance, is given by $4pqd^2 = \left[n_R^{\text{eff}}\right]^2$. We will now use this result. Suppose strategy R^* is the highest scoring at a particular moment: the anti-correlated strategy $\overline{R^*}$ is therefore the lowest scoring at that same moment. In the limit of small m, the size of the strategy space is small. Each agent hence carries a considerable fraction of all possible strategies. Therefore, even if an agent picks $\overline{R^*}$ among his s strategies, he is also likely to have a high scoring strategy. Therefore, many agents will choose to use either R^* itself (if they hold it) or a similar one. Very few agents will have such a poor set of strategies that they are forced to use a strategy similar to $\overline{R^*}$. In this regime there are practically no anticrowds, and the crowds dominate. Therefore $n_R \sim N\,\delta_{R,R^*}$ and hence $n_R^{\text{eff}} \sim N\,\delta_{R,R^*}$. Hence the variance varies as $\sigma^2 \sim N^2$ and is larger than the independent agent limit of N, in agreement with the plots in Fig. 19. By contrast in the limit of large m, the strategy space is very large and agents will have a low chance of holding the same strategy. Even if an agent has several low-scoring strategies, the probability of his best strategy being strictly anti-correlated to another agent's best strategy (hence forming a crowd-anticrowd pair) is small. All the crowds and anticrowds will tend to be of size 0 or 1, implying that the agents act independently. This yields the coin-toss limit discussed above. In the intermediate m region where the minimum in the observed volatility exists, the size of the strategy space is relatively large. Hence some agents may get stuck with s strategies which are all low scoring at a particular timestep. They hence form anti-crowds. Considering the extreme case where the crowd and anti-crowd are of similar size, we have $n_R^{\text{eff}} \sim 0$ and hence the volatility is essentially zero. This is again consistent with the numerical results. The regime of small volatility will arise for small s since, in this case, the number of strategies available to each agent is small – hence some of the agents may indeed be forced to use a strategy which is little better than the worst-performing strategy $\overline{R^*}$. In other words, the cancellation effect of the crowd and anticrowd becomes most effective in this intermediate *m* region for small s. Increasing s should make this minimum less marked, as again observed numerically.

8. Quantitative description of the extreme behaviour in the crowded algorithm regime

Because of our interest in the instabilities created by the crowding of computer trading algorithms, we will focus on the crowded regime in which the number of active agents *N* is larger than the diversity of strategies in the game (i.e. the ratio $\eta < 1$ where $\eta = 2^{m+1}/N$ and where 2^{m+1} is the number of strategies in the Reduced Strategy Space (RSS). Because of the dense packing of agents onto each possible strategy in this regime, the details of the full strategy space become less important and hence the skeleton-like RSS can be used as a good approximation). As discussed in previous sections, this is the regime in which crashes and spikes are generated: Any given strategy *R* is likely held by many agents -- hence when *R* becomes the best strategy, an appreciably sized crowd of agents will then use it and hence will buy or sell at the same time. This generates an extreme change which is rather deterministic in nature. As this strategy is repeatedly used, its relative advantage diminishes and eventually the crowd is broken up – hence the extreme behavior stops. It turns out that we can describe the actual path taken by the price, both prior to and during a large change, in terms of the nodal weights of the strategy score (see Johnson et al. (2003) for details). The upshot is that nodal imbalances then generate the large changes that we observe. Interestingly, Satinover and Sornette (2008) introduced a similar concept concerning persistent behavior when studying this type of model. Indeed, the concept of persistence in transistions is closely related to these authors' explanation of why events are correlated, producing extreme chains of events; and also their more general explanation for pockets of predictability. It also illustrates why the more extreme events are actually more deterministic, and hence more predictable, and underlies Sornette's suggestion of the term 'dragon-kings' in preference to 'black swans' (Sornette (2009)).

Instead of describing both the initial drop (or rise) *and* the recovery as in Figs. 4(a) and 4(b), we can just concern ourselves here with the initial large drop or rise in price, as shown in Fig. 21. This is because the subsequent recovery may be either an endogenous reaction or may come from exogenous effects such as enforced trading rules (e.g. stopping selling) which are outside our simple model. In any case, the recovery portion can simply be seen as a similar process to the initial drop/rise but in reverse. The simplest type of single large movement which exhibits perfect nodal repetition, would be $\mu = 0,0,0,\dots$ in which all successive changes are in the *same* direction (i.e. repeated red transitions from node 0 to 0 in Fig. 18). We call this a 'fixed-node crash'. During the fixed-node crash, agents are likely to be deploying only a single strategy because only one strategy at a time is likely to have a score above the threshold for trading. These agents are thus behaving in a non-adaptive, more deterministic way.

Figure 21. Left: Example of a price crash without a recovery, as generated by our model with $m = 2$, within the Reduced Strategy Space (RSS). Right: The corresponding weights (with magnitudes indicated by the darkness of the tints) for each node μ, as a function of time through the lifetime of the crash. The global information at each timestep is indicated by the black square. The crash is preceded by abnormally high nodal weight magnitude on node 0 (darker tints) yielding the tendency to repeatedly visit node 0 once that node is initially hit. The overall crash incorporates fixed-node and cyclic-node crashes (i.e. it a composite of two types of crash, the fixed-node crash and the cyclic-node crash).

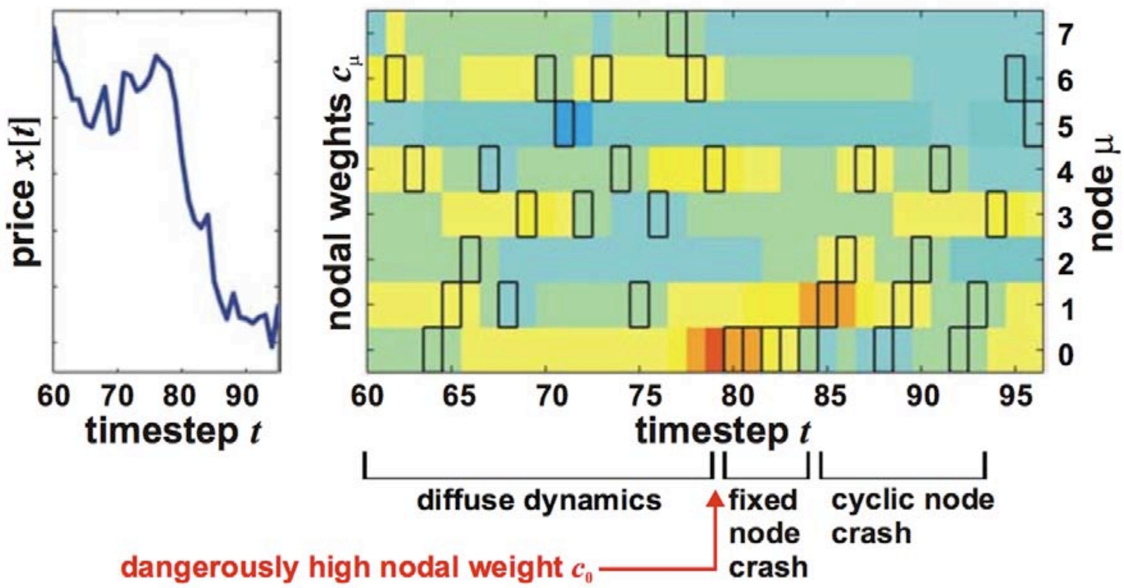

Interestingly, our model predicts various other possibilities for the behavior of the system during a crash, which is consistent with the fact that there are a range of possible forms and durations of the actual price-series during a real-world financial crash. For example, on the $m = 3$ DeBruijn graph in Fig. 18, the cycle $\mu = 0,0,1,2,4,0\dots$ has four out of the five transitions producing demands of the same sign (it repeats the outcomes on nodes $\mu = 1, 2, 4$ and produces opposite outcomes on node $\mu = 0$). We call this a 'cyclic-node crash'. Figure 21 shows an example which starts as a fixed-node crash and then subsequently becomes a cyclic-node crash. Cyclic-node crashes can be treated simply as interlocking fixed-node crashes which repeat themselves. Hence for clarity we focus on a single fixed-node crash. The presence of abnormally high nodal weights (particularly on a closed subset of connected nodes μ) will cause a large movement in the system *if* the system's trajectory hits any of these susceptible nodes in the global information space. We note that recoveries can indeed emerge spontaneously from our model (in particular for the $\eta < 1$ regime corresponding to very short durations as in Figs. 4(a) and 4(b)) without having to invoke external regulations or additional predatory algorithms. As the model updates μ at each timestep, it traces a trajectory around the network in Fig. 18. Each node acts like a coiled spring, in that the bigger weighting, the greater the tendency of the system to return to that node. Any trajectory that comprises mostly negative (positive) transitions will produce a large price drop (rise). Figure 21 shows these weightings and the overall model trajectory (green line) as it moves between nodes, expressed in their decimal representation with $\mu = \{000\}$ equivalent to 0 etc. Prior to the initial price drop in Fig. 21, there is a large positive weight (blue) on node 0 (i.e. at $\mu = \{000\}$ in Fig. 21). When the model's trajectory hits node 0, this large weight triggers repeated transitions back to node 0, like a spring uncoiling, producing a large number of consecutive negative price changes -- hence the large price drop.

Figure 22. The range of possible ‘flavours’ of fixed-node crash, in the subsecond computer trading regime, as system parameters are varied. In other words, by choosing a particular set of parameters concerning the distribution of strategy scores, different combinations of price and volume variation can arise during this particular type of extreme behavior, such that a given large change in price may have a range of different possible volume profiles over time. This variety in price-volume behaviours during a crash is consistent with what is observed in real markets.

!

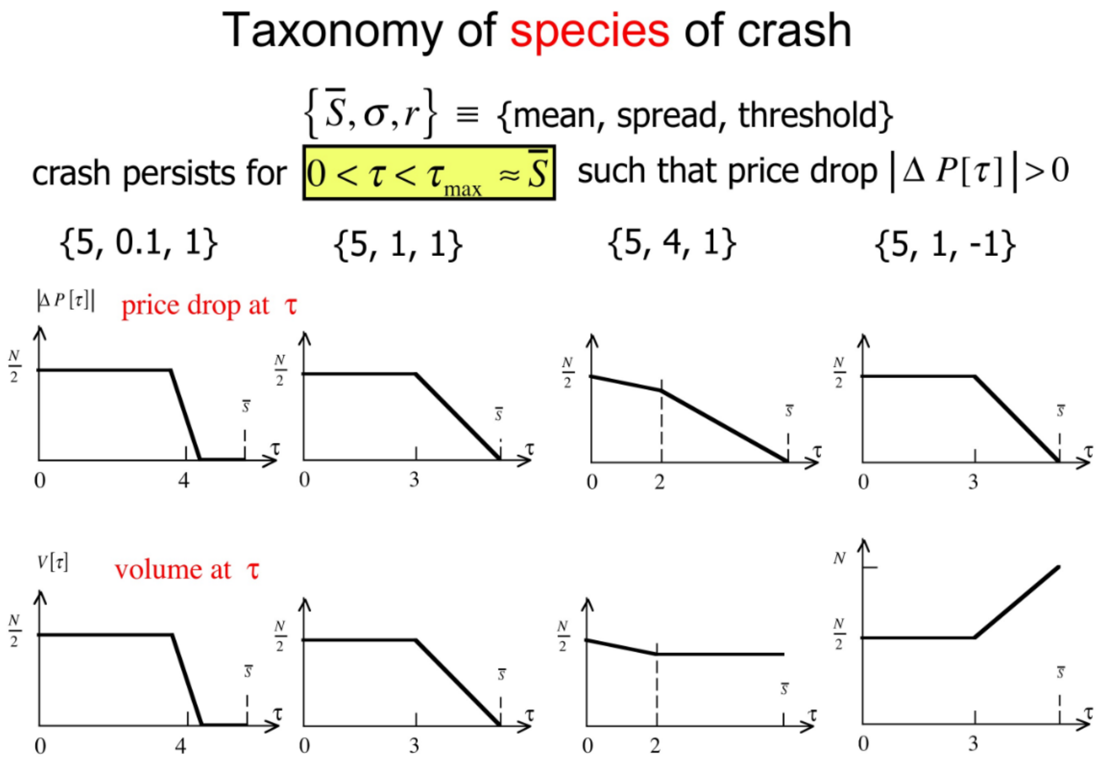

!

9. Consequences for next-generation risk management

How prices behave is not just an important point for academics or practitioners looking to make point predictions, it also has a direct impact on how risk is managed, and whether existing models of risk evaluation and derivative pricing can be relied upon. This is because well-established derivative pricing calculations such as Black-Scholes equation, assume certain characteristics of the time-series. If the actual market does not follow these assumptions, then the calculations are in principle wrong. Of course, this is known – and in reality the Black-Scholes approach is typically adapted by traders and risk-managers to try to account for these shortcomings. However, as we have seen, the subsecond machine phase is characterized by a very different distribution and price pattern, breaking the assumptions and making the requirement for a more general approach essential. Here we sketch out such an approach, which could be developed over the next decade as more information emerges concerning subsecond price dynamics. The attractive feature of our proposed formalism for managing risk on the subsecond scale, is that it does not treat the price changes as merely some perturbation away from a Gaussian-like form, nor does it treat time as continuous or following a discrete but regular clock-tick. Instead the format is general and can be adapted to a wide variety of price behaviours, whether they have already been observed in real price-series or await to be observed in the future.

The Black-Scholes equation (Bouchaud and Potters (2003)) showed how in theory it is possible to never lose any capital through writing an option (i.e a derivative). In particular, the variation of the option writer's wealth always remains zero: $\Delta W_T = 0$ and hence 'zero-risk'. The main underlying assumptions for this to be correct are: (1) Continuous time: continuous trading; (2) Efficient markets: no arbitrage; (3) Underlying assets follow a random walk. However, all three of these assumptions become inappropriate in the subsecond regime. In particular, in the context of hedging risk, the assumption of continuous time and hence continuous trading become highly suspect. In addition, the presence of transaction costs gives rise to a financial barrier to high-frequency trading: the greater the number of re-hedgings, the greater the cost. There will therefore be a trade off. However the extent and manner in which different suppliers of a contract will judge this risk, can be very different: after all, there are many ways of assessing risk and hence calculating an adequate compensatory 'risk premium'.

The equation for the variation of the option writer's wealth ΔW_T, in compact form, is $\Delta W_T = V_0 - V_T + H$ where H is the term corresponding to the gain or loss from hedging assets, V_0 is the option premium and V_T is the option payout. The variance of the wealth is:

$$\left\langle \Delta W_T^{\,2} \right\rangle - \left\langle \Delta W_T \right\rangle^2 = \begin{array}{l} \left(V_0^{\,2} + \left\langle V_T^{\,2} \right\rangle + \left\langle H^2 \right\rangle - 2V_0 \left\langle V_T \right\rangle - 2\left\langle V_T H \right\rangle + 2V_0 \left\langle H \right\rangle \right) - \\ \left(V_0^{\,2} + \left\langle V_T \right\rangle^2 + \left\langle H \right\rangle^2 - 2V_0 \left\langle V_T \right\rangle - 2\left\langle V_T \right\rangle \left\langle H \right\rangle + 2V_0 \left\langle H \right\rangle \right) \end{array}$$

where $\langle \ldots \rangle$ is our shorthand for averaging over all underlying asset price realizations. Cancelling, and using the fact that for unbiased increments of the underlying asset we have $\langle H \rangle = 0$, we get:

$$\begin{aligned} \left\langle \Delta W_T^{\,2} \right\rangle - \left\langle \Delta W_T \right\rangle^2 &= \left(\left\langle V_T^{\,2} \right\rangle - \left\langle V_T \right\rangle^2 \right) + \left(\left\langle H^2 \right\rangle - 2\left\langle V_T H \right\rangle \right) \\ &= \mathrm{R}_c + \left(\left\langle H^2 \right\rangle - 2\left\langle V_T H \right\rangle \right) = \mathrm{R} \end{aligned}$$

Therefore $\langle \Delta W_T \rangle = \lambda \sqrt{\mathrm{var}[\Delta W_T]} \;\Rightarrow\; V_0 - \langle V_T \rangle = \lambda \sqrt{\mathrm{R}} \;\Rightarrow\; V_0 = \langle V_T \rangle + \lambda \sqrt{\mathrm{R}}$, yielding an additive term to the option price, which is proportional to the standard deviation in the option writer's variation of wealth. One could use this equation to assess an option writer's degree of risk-aversion based on traded market option prices V_0. This gives an idea of how 'expensive' the option is: the higher the risk-aversion λ, the more the option will cost in excess of the 'fair' price $\langle V_T \rangle$. Most importantly, we have made no assumptions concerning the characteristics of the price-series, hence this approach is perfectly applicable in the subsecond computer trading regime – and may ultimately be used to develop a new generation of ultrafast derivatives to make the subsecond trading regime less risky.

Although we know of no other work that looks into such short timescale options, or even proposes them, this does not mean that such products will not come – after all, machines running ultrafast algorithms may themselves be programmed to hedge risk as they trade. More generally, such subsecond options can potentially provide new hedging opportunities for a wide range of market participants and market activities, making the global market more complete and potentially providing ultrafast insurance contracts for operations at the edge of human response times. Whether options at this subsecond level might eventually prevent or provoke fractures from propagating to other scales, remains to be explored. At the very least, it is an interesting academic exercise to consider such ultrafast options, and one that should ultimately relate to important pricing practices in the emerging subsecond market ecosystem.

10. Summary and outlook

The strategic advantage to a financial company of having a faster system than its competitors is currently driving a billion-dollar technological arms race to reduce communication and computational operating times down toward the physical limits of the speed of light – orders of magnitude below human response times. Given the markets' drive toward ever faster technologies, there is an urgent need to understand the new 'black swan' phenomena that might emerge, such as in Fig. 4.

We have here attempted to deepen our understanding of what happens in this ultrafast regime, first by analyzing state-of-the-art data from multiple stock and across multiple exchanges, and then presenting a simple yet highly non-trivial model which mimics a population of competitive trading agents (humans and/or machines). This model produced a number of features which are similar to those observed in the market data, most importantly the rapid proliferation of black swan events as the timescale moves below the timescale of human reaction times, and a rather abrupt change in the distribution of their size. These features are reminiscent of a phase transition in physical systems. The interpretation within the model is that this transition corresponds to the loss of human participants at subsecond timescales, due to their physical and mental limitations which prevent them from acting fast enough. This in turn produces a shrinkage of the available strategy space, which then leads to crowding by computer trading machines, and hence frequent and large price instabilites (Fig. 20, left-hand panel).

Of course, much remains to be explored concerning this 'strategy crowding' scenario. Indeed, although we presented it in terms of relatively simple computer algorithms acting inadvertently in unison, it may be that more sophisticated algorithms will be developed which are implementable in hardware. However the same argument holds: whether it be because of shared employees or simply shared or common ideas, the particular techniques implemented in the strategies will have a tendency to be replicated and used by many participants, hence yielding strategy crowding and observed behaviour which is characteristic of the subsecond black swan regime. Even if the existence of instabilities at the millisecond scale turn out to be a short-term phenomenon while methods to adopt hardware for sophisticated trading algorithms catch up, it is essential to understand their properties. Whatever the future market-wide setup is in terms of technologies, the subsecond trading space will become increasingly crowded given the strict lower cutoff of the physical speed of light, and hence future market dynamics will undoubtedly exhibit significant herding behaviour toward certain strategy types, as mimicked by our model.

We have also presented, in outline, an approach to quantifying the risk of writing a derivative in the subsecond computer trading world. This almost surely will become a huge area in the future: Indeed, one could envisage derivative (e.g. option) expiry times eventually becoming shorter than 1 second with the increase in market speed, such that a computer may end up selling an option with expiration time 1 second into the future. It may then trade 1000 times on the millisecond scale in order to hedge that option risk. At the moment, such derivatives do not exist and expiration times tend to be pegged to human timescales (e.g. 1 month), however there is no reason in principle why they should remain this way.

Looking forward, the burning question is whether the millisecond-scale instabilities discussed here (i.e. black swans) will ultimately just be a concern for subsecond trading, or if their causal knock-on effects will actually reverberate up to longer timescales on the second, minute, hour or even daily scale. Certainly there is a strong correlation between their proliferation and the financial crisis of 2008 (Fig. 4(c)). Further work needs to be done, involving more detailed and ultimately confidential trading datasets, to pin down the causality of this relationship. Turning back to the analogy with microfractures in aircraft structures, it is clear from the hard-learned lessons resulting from aviation disasters that effects on the micro and macro timescales cannot be safely assumed to decouple – and without a clearer understanding of this coupling, imposing regulatory constraints may prove counterproductive in ways which cannot be forseen. What is clear is that subsecond instabilities are a relatively new phenomena which has increased rapidly since 2006, and the possible connections to global market stability pose a tantalizing problem which will require careful unpacking on the academic and regulatory level.

References

Arthur, W.B. (1999) *Complexity and the Economy*. Science 284, 107.

Bouchaud, J-P. and Potters, M. (2003). *Theory of Financial Risk and Derivative Pricing* (Cambridge University Press)

Campbell, J.Y., Lo, A.W., and MacKinlay, A.C. (1996). *The Econometrics of Financial Markets* (Princeton University Press)

De Luca, M., Szostek, C., Cartlidge, J., and Cliff, D. (2011). Studies on interactions between human traders and Algorithmic Trading Systems. U.K. Government Foresight Project:

http://www.bis.gov.uk/assets/bispartners/foresight/docs/computer-trading/11-1232-dr13-studies-of-interactions-between-human-traders-and-algorithmic-trading-systems

Farmer, J.D. and Skouras, S. (2011). An ecological perspective on the future of computer trading. U.K. Government Foresight Project. Available at

bis.ecgroup.net/Publications/Foresight/ComputerTradinginFinancialMarkets.aspx

Gabaix, X., Gopikrishnan, P., Plerou, V., and Stanley, H.E. (2003). A Theory of Power-Law Distributions in Financial Market Fluctuations. Nature 423, 267.

Haldane, A. (2011). The Race to Zero. Bank of England media available at http://www.bankofengland.co.uk/publications/speeches/2011/speech509.pdf

Johnson, N.F., Jefferies, P. and Hui, P.M. (2003). Financial Market Complexity (Oxford University Press)

Liukkonen, T. (2009). Human Reaction Times as a Response to Delays in Control Systems. http://www.measurepolis.fi/alma/ALMA%20Human%20Reaction%20Times%20as%20a%20Response%20to%20Delays%20in%20Control%20Systems.pdf

Mantegna, R.N. and Stanley, H.E. (1995). Scaling Behaviour in the Dynamics of an Economic Index. Nature 376, 46.

May, R. and Haldane, A. (2011). Systemic risk in banking ecosystems. Nature 469, 351.

Perez, E. (2011). The Speed Traders: An Insider's Look at the New High-Frequency Trading Phenomenon That is Transforming the Investing World (McGraw-Hill, New York)

Popular Mechanics (2011) http://www.popularmechanics.com/technology/engineering/infrastructure/a-transatlantic-cable-to-shave-5-milliseconds-off-stock-trades

Preis, T., Schneider, J.J. and Stanley, H.E. (2011). Switching Processes in Financial Markets. Proceedings of the National Academy of Sciences 108, 7674.

Saariluoma, P. (1995). Chess players' thinking: a cognitive psychological approach (Routledge, New York) p.43.

Satinover, J. and Sornette, D. (2008). Cycles, determinism and persistence in agent-based games and financial time-series. Available at http://aps.arxiv.org/abs/0805.0428

Sornette, D. (2009). Dragon-Kings, Black Swans and the Prediction of Crises. Available at

http://arxiv.org/abs/0907.4290

Taleb, N.N. (2010). The Black Swan: The Impact of the Highly Improbable (Random House Trade, New York, 2010). 2nd Edition.

Wall Street Journal (2011). http://blogs.wsj.com/marketbeat/2011/06/14/wall-streets-need-for-trading-speed-the-nanosecond-age/

Foresight
1 Victoria Street
London SW1H 0ET
www.foresight.gov.uk

URN: 12/1060