



OPEN

Abrupt rise of new machine ecology beyond human response time

SUBJECT AREAS:

SCIENTIFIC DATA
NONLINEAR PHENOMENA
STATISTICAL PHYSICS
APPLIED PHYSICS

Neil Johnson¹, Guannan Zhao¹, Eric Hunsader², Hong Qi¹, Nicholas Johnson¹, Jing Meng¹ & Brian Tivnan^{3,4}

¹Physics Department, University of Miami, Coral Gables, Florida 33124, U.S.A., ²Nanex LLC, Evanston, Illinois, U.S.A., ³The MITRE Corporation, McLean, VA 22102, U.S.A., ⁴Complex Systems Center, University of Vermont, Burlington, VT 05405, U.S.A.

Received
19 November 2012

Accepted
22 August 2013

Published
11 September 2013

Correspondence and requests for materials should be addressed to N.J. (njohnson@physics.miami.edu)

Society's techno-social systems are becoming ever faster and more computer-orientated. However, far from simply generating faster versions of existing behaviour, we show that this speed-up can generate a new behavioural regime as humans lose the ability to intervene in real time. Analyzing millisecond-scale data for the world's largest and most powerful techno-social system, the global financial market, we uncover an abrupt transition to a new all-machine phase characterized by large numbers of subsecond extreme events. The proliferation of these subsecond events shows an intriguing correlation with the onset of the system-wide financial collapse in 2008. Our findings are consistent with an emerging ecology of competitive machines featuring 'crowds' of predatory algorithms, and highlight the need for a new scientific theory of subsecond financial phenomena.

As discussed recently by Vespignani¹, humans and computers currently cohabit many modern social environments, including financial markets^{1–25}. However, the strategic advantage to a financial company of having a faster system than its competitors is driving a billion-dollar technological arms race^{6–9,16–19} to reduce communication and computational operating times down to several orders of magnitude below human response times^{26,27} -- toward the physical limits of the speed of light. For example, a new dedicated transatlantic cable¹⁸ is being built just to shave 5 milliseconds (5 ms) off transatlantic communication times between US and UK traders, while a new purpose-built chip iX-eCute is being launched which prepares trades in 740 nanoseconds¹⁹ (1 nanosecond is 10^{-9} seconds). In stark contrast, for many areas of human activity, the quickest that someone can notice potential danger and physically react, is approximately 1 second^{26,27} (1 s). Even a chess grandmaster requires approximately 650 ms just to realize that she is in trouble^{26,27} (i.e. her king is in checkmate).

In this paper we carry out a study of ultrafast extreme events (UEEs) in financial market stock prices. Our study is inspired by the seminal works of Farmer, Preis, Stanley, Easley and Cliff and co-workers^{2,3,6–9} who stressed the need to understand ultrafast market dynamics. To carry out this research, we assembled a high-throughput millisecond-resolution price stream across multiple stocks and exchanges using the NANEX NxCore software package. We uncovered an explosion of UEEs starting in 2006, just after new legislation came into force that made high frequency trading more attractive². Specifically, our resulting dataset comprises 18,520 UEEs (January 3rd 2006 to February 3rd 2011) which are also shown visually on the NANEX website at www.nanex.net. These UEEs are of interest from the basic research perspective of understanding instabilities in complex systems, as well as from the practical perspective of monitoring and regulating global markets populated by high frequency trading algorithms.

Results

We find 18,520 crashes and spikes with durations less than 1500 ms in our dataset, with examples of each given in Fig. 1A (crash) and 1B (spike). We define a crash (or spike) as an occurrence of the stock price ticking down (or up) at least ten times before ticking up (or down) and the price change exceeding 0.8% of the initial price, i.e. a fractional change of 0.008. We have checked that our main conclusions are robust to variations of these definitions. In order to have a standardized measure of the size of a UEE across stocks, we take the UEE size to be the fractional change between the price at the start of the UEE, and the price at the last tick in the sequence of price jumps in a given direction. Since both crashes and spikes are typically more than 30 standard deviations larger than the average price movement either side of an event (see Figs. 1A and 1B), they are unlikely to have arisen by chance since, in that case, their expected number would be essentially zero whereas we observe 18,520.

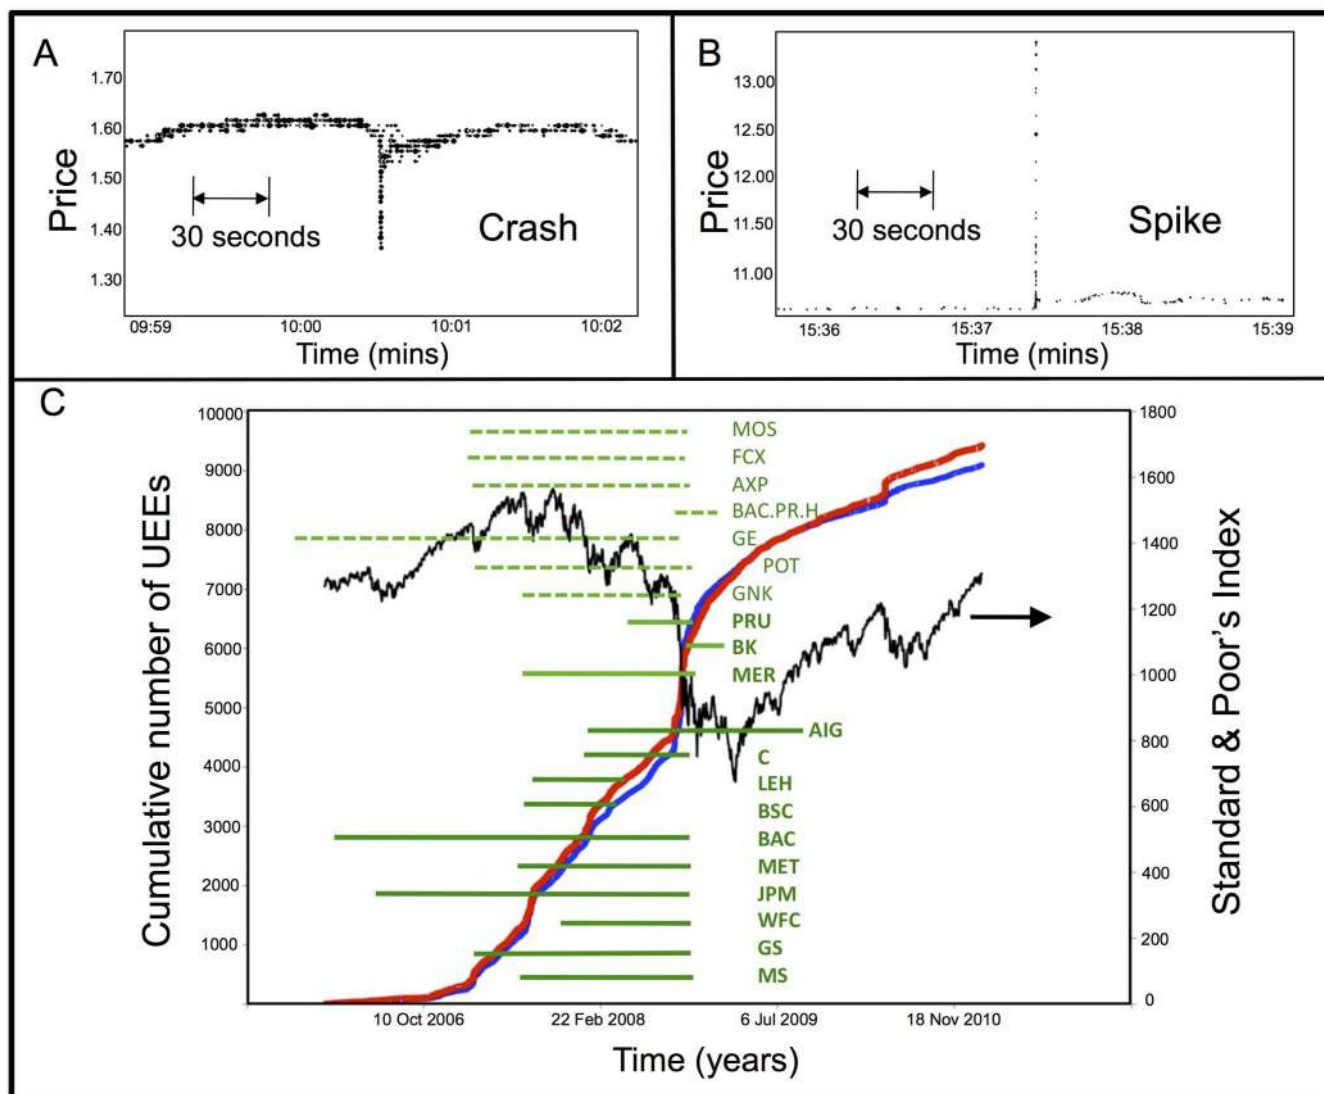


Figure 1 | Ultrafast extreme events (UEEs). (A) Crash. Stock symbol is ABK. Date is 11/04/2009. Number of sequential down ticks is 20. Price change is -0.22 . Duration is 25 ms (i.e. 0.025 seconds). The UEE duration is the time difference between the first and last tick in the sequence of jumps in a given direction. Percentage price change downwards is 14% (i.e. crash size is 0.14 expressed as a fraction). (B) Spike. Stock symbol is SMCI. Date is 10/01/2010. Number of sequential up ticks is 31. Price change is $+2.75$. Duration is 25 ms (i.e. 0.025 seconds). Percentage price change upwards is 26% (i.e. spike size is 0.26 expressed as a fraction). Dots in price chart are sized according to volume of trade. (C) Cumulative 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. Green horizontal lines show periods of escalation of UEEs. Non-financials are dashed green horizontal lines, financials are solid green. 20 most susceptible stock (i.e. most UEEs) are shown in ranked order from bottom to top, with Morgan Stanley (MS) having the most UEEs.

Figure 2 shows that as the UEE duration falls below human response times^{26,27}, the number of both crash and spike UEEs increases very rapidly. The fact that the occurrence of spikes and crashes is similar (i.e. blue and red curves almost identical in Fig. 1C and in Fig. 2) suggests UEEs are unlikely to originate from any regulatory rule that is designed to control market movements in one direction, e.g. the uptick regulatory rule for crashes^{16,17}. Their rapid subsecond speed and recovery shown in Figs. 1A and 1B suggests they are also unlikely to be driven by exogenous news arrival. We have also checked that using 'volume time' instead of clock time, does not simplify or unify their dynamics. The extensive charts at www.nanex.net, of which Figs. 1A and 1B are examples, show that the total volume traded within each UEE does not differ significantly from trading volumes during typical few-second market intervals, nor do the UEEs originate from one large but possibly mistaken trade.

The horizontal green lines in Fig. 1C show that the UEEs started appearing at different times in the past for individual stock, but then

escalated in the build-up to the 2008 global financial collapse (black curve). Moreover, these escalation periods tend to culminate around the 15 September bankruptcy filing of Lehman Brothers. Indeed, the ten stock with the most UEEs (solid green horizontal lines) are all major banks with Morgan Stanley (MS) first, followed by Goldman Sachs (GS). Figure 2 in the SI shows explicitly the escalation of UEEs in the case of Bank of America (BAC) stock. For each stock shown in Fig. 1C, the start and end times of the escalation period (i.e. horizontal green line) are determined by examining the local trend in the arrival rate of the UEEs. In determining these start and end times, we checked various statistical methods such as LOWESS and found them all to give very similar escalation periods to those shown in Fig. 1C. We also find that the occurrence of UEEs is not simply related to the daily volatility, price or volume (see SI Fig. 2 for the explicit case of BAC). Figure 1C therefore suggests that there may indeed be a degree of causality between propagating cascades of UEEs and subsequent global instability, despite the huge difference




Figure 2 | Number of UEEs as a function of UEE duration. The UEE duration is the time difference between the first and last tick in the sequence of jumps in a given direction. UEE crashes are shown as red curve, UEE spikes as blue curve. Since the clock time between ticks varies, two UEEs having the same number of ticks do not generally have the same durations.

in their respective timescales. Although access to confidential trade and exchange information is needed to fully test this hypothesis, at the very least Fig. 1C demonstrates a coupling between extreme market behaviours below the human response time and slower global instabilities^{2,5} above it, and shows how machine and human worlds can become entwined across timescales from milliseconds to months. We have also found that UEEs build up around smaller global instabilities such as the 5/6/10 Flash Crash: although fast on the daily scale, Flash Crashes are fundamentally different to UEEs in that Flash Crashes typically last many minutes ($\gg 1$ s) and hence allow ample time for human involvement. Future work will explore the connection to existing studies such as Ref. 28 of market dynamics immediately before and after financial shocks.

Having established that the number of UEEs increases dramatically as the timescale drops below one second, and hence drops below the human reaction time, we now seek to investigate how the character of the UEEs might also change as the timescale drops – and in particular, whether the distribution may become more or less akin to a power-law distribution. Power-law distributions are ubiquitous in real-world complex systems and are known to provide a reasonable description for the distribution of stock returns for a given time increment, from minutes up to weeks^{13–15}. Our statistical procedure to test a power-law hypothesis for the distribution of UEE sizes, and hence obtain best-fit power-law parameter values, follows Clauset et al.’s²⁹ state-of-the-art methodology for obtaining best-fit parameters for power-law distributions, and for testing the power-law distribution hypothesis on a given dataset. Following this procedure, we obtain a best estimate of the power-law exponent α , and a p -value for the goodness-of-fit, for the distribution of UEE sizes. Specific details of the implementation, including a step-by-step recipe and documented programs in a variety of computer languages, are given in Ref. 29.

Figure 3 shows a plot of the goodness-of-fit p -value, and the corresponding power-law exponent α , for the distribution of sizes of UEEs having durations above a particular threshold. As this duration threshold decreases, the character of the UEE size distribution exhibits a transition from a power-law above the limit of the human response time to a non-power-law below it – specifically, the goodness-of-fit p falls from near unity to below 0.1 within a small timescale range in Figs. 3B and 3C. This loss of power-law character at subsecond timescales suggests that a lower limit needs to be placed on the validity of Mandelbrot’s claim that price-changes exhibit approximate self-similarity (i.e. approximate fractal behavior and hence power-law distribution) across all timescales³⁰. It can be seen that the transition for crashes is smoother than for spikes: this may be because many market participants are typically ‘long’ the market¹⁶ and hence respond to damaging crashes differently from profitable spikes. Not only is the crash transition onset (650 ms) earlier in Fig. 3B than for spikes in Fig. 3C, it surprisingly is the same as the thinking time of a chess grandmaster, even though individual traders are not likely to be as attentive or quick as a chess grandmaster on a daily basis^{26,27}. This may be a global online manifestation of the ‘many eyes’ principle from ecology⁶ whereby larger groups of animals or fish may detect imminent danger more rapidly than individuals.

Figures 4 and 5 show further evidence for this transition in UEE size character as timescales drop below human response times. Figure 4 shows that the cumulative distribution of UEE sizes for the example of spikes, exhibits a qualitative difference between UEEs of duration greater than 1 second, where $p = 0.91$ and hence there is strong support for a power-law distribution, and those less than 1 second where $p < 0.05$ and hence a power-law can be rejected. A similar conclusion holds for crashes. Figure 5 shows the cumulative distribution of sizes for UEEs in different duration windows, with the distribution for the duration window 1200–1500 ms showing a marked change from the trend at lower window values. The following quantities that we investigated, also confirm a change in UEE character in this same transition regime: (1) a Kolmogorov-Smirnov two-sample test to check the similarity of the different UEE size distributions within different duration time-windows (see SI Fig. 5); (2) the standard deviation of the size of UEEs in a given window of duration (see SI Fig. 6); (3) the average and standard deviation in the number of price ticks making up the individual UEEs which lie in a given duration window (see SI Fig. 7); (4) a test for a lognormal distribution for UEE durations (see SI Fig. 8). Figure 9 of the SI confirms that using different binnings for the UEE durations does not change our main conclusions.

Discussion

Inspired by Farmer and Skouras’ ecological perspective⁶, we analyze our findings in terms of a competitive population of adaptive trading agents. The model is summarized schematically in Fig. 1 of the SI while Refs. 31, 40 and 41 provide full details and derivations of the quoted results below. Each agent possesses several ($s > 1$) strategies, but only trades at a given timestep if it has a strategy that has performed sufficiently well in the recent past. The common information fed back to the agents at each timestep is a bit-string encoding the m most recent price movements^{31–35}. The key quantity is $\eta = 2^{m+1}/N$ corresponding to the ratio of the number of different strategies (i.e. strategy pool size which is 2^{m+1} in our model^{31,33}) to the number of active agents N . For $\eta > 1$, there are more strategies than agents, which is consistent with having many human participants since individual humans have myriad ways of making decisions, including arbitrary guesswork, hunches and personal biases. Hence $\eta > 1$ is consistent with having many active human traders, which in turn is consistent with longer timescales (> 1 s) since this is where humans can think and act. The $\eta > 1$ output, illustrated in Fig. 6B (right panel), does indeed reproduce many well-known features of longer timescale price increments³¹. The chance that many

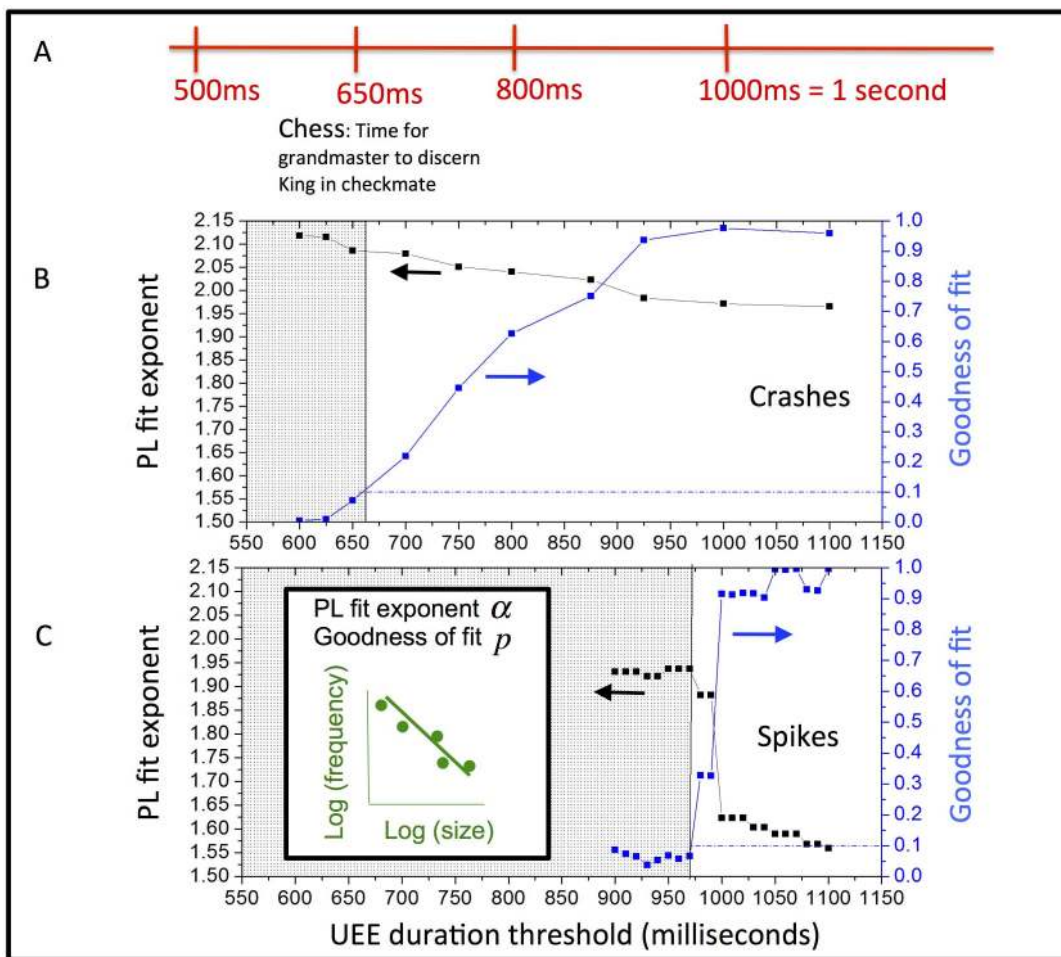


Figure 3 | Empirical transition in size distribution for UEEs with duration above threshold τ , as function of τ . (A) Scale of times. 650 ms is the time for chess grandmaster to discern King is in checkmate. Plots show results of the best-fit power-law exponent (black) and goodness-of-fit (blue) to the distributions for size of (B) crashes, and (C) spikes, as shown in the inset schematic.

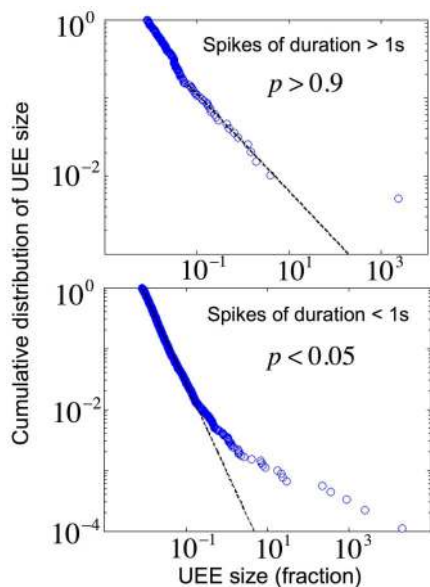


Figure 4 | Extent to which the cumulative distribution for UEE spikes follows a power-law, for the subset having durations greater than 1 second (upper panel) and less than 1 second (lower panel). For durations more than 1 second, there is strong evidence for a power-law (p -value is 0.912). For durations less than 1 second, a power-law can be rejected. Black line shows best-fit power-law.




Figure 5 | Cumulative distribution for UEE spikes with durations within a given millisecond range, having a size which is at least as big as the value shown on the horizontal axis.

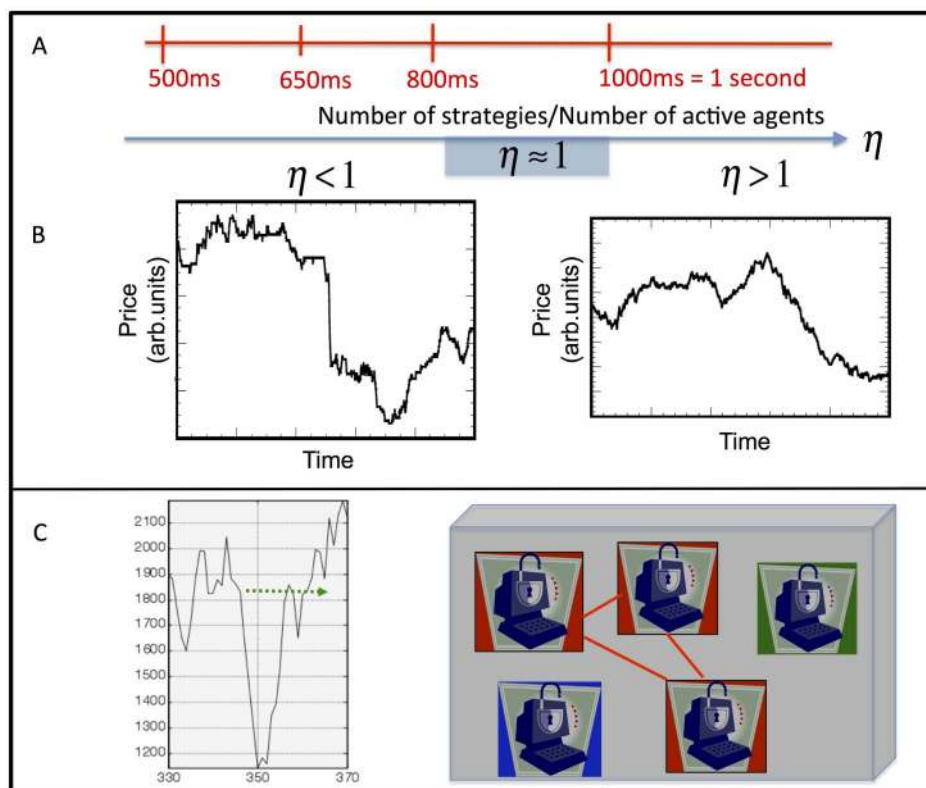


Figure 6 | Theoretical transition. (A) Timescales from Fig. 3A. (B) Our model's price output for the two regimes, using same vertical price scale. η is ratio of number of strategies to number of agents ($\eta = 2^{m+1}/N$). $\eta < 1$ implies more agents than strategies, hence frequent, large and abrupt price-changes as observed empirically for timescales < 1 s. $\eta > 1$ implies less agents than strategies, hence large changes are rare. (C) Large change with recovery from our model, similar to Fig. 1A on expanded timescale. Right panel shows schematic of our model: Machines in $\eta < 1$ regime unintentionally use same red strategy and hence form a crowd. Adding agents with different strategies (blue and green, schematic) prevents UEE (green dashed line indicates modified price trajectory).

agents simultaneously use the same strategy and submit the same buy or sell order, is small if $\eta > 1$, hence there are very few extreme price-changes -- exactly as observed in our data for > 1 s. Reducing η below 1 corresponds to reducing the strategy pool size below the number of agents, which is consistent with a market dominated by specific high-frequency trading algorithms. As the trading timescale moves into the subsecond regime, the number of pieces of information that can be processed by a machine decreases since each piece of information requires a finite time for manipulation (e.g. storage and recall), which is consistent with a reduction in m and hence a decrease in η since $\eta = 2^{m+1}/N$.

Remarkably, decreasing η continually in our model generates a visually abrupt transition in the output with frequent extreme price-changes now appearing (Fig. 6B, left panel), which is exactly what we observed in the data for < 1 s. $\eta < 1$ implies more than one agent per strategy on average: crowds of agents frequently converge on the same strategy and hence simultaneously flood the market with the same type of order, thereby generating the frequent extreme price-change events. Although it is quite possible that there are other models that could reproduce a *gradual* change in the instability as η decreases, the task of reproducing a visually abrupt transition as observed empirically in Fig. 3 (particularly Fig. 3C) is far harder. In addition, our model predicts (1) that the extreme event size-distribution in the ultrafast regime ($\eta < 1$) should not have a power law, exactly as we observe; (2) that recoveries as in Figs. 1A and 1B, can emerge endogenously in the regime $\eta < 1$ (see Fig. 6C, left panel), again as we observe; and (3) that extreme events can be diverted by momentarily increasing the strategy diversity. To achieve this latter effect, agents simply need to be added with complementary strategies -- shown as complementary colors in the right panel of Fig. 6C -- thereby partially

cancelling the machine crowd denoted in red. The fact that the actual model price trajectory can then bypass the potential extreme event (green dashed line in left panel of Fig. 6C) therefore offers hope of using small real-time interventions to mitigate systemic risk.

Although the simplicity of our proposed minimal model necessarily ignores many market details, it allows us to derive explicit analytic formulae for the scale of the fluctuations in each phase, and hence an indication of the risk, if we make the simplifying assumption that the number of agents trading each timestep is approximately N (see Refs. 31, 40 and 41 for details). For $\eta > 1$, the scale is given by $N^{\frac{1}{2}}(1 - 2^{-(m+1)}N)^{\frac{1}{2}}$ for general s , while for $\eta < 1$ it abruptly adopts a new form with upper bound $3^{-\frac{1}{2}}2^{-\frac{m}{2}}N(1 - 2^{-2(m+1)})^{\frac{1}{2}}$ and lower bound $3^{-\frac{1}{2}}2^{-\frac{(m+1)}{2}}N(1 - 2^{-2(m+1)})^{\frac{1}{2}}$ for $s = 2$. This predicted sudden increase in the fluctuation scale from being proportional to $N^{\frac{1}{2}}$ for $\eta > 1$, to proportional to N for $\eta < 1$, is consistent with the observed appearance of frequent UEEs at short timescales, and specifically the visually abrupt transition that we observe in Fig. 3.

More detailed investigation of the properties of UEEs, and the potential implications for financial market instability, will require access to confidential exchange data that was not available in the present study. However a remarkable new study by Cliff and Cartledge³⁶ provides some additional support for our findings. In controlled lab experiments, they found that when machines operate on similar timescales to humans³⁶ (longer than 1 s), the 'lab market' exhibited an efficient phase (c.f. few extreme price-change events in our case). By contrast, when the machines operated on a timescale faster than the human response time³⁶ (100 milliseconds) then the market exhibited an inefficient phase (c.f. many extreme price-change events in our case).



While our crowd model offers a plausible explanation of the observed transition in Fig. 3, we stress that our purpose in this paper was not to explain the details of the price changes during individual UEEs, nor was it to unravel the underlying market microstructure that might provoke or exacerbate such UEEs. A recent preprint by Golub et al.³⁷ claims that a majority of all UEEs carry the label of ISO (Inter-market Sweep Order³⁷). However this claim does not affect the validity of our findings. Moreover, Ref. 37 does not uncover or explain the visually abrupt transition that we observe in Fig. 3, nor does it invalidate our own crowd model explanation. Irrespective of the underlying order identities, every UEE is the result of a sudden excess buy or sell demand in the market, and our model provides a simple explanation for how sudden excess buy or sell demands are generated, not how they get fulfilled. Indeed it is a common feature of our model output that a large imbalance of buy or sell demand can suddenly appear, producing a UEE as observed empirically. We also note that Golub et al.³⁷ make several strong assumptions in their attempts to label the UEEs, each of which requires more detailed investigation since the resulting identifications are neither unique nor unequivocal. Whether the visually abrupt transition in Fig. 3 is a strict phase transition in the statistical physics sense, also does not affect the validity of our results. The extent to which UEEs were provoked by regulatory and institutional changes around 2006, is a fascinating question whose answer depends on a deeper understanding of the market microstructure along the lines started by Golub et al.³⁷. It may be that ISOs are particularly problematic, but this is still unclear because of the assumptions made in Ref. 37. Once this has been resolved, it should be possible to make definite policy recommendations based on our findings, as well as expanding the study to connect to systemic risk³⁸ and derivative operations³⁹.

Methods

The power-law analysis that we use to obtain our main result in Fig. 3, follows the state-of-the-art testing procedure laid out by Clauset et al.²⁹. Our accompanying crowd model considers a simple yet archetypal model of a complex system based on a population of agents competing for a limited resource with bounded rationality. This model has previously been shown to reproduce the main stylized facts of financial markets³¹. Its dynamics are based on the realistic notion that it is better to be a buyer when there is an excess of sellers or vice versa when in a financial market comprising agents (humans or machines) with short-term, high-frequency trading goals. The formulae given above for the scale of the fluctuations in each phase, are derived explicitly in Ref. 40, and also Refs. 31 and 41.

- Vespignani, A. Predicting the behaviour of techno-social systems. *Science* **325**, 425–428 (2009).
- Easley, D., Lopez de Prado, M. & O'Hara, M. The volume clock: Insights into the high-frequency paradigm. *Journal of Portfolio Management* **39**, 19–29 (2012).
- Easley, D., Lopez de Prado, M. & O'Hara, M. Flow toxicity and liquidity in a high-frequency world. *Review of Financial Studies* **25**, 1457–1493 (2012).
- Preis, T., Kenett, D. Y., Stanley, H. E., Helbing, D. & Ben-Jacob, E. Quantifying the behavior of stock correlations under market stress. *Sci. Rep.* **2**, 752 (2012).
- Rime, D., Sarno, L. & Sojli, E. Exchange rate forecasting, order flow and macroeconomic information. *Journal of International Economics* **80**, 72–88 (2010).
- Farmer, J. D. & Skouras, S. *An Ecological Perspective on the Future of Computer Trading*. U.K. Government Foresight Project (2011). Available at <http://www.bis.gov.uk/assets/foresight/docs/computer-trading/11-1225-dr6-ecological-perspective-on-future-of-computer-trading>. Last visited: 17-2-2013.
- Preis, T., Schneider, J. J. & Stanley, H. E. Switching processes in financial markets. *Proceedings of the National Academy of Sciences* **108**, 7674–7678 (2011).
- Haldane, A. The race to zero. Speech given at: *International Economic Association Sixteenth World Congress*, Beijing, China, 8 July 2011. Available at <http://www.bankofengland.co.uk/publications/speeches/2011/speech509.pdf>. Last visited: 17-2-2013.
- De Luca, M., Szostek, C., Cartlidge, J. & Cliff, D. *Studies on Interactions between Human Traders and Algorithmic Trading Systems*. U.K. Government Foresight Project (2011). Available at <http://www.bis.gov.uk/assets/bispartners/foresight/docs/computer-trading/11-1232-dr13-studies-of-interactions-between-human-traders-and-algorithmic-trading-systems>. Last visited: 17-2-2013.
- Schweitzer, F. et al. Economic networks: The new challenges. *Science* **325**, 422–425 (2009).
- Caldarelli, G., Battiston, S., Garlaschelli, D. & Catanzaro, M. Emergence of complexity in financial networks. *Lecture Notes in Physics* **650**, 399–423 (2004).

- Oltvai, Z. N. & Barabási, A. L. Life's complexity pyramid. *Science* **298**, 763–764 (2002).
- Mantegna, R. N. & Stanley, H. E. Scaling behavior in the dynamics of an economic index. *Nature* **376**, 46–49 (1995).
- Gabaix, X., Gopikrishnan, P., Plerou, V. & Stanley, H. E. A theory of power-law distributions in financial market fluctuations. *Nature* **423**, 267–270 (2003).
- Lux, T. & Marchesi, M. Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature* **397**, 498–500 (1999).
- Perez, E. *The Speed Traders: An Insider's Look at the New High-Frequency Trading Phenomenon That is Transforming the Investing World* (McGraw-Hill, New York, 2011).
- U.S. Securities and Exchange Commission. *Findings Regarding the Market Events of May 6, 2010*. Available at <http://www.sec.gov/news/studies/2010/marketevents-report.pdf>. Last visited: 17-2-2013.
- Pappalardo, J. New transatlantic cable built to shave 5 milliseconds off stock traders (2011). Available at: <http://www.popularmechanics.com/technology/engineering/infrastructure/a-transatlantic-cable-to-shave-5-milliseconds-off-stock-trades>. Last visited: 17-2-2013.
- Conway, B. Wall Street's need for trading speed: The nanosecond age. *The Wall Street Journal*. JUN 14, (2011). Available at: <http://blogs.wsj.com/marketbeat/2011/06/14/wall-streets-need-for-trading-speed-the-nanosecond-age/>. Last visited: 17-2-2013.
- Taleb, N. N. *The Black Swan: The Impact of the Highly Improbable* (Random House Trade, New York, 2010). 2nd Edition.
- Sornette, D. Dragon-kings, black swans and the prediction of crises. Available at <http://arxiv.org/abs/0907.4290> (2009).
- Johansen, A. & Sornette, D. Large stock market price drawdowns are outliers. *Journal of Risk* **4**, 69–110 (2001).
- Helbing, D., Farkas, I. & Vicsek, T. Simulating dynamical features of escape panic. *Nature* **407**, 487–490 (2000).
- Bouchaud, J. P. & Potters, M. *Theory of Financial Risk and Derivative Pricing* (Cambridge University Press, 2003).
- Cont, R. *Frontiers in Quantitative Finance: Credit Risk and Volatility Modeling* (Wiley, New York, 2008).
- Liukkonen, T. N. & Unit, K. Human reaction times as a response to delays in control systems – Notes in vehicular context (2009). Available at <http://www.measurepolis.fi/alma/ALMA%20Human%20Reaction%20Times%20as%20a%20Response%20to%20Delays%20in%20Control%20Systems.pdf>. Last visited: 17-2-2013.
- Saariluoma, P. *Chess Players' Thinking: A Cognitive Psychological Approach* (Routledge, New York, 1995), p. 43.
- Petersen, A. M., Wang, F., Havlin, S. & Stanley, H. E. Market dynamics immediately before and after financial shocks: Quantifying the Omori, productivity, and Bath laws. *Physical Review E* **82**, 036114–036125 (2010).
- Clauset, A., Shalizi, C. R. & Newman, M. E. Power-law distributions in empirical data. *SIAM Rev.* **51**, 661–703 (2009).
- Mandelbrot, B. *The Misbehaviour of Markets: A Fractal View of Market Turbulence* (Basic Books, New York, 2004).
- Johnson, N. F., Jefferies, P. & Hui, P. M. *Financial Market Complexity* (Oxford University Press, 2003), Chap. 4.
- Arthur, W. B. Complexity and the economy. *Science* **284**, 107–109 (1999).
- Challet, D., Marsili, M. & Zhang, Y. C. *Minority Games: Interacting Agents in Financial Markets* (Oxford University Press, Oxford, 2005).
- Johnson, N. F., Smith, D. M. D. & Hui, P. M. Multi-Agent Complex Systems and Many-Body Physics. *Europhysics Letters* **74**, 923–927 (2006).
- Satinover, J. & Sornette, D. Cycles, determinism and persistence in agent-based games and financial time-series (2008). Available at <http://arxiv.org/abs/0805.0428>. Last visited: 17-2-2013.
- Cartlidge, J., Szostek, C., Luca, M. D. & Cliff, D. Too fast - too furious: Faster financial – market trading agents can give less efficient markets. In *4th Int. Conf. Agents & Art. Intell* **2**, 126–135 (2012). Available at <http://www.cs.bris.ac.uk/~cszjpc/docs/cartlidge-icaart-2012.pdf>. Last visited: 17-2-2013.
- Golub, A., Keane, J. & Poon, S. High Frequency Trading and Mini Flash Crashes, e-print arXiv:1211.6667v1 (2013) at <http://arxiv.org/abs/1211.6667>. Last visited: 20-5-2013.
- Battiston, S., Puliga, M., Kaushik, R., Tasca, P. & Caldarelli, G. DebtRank: Too Central to Fail? Financial Networks, the FED and Systemic Risk. *Scientific Reports* **2**, 541–545 (2012).
- Battiston, S., Caldarelli, G., Georg, C. P., May, R. & Stiglitz, J. Complex Derivatives. *Nature Physics* **9**, 123 (2013).
- Johnson, N. F. & Hui, P. M. Crowd-Anticrowd Theory of Collective Dynamics in Competitive, Multi-Agent Populations and Networks. E-print arXiv:cond-mat/0306516v1 available at <http://arxiv.org/abs/cond-mat/0306516>. Last accessed 07-31-2013.
- Smith, D. M. D., Hui, P. M. & Johnson, N. F. Multi-Agent Complex Systems and Many-Body Physics. *Europhysics Letters* **74**, 923–927 (2006).

Acknowledgments

NJ (Neil Johnson) gratefully acknowledges support for this research from The MITRE Corporation and the Office of Naval Research (ONR) under grant N000141110451. The



views and conclusions contained in this paper are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the above named organizations, to include the U.S. government. We thank Amith Ravindar, Joel Malerba, Zhenyuan Zhao, Pak Ming Hui, Spencer Carran, David Smith, Michael Hart and Paul Jefferies for discussions surrounding this topic and help with assembling datafiles and parts of figures.

Author contributions

All authors participated in discussions of the research, its findings, and the content of the manuscript. NJ (Neil Johnson), GZ and BT wrote the manuscript. NJ (Neil Johnson), EH and BT designed the research. NJ (Neil Johnson) GZ, EH, HQ, NJ and BT analysed the empirical data. GZ, HQ and JM did the numerical calculations while NJ (Neil Johnson) completed the analytical derivations.

Additional information

Supplementary information accompanies this paper at <http://www.nature.com/scientificreports>

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Johnson, N. *et al.* Abrupt rise of new machine ecology beyond human response time. *Sci. Rep.* 3, 2627; DOI:10.1038/srep02627 (2013).



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported license. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/3.0>