

Anatomy of extreme events in a complex adaptive system

Paul Jefferies¹, David Lamper² and Neil F. Johnson¹

¹Physics Department, Oxford University, Oxford, OX1 3PU, U.K.

²Oxford Centre for Industrial and Applied Mathematics,
Oxford University, Oxford, OX1 3LB, U.K.

February 5, 2002

Abstract

We provide an analytic, microscopic analysis of extreme events in an adaptive population comprising competing agents (e.g. species, cells, traders, data-packets). Such large changes tend to dictate the long-term dynamical behaviour of many real-world systems in both the natural and social sciences. Our results reveal a taxonomy of extreme events, and provide a microscopic understanding as to their build-up and likely duration.

PACS: 87.23.Ge, 02.50.Le, 05.45.Tp, 87.23.Kg

Large unexpected changes or ‘extreme events’ (e.g. crashes in financial markets, or punctuated equilibria in evolution) happen infrequently, yet tend to dictate the long-term dynamical behaviour of real-world systems in disciplines as diverse as biology and economics, through to ecology and evolution. The ability to generate large internal, so-called *endogenous* changes is a defining characteristic of complex systems, and arguably of Nature and Life itself [1, 2, 3]. Such changes are manifestations of subtle, short-term temporal correlations resulting from internal collective behaviour. They seem to appear out of nowhere and have long-lasting consequences. To what extent can they ever be ‘understood’? Followers of the self-organized criticality view [3] would claim this question is naïve because of an inherent self-similarity in Nature: any large changes are simply magnified versions of smaller changes, which are in turn magnified versions of even smaller changes, and so on. Such self-similarity is presumed to underlie the power-law scaling observed in natural, social and economic phenomena [3]. However, there are reasons for believing that the largest changes may be ‘special’ in a microscopic sense [2]. Power-law scaling is only approximately true, and does not apply over an infinite range of scales. Apart from being atomistic at the smallest scale, a population of competing agents cannot cause any effect larger than the population size itself: in short, the largest changes will tend to ‘scrape the barrel’ in some way. Reference [2] quotes Bacon from *Novum Organum*: “Whoever knows the ways of Nature will more easily notice her deviations; and, on the other hand, whoever knows her deviations will more accurately describe her ways”.

This paper addresses the task of understanding, and eventually controlling, the large endogenous changes arising in a complex adaptive system comprising competing agents (e.g. species, cells, traders, data-packets). Our work reveals a taxonomy of large changes, and provides a quantitative microscopic description of their build-up and duration. Our results also provide insight into how a ‘complex systems manager’ might contain or control such extreme events.

We consider a generic complex system in which a population of N_{tot} heterogeneous agents with limited capabilities and information, repeatedly compete for a limited global resource. Our model was introduced in Ref. [4], and is a generalization of the El Farol bar problem and the Minority game, concerning a population of people deciding whether to attend a popular bar with limited seating [5]. At timestep t , each agent (e.g. a bar customer, or a market trader) decides whether to enter a game where the choices are option 1 (e.g. attend the bar, or buy) and option 0 (e.g. go home, or sell): N_0 agents choose 0 while N_1 choose 1. The ‘excess demand’ $D[t] = N_1 - N_0$ (which mimics price-change in a market) and number $V[t] = N_1 + N_0$ of active agents (which mimics volume of market orders) represent output variables. These two quantities fluctuate with time, and can be combined to construct other global quantities of interest for the complex system studied (e.g. summing the price-changes gives the current price). This model can reproduce statistical and dynamical features similar to those in a real-world complex adaptive system, namely a financial market [4], and exhibits the crucial feature of seemingly spontaneous large changes of variable duration [4, 6]. The resulting time-series appears ‘random’ yet is non-Markovian,

with subtle temporal correlations which put it beyond any random-walk based description. The temporal correlations of price-changes and volume, and their cross-correlation, are of intense interest in financial markets where so-called chartists offer a wide range of rules-of-thumb [7] such as ‘volume goes with price trend’. Although such rules are unreliable, the intriguing question remains as to whether there could *in principle* be a ‘science of charting’.

A subset $V[t] \leq N_{tot}$ of the population, who are sufficiently confident of winning, are active at each timestep. For $N_1 < N_0$ the winning decision is 1 and vice-versa, i.e. the winning decision is given by $H[-D[t]]$ where $H[x]$ is the Heaviside function. The global resource level is so limited, or equivalently the game is so competitive, that at least half the active population lose at each timestep [5]. The only global information available to the agents is a common bit-string ‘memory’ of the m most recent outcomes. Consider $m = 2$; the $P = 2^m = 4$ possible history bit-strings are 00, 01, 10 and 11, which can also be represented in decimal form: $\mu \in \{0, 1, \dots, P-1\}$. A strategy consists of a response, $a^\mu \in \{-1, 1\}$ to each possible bit-string μ , $a^\mu = 1 \Rightarrow$ option 1, and $a^\mu = -1 \Rightarrow$ option 0. Hence there are $2^P = 16$ possible strategies. The heterogeneous agents randomly pick s strategies each at the outset, and update the scores of their strategies after each timestep with the reward function $\chi[D] = \text{sgn}[-D]$, i.e. +1 for choosing the minority action, -1 for choosing the majority action. Agents have a time horizon T over which strategy points are collected, and a threshold level r which mimics a ‘confidence’. Only strategies having $\geq r$ points are used, with agents playing their highest scoring strategy. Agents with no such strategy become temporarily inactive [4]. We focus on the regime where the number of strategies in play is comparable to the total number available, since this yields seemingly random dynamics with occasional large movements [4, 6]. The coin-tosses used to resolve ties in decisions (i.e. $N_0 = N_1$) and active-strategy scores, inject stochasticity into the game’s evolution. Reference [8] showed that a simplified version of this system in the limit $r \rightarrow -\infty$ and $T \rightarrow \infty$, can be usefully described as a stochastically disturbed deterministic system. We are interested in the dynamics of large changes, and adopt the approach and terminology of Ref. [8]. Averaging over our model’s stochasticity yields a description of the game’s deterministic dynamics via mapping equations for the strategy score vector $\underline{S}[t]$ and global information $\mu[t]$. For $s = 2$ the deterministic dynamics are given exactly by the following equations:

$$\begin{aligned} \underline{S}[t] &= \underline{S}[0] - \sum_{i=t-T}^{t-1} \underline{a}^{\mu[i]} \text{sgn}[D[i]], \\ \mu[t] &= 2\mu[t-1] - PH[\mu[t-1] - P/2] + H[D[t-1]]. \end{aligned} \quad (1)$$

The corresponding demand function is given by

$$D[t] = \sum_{R=1}^{2P} a_R^{\mu[t]} H[S_R - r] \sum_{R'=1}^{2P} (1 + \text{sgn}[S_R[t] - S_{R'}[t]]) \Psi_{R,R'},$$

where $\underline{\Psi}$ is the symmetrized strategy allocation matrix which constitutes the

quenched disorder present during the system's evolution [8]. Elements $\Psi_{R,R'}$ enumerate the number of agents holding both strategy R and R' . The volume $V[t]$ is given by the same expression as $D[t]$ replacing $a_R^{\mu[t]}$ by unity.

Large changes such as financial market crashes, seem to exhibit a wide range of possible durations and magnitudes making them difficult to capture using traditional statistical techniques based on one or two-point probability distributions [2]. A common feature, however, is an obvious trend (i.e. to the eye) in one direction over a reasonably short time window: we use this as a working definition of a large change. In fact, all the large changes discussed here represent $> 3\sigma$ events. In both our model and the real-world system, these large changes arise more frequently than would be expected from a random-walk model [1, 2]. Our model's dynamics can be described by trajectories on a de Bruijn graph [8]: see Fig. 1 for $m = 3$, with a transition incurring an increment to the score vector \underline{S} . There are P orthogonal increment vectors \underline{a}^μ , one for each node μ . Setting the initial scores $\underline{S}[0] = \underline{0}$, the strategy score vector in Eq. (1) can be written exactly as:

$$\underline{S}[t] = c_0 \underline{a}^0 + c_1 \underline{a}^1 + \dots + c_{P-1} \underline{a}^{P-1} = \sum_{j=0}^{P-1} c_j \underline{a}^j$$

where c_j represents the *nodal weights* for history node $\mu = j$. The nodal weights enumerate the number of negative return transitions from node μ minus the number of positive return transitions, in the time window $t - T \rightarrow t - 1$. High absolute nodal weight implies persistence in transitions from that node i.e. persistence in $D|\mu$. Large changes will occur when connected nodes become persistent. The simplest type of large movement exhibiting perfect nodal persistence would be $\mu = 0, 0, 0, 0, \dots$ in which all successive price changes are in the *same* direction. We call this a ‘fixed-node crash’ (or rally). However, there are many other possibilities reflecting the wide range of forms and durations of the large change. For example, on the $m = 3$ de Bruijn graph in Fig. 1 the cycle $\mu = 0, 0, 1, 2, 4, 0, \dots$ has four out of the five transitions producing price-changes of the same sign (it is persistent on nodes 1, 2, 4 and antipersistent on node 0). We call this a ‘cyclic-node crash’ (or rally). Figure 2 illustrates a large change which starts as a fixed-node crash then subsequently becomes a cyclic-node crash. Cyclic-node crashes can be treated simply as interlocking fixed-node crashes, hence for clarity we focus here on a single fixed-node crash (or rally). For the parameter ranges of interest, the choice about whether a strategy is played by an agent is more determined by whether that strategy's score is above the threshold, than whether it is their highest-scoring strategy [9]. This is because agents are only likely to have at most one strategy whose score lies above the threshold for confidence levels $r \geq 0$. Making the additional numerically-justified approximation of small quenched disorder (i.e. the variance of the entries in the strategy allocation matrix $\underline{\Psi}$ is smaller than their

mean for the parameter range of interest [8]), the demand and volume become:

$$D[t] = \frac{N}{4P} \sum_{R=1}^{2P} a_R^{\mu[t]} \operatorname{sgn}[S_R[t] - r], \quad (2)$$

$$V[t] = \frac{N}{2} + \frac{N}{4P} \sum_{R=1}^{2P} \operatorname{sgn}[S_R[t] - r]. \quad (3)$$

Suppose persistence on node $\mu = 0$ starts at time t_0 . How long will the resulting crash last? To answer this, we decompose Eq. (2) into strategies which predict 1 at $\mu = 0$, and those that predict 0. We first consider the particular case where the node $\mu = 0$ was *not* visited during the previous T timesteps, hence the loss of score increment from time-step $t - T$ will not affect $\underline{S}[t]$ on average. At any later time $t_0 + \tau$ during the crash, (i.e. $\mu = 0$) Eqs. (2) and (3) are hence given by:

$$D[t_0 + \tau] = -\frac{N}{4P} \left\{ \sum_{R \ni a_R^\mu = -1} \operatorname{sgn}[S_R[t_0] - r - \tau] - \sum_{R \ni a_R^\mu = 1} \operatorname{sgn}[S_R[t_0] - r + \tau] \right\}, \quad (4)$$

$$V[t_0 + \tau] = \frac{N}{2} + \frac{N}{4P} \left\{ \sum_{R \ni a_R^\mu = -1} \operatorname{sgn}[S_R[t_0] - r - \tau] + \sum_{R \ni a_R^\mu = 1} \operatorname{sgn}[S_R[t_0] - r + \tau] \right\}.$$

$|D[t_0 + \tau]|$ decreases as the persistence time τ increases, and hence the crash ends at time $t_0 + \tau_c$ when the right-hand side of Eq. (4) changes sign. The persistence time or ‘crash-length’ τ_c is thus given by the mean of the scores of the strategies predicting 0, i.e. $\tau_c = \overline{S}_{R \ni a_R^\mu = -1}[t_0] = -c_0[t_0]$. In the more general case where the node $\mu = 0$ *was* visited during the previous T timesteps, τ_c is given by the largest τ value which satisfies:

$$\tau = - \left(c_0[t_0] + \sum_{\{t'\}} \operatorname{sgn}[D[t']] \right)$$

where $\{t'\} \ni (\mu[t'] = 0 \cap t_0 - T \leq t' \leq t_0 + \tau - T)$. Assume that the scores have a near-Normal distribution, i.e. $S_{R \ni a_R^\mu = -1}[t_0] \sim \mathcal{N}[\overline{S}_{-1}, \sigma]$ as in Fig. 3a. For each strategy R there exists an anticorrelated strategy \overline{R} and hence $S_R[t] = -S_{\overline{R}}[t]$ for all t . Consequently, prior to a crash, the score distribution tends to split into two halves as indicated schematically in Fig. 3a. The expected demand (and volume) during the crash are then:

$$\begin{aligned} \langle D[t_0 + \tau] \rangle &\propto \left(\operatorname{erf} \left[\frac{c_0[t_0] + r + \tau}{\sqrt{2}\sigma} \right] - \operatorname{erf} \left[\frac{-c_0[t_0] + r - \tau}{\sqrt{2}\sigma} \right] \right) \\ \langle V[t_0 + \tau] \rangle &\propto \left(2 - \operatorname{erf} \left[\frac{c_0[t_0] + r + \tau}{\sqrt{2}\sigma} \right] - \operatorname{erf} \left[\frac{-c_0[t_0] + r - \tau}{\sqrt{2}\sigma} \right] \right) \end{aligned}$$

These forms are illustrated in Figure 3b. As the spread in the strategy score distribution is increased, the dependence of $\langle D \rangle$ and $\langle V \rangle$ on the parameters τ and r becomes weaker and the surfaces flatten out leading to a smoother drawdown, as opposed to a sudden severe crash. As the parameters \bar{S}_{-1}, σ, r are varied, it can be seen that the behaviour of the demand and volume during the crash can exhibit markedly different qualitative forms [9], yielding a *taxonomy of different species of large change even within the same single-node family*. This result could explain why financial market chartists' rules-of-thumb [7], such as 'volume goes with price trend', are far too simplistic.

We now turn to the important practical question of whether history will repeat itself, i.e. given that a crash has recently happened, is it likely to happen again? If so, is it likely to be even bigger? Suppose the system has built up a negative nodal weight for $\mu = 0$ at some point in the game (see Fig. 4a). It then hits node $\mu = 0$ at time t_0 producing a crash (Fig. 4b). The nodal weight c_0 is hence restored to zero (Fig. 4c). In this model the previous build-up is then forgotten because of the finite T score window, hence c_0 becomes positive (Fig. 4d). The system then corrects this imbalance (Fig. 4e), restoring c_0 to 0. The crash is then forgotten, hence c_0 becomes negative (Fig. 4f). The system should therefore crash again - however, a crash will *only* re-appear if the system's trajectory subsequently returns to node $\mu = 0$. Interestingly, we find that the *disorder* in the initial distribution of strategies among agents (i.e. the quenched disorder in $\underline{\Psi}$) can play a deciding role in the issue of crash 'births and revivals' since it leads to a slight bias in the outcome, and hence the subsequent transition, at each node. When $c_{\mu[t]} = 0$ (see Fig. 4c), it follows that $\text{sgn}[D[t]]$ is more likely to be equal to $\text{sgn}[\underline{a}^{\mu[t]} \cdot \underline{x}]$ where $\underline{x} = \sum_{R'=1}^{2P} \underline{\Psi}_{R'}$ is a strategy weight vector with x_R corresponding to the number of agents who hold strategy R [9]. The quenched disorder therefore provides a crucial bias for determining the future trajectory on the de Bruijn graph when the nodal weight is small, and hence can decide whether a given crash recurs or simply disappears. The quenched disorder also provides a *catalyst* for building up a very large crash.

Our work opens up the study of how a 'complex-systems-manager' might use this information to control the long-term evolution of a complex system by introducing, or manipulating, such large changes. As an example, we give a quick three-step solution to prevent large changes: (1) use the past history of outcomes to build up an estimate of the score vector $\underline{S}[t]$ and the nodal weights $\{c_{\mu[t]}\}$ on the various critical nodes, such as $\mu = 0$ in the case of the fixed-node crash. (2) Monitor these weights to check for any large build-up. (3) If such a build-up occurs, step in to prevent the system hitting that node until the weights have decreased.

References

- [1] D. Sornette, *Physics World* **12**, 57 (1999); P.V. McClintock, *Nature* **401**, 23 (1999); R.G. Palmer, W.B. Arthur, J.H. Holland, B. LeBaron and P. Tayler, *Physica D* **75**, 264 (1994); R. Mantegna and H. Stanley, *Econophysics* (Cambridge University Press, Cambridge, 2000); J. Bouchaud and M. Potters, *Theory of Financial Risks* (Cambridge University Press, Cambridge, 2000); B. Huberman, P. Pirolli, J. Pitkow, and R. Lukose, *Science* **280**, 95 (1998); M. Nowak and R. May, *Nature* **359**, 826 (1992); T. Lux and M. Marchesi, *Nature* **397**, 498 (1999).
- [2] A. Johansen and D. Sornette, preprint xxx.lan.gov/cond-mat/0010050.
- [3] P. Bak, *How Nature Works* (Cambridge University Press, Cambridge, 1998); P. Ormerod, *Financial Times* (19 February 2001).
- [4] N.F. Johnson, M. Hart, P.M. Hui, and D. Zheng, *Int. J. Theo. Appl. Fin.* **3**, 443 (2000); P. Jefferies, N.F. Johnson, M. Hart, and P.M. Hui, *Eur. J. Phys. B* **20**, 493 2001.
- [5] W. Arthur, *American Economic Review* **84**, 406 (1994); see the papers of D. Challet, M. Marsili et al. at www.unifr.ch/econophysics/minority; N.F. Johnson, P.M. Hui, D. Zheng, and C.W. Tai, *Physica A* **269**, 493 (1999).
- [6] D. Lamper, S. Howison, and N.F. Johnson, *Phys. Rev. Lett.* **88**, 017902 (2002).
- [7] A. Blair, *Guide to Charting* (Pitman Publishing, London, 1996).
- [8] P. Jefferies, M. Hart, and N.F. Johnson, *Phys. Rev. E* **65**, 016105 (2002).
- [9] Further details available from the authors via www.maths.ox.ac.uk/~lamper.

Figure Captions

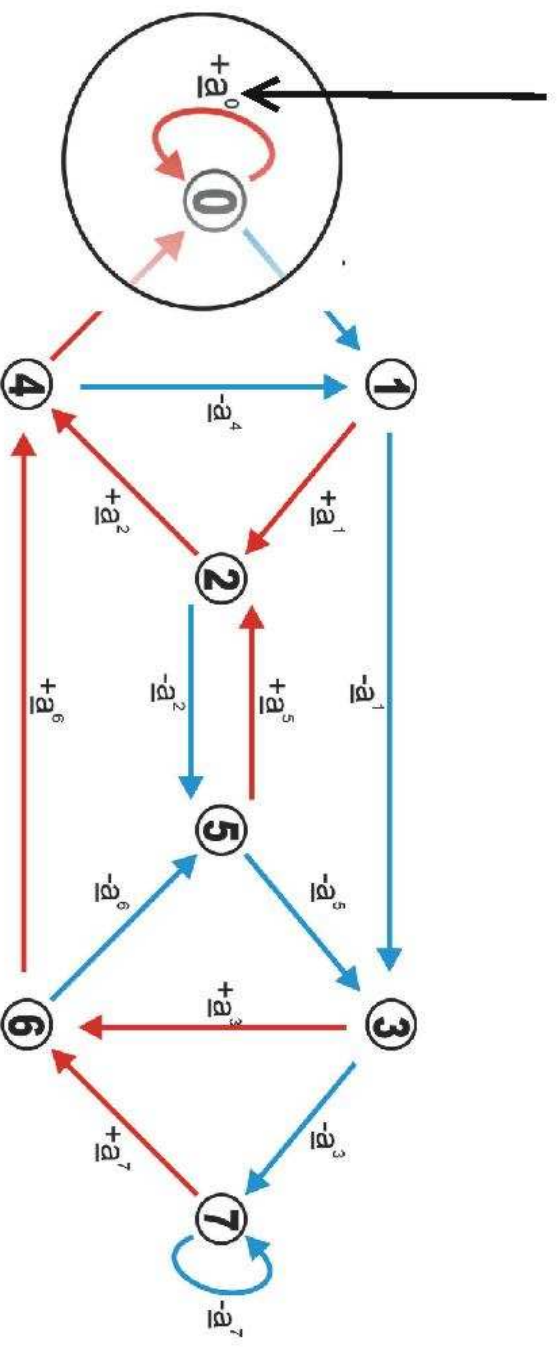
Figure 1: Dynamical behaviour of the global information is described by transitions on the de Bruijn graph. Graph for population of $m = 3$ agents. Blue transitions represent positive demand D , red transitions represent negative demand.

Figure 2: Dynamical behaviour of complex system (e.g. price $P[t]$ in financial market) described by evolution of nodal weights c_μ . History at each timestep indicated by black square. Large change preceded by abnormally high nodal weight. Large change incorporates fixed-node and cyclic node crashes

Figure 3: (a) Schematic representation of strategy score distribution prior to crash. Arrows indicate subsequent motion during crash period. (b) Plots of expected demand and volume during crash period showing range of different possible behaviour as system parameters are varied.

Figure 4: Representation of how large changes can recur due to finite memory of agents. Grey area shows history period outside agents' memory. Example shows recurring fixed-node crash at node $\mu = 0$.

transition incurs an increment
to the score vector \underline{S} :



Stable behavior: path with all transitions equally visited

e.g. $0 \rightarrow 0 \rightarrow 1 \rightarrow 3 \rightarrow 6 \rightarrow 5 \rightarrow 3 \rightarrow 7 \rightarrow 7 \rightarrow 6 \rightarrow 4 \rightarrow 1 \rightarrow 2 \rightarrow 5 \rightarrow 2 \rightarrow 4 \rightarrow \dots$

Crash: path with many negative (positive) return transitions

e.g. $0 \rightarrow 0 \rightarrow 0 \rightarrow 0 \rightarrow \dots$ or $2 \rightarrow 4 \rightarrow 0 \rightarrow 0 \rightarrow 1 \rightarrow 2 \rightarrow 4 \rightarrow 0 \rightarrow \dots$

fig2





fig3a




fig3b

