



ELSEVIER

Journal of Economic Dynamics & Control
000 (2001) 000–000JOURNAL OF
Economic
Dynamics
& Control

www.elsevier.nl/locate/econbase

Self-organized criticality in evolutionary systems with local interaction ☆

Alex Arenas^a, Albert Díaz-Guilera^b, Conrad J. Pérez^b,
Fernando Vega-Redondo^{c,d,*}

^aDepartament d'Enginyeria Informàtica, Universitat Rovira i Virgili, Tarragona, Spain

^bDepartament de Física Fonamental, Universitat de Barcelona, Spain

^cDepartamento de Fundamentos del Análisis Económico and Instituto Valenciano de Investigaciones Económicas Universidad de Alicante, Spain

^dDepartament d'Economia i Empresa Universitat Pompeu Fabra, Barcelona, Spain

Received 10 April 2000; accepted 19 March 2001

Abstract

This paper studies a stylized model of local interaction where agents choose from an ever increasing set of vertically ranked actions, e.g. technologies. The driving forces of the model are infrequent upward shifts ('updates'), followed by a rapid process of local imitation ('diffusion'). Our main focus is on the regularities displayed by the long-run distribution of diffusion waves and their implication on the performance of the system. By integrating analytical techniques and numerical simulations, we come to the following two main conclusions. (1) If non-coordination costs are sufficiently high, the system behaves *critically*, in the sense customarily used in physics. (2) The performance of the system is optimal at the frontier of the critical region. Heuristically, this may be interpreted as an indication that (performance-sensitive) evolutionary forces induce the system to be placed 'at the edge of order and chaos'. © 2001 Published by Elsevier Science B.V.

JEL classification: C72; O31

Keywords: Self-organization; Criticality; Local interaction; Technological change

☆ This work has been supported by the Spanish Ministry of Education, CICYT Project nos. 96-0168 and 97-0131.

* Corresponding author. Facultad de Económicas, Universidad de Alicante, 03071 Alicante, Spain. Tel.: 34-965-90-3616; fax: 34-965-90-3685.

E-mail address: vega@merlin.fae.ua.es (F. Vega-Redondo).

1 1. Introduction

3 Social and economic change typically comes in ‘waves’ or ‘avalanches’
 with seemingly little intertemporal structure. Its complex dynamics are partly
 a result of the following two features: (a) the stimulus for change spreads
 5 throughout the population by cumulative local interaction channeled along
 a social network; (b) the payoff incentives that govern individual behavior
 7 are subject to considerations of (local) compatibility and/or coordination with
 neighbors. In this paper, we focus on a simple model displaying these two
 9 features. To fix ideas, we shall often propose a technological interpretation
 of the model and conceive its dynamics as an unbounded process of tech-
 11 nological change. Its motion may then be seen as fueled by two forces:
 (i) infrequent upward perturbations, which are payoff independent and will
 13 be simply referred to as ‘updates’; (ii) relatively rapid payoff-responsive ad-
 justment, which will be known as ‘diffusion’. Thus, for the sake of focus, we
 15 study a context where diffusion is the only *purposeful* activity, reacting to a
 stochastic process of gradual ‘invention’ that is fully *exogenous*.

17 In our stylized framework, it turns out that different technological scenarios
 can be solely characterized by the value of a certain cost related parameter
 19 $k \in \mathbb{R}_+$. This parameter reflects the relative costs of upward versus lower
 non-coordination with neighbors, higher values of k indicating a situation
 21 where the former costs are relatively larger than the latter. In this context,
 our main conclusions can be summarized as follows. First, we find that there
 23 is a certain threshold for k which (roughly) separates the region where the
 population evolves in almost perfect coordination from the region where a
 25 wide range of technological heterogeneity tends to persist over time. More-
 over, within the latter region, the system behaves in a *critical* manner, the
 27 concept of criticality being understood here as in modern physics (Bak et al.,
 1987). In particular, it refers to the fact that the distribution over ‘avalanche
 29 sizes’ (i.e. the range of diffusion waves) generated by the recurrent updates
 obeys a power law. This fact has interesting theoretical implications. For
 31 example, it indicates that the system displays no characteristic scales (or mag-
 nitudes), the trade-off between avalanche size and its corresponding empirical
 33 frequency remaining constant throughout (in terms of proportional changes).¹
 But perhaps more importantly for our present purposes, another useful impli-
 35 cation of criticality pertains to the extreme simplicity displayed by long-run
 aggregate behavior under those circumstances. Such simplicity will render it

¹ Such a constant tradeoff can also be understood as an indication that local stimuli have
 a relatively large (i.e. non-local) range of influence. (For example, this would not occur if
 the decay of avalanche frequency with respect to size were exponential.) Traditionally, it was
 thought that physical environments had to be *accurately* tuned to some critical state in order to
 obtain such long-range effects. Since the work of Bak et al. (1987), such criticality is known
 to be attainable without any external fine-tuning (i.e. through self-organization).

1 tractable to carry out an analytical study of the model, thus shedding light
on some important issues. In particular, our main focus here will be on how
3 the long-run performance of the system depends on k , the key parameter of
the model.

5 In this latter respect, our main conclusion is quite stark: under some
natural measure of ‘performance’ (see later for details), the behavior of the
7 system is optimized within a rather *thin* region for k where the system has
already abandoned the synchronization region (a region that physicists of-
9 ten call supercritical) and has barely entered the critical region. Borrowing a
well-known phrase of Kaufman (1993), we may interpret this conclusion as
11 a tangible embodiment of the tenet that many dynamical systems have their
performance optimized at the thin ‘edge of order and chaos’.

13 The reason why *some* heterogeneity must be conducive to better perfor-
mance is not difficult to comprehend: only if some inter-agent differences arise
15 along the process may diffusion have a significant role to play. However, to
understand why optimality should be attained at the ‘brink’ where such het-
17 erogeneity is about to recede is a more subtle issue. In essence, we shall
find it to be a consequence of the following intriguing feature displayed by
19 critical behavior: across different avalanche sizes *and* values of k , there is a
constant (proportional) effect of avalanche size on the induced technological
21 advance.

In this paper, our approach to studying these phenomena relies on a syn-
23 ergic use of both numerical simulations and analytical techniques. On the
one hand, we resort to numerical simulations to obtain some regularities (e.g.
25 the verification of power laws) on which to build our ensuing formal anal-
ysis. This analysis is then used to gain a theoretical understanding of some
27 of the key conclusions also obtained from numerical simulations (e.g. those
concerning the optimal performance of the system). Overall, it is such a
29 ‘symbiosis’ of mathematical analysis and numerical results that will provide
us with most of the insights obtained on a dynamical system that, because of
31 its huge complexity, would be hardly tractable otherwise (e.g. by an exclusive
use of analytical techniques).

33 This paper owes much of its original inspiration to the booming literature
on self-organized criticality developed in physics over the last decade. (The
35 seminal articles are Bak et al. (1987, 1988), whereas a good recent survey can
be found in the monograph by Jensen (1998).) This field of research started
37 with the study of simple *sandpile-like* setups, where many of the essential
ideas and insights first originated. Subsequently, there has been a wide variety
39 of different contexts (biological evolution, the time pattern of earthquakes,
the functioning of the brain, the dynamics of traffic jams, etc.) where simple
41 adjustment rules have been seen to generate self-organized critical behavior
(see Bak (1996) for an informal survey of applications). More in line with
43 our present concerns, critical behavior has also been observed in the way

1 ‘information’ diffuses through social systems. For example, Redner (1998) has
 3 showed that the ‘waves’ of academic quotation display power distributions,
 and similar conclusions have been found for how connections are distributed
 5 in the internet (Faloutsos et al., 1999) or in the World-Wide Web (Albert
 et al., 1999 or Huberman and Adamic, 1999). In economics and game the-
 7 ory, these ideas have received only little attention. However, they have been
 applied to the study of business cycles by Scheinkman and Woodford (1994),
 economic geography by Krugman (1996), or games displaying strategic sub-
 9 stitutabilities (i.e. ‘anti-coordination’ games) by Agliardi (1998, Chapter 6).
 Our essential change of focus over this work is that our objective here is not
 11 only positive (describing long-run regularities of complex dynamic systems)
 but also normative, i.e. linking the particular form of these regularities to the
 13 performance of the system. In spirit, therefore, our perspective is quite akin to
 that held by those researchers (see Arthur et al. (1997) for a wide collection
 15 of representative work) whose objective has been to understand the interplay
 between complexity and performance in large socio-economic systems. This
 17 is also the concern of a previous paper of ours (Arenas et al., 2000), where
 some preliminary results along these lines were originally reported.

19 Our model has also borrowed some important features from the recent
 evolutionary literature on learning in games, initiated by the seminal pa-
 21 pers of Kandori et al. (1993) and Young (1993). In particular, we pursue
 the same methodology of integrating selection adjustment (here, myopic best
 23 response) with occasional perturbations. More specifically, the present model
 is akin to those of Blume (1993), Ellison (1993), or Young (1998) where,
 25 as in our case, a topology of local interaction is introduced and each agent
 is constrained to playing the game with her immediate neighbors. The game
 27 being played, on the other hand, bears some key similarities to the so-called
 minimum-effort game (see Bryant, 1983; van Huyck et al., 1990 or Craw-
 29 ford, 1991), which has also been widely studied in the equilibrium-selection
 literature. Our approach, however, displays two essential differences with this
 31 literature: (i) the setup involves a ‘changing game’ with an unbounded set of
 potential actions; (ii) the concern is not one of equilibrium selection but of
 33 dynamic performance.

In fact, it is precisely the ever-changing nature of the (coordination) game
 35 that allows us to interpret the dynamics induced as a process of technological
 change carried out under local complementarities (recall our former discus-
 37 sion). This then brings to mind two vast strands of related research: the
 literature concerned with the so-called network externalities and that studying
 39 technological change and growth. We address each of them in turn.

Network externalities have been a major field of study in the theory of
 41 industrial organization at least since Farrel and Saloner (1985) and Katz and
 Shapiro (1985) wrote their seminal papers on the subject—see Economides
 43 (1996) for a recent survey. The primary aim of this literature has been to

understand, from a strategic viewpoint, the considerations that might promote or deter technological adoption when the payoffs to it are directly dependent on the decisions of others; specifically, that is, on the number of other producers (or consumers) producing (or buying) the new good.

On the other hand, concerning the relationship between technological change and economic growth, this issue has regained a central role in recent times due to the importance accorded to it by the so-called New Growth Theory. A good case in point is the Schumpeterian model of growth through ‘creative destruction’ developed by Aghion and Howitt (1992), where innovations (that improve the whole production activities in the economy) arrive stochastically through investment in R & D geared towards the enjoyment of (temporary) monopoly gains. This model was enriched by Grossman and Helpman (1991) by integrating the technological-ladder approach proposed by Aghion and Howitt with the multi-sectorial features of a model formulated by Sergestrom et al. (1990). In such an enriched Schumpeterian model, technological ladders are climbed in each separate sector by specific R & D expenditures targeted at each of them. In every sector, therefore, a process of creative destruction unfolds, which forever pushes the economy’s average technological level upwards over time.

Heuristically, one could interpret our model as reflecting a stylized process of technological change that merges the Schumpeterian features displayed by the Grossman-Helpman approach with the technological complementarities studied by the aforementioned literature on network externalities. In contrast with the former, innovation is taken here to be purely random (i.e. not the outcome of purposeful and forward-looking agents). And, in contrast with the latter, the technological complementarities are studied dynamically, thus amounting to an unbounded process of technological growth through (purposeful, albeit myopic) diffusion.

The rest of the paper is organized as follows. Section 2 describes the framework. Section 3 presents the dynamics. Section 4 presents the numerical simulations: while its Section 4.1 focuses on the regularities displayed by the induced limit distributions, Section 4.2 is concerned with identifying interesting conclusions regarding issues of self-organization and long-run performance. Section 5 pursues a task of synthesis. That is, it builds upon the regularities found in Section 4.1 to provide an analytical explanation for the results obtained in Section 4.2. Finally, Section 6 closes the paper with a summary and a discussion of some issues left for future research.

2. The framework

We consider n agents, each of them occupying a particular node in a one-dimensional boundariless lattice (i.e. a ring). Time is discrete. At every

1 $t = 0, 1, 2, \dots$, each agent $i \in N \equiv \{1, 2, \dots, n\}$ adopts a certain action $a_i(t) \in \mathbb{R}_+$ that, for concreteness, may be conceived as the technology level she
 3 currently uses. Every agent i is assumed to interact with the individual to the right and to the left of her. For simplicity, these neighbors are taken to be
 5 those agents with adjacent indices, $i + 1$ and $i - 1$, where 0 and $n + 1$ are respectively interpreted as n and 1. Out of each of her two interactions, agent
 7 i obtains corresponding payoffs, $\psi(a_i(t), a_{i+1}(t))$ and $\psi(a_i(t), a_{i-1}(t))$, where $\psi: \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}$ is a fixed and common *payoff function*.

9 We shall postulate that the payoff function may be written in the following way:

$$\psi(a, a') = f(a) - g(a, a') \quad (1)$$

11 for some function $f: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ which is unboundedly increasing (i.e. $\lim_{a \rightarrow \infty} f(a) = \infty$) and a function $g: \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ which is bounded and satisfies

$$g(a, a') > 0 \Leftrightarrow a \neq a'. \quad (2)$$

13 For any given action $a \in \mathbb{R}_+$, $f(a)$ may be viewed as the payoff ceiling for this action. In view of (1) and (2), this maximum payoff is attained in any
 15 given interaction with someone playing a' if, and only if, $a = a'$. Otherwise, there are some bounded ‘incompatibility costs’ that detract from the base
 17 payoff. By way of illustration, we may think of every two neighboring agents as involved in the completion of a certain joint project, for which dissimilarity
 19 of actions (or technological levels) leads to some waste of resources.

In general, the incompatibility costs incurred by any given agent may arise
 21 from two alternative sources: (i) the agent is too advanced relative to her neighbors; or (ii) she is too backward. In either case, the induced effects
 23 may be of different significance. To account for this possibility, we posit:

$$g(a, a') = \begin{cases} \gamma_1(a' - a) & \text{if } a' > a, \\ \gamma_2(a - a') & \text{if } a > a', \end{cases} \quad (3)$$

where the functions $\gamma_1, \gamma_2: \mathbb{R}_{++} \rightarrow \mathbb{R}_{++}$ reflect, respectively, the negative pay-
 25 off consequences of being more or less (technologically) advance than one’s partner. For simplicity, these two functions will be taken to be ‘scaled sym-
 27 metric counterparts’ in the following sense: there are positive parameters k_1 and k_2 such that, for all $x \in \mathbb{R}_{++}$,

$$\frac{1}{k_1} \gamma_1(x) \equiv \frac{1}{k_2} \gamma_2(x). \quad (4)$$

29 Intuitively, the larger (smaller) k_1 is as compared to k_2 the more (less) detri-
 mental it is to be more advanced than one’s partner as opposed to being
 31 more backwards. As it turns out, only the difference $k \equiv k_1 - k_2$ will play a relevant role in the analysis.

Thus, heuristically, one may conceive the context described as reflecting a situation where every two neighboring agents, i and $i + 1$, are involved in a coordination game that, if faced in isolation, would induce both of them to choose the same action. Nevertheless, the key feature of our approach is that every player i must choose a *common* action in each of the games she plays. Therefore, the games played by agent i with $i - 1$ and $i + 1$ are not independent, i.e. cannot be treated in isolation. This is precisely the assumption that renders the model interesting, and is akin to that posited by the received evolutionary literature concerned with equilibrium selection in coordination games (recall the Introduction). Unlike this literature, however, we contemplate an unbounded ladder of possible actions where players may coordinate. This provides the basis for the rich adjustment dynamics that will be seen to arise along the induced ‘ever-moving game’. A formal description of this dynamics is undertaken in the next subsection.

3. The dynamics

Within the basic framework just introduced, we posit an adjustment dynamics displaying the following two components:

Diffusion: When receiving a revision opportunity, each agent i behaves ‘myopically’ and adopts an action that is a best response to what her neighbors are currently doing. Such a diffusion component of the process is taken to operate in a relatively fast manner.

Updates: Occasionally, an agent is subject to an exogenous perturbation (that may be conceived as an innovation) which shifts her action (technology level) upward by some randomly chosen amount. This component of the process is to be thought of as relatively slow.

For the sake of tractability, the diffusion and update processes are formally decoupled. That is, we postulate that the updates only perturb (‘punctuate’) the system when the diffusion has reached a standstill. Thus, in between any two consecutive updates, the process is supposed to have enough time to reach a point where no agent wants to revise her action any further. This is a convenient but extreme assumption made by the whole literature on self-organized criticality (recall the Introduction), since it allows for sharp definitions of the key notions of wave or avalanche (see below). However, the same long-run qualitative behavior of the model would be observed if the two dynamics (diffusion and updates) were genuinely integrated but the former happened to be much faster than the latter. In this sense, therefore, such a less extreme situation can be seen as suitably approximated by a context where both dynamics are formally separated.

1 Next, we describe precisely each of those two components of the dynamics:
 first the slow updates, then the fast diffusion.

3 3.1. Updates

5 Updates are indexed by $t = 1, 2, \dots$. As explained, after any given update,
 there is a diffusion wave that must reach a standstill before the next update
 occurs. Denote by $a(t) \equiv [a_i(t)]_{i=1}^n$ the profile displayed by the population
 7 once the diffusion phase triggered by update t has come to a halt. Then,
 update $t + 1$ is taken to operate as follows.

9 A single agent $i(t + 1)$ is randomly chosen to have her technological level
 subject to an upward shift. Specifically, her new technological level becomes
 11 $a_{i(t+1)}(t) + \tilde{\sigma}$ where $\tilde{\sigma}$ is a i.i.d. random variable, distributed on a finite in-
 terval $[0, v]$ according to some continuous density φ . These updates play the
 13 role of exogenous (i.e. unmodelled) perturbations that shift the technologi-
 cal level of a particular agent upwards. They may be provided with several
 15 motivations, e.g. payoff shocks or population renewal. However, our preferred
 interpretation is that of ‘innovation possibilities’ arising in conjunction with
 17 ‘optimistic expectations’.

19 More specifically, we propose to view an update as embodying some new
 option/idea received by the agent in question that, when pursued, this agent
optimistically expects to be followed by a sufficient upward adjustment by her
 21 neighbors. Notice that in view of the assumption that every update originates
 at an equilibrium state, there is a natural *asymmetry* between upward and
 23 downward changes in this respect. That is, even if an agent were to receive
 exogenously the *option* of moving to a lower action (i.e. a downward update),
 25 she will never choose to do so if she believes that others would only react
 (if at all) in the same direction. For, in this case, whatever expectations she
 27 might hold on the pattern of adjustment that could ensue, she can only lose
 by adopting a downward update.

29 As explained below (see (ii)), diffusion adjustments turn out to display
 the aforementioned monotonicity property, i.e. they always operate in the
 31 same direction as the update that triggers them. Thus, we may rely on the
 considerations just described to provide a heuristic motivation for an update
 33 as the *combination* of an innovation plus optimistic expectations. In this light,
 it is natural to assume that, once some optimistic expectations of this kind are
 35 in place, they should exhibit some persistence. That is, they need not always
 be downgraded (say, to static expectations²) if the expected response fails

² Static expectations (i.e. the belief that the actions currently adopted by others will remain in place in the next period) can be taken to underlie the myopic best-response dynamics posited below for the diffusion dynamics. Thus, in this light, an update can be regarded as an infrequent and ‘optimistic’ deviation of this state of affairs.

1 to materialize in the ensuing avalanche. Otherwise, updates would generally
 2 prove to be an insufficient ‘fuel’ for *sustained* advance, and our model would
 3 eventually become trapped in an uninteresting deadlock.³ For convenience,
 4 we shall introduce such persistence into the model quite starkly and simply
 5 assume that any player who undertakes an update towards some $a(t)$ at any
 6 t never adjusts her level below $a(t)$ thereafter. Formally, this is captured by
 7 defining, for each agent i and each t , a lower bound ς_i^t on her adjustments
 8 that is given by the last former update experienced by this player. That is,
 9 the action level attained at that last update is assumed to act as a floor on
 player i ’s future adjustments (see below for details).

11 3.2. Diffusion

After agent $i(t)$ has been perturbed at t , a diffusion (or adjustment) process
 13 ensues. Let us index the stages of this process by $q = 0, 1, 2, \dots$ and let $\alpha_i^t(q)$
 14 denote the action chosen by any agent i in stage q . With this notation in
 15 hand, the diffusion process may be described as follows.

At $q=0$, we have $\alpha_j^t(0) = a_j(t-1)$ for $j \neq i(t)$ and $\alpha_{i(t)}^t(0) = a_{i(t)}(t-1) + \sigma_t$,
 17 where σ_t is the realization of the random variable $\tilde{\sigma}$ at t . Subsequently, at ev-
 18 very $q=1, 2, \dots$, agents are randomly chosen to revise their action. Specifically,
 19 we postulate that any agent $i \in N$ who receives an adjustment opportunity at
 20 stage q chooses an action $\alpha_i^t(q)$ that maximizes (myopically) her payoffs under
 21 the following *double* constraint. On the one hand, she cannot surpass the cur-
 22 rent action ceiling prevailing in her own neighborhood. On the other hand, her
 23 new action cannot fall below the floor given by her own last update. The moti-
 24 vation for the first constraint is that any adjustment by player i which exceeds
 25 the maximum action level $\beta_i^t(q-1) \equiv \max\{\alpha_{i+1}^t(q-1), \alpha_i^t(q-1), \alpha_{i-1}^t(q-1)\}$
 26 should be conceived as an ‘innovation’ and thus restricted to enter the system
 27 as an update.⁴ The motivation for the second was already explained above
 when introducing our formulation of updates.

29 Formally, each agent $i \in N$ who receives a revision opportunity at q is
 taken to solve the following optimization problem:

$$\begin{aligned} \max_{\alpha} [\psi(\alpha, \alpha_{i+1}^t(q-1)) + \psi(\alpha, \alpha_{i-1}^t(q-1))] \\ \text{s.t. } \varsigma_i^t \leq \alpha \leq \beta_i^t(q-1). \end{aligned} \quad (5)$$

³ For example, in the specific context studied in Section 4, we find that if the process starts at a homogeneous profile *no* single update can produce a genuine avalanche (involving more than one agent) when $k \equiv k_1 - k_2$ is above the threshold value of 3.164. Consequently, under these conditions, the process would not grow at all if the updates did not exhibit some persistence.

⁴ Thus, implicitly, we suppose that new (technological) information only diffuses locally. We believe that this informational restriction can be substantially relaxed but have not yet conducted any detailed explorations of different alternatives.

1 The above problem determines whether the agent in question might be in-
 3 terested in revising her former action. Denote by $A_i^t(q)$ the set of solutions
 to the above problem, which is assumed non-empty (possibly including sev-
 5 eral actions, of which one may be the status-quo $\alpha_i^t(q-1)$). Furthermore,
 let $\mathcal{J}^t(q) \equiv \{i \in N: \alpha_i^t(q-1) \notin A_i^t(q)\}$ stand for the set of agents whose
 prior action is not optimal. Then, if $\mathcal{J}^t(q) \neq \emptyset$, we choose at random one
 7 individual $i \in \mathcal{J}^t(q)$ and make:

- (a) $\alpha_i^t(q) \in A_i^t(q)$, any choice in $A_i^t(q)$ with equal probability;
- 9 (b) $\alpha_j^t(q) = \alpha_j^t(q-1)$ for all other agents $j \neq i$.

Once (a) and (b) have been implemented, the process enters stage $q+1$.
 11 The payoff function $\psi(\cdot)$ will be assumed to guarantee the following two-fold
 property (see (8) for an example):⁵

- 13 (†) $\alpha_i^t(q) \geq \alpha_i^t(q-1)$ for all i and q ;
- (‡) there is some finite \bar{q} at which there is no agent left in a position to
 15 revise her action, i.e. $\mathcal{J}^t(\bar{q}) = \emptyset$.

Then, having reached the latter \bar{q} , we make $a_i(t) = \alpha_i^t(\bar{q})$.

17 The concatenation of diffusion phases and updates defines the dynamical
 system under consideration. For each t , the adjustment process that restores
 19 stability after the corresponding update is called a (technological) *avalanche*,
 where we use the usual term coined in the physics literature for this phe-
 21 nomenon. We shall be interested in quantifying the size $s(t)$ of each avalanche
 at t as follows:

$$s(t) \equiv \#\{i: a_i(t) \neq a_i(t-1)\}, \quad (6)$$

23 where $\#\{\cdot\}$ stands for the cardinality of the set in question. We shall also
 concern ourselves with the total advance triggered by the avalanche, as given
 25 by

$$H(t) \equiv \sum_{i=1}^n [a_i(t) - a_i(t-1)]. \quad (7)$$

As explained next, our numerical simulations show that these magnitudes,
 27 $s(\cdot)$ and $H(\cdot)$, display interesting long-run regularities.

4. Numerical analysis

29 Given the large dynamic complexity of our model, some of its key fea-
 tures have been determined through the performance of extensive numerical

⁵ In general terms, if the ceiling-payoff function $f(\cdot)$ is increasing and weakly convex (e.g. linear) and the cost functions $\gamma_i(\cdot)$ are increasing and strictly concave, this is enough to ensure that the desired property holds.

1 simulations. For concreteness, the numerical analysis reported here will focus
 2 on a simple scenario where the conditions required from our general model
 3 are satisfied in a specially transparent manner. Other alternative specifications
 4 consistent with the conditions posited in Sections 2 and 3 have been found
 5 to yield the same qualitative conclusions.

6 Essentially, the only two components of the model which need to be
 7 specified pertain to the payoff function $\psi(\cdot)$ and the update density $\varphi(\cdot)$.
 Concerning the payoff function, let it be given by:

$$\psi(a, a') = a - k_1(1 - \exp(-[a - a']_+)) - k_2(1 - \exp(-[a' - a]_+)) \quad (8)$$

9 for some $k_1, k_2 > 0$, where $[x]_+ \equiv \max\{0, x\}$. In terms of the general frame-
 work given by (1) and (3), this amounts to making

$$f(a) = a \quad (9)$$

$$\frac{\gamma_1(x)}{k_1} = \frac{\gamma_2(x)}{k_2} = 1 - \exp(-[x]_+). \quad (10)$$

11 This is obviously compatible with the required conditions—in particular,
 12 (2) and (4). Moreover, under suitable parameter conditions, (9) and (10)
 13 turns the bilateral situation faced by every two neighbors into a coordina-
 14 tion game which, locally around an equilibrium, behaves like the well-known
 15 minimum-effort game—a context widely discussed in the evolutionary and
 16 experimental literature (recall the Introduction). Specifically, the game then
 17 displays the following local property: if two individuals with ‘marginally
 18 different’ actions play the game, the one with the lower one obtains a higher
 19 payoff. To see this, note that if we consider two different actions, a and a' ,
 20 with $\delta \equiv a' - a > 0$, the function $h(\cdot)$ defined by $h(\delta) \equiv \psi(a', a) - \psi(a, a')$ has
 21 the following derivative at $\delta = 0$:

$$h'(0) = 1 - (k_1 - k_2).$$

22 Thus, if $k_1 - k_2 > 1$ (which is wholly within the range of interest to be con-
 23 sidered below), we have $\psi(a', a) - \psi(a, a') < 0$ provided that a' is marginally
 larger than a .

24 It is easy to check that the payoff function given in (8) guarantees that
 25 the diffusion dynamics satisfies (†) and (‡) above. On the other hand, a fur-
 26 ther convenient feature following from this formulation is that the decision
 27 problem described in (5) has its optimal solution necessarily lying in the
 28 set $\{\alpha_{i+1}^t(q-1), \alpha_i^t(q-1), \alpha_{i-1}^t(q-1)\}$. That is, the optimal action must be
 29 one of those currently chosen in the neighborhood of the player in question.
 30 Obviously, this will facilitate matters in what follows by allowing us to han-
 31 dle the optimization problem faced by each agent in an especially simple
 32 manner.
 33

1 Finally, concerning the stochastic updates, we shall postulate that they are distributed uniformly and independently on $[0, 1]$, i.e.

$$\varphi(x) = 1, \quad \forall x \in [0, 1]. \quad (11)$$

3 Much of our analysis will dwell on the effect of the incompatibility-cost parameters, k_1 and k_2 , on the long-run evolution of the process. In fact, to
5 understand the main issues involved, it turns out that only the difference $k \equiv k_1 - k_2$ needs to concern us. In view of (8), one can interpret k as
7 the cost difference resulting from ‘downwards incompatibility’ (i.e. being too advanced) as compared to that derived from ‘upwards incompatibility’ (i.e.
9 being too backwards).

To gain a heuristic understanding of why this cost difference should be the
11 key parameter, suppose that the system starts at a synchronized state at some t . That is, the previous state $[a_i(t-1)]_{i=1}^n$ satisfies $a_i(t-1) = \hat{a}$ for some
13 \hat{a} and all $i = 1, 2, \dots, n$. Now suppose that the magnitude of the ensuing t th update is Δ and let i_o be the particular agent affected by it. We may then
15 ask: When will this update again lead the system (once the corresponding diffusion phase has come to an end) into a new synchronized state at t ? Of
17 course, this will occur if the payoff to any agent $i \neq i_o$ of adopting $\hat{a} + \Delta$ when at least one of his neighbors has done so is higher than if he were to
19 remain at level \hat{a} . Thus, if we focus on the only non-trivial case where one neighbor still adopts \hat{a} , the relevant inequality is:

$$2(\hat{a} + \Delta) - k_1(1 - e^{-\Delta}) > 2\hat{a} - k_2(1 - e^{-\Delta}),$$

21 or equivalently:

$$k \equiv k_1 - k_2 < k^*(\Delta) \equiv \frac{2\Delta}{1 - e^{-\Delta}}. \quad (12)$$

If (12) holds, the update carried out by i_o will lead the system into a new
23 synchronized state at the *common* level $\hat{a} + \Delta$, i.e. $a_i(t) = \hat{a} + \Delta$ for all $i = 1, 2, \dots, n$. Otherwise, the update will introduce some ‘local heterogeneity’
25 (i.e. asynchronization) around i_o .

In general, as Δ ranges from 0 to 1 (i.e. the support of $\tilde{\sigma}$), the value
27 for $k^*(\Delta)$ increases from $\lim_{\Delta \downarrow 0} k^*(\Delta) = 2$ to $k^*(1) = 3.164$. Thus, in view of our preceding discussion, one may expect to find a transition from fully
29 synchronized behavior (i.e. what physicists often call *supercriticality*) for $k \leq 2$ to increasingly more heterogenous behavior as k grows above this
31 threshold. Indeed, such a transition will be clearly observed in our simulations, with sharply critical behavior (i.e. power laws) arising when k approaches
33 the higher threshold given by $k^*(1)$.

Next, we discuss our numerical simulations in some detail. We shall focus
in turn on the following issues:

35

- long-run distributions for the sizes of the avalanches;




Fig. 1. Distribution of avalanche sizes, $k = 2, 3, 3.5, 4, 6$; $n = 4096$ —double logarithmic scale (log frequency versus log size).

- 1 • long-run distributions for the (technological) advances induced by
- 3 • long-run relationship between size and advance across different avalanches.

4.1. Power laws

5 We have obtained the (empirical) avalanche size distributions for different
 7 values of k and different values of n , the latter being considered in order to
 9 check for possible scale effects. The outcome of these simulations is summa-
 11 rized in Figs. 1 and 2, where the relationship is depicted in doubly logarithmic
 13 scale.⁶

In Fig. 1, we show the size distribution for different values of k and a fixed
 population size $n = 4096$. We can notice the transition from a regime ($k \leq 2$)
 where all avalanches are system-size wide to a regime ($k \geq 3.5$) where the
 avalanche size sharply obeys a power-law distribution. As this transition un-
 folds (e.g. around $k = 3$),⁷ one observes that avalanches of different sizes do

⁶ Each point in the diagram gives the average log-frequency across all sizes in a corresponding interval. To identify graphically every such interval we take its mid-point (also in logarithmic scale).

⁷ For intermediate values of $k \in (2, 3)$, the relationships between frequency and size is found to be increasing, thus reflecting a gradual approach to the situation where *all* valances are of full system size as k decreases towards 2. Fig. 1 only shows the results for $k = 2, 3$ in order not to complicate the diagram with the depiction of non-critical behavior in the transition phase.




Fig. 2. Distribution of avalanche sizes, $k=4$; $n=1024, 2048, 4096, 8192, 16384, 32768$ —double logarithmic scale (log frequency versus log size).

1 occur, but they do not yet exhibit the clear-cut power-law regularities that
2 are the hallmark of criticality.

3 On the other hand, Fig. 2 shows that the power-law relationship between
4 size and frequency displayed for any given k in the critical region is inde-
5 pendent of population size. For concreteness, this is shown for $k=4$ and n
6 spanning three orders of magnitude, but similar diagrams are obtained for all
7 k in the critical region and even wider ranges in population size.

8 In Fig. 3, we show that analogous conclusions are obtained for the empiri-
9 cal distribution of advances across different avalanches. This distribution also
10 obeys a power law for values of $k \geq 3.5$, which is the (rough) threshold given
11 before for critical behavior to start manifesting itself fully in terms of size
12 distributions. This conclusion can be shown to be independent of population
13 size, much as this was also shown to be the case in Fig. 2 for avalanche
14 size.

15 The above results indicate that, as soon as downwards (relative) incompat-
16 ibility costs become significant, avalanche size and avalanche advance both
17 display a surprising regularity: the size- or advance-elasticity of the corre-
18 sponding long-run frequencies are constant throughout. In physics, this phe-
19 nomenon goes by the name of criticality, a term that points to the absence
20 of characteristic scales, i.e. the lack of prominent (relative) scales at which
21 the system behavior mainly takes place.




Fig. 3. Distribution of avalanche advances, $k = 3, 3.5, 4, 6$; $n = 4096$ —double logarithmic (log frequency versus log advance).

1 Given the common qualitative features displayed by size and advance dis-
 3 tributions, it is typically expected (see Jensen, 1998) that both magnitudes
 5 also display a power-like relationship between them. Indeed, this is confirmed
 7 by Fig. 4, where we have plotted avalanche size versus associated *average*
 9 advance for different values of k in the critical region ($k \geq 3.5$).

7 Fig. 4 also points to an additional important feature of our simulation
 9 results: throughout the critical region, the elasticity of total average advance
 11 with respect to avalanche size (i.e. the slope in double logarithmic scale) is
 13 independent of k , remaining essentially constant to a value slightly above 1
 15 (more precisely, at roughly 1.2). In contrast with previous observations (cf.
 17 Figs. 1–3), note that such a constant elasticity is not simply a feature that
 prevails for *given* k but, rather, is a characteristic that holds uniformly across
different (critical) values for k . The implications of this intriguing observation
 will be discussed below at some length.

15 Finally, it is worth mentioning that critical behavior is gradually lost for
 17 very large values of k .⁸ Specifically, as $k \rightarrow \infty$, any interaction between
 neighboring sites vanishes and one obtains a process of so-called random
 deposition, a well known process in the study of surface growth (cf. Barabasi

⁸ In our present context, the main distinguishing features of criticality start to fade away when $k \geq 30$.




Fig. 4. Average advance versus avalanche size, double logarithmic scale $k = 3.5, 4, 6, 9$; $n = 4096$.

1 and Stanley, 1995). As usual in these cases, the transition towards such a
 3 state of affairs involves an intermediate phase where the frequency decay is
 exponentially decreasing in avalanche size (i.e. it is much faster than the
 5 power decrease displayed in the critical region). For very large k , therefore,
 a disproportionate amount of avalanches are of very small size. Since this
 7 state of affairs displays few implications of interest for our purposes, we do
 not describe it here in any detail.

To recapitulate, it is useful to provide a formal summary of the main
 9 power-law regularities observed in our simulations as follows. There are some
 thresholds \underline{k} and \bar{k} defining the critical region (roughly, $\underline{k} = 3.5$ and $\bar{k} = 30$)
 such that if $\underline{k} < k < \bar{k}$:

11 P1 The long-run distribution of avalanche sizes s follows a power law of the
 13 form:⁹

$$P(s) \sim 1/s^\gamma \quad (13)$$

for some $\gamma > 0$, dependent of k but independent of population size n .

⁹ As standard, the symbol \sim signifies an asymptotic long-run equality between the two sides of the expression, modulo suitable constants.

- 1 P2 The long-run distribution of total advances H follows a power law of the form:

$$P(H) \sim 1/H^\beta \quad (14)$$

- 3 for some $\beta > 0$, dependent of k but independent of n .

- 5 P3 The relationship between avalanche sizes and corresponding advances follows:

$$H \sim s^\alpha, \quad (15)$$

where the exponent $\alpha > 1$ is independent of k and n .

- 7 The above conclusions indicate that, when k belongs to the critical region, the dynamic behavior of the system is in sharp contrast with that displayed
9 by the customary approaches to modelling social behavior in coordination setups. For example, in received models of evolutionary game theory (both
11 when interaction is assumed local as well as global), a population facing a coordination game is almost always to be found ‘synchronized’ (i.e. at a
13 homogenous equilibrium). Instead, in our case, the heterogeneous waves (and corresponding advances) realized along the process will typically induce the
15 rich diversity one often observes in real-world phenomena (e.g. concerning technological change).

17 4.2. Criticality and long-run behavior

- In view of P1–P3, it is natural to conjecture that some macroscopic variables (e.g. certain population averages) might exhibit interesting long-run regularities as well. Indeed, this is confirmed in what follows concerning two
21 interesting ‘summaries’ of the limit behavior of the system. The first one is a simple measure of population heterogeneity, often called the width of the
23 system. The second one captures a certain measure of its performance. We address each of them in turn.

- 25 In problems of surface growth in physics, it is common to quantify the roughness of an interface (what in our case could be conceived as the heterogeneity of the technological profile) by its so-called *width*. Restricting
27 attention to states reached once the diffusion process is complete, the width
29 of the system after update t is defined by

$$W(t) = \sqrt{\overline{a(t)^2} - \overline{a(t)}^2}, \quad (16)$$

where the overline denotes spatial average, i.e.

$$\overline{a(t)^q} \equiv \frac{1}{n} \sum_{i=1}^n (a_i(t))^q.$$

- 31 Of course, the width of the system at any given point in time simply coincides with the *standard deviation* of the action levels. It is, therefore, a




Fig. 5. Time evolution of the system width, $k = 4$; $n = 64, 128, 256, 512, 1024$ —double logarithmic scale (log width versus log (updates/ n)).

- 1 *global* statistic which would seem to abstract from the (local) spatial gra-
 - 3 dients that are the essential driving force of the model. Our interest in this
 - 5 magnitude, however, derives from the fact that when the system is critical,
 - 7 its width turns out to be closely related (i.e. is proportional, in a suitably
 - 9 extended fashion) to the spatial gradient functions—see Barabasi and Stan-
 - 11 ley (1995) for details. Thus, in this sense, the width of our critical system
 - 13 provides average ‘local’ information and may be regarded as a convenient
 - 15 (indirect) measure of ‘technological roughness’ for the induced profiles.
- Figs. 5 and 6 below depict (again in doubly logarithmic scale) the time evolution of the width of the system, as defined in (16), for two different values of k in the critical region and several population sizes (we have found the same qualitative behavior for other values of k in the critical region and larger population sizes). Since we are only interested in the *statistical* properties of the profile dynamics, the evolution of this magnitude is averaged over 1000 independent runs,¹⁰ the time periods reflecting the number of updates having materialized up to that point.

¹⁰ Note that, of course, the interface roughness is not a smooth increasing function of time for each independent run!




Fig. 6. Time evolution of the system width, $k = 7$; $n = 64, 128, 256, 512, 1024$ —double logarithmic scale (log width versus log (updates/ n)).

1 The main conclusion to be derived from the above diagrams is that, within
 2 the critical region, there exists a well-defined *long-run level of heterogeneity*
 3 (technological diversity) on which the system tends to settle as the process
 4 evolves and ‘self-organizes’. This is another interesting long-run regularity
 5 that follows from critical behavior. It indicates that, even though avalanches
 6 of all sizes (and, therefore, profiles of very different roughness) are to be
 7 observed over time, the long-run level of heterogeneity (either locally or
 8 globally defined) grows monotonically, on an average, to some well-defined
 9 magnitude. Naturally, this magnitude depends on the relevant parameters of
 10 the model, k and n . In this respect, the dependence on k is as one would
 11 expect: the larger k is, the larger is the long-run width and the slower (i.e.
 12 farther into the future) such a long-run level is attained. On the other hand,
 13 the dependence on n is equally clear-cut but perhaps less intuitive: population
 14 size affects positively long-run width but has no implications on the speed at
 15 which it is attained.

16 Now, we turn to what seems a richer and possibly more interesting impli-
 17 cation of criticality. It concerns the relationship between criticality and some
 18 suitable measure of the performance of the system. To fix ideas, suppose that
 19 the parameter k of our model defines a family of ‘technological systems’
 (computer architecture, communication protocols, etc.) that display identical

1 payoff possibilities but differ in their internal vertical compatibility. More
 3 specifically, assume that the payoff potential for each system grows along
 a common ‘performance ladder’ that is applicable only when agents coordi-
 5 nate on the same action. On the other hand, pertaining to non-coordinated
 situations, each system differs in the relative magnitudes of upward versus
 7 downward compatibility, as captured by their corresponding value of k in
 (8).

Suppose that the user population consists of individuals belonging to a cer-
 9 tain organization (say, a big firm) and every adjustment involves a certain
 (arbitrary) cost $c > 0$ that is independent of the extent of the change—for
 11 example, any adjustment might involve buying a new piece of equipment at
 a fixed cost. In this context, we may ask ourselves the following question:
 13 What is the cheapest way (i.e. cheapest technological system) by which the
 organization may eventually attain some pre-specified (average) technological
 15 level? Clearly, if the pre-specified level is high enough, such a cost minimiza-
 tion is essentially equivalent to a maximization of the following magnitude:

$$\rho = \lim_{T \rightarrow \infty} \rho(T) = \lim_{T \rightarrow \infty} \frac{\sum_{t=1}^T H(t)}{\sum_{t=1}^T s(t)}, \quad (17)$$

17 where $H(t)$ and $s(t)$ are defined by (6) and (7). For want of a better term,
 ρ will be called the *performance rate* of the system.

19 Consider now a related route to motivate the above measure of perfor-
 mance. Suppose that, within every time period, there is a fixed amount of
 21 resources that can be devoted to the upgrading of the actions (technologies)
 of the population, each such adjustment (purchase of new equipment) still
 23 requiring some fixed cost $c > 0$. Furthermore, assume that there are always
 enough candidates for upgrades (either as ‘adjustment’ or ‘updates’), but those
 25 that are geared towards matching a neighbor’s action (diffusion) always en-
 joy higher priority than those that are *not* (i.e. updates). In this context,
 27 maximizing the long-run time *rate* of technological change is equivalent to
 maximizing ρ , as defined in (17).

29 To facilitate the discussion, write $\rho(T, k)$ and $\rho(k)$ to reflect the dependence
 of $\rho(T)$ and ρ on k . Several interesting conclusions concerning these magni-
 31 tudes are obtained from our numerical simulations, as depicted in
 Figs. 7 and 8.

33 Concerning Fig. 7, the first observation one may readily make is that,
 for any given k , the long-run rate $\rho(k) \equiv \lim_{T \rightarrow \infty} \rho(T, k)$ is a well-defined
 35 magnitude since $\rho(T, k)$ is a (bounded) monotone function of T . This, how-
 ever, does not entail any new information by itself, but is just an indirect
 37 confirmation that the avalanche and advance distributions display long-run
 distributions with well-defined averages. For, in view of (17), note that the
 39 performance rate can be simply computed as the ratio of the average advance
 over the average avalanche size over a time path.




Fig. 7. Time evolution of performance rate ρ for different values of k and $n = 512$.




Fig. 8. Long-run performance rate ρ for different values of k and $n = 128, 512, 2048$.

1 There are, however, some additional and genuinely new observations arising from Figs. 7 and 8 that will underlie much of our ensuing discussion. For
3 future reference, it is useful to organize them into the following two (partially

overlapping) conclusions.

1

C1 The system's dynamic performance (as reflected by ρ) enjoys positive dependence on n only within (or very close to) the critical region—cf. Fig. 8. This implies that, for large n , the performance rate ρ is optimized within (or very close to) the critical region.

5

C2 The optimal value of k that maximizes $\rho(k)$ is independent of n and is located at the 'lower edge' of the critical region, i.e. in a narrow range around $k = 3.5$.

7

9 By C1, significant *scale effects* (i.e. improvements in performance due to population increases) are intimately associated to criticality, i.e. they only arise when, roughly, $k \in (3, 30)$. On the other hand, C2 indicates that the performance of the system is maximized around the point where such criticality starts to set in. This latter conclusion may be interpreted along the lines put forward by some authors (e.g. Kaufman, 1993) who have argued that the dynamic response of large interacting systems is often optimized at the brink where heterogeneity and 'disorder' (i.e. criticality in our case) just begins to manifest itself.

17

The former considerations suggest that our analysis may be provided with some normative interpretation. Again, to be specific, suppose that k is a policy variable that determines the 'technological system' to be applied in a certain organization (recall our former discussion). Then, if the underlying conditions are suitably approximated by the model (in particular, updates are infrequent relative to adjustment and the rate ρ is a relevant measure of performance), our results indicate that the optimal k is one that induces some 'barely critical' amount of inter-agent heterogeneity.

25

In a related vein, one may instead approach matters from an evolutionary viewpoint. Specifically, suppose that the underlying environment consists of a variety of different organizations that are characterized by alternative technological systems (i.e. idiosyncratic k 's), all of them initially co-existing as part of a heterogeneous population. Then, our conclusions suggest that if evolutionary forces impinge on the population composition among different types of organizations (say, because of performance-related differences in survival or through social imitation), the long-run situation that should eventually prevail is one where organizations are placed at the edge of criticality.

35

5. Theoretical analysis

37

The entire simulation results reported in Section 4 refer to the numerical scenario given by (8) and (11), which is a particular case of the general

context defined by (1) and (2)–(4). Similar regularities (i.e., in particular, P1–P3 and C1 and C2) arise in alternative scenarios consistent with these conditions. Given the vast complexity of the underlying process, it would have been very hard to establish those results by other than numerical methods. However, once established, we are now in a much better position to build partially upon them to aim at understanding their relationship analytically.

Specifically, our theoretical approach will start by *postulating* P1–P3, i.e. it will assume that the system is critical for a certain parameter range. Then, based on this assumption, we shall strive to *derive* analytically the conclusions C1 and C2. Overall, this theoretical exercise will improve our understanding of the important relationship between criticality and the system's dynamic performance.

Our first task is to provide an explicit expression for the performance rate ρ . For simplicity, we postulate that, for any given $k > 0$, the system is either in a critical, a supercritical, or a subcritical regime. In practice, of course, there must be a gradual transition from supercriticality to criticality and then subcriticality as k grows in the range $(0, \infty)$. For our present analytical purposes, however, that range is assumed *partitioned* into three subregions by thresholds \underline{k} and \bar{k} . Specifically, the interval $(0, \underline{k})$ is taken to define the supercritical subregion, the interval (\underline{k}, \bar{k}) the critical subregion, and the remaining (\bar{k}, ∞) the subcritical one.

The determination of ρ outside of the critical region is straightforward. On the one hand, when the system is supercritical ($k < \underline{k}$), avalanches are system-wide (i.e. $s = n$) and hence ρ is equal to $\bar{\sigma} \equiv E(\bar{\sigma})$. On the other hand, in the subcritical region ($k > \bar{k}$), all avalanches can be taken to be of (approximately) unit size and therefore the rate ρ is again equal to $\bar{\sigma}$.

Now, consider the critical region where $\underline{k} < k < \bar{k}$. Denoting by $P(H)$ and $P(s)$ the corresponding discrete empirical densities, let

$$\bar{H} \equiv \int HP(H) dH = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T H(t),$$

$$\bar{s} \equiv \int sP(s) ds = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T s(t),$$

define the long-run averages of H and s . Within the critical region, these densities are well-defined. Therefore, from (17), we may compute the performance rate as follows:

$$\rho = \lim_{T \rightarrow \infty} \frac{\bar{H}T}{\bar{s}T} = \frac{\bar{H}}{\bar{s}}. \quad (18)$$

- 1 Then, we may rely on P1 and P2 to conduct the following (approximate) derivations:¹¹

$$\rho \sim \frac{\int_1^n s^\alpha s^{-\gamma} ds}{\int_1^n s s^{-\gamma} ds}$$

- 3 which is easily seen to lead to

$$\rho \sim \frac{2 - \gamma}{\alpha - \gamma + 1} \frac{n^{1+\alpha-\gamma} - 1}{n^{2-\gamma} - 1}. \quad (19)$$

- 5 Clearly, for large n , the above value for ρ is crucially dependent on the exponents of n (in particular, on whether they are positive or negative). In this respect, three different cases may be considered, with respective values
7 for ρ that may be approximated (for large n) as follows:¹²

- (i) If $\gamma < 2$,

$$\rho \sim \frac{2 - \gamma}{\alpha - \gamma + 1} n^{\alpha-1},$$

- 9 (ii) If $2 < \gamma < \alpha + 1$,

$$\rho \sim \frac{\gamma - 2}{\alpha - \gamma + 1} n^{(\alpha-1)-(\gamma-2)},$$

- (iii) If $\gamma > \alpha + 1$,

$$\rho \sim \frac{\gamma - 2}{\gamma - \alpha - 1}.$$

- 11 By P1–P3, the parameters involved in the above expressions (in particular, γ and α) can be taken to be independent of n , the population size. Therefore,
13 for large systems (large n), it follows that the performance rate $\rho(k)$ will be maximized somewhere in the critical region, i.e. for $k \in (\underline{k}, \bar{k})$. For only in
15 this region does the performance of the system benefit from ‘scale effects’. Outside of it (that is, in either the supercritical or subcritical regions), the
17 performance rate is (approximately) equal to $\bar{\sigma}$, independently of n .

- The previous considerations provide the basis for understanding C1. To
19 understand C2, we now rely on P3, i.e. avalanche sizes and the induced advances are related through a power law that is not only independent of
21 n but also holds unchanged for all k (of course, as long as k remains in the critical region). This key feature of critical behavior has the following

¹¹ Of course, this computation must be viewed only as an approximation conducted under the implicit assumption that the support of the conditional distributions $P(H | s = \bar{s})$ displays a relatively narrow support—cf. Jensen (1998, p. 38).

¹² Close to the borders between the different regimes, limits must be handled with care since divergencies pertaining to vanishing terms occur and different limit operations do not commute. These technical issues notwithstanding, a detailed analysis of (19) will convince the reader that ρ displays the properties later required (e.g. given n and $\alpha > 1$, it induces a well-defined decreasing function of γ).

1 striking implication: the effect of k on ρ becomes *solely* channelled through its
 2 effect on γ (i.e. the steepness of the size distribution). Building upon this key
 3 observation, we are now in a position to understand why the maximization
 of ρ must indeed be achieved at the lower edge of the critical region.

5 To this end, restrict attention to the critical region, and conceive ρ as a
 function of α and γ , that is extended continuously from the interior of the
 7 different subregions (i)–(iii) for given n . Then, for small changes in k , we
 can symbolically write:

$$\begin{aligned}\frac{\Delta\rho}{\Delta k} &= \frac{\Delta\rho}{\Delta\gamma} \frac{\Delta\gamma}{\Delta k} + \frac{\Delta\rho}{\Delta\alpha} \frac{\Delta\alpha}{\Delta k} \\ &= \frac{\Delta\rho}{\Delta\gamma} \frac{\Delta\gamma}{\Delta k}\end{aligned}$$

9 since, by P3, we have $\Delta\alpha/\Delta k=0$ (i.e. α is unaffected by k within the critical
 region). Now, it is easy to check from (19) that $\Delta\rho/\Delta\gamma<0$. Combining this
 11 latter fact with $\Delta\gamma/\Delta k>0$ (i.e. the weight of small avalanches grows with k),
 it follows that $\Delta\rho/\Delta k<0$, i.e. ρ is a decreasing function of k . Obviously,
 13 this implies that the maximization of ρ must occur at the lower boundary
 of the critical region. As desired, therefore, this provides an (approximate)
 15 analytical basis for those features of our numerical simulations which were
 stated in C2.

17 In a sense, C2 may be interpreted as suggesting that optimal performance
 builds upon a rather delicate compromise between ‘order’ (synchronization
 19 or homogeneity, i.e. supercriticality) and ‘disorder’ (criticality). Heuristically,
 the underlying intuition for why *some* such balance should be expected to
 21 arise is not difficult to understand. On the one hand, if all individuals were to
 advance in step because of low incompatibility costs, individual adjustments
 23 would always be relatively small and no steep gradients could ever arise.
 Consequently, the overall pace of advance should be slow, individuals hardly
 25 taking advantage of the ‘scale economies’ that a large system would avail.
 But, on the other hand, if incompatibility costs were large, the scale effects
 27 impinging on overall advance that could be potentially afforded by a large
 system would be, again, not fully taken advantage of. In this second case,
 29 avalanches would typically be too small for any heterogeneity to be profited
 by its required complement: an effective process of diffusion.

31 The previous considerations provide quite a clear intuition for why opti-
 mal performance should require criticality, i.e. a suitable trade-off between
 33 homogeneity and heterogeneity in agents’ unfolding behavior. What seems
 substantially more subtle is the additional sharper conclusion that optimal-
 35 ity is to be expected at the lower edge of the critical region. As explained,
 this appears to be intimately associated to the intriguing ‘empirical’ (i.e. nu-
 37 merical) finding that criticality induces a fixed power relationship between
 avalanche sizes and the induced advances.

1 6. Summary and conclusion

3 This paper represents a first step in a research project where we plan
 5 to study in detail the relationship among complexity, optimality, and self-
 7 organization in systems composed by a large number of locally interacting
 9 entities. As suggested above, many social and economic systems may be
 modelled in this fashion. In particular, this approach seems particularly well
 suited to study technological evolution when the decisions adopted by the
 different entities (individual agents, firms, or even sectors) display local com-
 plementarities.

11 In the context of a simple model with these features, we have seen that
 13 if non-coordination costs exceed a certain threshold (but are not too large),
 the system self-organizes itself into a critical state and the sizes of diffu-
 sion waves are distributed according to a power law. In this case, moreover,
 the roughness of the population profile converges to a well-defined (average)
 magnitude, thus indicating that there is a specific degree of long-run popula-
 tion heterogeneity associated to each parameter (cost) configuration. In fact,
 it turns out that this ‘critical’ heterogeneity plays a crucial role in the perfor-
 mance of the system. Specifically, we have seen that, given a natural measure
 of performance (that admits an interpretation of either average adoption cost
 or rate of technological advance), the system behaves optimally within the
 critical region—or, more precisely, at the lower edge of this region. This sug-
 gests that inter-agent heterogeneity plays a crucial role in the evolution of the
 system. In other words, either too ‘orderly’ (synchronized) or too ‘chaotic’
 (non-coordinated) dynamics is detrimental to performance in that it imposes
 too frequent (and thus costly) or too rare (and therefore unduly staggered)
 an adjustment on agents’ actions.

27 In ongoing research, we are studying a number of extensions of the present
 29 model. One of them concerns higher dimensional setups (specifically, a two
 -dimensional torus), where each individual has more than two neighbors and
 therefore the diffusion paths may exhibit richer geometries. A second exten-
 sion involves studying irregular (but fixed)¹³ networks of the small-world
 variety (see Watts and Strogatz, 1998). In each of these contexts, the same
 qualitative conclusions found in the present model are essentially maintained.
 However, in the latter case (small-world networks), one obtains the expected
 result that, due to a relatively short expected path between any two agents,
 the parameter range where synchronous behavior tends to arise becomes sig-
 nificantly larger than in regular networks.

39 This paper has shown that quite a simple model of social, but local-based,
 interaction may produce persistent and wide heterogeneity in the induced

¹³ Note that, in order to preserve the key local structure of interaction that underlies our
 analysis, the network cannot change in an unrestricted (say, time-independent) fashion.

1 population dynamics. It is clear, however, that not all social networks can
2 be expected to display such a behavior. For example, complete networks,
3 which display no local structure, can only exhibit uninteresting ‘waves’ in
4 a coordination context analogous to that considered here. In such a context,
5 fixed-size avalanches of the order of system size (and only those) would
6 occur as the number of updates accumulated since the last avalanche come
7 to exceed a certain threshold.¹⁴

8 Naturally, the above point raises the question of what networks might be
9 conducive to critical behavior and, more importantly, whether such a class
10 of criticality-supporting networks could arise *endogenously* when the social
11 network is not fixed but may also co-evolve as dictated by agents’ own de-
12 cisions. The issue, in a sense, is analogous to that addressed by the most
13 recent developments of evolutionary game theory that focus on how the par-
14 allel co-evolution of *both* players’ links (or connections) and their actions
15 may affect the received analysis on equilibrium selection.¹⁵ Indeed, not only
16 the concerns but also some of the techniques used in this literature (e.g. those
17 employed to analyze perturbed Markov processes) would seem quite appli-
18 cable to the problem at hand. The study of this important topic is left for
19 future research.

Acknowledgements

21 We are grateful to Andreu Mas-Colell, Henrik J. Jensen, and three anony-
22 mous referees for very helpful comments.

References

- 25 Aghion, P., Howitt, P., 1992. A model of growth through creative destruction. *Econometrica*
60, 323–351.
- 26 Agliardi, E., 1998. *Positive Feedback Economies*. MacMillan, London.
- 27 Albert, R., Jeong, H., Barabasi, A.-L., 1999. Diameter of the World-Wide Web. *Nature* 401,
130–131.
- 28 Arenas, A., Pérez, C., Díaz-Guilera, A., Vega-Redondo, F., 2000. Self-organized evolution in
29 socio-economic environments. *Physical Review E* 61, 3466–3469.
- 30 Arthur, W.B., Durlauf, S.N., Lane, D.A., 1997. *The Economy as an Evolving System*, Santa
31 Fe Institute, Proceedings Vol. XXVII. Addison-Wesley, Reading, MA.
- 32 Bak, P., 1996. *How Nature Works: the science of self-organized criticality*. Copernicus,
33 New York.
- 34 Bak, P., Tang, C., Wiesenfeld, K., 1987. Self-organized criticality: an explanation of $1/f$ noise.
35 *Physical Review Letters* 59, 381.

¹⁴ For simplicity, we are implicitly assuming every update to be of the same *given* magnitude.

¹⁵ See, for example, Bala and Goyal (2000), Jackson and Watts (2000), or Goyal and Vega-Redondo (2000).

- 1 Bak, P., Tang, C., Wiesenfeld, K., 1988. Self-organized criticality. *Physical Review A* 38, 364.
- 3 Bala, V., Goyal, S., 2000. A non-cooperative model of network formation. *Econometrica* 68
(5), 1181–1231.
- 5 Barabasi, A.-L., Stanley, H.E., 1995. *Fractal Concepts in Surface Growth*. Cambridge University
Press, Cambridge.
- 7 Blume, L., 1993. The statistical mechanics of strategic interaction. *Games and Economic
Behavior* 5, 387–424.
- 9 Bryant, J., 1983. A simple rational expectations Keynes type model. *Quarterly Journal of
Economics* 98, 525–529.
- 11 Crawford, V., 1991. An evolutionary interpretation of van Huyck, Battalio, and Beil's
experimental results on coordination. *Games and Economic Behavior* 3, 25–59.
- 13 Economides, N., 1996. The economics of networks. *International Journal of Industrial
Organization* 16, 673–699.
- 15 Ellison, G., 1993. Learning, local interaction, and coordination. *Econometrica* 4, 1047–1073.
- 17 Faloutsos, M., Faloutsos, P., Faloutsos, C., 1999. On power-law relationships of the internet
topology, ACM SIGCOMM '99. *Computers and Communication Review* 29, 251–263.
- 19 Farrel, J., Saloner, J., 1985. Standardization, compatibility, and innovation. *Rand Journal of
Economics* 16, 70–83.
- 21 Goyal, S., Vega-Redondo, F., 2000. *Learning, network formation and coordination*. Mimeo,
Erasmus University, Rotterdam and Universidad de Alicante.
- 23 Grossman, G.M., Helpman, E., 1991. Quality ladders in the theory of growth. *Review of
Economic Studies* 58, 43–61.
- 25 Huberman, B.A., Adamic, L.A., 1999. Growth dynamics of the World-Wide Web. *Nature* 401,
131.
- 27 Jackson, M., Watts, A., 2000. On the formation of interaction networks in social coordination
games. Mimeo, Caltech.
- 29 Jensen, H.J., 1998. *Self-Organized Criticality*. Cambridge University Press, Cambridge.
- 31 Kandori, M., Mailath, G., Rob, R., 1993. Learning, mutations, and long-run equilibria in games.
Econometrica 61, 29–56.
- 33 Katz, M., Shapiro, C., 1985. Network externalities, competition, and compatibility. *American
Economic Review* 75, 424–440.
- 35 Kaufman, S., 1993. *The Origins of Order: Self Organization and Selection in Evolution*. Oxford
University Press, Oxford.
- 37 Krugman, P., 1996. *The Self-organizing Economy*. Blackwells, Cambridge, MA.
- 39 Redner, S., 1998. How popular is your paper: an empirical study of the citation distribution.
European Physical Journal B 4, 131–134.
- 41 Scheinkman, J., Woodford, M., 1994. Self-organized criticality and economic fluctuations.
American Economic Review 84, 417–421.
- 43 Sergestrom, P.S., Anant, T.C.A., Dinopoulos, E., 1990. A Schumpeterian model of the product
life cycle. *American Economic Review* 80, 1077–1091.
- 45 van Huyck, J.B., Battalio, R., Beil, R.O., 1990. Tacit coordination games, strategic uncertainty,
and coordination failure. *American Economic Review* 80, 234–248.
- Watts, D.J., Strogatz, S.H., 1998. Collective dynamics of 'small-world' networks. *Nature* 393,
440–442.
- Young, P., 1993. The evolution of conventions. *Econometrica* 61, 57–84.
- Young, P., 1998. *Diffusion in social networks*. Mimeo, John Hopkins University.