
ELSEVIER
Volume 34, Issue 8
August 2010
ISSN 0165-1889
JOURNAL OF
Economic
Dynamics
& Control
CONTENTS
J.S. Chadha and S. Holly, Macroeconomic models and the yield curve: An assessment of the fit 1343
Y.S. Kim and M. Lee, A model of debit card as a means of payment 1359
B. Unel, Firm heterogeneity, trade, and wage inequality 1369
A. Blume, J. Duffy and T. Temzelides, Self-organized criticality in a dynamic game 1380
R.A. Brecher, Z. Chen and E.U. Choudhri, A dynamic model of shirking and unemployment: Private saving, public debt, and optimal taxation 1392
M. Kumhof, On the theory of sterilized foreign exchange intervention 1403
E. Kováč, V. Vinogradov and K. Žigić, Technological leadership and persistence of monopoly under endogenous entry: Static versus dynamic analysis 1421
W. Huang, On the hidden hazards of adaptive behavior 1442
CONTENTS CONTINUED ON BACK COVER
www.elsevier.com/locate/jedc

Journal of Economic Dynamics & Control 34 (2010) 1380–1391

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

Self-organized criticality in a dynamic game

Andreas Blume [a], John Duffy [a,*], Ted Temzelides [b]

[a] Department of Economics, University of Pittsburgh, Pittsburgh, PA 15260, USA
[b] Department of Economics, Rice University, Houston, TX 77251, USA

ARTICLE INFO

Article history:
Received 19 October 2006
Accepted 9 March 2010
Available online 14 April 2010

JEL classification:
C72
C73
E32

Keywords:
Self-organization
Criticality
Local interaction
Power Law
Entry Game

ABSTRACT

We investigate conditions under which *self-organized criticality* (SOC) arises in a version of a dynamic *entry game*. In the simplest version of the game, there is a single location—a pool—and one agent is exogenously dropped into the pool every period. Payoffs to entrants are positive as long as the number of agents in the pool is below a critical level. If an agent chooses to exit, he cannot re-enter, resulting in a future payoff of zero. Agents in the pool decide simultaneously each period whether to stay in or not. We characterize the symmetric mixed strategy equilibrium of the resulting dynamic game. We then introduce local interactions between agents that occupy neighboring pools and demonstrate that, under our payoff structure, local interaction effects are necessary and sufficient for SOC and for an associated power law to emerge. Thus, we provide an explicit game-theoretic model of the mechanism through which SOC can arise in a social context with forward looking agents.

1. Introduction

Are large (aggregate) shocks to an economic system necessary for a noticeable change in equilibrium outcomes? One might suspect that this is the case since small independent shocks tend to cancel one another and, at least in large economic systems, they wash out without any aggregate implications. Alternatively, it is plausible that several social phenomena as diverse as the change in fashion and music trends, bank runs and other financial crises, investment in new technologies, or even recessions occur when tension from small shocks, gradually built into the system without notable aggregate implications, reaches a critical level. When such a critical state is reached, a small disturbance might have a disproportionately large effect. This phenomenon has been studied in the natural sciences in connection to avalanches, forest fires, earthquakes and even biological evolution and is referred to as *self-organized criticality* (SOC).

Models of SOC in the natural sciences involve the following distinct feature.[1] States that are "critical" typically constitute a small part of the state space. Nevertheless, the system "organizes itself" towards the situation in which a critical state is eventually reached. A useful expository example is that of a sand pile. Consider dropping one grain of sand at a time, say in a box. Most of the time such small temporary shocks will affect at most the neighboring area near where the grain fell. However, inevitably, the system will lead itself to a state where the addition of one grain leads to an avalanche. In fact, over a long time horizon, various statistics, including the distribution of sizes of avalanches, have been shown in simple models of sand piles to obey power laws.

* Corresponding author. Tel.: +1 412 648 1733; fax: +1 412 648 1793.
E-mail address: jduffy@pitt.edu (J. Duffy).

[1] For early reference see the seminal paper by Bak et al. (1988) as well as Bak and Chen (1991). For a summary of more recent research see Jensen (1998). See Bak and Sneppen (1993) for an application to evolutionary biology.

0165-1889/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.jedc.2010.04.002

Power laws are widely considered to be a signature of SOC and the search for SOC in various models typically turns into a search for power laws.[2] Such laws have important implications for the long run behavior of the system under study for at least two reasons. First, the frequency of events decreases slower than exponentially in their respective sizes. In other words, not only can local "perturbations" have large non-local effects, but these effects also appear with a relatively large frequency. Second, power laws are considered to suggest that both large and small events are the result of the same underlying mechanism at work.

Power laws might be created through other mechanisms that often assume some degree of bounded rationality.[3] This raises an important question. Is bounded rationality necessary for power laws to emerge? A novelty of our approach is that it offers an intuitive mechanism through which SOC and the associated power laws emerge in a purely game-theoretic setup with fully rational agents. Thus, we demonstrate that SOC might arise naturally in game-theoretic models and that this does not rely on restrictive assumptions on agents' rationality.

Models of self-organized criticality have been studied in economics.[4] However, as mentioned above, existing models involve restrictive assumptions on the behavior/rationality of agents. In this paper we study two versions of the model. In one version agents are myopic, while in the second version they are forward looking. We study the full rationality benchmark not necessarily because we believe that it is the best descriptive assumption, but rather because it allows us to explore an alternative economically meaningful mechanism as a potential source of SOC: the independent randomizations by individual players in a mixed strategy equilibrium of a dynamic game.

Our basic setup is that of a dynamic *entry game*. One way of motivating this game is by imagining that a finite number of firms decide whether to enter a market or not. If a single firm enters the market it enjoys monopoly profits. In general, profits are monotonically decreasing in the number of firms competing. If the number of firms in the market is large enough, a firm will be better off not entering, making an exogenous reservation profit instead. A second interpretation relates to computer science. In an environment where the number of potential users of a processor increases over time, each existing user might become worse off and "renege from the queue." This is because as the number of users increases, the work done for each individual user might decrease due to lower processing speeds.[5]

We assume a large number of identical agents, say around a swimming pool, and an infinite time horizon. In each period, one agent is chosen and is dropped into the pool. Agents like to stay in the pool only if it is not "too crowded." That is, payoffs are decreasing in the number of participants, becoming negative if this exceeds a fixed finite upper bound. The decision that each agent faces is whether to stay in the pool or to exit. If he exits, he earns a fixed utility, normalized to zero, in each future period, with no possibility of re-entering.

What are the likely outcomes in this model? If a few agents are in the pool, the unique best response is to stay in. However, the pool then gets increasingly "crowded" as one additional agent enters in each period. Eventually, a number of players is reached after which, assuming that everyone else stays in, the best response of an individual player is to exit. From that state on, the best action depends on what one believes the others will do. We restrict attention to strongly symmetric, Markov equilibrium outcomes, in which each agent stays in the pool with probability one before the pool reaches the critical size and with some probability less than one if the state, i.e., the number of people in the pool, is beyond the critical size. In equilibrium, each agent stays in the pool with the same probability, which decreases as the state increases.[6]

From the point of view of dynamics, which is the focus of our analysis, equilibrium behavior implies some interesting properties. First, as the state space is potentially infinite, there is a positive, if small, probability that the number of agents in the pool becomes arbitrarily large. Second, consider the ongoing "perturbation" of having one additional agent exogenously entering the pool in each period. At each state beyond the critical one, this leads to none, one, two, …, or all agents leaving the pool with positive probability. Thus, a small shock might eventually lead to a large effect. One question of interest is whether the empirical frequency of avalanches obeys a power law relationship. The answer, in the case of a single pool, seems to be no. The reason is that large events, in which many agents leave the pool simultaneously, are the results of many *iid* randomizations. Thus, even if the probability of exiting is high for each individual agent, the frequency of events in which a large number of agents exits decreases exponentially.

Intuitively, a power law is more likely to emerge in a model that involves an additional "local interaction" structure that is excluded in the one-pool example. To capture this effect, we study the model with n pools. All n pools are identical to the one-pool model but with the difference that there is a one-way externality between them. In each period, pool 1 opens first. After the exogenous inflow, agents in pool 1 decide whether to stay in the pool or exit. As before, exiting a pool implies leaving the game with a permanent future payoff of zero. Pool 2 opens next. After the exogenous inflow of one agent, agents in pool 2 decide whether to stay in their pool or exit.[7] After that, pool 3 opens, and so on. What connects two neighboring pools is that while the payoff to agents in pool 1 depends only on the number of agents staying in pool 1,

[2] See, for example, Bak (1996).

[3] In a social science context, interactions among agents at the micro level can "build up" towards power laws at the aggregate level. See, for example, LeBaron (2006) and Hommes (2006).

[4] See, for example, Bak et al. (1993), Scheinkman and Woodford (1994), Arenas et al. (2000, 2001, 2002) and Andergassen et al. (2006). Although the model we study here is very different, we will follow these papers in focusing on conditions under which the proposed model generates power laws.

[5] See, for example, Altman and Shimkin (1998) and Hassin and Haviv (2002, pp. 109–122). We thank Flavio Toxvaerd for pointing out this literature to us. Yet another interpretation relates to the "El Farol" dilemma, named after a popular bar in Santa Fe. See Arthur (1994) and Morgan et al. (1999). The story there concerns the decision of whether to visit a bar when it is individually optimal to do so only if the bar is not overcrowded.

[6] In what follows we shall ignore equilibria in which agents condition their exit decision on the number of periods they spend in the pool.

[7] Our results do not depend on the assumption that there is an exogenous flow of exactly one agent in every pool.

the payoff to agents in pool l, $l=2,\ldots,n$, has two components. It depends negatively on both the number of agents leaving pool $l-1$ and on the number of agents staying in pool l. We assume that the pools operate sequentially within each period and that the number of agents leaving pool $\{1,2,\ldots,l-1\}$ is known to agents in pool l before they make their choice about whether to stay or exit. The externality across neighboring pools is a tractable way of modeling local interaction effects. In the context of the processor use example, if a processor is overwhelmed by many users, resulting in lower speeds, some users might choose to employ another processor down the line instead, and so on. Our externality assumption provides a tractable way to model such local interaction effects across a sequence of "neighboring" processors. We numerically investigate two versions of the model. In one we assume that agents are myopic, while in the other we assume that they are forward looking. Our main finding is that local interaction effects lead to a power law relationship and that this does not appear to depend on the assumption about the agents' foresight.

The paper is organized as follows. Section 2 introduces the model with one pool, while Section 3 studies the model with many pools. In Section 4 we present our numerical results. A brief conclusion follows.

2. The model with one pool

Time is discrete and infinite. Our model involves an extensive form-game of complete information.[8] The set of agents is countably infinite. One possible interpretation of the model involves agents around one location, say, a pool. Initially, the pool is empty. One agent is randomly chosen and dropped into the pool in each period. Agents that are in the pool choose whether to stay (1) or exit (0). Once an agent exits, there is no possibility of re-entry. Formally, at each period of time, the set of agents is partitioned into three subsets: (i) agents that have never entered the pool; (ii) agents that are currently in the pool; and (iii) agents that have exited the pool in the past. We next describe the agents' action sets. Agents in subsets (i) and (iii) have a trivial (singleton) action set: $A=\{0\}$, that always results in a utility of 0. In each period, agents in subset (ii) have a binary action set $\{0,1\}$. Here, 1 stands for "stay in the pool," while 0 stands for "exit." We allow agents to mix, with $p_i \in \Delta(\{0,1\})$ denoting the probability with which agent i decides to stay in the pool in a given period. Agents' policies comprise infinite sequences of feasible actions. The state; i.e., the number of agents in the pool at the beginning of any given period prior to the exogenous entry is indicated by $x \in \{0,1,\ldots\}$. The period utility, as a function of x, is given by $u(x)$, where $u : Z_+ \to R$. In other words, u gives the payoff enjoyed in the current period by an agent that stays in the pool as a function of the number of other agents currently in the pool, x. We assume that u is strictly decreasing in x, that $u(1)>0$, and that there exists $\bar{x} \in Z_+$ such that $u(x)>0$, for all $x<\bar{x}$, $u(\bar{x}) \geq 0$, and $u(x)<0$, for all $x>\bar{x}$. We follow standard practice in game theory and assume that players discount future utility using the period discount factor $\delta \in [0,1)$.[9] Assuming that $\delta<1$ ensures that agents' payoffs remain bounded in what follows.

Hence, $p_i(x)$ denotes the probability that agent i decides to stay in the pool given that the state (the number of people in the pool) is x. Let $U(x,p)$ denote the value function for an agent. This function is formally defined below and it gives the expected current and future utility from staying in the pool when the current state is x and all other agents stay in the pool with probability p. The timing is as follows:

$$x_t \to \text{ period utility } \to \text{newcomer enters } \to \text{decisions } \to x_{t+1}. \tag{1}$$

In what follows, we will concentrate on strongly symmetric Markov equilibria; i.e., in equilibrium outcomes in which agents in any given period choose the same strategy in all histories and in which agents condition their strategies only on the current state, x. Clearly, for all states x such that $x<\bar{x}$ each agent will choose to stay, i.e., $p_i(x)=1$, for all $x<\bar{x}$, for all i. Notice that the expected utility to an agent depends on his beliefs about whether other agents will stay in the pool or not. These beliefs become relevant once $x>\bar{x}$. Each player's decision problem is characterized by a Dynamic Programming (Bellman) equation. The solution of this equation (the value function) is given by the expression below. The value function gives the equilibrium expected utility of staying in the pool as a function of the number of agents in the pool. In state x, this is given by[10]

$$U(x,p) = \left[\sum_{\tau=0}^{\bar{x}-x-1} \delta^{\tau} u(x+\tau+1)\right] + \delta^{\bar{x}-x} 0, \quad \text{if } x<\bar{x},$$

and

$$U(x,p) = \sum_{\tau=0}^{x-\bar{x}} \left\{ \binom{x}{\tau} p^{x-\tau}(1-p)^{\tau} u(x-\tau+1) \right\} + \sum_{\tau=x-\bar{x}+1}^{x} \left\{ \binom{x}{\tau} p^{x-\tau}(1-p)^{\tau} \sum_{\theta=0}^{\tau-x+\bar{x}-1} \delta^{\tau-x+\bar{x}-1-\theta} u(\bar{x}-\theta) \right\}, \quad \text{if } x \geq \bar{x}. \tag{2}$$

[8] For a formal study of such games see Fudenberg and Tirole (1991).

[9] For papers that use alternatives to discounting in environments similar to ours see, e.g., Sennott (1994), Federgruen (1978), Altman (1996) and Altman et al. (1997).

[10] As is standard in a game theory context, the value to an agent, U, depends on the strategies of the other agents as summarized by p. This is because these mixing probabilities determine the (current and future) number of agents in the pool, hence, the agent's current and future utilities, u. Since we only study symmetric equilibria, we will drop the index i in what follows. In addition, to economize in notation, we suppress the dependence of p on x.

The above expression is central in our paper and it is worth describing it in some detail here. The first line in (2) captures the fact that there is no exit before the state $\bar{x}$ is reached. In that case, all agents choose $p(x)=1$ and the size of the pool increases by the introduction of one newcomer per period. The second term in (2) captures the fact that, once $\bar{x}$ is eventually reached, agents will mix, staying in the pool with equal probability $p(x)<1$. As is standard, for a mixed strategy equilibrium we need that the agents mix with a probability that makes them indifferent between staying and leaving the pool. Thus, since exiting implies a continuation payoff of 0, staying in the pool must give the same expected value. In other words, we have that $U(x,p)=0$, if $x \geq \bar{x}$.

To summarize, agents stay in the pool with probability 1 for any $x<\bar{x}$, while they randomize choosing the mixed strategy $p(x)$ consistent with equilibrium for any $x \geq \bar{x}$. As an example, consider the case where $\bar{x}=3$ and $x=1$. Then, the above expression gives $U(x,p)=u(2)+\delta u(3)+\delta^2 0$. In other words, the agent already in the pool stays in and, joined by a newcomer, enjoys an immediate utility of $u(2)$. The two of them also stay and, joined by a newcomer, they each subsequently enjoy a discounted utility of $\delta u(3)$. Since the payoff of having four agents in the pool (three plus a newcomer) is assumed to be negative, agents are going to randomize when $\bar{x} \geq 3$. As is standard, the conditions for a mixed strategy equilibrium require that $p(4)$ is chosen so that the value function satisfies $U(4,p)=0$. Of course, as a result of the actual randomizations, it may happen that zero, one, two, three, or all four agents leave the pool. If, say, the realizations are such that everyone stays, in the next period agents will choose $p(5)$ so that $U(5,p)=0$, and so on.

Thus, solving for the equilibrium involves solving this system of equations for the unknown probabilities, p.[11] Having evaluated these probabilities, one can find the associated invariant distribution over pool sizes. In our simulations we adopt a fixed point approach. We jointly determine the mixing probabilities and the distribution of agents across states so that the mixing probabilities constitute best responses given the distribution and, at the same time, the best responses result in the anticipated distribution.

Our first goal is to establish the existence and uniqueness of a strongly symmetric Markov equilibrium. We have the following.[12]

Proposition 1. *There exists a unique strongly symmetric Markov equilibrium in which $p(x)=1$, for $x<\bar{x}$ and $p(x) \in (0,1)$, for $x \geq \bar{x}$.*

Proof. Clearly, $p(x)=1$, for $x<\bar{x}$. We need to demonstrate that for all $x \geq \bar{x}$, there exists a $p(x)$ such that $U(x,p)=0$. The payoff function is a polynomial of degree x and, therefore, continuous in p. Fix any $x \geq \bar{x}$ and assume that all agents but one leave the pool with probability 1. In that case, the value of the agent that stays in the pool is positive, i.e., $U(x,0)>0, \forall x \geq \bar{x}$. Similarly, for any $x \geq \bar{x}$, if all agents stay, the value of staying in the pool is negative, i.e., $U(x,1)<0, \forall x \geq \bar{x}$. Existence then follows from the Intermediate Value Theorem. Uniqueness follows by the strict monotonicity of the polynomial. □

Let K be the state of states, x, in which each agent's best response is to stay in the pool with probability 1 and let $P_{iK}^{(n)}$ denote the n-step transition probability from state i to a state belonging to set K.

Lemma 1. *There exists a $\underline{P}>0$ for which, for all n and i, there exists a $j \in K$ such that $P_{ij}^{(n)} \geq \underline{P}$.*

Proof. Recall that in the symmetric equilibrium $U(x,p)=0$, $\forall x \in K^c$. Therefore, since $u(x)<0$, for all $x \in K^c$, for the indifference condition to hold we need that, for all $i \in K^c$, there exists a $j \in K$ such that $P_{ij}>0$. Define $P^* = \min_{i \in K^c, j \in K}\{P_{ij}\}$. Since $u(x)$ is decreasing in x, we have that P^* is well defined and strictly positive. Further, we have that

$$P_{iK}^{(n)} = P[x_{m+n} \in K | x_m = i] = \sum_{j \in K} P_{ij}^{(n)} = \sum_{j \in K} \sum_{k_1 \ldots k_{n-1}} [P_{ik_1} P_{k_1 k_2} \ldots P_{k_{n-1} j}] = \sum_{k_1 \ldots k_{n-1}} \left\{ [P_{ik_1} P_{k_1 k_2} \ldots P_{k_{n-2} k_{n-1}}] \sum_{j \in K} P_{k_{n-1} j} \right\} \geq P^*. \tag{3}$$

The last inequality in the above expression follows since, for all k, $\sum_{j \in K} P_{kj} \geq P^*$ and, in addition, $\sum_{k_1 \ldots k_{n-1}} [P_{ik_1} P_{k_1 k_2} \ldots P_{k_{n-2} k_{n-1}}] = 1$. Hence, for all n and $i \in K^c$, there exists a $j \in K$ such that $P_{ij}^{(n)} \geq (1/\#(K))P^*$. To complete the proof, let $\underline{P} = (1/\#(K))P^*$. □

Proposition 2. *There exists a unique, stationary measure of the size of the pool.*

Proof. Let x_i stand for the state where there are i agents in the pool at the beginning of the period. Clearly, the range of $\{x_i\}$ is contained in Z_+. It is straightforward to verify that the countable state Markov chain associated with the process is irreducible and aperiodic. Let K be the set of states associated with a positive payoff (that is, $K=\{x : x \leq \bar{x}\}$) and let K^c be the

[11] In principle, this involves an infinite number of equations/unknowns since the state space is not finite. However, as the number of agents increases, the number of agents that exits the pool as a result of the randomizations in the mixed strategy equilibrium exceeds 1, the exogenous pool entry. This feature comes very handy in our simulations since it effectively allows us to deal with a bounded state space.

[12] Analogs of the Propositions in the remainder of this Section can be derived for any finite number of pools. While the intuition of the many-pool case is the same, the notation is heavily involved. Hence, we decided to include the details of the $n=1$ case in the text. The proofs for $n>1$ are available from the authors.

complement of K. Finally, let $P_{ij}^{(n)}$ give the probability of a transition from state i to state j after n steps. Theorem 8.7 in Billingsley (1979) can be used to establish that if a stationary distribution does not exist in our setup, then $\lim_{n\to\infty} P_{ij}^{(n)} = 0$, for all i and j. Lemma 1, however, implies that $\lim_{n\to\infty} P_{ij}^{(n)} > 0$, for at least some i and j. We conclude that a stationary distribution must exist. □

3. A numerical example

Here we consider simulations of the one-pool model based on a particular form for the utility function

$$u(x) = \frac{a}{(1+\exp(a(x-\bar{x})))} - b, \tag{4}$$

where a=10 and b=5, and $\bar{x} = 5$. This "inverse sigmoidal" function, illustrated in Fig. 1, was chosen over a step function, for the continuity it provides.

We set $\delta = 0.9$ ("forward looking") or $\delta = 0$ ("myopic"). For each case, we set $p(x)$=1 for $x < \bar{x}$ but for all $x \geq \bar{x}$, we substitute the above expression for $u(x)$ and solve for $p(x)$ by setting $U(x,p)$ equal to zero. Using the resulting values for $p(x)$, we conduct a simulation exercise in which one agent exogenously enters the pool each period and agents in the pool use the calculated probabilities to determine whether to remain in the pool or not. The outflow from the pool (i.e., the number of agents choosing to leave) in period t, if any, is our measure of the "avalanche" size for that period. The simulation is conducted for 1 million periods. Fig. 2 illustrates the log frequency of avalanche sizes, log($f(s)$), (vertical scale) against the log of their sizes, log(s) (horizontal scale) for both the myopic and rational cases.

We found that this picture is qualitatively consistent with a wide variety of specifications for $u(x)$, e.g., decreasing linear functions of x and of the discount factor, δ. While a straight line in the log–log plot would indicate a power law, the above relation suggests that the decay in avalanche frequencies is exponential in their size. In other words, the on-going shock of adding one agent at a time is not able to produce large effects with a high frequency.

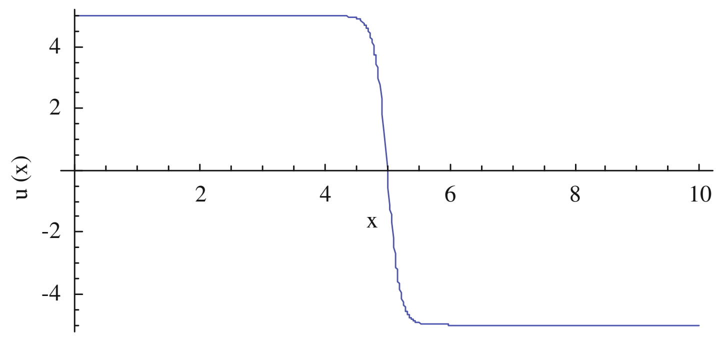

Fig. 1. Illustration of example utility function.

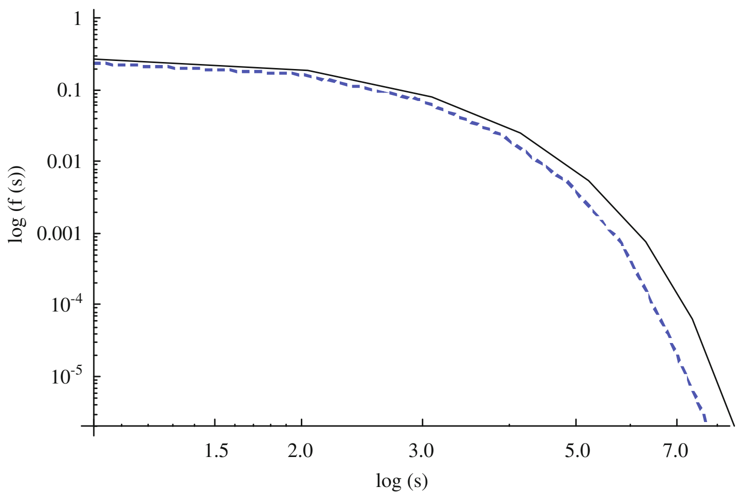

Fig. 2. Log of avalanche frequency, $f(s)$, against log of avalanche size, s (number leaving pool). Dashed line is $\delta = 0$ (myopic) case, solid line is $\delta = 0.9$, (forward-looking) case.

Table 1
Results from an OLS regression, $\log(f(s)) = \log\alpha + \beta\log(s)$, one pool case, $\delta = 0$ or 0.9.

1-pool δ value	$\hat{\log(\alpha)}$ (st. err.)	$\hat{\beta}$ (st. err.)	No. obs.	R^2	$R^2_{0.05}$
$\delta = 0$	6.47 (0.80)	−4.98 (1.24)	10	0.728	0.806
$\delta = 0.9$	6.50 (0.81)	−5.05 (1.13)	10	0.725	0.806

To get further at this issue, suppose the long-run frequency of avalanche sizes s, in the one-pool model did obey a power law of the form:

$$f(s) = \alpha s^{\beta},$$

over some range $(0,\bar{s})$. Taking logarithms, we have

$$\log(f(s)) = \log(\alpha) + \beta\log(s).$$

To assess whether this log–log relationship is linear, as would be the case if the frequency distribution obeyed a power law, we can run an ordinary least squares (OLS) regression of $\log(f(s))$ on a constant and $\log(s)$. Specifically we form the hypothesis:

H_0: the frequency distribution of avalanche sizes obeys a power law.
H_1: the frequency distribution of avalanche sizes does not obey a power law.

One approach would be to simply look at how close the OLS, R^2 goodness-of-fit measure is equal to 1, indicating a perfect linear relationship. But it is possible to go even further, as Gaudoin et al. (2003) show that the distribution of the R^2 statistic under H_0 is independent of α and β. They provide a table of critical values, R^2_θ such that one may reject H_0 at significance level θ if $R^2 < R^2_\theta$. In testing H_0 for the one pool case (and later for the multiple pool cases) we make use of this goodness-of-fit test.

OLS regression results for the one-pool model are reported in Table 1. We see that for both the $\delta = 0$ and 0.9 cases, we may reject (at the 5% level of significance) H_0, that the relationship between $f(s)$ and s obeys a power law. The explanation for this finding is simple. Consistent with the mixed strategy equilibrium, a large number of agents exiting the pool is the result of independent randomizations. Given the payoff structure we assume, this implies a very small probability that a large number of agents will independently decide to exit. This is true in both the case where agents are forward looking ($\delta > 0$) and where they are myopic ($\delta = 0$). Local interaction effects offer one way to overcome this effect. We introduce such effects next in the context of a model with multiple pools.

4. The model with *n* pools

The main setup is the same as in the one-pool model but we now assume that there are n pools. All pools are identical to the one-pool model but with the difference that there is a local one-way externality. In each period, pool 1 opens first. After the exogenous inflow, agents in pool 1 decide whether to stay or exit. Pool 2 opens next and agents in pool 2 decide whether to stay in their pool or exit. Exiting pool 2 also implies leaving the game with a permanent future payoff of zero. After that pool 3 opens, and so on. As before, the payoff to agents outside the pool is normalized to be 0.

What connects two neighboring pools is that we assume that the payoff to agents in pool l, l=2,…,n, depends negatively on two things: the number of agents leaving pool $l-1$ and on the number of agents staying in pool l. The pools operate sequentially within each period and the number of agents leaving pool $\{1,2,\ldots,l-1\}$ is known to agents in pool l before they make their choice about whether to stay or exit.[13]

Here we let x_t^l be the number of agents in pool l, $l \in \{1,2,\ldots,n\}$ at the beginning of period t. Let z_t^{l-1} be the number of agents leaving pool $l-1$ in period t. Thus, $x^{l-1}+1-z^{l-1}$ denotes the number of agents staying in pool $l-1$. As before, the discount factor is $\delta \in (0,1)$. The period payoff functions are given by $u^1(x^1)$ and $u^l(x^l,z^{l-1})$, for all l=2,…,n. The first payoff function, $u^1 : \mathbb{Z}_+ \to \mathbb{R}$ is assumed to be decreasing in x^1. Like before, we assume that $u^1(1) > 0$ and that there exists $\bar{x}^1 \in \mathbb{Z}_+$ such that $u^1(x^1) > 0$, for all $x^1 < \bar{x}^1$, $u^1(\bar{x}^1) \geq 0$, and $u^1(x^1) < 0$, for all $x^1 > \bar{x}^1$. The payoff functions for the other pools, $u^l : \mathbb{Z}_+ \times \mathbb{Z}_+ \to \mathbb{R}$ are assumed to be decreasing in both z^{l-1} and x^l. We assume that, for all z^{l-1}, $u^l(1,z^{l-1}) > 0$ and that there exists $\bar{x}^l(z^{l-1}) \in \mathbb{Z}_+$ such that $u^l(x^l,z^{l-1}) > 0$, if $x^l < \bar{x}^l(z^{l-1})$, $u^l(x^l,z^{l-1}) \geq 0$, if $x^l = \bar{x}^l(z^{l-1})$, and $u^l(x^l;z^{l-1}) < 0$, if $x^l > \bar{x}^l(z^{l-1})$. Finally, we assume that $\lim_{x^l \to \infty} u^l(x^l,z^{l-1}) = -\infty$, for all l=1,…,n.[14]

[13] One way to think about the negative externality is as a reduced-form way of capturing a situation in which the outflow in pool $l-1$ contains information about some change in fundamentals that, everything less being equal, makes staying in pool l less desirable.
[14] We ensure that these assumptions hold throughout our simulations.

Let $p_1(x^1)$ be the probability that an agent decides to stay in pool 1 if the state is $(x_t^1,\ldots,x_t^l)$, and define $(\mathbf{x}^l;\mathbf{z}^{l-1})\equiv(x^1,z^1;x^2,z^2;\ldots,x^{l-1},z^{l-1};x^l+1)$. Let $p_l(\mathbf{x}^l;\mathbf{z}^{l-1})$ be the probability that an agent decides to stay in pool l, l=2,…,n, in state $(\mathbf{x}^l;\mathbf{z}^{l-1})$. We concentrate on symmetric equilibria in which each agent stays in pool l with equal probability. Finally, let $p=\{p_l\}_{l=1,\ldots,n}$. The expected payoff of staying in pool 1 as a function of the number of agents in that pool, x^1, is identical to the one pool case in the previous section.

Turning to pool l, l=2,…,n, let $v^l(\mathbf{x}^l;\mathbf{z}^{l-1})$ denote the value function of an agent who decides to stay in pool l given that the state is $(\mathbf{x}^l;\mathbf{z}^{l-1})$. This is given by

$$v^l(\mathbf{x}^l;\mathbf{z}^{l-1})=\left\{\sum_{k=0}^{x^l}\binom{x}{l}k(p_l)^k(1-p_l)^{x^l-k}u^l(z^{l-1},k+1)+\delta Ev^l(\mathbf{x}'^l;\mathbf{z}'^{l-1})\right\}. \tag{5}$$

The values of $\mathbf{x}'^l$ and $\mathbf{z}'^{l-1}$ in the expression for $Ev^l(\mathbf{x}'^l;\mathbf{z}'^{l-1})$ are given by one of the following expressions. If $x^m-z^m+1<\overline{x}^m$; m=1,…,$l-1$, then

$$(\mathbf{x}'^l;\mathbf{z}'^{l-1})=(x^1-z^1+1,0;x^2-z^2+1,0;\ldots,x^{l-1}-z^{l-1}+1,0;k+1). \tag{6}$$

On the other hand, if $x^m-z^m+1\geq\overline{x}^m$, for some m, then $(\mathbf{x}'^l,\mathbf{z}'^{l-1})$ is given by

$$\sum_{k^l=0}^{x^l}\binom{x^l}{k^l}p^l(\mathbf{x}^l,\mathbf{z}^{l-1})^{k^l}(1-p^l(\mathbf{x}^l;\mathbf{z}^{l-1}))^{x^l-k^l}$$
$$\sum_{k^1=0}^{x^1+1}\binom{x^1+1}{k^1}p^1(x^1)^{k^1}(1-p^1(x^1))^{x^1+1-k^1}$$
$$\sum_{k^2=0}^{x^2+1}\cdots\sum_{k^{l-1}=0}^{x^{l-1}+1}\prod_{j=2}^{l-1}\binom{x^j+1}{k^j}p^j(\bullet)^{k^j}(1-p^j(\bullet))^{x^j+1-k^j}(k^1,x^1+1-k^1;\ldots,k^{l-1},x^{l-1}+1-k^{l-1};k^l+1), \tag{7}$$

where $p^j(\bullet)\equiv p^j(k^1,x^1+1-k^1,\ldots,k^{j-1},x^{j-1}+1-k^{j-1},x^j)$. The respective probabilities satisfy the following best response conditions:

$$p_1(x^1)\begin{cases}=1 & \text{if } v^1(x^1)>0,\\ \in\Delta\{0,1\} & \text{if } v^1(x^1)=0,\\ =0 & \text{if } v^1(x^1)<0,\end{cases} \tag{8}$$

and, for l=2,…,n,

$$p_l(\mathbf{x}^l;\mathbf{z}^{l-1})\begin{cases}=1 & \text{if } v^l(\mathbf{x}^l;\mathbf{z}^{l-1})>0,\\ \in\Delta\{0,1\} & \text{if } v^l(\mathbf{x}^l;\mathbf{z}^{l-1})=0,\\ =0 & \text{if } v^l(\mathbf{x}^l;\mathbf{z}^{l-1})<0.\end{cases} \tag{9}$$

By an analogous argument to that in the one-pool case we can demonstrate the following.

Proposition 3. *A strongly symmetric Markov equilibrium exists in which (i)* $p_1(x^1)$=1, *for* $x^1<\overline{x}^1$, *and* $p_1(x^1)\in(0,1)$, *for* $x^1\geq\overline{x}^1$; *(ii) for all* $l\in\{2,\ldots,n\}$, *and for all* $(\mathbf{x}^l;\mathbf{z}^{l-1})$, *there exists* $\tilde{x}^l$ *such that* $p_l(\mathbf{x}^l;\mathbf{z}^{l-1})$=1, *for* $x^l<\tilde{x}^l$, *and* $p_l(\mathbf{x}^l;\mathbf{z}^{l-1})\in(0,1)$, *for* $x^l\geq\tilde{x}^l$.

One can also prove an analog of Proposition 2 for the many pool setup. In the next Section we explore other properties of the model by performing numerical simulations. We are particularly interested in whether the multiple-pool version can give rise to a power law, and in whether this depends on the value of the discount factor, δ.

5. A numerical example with *n* pools

To illustrate the effect of introducing multiple pools and the resulting local interaction structure, we consider the same parametric utility function (4) as before. For pool 1, the utility function is exactly the same as in the earlier example. For all pools $l>1$, the utility function $u^l(y)$ has the same functional form as (4) but for these pools we assume, $y=x^l+\lambda z^{l-1}$. The parameter λ measures the relative strength of the local interaction effect. Both λ and the number of pools, n, are varied in the simulation exercise reported in this section. We found that our results for the n-pool case were largely insensitive to other model parameters, e.g., the parameters (a, b) of the utility function (4) or the discount factor δ; we therefore set the utility function parameter values equal to those used in the simulation exercise for the one-pool case (a=10, b=5, $\overline{x}=5$) and we set $\delta=0$ (myopic case only).

We repeat the exercise conducted earlier for the one-pool case, but with a total of one, two or three pools and in the latter two cases, varying λ from 0.0 to 1.5 in increments of 0.5. For each of these four cases, we ran a simulation of 1 million periods, where in each period, one agent enters each pool, and the agents in each pool must decide whether to remain or not using the appropriate probabilities for the state, beginning with the players in pool 1 and then moving sequentially to the players in the other pools. In the case of multiple pools, we use the sum of the outflows from all l pools in period t as our measure of the avalanche size in period t, i.e., $s_t=\sum_l z_t^l$.

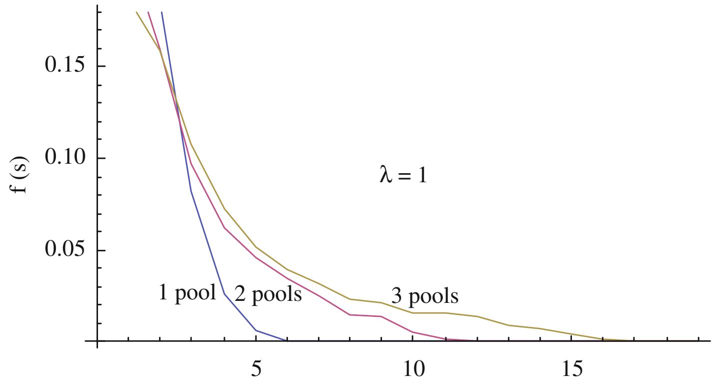

Fig. 3. Frequency of avalanche sizes $f(s)$, one-, two-, three-pool cases, 1 million simulations, $\lambda = 1$.

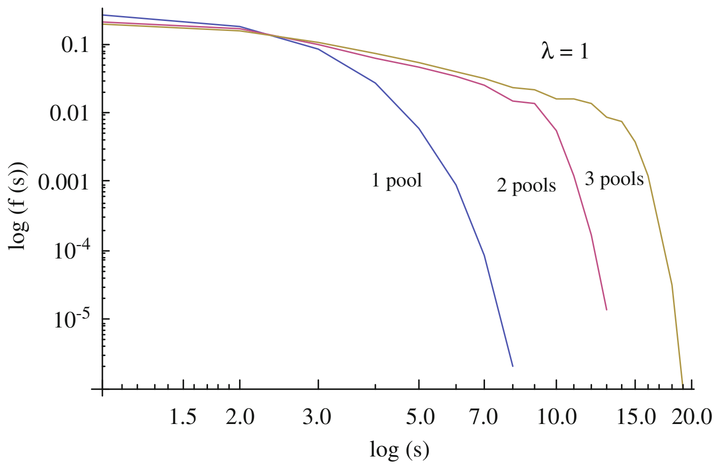

Fig. 4. Log($f(s)$) against log(s), one-, two-, three-pool cases, 1 million simulations, $\lambda = 1$.

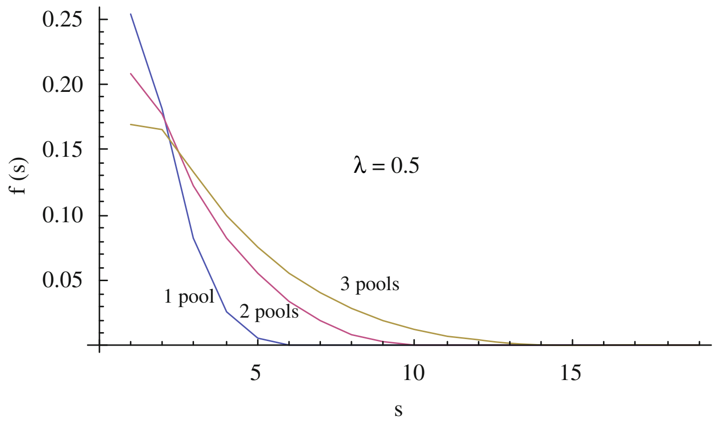

Fig. 5. Frequency of avalanche sizes $f(s)$, one-, two-, three-pool cases, 1 million simulations, $\lambda = 0.5$.

Figs. 3 and 4 report results from the numerical experiments where we have set $\lambda = 1$. Fig. 3 shows the frequency of avalanche sizes for the one-, two-, and three-pool cases and Fig. 4 shows a log–log plot of these same data. For the one-pool environment, we have seen before that there is no evidence for a power-law relationship between $f(s)$ and s. However, as the number of pools increases, Fig. 3 reveals that, a least over some range of s, there is an approximately linear relationship between $\log(f(s))$ and $\log(s)$, up to some cut-off, $\bar{s}$, after which we observe exponential decay.[15] Notice that this linearity in the log–log relationship precisely corresponds to a flattening of the tail in the frequency of distribution of avalanche sizes.

The next several Figs. 5–8 show that such self-organized behavior requires not only that there be several pools, but also that the interaction effect, as captured by the parameter λ, be sufficiently large. As λ is decreased from 1 to 0.5 and finally

[15] As Jensen (1998) notes, such crossover to exponential decay above a certain avalanche size is common in finite systems.

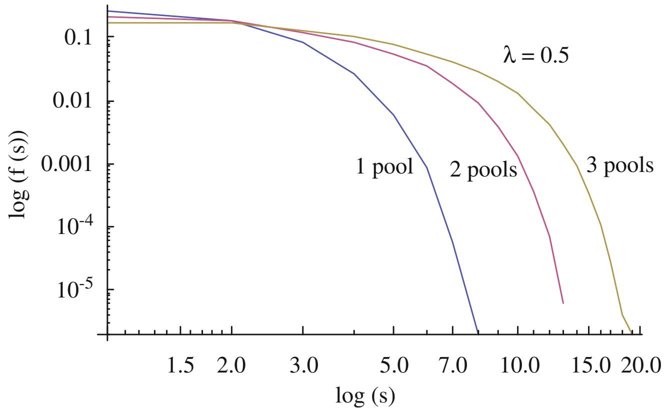

Fig. 6. Log($f(s)$) against log(s), one-, two-, three-pool cases, 1 million simulations, $\lambda = 0.5$.

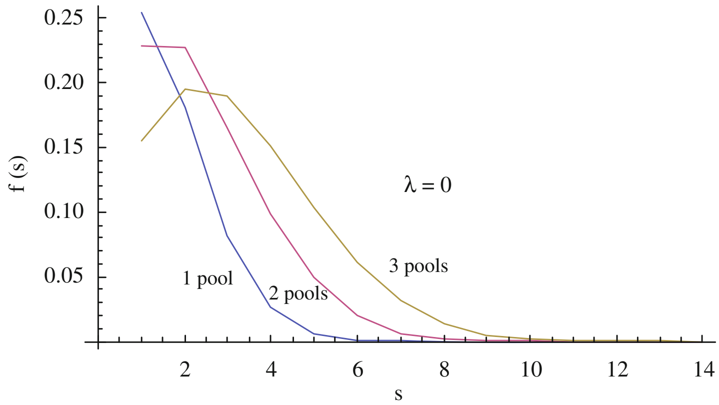

Fig. 7. Frequency of avalanche sizes $f(s)$, one-, two-, three-pool cases, 1 million simulations, $\lambda = 0$.

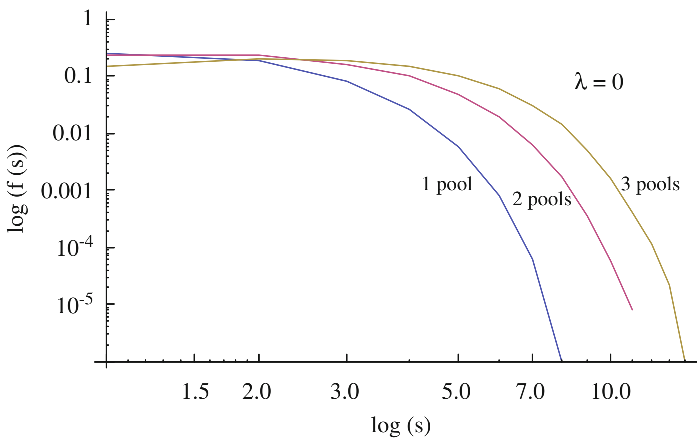

Fig. 8. Log($f(s)$) against log(s), one-, two-, three-pool cases, 1 million simulations, $\lambda = 0$.

to 0.0—so that interaction effects across pools are eventually eliminated—we see that the log–log plots for the multiple pool cases increasingly resemble that for the one-pool case, i.e., the power law signature of self-organized criticality gradually disappears.

Finally, we report results for the case where $\lambda = 1.5$ in Figs. 9 and 10. Values for λ in excess of 1 are not ruled out by our model and simply indicate that the outflow from pool $l-1$ has a greater impact on the utility of pool l participants than does the number of those remaining in pool l. The results for the $\lambda = 1.5$ case are similar to those found for the $\lambda = 1$ case; as the number of pools increases, there is some range of s, for which there is an approximately linear relationship between $\log(f(s))$ and $\log(s)$.

Table 2 repeats the regression analysis we performed for the 1-pool case (cf. Table 1) using the data from our simulation exercises as depicted in Figs. 3–10. For comparison with the 1-pool case we truncate our data sample to the first 10

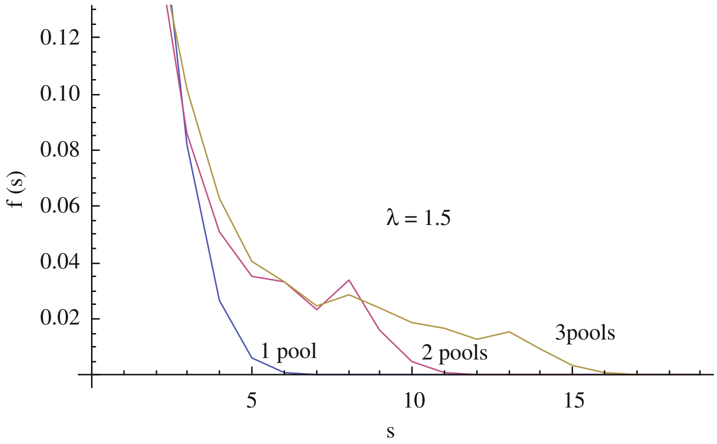

Fig. 9. Frequency of avalanche sizes $f(s)$, one-, two-, three-pool cases, 1 million simulations, $\lambda = 1.5$.

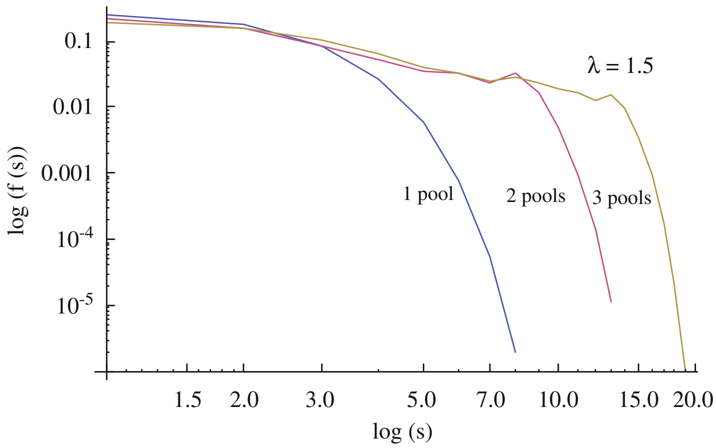

Fig. 10. Log($f(s)$) against log(s), one-, two-, three-pool cases, 1 million simulations, $\lambda = 1.5$.

Table 2
Results from an OLS regression, $\log(f(s)) = \log\alpha + \beta\log(s)$, two and three pool cases, $\lambda = 0$, 0.5, 1.0 and 1.5.

λ value	No. of pools	$\widehat{\log(\alpha)}$ (st. err.)	$\hat{\beta}$ (st. err.)	No. obs.	R^2	$R^2_{0.05}$
0.0	2	6.30 (0.56)	−3.30 (0.76)	10	0.693	0.806
0.0	3	5.80 (0.37)	−1.78 (0.51)	10	0.604	0.806
0.5	2	5.80 (0.29)	−2.02 (0.14)	10	0.755	0.806
0.5	3	5.50 (0.14)	−1.12 (0.19)	10	0.804	0.806
1.0	2	5.57 (0.14)	−1.47 (0.19)	10	0.882	0.806
1.0	3	5.45 (0.07)	−1.13 (0.11)	10	0.931	0.806
1.5	2	5.51 (0.15)	−1.39 (0.21)	10	0.843	0.806
1.5	3	5.42 (0.06)	−1.11 (0.09)	10	0.950	0.806

observations (avalanche sizes, s). Table 2 confirms the impression given by Figs. 3–10 that, for the range of avalanche sizes considered ($s = 1, 2, \ldots, 10$) the relationship between $\log f(s)$ and $\log(s)$ becomes increasingly linear—as reflected in the higher R^2 statistic—as (1) the number of pools is increased and (2) the value of λ is increased. Indeed, when $\lambda \geq 1$ and there are 2 or 3 pools, we cannot reject the null hypothesis that the frequency distribution of avalanche sizes obeys a power law on the basis of the R^2 goodness-of-fit criterion.

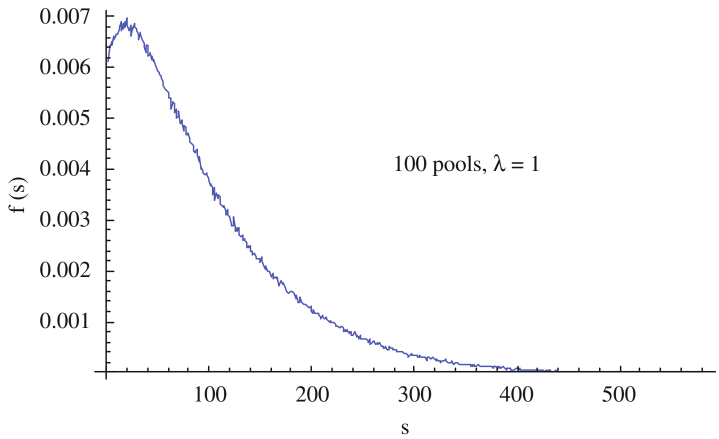

Fig. 11. Frequency of avalanche sizes $f(s)$, 100 pools, $\lambda = 1$, 1 million simulations.

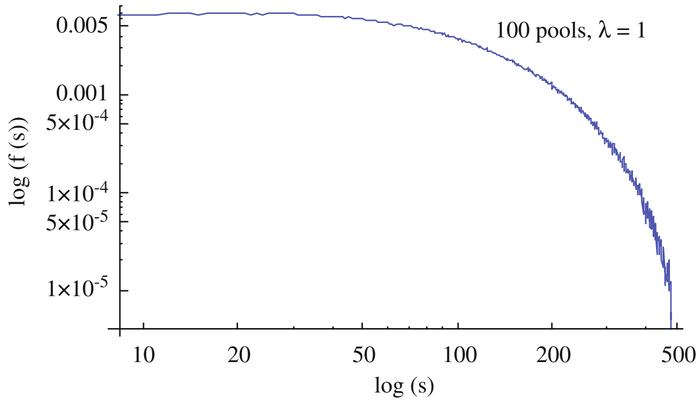

Fig. 12. Log($f(s)$) against log(s), 100 pools, $\lambda = 1$, 1 million simulations.

As a final exercise, we consider the case where the number of pools is an order of magnitude larger than we have considered so far. Specifically, we consider the case where $n=100$ and where $\lambda = 1$. Figs. 11 and 12 show the frequency of avalanche sizes and the log–log plot from 1 million simulations of the $n=100$ $\lambda = 1$ version of our model. We see in Fig. 11 that with 100 pools, the modal avalanche size shifts away from 1 (it is now 37), but the frequency distribution continues to be thickly right-tailed. The corresponding log–log plot in Fig. 12 reveals once again that for some range of avalanche sizes s, there is an approximately linear relationship between $\log(f(s))$ and $\log(s)$.

Using the first 100 avalanche sizes as observations (10 times our earlier sample size for the $n=1,2,3$ cases) we again run the regression of $\log(f(s))$ on $\log(s)$ and obtain an R^2 of 0.95. For a sample size of 100 observations, the critical value for Gaudoin et al.'s (2003) goodness-of-fit test, $R^2_{0.05}=0.93$. Thus we cannot reject H_0, that the relationship between $f(s)$ and s obeys a power law over the range studied in the 100-pool case.

6. Conclusion

We have demonstrated the existence and uniqueness of a strongly symmetric Markov equilibrium in a dynamic entry game. This equilibrium has agents remaining in the game with probability 1 before the number of participants reaches a certain threshold, and with some probability after the threshold is reached. We have explored the implications of this environment for the possibility of self organized criticality. Given our specification for the payoff structure, we have found evidence for power law regularities, indicative of SOC, only when the number of pools is sufficiently large and provided that there is a sufficiently strong local interaction effect between pools. This finding does not appear to depend on whether agents are myopic or forward looking. A main contribution of this paper is that we provide an explicit game-theoretic model of the mechanism through which SOC might arise in a social context. We believe that this mechanism can be useful in understanding related issues in dynamic setups that arise, for example, in oligopoly theory or in queueing systems.

SOC is hard to detect experimentally even in the controlled experiments of the natural sciences. An advantage of the proposed setup is that it is simple enough that it can be implemented in an experimental environment. In future work, we plan to study under what parameterizations human subjects will exhibit behavior consistent with SOC in such a controlled environment.

Acknowledgments

We thank Cars Hommes and two anonymous referees as well as Flavio Toxvaerd and audiences at the Society of Economic Dynamics meetings in Paris (2003), the Second World Congress of Game Theory in Marseille (2004), and the Santa Fe Institute for comments. The third author gratefully acknowledges support from the NSF through Grant SES 0517862. All errors are ours.

References

Altman, E., Shimkin, N., 1998. Individual equilibrium and learning in processor sharing systems. Operations Research 46, 776–784.
Altman, E., 1996. Non zero-sum stochastic games in admission, service and routing control in queueing systems. Queueing Systems 23, 259–279.
Altman, E., Hordijk, A., Spieksma, F.M., 1997. Contraction conditions for average and $\alpha-$discount optimality in countable state Markov games with unbounded rewards. Mathematics of Operations Research 22, 588–618.
Andergassen, R., Nardini, F., Ricottilli, M., 2006. Innovation waves, self-organized criticality and technological convergence. Journal of Economic Behavior and Organization 61, 710–728.
Arenas, A., Díaz-Guilera, A., Pérez, C.J., Vega-Redondo, F., 2002. Self-organized criticality in evolutionary systems with local interactions. Journal of Economic Dynamics and Control 26, 2115–2142.
Arenas, A., Díaz-Guilera, A., Guardiola, X., Llas, M., Oron, G., Pérez, C.J., Vega-Redondo, F., 2001. New results in a self-organized model of technological evolution. Advances in Complex Systems 4, 89–100.
Arenas, A., Díaz-Guilera, A., Pérez, C.J., Vega-Redondo, F., 2000. Self-organized evolution in socio-economic environments. Physical Review E 61, 3466–3469.
Arthur, B., 1994. Inductive reasoning and bounded rationality. American Economic Association Papers and Proceedings 84, 406–411.
Bak, P., 1996. How Nature Works: The Science of Self-Organized Criticality. Springer, New York.
Bak, P., Chen, K., 1991. Self-organized criticality. Scientific American, January, pp. 46–53.
Bak, P., Chen, K., Scheinkman, J., Woodford, M., 1993. Aggregate fluctuations from independent sectoral shocks: self-organized criticality in a model of production and inventory dynamics. Ricerche Economiche 47, 3–30.
Bak, P., Sneppen, K., 1993. Punctuated equilibrium and criticality in a simple model of evolution. Physical Review Letters 71, 4083–4086.
Bak, P., Tang, C., Wiesenfeld, K., 1988. Self-organized criticality. Physical Review A 38, 364–374.
Billingsley, P., 1979. Probability and Measure. Wiley, New York.
Federgruen, A., 1978. On *n*-person stochastic games with denumerable state space. Advances in Applied Probability 10, 452–471.
Fudenberg, D., Tirole, J., 1991. Game Theory. MIT Press, Cambridge, MA.
Gaudoin, O., Yang, B., Xie, M., 2003. A simple goodness-of-fit test for the power-law process, based on the Duane plot. IEEE Transactions on Reliability 52, 69–74.
Hassin, R., Haviv, M., 2002. To Queue or Not to Queue: Equilibrium Behavior in Queueing Systems. In: International Series in Operations Research Management Science, Springer, New York.
Hommes, C.H., 2006. Heterogeneous agent models in economics and finance. In: Tesfatsion, L., Judd, K.L. (Eds.), Handbook of Computational Economics, vol. 2. North-Holland, Amsterdam, pp. 1109–1186.
Jensen, H.J., 1998. Self-Organized Criticality: Emergent Complex Behavior in Physical and Biological Systems. Cambridge University Press, Cambridge, UK.
LeBaron, B., 2006. Agent-based computational finance. In: Tesfatsion, L., Judd, K.L. (Eds.), Handbook of Computational Economics, vol. 2. North-Holland, Amsterdam, pp. 1187–1233.
Morgan, D., Bell, A.M., Sethares, W.A., 1999. An Experimental Study of the El Farol Problem. University of Wisconsin-Madison, Working paper.
Scheinkman, J.A., Woodford, M., 1994. Self-organized criticality and economic fluctuations. American Economic Review Papers and Proceedings 84, 417–421.
Sennott, L.I., 1994. Nonzero-sum stochastic games with unbounded costs: discounted and average cost cases. Zeitschrift fur Operations Research 40, 145–162.