Memory and self-induced shocks in an evolutionary population competing for limited resources

Roland Kay* and Neil F. Johnson
Physics Department, Oxford University, Parks Road, Oxford, QX1 3PU, United Kingdom
(Received 21 December 2003; revised manuscript received 17 May 2004; published 5 November 2004)

We present a detailed discussion of the role played by memory, and the nature of self-induced shocks, in an evolutionary population competing for limited resources. Our study builds on a previously introduced multi-agent system [Phys. Rev. Lett. **82**, 3360 (1999)] which has attracted significant attention in the literature. This system exhibits self-segregation of the population based on the "gene" value p (where $0 \leqslant p \leqslant 1$), transitions to "frozen" populations as a function of the global resource level, and self-induced large changes which spontaneously arise as the dynamical system evolves. We find that the large, macroscopic self-induced shocks that arise are controlled by microscopic changes within extreme subgroups of the population (i.e., subgroups with "gene" values $p \sim 0$ and $p \sim 1$).

DOI: 10.1103/PhysRevE.70.056101 PACS number(s): 05.65.+b, 05.40.−a, 64.75.+g, 02.50.−r

I. INTRODUCTION

The dynamical behavior of a population of objects or "agents" (e.g., software or hardware modules, cellular organisms such as bacteria or viruses, human beings, animals) is of interest across a range of disciplines. Physics is arguably luckier than most disciplines in that the agents of interest (i.e., particles) do not adapt their behavior according to past failure, hence evolving new sets of rules as time progresses. Nor do the agents in question have any individual memory. Biological and social disciplines are not so lucky. Through a desire to develop a minimal model that could incorporate such features into a manageable yet nontrivial system, Arthur introduced the so-called "El Farol" bar problem, which concerns the repeated competition between bargoers to attend a popular bar with limited seating [1]. Challet and Zhang [2] subsequently introduced a binary version of this bar problem for the case where the amount of resource (e.g., number of seats) is just less than half the number of agents (e.g., possible attendees). This system is referred to as the minority game.

The minority game does not allow an agent to continuously evolve new strategies and hence explore the entire strategy space. The minority game is also essentially deterministic, apart from occasional coin tosses which are used to break ties in strategy scores. Furthermore, the resource level is set at just less than half the number of agents, so that there are always more losers than winners. To help overcome these limitations, Johnson *et al.* introduced a stochastic version of the minority game [3], which is subsequently referred to as the genetic model, in which an agent's strategy (characterized by a "gene" value p) can evolve indefinitely in time and is in principle allowed to access the entire space of strategies (i.e., all p values). The resulting genetic model has provoked much interest in the literature (for example, see Refs. [4–9]). As commented in the original paper of Johnson *et al.* [3] and confirmed by Burgos *et al.* [5,7], the self-segregation observed in the genetic model is insensitive to changes in an agent's memory length m.

In this paper we present a detailed discussion of the role of memory in the genetic model. We also explain the origin of the remarkable steplike structure in the global output time series as a function of the resource level, which was first observed by Johnson *et al.* in Ref. [4]. We then introduce (Sec. III) a variant of the genetic model in which the number of agents competing at a given time step is allowed to fluctuate. Because of the analogy with the grand canonical ensemble in physics, we shall refer to this model as the "grand canonical genetic model" (GCGM) By considering versions of the GCGM both with and without memory, we shall investigate the endogenous (i.e., self-induced) large changes that arise in the system. These large changes represent abrupt macroscopic "shocks," and occur with a greater probability than would be expected based on random walk statistics. We provide a detailed analysis of the mechanism that generates these large changes.

II. EFFECTS OF MEMORY IN THE GENETIC MODEL

Various papers [3–5,7] have made claims with regard to the role of memory in the genetic model. To date, though, no one has performed a detailed analysis of this problem. In this section we present such an analysis which involves comparing the behavior of the original model with that of a memoryless variant. The results presented here for the genetic model are reminiscent of earlier results for the minority game. In particular, Hart *et al.* [10] showed that a crowd-anticrowd theory which assumes *random* history provides a quantitative description of the time-averaged fluctuations in the minority game. Subsequently Cavagna [11] demonstrated numerically that the time-averaged fluctuations were indeed largely unaffected if the global history was replaced with randomly generated data.

Ceva and Burgos [5] explicitly investigated the role of memory in the genetic model; however, the results are restricted to a comparison of the gene value distributions in the minority case, in which the amount of resource is just less than half the number of agents. In contrast, in this section we shall treat the general case in which the amount of resource is unrestricted, and provide some theoretical analysis to ex-

*Electronic address: roland.kay@physics.ox.ac.uk

TABLE I. Parameters used to generate numerical data.

Number of agents	$N=501$	Memory length	$m=4$
Death score	$D=4$	Mutation range	$r=0.2$

plain the differences that we shall observe between the two models. We shall show in Secs. II G and II H that the observation that the gene value distributions are identical holds only in the special case considered in Ref. [5]. Furthermore, we shall show that the feedback introduced by the existence of memory does influence the behavior of the model in that it controls the time average of the prediction (Sec. II G) and introduces autocorrelations into the mean attendance time series (Sec. II I). By comparing and contrasting the memoryless genetic model with the original model, we shall be able to make some important observations about the memory's true significance. We shall also answer some important questions as to the extent to which memory is of benefit to the agents and the system as a whole.

Before going on to describe the genetic model, we note that all of the numerical results presented in this paper were obtained using the model parameters listed in Table I unless otherwise stated. Similarly, all time averages are taken over the period $10\,000<t<60\,000$, the first 10 000 time steps being neglected to allow any transients due to the initial conditions to die away.

A. Original genetic model

In this section we present expressions for some of the most basic quantities in the original genetic model, which includes the memory. In Sec. II B we shall consider the equivalent expressions in the memoryless variant. We start with a brief summary of the genetic model. Fuller details are given in Ref. [3].

The genetic model consists of a population of agents who must decide at every time step between two possible choices which we will label -1 and $+1$. We shall refer to the decision of an agent i as its *action* $a_{i;t}$. Each agent is defined by a *gene value* $p_{i;t}$, which can take any value $0\leqslant p_{i;t}\leqslant 1$. At each time step the model makes a *prediction* h_t of the outcome of the time step available to all of the agents. Each agent chooses its action to be equal ($a_{i;t}=+h_t$) or opposite ($a_{i;t}=-h_t$) to this prediction with probability $p_{i;t}$ or $1-p_{i;t}$, respectively. h_t is calculated based on a global memory that the model maintains of the outcomes of the previous m time steps and the assumption that patterns that have occurred in the time series of these outcomes in the past will recur in the future. The parameter m is known as the *memory length*. For example, if $m=2$ then at time t the memory might contain the following entries: $(-1,-1;+1)$, $(-1,+1;-1)$, $(+1,-1;-1)$, and $(+1,+1;+1)$. This signifies that the last time the outcome of two consecutive time steps was -1 the outcome of the next time step was $+1$. Similarly the last time the sequence $-1+1$ occurred, it was followed by a -1. If at time t the outcome of the previous two time steps was $+1$ and then -1 then, in this case, $h_t=-1$. The outcome of each time step is determined based on the actions of all of the agents. From now on we shall refer to the outcome as the *global action* at time t, A_t.

Agent gene values are not constant with time. Each agent maintains a record of its *score*, $s_{i;t}$, which determines when it changes its gene value. At every time step, $s_{i;t}$ increases by one unit if $a_{i;t}=+A_t$ and decreases otherwise. If $s_{i;t}=-D$ then the agent *mutates*. The parameter D is known as the *death score*. When an agent mutates it chooses a new gene value at random from a range of values of width $2r$ centered on the old gene value. The parameter r is known as the *mutation range*.

There are two opposite definitions of the global action A_t used in the literature. For example, Ref. [4] defines A_t, by analogy with Zhang and Arthur's bar model, to be the state of the bar at time t. Thus, $A_t=+1$ would denote an *overcrowded* bar and the optimal action of each agent would be to stay at home (i.e. $a_{i;t}=-1$). However, in this paper we shall adopt the convention of Ref. [3] whereby A_t represents the optimal decision of each agent at time t. As defined in Ref. [4], the global action A_t is given, in terms of a model parameter l which can take values $0\leqslant l\leqslant 1$, by

$$A_t=\begin{cases}+1, & n_t^{+1}\leqslant Nl,\\ -1, & n_t^{+1}>Nl,\end{cases} \tag{1}$$

where

$$n_t^{+1}=\frac{1}{2}\left(\sum_i a_{i;t}+N\right). \tag{2}$$

In other words, n_t^{+1} is the number of agents for which $a_{i;t}=+1$. We shall refer to l as the *resource level*.

Let $\{p_{i;t}\}_i$ denote the set of values of $p_{i;t}$ for all agents at time t. As Lo *et al.* state [12], the ensemble average number of agents following the prediction $\langle n_t^{+h_t}\rangle$ is given by

$$\langle n_t^{+h_t}\rangle=N\overline{p_t} \text{ where } \overline{p_t}=\frac{1}{N}\sum_{i=1}^{N} p_{i;t}. \tag{3}$$

In equilibrium, where the population evolves such that $\overline{p_t}$ is approximately constant, the ensemble average $\langle n_t^{+h_t}\rangle$ and the time average $\langle n_t^{+h_t}\rangle_t$ will coincide. Thus,

$$\langle n_t^{+h_t}\rangle_t\approx N\langle\overline{p_t}\rangle_t. \tag{4}$$

B. Memoryless genetic model

In this section we introduce a memoryless variant of the original genetic model. In contrast to the original model described above, where the gene value of the ith agent $p_{i;t}$ gives the probability of it choosing to follow the prediction ($a_{i;t}=+h_t$), in the memoryless model $p_{i;t}$ gives the probability that $a_{i;t}=+1$ directly. With this modification the prediction h_t and hence the global memory that produce it become redundant and can be removed from consideration. The agents in this variant are memoryless, by which we mean that their actions $a_{i;t}$ at time t are independent of the state of the model at earlier times. $a_{i;t}$ is dependent only on $p_{i;t}$. An equivalent way of considering this is to take $h_t=+1\,\forall\, t$. The global action A_t and the number of agents attending the bar n_t^{+1} are

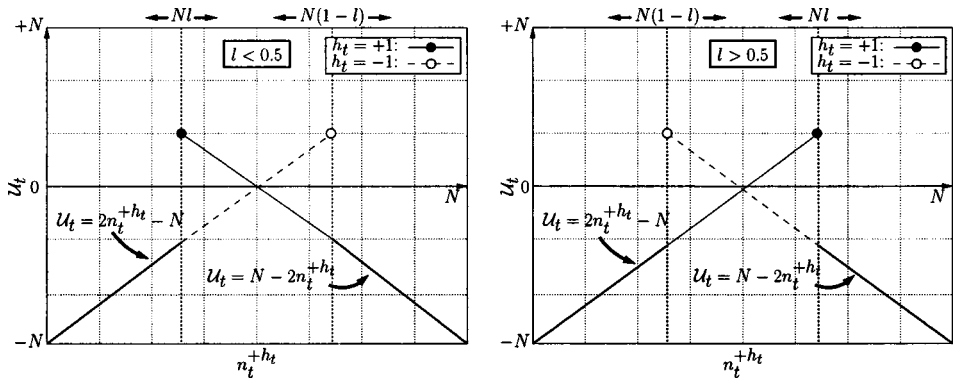

FIG. 1. $\mathcal{U}_t$ as a function of $n_t^{+h_t}$ in the original genetic model with memory. The dashed lines indicate $\mathcal{U}_t$ when $h_t=-1$, the narrow lines $\mathcal{U}_t$ when $h_t=+1$. The bold lines represent those part of the lines defined by $\mathcal{U}_t$ that are invariant under $h_t \rightarrow -h_t$. The black and white circles represent the values of $n_t^{+h_t}$ at which $\mathcal{U}_t$ is a maximum when $h_t=+1$ and -1, respectively.

unchanged and so are given by Eqs. (1) and (2) as before. However, Eq. (4) becomes

$$\langle n_t^{+1} \rangle_t \approx N \langle \overline{p_t} \rangle_t. \tag{5}$$

C. Comparison of the performance of the original and memoryless models

We shall now compare the performance of the two models. In order to quantify performance we define $\mathcal{U}_t$ to be the total number of points scored by all of the agents at time t. Therefore we are considering the performance of the system as a whole, rather than that of individual agents. If we consider the models to be analogous to an economic system then the question that we are investigating becomes to what extent the agents in this system can exploit the potential wealth available to them as a population. Note that in Ref. [6] Burgos *et al.* treat the memoryless genetic model in terms of a cost function given by the second moment of n_t^{+1} with respect to Nl. However, this cost function is symmetric in that it assigns an equal cost to deviations of n_t^{+1} from Nl of opposite signs. As we shall see, $\mathcal{U}_t$ is not symmetric about Nl and hence can distinguish between positive and negative deviations.

1. Original genetic model

First we shall derive expressions for $\mathcal{U}_t$ in the model with memory. Later we will see how these expressions are modified in the absence of the memory. From Eq. (1) the condition that $A_t=+1$ is

$$n_t^{+h_t} \leqslant Nl \quad \text{if } h_t = +1,$$

$$n_t^{+h_t} \geqslant N(1-l) \quad \text{if } h_t = -1. \tag{6}$$

Now consider the total number of points scored by the agents, $\mathcal{U}_t$. Agents for which $a_{i;t}=+A_t$ will gain one point whereas agents for which $a_{i;t}=-A_t$ will lose. If $h_t=+1$ and $A_t=+1$ it will be the $n_t^{+h_t}$ agents who choose to follow h_t who will gain. If $A_t=-1$, then the $N-n_t^{+h_t}$ agents who choose to refute h_t will gain. Thus, using Eq. (6), we have, for $h_t=+1$,

$$\mathcal{U}_t(n_t^{+h_t}) = \begin{cases} 2n_t^{+h_t} - N & \text{if } n_t^{+h_t} \leqslant Nl, \\ -2n_t^{+h_t} + N & \text{if } n_t^{+h_t} > Nl. \end{cases} \tag{7}$$

When $h_t=-1$ and $A_t=+1$ it will be the $N-n_t^{+h_t}$ agents who choose to refute the prediction who will gain, and vice versa for $A_t=-1$. Thus, for $h_t=-1$ the above expression becomes,

$$\mathcal{U}_t(n_t^{+h_t}) = \begin{cases} -2n_t^{+h_t} + N & \text{if } n_t^{+h_t} \geqslant N(1-l), \\ 2n_t^{+h_t} - N & \text{if } n_t^{+h_t} < N(1-l). \end{cases} \tag{8}$$

The expressions in Eqs. (7) and (8) are plotted in Fig. 1. The black and white circles represent the value of $n_t^{+h_t}$ at which $\mathcal{U}_t$ is a maximum for $h_t=+1$ and $h_t=-1$, respectively. We shall call the value of $n_t^{+h_t}$ at which $\mathcal{U}_t$ is a maximum the *optimal* value and denote it by $n_{t,\mathrm{opt}}^{+h_t}$. From Fig. 1 we can see that $n_{t,\mathrm{opt}}^{+h_t}$ is given as follows.

For $l<0.5$,

$$n_{t,\mathrm{opt}}^{+h_t} = \begin{cases} Nl+1 & \text{if } h_t = +1, \\ N(1-l)-1 & \text{if } h_t = -1. \end{cases} \tag{9}$$

For $l>0.5$,

$$n_{t,\mathrm{opt}}^{+h_t} = \begin{cases} Nl & \text{if } h_t = +1, \\ N(1-l) & \text{if } h_t = -1. \end{cases} \tag{10}$$

The most important feature of this equation to recognize is that, in general, there is no unique value of $n_{t,\mathrm{opt}}^{+h_t}$ independent of t. Only if $h_t=+1$ or $-1 \forall t$ would such a unique solution exist.

We note in passing that since $N=n_t^{+1}+n_t^{-1}$ the optimal value of n_t^{+1} in the original model is

$$n_{t,\mathrm{opt}}^{+1} = Nl. \tag{11}$$

In contrast to $n_{t,\mathrm{opt}}^{+h_t}$, the optimal value of n_t^{+1} is independent of t. Nevertheless, the most important quantity for the analysis that we present here is $n_{t,\mathrm{opt}}^{+h_t}$ in the case of the original model since, as we shall see later, it is the value of $n_t^{+h_t}$ that the agents can directly control and not that of n_t^{+1}.

Reference [4] demonstrated the existence of so-called *frozen* regimes which exist when l lies outside the region bounded by two critical values, which we shall label here l_{c1} and l_{c2}. These regimes were described as *quenched* by Bur-

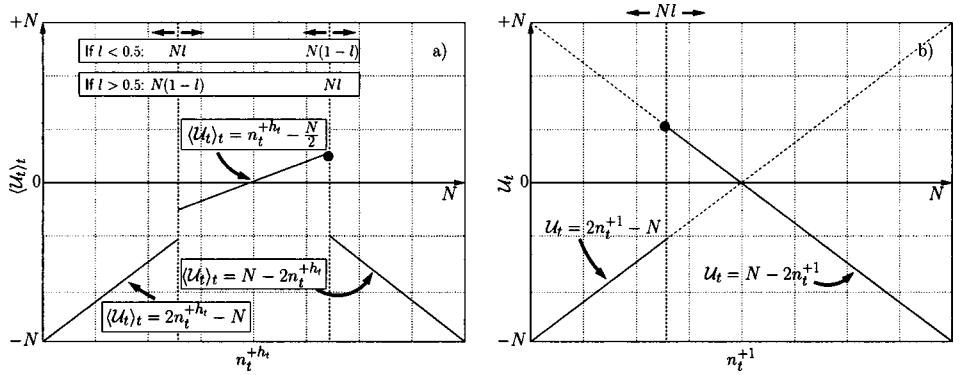

FIG. 2. $\langle \mathcal{U}_t(n_t^{+h_t})\rangle_t$ in (a) the original genetic model with memory and (b) the memoryless model.

gos *et al.* in Ref. [7]. The frozen regime obtains when $l<l_{c1}$ or $l>l_{c2}$. The behavior of the genetic model in these regimes is well understood [4,7]; therefore we shall restrict ourselves to a consideration of the dynamic regime. In what follows, unless otherwise stated, l takes values in the interval $l_{c1}<l<l_{c2}$. Therefore, inequalities such as $l<0.5$ should be taken as shorthand for $l_{c1}<l<0.5$.

Compared with the variation of h_t, agent mutation is a slow process. The value of h_t changes on a time scale of $\Delta t\sim 1$ while agents mutate on a time scale $\Delta t \gg D$. The result of this is that we should not expect the agents to be sensitive to the instantaneous value of $\mathcal{U}_t$ given in Eqs. (7) and (8). They will be sensitive to only the time average $\langle \mathcal{U}_t\rangle_t$. Translated to the conventions used here, Ref. [4] found that

$$\langle h_t\rangle_t = \begin{cases} +0.5 & l>0.5, \\ -0.5 & l<0.5. \end{cases} \tag{12}$$

Therefore, for $l<0.5$, $h_t=+1$ for a fraction 0.25 of the time steps whereas the fraction is 0.75 for $l>0.5$. Thus we can calculate the following expression for $\langle \mathcal{U}_t(n_t^{+h_t})\rangle_t$.

For $l<0.5$,

$$\langle \mathcal{U}_t\rangle_t = \begin{cases} 2n_t^{+h_t}-N, & n_t^{+h_t}\leqslant Nl, \\ n_t^{+h_t}-\dfrac{N}{2}, & Nl<n_t^{+h_t}<N(1-l), \\ -2n_t^{+h_t}+N, & n_t^{+h_t}\geqslant N(1-l). \end{cases} \tag{13}$$

For $l>0.5$,

$$\langle \mathcal{U}_t\rangle_t = \begin{cases} 2n_t^{+h_t}-N, & n_t^{+h_t}<N(1-l), \\ n_t^{+h_t}-\dfrac{N}{2}, & N(1-l)\leqslant n_t^{+h_t}\leqslant Nl, \\ -2n_t^{+h_t}+N, & n_t^{+h_t}>Nl. \end{cases} \tag{14}$$

This expression is plotted in Fig. 2(a). From the figure we can see that the optimal value of $n_t^{+h_t}$ that maximizes $\langle \mathcal{U}_t(n_t^{+h_t})\rangle_t$ is given by

$$n_{t,\mathrm{opt}}^{+h_t} = \begin{cases} N(1-l)-1, & l<0.5, \\ Nl, & l>0.5. \end{cases} \tag{15}$$

Note that, although we have assumed the values given in Eq. (12) for $\langle h_t\rangle_t$, the values of $n_{t,\mathrm{opt}}^{+h_t}$ given in the above equation in fact depend only upon the signs of the values of $\langle h_t\rangle_t$. The significance of this will become apparent in Sec. II G.

Let $\Pi_{i;t}$ be the probability that $n_t^{+h_t}=i$. Lo *et al.* [12] demonstrated that $\Pi_{i;t}$ will be approximately Gaussian with mean $\mu=\langle n_t^{+h_t}\rangle=N\overline{p_t}$ and standard deviation $\sigma=\sqrt{\Sigma_i p_{i;t}(1-p_{i;t})}$. From Eq. (15) it follows that the optimal from of $\Pi_{i;t}$ will obtain if $\overline{p_t}$ is as given by the following equation and $\sigma=0$:

$$\overline{p}_{t,\mathrm{opt}} = \begin{cases} 1-l-\dfrac{1}{N}, & l<0.5, \\ l, & l>0.5. \end{cases} \tag{16}$$

The term $1/N$ results from the asymmetry of the condition in Eq. (1) which determines A_t in terms of n_t^{+1}. For $l<l_{c2}$ in the case of $N\gg 1$ considered here $1/N\ll 1-l$ and so we can neglect this term. $\sigma=0$ if $\{p_{i;t}\}_i$ contains only the values 0 and 1. Let $P(x)$ be the distribution of $\{p_{i;t}\}_i$ such that $NP(x)dx$ is the probability that, if an agent i is chosen at random from the set $\{p_{i;t}\}_i$ then $x\leqslant p_{i;t}\leqslant x+dx$. The optimal form for $P(x)$ is then

$$P(x) = (1-\overline{p}_{t,\mathrm{opt}})\,\delta(x) + \overline{p}_{t,\mathrm{opt}}\,\delta(1-x). \tag{17}$$

This represents a distribution that is zero everywhere except for peaks at $x=0$ and $x=1$, the relative heights of the peaks being such that $\mu=\overline{p}_{t,\mathrm{opt}}$. Burgos *et al.* [6] derived a similar expression for $P(x)$ in the case of the memoryless model using their symmetric cost function.

It is well known [3,13,14] that in the long time limit where $t\to\infty$ the population of agents evolves such that $P(x)$ is strongly peaked about $x=0$ and $x=1$ and $\overline{p_t}$ takes the value given by Eq. (16). Although the agents never manage to achieve a form such that the standard deviation $\sigma(n_t^{+h_t},t)$ is exactly zero, they do approach the optimal distribution represented by Eq. (17). Thus, we can see that the population of agents is capable of evolving such that $\Pi_{i;t}$ is close to its optimal form and the time series of values of $n_t^{+h_t}$ contains values clustered around the optimal value $n_{t,\mathrm{opt}}^{+h_t}$ given by Eq. (15).

2. Memoryless genetic model

In this section we shall see how the above analysis applies to the memoryless variant of the model. $\mathcal{U}_t$ in the memoryless model is given by Eq. (7) above. Thus,

$$\mathcal{U}_t(n_t^{+1}) = \begin{cases} 2n_t^{+1} - N & \text{if } n_t^{+1} \leqslant Nl, \\ -2n_t^{+1} + N & \text{if } n_t^{+1} > Nl. \end{cases} \tag{18}$$

Equation (18) is plotted in Fig. 2(b). Once again the black circle represents the value of n_t^{+1} at which $\mathcal{U}_t$ is a maximum. From the figure we can see that the optimal value of n_t^{+1} is now

$$n_{t,\text{opt}}^{+1} = \begin{cases} Nl + 1 & \text{if } l < 0.5, \\ Nl & \text{if } l > 0.5. \end{cases} \tag{19}$$

Unlike the original genetic model, this optimal value of n_t^{+1} is independent of t. As before, the optimal form of $\Pi_{i;t}$ will be that for which $\sigma = \sigma(n_t^{+1}, t) = 0$ and $\mu = \langle n_t^{+1} \rangle_t = N\overline{p_t}$ with $\overline{p_t}$ given by

$$\overline{p}_{t,\text{opt}} = \begin{cases} l + \dfrac{1}{N} & \text{if } l < 0.5, \\ l & \text{if } l > 0.5. \end{cases} \tag{20}$$

Note that, from Eqs. (16) and (20), this will mean that the gene value distributions of the agents for $l < 0.5$ in the original and memoryless models will not be identical, but will be related by the transformation $\{p_{i;t}\}_t \to \{1 - p_{i;t}\}_t$.

3. Direct comparison of the models

In this section we shall compare the value of $\mathcal{U}_t(n_{t,\text{opt}}^{+h_t})$ in the original genetic model with that of $\mathcal{U}_t(n_{t,\text{opt}}^{+1})$ in the memoryless genetic model in order to establish what effect the memory has on the performance of the model. In both the original genetic model with memory and the memoryless variant, the agents are rewarded based on the value of n_t^{+1}. This is because the value of the global action is determined from the condition on n_t^{+1} in Eq. (1) and an agent i gains or loses one point depending on whether $a_{i;t} = \pm A_t$. There is, however, one difference between the two models that will be extremely important in what follows. The population of agents can control $\Pi_{i;t}$ through their effect on $\overline{p_t}$ and $\sigma = \sqrt{\Sigma_i p_{i;t}(1 - p_{i;t})}$. In the memoryless model $\Pi_{i;t}$ represents the probability distribution for n_t^{+1}, whereas in the original model $\Pi_{i;t}$ represents the distribution function for $n_t^{+h_t}$. The result of this is that in the memoryless model the population of agents can directly control the values that occur in the time series of n_t^{+1} whereas in the original model they can control only the values of $n_t^{+h_t}$. In the latter case n_t^{+1} will also depend on the value of h_t over which the agents have no direct control.

From Eqs. (13), (14), and (18) it follows that the maximum values of $\mathcal{U}_t$ [which obtain at $n_{t,\text{opt}}^{+h_t}$ and $n_{t,\text{opt}}^{+1}$ given by Eqs. (15) and (19)] are, for the model with memory,

$$\mathcal{U}_t(n_{t,\text{opt}}^{+h_t}) = \begin{cases} \dfrac{N}{2}(1 - 2l) - 1 & \text{if } l < 0.5, \\ \dfrac{N}{2}(2l - 1) & \text{if } l \geqslant 0.5, \end{cases} \tag{21}$$

and for no memory,

$$\mathcal{U}_t(n_{t,\text{opt}}^{+1}) = \begin{cases} N(1 - 2l) - 2 & \text{if } l < 0.5, \\ N(2l - 1) & \text{if } l \geqslant 0.5. \end{cases} \tag{22}$$

Thus the optimal value of $\mathcal{U}_t$ in the original model is exactly half that achieved by the memoryless model. As we suggested above, the reason for this is because in the original model n_t^{+1} is a function of both h_t and $\Pi_{i;t}$. The instantaneous optimal value of $n_t^{+h_t}$ will therefore depend on h_t [see Fig. 1 and Eq. (9)]. Note that the value of $n_t^{+h_t}$ that maximizes the time average $\langle \mathcal{U}_t(n_t^{+h_t}) \rangle_t$, $n_{t,\text{opt}}^{+h_t}$ given by Eq. (15), will always be one of the instantaneous optimal values given in Eq. (9). Thus the agents cannot increase $\mathcal{U}_t$ by varying $\overline{p_t}$. They adopt the value of $\overline{p_t}$ that is optimal for the most common value of h_t, but they must pay the penalty when h_t takes the opposite value. In contrast, in the memoryless model the instantaneous optimal value of n_t^{+1} in Eq. (19) is independent of t. Thus, by evolving such that $\overline{p_t} = \overline{p}_{t,\text{opt}}$ the agents can ensure that n_t^{+1} is close to the optimal value at each time step.

D. Analytical expressions for $\langle n_t^{+1} \rangle_t$ and $\sigma(n_t^{+1}, t)$

We can use the same method that we used to derive the expression for $\langle \mathcal{U}_t(n_t^{+h_t}) \rangle_t$ in Eqs. (13) and (14) to obtain expressions for $\langle n_t^{+1} \rangle_t$ and the standard deviation $\sigma(n_t^{+1}, t)$ of the n_t^{+1} time series. From Eq. (3),

$$\langle n_t^{+h_t} \rangle = N\overline{p_t}. \tag{23}$$

This leads to the following expression for $\langle n_t^{+1} \rangle$:

$$\langle n_t^{+1} \rangle = \begin{cases} N(1 - \overline{p_t}) & \text{if } h_t = -1, \\ N\overline{p_t} & \text{if } h_t = +1. \end{cases} \tag{24}$$

Taking the time average in exactly the same way as in Sec. II C 1 in equilibrium where $\overline{p_t}$ is approximately constant yields expressions for $\langle n_t^{+1} \rangle_t$ and $\langle (n_t^{+1})^2 \rangle_t$:

$$\langle n_t^{+1} \rangle_t = \begin{cases} \dfrac{N}{4}(3 - 2\langle \overline{p_t} \rangle_t) & \text{if } l < 0.5, \\ \dfrac{N}{4}(1 + 2\langle \overline{p_t} \rangle_t) & \text{if } l > 0.5, \end{cases} \tag{25}$$

$$\langle [n_t^{+1}]^2 \rangle_t = \begin{cases} \dfrac{N^2}{4}(4\langle \overline{p_t} \rangle_t^2 - 6\langle \overline{p_t} \rangle_t + 3) + \sigma^2 & \text{if } l < 0.5, \\ \dfrac{N^2}{4}(4\langle \overline{p_t} \rangle_t^2 - 2\langle \overline{p_t} \rangle_t + 1) + \sigma^2 & \text{if } l > 0.5, \end{cases} \tag{26}$$

where σ is the standard deviation of $\Pi_{i;t}$ introduced in Sec. II A, which will be of order unity. If we assume that the agents adopt the optimal distribution in Eq. (17) then we can take $\sigma = 0$. We now obtain an expression for $\sigma(n_t^{+1}, t)$ as follows:

$$[\sigma(n_t^{+1}, t)]^2 = \langle [n_t^{+1}]^2 \rangle_t - [\langle n_t^{+1} \rangle_t]^2 = \frac{3N^2}{4}\left(\langle \overline{p_t} \rangle_t - \frac{1}{2} \right)^2 + \sigma^2. \tag{27}$$

Taking $\sigma = 0$ gives

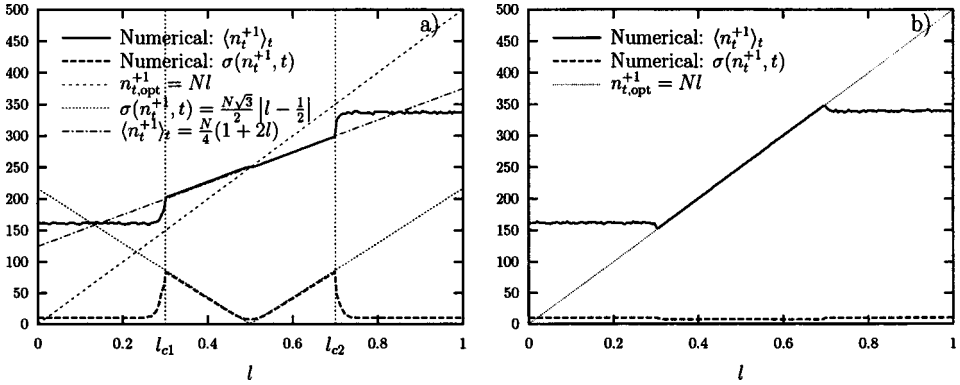

FIG. 3. (a) Numerical results for $\langle n_t^{+1}\rangle_t$ and $\sigma(n_t^{+1},t)$ as a function of the resource level l in the original genetic model. We also include lines that represent $n_{t,\mathrm{opt}}^{+1}=Nl$ from Eq. (11) and the analytical expressions for $\langle n_t^{+1}\rangle_t$ and $\sigma(n_t^{+1},t)$ presented in Eq. (32). (b) $\langle n_t^{+1}\rangle_t$ and $\sigma(n_t^{+1},t)$ as a function of the resource level l in the memoryless genetic model.

$$\sigma(n_t^{+1},t)=\frac{N\sqrt{3}}{2}\left|\langle\overline{p_t}\rangle_t-\frac{1}{2}\right|. \tag{28}$$

Reference [4] used a mean-field approximation to derive expressions for n_t^{+1} and $\sigma(n_t^{+1},t)$ in terms of $\langle h_t\rangle_t$ and $\langle\overline{p_t}\rangle_t$. In particular,

$$\langle n_t^{+1}\rangle_t=\begin{cases}\dfrac{N}{4}(3-2\langle\overline{p_t}\rangle_t) & \text{if } l<0.5,\\ \dfrac{N}{4}(1+2\langle\overline{p_t}\rangle_t) & \text{if } l>0.5,\end{cases} \tag{29}$$

$$\sigma(n_t^{+1},t)=\frac{N\sqrt{3}}{2}\left|\langle\overline{p_t}\rangle_t-\frac{1}{2}\right|. \tag{30}$$

Note that we have substituted for $\langle h_t\rangle_t$ in the expressions of Ref. [4] with the values of $\langle h_t\rangle_t$ that obtain for $l_{c1}<l<l_{c2}$ ($\langle h_t\rangle_t=0.75$ for $l<0.5$ and $\langle h_t\rangle_t=0.25$ for $l>0.5$ [4]). Thus, we can see that the expressions that we derived in Eqs. (25) and (28) are consistent with those obtained in Ref. [4].

E. Numerical results

In this section we present numerical data which support the analytical results that we presented in the previous sections. Reference [4] investigated the behavior of $\langle n_t^{+1}\rangle_t$ and $\sigma(n_t^{+1},t)$ as a function of the resource level l. Further work was done with regard to the memoryless genetic model and general l by Burgos *et al.* in Refs. [6,7]. Figure 3(a) recalls the results of Ref. [4]. We can clearly see the dynamic and frozen regimes for $l_{c1}<l<l_{c2}$ and $l<l_{c1}$, $l>l_{c2}$ respectively. Figure 3(b) shows $\langle n_t^{+1}\rangle_t$ and $\sigma(n_t^{+1},t)$ in the memoryless variant. Figure 4 shows $\langle\overline{p_t}\rangle_t$ as a function of l in the original and memoryless models and $\langle h_t\rangle_t$ in the original model. Equivalent results for $\langle\overline{p_t}\rangle_t$ and $\langle h_t\rangle_t$ for the original model were also presented in Ref. [4].

First of all note that in Fig. 4 $\langle\overline{p_t}\rangle_t$ lies to a very good approximation on the following lines:

$$\text{memory: } \langle\overline{p_t}\rangle_t=\begin{cases}1-l & l<0.5,\\ l & l>0.5,\end{cases} \qquad \text{no memory: } \langle\overline{p_t}\rangle_t=l. \tag{31}$$

This confirms that the population of agents is capable of evolving to achieve the optimal values of $\overline{p_t}$ given in Eqs. (16) and (20).

In Sec. II D we presented expressions in Eqs. (25) and (28)–(30) for $\langle n_t^{+1}\rangle_t$ and $\sigma(n_t^{+1},t)$. If we substitute for $\langle\overline{p_t}\rangle_t$ in these equations with the optimal values from Eq. (16), neglecting the term $1/N$, we obtain the following analytical equations for $\langle n_t^{+1}\rangle_t$ and $\sigma(n_t^{+1},t)$ for $l_{c1}<l<l_{c2}$:

$$\langle n_t^{+1}\rangle_t=\frac{N}{4}(1+2l), \quad \sigma(n_t^{+1},t)=\frac{N\sqrt{3}}{2}\left|l-\frac{1}{2}\right|. \tag{32}$$

In Fig. 3(a) we show these analytic expressions together with the numerical data and $n_{t,\mathrm{opt}}^{+1}=Nl$ from Eq. (11). We can see that $\langle n_t^{+1}\rangle_t$ deviates from the optimal value of Nl for $l_{c1}<l<0.5$ and $0.5<l<l_{c2}$ as pointed out in Ref. [4]. We now know, from Sec. II C 3, that the reason for this is that the population of agents can only control $n_t^{+h_t}$ directly and not n_t^{+1}. Thus their performance is reduced by the action of h_t. We can see that for $l_{c1}<l<l_{c2}$ $\langle n_t^{+1}\rangle_t$ instead lies on the

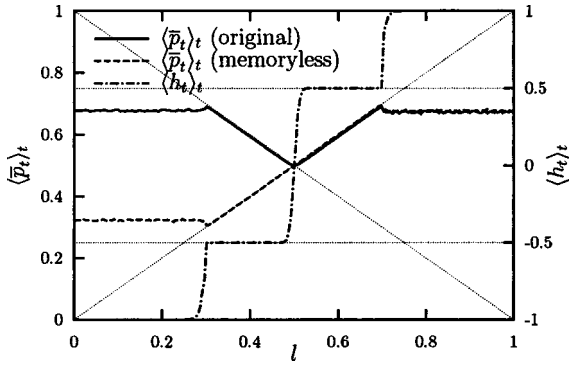

FIG. 4. Numerical results for $\langle\overline{p_t}\rangle_t$, in the original and memoryless models, and $\langle h_t\rangle_t$, in the original model, as a function of l. The diagonal lines represent the functions $\langle\overline{p_t}\rangle_t=l$ and $\langle\overline{p_t}\rangle_t=1-l$.

line defined by Eq. (32).[1] In Fig. 3(a) we also see that the numerical data for $\sigma(n_t^{+1},t)$ agree to a good approximation with the expressions in Eq. (32). The analytic expression deviates from the numerical data in the vicinity of $l=0.5$. The reason for this is that when we derived the expression for $\sigma(n_t^{+1},t)$ in Eq. (28) we assumed that the agent gene value distribution is as given by Eq. (17) and so $\sigma=0$. In fact $\sigma \neq 0$ and near $l=0.5$ the σ^2 term in Eq. (27) dominates. Therefore $\sigma(n_t^{+1},t)$ does not go to zero as predicted.

In Fig. 3(b) we see that, as predicted by Sec. II C 3, $\langle n_t^{+1}\rangle_t$ in the memoryless model does lie on the optimal line defined by $n_{t,\mathrm{opt}}^{+1}=Nl$ for $l_{c1}<l<l_{c2}$. We can also see that $\sigma(n_t^{+1},t)$ in the memoryless model is much lower than in the original model. The large value of $\sigma(n_t^{+1},t)$ in the original model results from the fact that n_t^{+1} is a function of both $n_t^{+h_t}$, determined via $\Pi_{i;t}$ by the distribution of agent gene values $P(x)$, and h_t. In the memoryless model n_t^{+1} is a function of $P(x)$ only, which in equilibrium will be approximately constant in form. The small remaining fluctuations are due to the fact that the agent population does not achieve the ideal distribution of Eq. (17). We can therefore say that the memoryless model is efficient in accessing the available resources.

F. Generation of the prediction from an exogenous source

In Sec. II C we showed that the effect of the prediction h_t is to reduce the agents' performance via its effect on n_t^{+1}. This being the case, we should expect that the effect of the prediction on the model would be no different from that of an exogenous source provided that the value of $\langle h_t\rangle_t$ is preserved. In this section we shall check this by comparing the behavior of the original model with a different memoryless variant. In this variant the prediction h_t will be generated by a random source, external to the model, rather than taking the value $+1\,\forall\, t$. We shall let $\mathcal{R}_\alpha$ represent the output of such a random exogenous source, which contains only the two values -1 and $+1$ and for which α is the time average, $\langle \mathcal{R}_\alpha\rangle_t=\alpha$. We represent the binary sequence generated by the memory for h_t in the original genetic model by $\mathcal{S}$.

Figure 5 shows numerical results for $\langle n_t^{+1}\rangle_t$ in the memoryless model with h_t given by the exogenous sources $\mathcal{R}_{0.5}$ and $\mathcal{R}_{1.0}$. The results for the original model, $\{h_t\}_t=\mathcal{S}$, are included for comparison. The results for $\{h_t\}_t=\mathcal{R}_{1.0}$ duplicate those presented in Fig. 3(b) since $\{h_t\}_t=\mathcal{R}_{1.0}$ is equivalent to $h_t=+1\,\forall\, t$. In other words taking $\{h_t\}_t=\mathcal{R}_{1.0}$ is exactly equivalent to the memoryless model that we considered in previous sections. In Fig. 5 the data produced using $\{h_t\}_t=\mathcal{R}_{0.5}$ for $\langle n_t^{+1}\rangle_t$ and $\sigma(n_t^{+1},t)$ agree with those from the original model for $l_{c1}<l<l_{c2}$. For $l<l_{c1}$ and $l>l_{c2}$ the data from the original model switch to agree with those from the memoryless model with $\{h_t\}_t=\mathcal{R}_{1.0}$ corresponding to the value of $\langle h_t\rangle_t$ from the original model in these regions. Note that there is no need to consider $\mathcal{R}_{-0.5}$ and $\mathcal{R}_{-1.0}$. The lack of physical significance attached to the labeling of the states of h_t means that the model behaves equivalently for $\{h_t\}_t=\mathcal{R}_{\pm\alpha}$.

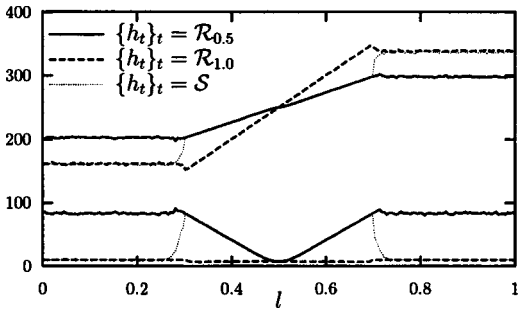

FIG. 5. Numerical results for $\langle n_t^{+1}\rangle_t$ and $\sigma(n_t^{+1},t)$ in the memoryless model using exogenous sources $\mathcal{R}_{0.5}$ and $\mathcal{R}_{1.0}$ for h_t. The results for $\langle n_t^{+1}\rangle_t$ and $\sigma(n_t^{+1},t)$ in the original model, where h_t is determined by the memory ($\{h_t\}_t=\mathcal{S}$), are included for comparison. Each pair of lines shows $\langle n_t^{+1}\rangle_t$ and $\sigma(n_t^{+1},t)$ for $\{h_t\}_t$ given as indicated.

These results confirm that the original genetic model and the memoryless genetic model with h_t taken from an exogenous source can be regarded as equivalent when considering $\langle n_t^{+1}\rangle_t$. In contrast, we shall see in Sec. III that this does not apply when considering higher moments.

G. The values of $\langle h_t\rangle_t$

So far we have treated the values of $\langle h_t\rangle_t$ that obtain for $l_{c1}<l<0.5$ and $0.5<l<l_{c2}$ as values to be derived empirically by numerical simulation. Now we shall discuss the theoretical reasons for their observed values.

Lo [14] has presented a theory that predicts, using our conventions, the following values for $\langle h_t\rangle_t$:

$$\langle h_t\rangle_t=\begin{cases}-1 & \text{for } l<l_{c1},\\ -\dfrac{1}{3} & \text{for } l_{c1}<l<0.5,\\ +\dfrac{1}{3} & \text{for } 0.5<l<l_{c2},\\ +1 & \text{for } l_{c2}<l.\end{cases} \tag{33}$$

However, numerical simulation robustly yields values of $\langle h_t\rangle_t\approx\pm\frac{1}{2}$ in the dynamic regime. In what follows we briefly recall Lo's analysis with the addition of some observations that explain why the numerical and analytical results differ. Note that as we pointed out in Sec. II C 1 the absolute values that obtain for $\langle h_t\rangle_t$ in the dynamic regime are not important for the theory that we present here. As long as $\langle h_t\rangle_t<0$ for $l_{c1}<l<0.5$ and $\langle h_t\rangle_t>0$ for $0.5<l<l_{c2}$ everything that we have said about $\overline{p}_{t,\mathrm{opt}}$ will remain unchanged. Only the magnitude of the relative performance of the original and memoryless genetic models depends upon the values taken by $\langle h_t\rangle_t$.

Lo's analysis [14] hinges on the observation that

$$\langle A_t\rangle_t=\langle h_t\rangle_t. \tag{34}$$

It is then easy to show that $\langle A_t\rangle_t$ is given by

[1]Note that in fact $\langle n_t^{+1}\rangle_t$ lies slightly closer to $N/2$ than Eq. (32). This is due to the deviation depicted in Fig. 6(a) which we shall discuss in Secs. II G and II H.

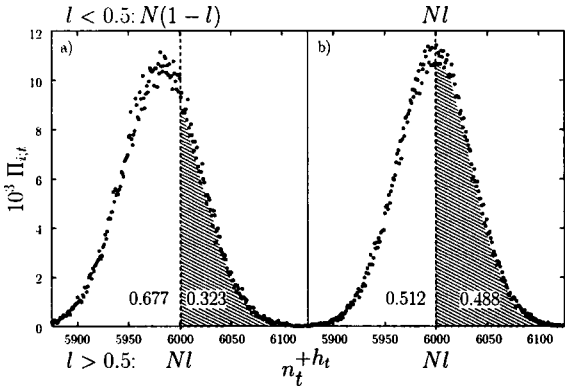

FIG. 6. Numerical data for $\Pi_{i;t}$ in (a) the original model with memory and (b) the memoryless model. The numbers under the curve represent the sum of the values of the data points lying on either side of the cutoff indicated by the dashed line. Note that in (a) this cutoff is at $n_t^{+h_t}=Nl$ if $l>0.5$ and $n_t^{+h_t}=N(1-l)$ otherwise. The model parameters were as given in Table I except that $n=10\,000$ and the value of $n_t^{+h_t}$ was sampled over the interval $2000<t<102\,000$.

$$\langle A_t\rangle_t=\left(\frac{\langle h_t\rangle_t+1}{2}\right)(2I_1-1)+\left(\frac{1-\langle h_t\rangle_t}{2}\right)(2I_2-1), \quad (35)$$

where

$$I_1=\sum_{i=0}^{Nl}\Pi_{i;t}, \quad I_2=\sum_{i=N(1-l)}^{N}\Pi_{i;t}. \quad (36)$$

Finally, using Eq. (34) gives

$$\langle h_t\rangle_t=\frac{I_1+I_2-1}{I_2-I_1+1}. \quad (37)$$

Assuming that $\mu=n_{t,\mathrm{opt}}^{+h_t}$ from Eq. (15), we can see that I_1 and I_2 will be given by

$$I_1=\begin{cases}0 & \text{if } l<0.5,\\ \frac{1}{2} & \text{if } l>0.5,\end{cases} \quad I_2=\begin{cases}\frac{1}{2} & \text{if } l<0.5,\\ 1 & \text{if } l>0.5,\end{cases} \quad (38)$$

since the summations cover the entire peak of the Gaussian, exactly half of it, or none at all. From Eqs. (37) and (38) we obtain values for $\langle h_t\rangle_t$ in the dynamic regime of

$$\langle h_t\rangle_t=\begin{cases}-\frac{1}{3} & \text{if } l<0.5,\\ +\frac{1}{3} & \text{if } l>0.5,\end{cases} \quad (39)$$

which agree exactly with the values found by Lo in Ref. [14].

The values given in Eq. (38) depend on the assumption that $\langle n_t^{+h_t}\rangle_t$ is given by the optimal value of Eq. (15). In fact, as we see in Fig. 6(a), numerical simulation using a large number of agents reveals that this is not the case. In the figure $n_{t,\mathrm{opt}}^{+h_t}-\langle n_t^{+h_t}\rangle_t\approx 0.002N$ corresponding to $\overline{p_{t,\mathrm{opt}}}-\langle\overline{p_t}\rangle_t \approx 0.002$. Careful observation of Fig. 4 reveals that $\langle\overline{p_t}\rangle_t$ always lies slightly closer to 0.5 than the value predicted by Eq. (16) in the original model. As we can see from Fig. 6(a), this changes the values of the summations in Eq. (38) as follows:

$$I_1=\begin{cases}0 & \text{if } l<0.5,\\ 0.677 & \text{if } l>0.5,\end{cases} \quad I_2=\begin{cases}0.323 & \text{if } l<0.5,\\ 1 & \text{if } l>0.5,\end{cases} \quad (40)$$

which by Eq. (37) yields the following values for $\langle h_t\rangle_t$:

$$\langle h_t\rangle_t=\begin{cases}-0.511 & \text{if } l<0.5,\\ +0.511 & \text{if } l>0.5.\end{cases} \quad (41)$$

These figures agree much more closely with the numerically observed values of $\langle h_t\rangle_t=\pm 0.5$ than the values obtained in Eq. (38).

We can therefore see that the values of $\langle h_t\rangle_t=\pm 0.5$ that obtain in the dynamic regime result from the fact that $\overline{p_t}$ deviates slightly from the optimal value of Eq. (16). In the absence of this deviation we would indeed obtain the values given in Eq. (33) for $\langle h_t\rangle_t$. We can see that the memory does indeed have a nontrivial role to play in the genetic model. The feedback between the prediction h_t and the global action A_t controls the values of $\langle h_t\rangle_t$ that obtain in the dynamic regime $l_{c1}<l<l_{c2}$.

H. Deviation of $\overline{p_t}$ from its optimal value

In Sec. II C when we calculated the optimal value of $\overline{p_t}$ in Eq. (16), we implicitly assumed that the gene value distribution is given by Eq. (17) and that therefore $\sigma=0$. If σ is nonzero then the values of n_t^{+1} will be spread about $\langle n_t^{+1}\rangle_t$. Considering the asymmetric nature of the discontinuity in $\mathcal{U}_t$ at its maximum value in Fig. 2(a), we might expect the agents to evolve such that $\overline{p_t}<\overline{p_{t,\mathrm{opt}}}$ in order to increase their total earnings. However, the asymmetry at the discontinuity in the memoryless model is greater than in the original. Thus, we would also expect the deviation of $\overline{p_t}$ to be greater in the memoryless model.

Figure 6(b) shows numerical data for $\Pi_{i;t}$ in the memoryless model. Any deviation in the figure is too small to be observable (see also Fig. 4). Therefore, we conclude that a different effect, one most probably related to the variation of the prediction h_t, must be responsible for the deviation observed in the original model.

I. Autocorrelation of n_t^{+1}

In this section we consider another property of the genetic model: the autocorrelation of the time series n_t^{+1}. First of all we shall present numerical data contrasting the behavior of the autocorrelation in the original and memoryless models. We shall then consider a Markovian analysis of the original genetic model which explains the behavior of the autocorrelation of n_t^{+1} observed therein in the simplest case of $m=1$.

1. Numerical results

We define the autocorrelation $\mathcal{C}_\tau(x_t)$ of a time series x_t as follows:

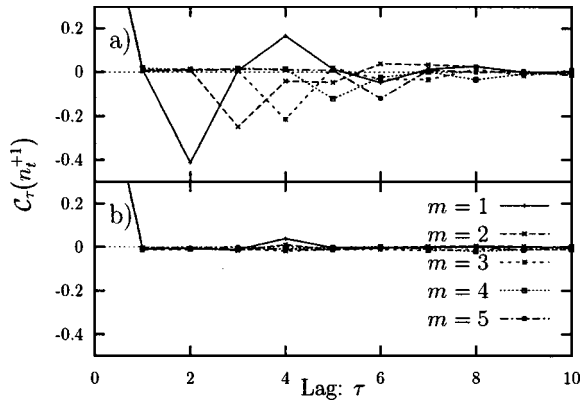

FIG. 7. Autocorrelation of the n_t^{+1} time series for the range of values of m shown. Model parameters $N=1000$, (a) $l=0.4$, and (b) $l=0.5$.

$$\mathcal{C}_\tau(x_t)=\frac{\langle x_t x_{t+\tau}\rangle_t-\langle x_t\rangle_t^2}{\langle x_t^2\rangle_t-\langle x_t\rangle_t^2}, \tag{42}$$

where τ gives the lag time.

Figure 7 shows the autocorrelation of the attendance time series $\mathcal{C}_\tau(n_t^{+1})$ for $0\leqslant\tau\leqslant 10$ and $1\leqslant m\leqslant 5$. We can see from the figure that for $l=0.5$ there is no significant correlation for $\tau>0$. For $l\neq 0.5$, in contrast, $\mathcal{C}_\tau(n_t^{+1})$ as a function of τ has a clear structure. The magnitude of $\mathcal{C}_\tau(n_t^{+1})$ is nonzero for $\tau=(m+1)i$, where $i=1,2,3,\ldots n_t^{+1}$ and $n_{t+\tau}^{+1}$ are anticorrelated for odd values of i and correlated for even values. Figure 8 shows $\mathcal{C}_\tau(n_t^{+1})$ for the same values of τ and l depicted in Fig. 7. However, this time h_t was derived from a random exogenous source rather than the memory. It is clear from a comparison of Figs. 7 and 8 that the structure observed in $\mathcal{C}_\tau(n_t^{+1})$ in Fig. 7 is present only in the original model with memory. Plotting the autocorrelation of the prediction time series h_t yields a graph identical in form to Fig. 7. From this we can see that the structure present in Fig. 7 derives from structure present in the prediction time series.

Thus, one of the functions of the memory is to introduce nonzero autocorrelations into the prediction time series h_t. These will clearly not be present in either of the memoryless models in which either $h_t=+1\,\forall\, t$ or $h_t=\mathcal{R}_\alpha$. In Ref. [5] Burgos and Ceva found that, for $l=0.5$, the gene value distribution functions in the models with and without memory (taking $h_t=+1\,\forall\, t$) are equivalent. Further, in Ref. [7] Burgos *et al.* state that "a simplified version of the model, *that makes no use of memory* is indistinguishable from the original formulation." We agree that, in the restricted case of $l=0.5$ considered in Ref. [5], the gene value distribution functions and the autocorrelation functions of the n_t^{+1} time series are indistinguishable. This does not, however, represent a rigorous proof of equivalence. In contrast, the numerical data presented in Fig. 7 demonstrate that, for $l\neq 0.5$, the two formulations *are* distinguishable. While for $h_t=\mathcal{R}_\alpha$ the memoryless model yields equivalent results for $\langle n_t^{+1}\rangle_t$, it does not do so for higher moments. Therefore, any analysis in which possible autocorrelations in the n_t^{+1} time series could be a factor could not be conducted using the memoryless model. Also, the deviation of $\overline{p_t}$ from the optimal value of Eq. (16) (discussed in Secs. II G and II H) means that the gene value distribution functions are also not strictly indistinguishable accept at $l=0.5$. Furthermore, Fig. 3 demonstrates that if h_t in the memoryless model is taken to be constant the two models are also distinguishable, for $l\neq 0.5$, in terms of $\langle n_t^{+1}\rangle_t$.

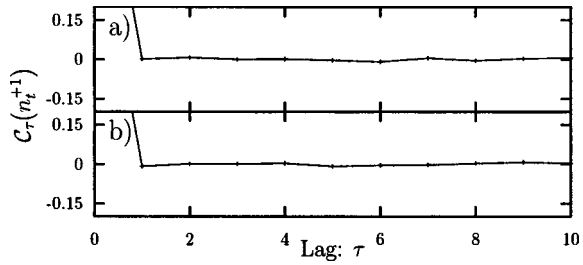

FIG. 8. Autocorrelation of the n_t^{+1} time series with h_t given by a random exogenous source. Model parameters $N=1000$, (a) $l=0.4$, $h_t=\mathcal{R}_{-0.5}$, and (b) $l=0.5$, $h_t=\mathcal{R}_{0.0}$.

2. Markovian analysis

In this section we will present a Markovian analysis of the action of the memory in the genetic model. This analysis must be performed separately for each value of m of interest since each leads to a distinct state space. Here we shall present the analysis for $m=1$ only. Treatment of higher values of m is possible, although cumbersome, since they lead to state spaces that are too large to be treated conveniently by hand.

The first stage of our analysis is to define a convention for labeling the states of the memory. Each state label must define the values of h_t, A_t, and the state of the memory. This is the minimum set of information needed to calculate the state transition probabilities. For $m=1$ the memory will contain $2^{m=1}=2$ entries corresponding to the two possible histories $A_{t-1}=-1$ and $A_{t-1}=+1$. We shall label these entries m_t^{-1} and m_t^{+1}, respectively. Each state therefore comprises four attributes each of which can take values ± 1. There are therefore 16 possible states which we denote using the shorthand notation described in Fig. 9(b). We label these states by analogy with the binary number system as shown in Fig. 9(a). Note that h_t gives the prediction made available to the agents at time t whereas m_t^{-1} and m_t^{+1} represent the state of the memory once it has been updated. Therefore, the prediction agrees with one of the states of the memory at time $t-1$ and not at time t.

For the sake of clarity we shall consider the case of $l<0.5$. By symmetry the results that we derive will also apply to $l>0.5$. From Eqs. (6) and (40) and the definition of $\Pi_{i;t}$, we have the following expression for the probability that $A_t=+1$:

$$P[A_t=+1]=\begin{cases} I_2=\alpha & \text{if } h_t=-1, \\ I_1=0 & \text{if } h_t=+1, \end{cases} \tag{43}$$

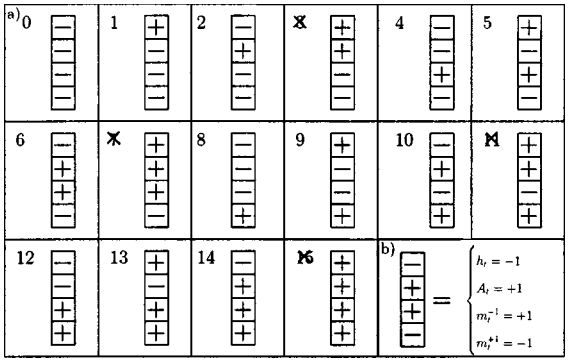

FIG. 9. (a) The labels used to index the 16 possible states. The states that are crossed out are those in which $h_t=+1$ and $A_t=+1$. These can be omitted due to Eq. (43). (b) A key for the graphical representation of each state.

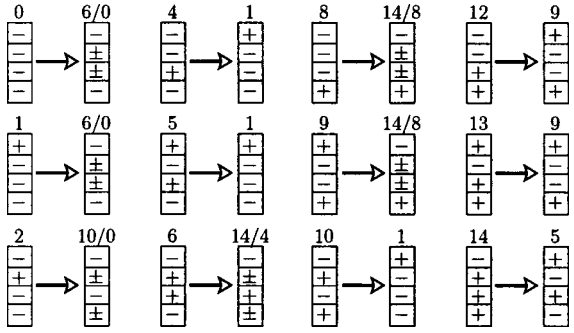

FIG. 10. List of the one-step transitions that the model can make from each of the states listed in Fig. 9.

where I_1 and I_2 are given by Eq. (36) and α represents the numerical value of I_2 in Eq. (40). The result of this is that if the prediction $h_t=+1$ then the global action $A_t=-1$ with probability $P=1$. We can therefore discard states 3, 7, 11, and 15 from consideration. In each of these states $h_t=+1$ and $A_t=+1$ and thus they will never be visited by the model.

Using Eq. (43) we derive the state transitions shown in Fig. 10. We can see that each state makes a transition either to one other state with probability $P=1$ or to one of two possible states with probabilities $P=\alpha$ and $P=1-\alpha$, respectively. This information can be used to further simplify the state space. We need not consider states that have no inward transitions or states that have inward transitions only from states that we have removed. This allows us to remove states 2, 8, 9, 10, 12, and 13 from consideration. The model will not visit these states once initial transients have died away. This reduces our state space from sixteen states to only six: 0, 1, 4, 5, 6, and 14.

Figure 11 shows the state transition diagram corresponding to Fig. 10 in the simplified state space described above. From this we can form the following Markov transition matrix, in which the remaining states are arranged in numerical order:

$$\mathbf{T} = \begin{pmatrix} 1-\alpha & 1-\alpha & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1-\alpha & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ \alpha & \alpha & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \alpha & 0 \end{pmatrix} \tag{44}$$

which defines the stationary Markov chain.

From Eq. (44) we can derive the n-step autocorrelation functions $\mathcal{C}_n(h_t)$. Recall the oscillatory dependence of $\mathcal{C}_\tau(h_t)$ on τ [see Fig. 7(a) and Sec. II I 1]. If $\mathcal{C}_1(h_t)=0$ and $\mathcal{C}_2(h_t)=-\epsilon$ (where ϵ is of order 0.1) then we would expect to observe the dependence depicted in the figure. If h_t and h_{t+2} are anticorrelated and $|\mathcal{C}_2(h_t)|<1$ then we should expect h_t and h_{t+4} to be correlated and $|\mathcal{C}_4(h_t)|<|\mathcal{C}_2(h_t)|$, and so on. Thus the form of $\mathcal{C}_\tau(h_t)$ for $m=1$ depends only on the values of $\mathcal{C}_1(h_t)$ and $\mathcal{C}_2(h_t)$.

First we calculate the stationary state $\mathbf{s}$ of the Markov chain defined by Eq. (44). This gives

$$\mathbf{s} = \frac{1}{(1+\alpha)^2}\begin{pmatrix} 1-\alpha \\ \alpha \\ \alpha(1-\alpha) \\ \alpha^2 \\ \alpha \\ \alpha^2 \end{pmatrix}. \tag{45}$$

From Eq. (42), the one-step autocorrelation $\mathcal{C}_1(h_t)$ is given by

$$\mathcal{C}_1(h_t) = \frac{\langle h_t h_{t+1}\rangle_t - \langle h_t\rangle_t^2}{\langle h_t^2\rangle_t - \langle h_t\rangle_t^2}. \tag{46}$$

Therefore we need to calculate values for $\langle h_t^2\rangle_t$, $\langle h_t\rangle_t$, and $\langle h_t h_{t+1}\rangle_t$. Since $h_t=\pm 1$, $\langle h_t^2\rangle_t=1$. In order to calculate values for the other two averages it is necessary to form the vectors $\mathbf{h}_0$ and $\mathbf{h}_1$ which, respectively, give the values of h_t and $h_t h_{t+1}$ for each of the states in the simplified state space. The elements of $\mathbf{h}_0$ can be read directly from Fig. 9. This gives

$$\mathbf{h}_0 = (-1+1-1+1-1-1). \tag{47}$$

Taking the scalar product of $\mathbf{h}_0$ with $\mathbf{s}$ gives

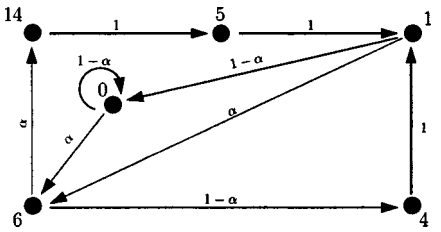

FIG. 11. State transition diagram corresponding to the transitions depicted in Fig. 10 in the simplified state space. The arrow labels give the transition probabilities.

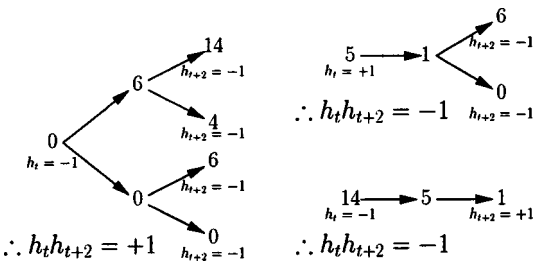

FIG. 12. Example of how to derive the values of $h_t h_{t+2}$ for each of the states in the simplified state space.

$$\langle h_t \rangle_t = \mathbf{h}_0 \cdot \mathbf{s} = \frac{\alpha - 1}{\alpha + 1}. \tag{48}$$

In Sec. II G we used an alternative method, not restricted to $m=1$, to derive the expression for $\langle h_t \rangle_t$ in Eq. (37). Substituting $I_1 = 0$ and $I_2 = \alpha$ in Eq. (37) yields Eq. (48) above. Thus, the analyses presented here and in Sec. II G are consistent.

We can derive $\mathbf{h}_1$ by multiplying the values of h_t and h_{t+1} in Fig. 10. This gives

$$\mathbf{h}_1 = (+1 - 1 - 1 + 1 + 1 - 1). \tag{49}$$

Once again taking the scalar product with $\mathbf{s}$ gives

$$\langle h_t h_{t+1} \rangle_t = \mathbf{h}_1 \cdot \mathbf{s} = \left(\frac{\alpha - 1}{\alpha + 1} \right)^2. \tag{50}$$

Therefore, from Eqs. (46), (48), and (50) we have

$$\mathcal{C}_1(h_t) = 0. \tag{51}$$

Thus, we expect h_t and h_{t+1} to be uncorrelated, which is in agreement with Fig. 7(a).

The two-step autocorrelation $\mathcal{C}_2(h_t)$ is given by

$$\mathcal{C}_2(h_t) = \frac{\langle h_t h_{t+2} \rangle_t - \langle h_t \rangle_t^2}{\langle h_t^2 \rangle_t - \langle h_t \rangle_t^2}. \tag{52}$$

Therefore in order to calculate the two-step autocorrelation function $\mathcal{C}_2(h_t)$ we must calculate $\langle h_t h_{t+2} \rangle_t$. As before we must form the vector $\mathbf{h}_2$ corresponding to the values of $h_t h_{t+2}$ for each of the states in the simplified space. This can be done following the method described in Fig. 12, which gives

$$\mathbf{h}_2 = (+1 - 1 + 1 - 1 - 1 - 1). \tag{53}$$

Taking the scalar product with $\mathbf{s}$ gives

$$\langle h_t h_{t+2} \rangle_t = \mathbf{h}_2 \cdot \mathbf{s} = \frac{(1 - 3\alpha)(1 + \alpha)}{(1 + \alpha)^2}. \tag{54}$$

From Eqs. (48), (52), and (54) we have

$$\mathcal{C}_2(h_t) = -\alpha. \tag{55}$$

Therefore, from Eqs. (51) and (55) we have that

$$\mathcal{C}_\tau(h_t) = 1, 0, -\alpha, 0, \alpha^2, 0, -\alpha^3, \ldots, \quad \tau = 0, 1, \ldots. \tag{56}$$

Note that as a result of this the autocorrelation of h_t at $m=1$ gives a direct measurement of the value of α. This provides a much more convenient measure of α than computing the summation for I_2 in Eq. (36), as was done in Sec. II G.

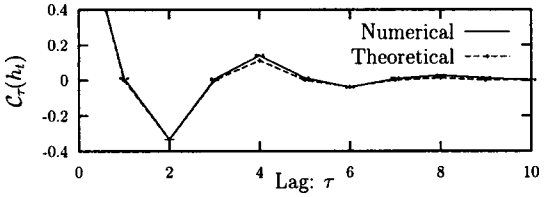

FIG. 13. Comparison of the theoretical result of Eqs. (51) and (55) with numerical results. The error bars show one standard deviation on the mean over an ensemble of five separate data sets.

Figure 13 shows a comparison of the theoretical values of $\mathcal{C}_\tau(h_t)$, obtained with the numerical value of $\alpha = 0.323$ from Sec. II G, with numerical values. We can see from Fig. 13 that Eq. (56) correctly predicts the form of $\mathcal{C}_\tau(h_t)$ although it does underestimate the correlation for $\tau = 4$. One possible explanation of this slight deviation is that in the analysis above we have assumed that α is a constant. However, the value of α depends very sensitively on the form of $\Pi_{i;t}$. Therefore fluctuations in the gene value distribution of the agents could cause significant fluctuations in α. Markovian analysis of the $l = 0.5$ case confirms that $\mathcal{C}_\tau(h_t) = 0$ for $\tau > 1$, in agreement with the numerical results of Sec. II I 1.

J. Summary of memory characteristics

We conclude from the results presented in this section that the genetic model performs better in the absence of memory. By this we mean that the average total number of points scored per time step when $h_t = +1 \, \forall \, t$ is twice that when h_t is determined by the global memory. This was established semianalytically in Sec. II C and supporting numerical data were presented in Sec. II E. We showed in Sec. II C 3 that the reason for this reduction in the performance in the presence of memory is that the agents cannot directly control the distribution of values for n_t^{+1} as they can in the memoryless model, because n_t^{+1} is now also a function of h_t. In Sec. II F we presented numerical data to show that the values of $\langle n_t^{+1} \rangle_t$ and $\sigma(n_t^{+1}, t)$ can be reproduced if the prediction h_t is taken from an exogenous source providing that the value of $\langle h_t \rangle_t$ is preserved. Thus, the feedback between the global action of the agents and h_t is of no benefit to the population of agents. The only function of this feedback is to regulate the value of $\langle h_t \rangle_t$.

We also investigated the values observed for the time average of the prediction $\langle h_t \rangle_t$ as a function of the resource level l. We showed that the values of $\langle h_t \rangle_t = \pm 0.5$ that obtain in the dynamic regime of $l_{c1} < l < l_{c2}$ are due to the deviation observed between the mean of the agent gene value distribution $P(x)$ and the optimal value predicted by Eq. (20).

In Sec. II I we showed that the form of the autocorrelation function of the n_t^{+1} time series at $l \neq 0.5$ occurs as a result of the cycles in state space performed by the memory. Finally, we demonstrated that the two-step autocorrelation of the prediction $\mathcal{C}_2(h_t)$ can be used to provide a direct measurement of the deviation described above.

III. SELF-INDUCED SHOCKS IN THE GRAND CANONICAL GENETIC MODEL

In this section we move on to consider the important practical property of self-induced shocks, otherwise known as endogenous large changes (ELCs). Such large changes are arguably a defining characteristic of complex systems, yet there is no rigorous quantitative description of such events in real-world realizations of complex systems (for examples, see Refs. [15–22]).

A. Introduction

As we show here, the genetic model can be generalized in a straightforward way to produce a model system that demonstrates such large changes. The extent to which these large changes are insensitive to the memory then provides a useful tool for analyzing the microscopic causes underlying these large changes. In particular, we introduce an extension of the genetic model in Sec. II, in which the number (or "volume") of active agents is a time-dependent quantity. By analogy with systems with variable particle number in statistical mechanics, we shall refer to this variant of the genetic model as the grand canonical genetic model (GCGM). This extension is analogous to various extensions of Challet and Zhang's minority game [2] with variable particle number [15,16,21–25], generally known as the grand canonical minority game (GCMG).

One particular application might be to financial markets, where large changes are called crashes or drawdowns. However, the genetic model does not directly yield a price time series. Therefore, if we want to consider the effect of endogenous large changes on price, we must derive one from fundamental observables such as $n_t^{\pm 1}$. By definition the threshold value of $n_t^{+1}=Nl$ corresponds to the state in which the volume of the item that is being traded which is available for sale is equal to the demand. Therefore, the more general case of $l \neq 0.5$ represents a system in which the quantities in which an item is bought and sold are not equal. The excess demand is then given by

$$\Delta = \left(\frac{l}{1-l}\right) N_{\text{buy}} - N_{\text{sell}}. \tag{57}$$

If we let the action of an agent $a_{i;t}=-1,+1$ represent choosing to *buy* or *sell*, respectively then we obtain the following expression for the price π_{t+1} at time $t+1$ in terms of the price at t:

$$\pi_{t+1} = \pi_t + \frac{1}{\lambda}\left[\left(\frac{l}{1-l}\right) n_t^{-1} - n_t^{+1}\right], \tag{58}$$

where λ is known as the *market depth* and determines the magnitude of the change in price caused by a unit change in Δ. Different expressions for the price π_t in terms of the excess demand Δ have been discussed in the literature (see, for example, Refs. [26,27]). The linear expression in Eq. (58) represents the simplest of these and is not as realistic as expressions in higher powers of Δ. Nevertheless, it is more than adequate for the illustrative purposes for which we shall need it. Since in all that follows the units of the price π_t are arbitrary we will take $\lambda=1$.

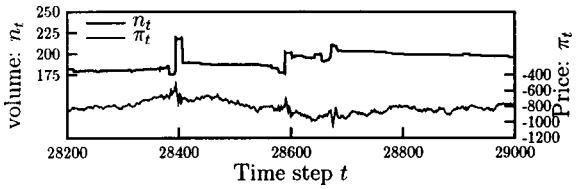

FIG. 14. Example of an endogenous large change of the volume in the GCGM with accompanying price time series. Model parameters: $N=501$, $m=3$, $r=0.2$, $l=0.5$, and $T=12$.

An example of a large volume change observed in the GCGM is given in Fig. 14. We can see that large changes occur in the volume accompanied by large price movements. However, the behavior of the volume in the figure is qualitatively different from that observed by Johnson *et al.* in the GCMG [21]. Reference [21] reported two distinct types of behavior. For traders with a long memory the volume was observed to be continuously fluctuating with occasional particularly large fluctuations which were not instantaneous; much like the "drawdowns" and "drawups" discussed by Sornette and Johansen in Ref. [28]. For traders with short memories the volume was frequently zero with occasional large instantaneous spikes and corresponding instantaneous price movements. In contrast, in Fig. 14 the volume exhibits small fluctuations and occasional instantaneous changes which are accompanied by periods of large fluctuations in the price. We shall see that the behavior that we observe in these figures is typical of the behavior of the volume in the GCGM.

B. The grand canonical genetic model

The genetic model represents an abstract model of a population competing for a limited resource and as such it often discussed in the context of financial markets [3,4,8,9,12,13,29]. In such a context, however, it does not seem realistic that the agents trade at every time step. A real market trader would also have the option of withdrawing from the market and returning when s/he felt confident of a successful outcome. In order to model this extra degree of freedom, and in keeping with the work of Johnson *et al.* in Ref. [21], we shall introduce an extension of the genetic model in which the agents are free to opt in and out of the game.

As in the original genetic model described in Sec. II and Ref. [3] there are N agents participating in the model. However, unlike the original model they do not all play at each time step. At any time t there will be two populations of agents, an active population and an inactive one. Active agents participate in exactly the same way as they do in the original model. In contrast, an inactive agent i will continue to make its choice as if it were participating; however, it will not be considered when calculating the global action and its score will not be updated. We can imagine that inactive agents represent traders who do not make a trade at time t. Instead they make a prediction of whether their best choice would have been to follow the prediction $a_{i;t}=+h_t$ or to refute it $a_{i;t}=-h_t$. Since such a trade is *virtual* in that the trader

does not act on it, it has no effect on the market and the trader is protected from losing, or even making, money. The only effect of such a virtual trade is that the agent reevaluates its confidence level depending on whether the trade would have been successful or not. These *virtual trades* are analogous to the *virtual points* earned by strategies in the minority game [2]. They allow the agents to keep track of their potential performance without in any way representing an agent's wealth.

These modifications necessitate some minor changes to the basic equations of the genetic model that we introduced in Sec. II. Let the activity of an agent i be given by $z_{i;t}$. If $z_{i;t}=0$ then agent i belongs to the population of inactive agents, and vice versa for $z_{i;t}=1$. Equation (1) for the global action A_t now becomes

$$A_t = \begin{cases} +1 & n_t^{+1} \leqslant l n_t, \\ -1 & n_t^{+1} > l n_t, \end{cases} \quad \text{where } n_t = \sum_i z_{i;t}. \tag{59}$$

Inactive agents become active and active agents become inactive according to their performance in the recent past. We defined a model parameter known as the *confidence interval* T. An inactive agent will become active if it would have won for T consecutive time steps; in other words, an agent i for which $z_{i;t}=0$ will activate, $z_{i;t+1}=1$, if $a_{i;\tau}=+A_\tau$ for $t-T<\tau \leqslant t$. In the same way, an active agent will become inactive if it loses for T consecutive time steps. In order to control the activation and deactivation of agents we assign each agent a quantity with we shall call its virtual points $v_{i;t}$, in keeping with the virtual points allocated to strategies in the minority game. For an active agent $v_{i;t}$ is increased each time the agent loses and is reset to zero if it wins. Thus, for active agents $v_{i;t}$ is the number of consecutive time steps for which the agent has lost. For an inactive agent $v_{i;t}$ is increased each time that the agent would have won and is reset to zero each time it would have lost. The updating rules for $v_{i;t}$ can be summarized as follows.

If $z_{i;t}=0$,

$$v_{i;t+1} = \begin{cases} 0 & \text{if } a_{i;t} = -A_t, \\ v_{i;t}+1 & \text{if } a_{i;t} = +A_t. \end{cases} \tag{60}$$

If $z_{i;t}=1$,

$$v_{i;t+1} = \begin{cases} v_{i;t}+1 & \text{if } a_{i;t} = -A_t, \\ 0 & \text{if } a_{i;t} = +A_t. \end{cases} \tag{61}$$

The rules for agent activation and deactivation are then, if $z_{i;t}=0$,

$$z_{i;t+1} = \begin{cases} 0 & \text{if } v_{i;t} < T, \\ 1 & \text{if } v_{i;t} = T; \end{cases} \tag{62}$$

if $z_{i;t}=1$,

$$z_{i;t+1} = \begin{cases} 1 & \text{if } v_{i;t} < T, \\ 0 & \text{if } v_{i;t} = T. \end{cases} \tag{63}$$

The expression in Eq. (2) for the number of agents for which $a_{i;t}=+1$ becomes

$$n_t^{+1} = \frac{1}{2}\sum_i z_{i;t}(a_{i;t}+1). \tag{64}$$

In this section we shall define $\overline{p_t}$ to be the mean gene value of the *active* agent population. Thus, $\overline{p_t}$ is given by

$$\overline{p_t} = \frac{1}{n_t}\sum_i z_{i;t} p_{i;t}. \tag{65}$$

Therefore, the expression in Eq. (3) for the mean number of agents following the prediction h_t becomes

$$\langle n_t^{+h_t} \rangle = \sum_i z_{i;t} p_{i;t} = n_t \overline{p_t}. \tag{66}$$

C. Price time series

As we stated in Sec. III A the cutoff, now given by $n_t^{+1}=n_t l$, in the genetic model, defined by Eq. (59), is by definition the state in which the excess demand $\Delta=0$. From Eq. (58) we can see that the price change $\Delta\pi=\pi_t-\pi_{t-1}$ is positive and negative for $n_t^{+1}<n_t l$ and $n_t^{+1}>n_t l$, respectively. By comparison with Eq. (59) we can see that the condition that determines the sign of the price change at time t is the same as that which determines the global action A_t; with the exception that the equality in Eq. (59) gives a price change of zero. With reference to Eq. (6) we have the condition that $\Delta\pi \geqslant 0$ in terms of $n_t^{+h_t}$:

$$\begin{aligned} n_t^{+h_t} &\geqslant n_t(1-l) \quad \text{if } h_t = -1, \\ n_t^{+h_t} &\leqslant n_t l \qquad \text{if } h_t = +1. \end{aligned} \tag{67}$$

From this it follows that the probability $P[\Delta\pi \geqslant 0]$ that the price rises or remains the same at time t is given by

$$P[\Delta\pi \geqslant 0] = \begin{cases} I_2 & \text{if } h_t = -1, \\ I_1 & \text{if } h_t = +1, \end{cases} \tag{68}$$

where I_1 and I_2 represent the summations defined by Eq. (36) with the substitution $N \rightarrow n_t$. We can see from Eqs. (40) and (68) that for $l<0.5$ $P[\Delta\pi \geqslant 0]=0$ for $h_t=+1$ and for $l>0.5$ $P[\Delta\pi \geqslant 0]=1$ for $h_t=-1$. The result of this is that, in the dynamic regime of the original GCGM with memory, one of the values that h_t can take will cause the price to rise or fall with probability $P=1$. Note that in the cases of $l<0.5, h_t=-1$ and $l>0.5, h_t=+1$ we do not expect that $P[\Delta\pi \geqslant 0]=0.5$, as might be expected. Recall that, as we discussed in Sec. II G, $\overline{p_t}$ deviates from the optimum values given in Eq. (16) in the genetic model. Furthermore, because the GCGM is frequently perturbed by ELCs it does not settle into equilibrium in the same way as the genetic model and so $\overline{p_t}$ is more variable although, as we shall see, it does remain close to the values given by Eq. (16) in the periods between ELCs.

D. Endogenous large changes in the GCGM

As we shall see the behavior of the model in this regime is rich and complex. For this reason we shall initially consider a simplified memoryless variant of the GCGM which is analogous to the memoryless genetic model considered in Sec. II. Subsequently we shall consider how the inclusion of memory affects the behavior of the full GCGM.

ELCs in the GCGM result from a combination of two factors, both of which must be present if a large change is to occur. We shall see later that the capacity of the model to undergo an ELC depends sensitively on the distribution of gene values $P(p)$. However, a suitable $P(p)$ is not a sufficient condition for an ELC to occur. It is also necessary for a particular pattern to occur in the global action time series. We can think of this pattern as a *trigger* that initiates the ELC, but only if $P(p)$ is in a susceptible state. Furthermore, we shall see that the natural evolution of the model causes $P(p)$ to evolve toward the most susceptible state while ELCs move $P(p)$ toward the state in which it is least susceptible. Thus, rather than settling into equilibrium like the original genetic model, the evolution of the GCGM is characterized by a cyclic behavior: $P(p)$ periodically evolving toward a more susceptible state until its progress is reversed by an ELC.

1. Susceptibility of $P(p)$

ELCs like the one illustrated in Fig. 14 are the result of highly correlated behavior of the agents. By this we mean that a significant number of agents activate or deactivate at the same time step. This implies highly correlated behavior since in order to do so the actions of all of the agents involved must be identical for the T preceding time steps. It is initially surprising that such a high degree of correlation could arise in the GCGM because, unlike minority game agents, GCGM agents make their decisions stochastically. The probability of coincidence between the actions of a large group of agents will usually be very small. However, there are two groups of agents in the model whose behavior is well correlated. These two groups are those agents whose gene values lie within a certain small range δ of 0 and 1. We shall call agents belonging to these groups *zero agents* and *one agents* respectively. The degree of correlation of these agents is a decreasing function of δ, being a maximum for $\delta=0$. It is easy to show that if we set an upper limit on the fraction f_d of zero and one agents whose actions are not perfectly correlated over a period of T time steps, δ is given by

$$\delta = 1 - (1 - f_d)^{1/T}. \tag{69}$$

The precise value chosen for δ is not important since it serves only to give a measure of the population of zero and one agents. Therefore it is more convenient to choose a fixed value for δ which gives rise to values of f_d that lie within an acceptable range, rather than choosing a different δ for each value of T. In all that follows we shall take $\delta=0.02$ which gives $f_d < 0.33$ for $T \leqslant 20$.

In short, we see that the probability of highly correlated agent behavior increases rapidly as the number of zero and one agents increases. Thus, gene value distribution functions $P(p)$ which are biased to favor agents with gene values near $p=0.0$ and $p=1.0$ will be the most susceptible to ELCs.

Now that we have considered what forms of $P(p)$ are most susceptible we will discuss how $P(p)$ evolves in the GCGM. In the periods between the ELCs the number of active agents n_t given by Eq. (59) is a slowly varying function of time. As we remarked above, correlations between the behavior of large numbers of agents are expected to be rare, and so n_t will fluctuate slowly with time as individual agents activate and deactivate. Hence the results of Refs. [4,13] can be applied to the GCGM in these periods. References [4,13] describe how in the genetic model the agents self-segregate into two populations having low and high gene values—these two populations can be thought of as a "crowd" and an "anticrowd". Thus, although we have yet to consider what the effect of an ELC will be on $P(p)$, we can see that after such an event $P(p)$ will evolve continuously toward the extremized distribution described by Ref. [13]. From our discussion above we know that it is this type of extremized gene value distribution that is most susceptible to ELCs.

2. Triggers in the global action time series

We saw above that it is only the zero and one agents that can participate in the highly correlated behavior necessary for an ELC. Therefore, in order to think about what patterns in the global action time series might induce an ELC it is necessary to consider these agents. If at time step t $A_t=+h_t$ then each zero agent will lose while each one agent will win. If $A_t=+h_t$ for T consecutive time steps then immediately following the Tth time step a fraction $1-f_d$ of the active zero agents will deactivate while the same fraction of the inactive one agents will activate. Similarly, if $A_t=-h_t$ for T time steps then a fraction $1-f_d$ of the inactive zero agents will activate and a fraction $1-f_d$ of the active one agents will deactivate. Thus we can see that sequences of time steps in which $A_t=+h_t$ or $A_t=-h_t$ for T time steps will be important for the correlated agent activations and deactivations that make up an ELC.

Since such sequences are important in the occurrence of ELCs it would be useful to have an expression for the probability that they will occur. The first step is to derive expressions for the probability that $A_t=\pm h_t$. We shall see later that these expressions are important in their own right. From Eq. (6), substituting n_t for N, and the definition of $\Pi_{i;t}$ we have the following expressions for the probability that $A_t=\pm h_t$:

$$P[A_t=-h_t]=\begin{cases}\displaystyle\sum_{i=(1-l)n_t}^{n_t}\Pi_{i;t} & \text{if } h_t=-1,\\ \displaystyle\sum_{i=ln_t+1}^{n_t}\Pi_{i;t} & \text{if } h_t=+1,\end{cases}$$

$$P[A_t=+h_t]=\begin{cases}\displaystyle\sum_{i=0}^{(1-l)n_t-1}\Pi_{i;t} & \text{if } h_t=-1,\\ \displaystyle\sum_{i=0}^{ln_t}\Pi_{i;t} & \text{if } h_t=+1.\end{cases} \tag{70}$$

We shall see later that $n_t \gg 1$. Therefore we can use the continuous approximation of $\Pi_t(x)$ for $\Pi_{i;t}$. As demonstrated in

Ref. [12], $\Pi_t(x)$ will be a Gaussian. In the GCGM only active agents contribute to the mean μ and variance σ^2 of $\Pi_t(x)$. Therefore

$$\Pi(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2 \right], \tag{71}$$

where

$$\mu = n_t \overline{p_t}, \quad \sigma^2 = \sum_i z_{i;t} p_{i;t}(1 - p_{i;t}). \tag{72}$$

In the continuous approximation we have

$$P[A_t = -h_t] \approx \begin{cases} \displaystyle\int_{(1-l)n_t - 1/2}^{+\infty} \Pi_t(x)dx & \text{if } h_t = -1, \\ \displaystyle\int_{ln_t + 1/2}^{+\infty} \Pi_t(x)dx & \text{if } h_t = +1, \end{cases}$$

$$P[A_t = +h_t] \approx \begin{cases} \displaystyle\int_{-\infty}^{(1-l)n_t - 1/2} \Pi_t(x)dx & \text{if } h_t = -1, \\ \displaystyle\int_{-\infty}^{ln_t + 1/2} \Pi_t(x)dx & \text{if } h_t = +1, \end{cases} \tag{73}$$

where we have used the fact that since σ is of order unity the integrands in Eq. (73) will approximately vanish for $x < 0$ and $x > n_t$. This has allowed us to replace lower limits of 0 with $-\infty$ and upper limits of n_t with $+\infty$.

Finally we can express Eq. (73) in terms of erf functions as, for $h_t = \pm 1$,

$$P[A_t = -h_t] \approx \frac{1}{2}\left[1 \mp \mathrm{erf}\left(\frac{\frac{1}{2} \mp n_t \lambda_t^{\pm}}{\sigma\sqrt{2}} \right) \right],$$

$$P[A_t = +h_t] \approx \frac{1}{2}\left[1 \pm \mathrm{erf}\left(\frac{\frac{1}{2} \mp n_t \lambda_t^{\pm}}{\sigma\sqrt{2}} \right) \right], \tag{74}$$

where for convenience we have defined $\lambda_t^{\pm}$ as follows:

$$\lambda_t^- = \overline{p_t} - (1 - l), \quad \lambda_t^+ = \overline{p_t} - l. \tag{75}$$

The erf function is defined as

$$\mathrm{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt. \tag{76}$$

We can see from Eq. (74) that the expressions for the probability that $A_t = \pm h_t$ depend upon the values taken by h_t. For this reason it will not be possible in general to derive a simple expression for the probability that $A_t = \pm h_t$ for T consecutive time steps. Such an expression would depend upon the realization of the prediction time series during the specific T time steps under consideration. In the next section, however, we shall see that a simple expression can be derived in the case of the memoryless model.

Reference [30] presented an analysis of crashes in financial markets such as the one that occurred on the NASDAQ in April 2000. The authors propose that such crashes result from speculative bubbles in which large numbers of traders share the same unrealistic expectations of the future performance of the companies in question. These bubbles eventually burst, apparently in response to some event which acts as a trigger. We can draw a broad qualitative analogy between this and the GCGM. The extremized gene value distribution discussed in Sec. III D 1 corresponds to a state of the model in which large numbers of agents share the same unrealistic expectation that the global action will be equal to or the opposite of the prediction. It is this *speculative* distribution that is most susceptible to triggers that occur from time to time in the global action time series.

E. ELCs in the memoryless GCGM

We form the memoryless GCGM from the full model described in Sec. III B by taking $h_t = +1 \, \forall \, t$ in exactly the same way as we did in Sec. II. Our discussion of the susceptibility of the gene value distribution function $P(p)$ in Sec. III D 1 applies equally to the memoryless and the full GCGM. The equivalent of the patterns of T time steps in which $A_t = -h_t$ or $A_t = +h_t$ are those in which $A_t = -1$ or $A_t = +1$. In the memoryless case, the expressions in Eq. (74) become

$$P[A_t = \pm 1] \approx \frac{1}{2}\left[1 \pm \mathrm{erf}\left(\frac{\frac{1}{2} - n_t \lambda_t^+}{\sigma\sqrt{2}} \right) \right]. \tag{77}$$

This expression is no longer dependent on h_t, since $h_t = +1 \, \forall \, t$. Therefore, in the case of the memoryless GCGM, we can derive the following simple expression for the probability Λ_t that $A_t = \pm 1$ for T consecutive time steps:

$$\Lambda_t = \frac{1}{2^T}\left[1 - \mathrm{erf}\left(\frac{\frac{1}{2} - n_t \lambda_t^+}{\sigma\sqrt{2}} \right) \right]^T + \frac{1}{2^T}\left[1 + \mathrm{erf}\left(\frac{\frac{1}{2} - n_t \lambda_t^+}{\sigma\sqrt{2}} \right) \right]^T. \tag{78}$$

Figure 15 shows Λ_t given by Eq. (78) and the average waiting time, given by Λ_t^{-1}.

There are a couple of points to notice in Fig. 15. First of all the minimum of Λ_t does not occur at $\lambda_t^+ = 0$ but at a value of $\lambda_t^+ = 1/2n_t$. This results from the fact that in the case of $n_t^{+1} = ln_t$ our model tie breaks by declaring the global action $A_t = +1$. The most obvious feature of Fig. 15 is that Λ_t increases rapidly with increasing λ_t^+. This means that the probability of a trigger sequence occurring in the global action time series A_t increases with the deviation of $\overline{p_t}$ from l. Therefore we can consider the value of λ_t^+ to be controlling the probability that a trigger sequence will occur. Furthermore, note that Λ_t never vanishes and so the average waiting time never goes to infinity. Crucially, this means that regardless of the value of $\overline{p_t}$ there is always a nonzero probability that a trigger sequence will occur.

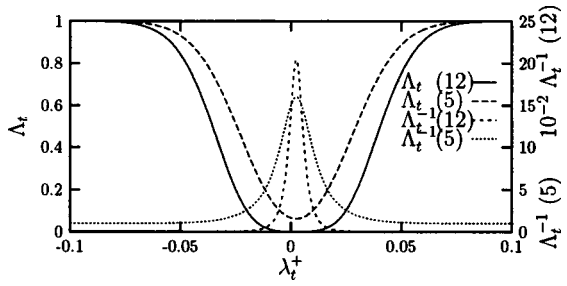

FIG. 15. The probability of occurrence Λ_t and the average waiting time Λ_t^{-1} for a pattern of T consecutive times steps in which $A_t=+1$ or -1. In both cases results are given for $T=5$ and 12 (given by the numbers in parentheses). We have taken $n_t=205$ and $\sigma_t=4.9$ which correspond to their mean values over the period described by Fig. 16 below. Note that, in the case of $T=12$, the right hand axis is scaled by a factor 10^{-2}.

1. Price time series in the memoryless GCGM

One of the advantages of the memoryless GCGM is that the dynamics of the price time series are particularly simple. We saw in Sec. III C that the probability of a price fall at time t is given by the same condition that the global action $A_t=-1$. Therefore, $P[\Delta\pi<0]=P[A_t=-1]$ given by Eq. (77). Thus the price will fall with probability $P<0.5$ if $\lambda_t^+<0$ and with probability $P>0.5$ if $\lambda_t^+>0$.

2. Example ELC

Now that we have introduced the memoryless GCGM we shall consider a specific example of an ELC. This will allow us to see how the elements discussed in Sec. III D are involved in ELCs in the GCGM. In order to do this, however, we need to introduce one further quantity which we shall call the *prediction performance* and denote by η_t. At time step t the value of η_t gives the number of previous consecutive time steps at which the prediction was the same as the global action $A_t=+h_t$. In the memoryless GCGM the value of η_t gives the number of consecutive time steps preceding t at which $A_t=+1$. The reason that η_t is useful is that during one of the so-called trigger sequences that we discussed in Sec. III D 2 η_t will become large and, thus, it can be used to identify these events.

Figure 16 shows the values of the global action A_t and the prediction performance η_t for a period of time in which an ELC occurs in the memoryless GCGM with $l=0.4$ and $T=12$. In order to show the behavior of the one and zero agents that we discussed in Sec. III D 1 we have also included the virtual points $v_{i;t}$ of an agent i which is inactive and has $p_{i;t}=1.0$ at the beginning of the time period shown. Since the actions of zero and one agents are anticorrelated (as we saw in Sec. III D 1) the virtual points of inactive one agents and active zero agents will always be the same. The same applies to the virtual points of active one and inactive zero agents. For this reason it is necessary to give $v_{i;t}$ for only one of these four groups in Fig. 16 since from this we can infer the virtual points of the others. Figure 16(b) also shows the probability $P[A_t=-1]$ that the global action $A_t=-1$ at each time step, given in terms of λ_t^+ by Eq. (77). Figure 16(c) shows the deviation of $\overline{p_t}$ from l, λ_t^+, over the same time period. The quantities depicted in the figure are those that play an important role in the mechanism that causes ELCs. Later, in Fig. 18, we shall demonstrate what effect the ELC described here has on external *observables* such as the price and the volume. In the following paragraphs we describe the significant features in Fig. 16. The paragraphs labeled $a-e$ correspond to the identically labeled time intervals in the figure.

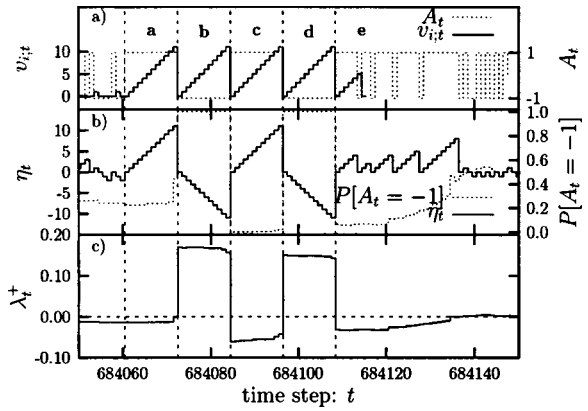

FIG. 16. An example of an ELC in the *memoryless* GCGM with $N=500$, $l=0.4$, $T=12$, $r=0.2$, and $m=3$. (a) The global action A_t and the virtual points $v_{i;t}$ of an agent i which initially is inactive and has a gene value of $p_{i;t}=1.0$. (b) The prediction performance η_t and the probability $P[A_t=-1]$ that the global action is -1. (c) The deviation of $\overline{p_t}$ from l, λ_t^+.

a. This sequence of T time steps in which $A_t=+1$ provides the trigger sequence discussed in Sec. III D 2. At each time step during this period the virtual points $v_{i;t}$ of the inactive one agents and active zero agents increases. When the model reaches the final time step in this period $v_{i;t}=T=12$ for both inactive one agents and active zero agents. The inactive one agents will then activate while the active zero agents will deactivate. The virtual points of these agents will then be reset to 0. Unless the numbers of one agents activating and zero agents deactivating are approximately equal, this correlated behavior will lead to a step change in the volume like the one depicted in Fig. 14. Only active agents contribute to $\overline{p_t}$ [see Eq. (65)] and so this instantaneous loss of zero agents and gain of one agents causes $\overline{p_t}$, and therefore λ_t^+, to undergo a step increase. We can see this clearly in Fig. 16(c).

b. Throughout the period of T time steps labeled b $\lambda_t^+\approx 0.16$. Equation (77) gives $P[A_t=-1]$ in terms of λ_t^+. The mean values of n_t and σ_t over the time period described by Fig. 16 are $\overline{n_t}=205$ and $\overline{\sigma_t}=4.9$. By substituting these values into Eq. (77) we can see that for $\lambda_t^+<-0.05$ and $\lambda_t^+>0.05$ $P[A_t=-1]\approx 0$ and 1, respectively. Therefore, if the magnitude of λ_t^+ exceeds 0.05 the model becomes quasideterministic at time step t. Thus, for the period labeled b, $A_t=-1$. The effect of this on the zero and one agents is exactly the opposite of that of period a; the virtual points of the inactive zero agents and the active one agents now increase at each time step. Once again $v_{i;t}=T=12$ for both these populations at the end of period b and so the inactive zero agents activate while the active one agents deactivate. Note that the inactive zero agents activating at the end of period b are not just those

that deactivated at the end of period a. Any zero agents that were inactive at the start of period a would have been unaffected by the trigger sequence; however they now activate along with those that previously deactivated. The result of this is that λ_t^+ does not return to the value of $\lambda_t^+ \approx 0.4$ that it had at the end of period a. That would have returned the model to its nondeterministic state. Instead, however, $\lambda_t^+ \approx -0.06$ and so at the start of period c $P[A_t=-1]=0$.

c. This period of T time steps in which $A_t=+1$ is identical to period a in terms of its effect on the one and zero agents. Thus, at the end of period c the inactive one agents reactivate and the active zero agents deactivate, resulting in the same step increase in λ_t^+ that we observed before. The important difference between periods a and c is that the first occurred stochastically (there was a significant nonzero probability that $A_t=\pm 1$ at each time step) whereas period c occurs quasideterministically (the probability that $A_t=+1$ for the T time steps in period c is $P \approx 1$).

d. Similarly, period d is identical to period b except that at the end $\lambda_t^+=-0.03$. Once again from Eq. (77) we can see that this gives $P[A_t=-1]>0$.

e. $P[A_t=-1]$ is no longer ≈ 0 and so the model returns to its usual stochastic behavior, A_t taking values -1 and $+1$ probabilistically. This represents the end of the ELCs since there is now no mechanism for the synchronized activation and deactivation that occurs during periods $a-d$. The model now returns slowly to the *equilibrium* state, in other words, the state that it is in once transients due to any ELCs have died away.

From the analysis that we have presented above we might have expected that the periodic synchronized activations and deactivations that we described above would continue indefinitely and that the model would never return to the stochastic state. One question that we did not address above, however, is that of how the model manages to break out of the deterministic behavior that it exhibits during the ELC. We have considered the effect of a period of T time steps in which $A_t=\pm 1$ on the zero and one agents in terms of agents activation and deactivation. However, we have not considered agent mutation. We shall see in what follows that it is agent mutation that allows the model to return to the stochastic state. During periods $b-d$ the one agents are active only at time steps at which $A_t=-1$ and the zero agents are active only when $A_t=+1$. Because of this the scores of zero and one agents are decreasing functions of time. Their scores are fixed when they are inactive and when they are active their individual actions are the inverse of the global action: $a_{i;t}=-A_t$. While these agents are inactive the model is favorable to them and so after T time steps they reactive. However, because the behavior of all these agents is so highly correlated, in doing so they change the dynamic of the model so that it is no longer favorable. This has a clear analogy with the phenomenon of *market impact* in economic systems.

Since the scores of the zero and one agents are decreasing functions of time during the ELC, the scores of these agents will rapidly reach the death score $s_{i;t}=-D$ at which they mutate. If $r \gg \delta=0.02$ [defined by Eq. (69)] then with a very high probability of $(r-\delta)/r$ a mutating zero or one agent will mutate to a gene value of $p_{i;t}>\delta$ or $p_{i;t}<1-\delta$, respectively.

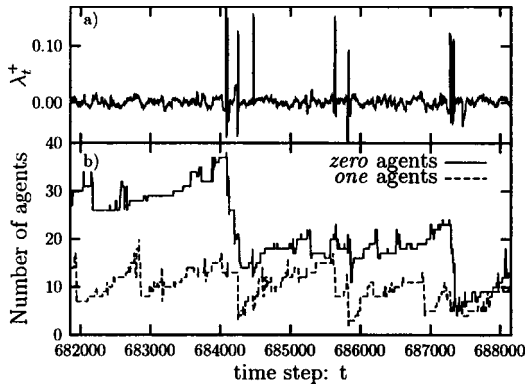

FIG. 17. (a) The evolution of the numbers of agents for which $p_{i;t}<\delta$ (zero agents) and $p_{i;t}>1-\delta$ (one agents) over a time period that includes that depicted in Fig. 16. (b) λ_t^+ over the same time period. The occurrence of ELCs is indicated by spikes in the λ_t^+ time series. The first ELC identifiable is that depicted in Fig. 16.

The result of this is that the population of zero and one agents that participate in the synchronized activations and deactivations steadily decreases throughout the ELC. Eventually there are no longer enough of these agents to maintain λ_t^+ at a magnitude greater than 0.05 and so $P[A_t=-1] \neq 0$ or 1. The model then returns to the stochastic state.

Note that the example depicted in this section does to some extent represent an idealized case. It is not guaranteed that the oscillations in λ_t^+ are of a sufficiently large magnitude that $P[A_t=-1]$, given by Eq. (77), takes only values ≈ 0 and ≈ 1. Therefore, some of the periods of T time steps in which A_t is consistently -1 or $+1$ will not occur deterministically. For this reason short interjected periods in which A_t is not consistent can occur between such periods as a to d in Fig. 16. The occurrence of ELCs in the model is robust against such stochastic fluctuations. However, during such an interjected time period agent mutation will act to bring λ_t^+ closer to 0.0. Therefore the longer such a period is the lower the probability that the ELC will continue.

3. Summary of ELCs in the memoryless GCGM

In this section we shall bring together the elements that we have introduced so far in order to give a broad overview of ELCs in the memoryless model. We demonstrated in Sec. III D 1 that in what we now call the stochastic state the agents migrate toward gene values $p=0.0$ and $p=1.0$. Therefore, while the model is in the stochastic state the number of zero and one agents increases. This increases the susceptibility of the gene value distribution $P(p)$ to any trigger sequences that might occur and also increases the duration of the next ELC. If a trigger sequence occurs in the evolution of the model then if there are enough zero and one agents an ELC will take place as described in the previous section. One of the effects of the ELC is to reduce the numbers of zero and one agents. This decreases the probability of a subsequent ELC occurring.

In order to make this clear we show in Fig. 17(b) the numbers of zero and one agents during the time period leading up to and after that depicted in Fig. 16. Figure 17(a)

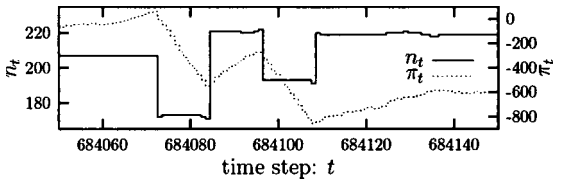

FIG. 18. Evolution of the number of active agents n_t (also know as the *volume*) and the price π_t [defined by Eq. (58)] over the same time period described by Fig. 16.

shows λ_t^+ over the same time period. As we saw in Sec. III E 2, ELCs can be identified by the spikes that occur in the λ_t^+ time series. We can see in Fig. 17 that, while the ELC that we examined in Sec. III E 2 causes a large decrease in the numbers of zero and one agents, significant numbers remain and so the initial ELC is followed by several smaller ones. The most important feature of Fig. 17 to notice, however, is that sudden decreases in the numbers of zero and one agents that occur at each ELC and the steady increase that these quantities exhibit in the periods between ELC.

So far, in order to understand the mechanism that leads to ELCs, we have concentrated on quantities that are internal to the model. In the economic analogy these would correspond to quantities whose values would be extremely hard to quantify, for example, the confidence of, or strategies adopted by, traders. However, as we remarked in Sec. III A, ELCs in the GCGM also affect quantities that are directly observable and quantifiable such as the *volume* and the *price*. We have plotted in Fig. 18 the volume and the price over the same time period described by Fig. 16. We can see from Fig. 18 that the oscillatory activation and deactivation of the zero and one agents that we described in Sec. III E 2 leads to corresponding oscillations in the volume. Note that the oscillatory nature of ELCs will lead to *volatility clustering* since each ELC contributes several large changes in the volume.

We saw in Sec. III E 1 that the probability that the price falls at time t $P[\Delta\pi_t<0]$ is equal to $P[A_t=-1]$ which is in turn given in terms of λ_t^+ by Eq. (77). Thus, we can see from the plot of $P[A_t=-1]$ in Fig. 16(a) that the oscillations in λ_t^+ will give rise to alternate periods in which the price rises and falls, as we see in Fig. 18. In terms of λ_t^+ the ensemble average excess demand is given by:

$$\langle\Delta\rangle = -\frac{n_t\lambda_t^+}{1-l}. \tag{79}$$

Thus we can see that the magnitude of the price change at each time step is proportional to $n_t|\lambda_t^+|$. Therefore, the overall fall in price depicted in Fig. 18 is due to the fact that the magnitude of the positive excursions of λ_t^+ during the ELC exceeds that of the negative excursions.

F. ELCs in the full GCGM: An idealized case

Before we examine some examples of numerically observed ELCs in the full GCGM we shall consider a theoretically idealized case. In this section we will assume that in equilibrium $\langle\overline{p_t}\rangle_t$ is equal to the optimal value $\overline{p}_{t,\mathrm{opt}}$ given by Eq. (16) of $1-l$ for $l<0.5$ and l for $l>0.5$. During the ELC

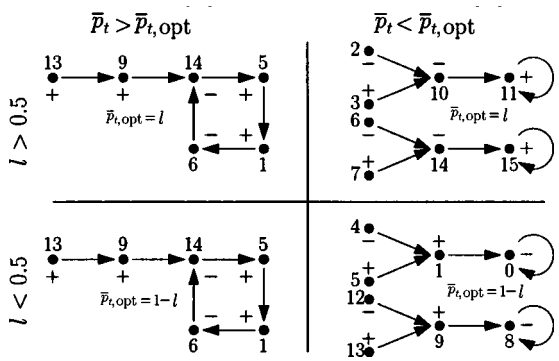

FIG. 19. Markovian transition diagrams for ELCs in the full GCGM. The state labels are as defined by Fig. 9 and ± signs give the value of h_t in each state.

we will assume that $\overline{p_t}$ oscillates between values that are greater and less than the equilibrium value by a magnitude greater than σ_t/n_t but less than $|2\overline{p_t},\mathrm{opt}-1|$ such that $P[A_t=-h_t]$, given by Eq. (74), takes only the values 0 and 1. This yields the following values for $P[A_t=-h_t]$.

For $l<0.5$,

$$\overline{p_t}<\overline{p}_{t,\mathrm{opt}}:\ P[A_t=-h_t]=\begin{cases}0 & \text{if } h_t=-1,\\ 1 & \text{if } h_t=+1,\end{cases}$$

$$\overline{p_t}>\overline{p}_{t,\mathrm{opt}}:\ P[A_t=-h_t]=\begin{cases}1 & \text{if } h_t=-1,\\ 1 & \text{if } h_t=+1.\end{cases} \tag{80}$$

For $l>0.5$,

$$\overline{p_t}<\overline{p}_{t,\mathrm{opt}}:\ P[A_t=-h_t]=\begin{cases}1 & \text{if } h_t=-1,\\ 0 & \text{if } h_t=+1,\end{cases}$$

$$\overline{p_t}>\overline{p}_{t,\mathrm{opt}}:\ P[A_t=-h_t]=\begin{cases}1 & \text{if } h_t=-1,\\ 1 & \text{if } h_t=+1.\end{cases} \tag{81}$$

As we saw in Sec. III D 2 it is sequences of T time steps in which $A_t=-h_t$ or $A_t=+h_t$ which act as the triggers for ELCs in the full GCGM. We can see from Eqs. (80) and (81) that $A_t=-h_t$ with probability $P=1$ when $\overline{p_t}>\overline{p}_{t,\mathrm{opt}}$. Thus, $\overline{p_t}>\overline{p}_{t,\mathrm{opt}}$ leads to the sequences of time steps which have the same effect on the zero and one agents as the sequences of time steps in which $A_t=-1$ that we saw in the memoryless GCGM. The situation when $\overline{p_t}<\overline{p}_{t,\mathrm{opt}}$ is more complicated. We can see from the above expressions that, for $l<0.5$, $A_t=-1$ while, for $l>0.5$, $A_t=+1$. Thus, for $\overline{p_t}<\overline{p}_{t,\mathrm{opt}}$ we expect sequences of time steps in which $A_t=-1$ and $A_t=+1$, respectively. However, it is not immediately apparent that $A_t=+h_t$ as we might expect.

By application of the same Markovian analysis that we used in Sec. II I 2 to the $m=1$ case, we can derive the state transition diagrams given in Fig. 19. The state labels are as defined by Fig. 9. We can see from Fig. 19 that when $\overline{p_t}<\overline{p}_{t,\mathrm{opt}}$ the transition diagrams each contain two attractor states in which $h_t=-1$ and $+1$ for $l<0.5$ and $l>0.5$, respectively. Thus, when the model is in these states $P[A_t=-h_t]=0$ and therefore $A_t=+h_t$. The reason that they can be divided into two congruent subdiagrams is that in each case, as

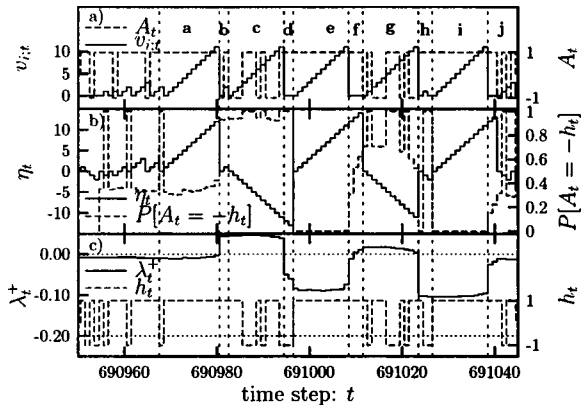

FIG. 20. An example of an ELC in the *full* GCGM with $N=500$, $l=0.6$, $T=12$, $r=0.2$, and $m=3$. (a) The global action A_t and the virtual points $v_{i;t}$ of an agent i which initially is inactive and has a gene value of $p_{i;t}=1.0$. (b) The prediction performance η_t and the probability $P[A_t=-h_t]$ that the global action is $-h_t$. (c) The deviation of $\overline{p_t}$ from l, λ_t^+, and the prediction h_t.

we remarked above, A_t takes only a single value. Therefore, the value of the memory bit that corresponds to the opposite value of A_t has no significance. This leads to the twofold state degeneracy that we observe; states that differ only in the value of this attribute are equivalent. Another feature to note in Fig. 19 is that, depending on the state that the model is in when it changes from $\overline{p_t}>\overline{p_{t,\mathrm{opt}}}$ to $\overline{p_t}<\overline{p_{t,\mathrm{opt}}}$, it may take several time steps to reach the attractor. Thus, unlike in the case of the memoryless GCGM, we should not expect oscillations with period T. There will likely be interjected time steps while the model finds the attractor.

We can see that this analysis will apply to the case of general m by considering Eq. (80). For $\overline{p_t}<\overline{p_{t,\mathrm{opt}}}$, $A_t=-1$ or $+1$ consistently. At the first time step after the activation and deactivation of one and zero agents the history will contain a mixture of -1's and $+1$'s. However, it is clear that after m $+1$ time steps it will contain only -1's and $+1$'s for $l<0.5$ and $l>0.5$, respectively, and the memory bit corresponding to this history will also take the same value. These states in which the history is $\{-1,-1,\ldots,-1\}$ and $h_t=-1$ or $\{+1,+1,\ldots,+1\}$ and $h_t=+1$ correspond to the attractor states in Fig. 19.

1. Example ELC

Now we shall look at a numerical example. Figure 20 gives the values of the global action A_t, the prediction h_t, the prediction performance η_t, λ_t^+ and $P[A_t=-h_t]$ given by Eq. (74). Once again we have also included the virtual points $v_{i;t}$ of an agent i which is inactive and has $p_{i;t}=1.0$ at the beginning of the time period shown. Figure 20 is equivalent to Fig. 16 except that we have additionally included the value of the prediction h_t. Note that in Fig. 20(c) the dotted lines indicate $\lambda_t^+=\overline{p_{t,\mathrm{opt}}}$ and $\lambda_t^+=1-\overline{p_{t,\mathrm{opt}}}$ which, in this case, correspond to $\lambda_t^+=0$ and $\lambda_t^-=0$, respectively. The paragraph labels below correspond to the labels in Fig. 20(b).

a. Period *a*, in which $A_t=+h_t$, provides the trigger that causes the activation and deactivation of zero and one agents.

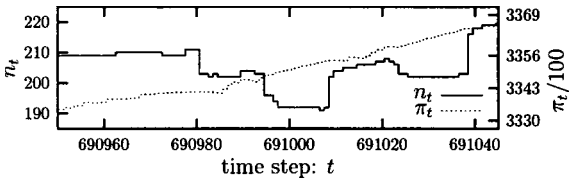

FIG. 21. Evolution of the number of active agents n_t (also known as the volume) and the price π_t [defined by Eq. (58)] over the same time period described by Fig. 20.

b,*c*. The (de) activation that is the result of *a* causes a step increase in λ_t^+ as we expect. However, the magnitude of λ_t^+ during *b* and *c* is not sufficiently large that $P[A_t=-h_t]=1$ when $h_t=+1$. In fact, as we see from Fig. 20(b), $P[A_t=-h_t]\approx 0.9$ when $h_t=+1$. Thus we can see that in this regard the realization of an ELC described here is not ideal in the sense discussed in the previous section. It is because $P[A_t=-h_t]\neq 1$ that for the second time step of *b* $A_t=+h_t$ resulting in these two interjected time steps. Throughout period *c* $A_t=-h_t$ as expected.

d,*e*,*f*. The interjected period *d* corresponds to the model finding one of the attractor states in which $A_t=+h_t$. During period *e*, $A_t=+h_t$ once again. However, after the reactivation and deactivation at the end of *e* $\lambda_t^+\approx 0$. Therefore, by chance $A_t=+h_t$ for the next three time steps as well, resulting in the interjected period *f*. This extra long period in which $A_t=+h_t$ allows some zero and one agents who had failed to (de)activate during the T time steps *e* to do so. Thus, λ_t^+ increases and so therefore does $P[A_t=-h_t]$ (for $h=+1$).

g,*h*,*i*. During *g*, $A_t=-h_t$ despite the fact that $P[A_t=-h_t]\approx 0.7$ (for $h=+1$). *h* once again corresponds the model finding the attractor state. *i* represents the final period in which $A_t=+h_t$ before λ_t^+ returns to approximately the equilibrium value and the ELC comes to an end.

Note that the periods like *c* and *g*, in which the magnitude of λ_t^+ is not sufficiently large that $P[A_t=-h_t]=1$ when $h_t=+1$, but in which nevertheless $A_t=-h_t$ for T or more time steps, occur with a much greater probability than they do in the memoryless model. The reason for this is that $P[A_t=-h_t]=1$ when $h_t=-1$, unless the magnitude of the oscillations in $\overline{p_t}$ is so great that $\lambda_t^-\approx 0$. Therefore, for any time steps during periods like *c* and *g* for which $h_t=-1$, $A_t=-h_t$ with probability $P=1$. We can see this clearly in the plot of $P[A_t=-h_t]$ in Fig. 20(b).

In Fig. 21 we show once again the volume n_t and the price π_t over the same period described by Fig. 20. We can clearly see the effect of the discussion in Sec. III C. Each time that $h_t=-1$ the price rises with probability $P=1$. Furthermore, the result of the deviation of the equilibrium value of $\overline{p_t}$ from $\overline{p_{t,\mathrm{opt}}}$ (see Sec. II G) is that it is more probable that the price will rise rather than fall when $h_t=+1$. These two effects ensure that in equilibrium (between ELCs) the price π_t is an increasing function of time. We can see in Fig. 21 that one of the effects of the ELC described in this section is to halt and even briefly reverse this continuous price rise.

The fundamental reason for this behavior is that, as we remarked in Sec. II, the agents in the GCGM are unable to

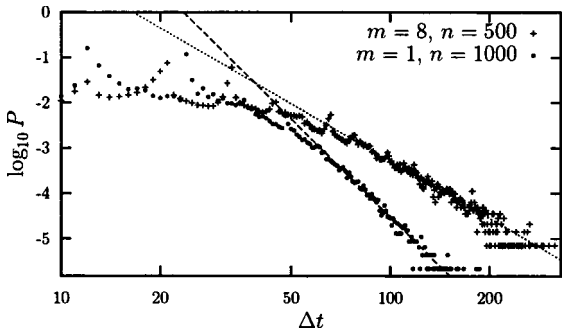

FIG. 22. Numerical distribution functions for the probability P of an ELC of duration Δt occurring in the GCGM with memory. The sets of data points correspond to the values of m and N indicated. The dashed lines represent inverse power law distributions with exponents −4.2 and −7.2.

control directly whether they will buy or sell at each time step. They can control only the probability with which they will follow the prediction h_t. As we saw in Sec. II the optimum behavior for the agents is to evolve such that the excess demand will be zero when h_t takes its most common value ($h_t=+1$ in the case of $l>0.5$). Therefore, in the asymmetric case of $l\neq 0.5$, the magnitude of the excess demand Δ will be large when h_t takes the opposite value. This represents the agents mistakenly believing the prediction, which in turn leads to an excess of buyers or sellers for $l<0.5$ and $l>0.5$, respectively.

Johansen and Sornette have provided evidence [28,30] that large price changes in financial markets are *outliers*. By this it is meant that the frequency with which large price changes occur cannot be predicted using the distribution of smaller price changes. From the results presented in this section, we can see that the large changes that occur in the volume during ELCs are also outliers in this sense. The distribution of volume changes between ELCs is such that changes of the magnitude observed during an ELC occur with a very small probability. As we saw in Sec. III D a different mechanism [i.e., that of the susceptibility of $P(p)$ and the occurrence of triggers in A_t] is responsible for the occurrence of ELCs, which therefore occur with a much greater probability.

G. ELC duration distribution

Finally in this section, we shall take a brief look at the distribution function of the duration Δt of ELC. Numerical simulations using the memoryless model suggest that the distribution function is of the form of a Gaussian. Figure 22 shows this distribution function in the full GCGM as well as straight line fits for $\Delta t \gtrapprox 60$. The figure suggests that, while short ELCs may follow a Gaussian duration distribution, longer changes are better fitted by an inverse power law.

Bak, Tang, and Wiesenfeld [31,32] introduced the theory of self-organized criticality which seeks to explain the occurrence of inverse power law distributions in nature. Reference [32] describes two- and higher dimensional systems, which are shown to self-organize to a critical point where the distribution of the magnitude and duration of *shocks* (called avalanches in Ref. [32]) follows an inverse power law with exponent ≈ -1.

If the duration distribution were found to display power law behavior, this might indicate that the theory of self-organized criticality was applicable to the GCGM. However, the measured exponents of −4.2 and −7.2 are very different from the value of −1 reported in Refs. [31,32]. Furthermore, the data in Fig. 22 are not sufficient to draw such a conclusion, and obtaining reliable data over a range of values of Δt sufficient to do so is prohibitively computationally intensive. This point requires further investigation.

H. Summary of ELC characteristics

We have seen how ELCs can occur in the GCGM as a result of the susceptibility of the gene value distribution $P(p)$ to triggers in the global action time series A_t. We also saw that between ELCs the self-segregation of the agents increases the susceptibility of $P(p)$ while this process is reversed during an ELC. Furthermore, we saw that an ELC in the memoryless GCGM leads to approximately periodic oscillations in a derived price time series. In contrast, the price time series in the full GCGM is a divergent quantity resulting from the inability of the agents to control $n_t^{+h_t}$ directly. In both models ELCs lead to approximately periodic volume n_t oscillations.

IV. CONCLUSIONS

In this section we provide a brief summary of the main results presented in this paper. For more details we refer the reader to Secs. II J and III H. In Sec. II C we showed that a simplified genetic model that made no use of memory is more efficient (for $l\neq 0.5$) in accessing the available resources than the original genetic model. Furthermore, in Sec. II E, we demonstrated that the reason for this is that in the memoryless model the agents can control n_t^{+1} directly, whereas in the original model n_t^{+1} is also a function of h_t. In Sec. II F we showed that, if h_t is generated by a random exogenous source with an appropriate mean value, then in terms of $\langle n_t^{+1}\rangle_t$ and $\sigma(n_t^{+1},t)$ the two models are indistinguishable. We noted that without such an external source the two models are distinguishable in terms of these two quantities for $l\neq 0.5$.

In Secs. II G and II H we demonstrated that the values taken by $\langle h_t\rangle_t$ are a direct result of the deviation of $\overline{p_t}$ from the optimal value of Eq. (16), but that this does not result from the nonzero standard deviation of the gene value distribution.

Section III showed that the memory introduces nonzero autocorrelations into the h_t and n_t^{+1} time series. Thus, for $l\neq 0.5$, the original and memoryless models are distinguishable even if h_t is derived from a random exogenous source, rather than being constant. Section III B provided a Markovian analysis that showed how the autocorrelation function of h_t resulted from the cycles performed by the model in the state space of the memory.

In Sec. III we introduced a different version of the genetic model (GCGM) in which the number of active agents was no

longer fixed. Section III E provided a detailed description of how a memoryless GCGM undergoes large self-induced shocks and discussed the effect of such shocks on external observables such as the volume and price. Section III F extended this discussion to the full GCGM including memory.

V. FUTURE DIRECTIONS

This paper leaves open several interesting questions that we hope will be addressed by future work. First of all in Sec. III we considered only the case in which the death score D is less than the confidence interval T. The result of this is that agents mutate over a shorter time scale than the period of oscillation of an ELC. We would expect that values of $D>T$ would lead to ELCs that persisted for many more periods.

We have taken the ratio of the number of points gained by an agent when $a_{i;t}=+A_t$ to those lost when $a_{i;t}=-A_t$ to be unity. Hod and Nakar [8] demonstrated that for values of this *prize-to-fine* ratio $R<1$ the self-segregation of the gene value distribution [3] is replaced by clustering behavior in which the agents tend to evolve toward $p_{i;t}=0.5$ in equilibrium (see also Refs. [33,34]). We saw in Sec. III D 1 that this clustered gene value distribution is not susceptible to the trigger sequences that cause ELCs. Further work is required to establish whether reducing the prize-to-fine ratio suppresses the occurrence of ELCs.

Most importantly, we showed in Secs. II G and II H that for the original genetic model in equilibrium $\overline{p_t}$ deviates from the optimal value given in Eq. (16). This apparently small deviation is important since it determines the values taken by $\langle h_t \rangle_t$ and the magnitude of the autocorrelations observed in the h_t time series. We have shown that since this effect is not present in the memoryless genetic model it does not result from the finite standard deviation of the gene value distribution, but must instead result from the action of the memory. We hope that future work will provide clarification of this point.

Finally, more work is required to establish to what extent, if any, the ELC duration distribution function described in Sec. III G represents an inverse power law and to clarify the effect of memory on this function.

ACKNOWLEDGMENTS

We are extremely grateful to P. M. Hui and T. S. Lo (Chinese University of Hong Kong) for detailed discussions, and for sharing their results with us. R.K. is supported by the EPSRC.

[1] W. B. Arthur, Am. Econ. Rev. **84**, 406 (1994).
[2] D. Challet and Y.-C. Zhang, Physica A **246**, 407 (1997).
[3] N. F. Johnson, P. M. Hui, R. Johnson, and T. S. Lo, Phys. Rev. Lett. **82**, 3360 (1999).
[4] N. F. Johnson, D. J. T. Leonard, P. M. Hui, and T. S. Lo, Physica A **283**, 568 (2000).
[5] E. Burgos and H. Ceva, Physica A **284**, 489 (2000).
[6] E. Burgos, H. Ceva, and R. P. J. Perazzo, Phys. Rev. E **65**, 036711 (2002).
[7] E. Burgos, H. Ceva, and R. P. J. Perazzo, Phys. Rev. E **64**, 016130 (2001).
[8] S. Hod and E. Nakar, Phys. Rev. Lett. **88**, 238702 (2002).
[9] K. Chen, B.-H. Wang, and B. Yuan, e-print cond-mat/0305673.
[10] N. F. Johnson, M. Hart, and P. M. Hui, Physica A **269**, 1 (1999).
[11] A. Cavagna, Phys. Rev. E **59**, R3783 (1999).
[12] T. S. Lo, P. M. Hui, and N. F. Johnson, Phys. Rev. E **62**, 4393 (2000).
[13] P. M. Hui, T. S. Lo, and N. F. Johnson, Physica A **288**, 451 (2000).
[14] T. S. Lo, MS thesis, Physics Department, Chinese University of Hong Kong, 2000.
[15] I. Giardina and J.-P. Bouchaud, e-print cond-mat/0206222.
[16] D. Challet, M. Marsili, and Y.-C. Zhang, Physica A **294**, 514 (2001).
[17] V. M. Eguíluz and M. G. Zimmermann, Phys. Rev. Lett. **85**, 5659 (2000).
[18] M. L. Hart, D. Lamper, and N. F. Johnson, Physica A **316**, 649 (2002).
[19] P. Jefferies, D. Lamper, and N. F. Johnson, Physica A **318**, 592 (2003).
[20] P. Jefferies and N. Johnson, e-print cond-mat/0207523.
[21] N. F. Johnson, M. Hart, P. M. Hui, and D. Zheng, Int. J. Theor. Appl. Finance **3**, 443 (2000).
[22] D. Lamper, S. D. Howison, and N. F. Johnson, Phys. Rev. Lett. **88**, 017902 (2002).
[23] P. Jefferies, M. Hart, P. Hui, and N. Johnson, Eur. Phys. J. B **20**, 493 (2000).
[24] F. Slanina and Y.-C. Zhang, Physica A **272**, 257 (1999).
[25] I. Giardina, J.-P. Bouchaud, and Mézard, Physica A **299**, 28 (2001).
[26] J.-P. Bouchaud and R. Cont, Eur. Phys. J. B **6**, 543 (1998).
[27] Y.-C. Zhang, Physica A **269**, 30 (1999).
[28] D. Sornette and A. Johansen, J. Risk **4**(2), 69 (2001).
[29] T. S. Lo, S. W. Lim, P. M. Hui, and N. F. Johnson, Physica A **287**, 313 (2000).
[30] A. Johansen and D. Sornette, Eur. Phys. J. B **17**, 319 (2000).
[31] P. Bak, C. Tang, and K. Wiesenfeld, Phys. Rev. Lett. **59**, 381 (1987).
[32] P. Bak, C. Tang, and K. Wiesenfeld, Phys. Rev. A **38**, 364 (1987).
[33] E. Burgos, H. Ceva, and R. Perazzo, Phys. Rev. Lett. **91**, 189801 (2003).
[34] S. Hod and E. Nakar, Phys. Rev. Lett. **91**, 189802 (2003).