

## Dynamics of computational ecosystems

J. O. Kephart, T. Hogg, and B. A. Huberman

*Xerox Palo Alto Research Center, Palo Alto, California 94304*

(Received 24 October 1988)

Recently, Huberman and Hogg [in *The Ecology of Computation*, edited by B. A. Huberman (North-Holland, 1988), pp. 77–115] analyzed the dynamics of resource allocation in a model of computational ecosystems which incorporated many of the features endemic to large distributed processing systems, including distributed control, asynchrony, resource contention, and cooperation among agents and the concomitant problems of incomplete knowledge and delayed information. In this paper we supplement an analysis of several simple examples of computational ecosystems with computer simulations to gain insight into the effects of time delays, cooperation, multiple resources, inhomogeneity, etc. The simulations verify Huberman and Hogg's prediction of persistent oscillations and chaos, and confirm the Ceccatto-Huberman [Proc. Natl. Acad. Sci. U.S.A. **86**, 3443 (1989)] prediction of extremely long-lived metastable states in computational ecosystems. Extending the analysis to inhomogeneous systems, we show that they can be more stable than homogeneous systems because agents with different computational needs settle into different strategic niches, and that overly clever local decision-making algorithms can induce chaotic behavior.

### I. INTRODUCTION

The emergence of distributed parallel processing in large computer networks<sup>1</sup> has opened up an interesting new frontier of computer research. The growing interconnection of diverse processors in networks allows for a loosely coupled form of concurrency with complex interdependencies, leading to self-regulating computational entities very different in nature from their individual components. Competition and cooperation abound in the form of resource contention and the sharing of information from databases, sensors, and other knowledge sources, and even a sort of reproductive behavior can be achieved by spawning remote processes. Given the many characteristics it shares with biological and social organizations, one may regard such a collection of interacting computational agents as a *computational ecosystem*.

Some systems which are not necessarily distributed among many computers can also be thought of as computational ecosystems. Robots, monitoring systems,<sup>2</sup> process schedulers for integrated circuit fabrication,<sup>3</sup> and other systems which interact directly with the physical world through sensors or motors must respond in real time to information which is not only constantly changing, but often inconsistent and incomplete. This is due to inherent limitations in the accuracy and interpretation of sensory signals and in the time available to obtain or process information from other components of the system. One approach to these problems is to organize the system as a loosely coupled collection of agents which specialize in various strategies for dealing with different contingencies and compete among themselves for resources with which to address the overall goals of the system.

Just as the computation within a computational ecosystem is distributed, so must be the resource allocation. Since the composition of such a system continually

evolves in unforeseen directions, a central controller cannot be kept up to date about the state of the system, resulting in delayed responses to new opportunities. Even more importantly, the system must continue to operate even if a few machines or agents, including any central one, fail. The necessity for local decisions about resource management and communication in order to ensure robustness to change and failure raises a number of basic issues concerning the behavior and, ultimately, the design of computational ecosystems. In particular, it is vital to understand how the overall behavior of a group of cooperating agents depends upon that of the individuals which comprise it, and what implications this relationship has for system design principles and heuristics. This issue is also of central importance to proposed problem-solving strategies in distributed artificial intelligence systems.<sup>4</sup>

Some qualitative insight into this and related issues has been provided by exploiting superficial analogies between computational network architectures and various human and natural organizational structures, such as groups of human experts,<sup>4</sup> the scientific community,<sup>5</sup> economic markets,<sup>6,7</sup> the Society of Mind,<sup>8</sup> and biological ecosystems.<sup>9</sup> Such analogies have led to a number of proposals for the design of computational ecosystems. For example, one method for deciding among conflicting messages is due-process reasoning,<sup>4</sup> which entails procedures very similar to those used to validate scientific claims in the face of possibly conflicting results. In analogy to economic markets, limited resources such as memory and processor time can be allocated in a distributed fashion by allowing processes to bid for them, resulting in a more efficient system than can be achieved using standard queueing techniques.<sup>5</sup> A further improvement of this approach allows several bidding strategies for the processes, leading to the interesting question of the existence of evolutionarily stable equilibria<sup>10</sup> for the system as a whole.

Recently, Huberman and Hogg<sup>11</sup> analyzed the dynamical behavior of computational ecosystems from a more quantitative perspective. They found that such systems can display a panoply of behavioral regimes which, depending on particular system parameters, are characterized by fixed points, oscillations, or even chaos. In this paper we elaborate their theory, compare it to computer simulations, and show that its predictions are generally quite accurate when there are at least a few hundred agents in the system. We also present a number of new properties of computational ecosystems. These include our observation that the presence of several different delays in a system do not greatly affect its tendency to oscillate, that enhancing the decision-making abilities of some of the individual agents can actually decrease overall system performance and lead to chaotic behavior, and that systems can remain in nonoptimal metastable states for extremely long periods of time before escaping to the globally optimum state, in agreement with theoretical predictions.<sup>12</sup> After studying the dependence of the system's behavior upon several different parameters, we are able to offer a number of general heuristics which ought to be helpful in the design of such systems. We show that undesirable oscillatory behavior can be reduced or eliminated by purposely introducing randomness into the decision procedures used by the agents, or by introducing agents into the system which use different decision procedures.

Specifically, in Sec. II we discuss our basic model, which incorporates many of the characteristic features of computational ecosystems. In Sec. III we present various aspects of the model's behavior and compare the results of theory and simulations. Finally, we discuss implications and possible extensions of our work in Sec. IV.

## II. MODEL OF COMPUTATIONAL ECOSYSTEMS

To investigate the behavior of computational ecosystems, we consider a particular model which incorporates the essential features described above. These include distributed control, asynchrony in execution, resource contention, and cooperation among agents, along with the concomitant problems of incomplete knowledge and delayed information. We suppose that a large number of tasks are to be performed on a network of interconnected computers. The tasks, which could be generated continually by user requests or as spawned processes, are managed by active agents, which are responsible for choosing among various computational resources to perform the task. In the simplest case, these are hardware resources, such as execution time on a computer or use of a communication line, but more generally they could include use of specific software packages or access to information in various databases.

The basic model considered in this paper,<sup>11</sup> illustrated in Fig. 1, consists of  $A$  agents which are free to choose among  $R$  resources according to the *perceived* (i.e., not necessarily correct) payoff for using each resource. The payoff is related to actual computational measures of performance, such as time to complete the task, accuracy of the solution, amount of memory used, etc. Competition and cooperation among the agents are taken into account

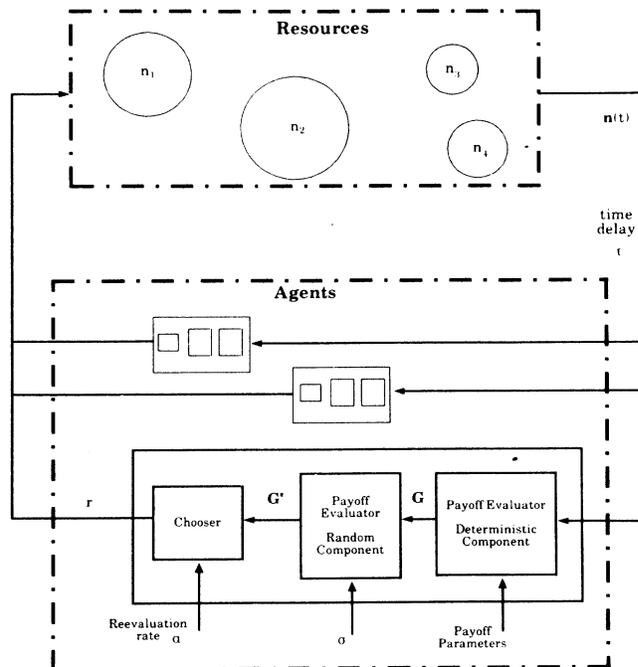


FIG. 1. Model of a computational ecosystem with  $n_i$  agents using the  $i$ th resource. The circles denote computational resources, and the solid rectangles, agents in the system whose behavior is determined by three computations. Lowest solid rectangle is expanded to show then in detail.

by allowing the payoff for using a particular resource to depend on the number of agents already using it. For example, in a purely competitive environment, the payoff for using a particular resource  $r$  would decrease monotonically with the fraction of agents already using it,  $f_r(t)$ . Alternatively, the agents using a resource could assist one another in their computations, as might be the case if the overall task could be decomposed into a number of subtasks. In this case, the payoff might increase as more agents used that resource. Each agent evaluates the payoff associated with each resource asynchronously at an average rate  $\alpha$  and switches to the resource with the highest payoff. To account for the fact that its information about the current state of the system can be somewhat imperfect and delayed, we add a normally distributed quantity with zero mean and standard deviation  $\sigma$  to each payoff, and delay the information available to each agent by a time  $\tau$ .

Due to the variation in reevaluation times and the uncertainty of information, the dynamics of the system must be described probabilistically. The system's behavior is fully specified by the probability distribution functional which describes the probability of obtaining a particular evolution of the system. Unfortunately, evaluating it in realistic situations is extremely difficult. However, it is possible to find coarser-grained quantities which provide clear and reasonably complete characterizations of the overall system dynamics, and are also amenable to theoretical analysis and computer simulation.

Theoretical analysis<sup>11</sup> of the model, described in some

detail in Sec. III, leads to a differential-delay equation for the time-evolution of the average of the vector  $f(t)$ , whose components specify the fraction of agents using each resource at time  $t$ . The average behavior of various quantities of interest which depend upon  $f(t)$  (e.g., correlations, oscillation periods, and total system performance) can be evaluated with a mean-field approximation. Correspondingly, our simulation is a straightforward, event-driven implementation of the model. In this case, the various quantities of interest are evaluated over a single simulation run and the result averaged over several runs.

Theoretical analysis and computer simulation are complementary in at least two important respects. The approximations of the theory, which involve replacing averages of functions by functions of averages (mean-field approximation), and those of the simulation, which are related to the finite sampling of the probability-distribution functional, are completely different in nature, so that agreement between the two is a clear indication of correct conclusions about the behavior of the model. In addition to corroborating the predictions of mean-field theory, the simulations enable finite-size and fluctuation effects to be measured.

The very differences that make the two techniques complementary also create certain subtleties in their comparison. For example, in a system with sustained oscillations, the period and amplitude are trivial to extract from the theoretical predictions of  $\langle f(t) \rangle$ . However, an average of  $f(t)$  over several simulation runs yields an  $\langle f(t) \rangle$  which initially oscillates with the same period and amplitude, but eventually decays and settles to some fixed value. The reason is that, while any one simulation run bears a strong qualitative resemblance to that predicted by the theory, random-phase drifts between different simulation runs cause  $\langle f(t) \rangle$  to settle to a fixed value on a time scale which depends on the number of agents. Thus, in the simulations, oscillation periods and amplitudes cannot be extracted from  $\langle f(t) \rangle$ . However, as described in Appendix B, they *can* be extracted from correlation functions measured within a single simulation run and then averaged over several runs. As we shall see in Sec. III, the results of this method agree well with theoretical predictions.

### III. BEHAVIOR OF COMPUTATIONAL ECOSYSTEMS

In this section, we present a number of behavioral phenomena which can be expected in computational ecosystems of the type described above, starting with the simplest systems and gradually working towards more complicated ones. Theoretical analysis and simulations are used to determine the conditions under which various types of behavior occur. First, in Sec. III A, we examine a simple system of identical agents which compete for two resources. Next, in Sec. III B, we introduce time delays into the system and show that they can induce persistent oscillations. Then, we complicate the system still further by making the agents somewhat cooperative (Sec. III C) or by introducing more resources (Sec. III D), finding that under such circumstances the behavior can become chaotic. In Sec. III E, we remove the restriction

that all agents behave identically and show that inhomogeneity can increase the stability of the system. Finally, we explore the consequences of introducing more sophisticated algorithms for choosing resources (Sec. III F) and discuss issues of metastability in systems with more than one mode of behavior (Sec. III G).

#### A. Exact solution for zero delay

Consider a system with two resources and  $A$  identical agents. In this case, the probability distribution  $P_i(t)$  for  $i$  agents to be using resource 1 at time  $t$  evolves according to<sup>11</sup>

$$\frac{d\mathbf{P}(t)}{dt} = \alpha \underline{M}\mathbf{P}(t), \quad (1)$$

where  $\alpha$  is the rate at which each agent reevaluates its choice, and  $\underline{M}$  is the tridiagonal matrix derived in Ref. 11 and which is given by

$$M_{ij} = \begin{cases} (A-j)\rho(j), & j=i-1 \\ -[(A-j)\rho(j)+j(1-\rho(j))], & j=i \\ j(1-\rho(j)), & j=i+1, \end{cases} \quad (2)$$

and  $\rho(i)$  is the probability that an agent will choose resource 1 if  $i$  agents are already using it.

In general,  $\rho$  might be an arbitrary function of  $i$ , which could be thought of as characterizing an arbitrary deterministic or probabilistic classifier. For example, if, for all  $0 \leq i \leq A$ ,  $\rho(i)$  is equal to either 1 or zero,  $\rho(i)$  corresponds to a deterministic classifier. Otherwise, the decision regions are somewhat fuzzy, the most extreme example of which is  $\rho(i) = \frac{1}{2}$  for all  $i$ . However, in the model of Fig. 1, the form of  $\rho(i)$  is restricted somewhat because it is determined by payoffs. In terms of  $G_1(i)$  and  $G_2(i)$ , the payoffs for using resources 1 and 2, and the uncertainty  $\sigma, \rho$  is given by<sup>11</sup>

$$\rho(i) = \frac{1}{2} \left[ 1 + \operatorname{erf} \left[ \frac{G_1(i) - G_2(i)}{2\sigma} \right] \right], \quad (3)$$

where  $\operatorname{erf}(x)$  denotes the error function of  $x$ .

In order to facilitate comparison of systems with different numbers of agents, it is convenient to express the payoffs and the preference probability  $\rho$  in terms of the *fraction* of agents using resource 1, introduced earlier:  $f \equiv i/A$ . With this reinterpretation, Eq. (3) remains valid if  $f$  is substituted for  $i$ .

If the agents are in competition with one another, the payoffs decrease monotonically with their number, as illustrated in Fig. 2(a), because each resource becomes saturated as it is used by more agents. The corresponding  $\rho(f)$  is displayed in Fig. 2(b) for two values of the uncertainty,  $\sigma = 0$  and 0.125.

When there is no uncertainty, it is intuitively clear that the system will tend towards an equilibrium  $f_0$  in which the two payoffs are equal. If  $f < f_0$ , then  $G_1(f) < G_2(f)$ , so the next agent to reevaluate its choice will choose resource 1. Therefore,  $f$  will either stay the same or increase. Similarly, if  $f > f_0$ ,  $f$  will either decrease or stay the same. A simple extension of this reasoning shows

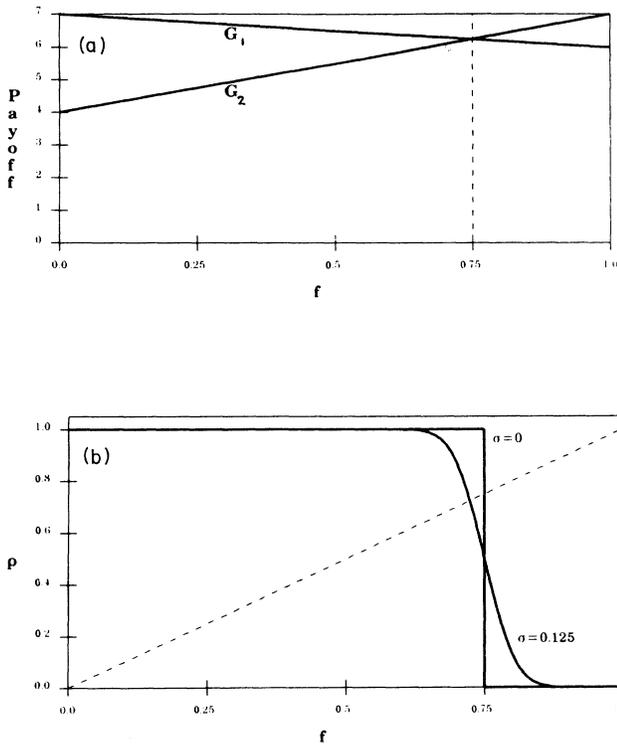


FIG. 2. (a) Monotonically decreasing payoffs  $G_1 = 7 - f_1$ ,  $G_2 = 7 - 3f_2$ . (b) Preference probability function  $\rho(f)$  for these payoffs for two values of the uncertainty parameter  $\sigma = 0$  and 0.125. The dashed line is  $\rho = f$  and its intersection with the  $\rho(f)$  curves gives the equilibrium values,  $f_0 = 0.75$  and 0.724, respectively.

that, for  $\sigma > 0$ , the system will tend towards an equilibrium value given by the solution to

$$f_0 = \rho(f_0). \quad (4)$$

This is confirmed by Fig. 3, which plots the mean and standard deviation of the probability distribution  $P(f, t)$

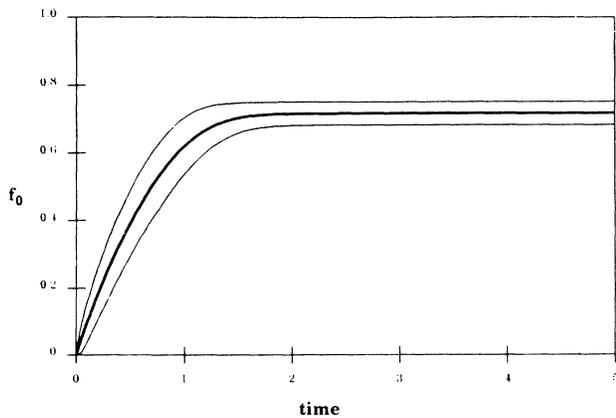


FIG. 3. Mean and standard deviation of  $P(f, t)$  calculated from Eq. (1) using the payoffs of Fig. 2,  $\sigma = 0.125$ , and  $A = 25$ . Initially, all of the agents are using resource 2. The dark line shows the mean value of  $f$ ; the lighter lines which flank it are one standard deviation from the mean.

given by the solution to Eq. (1). As  $t \rightarrow \infty$ , the mean of  $P(f, t)$  monotonically approaches a value close to that given by Eq. (4).

By setting the left-hand side of Eq. (1) equal to zero, and rewriting Eq. (2) and  $\rho$  in terms of  $f$ , one can show that, as  $A \rightarrow \infty$ , the equilibrium distribution  $P(f, t \rightarrow \infty)$  approaches a Gaussian with mean  $f_0$  given by Eq. (4) and variance

$$s^2 = \frac{1}{A} \frac{f_0(1-f_0)}{1-\rho'(f_0)}. \quad (5)$$

These results are compared with the actual values obtained from the eigenvectors of  $\underline{M}$  in Fig. 4. As can be seen, they are accurate to within a few percent when the number of agents exceeds a few dozen.

Since  $P(f, t)$  is arbitrarily narrow for sufficiently large  $A$ , the most important aspect of the distribution is its average  $\langle f \rangle$ . By making a mean-field approximation in which  $\langle \rho(f) \rangle$  is replaced by  $\rho(\langle f \rangle)$ , Eq. (1) becomes

$$\frac{d\langle f \rangle}{dt} = -\alpha[\langle f \rangle - \rho(\langle f \rangle)]. \quad (6)$$

Since each agent interacts with each of the other agents through the global variable  $f$ , the interactions in the system are effectively infinitely ranged, a situation in which mean-field theory is expected to work very well.

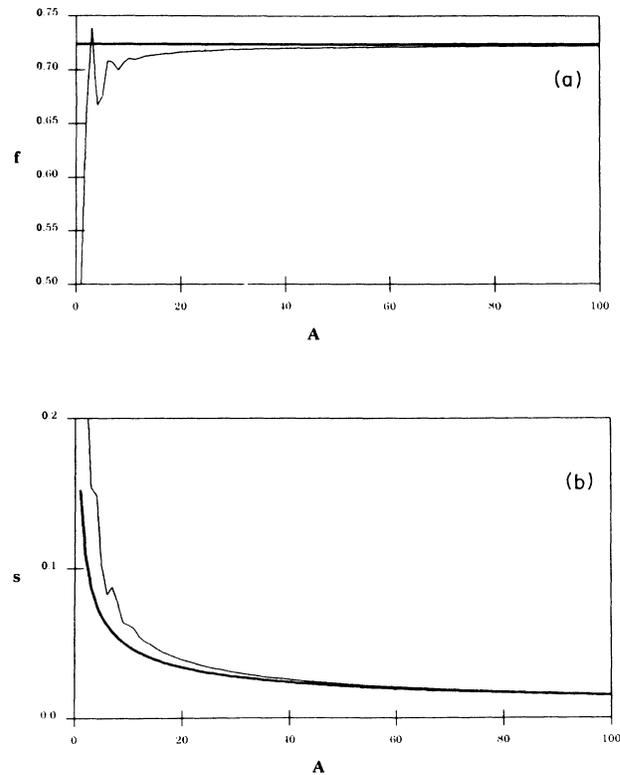


FIG. 4. (a) Mean and (b) standard deviation of the distribution  $P(f, t \rightarrow \infty)$  vs the number of agents  $A$ . The light curves are calculated from Eq. (1). The dark lines are calculated from Eqs. (4) and (5), which are probably accurate in the limit  $A \rightarrow \infty$ .

**B. Time delays**

Agents in a computational ecosystem will often base decisions on information which is no longer current. In this section we show how the presence of time delays in a system can lead to oscillations and instabilities. As an example, we consider a system with the same payoffs and uncertainty as in Fig. 3, but with a nonzero time delay  $\tau$ . Figure 5(a) shows two typical simulation runs with 200 agents, one with  $\beta \equiv \alpha\tau = 0.04$ , and the other with  $\beta = 0.40$ . For  $\beta = 0.04$ , the simulation looks very similar to Fig. 3; the system evolves essentially monotonically towards the fixed value  $f_0 = 0.724$  predicted by the theory, with some small fluctuations due to the finite number of agents. However, for  $\beta = 0.40$ , the fixed point is unstable, and  $f(t)$  exhibits oscillations which fluctuate slightly in phase and amplitude.

In order to understand this behavior, we would like to analyze the system in a manner analogous to that of Sec. III A. A phenomenological approach leads to the following form:<sup>11</sup>

$$\frac{df(t)}{dt} = -\alpha[f(t) - \rho(f(t - \tau))], \quad (7)$$

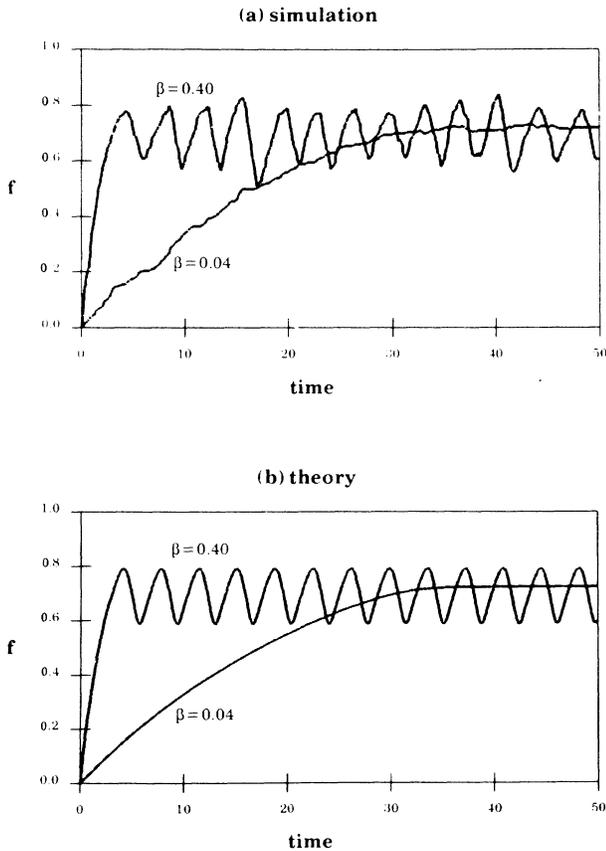


FIG. 5. (a) Simulated (200 agents) and (b) theoretical  $f(t)$  for a system of purely competitive agents with the same payoffs and uncertainty as in Fig. 3. In each case, results are displayed for two different values of  $\beta \equiv \alpha\tau$ : 0.04 and 0.40. The initial conditions are  $f(t) = 0$  for  $-\tau < t < 0$ . Time is measured in units of  $\tau$ .

where, for the sake of legibility, we have dropped the explicit angle brackets on  $\langle f \rangle$ . This equation can also be obtained from a master equation.<sup>13</sup> In what follows we confirm the validity of this equation by extensive computer simulations.

Figure 5(b) displays the numerical solutions of Eq. (7) for the same values of  $\beta$  as in Fig. 5(a). The agreement between the theory and simulation in Fig. 5 is excellent. For  $\beta = 0.04$ , the theoretical solution tends toward the correct asymptote at nearly the same rate as the simulation, and for  $\beta = 0.40$  the period and amplitude of the oscillation and the average value of  $f$  are in agreement to within a few percent.

The conditions under which oscillations and instabilities occur can be obtained by linearizing Eq. (7) in the neighborhood of the fixed point. Defining the variation around the fixed point  $\delta(t) = f(t) - f_0$ , we obtain

$$\frac{d\delta(t)}{dt} = -\alpha[\delta(t) - \gamma\delta(t - \tau)], \quad (8)$$

where  $\gamma = \rho'(f_0)$ . The behavior of solutions to this equation can be characterized by making the substitution<sup>14</sup>  $\delta(t) = e^{a\xi t}$ , which yields

$$\gamma e^{-\beta\xi} - \xi - 1 = 0, \quad (9)$$

where  $\beta \equiv \alpha\tau$ . In general, Eq. (9) has infinitely many complex solutions  $\xi = \xi_r + i\xi_i$ . The asymptotic behavior of the solution to Eq. (8) can be determined from the roots of Eq. (9) with the largest real parts.

The general behavior of Eq. (8) for various regimes of  $\gamma$  and  $\beta$  is represented in the phase diagram of Fig. 6. The parameters corresponding to the two curves in Fig. 5(b) are denoted by asterisks. In the shaded region, there is at least one solution to Eq. (9) with  $\xi_1 > 0$ , so the fixed point  $f_0$  is unstable. For  $\gamma < -1$ , the boundary between stability and instability depends upon  $\beta$ . We shall henceforth refer to this critical threshold for instability as  $\beta_2$ . Within the stable region, the light curve indicates the

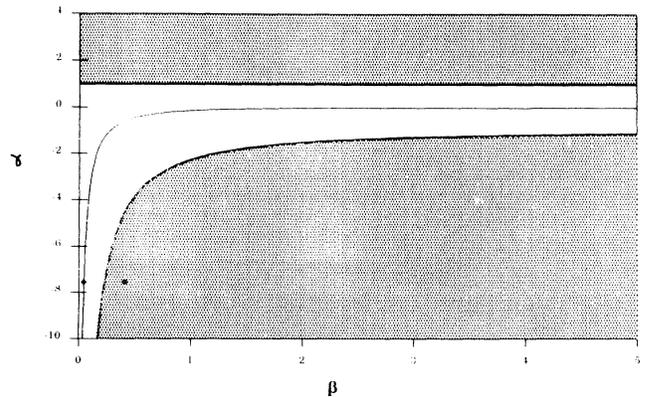


FIG. 6. Stability and oscillation regimes as a function of  $\beta$  and  $\gamma \equiv \rho'(f_0)$  for an arbitrary system with two resources. Shaded area represents the unstable region. Within the stable region (unshaded), the light curve is the boundary below which nonoscillatory decay to the fixed point is not possible; instead, all solutions exhibit damped oscillations. Asterisks indicate the parameters used in Fig. 5.

boundary below which the solution to Eq. (8) exhibits damped oscillations about the fixed point; this critical threshold for oscillation will be referred to as  $\beta_1$ . Note that, for all  $\gamma$ ,  $\beta_1 < \beta_2$ ; i.e., the threshold for oscillation is always less than that for instability. Analytic expressions for the critical values  $\beta_1$  and  $\beta_2$  can be derived from Eq. (9):

$$\beta_1 = -\gamma^{-1} e^{-(1+\beta_1)} \approx 1/(|\gamma|e-1) \text{ when } |\gamma| \gg 1/e \quad (10)$$

and

$$\beta_2 = \frac{\cos^{-1}(1/\gamma)}{(\gamma^2-1)^{1/2}}, \quad (11)$$

where the principle value of the arccos is to be taken.

For the payoffs of Fig. 2 and  $\sigma=0.125$ , Eqs. (10) and (11) yield the oscillation and instability thresholds  $\beta_1=0.0464$  and  $\beta_2=0.2271$ . The decay rate  $\tau_r \equiv 1/|\alpha\xi_r|$  and the oscillation frequency  $\nu \equiv \alpha\xi_i/2\pi$  can be calculated as a function of  $\beta$  from Eq. (9) for  $\beta < \beta_2$ . For  $\beta > \beta_2$ , the fixed point is unstable, and the amplitude of oscillation grows until it is limited by nonlinearities in  $\rho(f)$ . In this unstable regime, Eq. (7) can be solved numerically to obtain  $\nu$  and  $s$ , the rms amplitude of the oscillations about the average value of  $f$ . In order to compare these quantitative theoretical predictions with the results obtained by simulation, we calculate the autocorrelation function  $C(t')$  defined in Appendix B. Estimates of  $s$ ,  $\xi_r$ , and  $\xi_i$  (and hence  $\tau_r$  and  $\nu$ ) are extracted from  $C(t')$  by methods also described in that appendix.

Figures 7, 8, and 9 compare the theoretical values of  $\tau_r$ ,  $\nu$ , and  $s$  to those obtained from simulations with different numbers of agents. In all cases, the agreement between theory and simulation is best when there are more agents in the simulation.

For  $\beta \ll \beta_1$ , the system takes a very long time to reach

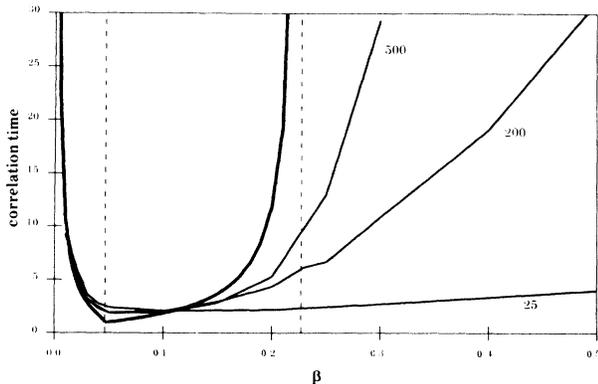


FIG. 7. Correlation time  $\tau_r$  vs  $\beta$  for the conditions of Fig. 5. The theoretical mean-field results (dark curve) are compared to results obtained from simulations of systems with 25, 200, and 500 agents as indicated in the figure. The oscillation and stability thresholds  $\beta_1=0.0464$  and  $\beta_2=0.2271$  are indicated by dashed vertical lines. For  $\beta < \beta_2$ ,  $\tau_r$  is the relaxation time, while for  $\beta > \beta_2$  it is the phase coherence time of oscillations about the fixed point.

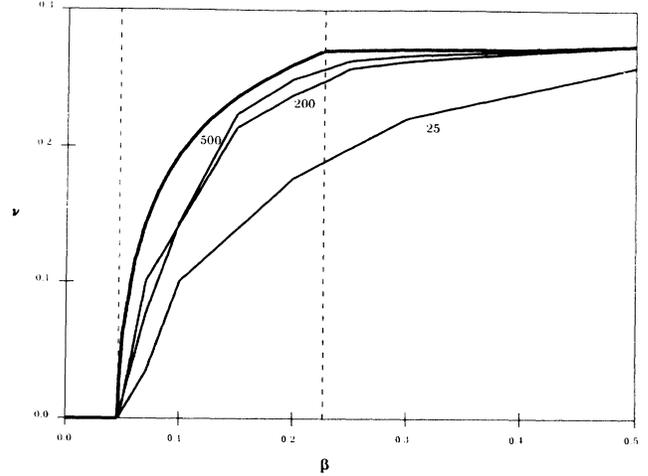


FIG. 8. Oscillation frequency  $\nu$  vs  $\beta$  for the situation described in Fig. 7. Discontinuities in the slope of the theoretical  $\nu(\beta)$  occur at the oscillation and stability thresholds  $\beta_1$  and  $\beta_2$ . For sufficiently large  $\beta$ , beyond the range plotted here,  $\nu$  approaches  $\frac{1}{2}$  asymptotically (e.g., at  $\beta=10$ ,  $\nu$  is 0.46).

equilibrium, and does so without oscillation. For  $\beta \approx \beta_1$ , relaxation to equilibrium is fastest, taking just a few delay times. This may be an advantageous regime in which to operate because of the rapid, yet stable, response of the system. Figure 8 confirms that the onset of decaying oscillations occurs at  $\beta = \beta_1$ . As  $\beta$  increases towards  $\beta_2$ ,  $\nu$  increases towards its asymptote of  $\frac{1}{2}$ , and the relaxation time  $\tau_r$  increases. The theoretical values of  $d\nu/d\beta$  and  $ds/d\beta$  are both discontinuous at  $\beta_2$ . Since the oscillation amplitude grows suddenly beyond  $\beta_2$ , the linearized analysis in the neighborhood of the fixed point fails to provide an accurate description of the global behavior of  $f(t)$  for  $\beta > \beta_2$ .

The mean-field theory predicts that  $\tau_r$  becomes infinite at  $\beta = \beta_2$  and remains infinite for  $\beta > \beta_2$ , i.e., the system

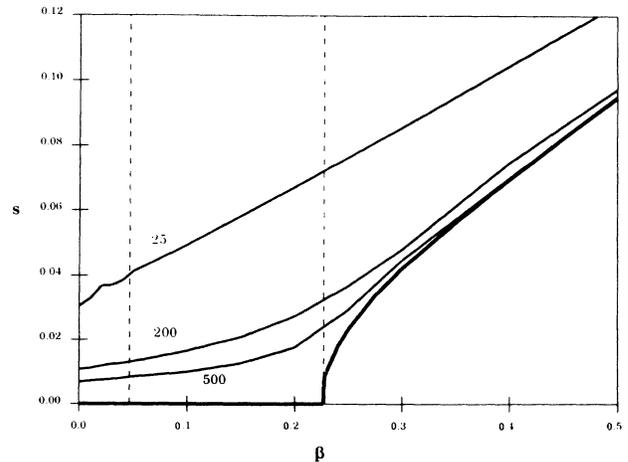


FIG. 9. rms oscillation amplitude  $s$  vs  $\beta$  for the situation described in Fig. 7. Discontinuities in the slope of the theoretical  $s(\beta)$  occur at the oscillation and stability thresholds  $\beta_1$  and  $\beta_2$ .

exhibits coherent oscillations that do not decay in time. In the simulations, however,  $\tau_r$  is always finite, being larger when there are more agents. The finite value of  $\tau_r$  in the simulations cannot be attributed to a decay in the oscillation amplitude. For example, the oscillation amplitude of the  $\beta=0.4$  simulation run in Fig. 5(a) does not appear to decay at all within  $50\tau$ , whereas according to Fig. 7,  $\tau_r$  is only  $19.1\tau$ . Figure 9 further confirms that the oscillation amplitude in the simulations does not decay when  $\beta > \beta_2$ . In this regime, the magnitude of the fluctuations is slightly greater than the rms amplitude of the coherent oscillations predicted by the mean-field theory, provided that the number of agents is sufficiently large. The remaining component of  $s$  can be attributed to random statistical fluctuations which scale approximately as  $1/\sqrt{A}$ . [Note that, according to Eq. (4), the magnitude of the random statistical fluctuations is *exactly* proportional to  $1/\sqrt{A}$  in the limit as  $\beta \rightarrow 0$ .] Therefore, we conclude that, in the simulations,  $\tau_r$  for  $\beta > \beta_2$  is to be interpreted as the phase coherence time for persistent oscillations.

Now we shall consider the effect of modifying the parameter  $\sigma$ , which reflects a number of different sources of uncertainty, e.g., in the information about resource allocation, the payoff function, and perhaps some additional randomness in the decision procedure, which is intentionally introduced by the system designer for reasons which we shall soon elucidate. Figure 10, which is essentially a transformed version of Fig. 6, illustrates how the onsets of damped and persistent oscillations depend upon  $\sigma$ . If  $\sigma$  is greater than the critical value  $\sigma_2$ ,  $\rho'(f_0) > -1$ , so that, according to Eq. (11) and Fig. 6,  $\beta_2$  becomes infinite. Thus, by intentionally increasing the randomness in the decision procedure used by each agent, persistent oscillations can be eliminated entirely. This is reasonable, since, in the limit as  $\sigma \rightarrow \infty$ , agents make completely random decisions independent of the payoffs, and the system set-

ties into the equilibrium  $f_0 = \frac{1}{2}$ . Figure 11 demonstrates that elimination of persistent oscillations can result in a dramatic improvement in the system performance (as defined by the sum of the payoffs of the individual agents). For fixed  $\beta$ , the best possible performance is obtained when  $\sigma$  is just enough to suppress persistent oscillations. The increase in performance can be particularly impressive for large reevaluation rates—from worse than that for completely random decisions to almost as good as that for perfect decisions (i.e., decisions made in the absence of uncertainty and delays). Figures 10 and 11 show that this system with an arbitrarily large reevaluation rate can always be made to operate at greater than 76% of the optimal performance, provided that the intrinsic uncertainty is less than  $\sigma_2$ .

### C. Cooperative agents

In Sec. III B it was assumed that the agents were completely independent of one another and, therefore, preferred resources solely on the basis of their capacity. However, in some cases an agent can benefit from results being generated by others. For example, agents searching through a common database could leave a record of discovered relations which are useful to others. In this case, if communication overhead costs between agents using different resources are significant, an agent might prefer to choose a resource that is simultaneously being used by others rather than one that is relatively unused. In this case, the payoff for using a resource could be maximal somewhere in the interval  $0 < f < 1$ .

Figure 12(a) displays payoff functions which are identical to those of Fig. 2(a) in the neighborhood of the stable fixed point, but with a quadratic maximum in  $G_1(f)$ . As displayed in Fig. 12(b),  $\rho(f)$  for  $\sigma=0.125$  has a stable fixed point at  $f_0=0.700$ . Thus, by design, the behavior of this system is essentially the same as that of the com-

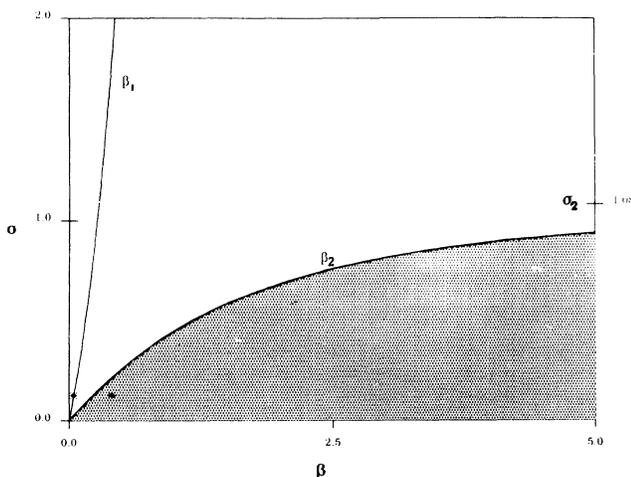


FIG. 10. Oscillation and stability thresholds  $\beta_1$  and  $\beta_2$  as a function of the uncertainty  $\sigma$ , given the payoffs of Fig. 2. As in Fig. 6, the shaded area represents the unstable region. The asterisks indicate the parameters used in Fig. 5.

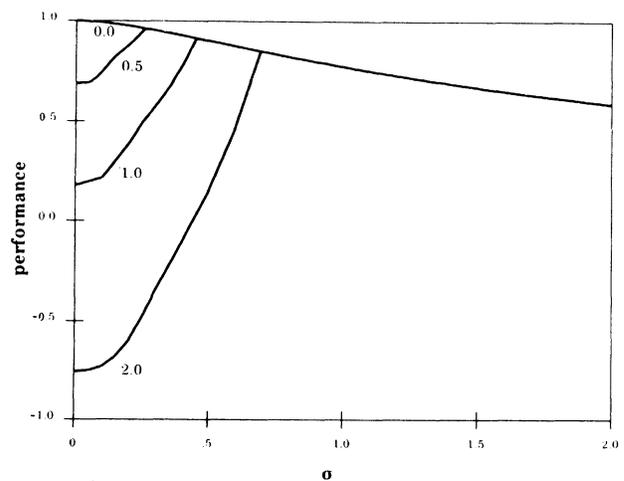


FIG. 11. Dependence of system performance upon the uncertainty for several different values of  $\beta$  as indicated in the figure. The performance is calculated by adding time-averaged payoffs for all agents and normalizing to 1 for perfect information ( $\sigma=0$  and  $\tau=0$ ), and 0 for completely random decisions ( $\sigma \rightarrow \infty$  and  $\tau=0$ ).

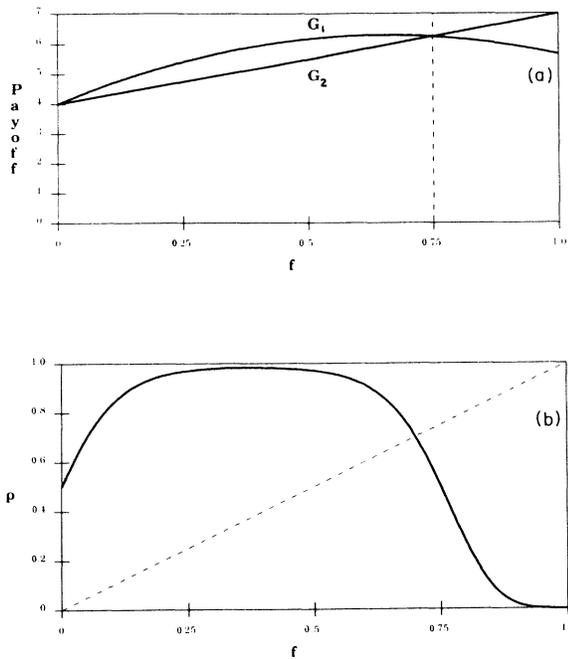


FIG. 12. (a) Cooperative payoffs:  $G_1 = 4 + 7f_1 - 5.333f_1^2$ ,  $G_2 = 7 - 3f_2$ . (b) Preference probability function  $\rho(f)$  for these payoffs for  $\sigma = 0.25$ . The equilibrium value is  $f_0 = 0.700$  as indicated by the intersection with the dashed line  $\rho = f$ .

petitive case of Fig. 2 for  $\beta \leq \beta_2$ .

However, the notch in  $\rho(f)$  in the vicinity of  $f = 0$  in Fig. 12(b) leads to new types of behavior in the regime  $\beta \gg \beta_2$ , such as period-doubling bifurcations and chaos. As displayed in Fig. 13, numerical solution of Eq. (7) reveals that the frequency of the persistent oscillations in  $f(t)$  increases initially towards  $\nu = \frac{1}{2}$ . However, at the critical value  $\beta_d = 2.93$ ,  $\nu$  is suddenly reduced by a factor of 2. A typical example of the behavior of  $f(t)$  in this regime is displayed in Fig. 14(a). As  $\beta$  is increased further,  $\nu$  resumes its upward climb, but it is soon interrupted

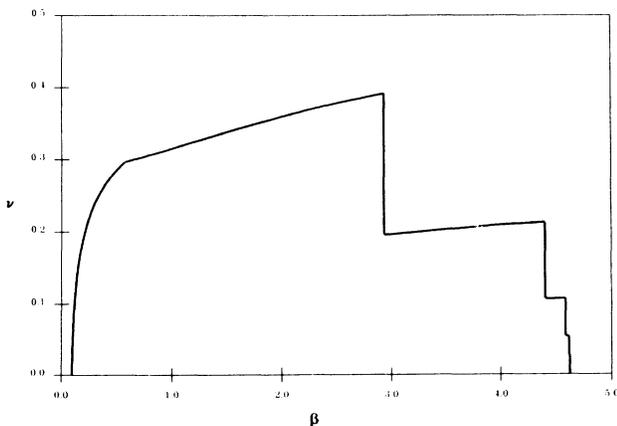


FIG. 13. Oscillation frequency  $\nu$  vs  $\beta$  for the parameters of Fig. 12(b). The sequence of bifurcations starts at  $\beta_d = 2.93$  and culminates in chaotic behavior at the critical value  $\beta_c = 4.63$ . Beyond this value lies a chaotic regime which is punctuated by numerous windows of periodic behavior.

again by another bifurcation. The infinite sequence of bifurcations accumulates at  $\beta_c = 4.63$ , at which point  $f(t)$  becomes chaotic. For  $\beta > \beta_c$ , there are numerous windows of  $\beta$  in which there are oscillations of period 5, 7, 8, 10, etc. For a fuller exposition of period doubling and chaos in solutions to differential-delay equations, along with a discussion of Lyapunov exponents, dimensions, and entropy, see Ref. 15.

Period-doubling bifurcations can also be observed in simulations of this system. As shown in Fig. 14(b), a typical simulation of 400 agents with the same parameters as in Fig. 14(a) displays a noticeable alternation in the successive maxima of  $f(t)$  which is qualitatively similar to that predicted by the theory. The excellent quantitative agreement between theory and simulation is demonstrated in Fig. 15, which compares the corresponding auto-correlations. The presence of a subharmonic at twice the fundamental period of oscillation is clearly visible. However, the rest of the period-doubling bifurcation sequence is usually not observed in the simulations; instead, there appears to be a direct transition from period-two oscillation to chaos. This loss of the fine period-doubling structure in the route to chaos is typically observed in noisy systems.<sup>16</sup> The existence of a bifurcation gap shows that fluctuation corrections act as an external random force

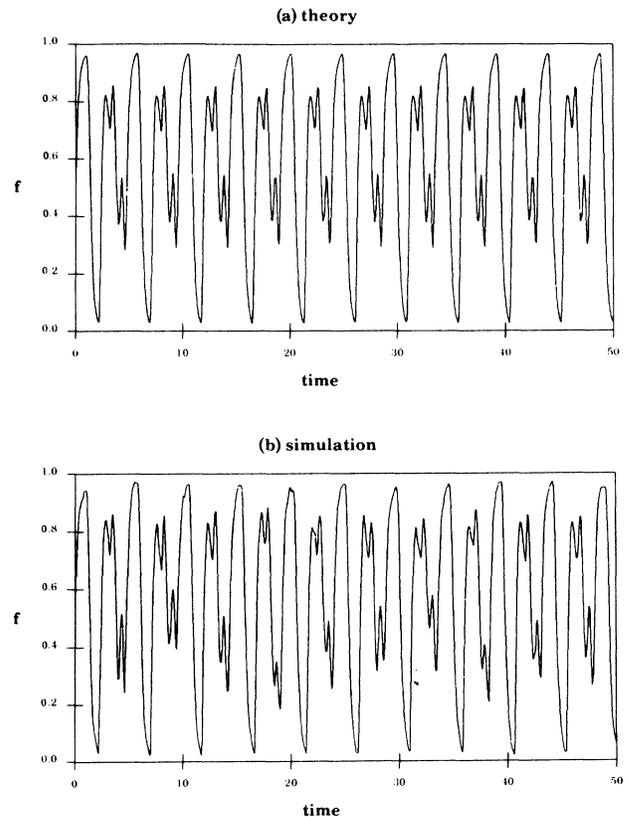


FIG. 14. (a) Theoretical and (b) simulated (400 agents)  $f(t)$  for a system of partially cooperative agents with the same parameters as in Fig. 12(b) and  $\beta = 4.0$  (period-doubling regime). The initial conditions are  $f(t) = 0.5$  for  $-\tau < t < 0$ . Time is measured in units of  $\tau$ .

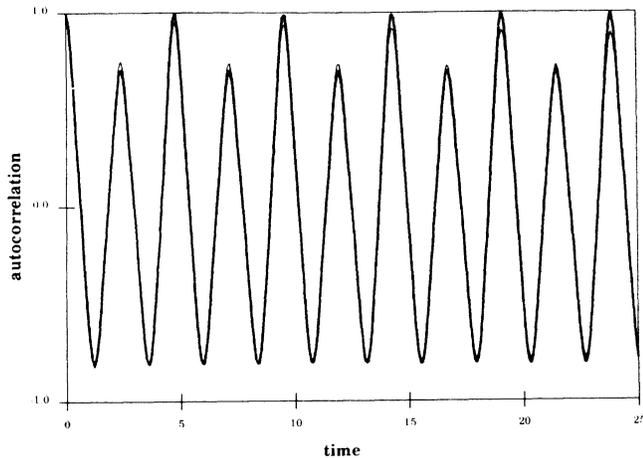


FIG. 15. Theoretical (dark curve) and simulated (light curve) autocorrelations calculated from the  $f(t)$  displayed in Fig. 14. The doubled period is clearly visible. The slight decay in amplitude of the autocorrelation obtained from simulation is due to phase incoherence arising from the finite number of agents.

on the mean-field dynamics which decreases in magnitude as the number of agents increases.

For  $\beta > \beta_c$ , simulations exhibit many of the characteristics of chaotic behavior. Figure 16(b), which displays a simulation of 400 agents using the same parameters as were used in Fig. 15, appears to be much more random and unpredictable. It bears a strong resemblance to its mean-field counterpart in Fig. 16(a). More quantitatively, the decreased predictability of the system in the chaotic regime is reflected by a dramatic drop in the correlation time  $\tau_r$  when  $\beta > \beta_c$ . As illustrated in Fig. 17,  $\tau_r$  is reduced to about  $40\tau$  in the mean-field case and  $15\tau$  for a simulation of 400 agents when  $\beta = 6.0$ , compared to infinity and  $110\tau$ , respectively, when  $\beta = 4.0$  (Fig. 15).

In Sec. III B we found that persistent oscillations could be eliminated by deliberately increasing the randomness of a decision procedure. Chaotic behavior can be eliminated by the same technique. Figure 18, an extension of Fig. 10(a) to cooperative agents, shows how the critical values  $\beta_1$ ,  $\beta_2$ ,  $\beta_d$ , and  $\beta_c$  depend upon  $\sigma$ . Interestingly, the boundaries for bifurcation and chaos double back on themselves. Consequently, as one increases the uncertainty from zero while holding the reevaluation rate fixed, it is possible to pass from a region of simple period-one oscillation through the full bifurcation sequence to the chaotic regime, then back in reverse order to simple oscillation, and, finally, stability. The critical values of  $\sigma_2$ ,  $\sigma_d$ , and  $\sigma_c$ , for which  $\beta_2$ ,  $\beta_d$ , and  $\beta_c$  become infinite, can be determined by noting that Eq. (7) is equivalent to a discrete map from the unit interval onto itself when  $\beta$  is set to infinity:

$$f(t) = \rho(f(t - \tau)). \quad (12)$$

Since Eq. (12) involves so much less computation than Eq. (7), it is useful for determining  $\sigma_2$ ,  $\sigma_d$ , etc., and for obtaining a good estimate of the form of  $f(t)$  when  $\beta$  is large. In many cases,  $f(t)$  converges in mean square to a series of plateaus of duration  $\tau$  corresponding to the discrete series given by Eq. (12).

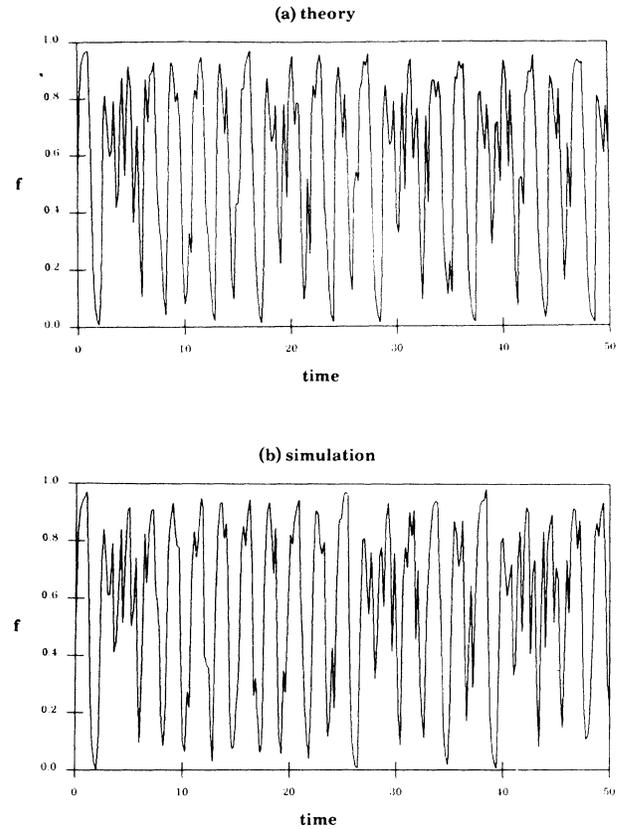


FIG. 16. (a) Theoretical and (b) simulated (400 agents)  $f(t)$  for a system of partially cooperative agents with the same parameters as in Fig. 12(b) and  $\beta = 6.0$  (chaotic regime). The initial conditions are  $f(t) = 0.5$  for  $-\tau < t < 0$ . Time is measured in units of  $\tau$ .

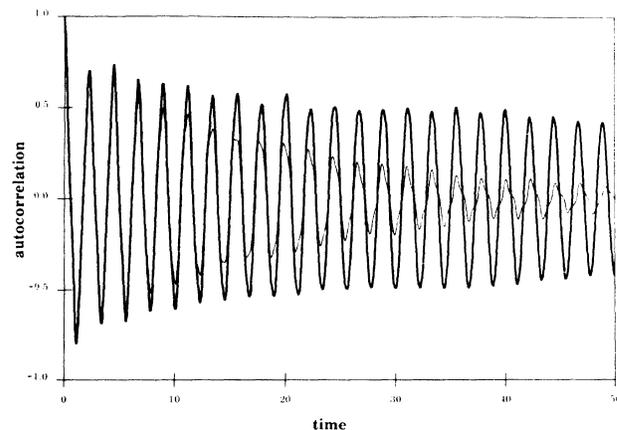


FIG. 17. Theoretical (dark curve) and simulated (light curve) autocorrelations calculated from the chaotic  $f(t)$  displayed in Fig. 16. The decreased predictability of the system's behavior in the chaotic regime is reflected by the relatively strong decay in oscillation amplitude of the autocorrelation function.

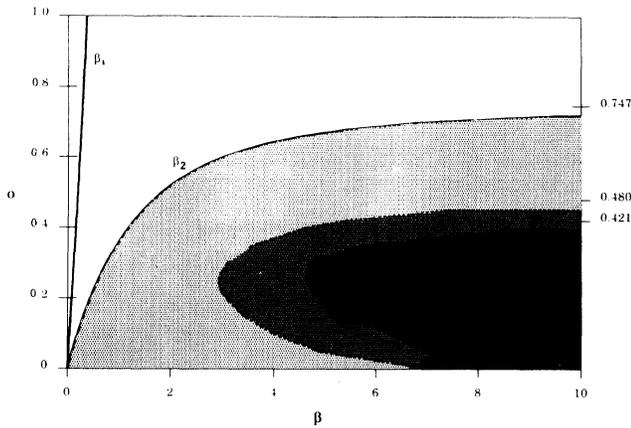


FIG. 18. Behavior phase diagram in terms of  $\beta$  and  $\sigma$  for payoffs of Fig. 12. As in Fig. 10, the unshaded area represents stable behavior. Light gray, period-1 oscillations; dark gray, bifurcations; black, chaotic regime. Windows of periodicity within the chaotic region are not shown. The values of  $\sigma$  at which instability, bifurcation, and chaos are eliminated for arbitrarily large  $\beta$  are indicated on the right-hand axis.

#### D. Multiple resources

In most systems of interest, agents will have many possible resources from which to choose. In order to investigate the consequences of having more than two resources in the system, we generalize Eq. (7) to include  $R$  resources by defining  $f_r(t)$  to be the fraction of agents using resource  $r$  at time  $t$ . Note that, since  $f_R = 1 - \sum_{r=1}^{R-1} f_r$ , it is sufficient to consider the  $(R-1)$ -dimensional vector  $\mathbf{f}$ . Equation (7) generalizes to

$$\frac{d\mathbf{f}(t)}{dt} = -\alpha[\mathbf{f}(t) - \rho(\mathbf{f}(t - \tau))], \quad (13)$$

where  $\rho_r$  is the probability that an agent will select resource  $r$  when it performs its evaluation. An expression for  $\rho_r$  in terms of the payoffs and the uncertainty is easily derived from the more general expression given in Appendix A. As in the scalar case considered previously, this differential-delay equation can give rise to oscillations and chaotic behavior.

As a simple example, we shall consider a system of purely competitive agents vying for three resources. The dynamics of the system is described in terms of the two components  $f_1(t)$  and  $f_2(t)$ . Figure 19 displays the decision regions for a particular set of payoffs which decrease monotonically with resource usage. When  $\sigma > 0$ , the boundaries between these regions become somewhat fuzzy, and the equilibrium point drifts away from the optimal position at  $(\frac{3}{23}, \frac{5}{23})$  towards  $(\frac{1}{3}, \frac{1}{3})$  (just as in Fig. 2 the fixed point drifts from  $f_0 = 0.75$  towards  $f_0 = 0.50$  as  $\sigma$  is increased).

As  $\beta$  is increased while  $\sigma$  is held fixed, the behavior of the system passes through the same stages as were ob-

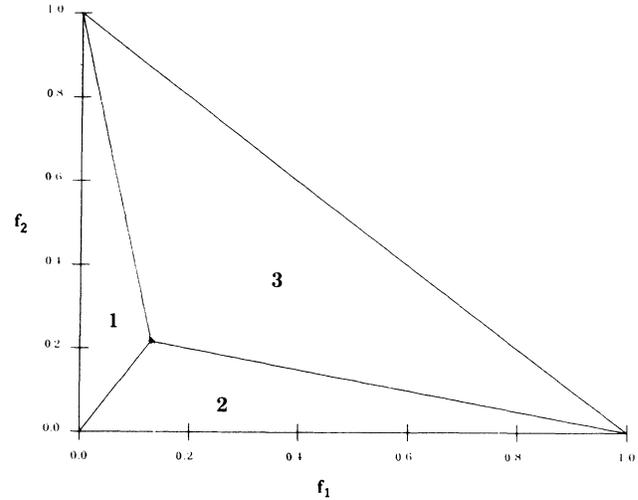


FIG. 19. Decision regions for a system of purely competitive agents vying for three resources. The payoffs used by each agent are  $G_1 = 12 - 10f_1$ ,  $G_2 = 12 - 6f_2$ , and  $G_3 = 12 - 2f_3$ . The decision regions are labeled according to the resource yielding the highest payoff when the state of the system is within that region, assuming that  $\sigma = 0$ . As in Fig. 2, the boundaries between decision regions become fuzzy when  $\sigma > 0$ . The asterisk marks the equilibrium point to which the system will settle in the absence of delays.

served for the two-resource case in Sec. III B, with a frequency that decreases towards the asymptotic value of  $1/(2\tau)$ . The oscillation and instability thresholds  $\beta_1$  and  $\beta_2$  can be determined from Eqs. (10) and (11). For even larger values of  $\beta$ , the three-resource system can exhibit the period-doubling and chaotic behavior that were described in Sec. III C. It is interesting that these types of behavior are observed only for partially cooperative agents in a two-resource system, whereas in a three-resource system they occur even when the agents are purely competitive.

The double-loop limit cycle in the phase-plane portrait of Fig. 20(a) illustrates period-doubling behavior in the system described in Fig. 19. Figure 20(a) is generated by solving Eq. (13) with  $\beta = 0.25$  and  $\sigma = 0.125$ , starting from the initial condition  $(f_1, f_2) = (\frac{1}{3}, \frac{1}{3})$  for  $-\tau < t < 0$ , letting the transient behavior die out, and then plotting  $f_2(t)$  versus  $f_1(t)$ . Provided that there are enough agents, period doubling can be observed in simulations as well. Figure 20(b) is a noisy simulacrum of Fig. 20(a), which was obtained by running a simulation with 300 agents for the same parameters. The general shape and double-loop behavior of the attractor in Fig. 20(a) is still perceptible.

When  $\beta$  is increased further to 0.50, the phase-plane portrait is a chaotic attractor, as illustrated in Fig. 21. The statistical fluctuations in a simulation would almost certainly obscure the chaotic nature of the attractor unless the number of agents was drastically increased or the parameters were changed. As was the case for the partially cooperative agents of Sec. III C, the chaotic regime is punctuated by windows of periodic behavior, a fact which will be made use of in Sec. III G.

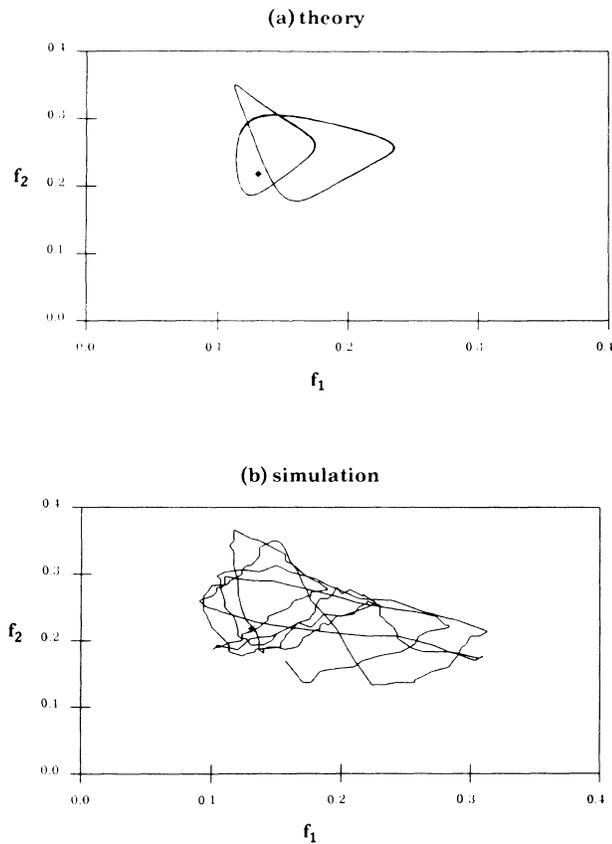


FIG. 20. (a) Theoretical and (b) simulated (300 agents) phase-plane portraits of the system described in Fig. 19 for  $25\tau < t < 50\tau$ , with  $\beta=0.25$  and  $\sigma=0.125$ . The initial conditions are  $(f_1, f_2) = (\frac{1}{3}, \frac{1}{3})$  for  $-\tau < t < 0$ . The asterisk marks the optimal equilibrium points to which the system would converge in the absence of delays.

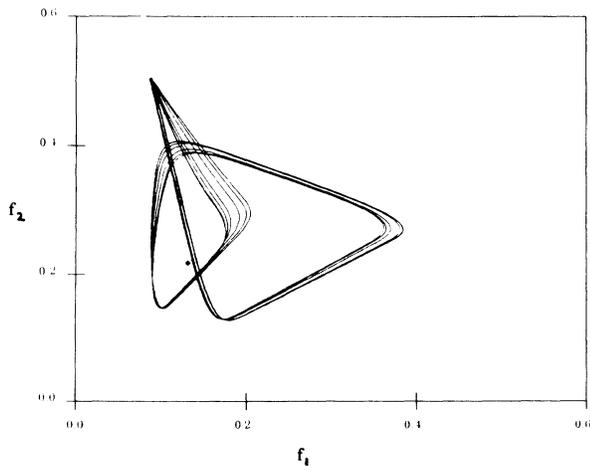


FIG. 21. Chaotic attractor observed in a system with parameters identical to those of Fig. 20, except that  $\beta=0.50$ . The attractor is plotted for  $100\tau < t < 150\tau$ , which ensures that all initial transients have died out.

### E. Inhomogeneous agents

Computational ecosystems are likely to contain agents engaged in a large variety of tasks with differing computational needs. This will lead to corresponding differences in how various resources are valued. For example, other factors being equal, an agent performing a numerical simulation will assign more value to a machine with floating-point hardware than would one searching through a large database. Other factors which might differ among the various agents include the sophistication of the payoff evaluation strategies, the relative preference for speed versus accuracy of results, and the time delay.

Such differences in preferences or other characteristics among the agents can be included in our model by supposing that there are  $S$  different species of agents, each with a different payoff function  $G_{rs}$ , which describes how much agents of species  $s$  value resource  $r$ . The generalization of the mean-field equation to inhomogeneous agents competing for multiple resources, given in Appendix A, has the same form as Eq. (13), except that now  $\mathbf{f}$  and  $\boldsymbol{\rho}$  are to be interpreted as doubly indexed vectors.

We will now consider two simple examples of such systems. In both cases, there are two resources and two species which are purely competitive, i.e., the payoffs decrease monotonically with resource usage.

First, we consider a system in which the two species have different preferences as embodied in their payoff functions, and all other parameters are identical. For both species, the payoffs are the same as in Fig. 2, except that the slope of  $G_2$  is modified such that species 1 and species 2 have greater and lesser preferences for resource 2, respectively.

With these payoffs and an uncertainty of  $\sigma=0.125$ , homogeneous systems consisting entirely of either species 1 or species 2 possess instability thresholds of  $\beta_2=0.277$  and  $\beta_2=0.200$ , respectively. However, as illustrated in Fig. 22, a system consisting of an even mixture of the two species is stable even when  $\beta=0.5$ . In the inhomogeneous system, species 1 tends towards an equilibrium in which the fraction of such agents using resource 1 is  $f_0=0.372$ , whereas species 2 tends towards  $f_0=0.999$ . Thus each species has its own niche, species 2 preferring resource 2 and *vice versa*. Interestingly, in the corresponding homogeneous systems, the fixed points are  $f_0=0.645$  for species 1 and  $f_0=0.774$  for species 2, so that both species prefer resource 1.

In order to provide a better understanding of the increased stability of inhomogeneous systems, the dependence of the instability threshold  $\beta_2$  upon the fraction  $g$  of agents of species 1 is shown in Fig. 23. If a homogeneous system of one species is contaminated by the introduction of some agents of the other species, the instability threshold increases, eventually reaching a maximum of  $\beta_2=0.92$  at  $g_0=0.35$ . If the uncertainty is increased to  $\sigma=0.25$ , the maximum shifts to an entirely different position. The oscillation threshold  $\beta_1$  mimics the behavior of  $\beta_2$ , peaking at the same percentage  $g$  as  $\beta_2$  for both values of  $\sigma$ . In no case does the system exhibit period-doubling or chaotic behavior. As demonstrated in Fig. 24, the system performance peaks near its maximum pos-

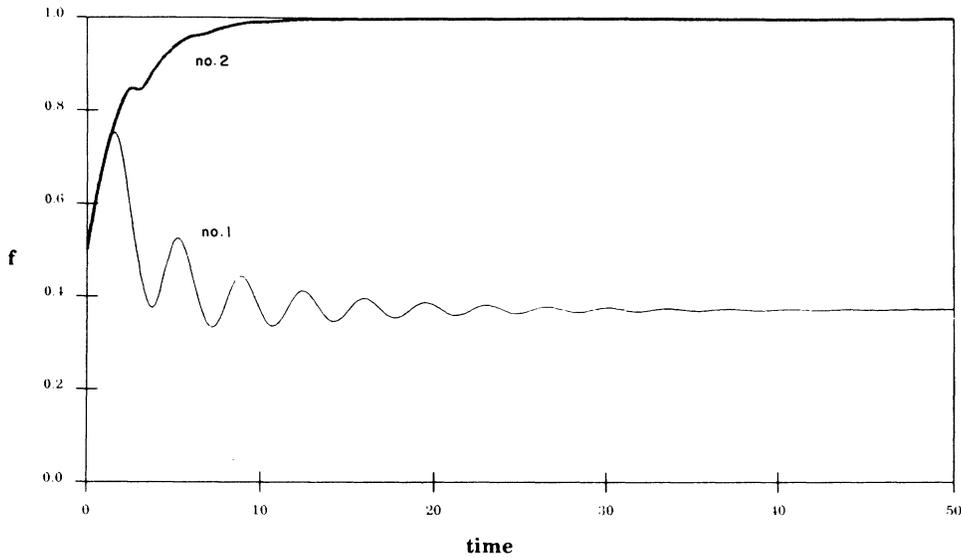


FIG. 22. Fraction of agents of species 1 (light curve) and species 2 (dark curve) using resource 1 as a function of time in an inhomogeneous system with two resources and two species of agents. Both species use  $\tau=1.0$  and  $\sigma=0.125$ , and the same payoff for using resource 1 as in Fig. 2:  $G_1=7-f_1$ . For species 1,  $G_2=7-2f_2$ , whereas for species 2,  $G_2=7-4f_2$ .  $\beta=0.5$ .

sible value in the vicinity of  $g_0$ , where the system is stable.

In Sec. III B we found that increasing the uncertainty in a homogeneous system always led to increased stability (albeit at the possible expense of decreased performance). However, Fig. 23 shows that, for certain inhomogeneous mixtures, increasing the uncertainty can actually de-

crease the stability of the system.

In the second example, all of the agents use the same payoffs as given in Fig. 2, but species 1 and 2 have unequal delay times  $\tau_1$  and  $\tau_2$ . Simplifying Eq. (A6), we find that the exponential on the right-hand side of Eq. (9) is replaced by an average of exponentials of the same form weighted by the relative proportion of agents of each species:

$$\xi = g\gamma e^{-\alpha\tau_1\xi} + (1-g)\gamma e^{-\alpha\tau_2\xi} - 1. \quad (14)$$

This tends to make the system behave much as though it had an effective delay given by the weighted average of  $\tau_1$  and  $\tau_2$ . As we have verified by numerical solution of Eq. (A1), the coupling between the two species results in a well-defined oscillation frequency which is midway be-

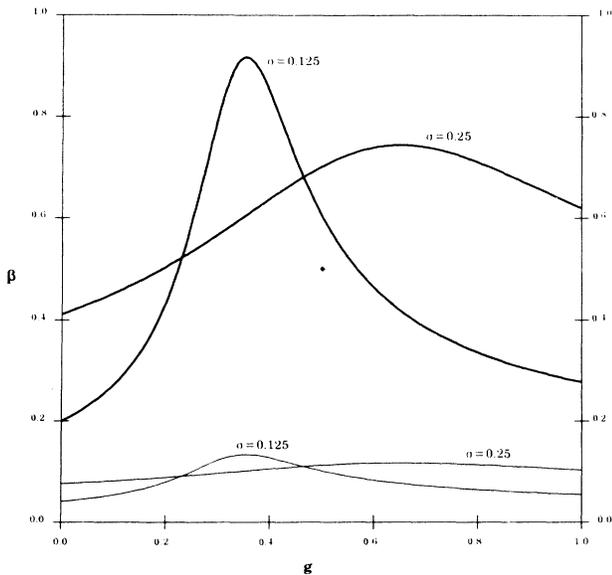


FIG. 23. Oscillation and stability thresholds  $\beta_1$  (light curve) and  $\beta_2$  (dark curve) as a function of the fraction of agents  $g$  of species 1 in the inhomogeneous system described in Fig. 22. The curves, which are calculated by solving the mean-field Eq. (A6) of Appendix A, are displayed for two different values of the uncertainty:  $\sigma=0.125$  and  $\sigma=0.25$ . The parameters of Fig. 22 are indicated by an asterisk.

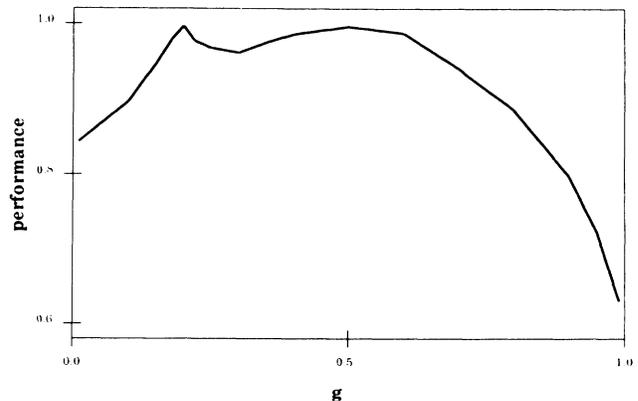


FIG. 24. Performance vs  $g$  for the inhomogeneous system of Fig. 22. For each value of  $g$ , performance is normalized to zero for purely random decisions and unity for perfect decisions (those which would be made in the absence of delays or uncertainty).

tween those of the corresponding homogeneous systems. As illustrated in Fig. 25, the fractions of agents of species 1 and 2 using resource 1 oscillate with the same frequency. Even though  $\tau_1$  and  $\tau_2$  differ by 40%, the phase shift between the two species in Fig. 25 is less than 0.6 rad—small enough so that the overall oscillation amplitude of the system is still quite large. This suggests that the presence of multiple delays in a system does not greatly reduce the susceptibility of that system to the large oscillations described in this and preceding sections.

### F. Smart agents

In preceding sections we have shown that the nonlinear dynamics of interacting agents can lead to oscillations and chaos, which tend to reduce overall system performance. We have also seen that performance could be increased if the agents deliberately used a probabilistic decision procedure. In this case, an agent sometimes chooses a resource other than the one which appears to be best in light of the information available to it at the time. An alternative approach is to provide agents with more sophisticated, deterministic methods for deciding among resources.

As an example, we reconsider the system discussed in Sec. III B, in which purely competitive agents vie for two resources. The normal agents use the most recently available information in their decisions. When this information is uncertain and delayed, these agents make suboptimal decisions, leading to global oscillations with a well-defined period and amplitude. By contrast, we now introduce “smart” agents which try to take advantage of these simple oscillations to better estimate current resource utilization. Specifically, a smart agent continually monitors the oscillation period, estimating it as the time between successive maxima in  $f(t)$ . Knowing this period, the agent can estimate current utilization by its value at some number of periods in the past. Thus it will (presumably) have more accurate information upon which to base its decision.

To assess the efficacy of this technique, we consider an

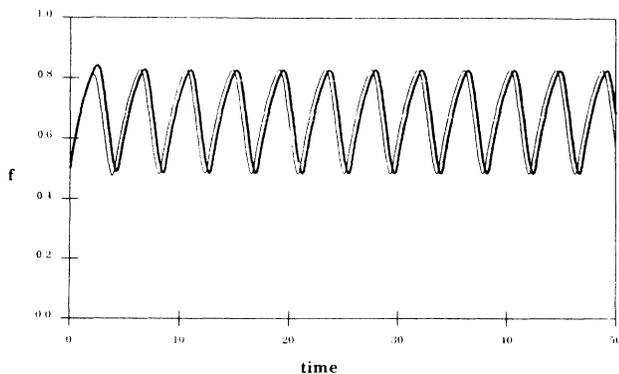


FIG. 25. Fraction of agents of species 1 (light curve) and species 2 (dark curve) using resource 1 as a function of time in an inhomogeneous system with two resources and two species of agents. Both species use the payoffs of Fig. 2 and  $\sigma=0.125$ , and their respective time delays are  $\tau_1=1.0$  and  $\tau_2=1.4$ .  $\beta=0.5$ .

inhomogeneous system consisting of smart and normal agents. As shown in Fig. 26(a), a system containing 10% smart agents and 90% normal ones becomes somewhat more stable after the smart agents determine the oscillation period (which takes a few periods). The amplitude of the oscillations is reduced by 24%, and the period is decreased by about 6%. Encouraged by this success, one might be tempted to make all of the agents smart. However, as shown in Fig. 6(b), this leads to very large and complex oscillations, and the overall system performance is degraded substantially below that of a system consisting of all normal agents. For other mixtures of smart and normal agents, the behavior of the system can even be chaotic. There does not appear to be any regularity in the dependence of the qualitative behavior upon the percentage of smart agents.

Figure 27 reveals that, when only 10% of agents are smart, they operate at a performance level (defined in Sec. III B) of approximately 100%, while the performance level of the normal agents doubles to nearly 50%—a sort of trickle-down effect. However, when the proportion of smart agents exceeds 80%, their performance can be substantially lower than that of the normal agents. For this particular system, the best overall system performance is obtained when approximately 20–50% of the agents are smart.

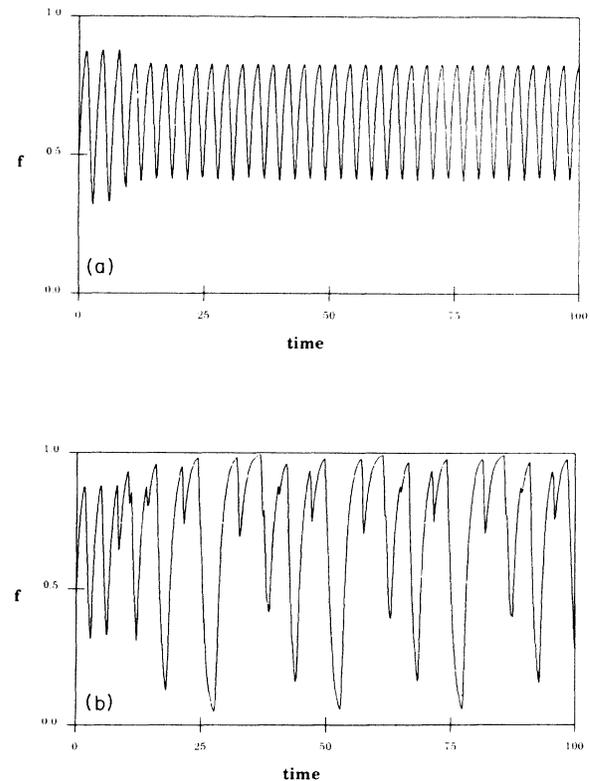


FIG. 26. Fraction of agents using resource 1 as a function of time in a two-resource, purely competitive system with smart agents. All system parameters are the same as in Fig. 5, except that  $\beta=1.0$ , and the initial conditions are  $f(t)=0.5$  for  $-\tau < t < 0$ . (a) 10% of the agents are “smart.” (b) 100% of the agents are “smart.”

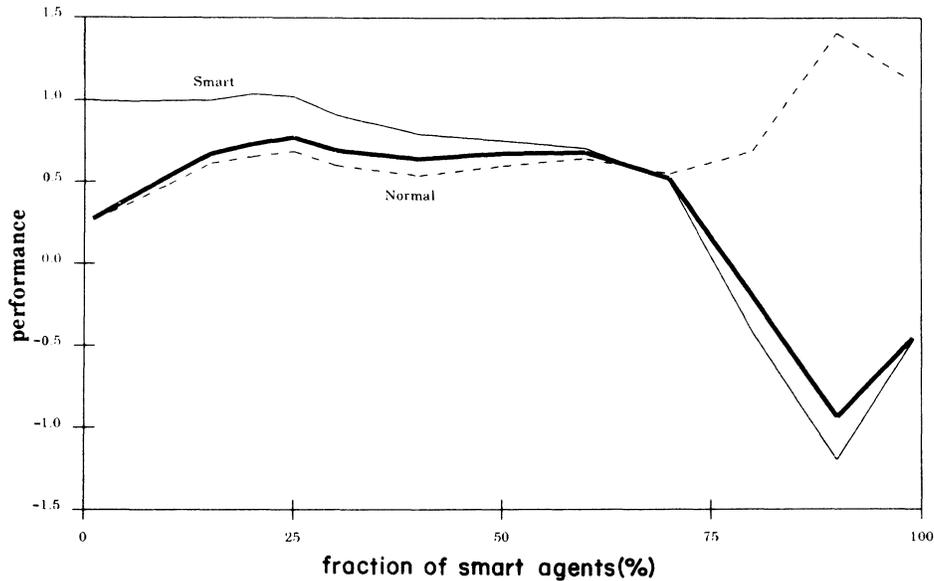


FIG. 27. Performance of a system with the same parameters as in Fig. 26 as a function of the percentage of smart agents. The dark curve is overall system performance. The performances of the smart- and normal-agent populations are given by the light solid curve and the dashed curve, respectively.

The problem with the smart agents is that their technique for detecting periodicity is based on the simple oscillations that occur when all of the other agents in the system are normal. This technique is adequate when only 10% of the agents are smart, but, when too many smart agents are present, the increased complexity of the dynamics causes this technique to give horribly inaccurate estimates of the oscillation period.

One might try to ameliorate this problem by making the agents even smarter, either by giving them more sophisticated ways to extrapolate from the available historical data or by providing them with knowledge of the intelligence level of the other agents in the system. This raises some important issues and potential problems. First, computational cost places a bound on the sophistication of any decision-making algorithm. Second, as we have just seen, an increment in intelligence may, if adopted by a sufficiently large fraction of the agents, alter the behavior of the system in such a way as to render the extra sophistication useless or even harmful. Finally, if an agent has the ability to predict the system's future behavior based upon knowledge of the decision-making algorithms of the other agents, the fact that *other* agents must predict *its* decisions could lead to an undecidable infinite regression or other paradoxes of rationality.<sup>17</sup> Even if an agent's choice of resource is decidable, the solution may be suboptimal in much the same way as in the prisoner's dilemma<sup>17,18</sup> or the tragedy of the commons.<sup>19</sup> An amalgam of analysis, simulation, and game theory might yield a great deal of insight into these issues.

### G. Punctuated equilibria and metastability

As a computational ecosystem evolves, not only will the agents adjust their choices of resources, but changes

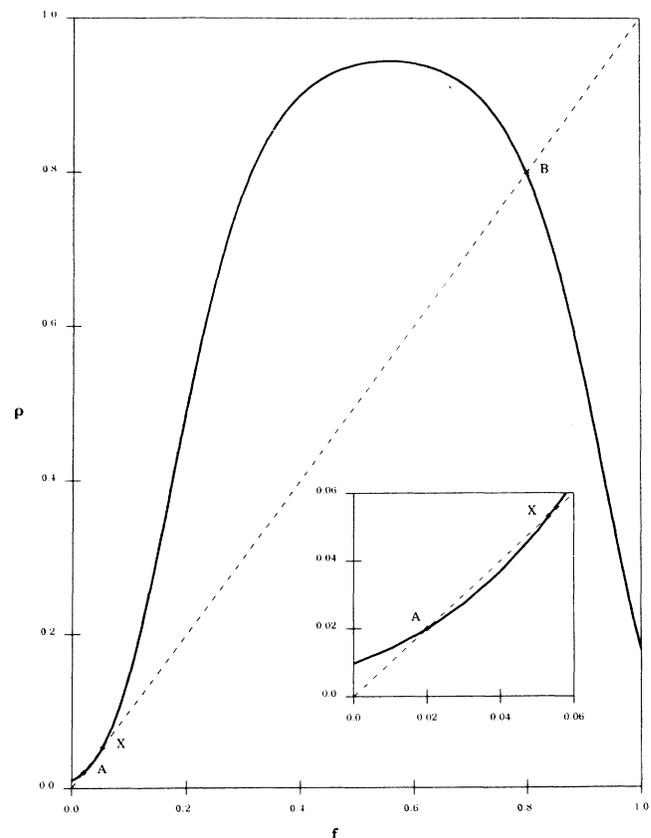


FIG. 28.  $\rho(f)$  for system of partially cooperative agents vying for two resources, with payoffs  $G_1 = 3.35 + 12f_1 - 10f_1^2$ ,  $G_2 = 6 - f_2^2$ , and uncertainty  $\sigma = 0.5$ . The inset shows the values for small  $f$ . The fixed points are indicated at the intersection of the curve with the dashed line  $\rho = f$ : stable equilibria  $f_A = 0.020$  (point A) and  $f_B = 0.799$  (point B); unstable fixed point  $f^* = 0.053$  (point X).

will be imposed by the external environment. For instance, new machines and databases could be added to the network, changing the relative payoffs of the resources. How readily can the systems adapt to such changes?

Agents using simple competitive payoffs can determine the benefit of any new resource individually and move to utilize it accordingly. However, if the agents are cooperative (e.g., each agent is unwilling to use a new resource unless many others are using it) or the transition overhead is high, the transition to the new state can be hampered by an "optimality barrier."<sup>12</sup> If the system possesses more than one attractor, an environmental change may alter the payoffs, changing the relative depths of the corresponding minima in the optimality surface so as to produce a new optimal state. However, the system might stay in the vicinity of its original state if it is unable to surmount its local barrier.

Our simulations reveal that statistical fluctuations can induce the system to jump from one attractor to another, although the time between jumps grows rapidly with the number of agents. To illustrate this behavior, we present

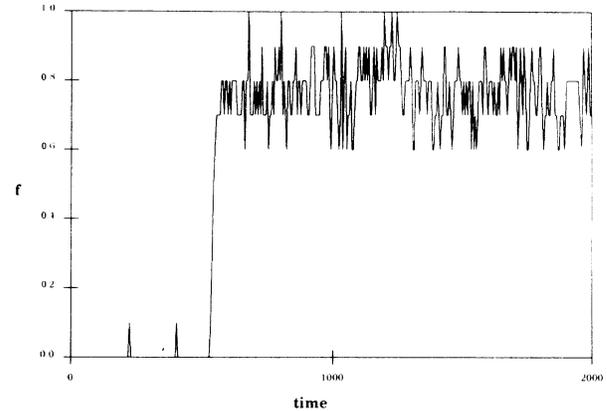


FIG. 29. Punctuated equilibrium in a simulation of ten agents with parameters as given in Fig. 28 and  $\beta=0.1$ . After  $530\tau$ , a random statistical fluctuation induces the system to hop over the energy barrier at  $f^*$ , thereby escaping the metastable equilibrium  $f_A$  and settling into the globally optimal equilibrium  $f_B$ . The width of the transition region is less than  $20\tau$ .

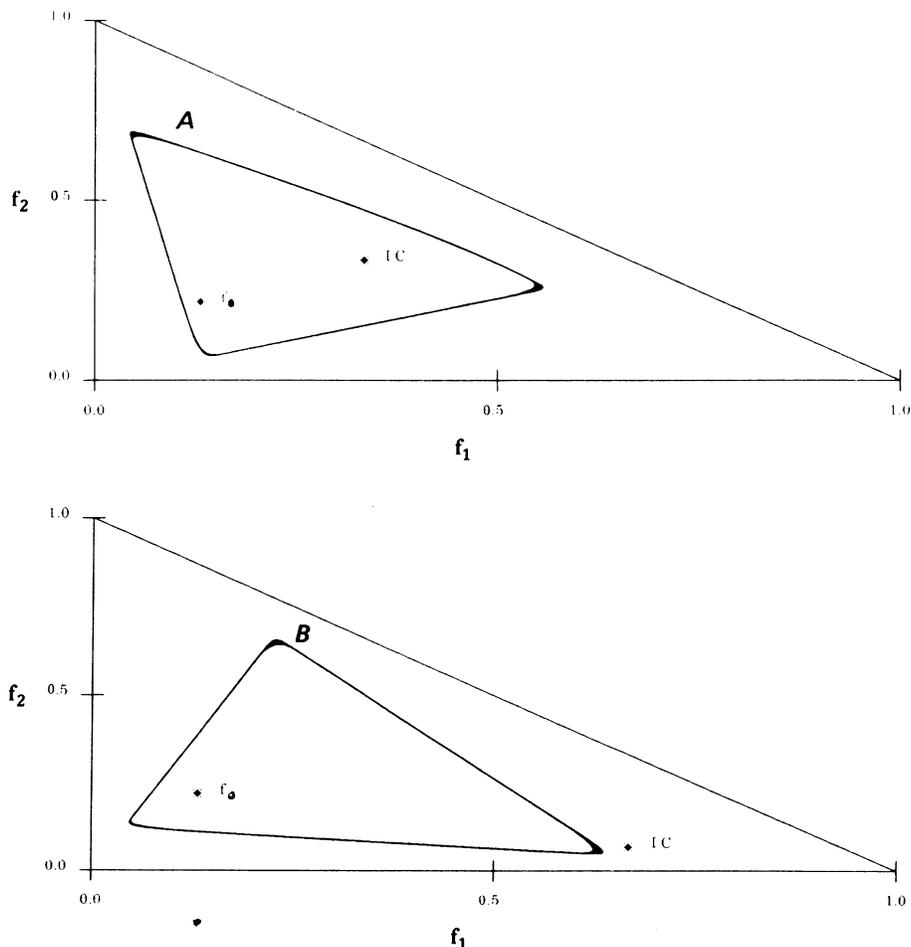


FIG. 30. Two different mean-field limit cycles for the competitive, three-resource system of Fig. 19 with  $\beta=1.0$  and  $\sigma=0.1$ . Initial conditions completely determine whether the system ends up in limit cycle  $A$ , which is followed in a clockwise direction as  $t$  increases; or  $B$ , which is followed in a counterclockwise direction as  $t$  increases. Typical initial conditions which lead to each limit cycle are indicated in the figure [assuming  $f(t)$  is constant during the interval  $-\tau < t < 0$ ].

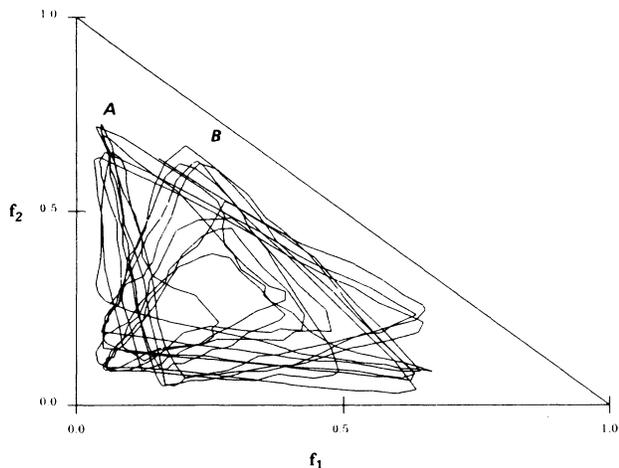


FIG. 31. Punctuated limit cycles in a simulation of 300 agents using the parameters of Fig. 30. The system is quickly captured into a noisy version of limit cycle  $A$ . Near  $t = 172\tau$ , there is a sudden transition taking less than  $2\tau$  to a noisy version of limit cycle  $B$ . In this simulation run, which lasted for  $1000\tau$ , the system flipped its state three times.

simulations of two different systems with multiple final states. In the first example, we increase the depth of the cooperative notch in  $\rho(f)$  of Fig. 12 to obtain two stable equilibrium points at  $f_A$  and  $f_B$ , and an unstable one at  $f_X$ , as illustrated in Fig. 28. In this case, the theory predicts that, for  $\beta < \beta_2$ , all initial conditions  $f_0 < f_X$  will flow into the fixed point  $f_A$ , while all initial conditions  $f_0 > f_X$  will flow into  $f_B$ . However, if there is a finite number of agents, it is possible for statistical fluctuations to jog a system that is stuck in  $f_A$ , into  $f_B$  or *vice versa*. Figure 29 shows a typical simulation run with ten agents, in which the system is stuck in the locally optimal metastable state  $f_A$  for some time before it makes a very sudden transition to the globally optimal state  $f_B$ . According to the theory developed by Ceccato and Huberman,<sup>12</sup> the amount of time spent in the metastable state is exponential in the number of agents, while the transition time from one state to another is logarithmic in the number of agents.

This type of metastable behavior is also observed in the three-resource system of Sec. III D. With the same payoffs as in Fig. 19,  $\beta = 1.0$  and  $\sigma = 0.1$ , there are two different single-loop limit cycles to which the system can be attracted, depending upon the initial conditions. These limit cycles are illustrated in Fig. 30. In a simulation with 300 agents, the system switches spontaneously between the two limit cycles, as illustrated in Fig. 31. The metastability of nonoptimal choices and the suddenness of the transition between attractors are reminiscent of evolutionary models of “punctuated equilibria.”<sup>20,21</sup>

Since the adaptability of a system in such a metastable situation is generally poor, additional mechanisms are needed to enhance the transition rate. In the case of cooperative agents and/or multiple resources, the adaptability of the system might be increased dramatically by giving the agents the capability to negotiate with one

another<sup>22</sup> or anticipate one another’s choices. Similarly, if transition overhead costs are preventing the system from reaching a more optimal state, it might be more appropriate for the agents to forecast expected payoffs over some future time interval rather than just trying to improve their instantaneous payoffs.

#### IV. CONCLUSION

Using both theoretical analysis and simulations, we have studied in detail the dynamics of a model which captures the essential features of computational ecosystems. We have illustrated and quantified the dependence of the model’s behavior upon delays and uncertainty in information, the degree of cooperation and competition among the agents, and the degree of inhomogeneity in the system. In addition, we have determined the quantitative conditions under which a system of cooperating agents can function acceptably well in the absence of central control. In general, the predictions of the theory of Huberman and Hogg<sup>11</sup> are validated by computer simulations, particularly in the limit of large numbers of agents. These experiments reveal that, in most cases, statistical fluctuations (neglected in the theory) do not shift the boundaries between various behavioral regimes; they merely blur them. One notable exception is the phenomenon of punctuated equilibria, which depends critically on the presence of fluctuations, and can only be explained by theories which explicitly incorporate them.<sup>12,20</sup> The basic lesson that emerges from this study is that the same complex global behavior observed in many physical systems can arise in computational ecosystems from simple interactions among programs. Given that slight complications introduced into the model tend to yield new behavioral phenomena, it is likely that real systems would exhibit behavior at least as exotic as that illustrated here.

A number of system design principles can be deduced from our work. For example, as illustrated in Fig. 7, the system’s responsiveness will be optimal if the decision rate is adjusted so that the system is critically damped (i.e.,  $\beta = \beta_1$ ). We have investigated three different heuristics for counteracting the strongly deleterious effect that delays have upon system performance. In Sec. III B, we found that a judicious increase in the randomness of the decision-making process of the individual agents can greatly ameliorate the problem of oscillations (Fig. 11). In Sec. III E we noted that an inhomogeneous system consisting of agents with different computational needs can be considerably more stable than a homogeneous system because the agents automatically create their own niches. As a result, the system performance can be optimized by purposely adjusting the heterogeneity of the system (Fig. 24). Finally, in Sec. III F we investigated the effect of using deterministic control algorithms which are more intelligent from the standpoint of an individual agent, and found that over-proliferation of that algorithm in the population of agents can actually decrease the overall system performance (Fig. 27).

An important attribute of computational ecosystems which is readily suggested by the analogy to biological

ecosystems is their ability to adapt to changing environments. The brief consideration of this issue in Sec. III G exposed a potentially serious problem: the lifetime of suboptimal metastable states is extremely long. Possible mechanisms for enhancing the transition rate (and hence the adaptability), which include negotiation among agents and forecasting of future payoffs, offer an interesting area for future research. It is also necessary to study, not just the evolution of the system for *fixed* payoffs, but the evolution of the individual agents on a slower time scale during which the agents modify their decision procedures so as to maximize their individual efficiencies. We must seek feedback mechanisms which can give computational ecosystems the same evolutionary capability as their biological counterparts. Such a study should also provide insight into the tradeoffs between a system's ability to adapt to a wide range of possible environments and its performance in a particular environment.

As has been suggested by several authors,<sup>6,7</sup> the most successful architectures might be founded upon economic rather than biological principles. The passive resources of our model could be replaced by ones which participate in the control of the system by auctioning their wares to the agents, resulting in a dynamical alteration of the payoffs. This may be particularly important in systems in which the distinction between servers and clients becomes blurred or nonexistent. For example, a resource trying to satisfy one agent's computational demand might act as a client in consulting with another agent that is knowledgeable about the relative merits of different algorithms,<sup>6</sup> or an agent that has derived some interesting result may serve as a resource for other agents. Smart

agents would be interesting to study in an economic context, particularly if they are endowed with some ability to anticipate the actions of their fellow agents. It would be of great interest and value to extend our model to encompass distributed-system architectures modeled after economic markets in order to determine whether they might possess some inherent fundamental advantages over the model we have considered in this work.

#### ACKNOWLEDGMENTS

This work was supported in part by the U.S. Office of Naval Research, Contract No. N00014-82-0699.

#### APPENDIX A: INHOMOGENEOUS AGENTS AND MULTIPLE RESOURCES

For the general case of inhomogeneous agents competing for multiple resources, the delayed-time mean-field Eq. (7) generalizes to

$$\frac{d\mathbf{f}(t)}{dt} = -\alpha[\mathbf{f}(t) - \rho(\mathbf{f}^*(t))], \quad (\text{A1})$$

where  $\mathbf{f}$ ,  $\mathbf{f}^*$ , and  $\rho$  denote the doubly-indexed vectors  $f_{rs}$ ,  $f_{rs}^*$ , and  $\rho_{rs}$ , respectively. The index  $r$  runs from 1 to  $R-1$ , and  $s$  runs from 1 to  $S$ , where  $R$  and  $S$  are the total number of resources and species, respectively. The components of  $\mathbf{f}^*$  are given by

$$f_{rs}^*(t) = f_{rs}(t - \tau_{rs}), \quad (\text{A2})$$

where  $\tau_{rs}$  is the time delay between resource  $r$  and an agent of species  $s$ . The preference probabilities  $\rho_{rs}(\mathbf{f}^*)$  are given in terms of the species-specific payoffs  $G_{rs}(\mathbf{f}^*)$  by

$$\rho_{rs}(\mathbf{f}^*) = \frac{1}{\sqrt{2\pi}\sigma_s} \int_{-\infty}^{\infty} dx e^{-x^2/2\sigma_s^2} \prod_{r'(\neq r)} \frac{1}{2} \left[ 1 + \operatorname{erf} \left[ \frac{x + G_{rs}(\mathbf{f}^*) - G_{r's}(\mathbf{f}^*)}{\sqrt{2}\sigma_s} \right] \right], \quad (\text{A3})$$

where  $\sigma_s$  is the uncertainty for agents of species  $s$ .

In the neighborhood of the fixed point given by  $\rho(\mathbf{f}_0) = \mathbf{f}_0$ , the variation  $\delta(t) = \mathbf{f}(t) - \mathbf{f}_0$  obeys

$$\frac{d\delta(t)}{dt} = -\alpha[\delta(t) - \underline{\gamma}\delta^*(t)], \quad (\text{A4})$$

where  $\underline{\gamma}$  is the Jacobian

$$\gamma_{rsr's'} \equiv \frac{\partial \rho_{rs}}{\partial f_{r's'}}(\mathbf{f}_0) \quad (\text{A5})$$

and the star notation is defined as in Eq. (A2). The substitution  $\delta(t) = \mathbf{c}e^{\alpha \xi t}$  in Eq. (A4) reduces it to a system of self-consistent eigenvalue equations of order  $(R-1)S$ :

$$[\underline{M}(\xi) - \xi - 1]\mathbf{c} = 0, \quad (\text{A6})$$

where the matrix  $\underline{M}$  is given by

$$M_{rsr's'}(\xi) \equiv \gamma_{rsr's'} e^{-\alpha \tau_{r's'} \xi} - \gamma_{rsR_s} e^{-\alpha \tau_{R_s} \xi}. \quad (\text{A7})$$

Note that Eqs. (A6) and (A7) reduce to Eq. (9), provided that  $\rho$  is redefined to be a function of the single variable  $f_1$ , which is obtained from  $\rho_1(f_1, f_2)$  by substituting  $1 - f_1$  for  $f_2$ .

#### APPENDIX B: CORRELATION FUNCTIONS

In order to compare the theoretical results with simulations, it is useful to calculate the autocovariance and autocorrelation functions  $V(t')$  and  $C(t')$ , which, for the simplest case of two resources and one species of agent, are defined by

$$V(t') = V(0)C(t') = \left\langle \frac{1}{T_f - T_i} \int_{T_i}^{T_f} dt f(t) f(t - t') \right\rangle - \bar{f}^2, \quad (\text{B1})$$

where

$$\bar{f} = \left\langle \frac{1}{T_f - T_i} \int_{T_i}^{T_f} dt f(t) \right\rangle, \quad (\text{B2})$$

and  $T_i$  and  $T_f$  are sufficiently large to ensure that initial transients have died out and an adequate sample of  $f(t)$  has been taken. In the mean-field theory, the autocorrelation is approximated by replacing the integrand of Eq. (B1) by  $\langle f(t) \rangle \langle f(t - t') \rangle$ , with the result (valid for  $\beta < \beta_2$ )

$$C(t') = e^{-\alpha \xi_r t'} \cos(\alpha \xi_i t'),$$

$$V(0) = 0.$$
(B3)

In order to compare these predictions with simulation results, we calculate the autocorrelation function  $C(t')$  with  $T_i > 4/\alpha \xi_r$ , and  $T_f > T_i + 4/\alpha \xi_r$  [which is sufficient to ensure that  $C(t')$  is virtually independent of  $T_i$  and  $T_f$ ] and average it over several simulation runs. The result is then fitted to a decaying sinusoid in the range  $0 \leq t' \leq 1.61\tau_r$  to extract estimates of  $\xi_r$  and  $\xi_i$ , from which  $\tau_r$  and  $\nu$  are calculated. The standard deviation  $s$

of the fluctuations about the average value of  $f$  is given by  $\sqrt{V(0)}$ . If  $\beta > \beta_2$ , the fixed point  $f_0$  is unstable, and Eq. (B3) is not strictly valid. In this case, the theoretical value of  $\xi_i$  is obtained by direct measurement of the period of the nonlinear oscillations in  $f(t)$ .

In more complicated systems with more resources or more species of agents, the system's behavior can be analyzed by calculating cross covariances and cross correlations in a manner similar to that described above. In this case, all of the cross correlations which describe the system have the same decay rate and frequency—that determined by the eigenvalue with the largest (least negative) real part.

<sup>1</sup>P. Brinch Hansen, *Commun. ACM* **21**, 11, 934 (1978).

<sup>2</sup>M. Factor and D. Gelernter, Yale University Report (1988) (unpublished).

<sup>3</sup>J. Harris (unpublished).

<sup>4</sup>R. Davis and R. G. Smith, *Artif. Intell.* **20**, 63 (1983).

<sup>5</sup>W. A. Kornfeld and C. W. Hewitt, *IEEE Trans. Syst. Man Cybern.* **SMC-11**, 24 (1981).

<sup>6</sup>M. S. Miller and K. E. Drexler, in *The Ecology of Computation*, edited by B. A. Huberman (North-Holland, Amsterdam, 1988).

<sup>7</sup>T. W. Malone, R. E. Fikes, K. R. Grant, and M. T. Howard, in *The Ecology of Computation*, edited by B. A. Huberman (North-Holland, Amsterdam, 1988).

<sup>8</sup>M. Minsky, *The Society of Mind* (Simon and Schuster, New York, 1986).

<sup>9</sup>C. J. Krebs, *Ecology* (Harper and Row, New York, 1972).

<sup>10</sup>J. M. Smith, *Evolution and the Theory of Games* (Cambridge University Press, Cambridge, England, 1982).

<sup>11</sup>B. A. Huberman and T. Hogg, in *The Ecology of Computation*, edited by B. A. Huberman, (North-Holland, Amsterdam,

1988).

<sup>12</sup>H. A. Ceccato and B. A. Huberman, *Proc. Natl. Acad. Sci. U.S.A.* **86**, 3443 (1989).

<sup>13</sup>H. A. Ceccato (unpublished).

<sup>14</sup>R. Bellman and K. L. Cooke, *Differential-Difference Equations* (Academic, New York, 1963).

<sup>15</sup>J. D. Farmer, *Physica* **4D**, 366 (1982).

<sup>16</sup>J. P. Crutchfield and B. A. Huberman, *Phys. Lett.* **77A**, 407 (1980).

<sup>17</sup>R. A. Sorensen, *Synthesis* **63**, 157 (1985).

<sup>18</sup>R. Axelrod, *The Evolution of Cooperation* (Basic, New York, 1985).

<sup>19</sup>G. Hardin, *Science* **162**, 1243 (1968).

<sup>20</sup>R. Lande, *Proc Natl. Acad. Sci. U.S.A.* **82**, 7641 (1985).

<sup>21</sup>C. M. Newman, J. E. Cohen, and C. Kipnis, *Nature* **315**, 400 (1985).

<sup>22</sup>J. S. Rosenschein and M. R. Genesereth, in *The Ecology of Computation*, edited by B. A. Huberman (North-Holland, Amsterdam, 1988).