DYNAMICS WITH EXPECTATIONS

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

By

Natalie Sandrine Glance

June 1993

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Benardo A. Huberman (Principal Advisor)
(Applied Physics)

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Alexander Fetter (Physics)

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

William A. Little (Physics)

Approved for the University Committee on Graduate Studies:

Abstract

The study of dynamics with expectations is relevant whenever the entities or “agents” comprising a complex system take into account possible future states when making decisions in the present. Goal-seeking agents form plans to achieve their ends based on their knowledge of causes, their memory of the past, their predictions for the future, as well as on their biases and their beliefs. A dynamical formulation of agent interactions requires a different approach when the behavior of the entities depends on their beliefs and knowledge as well as on the rules governing the evolution of the system. The agents’ predictions for the future must enter explicitly into the dynamical rules governing the evolution of the system.

This thesis examines the behavior of two dynamical systems in particular. The first, referred to as a computational ecosystem, is composed of a collection of processes, or agents, that compete among themselves for limited resources. The analytical tools of nonlinear dynamics and the modern-day aid of numerical integration via computer are employed to map out the behavioral regimes of the system first when the external environment changes over time, and second when the individual agents incorporate expectations about the future into their decisions.

In the second part, a system composed of interacting agents whose individual goals conflict with the overall good of the group is studied both analytically and through computer simulations. It is shown how spontaneous cooperation can be achieved if the agents take into account their expectations about the future when making their choices in the present. Further, the system can remain trapped in a metastable state for long periods of time; eventually, rare large fluctuations occur due to uncertain information that can cause a transition to the global equilibrium. The effects of ultrametric group structure and diversity on the dynamics of cooperation are also elucidated. Specifically, it is seen how the interplay between evolving group structure and the number of agents cooperating introduces new dynamical pathways to cooperation.

Acknowledgments

I am very grateful to my advisor, Bernardo Huberman, for his guidance and enthusiasm in matters scientific and his thoughtfulness and concern in matters personal. I feel lucky to count myself among the few who find in their advisor both a mentor and a friend.

To the others at Xerox PARC who have enriched my years here and contributed to my growth as both a researcher and a person, I also extend my thanks: in particular, to Scott Clearwater, Baldo Faieta, Tad Hogg, Erik Lumer, Brian Williams and Colin Williams. Many thanks as well go to my friends in the real world, those kindred spirits who have laughed with me, listened, understood, explained, and shared with me their lives. And, thank you David, my husband, for our love and our life together.

Finally, I thank my parents, for fostering in me a thirst for knowledge, the confidence to pursue my dreams, and the wherewithal to achieve them.

Contents

List of Tables

List of Figures

1 Introduction

It is a long tradition in physics to proceed from the study of very simple systems to more complex ones. In the case of dynamics, one can think of the transition from simple two–body problems to the many-body problem in celestial mechanics or condensed matter physics. Recent developments in nonlinear dynamics have led to novel insights into the dynamics of systems with many degrees of freedom (Schuster 1988). The techniques of statistical mechanics (Reichl 1987) and computer simulations also provide tools with which to study these systems (Koonin 1990).

Over the past decade or two a physics of complex systems has sprung up from these roots which extends the traditional domain of physics to include the study of large systems made up of many simple interacting units. This arm of research builds upon the large base of ideal models of physical systems developed to shed light on the global properties that result from the sum of local interactions. For example, the Ising model (Ising 1925), a simple conception of magnetic systems, led to an understanding of how short-range interactions among particles with spin leads to an overall magnetic moment. The ideal gas is another type of simple model in which the abstraction of the individual elements and their interaction makes tractable the derivation of global properties. From the idealized properties of molecules in such a gas, macroscopic properties are derived which describe real gases well within certain constraints. For these models, the techniques of statistical mechanics, in which a description of the typical behavior and interactions of a system's elements is used to determine overall behavior, prove invaluable. The macroscopic properties of systems with degrees of freedom on the order of Avogrado's number can be derived if the interactions are simple enough.

Working from this paradigm, researchers have tackled the questions of how self-organized criticality applies to earthquakes or landforms (Bak *et. al.* 1988), how randomness and disorder can result in general properties (Grimmet 1957, Broadbent and Hammersly 1957, Abrahams *et. al.* 1979), and how memory and learning might emerge from collections of "neurons" (Little 1975, Hopfield 1982, Amit *et. al.* 1985) In recent years, an inquiry into the physics of computational systems has also developed. A review of the perspective gained by a physics approach can be found in (Huberman and Hogg 1987). Percolation-like phase transitions in the topological structure of the search space of large algorithms have been found and demonstrated empirically for several different search processes (Huberman and Hogg 1987, Langton 1991, Cheeseman *et. al.* 1990, Williams and Hogg 1993). Lumer (1990) has elucidated the ultradiffusive dynamics of distributed resource allocation in hierarchical systems. In addition, Clearwater, Hogg, and Huberman (1991, 1992) have shown how cooperative interactions yield super-linear speed-up in distributed problem-solving performance. The work presented in this thesis also makes a contribution to this area.

In fact, there are a variety of interesting dynamical problems in the biological and social, as well as computational, arenas that physics' techniques are uniquely able to treat. In particular, when these dynamical systems can be modeled as collections of agents interacting according to a set of laws or rules in the game-theoretical sense (von Neumann and Morgenstern 1944, Rasumsen 1989), then their dynamics can be studied using several methods. On the one hand, a statistical mechanical treatment of the system yields a closed form description of average behavior; on the other, computer simulations provide an "experimental" playing field for observing and understanding the system's behavior.

1.1 Dynamics with Expectations

In many computational, economic, or biological systems, the comprising entities or "agents" have additional non-Newtonian properties. The agents may form plans, or be programmed to plan ahead. As a result, their actions appear to be (or are) intentional. For example, beliefs about the future could be programmed into the design of computational agents or evolve genetically in biological ones. Alternatively, social agents might form expectations because of their intentional nature. This property of intentionality is absent in the objects that are the purview of physics and chemistry.

The agents make plans to achieve their goals based on their knowledge of causes, their memory of the past, their predictions for the future, as well as on their biases and their beliefs.

Such agents may also be capable of learning, possibly by adopting strategies that maximize their performance. Or they may simply appear to learn, if agent types evolve over generations, according to some fitness function. Along these lines, a computational analogue of genetic evolution called genetic algorithms (Holland 1975) has become a common optimization technique in which a population of solution "agents" are bred until a satisfactory solution is reached. In this way, through learning and planning, agents can to some extent adapt to the changing environment made up of the other agents and external events. Examples of cases in which individuals' expectations of others' behavior plays an important role in the economic and computational arenas include stock trading, currency markets, and resource contention.

A dynamical formulation of agent interactions requires a different approach when the behavior of the entities depends on their beliefs and knowledge as well as on the rules governing the interaction. The agents' predictions for the future must enter explicitly into the dynamical rules governing the evolution of the system. In general, the agents' expectations for the future can be recast as a prescription on how to act in the present based on imperfect present and past knowledge in a way that takes their beliefs into account implicitly.

1.2 Overview

The study of dynamics with expectations is then relevant whenever the agents comprising a complex system, be they computational processes or biological entities, take into account possible future states when making decisions in the present. Two systems in particular are the subject of study of this thesis. The first, established under the rubric computational ecosystem by Huberman and Hogg (1988), consists of a loosely coupled collection of agents which compete among themselves for resources according to specialized strategies. Specifically, each agent chooses among different resources according to its perceived payoff for using each resource, which depends on the number of agents already using it. Expectations come into play if agents use past and present global behavior in estimating the expected future payoff for each resource. An extended dynamical formulation of computational ecosystems is necessary to understand the effect of these expectations on the global performance of the system.

A second context in which the effect of agents' expectations becomes important is the question of how spontaneous cooperation in a group can be achieved through individually rational decisions. This problem of collective action arises from the conflict between individual benefit and collective good (Olson 1965). Such dilemmas underlie many ongoing collective action problems, such as the functioning of large organizations (Bendor and Mookherjee 1987) and the mobilization of political movements (Hardin 1972, Oliver 1988). In a general social dilemma, a group attempts to obtain a common good in the absence of a central authority. Each agent has two choices: either to contribute to the common good, or to shirk and free ride on the work of others. The logic behind the decision to cooperate or not changes when the interaction is ongoing since future expected utility gains join present ones in influencing the rational individual's decision. This indicates the importance of including expectations in a dynamical description of ongoing collective action in a group of social agents.

Another aspect which comes to the forefront when studying the dynamics of collective action is the organizational structure of the group. This structure emerges from the pattern of interdependencies among agents. In a fluid structure, the pattern of interactions can vary widely over time, since the sum of small local changes in the structure of a group results in overall broad restructurings. If individual decisions to locally alter the structure depend on the perceived benefit as opposed to the actual benefit, then expectations once again play a role.

A dynamical model of collective action that includes expectations may then help answer the following questions: if agents make decisions on whether or not to cooperate on the basis of imperfect information about the group activity, and incorporate expectations on how their decision will affect other agents, can overall cooperation be sustained for long periods of time? Moreover, how do expectations, group size, and diversity affect cooperation? And lastly, which kinds of organizations are most able to sustain ongoing collective action, and how might such organizations evolve over time? These questions are relevant to both social science and to distributed artificial intelligence (Gasser and Huhns 1989), where instances of negotiation and cooperative problem solving imply the existence of intentional agents (Davis and Smith 1983, Zlotkin and Rosenschein 1991, Osawa and Tokoro 1991). In addition, the emergence of market-like computational systems that have no global controls (Malone *et. al.* 1988, Waldspurger 1992) raises the issue of having computational processes that could free ride on the work of others (Miller and Drexler 1988a,b).

The format of the thesis proceeds as follows. In Chapter 2, I examine the behavior of computational ecosystems in a changing environment. The emphasis is on the dynamics of such a system and on how well agents without learning or predictive capabilities can track the changes in the environment through elementary decision rules. In Chapter 3, the agents' strategies are extended to take into account their expectations of how the system's behavior will evolve into the future. In order to separate the effects of a changing environment and agent expectations, the environment is taken to be fixed.

In the second part of the thesis, I turn to the problem of how ongoing collective action can arise and be sustained in groups of economic or computational agents. In Chapter 3, a dynamical model of collective action is developed and the possibilities for cooperative behavior are elucidated. In Chapter 4, I extend the model to allow for a diversity of beliefs among the agents and find how the dynamical regimes change. In Chapter 5, I examine the effect of ultrametric group structure on the dynamics of cooperation. Specifically, I study how the interplay between evolving group structure and levels of cooperation can yield new regimes of dynamical behavior.

Finally, in the appendix, I elaborate on one particular aspect of how computer simulations are performed: the choice between synchronous *versus* asynchronous updating of agent states. This argument was laid out by my advisor, Bernardo Huberman, and myself in response to a paper by May and Nowak (1992) in particular, and to a tendency in the theoretical biology and artificial life communities to tacitly assume discrete dynamics by performing synchronous simulations, in general. The contents have been submitted for publication (Huberman and Glance 1993b).

As concerns the rest of the thesis, the chapters have all been published or are in the process of being published, although in somewhat different form. In particular, Chapter 2 corresponds roughly to (Glance and Huberman 1991), Chapter 3 to (Glance and Huberman 1992a, 1992b), Chapter 4 to (Glance and Huberman 1992c), Chapter 5 to (Huberman and Glance 1993a) and Chapter 6 to (Glance and Huberman 1993a).

PART I

COMPUTATIONAL ECOSYSTEMS

2 Changing Environments

2.1 Introduction

Social conventions and the cycles of nature create periodic variations in the prices of commodities and resources. The use of computers in the workplace, for example, tends to be high during the day but low at night, and fluctuates similarly between workdays and weekends. The availability of the machines varies in tandem in a way beyond the individual's control. Likewise, the variation in price of fresh produce through the year is caused by the fluctuation of seasonal outputs in agriculture, and the familiar cycles of electricity consumption and phone utilization lead to pricing structures that exploit their periodicity.

Determining the global response to periodic fluctuations of a large system consisting of many autonomous agents making local decisions is generally much harder than predicting their influence on isolated individuals. The aggregate behavior of a system of agents, however, determines important characteristics such as system performance and adaptability to a changing environment. As distributed systems arise in many different contexts including computation, biological ecosystems (Krebs 1972), economies, and social structures (Simon 1962), understanding aggregate behavior under varying conditions may indicate better designs.

A computational ecosystem (Huberman and Hogg 1988) is a type of distributed computation in which a set of independent agents vie for limited resources. These ecosystems embody several characteristics of distributed systems including distributed control, asynchrony in execution, resource contention, and cooperation among agents. They also suffer from the concomitant problems of incomplete knowledge and delayed information. The agents choose between resources based on their perceived payoffs. In their analysis of the dynamical behavior of computational ecosystems, Huberman and Hogg found that the systems display several behavioral regimes which, depending on particular system parameters, are characterized by fixed points, oscillations, and even chaos. Kephart *et. al.* (1989) use computer simulations to verify that the theoretical predictions are accurate when there are at least a few hundred agents in the system.

The work on computational ecosystems assumes resource utilities that depend endogenously on the behavior of the agents. In this chapter, we will modify the model to include time-dependent variations in the environment, an exogenous variable. The effect of a changing environment is modeled by a sinusoidal time-fluctuating component in the resource payoffs. We assume that the agents are not able to perceive or predict patterns in either the behavior of the other agents or in the external environment. The effect of predictive computational agents is studied elsewhere (Kephart 1990).

In general, environmental changes are often uncontrollable and unpredictable, not periodic. Allowing the payoff functions to vary erratically would produce results which would be very difficult to interpret[1]. However, a simple sinusoidal modulation should capture many of the pertinent effects. From analogy with driven oscillators (d'Humieres *et. al.* 1982), it is expected that the interplay between the frequency of the global utility variation and those of an intrinsically nonlinear system leads to rich dynamical behavior.

Studying computational ecosystems in a changing environment also sheds light on the issue of adaptability. What are the characteristics of an adaptable system, and is it possible to predict its adaptability from knowledge of the underlying dynamics? These questions are important since the existence of a stable equilibrium for an isolated computational ecosystem does not at all imply that the system is adaptive. Instead it is possible that the system's very stability may prevent it from adjusting to time-dependent constraints.

In the following section, we describe the model of computational ecosystems in a changing environment. In Section 2.3, we analyze the underlying differential equations to determine the effect of time-modulation on dynamical stability in various limits for regimes in which the system exhibits some ability to track the fluctuating environment. In Section 2.4, we map out the remaining behavioral regimes of the system, some of which display severe lack of tracking, using numerical integration. Finally, in Section 2.5, we address the question of adaptability, defining

---

[1] Also, note that a simple random zero-mean uncorrelated fluctuation is already included in the theory to model information uncertainty.

a measure of performance which is used to describe how well a set of agents and resources is able to adapt to a time-varying environment.

2.2 The Model

In the model of computational ecosystems, *N* agents are free to choose among *R* resources according to the *perceived* (*i.e.*, not necessarily correct) expected payoff for using each resource. These payoffs are actual computational measures of the resources's performance, such as the time required to complete a task, accuracy of the solution, amount of memory required, *etc*. Competition and cooperation among the agents are taken into account 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 *i* would decrease monotonically with the fraction of agents, $f_i(t)$, already using it. Alternatively, the agents using the same resource could assist one another, and 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 α and switches to the resource with the highest payoff.

In addition, the information available to the agents may be imperfect and delayed in time. Because of imperfect information, each agent's perceived payoff differs somewhat from the actual value. The difference increases when there is more uncertainty. This simple model of uncertainty captures the effect of many sources of errors including some kinds of program bugs, the incorrect evaluation of choices by heuristics, errors in communicating the load on various machines and mistakes in interpreting sensory data. Secondly, information delays cause each agent's knowledge of the state of the system to be somewhat out of date. To account for the fact that an agent's 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 σ to each (instantaneous) payoff, and delay the information available to each agent by a time τ.

In this chapter, we consider a simple ecosystem composed of identical agents competing for two resources. For this case, the time-evolution of the fraction of agents using resource 1 is described within the mean-field approximation by the differential-delay equation (Bellman and

Cooke 1963)

$$\frac{df_1(t)}{dt} = -\alpha[f_1(t) - \rho(f_1(t-\tau))] \tag{2.1}$$

where $\rho(f)$ is the probability that a single agent will choose resource 1 if a fraction f of the agents is already using it, α is the average rate at which agents reevaluate their choices, and τ is the delay (Huberman and Hogg 1988). In terms of the density-dependent payoff functions for using resources 1 and 2, $G_1(f_1(t))$ and $G_2(f_1(t))$, and the uncertainty σ characterizing the size of the errors in the agents' perceptions, $\rho(f)$ is given by

$$\rho(f_1(t)) = \frac{1}{2}\left[1 + \mathrm{erf}\left[\frac{G_1(f_1(t)) - G_2(f_1(t))}{2\sigma}\right]\right] \tag{2.2}$$

The fraction of agents using resource 2 is then simply $1 - f_1(t)$.

The nature of the agents' interactions determines the functional dependence of the payoff functions on the fraction of agents using each resource. In a purely competitive environment, the payoff for using a particular resource decreases as more agents make use of it. Alternatively, the agents using a resource could assist one another in their computations. Such might be the case if the overall task was decomposable into subtasks. If these subtasks communicate extensively to share partial results, the agents might be better off sharing the same computer rather than running more rapidly on separate machines but then being limited by slow communications. As another example, agents using a particular database could leave index links that are useful to others. In such cooperative situations, the payoff of a resource increases as more agents use it, at least until it became too crowded. A more general scenario combines the two effects of competition and cooperation, giving rise to a payoff function which at first increases as the number of agents using the resource increases but eventually saturates. The simplest example of this is a payoff for using a resource which reaches a maximum somewhere in the interval $0 < f < 1$ as apposed to a purely competitive payoff which decreases monotonically with increasing f. For simplicity, we refer to the cooperative/competitive scenario as the cooperative case.

The behavior of computational ecosystems for two different choices of payoff functions corresponding to competitive and cooperative systems has been studied in detail in (Kephart *et. al.* 1989). The competitive payoffs considered were given by

$$G_1 = 7 - f_1;$$
$$G_2 = 7 - 3f_2. \tag{2.3}$$

For this situation, the dynamics of the undriven system is well-understood. Given the uncertainty σ, $\beta\ (= \alpha\tau)$ is varied to obtain three types of behavior: (1) exponential decay to the fixed point, f_0, given by $\rho(f_0) = f_0$; (2) damped oscillations about f_0; and (3) persistent oscillations. The behavioral phase diagram for this system given in Fig. 2.1(a).

The case of a system of cooperating agents with the following payoff functions was also studied:

$$G_1 = 4 + 7f_1 - 5.333f_1^2$$
$$G_2 = 7 - 3f_2. \tag{2.4}$$

The payoffs were chosen to have the same fixed point as for the competitive case when there is no uncertainty. This combination of agents and resources exhibits stability, instability, bifurcation, and chaos for different choices of β and σ (see the behavioral phase diagram given in Fig. 2.1(b)).

In the previous examples, G_1 and G_2 depend only on f_1 and f_2, and only indirectly upon time. To investigate the behavior of this model within a changing environment, we now take the utility of one of the resources to vary sinusoidally in time. This could model, for example, the effect of diurnal variations in the load on a network of computers. The variation in load capacity is considered to be an accumulation of effects external to the ecosystem. Specifically, we introduce an explicit time-dependence for the utility of resource 1, say, by taking

$$G_1(f_1) \rightarrow G_1(f_1) + A \sin \omega t \tag{2.5}$$

with *A* the amplitude and ω the frequency of the sinusoidal variation in the payoff of resource 1. In this way, changes in the resource's base capacity are taken to be independent of the agents' actions. Note that since the payoffs for using the two resources enter into the equations describing the evolution of the ecosystem only through their difference, $G_1 - G_2$, time-modulating the payoff of one or the other leads to identical results. This does not remain true if there are more than two resources.

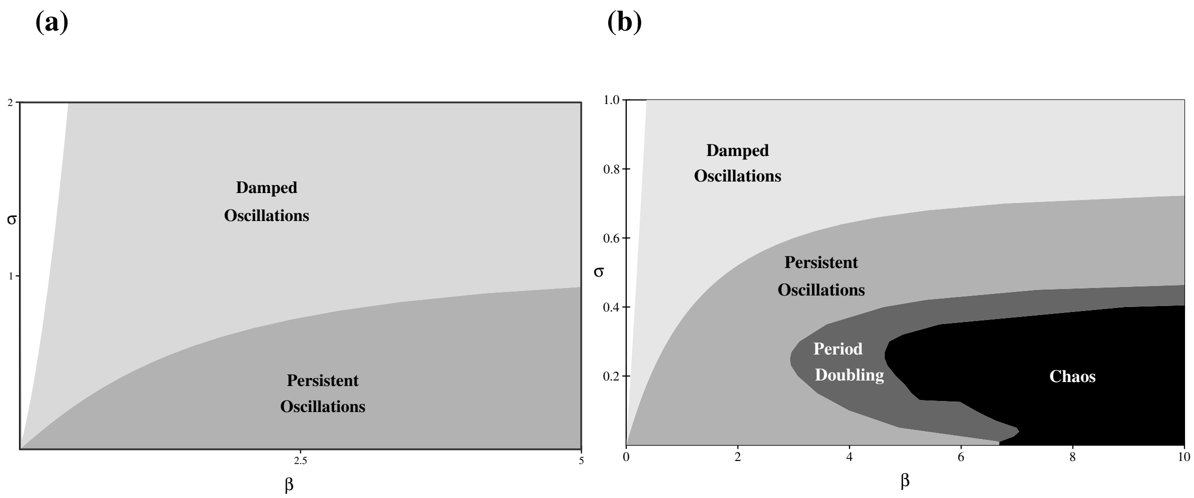

Figure 2.1: *Behavioral phase diagram* as a function of $\beta = \alpha\tau$ and uncertainty parameter σ for an undriven system of identical agents with the competitive payoffs $G_1 = 7 - f_1$ for resource 1 and $G_2 = 7 - 3f_2$ for resource 2 in (a); with the cooperative payoffs $G_1 = 4 + 7f_1 - 5.333f_1^2$ and $G_2 = 7 - 3f_2$ in (b). In the unshaded regions, the system exhibits overdamped (non-oscillatory) relaxation to the fixed point. The chaotic region in (b) is punctuated by numerous windows of periodic behavior (not shown). Results were obtained by analysis and by integrating Eq. 2.1 numerically.

Lastly, we remark on a notational convenience adopted for the remainder of the chapter. In order to make explicit the functional dependence of the probability ρ on the driving term $u \equiv A \sin \omega t$, we will refer to the probability ρ by $\rho(f, u)$ or, alternatively, $\rho(f, A)$, when the frequency is irrelevant to the discussion.

2.3 Tracking the Utility Cycle

For a wide range of parameter values, a computational ecosystem in a static environment evolves to a fixed equilibrium, regardless of the initial conditions (Kephart *et. al.* 1989). How will such a system behave if the environment fluctuates cyclically over time? In particular, how well will the aggregate behavior of the agents track the changing phase of the fluctuation? In this section, we will derive the behavior for the limiting cases of small and large amplitudes as well as low and high frequencies. The derivations are supplemented by examples of the time evolution of the system obtained through numerical integration of Eq. 2.1.

2.3.1 Small amplitudes

In the limit of very small amplitude oscillations in the resource utilities (*i.e.* A much smaller than the typical difference between the payoffs G_1 and G_2), we expect the effects of the driving term to vanish. As a result, we consider the solution to Eq. 2.1 in which the system oscillates at the driving frequency about the equilibrium value f_0, which is given by

$$f_0 = \rho(f_\circ, A = 0) \tag{2.6}$$

We can obtain the conditions under which such oscillations are stable by linearizing Eq. 2.1 in the neighborhood of the fixed point f_0 and about $A = 0$. Defining the variation around the fixed point as $\delta(t) = f(t) - f_0$, we obtain

$$\frac{d\delta}{dt} + \alpha[\delta(t) - \gamma\delta(t-\tau)] = \alpha\eta A \sin\omega(t-\tau) \tag{2.7}$$

where $\gamma = \rho_f(f_0, A = 0)$ and $\eta = \rho_u(f_0, A = 0)$ are the partial derivatives of ρ with respect to f and u, the latter defined by $u = A\sin\omega t$. The solution to this inhomogeneous differential–delay equation (Bellman and Cooke 1963) consists of the sum of the particular solution, $\delta(t) = B\sin(\omega t + \phi)$, where both B and ϕ depend in a complicated fashion upon A, α, τ, and σ, and the solution to the homogeneous equation $\frac{d\delta}{dt} + \alpha[\delta(t) - \gamma\delta(t-\tau)] = 0$. Since the particular solution remains small (*i.e.*,$B \to 0$ as $A \to 0$), the stability of Eq. 2.7 is determined by the behavior of the homogeneous equation, as previously studied in the context of time-independent resource utilities.

The results of that analysis were discussed in the previous section and displayed in Fig. 2.1, which shows the oscillation and stability thresholds for two different sets of payoffs (competitive and cooperative) for an undriven computational ecosystem. Since a system in a changing environment is "stable" for the same parameter values[2], we can apply these previous stability results. Specifically, for the same parameters that the undriven system relaxes to a fixed point, the driven one tracks the time-dependent variation with some phase shift dependent on both α and τ. The boundaries delimiting this behavioral regime correspond to the stability boundaries of the undriven system.

[2] The term "stability" is used in the context of driven computational ecosystems to indicate behavioral oscillations at the frequency of the driving term.

This tracking behavior is shown in Fig. 2.2: the fraction of agents $f(t)$ using resource 1 ($\equiv f_1(t)$) exhibits small amplitude oscillations at the driving frequency with a phase lag. (Undriven, the system relaxes to a fixed point.) Here, as elsewhere, the evolution of the system's behavior is obtained by numerical integration of Eq. 2.1. In this case the phase lag between the driving signal and the response of the system is 1.45 time units, significantly smaller than the delay, which equals three time units. Since the phase lag can vary from 0 to 2π (depending on the value of the delay and the reevaluation rate), this result indicates that the system is only partially able to track the time variation in the relative payoffs between the resources.

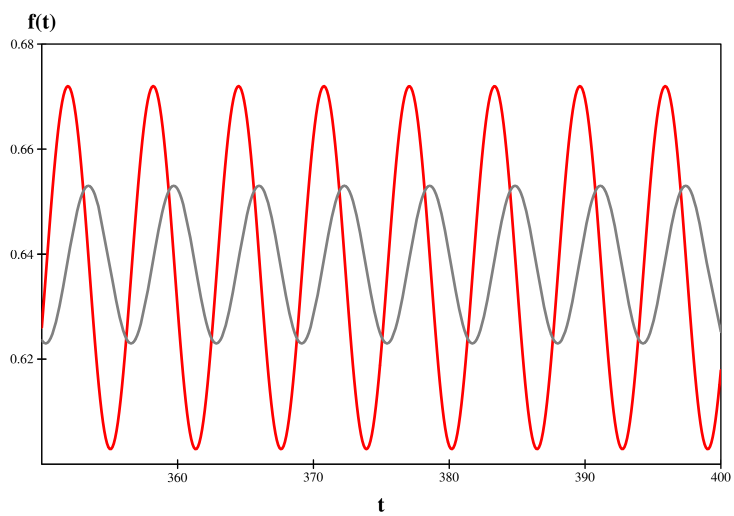

Figure 2.2: *Low amplitude driving.* Dark curve shows fraction of agents $f(t)$ using resource 1 ($\equiv f_1(t)$) with competitive payoffs $G_1 = 7 - f_1$ and $G_2 = 7 - 3f_2$ when driven at frequency ω=1, small amplitude A=0.1, for parameters α=1, τ=3, σ=0.9. Notice that $f(t)$ remains close to the fixed point of the undriven system, $f_0 = 0.64$. The light curve shows the phase of driving term — its scale is unrelated to amplitude of driving term.

2.3.2 High amplitudes

In the limiting case of high amplitude, Eq. 2.1 can be solved exactly to zeroth order. If A is much greater than the typical difference between the payoffs G_1 and G_2, the value of $\rho(f, u = A\sin\omega t)$ is determined to zeroth order by the sign of the sine function. Thus, when $\sin(\omega t)$ is positive $\rho = 1$, and $f(t)$ relaxes exponentially towards one, whereas when $\sin(\omega t)$ becomes negative $\rho = 0$, and $f(t)$ decays exponentially to zero. However, since ρ is evaluated at $t - \tau$, not at t, the switching between growth and decay will lag the driving term by an amount equal to the delay τ. While this analysis is only a zeroth order solution, numerical integration of Eq. 2.1 verifies its accuracy. Fig. 2.3 shows the typical response of the system when there is a large amplitude modulation of the utilities for a system with two resources. Here $A = 100$, $\omega = 0.5$ ($\tau = 1$, $\alpha = 1$, $\sigma = 0.5$) and the oscillation lags the driving frequency by τ, after an initial transient.

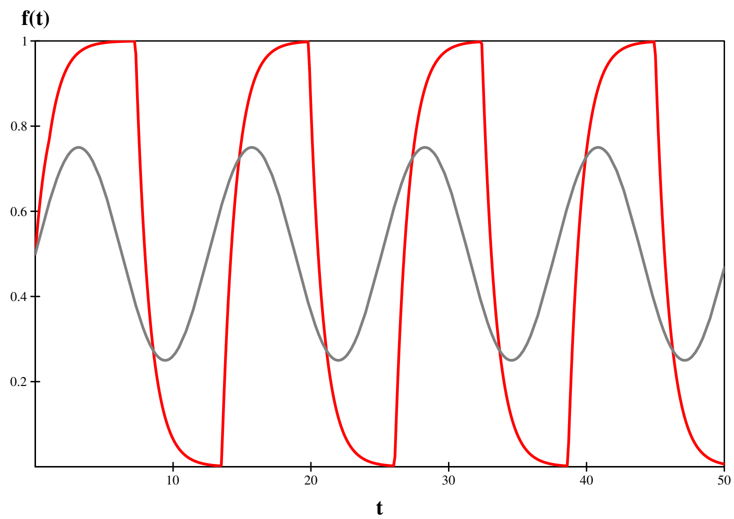

Figure 2.3: *High amplitude driving.* Example of high amplitude modulation of the resource utilities. Dark curve shows fraction of agents $f(t)$ using resource 1 with competitive payoffs $G_1 = 7 - f_1$ and $G_2 = 7 - 3f_2$ when the system is driven at frequency $\omega = 0.5$, high amplitude $A = 100$, for parameters $\tau = 1$, $\alpha = 1$, and $\sigma = 0.5$. Light curve shows phase of driving term. Again, its scale is unrelated to amplitude of driving term.

In closing, we mention that as in the previous case of low amplitudes, $f(t)$ can oscillate 180° out of phase with the driving term (if $\tau = \pi/\omega$), resulting in a loss of tracking. We have observed this behavior in response to high amplitude driving everywhere in the behavioral phase diagram of the computational ecosystem (*i.e.* for all values of uncertainty, delay, and reevaluation rate).

2.3.3 Intermediate amplitudes

Next we examine the behavior of the system when the resource utilities are modulated with intermediate amplitude. The results obtained for small and high amplitudes lead us to expect that the system will oscillate at the same frequency as the driving term and with some phase shift for those parameters corresponding to the stable regime of the undriven system. Indeed, for both high and low amplitudes the stable region of the undriven system maps onto oscillatory behavior at the driving frequency with some phase lag. Therefore, the system is unable to exactly track the external variation in the environment even in parameter regions corresponding to stable behavior for the undriven system. Note that in the high amplitude limit, the phase lag is determined by τ, while for very low amplitudes the phase lag depends both on α and τ in a non-trivial way. Thus, in general, we expect that the dynamics of the system will depend on both α and τ independently, not simply on their product, unlike the case of the undriven system whose behavior can be characterized by the parameter $\beta = \alpha\tau$.

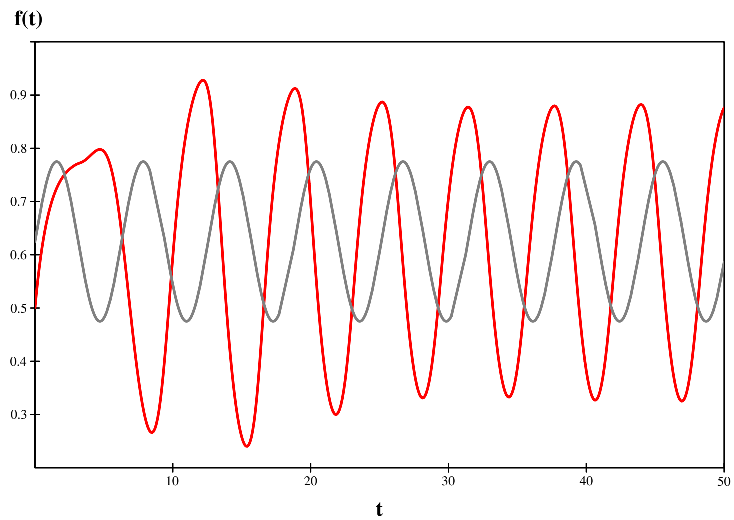

Figure 2.4: *Intermediate amplitude driving.* Competitive system of agents in a changing environment within parameter region corresponding to damped oscillation region of undriven system. Dark curve shows fraction of agents $f(t)$ using resource 1 for competitive payoffs $G_1 = 7 - f_1$ and $G_2 = 7 - 3f_2$ when driven at frequency $\omega = 1$, intermediate amplitude $A = 1$, for parameters $\alpha = 1$, $\tau = 3$, and $\sigma = 0.9$. Light curve shows phase of driving term (scale is unrelated to amplitude of driving).

Fig. 2.4 shows the fraction of agents using resource 1 *vs.* time when the relative payoffs of the two resources are modulated at a frequency $\omega = 1$ with amplitude $A = 1$. This is shown superimposed upon a schema of the sinusoidal driving term to make apparent the phase difference between them. The phase difference in this case equals 1.55 time units. Notice that at first the phase lag between the driving term and the behavior of the agents is about one delay period (which is only to be expected since it takes the system one delay period to observe the signal); as the system settles into its equilibrium behavior, however, the lag decreases to the value stated above.

2.3.4 The low frequency limit

We can better understand the behavior of the system at intermediate amplitudes by studying the low and high frequency limits. In the limit of very low frequency time-modulation of one of the resource payoffs, some analytic results can be obtained. As long as the period of driving is much longer than both the delay and $1/\alpha$, *i.e.* $1/\omega \gg 1/\alpha$ and $1/\omega \gg \tau$, the behavior over time intervals small compared with the period of oscillation of the payoff difference corresponds to the behavior found for the undriven case. Over longer time scales (on the order of $1/\omega$), the small time scale behavior is simply modulated by the low frequency driving. This can be shown by linearizing Eq. 2.1 about the equilibrium point $f_0(t)$, which is allowed to vary slowly in time, in the spirit of an adiabatic approximation. This equilibrium is given by

$$f_0(t) = \rho(f_0(t-\tau), A\sin\omega(t-\tau)) \tag{2.8}$$

where $\tau \ll 1/\omega$. Since $f_0(t)$ changes slowly, we use the approximation $f_0(t) \approx f_0(t-\tau)$ when solving for $f_0(t)$.

Defining the variation around the time-dependent equilibrium value as $\delta(x) = f(x) - f_0(t)$ and using x as a rescaled time variable describing intervals much smaller than $1/\omega$, we obtain a differential equation describing the evolution of the variation $\delta(x)$ about the fixed point. This equation is valid for time scales short compared with $1/\omega$. Thus, we recover the familiar equation

$$\frac{d\delta}{dx} + \alpha[\delta(x) - \gamma(t)\delta(x-\tau)] = 0 \tag{2.9}$$

where the time dependent variable γ is given by $\gamma(t) = \rho_f(f_0(t))$ and changes slowly.

This result indicates that the behavioral regimes of the system are well-described by the behavioral phase diagrams given for the undriven cases; however, the boundaries separating the different regimes shift in response to the driving.

If we choose a point in the phase diagram close to the stability threshold of the undriven system, we indeed find instances in which a low frequency modulation of the payoff difference repeatedly takes the system across the boundary separating two different behavioral regimes.

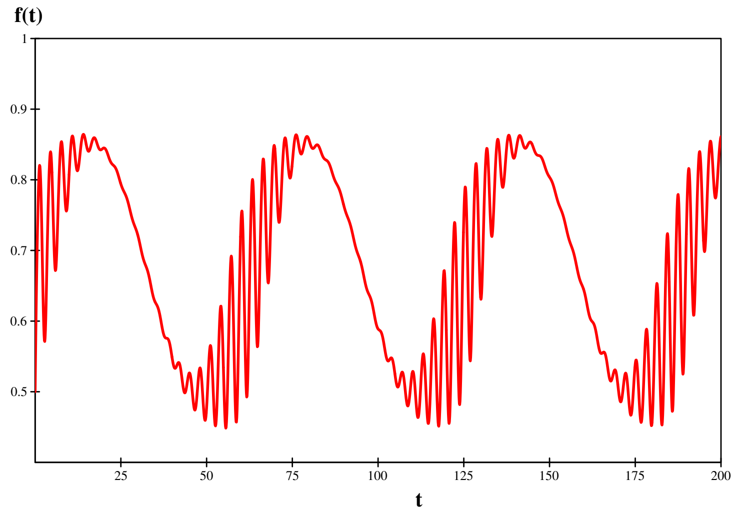

Figure 2.5: *Low frequency driving.* Shifting of the boundary between two different behavioral regimes occurs when system is driven at low frequencies. The figure shows the fraction of agents $f(t)$ using resource 1 for competitive payoffs $G_1 = 7 - f_1$ and $G_2 = 7 - 3f_2$ when driven at frequency $\omega = 0.1$, amplitude $A = 1$, for parameters $\alpha = 1$, $\tau = 1$, and $\sigma = 0.4$.

Fig. 2.5 portrays one such example. In this case, the choice of σ, α, and τ places the system right below the oscillation threshold for the undriven system of Fig. 2.1. Thus, a low frequency modulation of the payoffs changes the difference in payoffs, $G_1(t) - G_2(t)$, and shifts the boundary downwards. This shift continues lasts enough that the behavior of the system stabilizes until the next upward shift of the boundary causes oscillations to appear anew. This example potentially describes the behavior of a system which is stable at night when network use is conventionally light and inexpensive, but unstable during the day when use is heavy.

2.3.5 The high frequency limit

The case of high frequencies poses some interesting questions on how well the computational ecosystem model aligns with real distributed systems. If the environment varies at a rate much greater than the agents' reevaluation rate, do the agents see an average of the variation over time scales comparable to $1/\alpha$ or do they observe fluctuations that look like correlated noise? For the case of a sinusoidal modulation, the former assumption implies that the time-variation is not seen by the agents (zero mean), while the latter indicates that the resulting behavior might correspond to that of an undriven system with a renormalized value of the uncertainty.

In order to address this issue, we must reexamine how the agents obtain information about the resources. While the agents perform their evaluations asynchronously and independently, they could either accumulate information over discrete intervals of time or obtain it instantaneously. In our model, the agents are assumed to act in the latter fashion. Thus, in the mean-field limit of an infinite number of agents, during any short time interval many agents are gathering data about the various resources. The very rapid modulation of the payoffs then affects the system through its root–mean–square value, not through its average value. This remains true as the number of agents becomes finite, as long as there are a significant number of agent evaluations during one period of the very rapid modulation. Thus, we expect the mean-field equation to remain valid in the high-frequency limit of $\omega \gg \alpha$ if $N\alpha \gg \omega$, where N is the total number of agents.

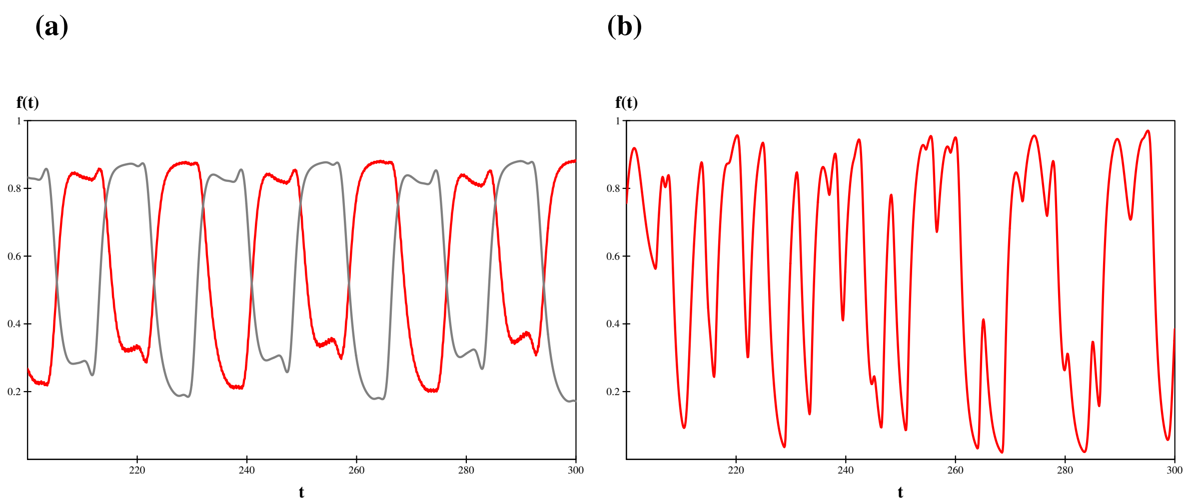

Figure 2.6: *High frequency driving.* Cooperative system of agents with resource payoffs $G_1 = 4 + 7f_1 - 5.333f_1^2$ and $G_2 = 7 - 3f_2$, α=1 and τ=8. In (a) the dark curve shows the fraction of agents $f(t)$ using resource 1 when the system is driven at frequency $\omega = 100$ with $\sigma = 0.2$. The light curve shows the behavior of the undriven system when the uncertainty level is increased to $\sigma - 0.438$. (b) The undriven system is chaotic for $\sigma = 0.2$, which is the same value of the uncertainty as in (a).

Numerical integration verifies the prediction that high-frequency modulation of the payoff utilities causes the system to behave as if the uncertainty were increased and the payoffs unmodulated, an effect which is independent of the stability of the undriven system. Thus, the overall effect of the high-frequency modulation is to stabilize the behavior of the agents in a way similar to the effect of increasing the uncertainty. The oscillation in the system's resource use is reduced in magnitude, fixed point behavior falls off from optimal values and chaotic systems become oscillatory. In addition, the amount by which the uncertainty must be increased in an undriven system to obtain similar behavior is a function of the amplitude of the modulation. For example, in Fig. 2.6(a) the effect of modulating the payoffs at frequency $\omega = 100$ with amplitude 0.75 is seen to be very similar to that of increasing the uncertainty from $\sigma = 0.2$ to $\sigma = 0.438$ in the undriven system. In this example, the resource payoffs are given by $G_1 = 4 + 7f_1 - 5.333f_1^2$ and $G_2 = 7 - 3f_2$, α=1 and τ=8. Either high-frequency modulation or the specified increase in the uncertainty take the system from a chaotic behavior to period-doubling. Fig. 2.6(b) shows the chaotic behavior of the system when the uncertainty is 0.2 and the payoffs are not modulated.

2.4 Loss of Tracking

In this section, we now consider the parameter regions in which the undriven computational ecosystem exhibits unstable behavior, such as oscillations, period-doubling and chaos. We will find that the interplay between the frequency of the global utility variation and those of the nonlinear system can lead to complex dynamical behavior not observed in the undriven system. A loss of tracking ability generally accompanies this complex behavior.

2.4.1 Quasiperiodicity and period-doubling

First, we study the effect of driving on a system whose natural undriven behavior is oscillatory (and describable by a single frequency). The resource payoffs used here are the competitive payoffs given in Section 2.3 with a sinusoidal term. The behavior of the system is obtained through numerical integration since closed—form analytical results elude derivation.

The simplest types of behavior correspond to the low and high amplitude limits and the low frequency limits. At very low frequencies, $\omega \ll \alpha$ and $\omega \ll 1/\tau$, there occurs beating of the natural frequency with the slowly changing driving signal and the system tracks successfully with negligible lag. For high frequencies and high amplitudes, the results of Section 2.3 hold: high frequency modulation has the same effect on the behavior of the agents as increasing their uncertainty, where the increase in σ depends on the amplitude; for high amplitudes, the system exhibits periodic behavior at the frequency of the time-modulation with a phase lag equal to the delay τ.

For low amplitudes, numerical integration of Eq. 2.1 shows that the natural frequency of the system takes over and the aggregate behavior of the agents fails completely to track the time-modulation. Surprisingly, these two limits of high and low amplitude appear to almost exhaust the range of possible behaviors, and there is only a small interval of amplitudes for given ω, α, τ, and σ which do not lead to either the natural frequency or the driving frequency taking over when the two frequencies are of the same order of magnitude.

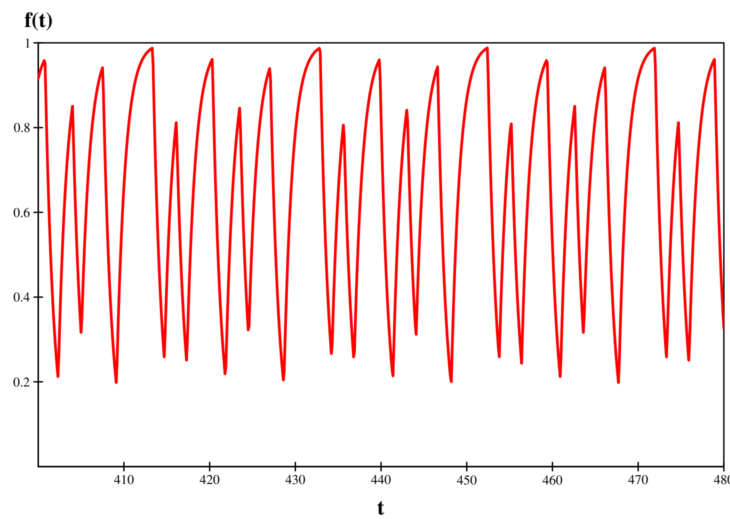

Figure 2.7: *Period-3 of the driving frequency.* Figure shows fraction of agents $f(t)$*f*(*t*) using resource 1 for competitive payoffs $G_1 = 7 - f_1$ and $G_2 = 7 - 3f_2$ when driven at frequency $\omega = 0.965$, amplitude $A = 1$, for parameters $\alpha = 1$, $\tau = 1$, and $\sigma = 0$.

However, when the strengths of the two signals, *i.e.* the natural oscillation of the system and the time-dependent modulation, are comparable in amplitude, the dynamics is complex and the agents cannot track successfully. For certain choices of amplitude and frequency in this regime, we find period doubling, periods 3, 5, 7, . . . of the driving frequency, and quasiperiodicity, phenomena which are commonly encountered in other nonlinear systems (Schuster 1988, Bai-Lin 1990). Fig. 2.7 shows an example of period 3 of the driving frequency. The choice of driving frequency, $\omega = 0.965$ is almost half of the natural frequency of the system. When the relative

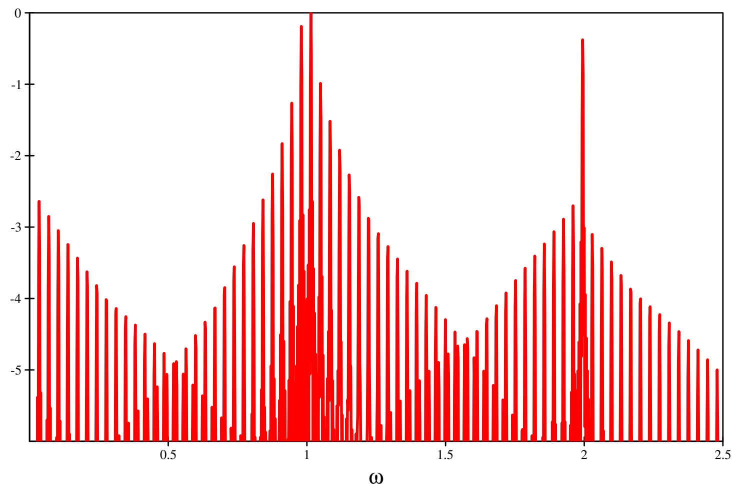

Figure 2.8: *Power spectrum of quasiperiodic behavior* shown on a log (base 10) scale. System driven at frequency $\omega = 1.9945$ with amplitude $A = 1$, $\alpha = 1$, $\tau = 2$, and $\sigma = 0.3$ for competitive payoffs $G_1 = 7 - f_1$ and $G_2 = 7 - 3f_2$. Natural frequency of the undriven system is 1.1400. Spectrum was obtained from a time series obtained by integrating Eq. 2.1 over $t = 0$ to 10,000, neglecting initial transients.

payoffs are modulated at this frequency, the natural frequency locks onto the second harmonic of the driving frequency and a subharmonic at $\omega/3$ appears.

A more common type of behavior is quasiperiodic oscillation resulting from the mixing of two irrationally related frequencies close in magnitude and in strength. Specifically, the system's behavior is governed by both the natural frequency of the undriven system and the driving frequency, which in general are not rationally related. Although in most cases the system will lock onto one or the other of the two frequencies, for some choices of the amplitude A, both frequencies become important and quasiperiodicity results. Fig. 2.8 shows the power spectrum for a quasiperiodic time series. Small shifts in the driving frequency lead to small shifts in the ratio of the two strongest frequencies near one, strongly indicating that the behavior is indeed quasiperiodic. Since it is impossible to distinguish numerically between a quasiperiodic series and a periodic series, we cannot make a more exact statement.

Scattered through the parameter space of interest are windows where clean subharmonics of the driving frequency are seen. We have observed period-n of the driving term for n from 1 to 8 as any one of the parameters is varied and deduce the presence of even higher periods. For example, for fixed values of the amplitude, frequency, α, and σ ($A = 1$, $\omega = 3$, $\sigma = 0.3$, and

$\alpha = 2.0\alpha$=2.0), the system passes successively through periods 1, 2, 3, 4, 5, 6 as τ is increased from zero to six. Clean period-*n*'s, however, are observed only at certain values of τ; at these values, the natural frequency of system is very close to ω/n, and the system successfully locks onto the *n*th subharmonic when driven. Finally, we add that for this choice of payoffs with two resources and a set of identical agents, no chaotic behavior (as characterized by broadband spectra with peaks at the driving frequency and its harmonics) is observed, although it may still be lurking in some obscure unexplored region of the parameter space.

2.4.2 Transient tracking

We now study the tracking ability of a computational ecosystem with the cooperative payoffs of Eq. 2.4, a highly non-linear system. Specifically, we examine the system for those parameters that yield chaotic behavior when the environment is fixed in time.

In general, numerical integration shows that utility modulation at frequencies on the order of the reevaluation rate α quells the naturally chaotic tendencies of the system. Further, the behavior of the agents exhibits a strong frequency component at the driving frequency (as usual, the stronger we drive, the larger the effect). We hypothesize, then, that a system whose natural tendency is to behave chaotically can on average track the time-dependence in the relative payoffs at least as well as a system which, undriven, oscillates at a single or even several frequencies (barring the unlikely possibility that the driving frequency happens to correspond exactly to the natural frequency).

However, in addition to the behavior described above, we find that when the system is driven at frequencies greater than α, the transient behavior of the system displays superior tracking. Fig. 2.9 demonstrates the transient stabilizing effect of driving at frequency three with amplitude one for choices of σ, α, and τ which lead to chaotic behavior in the undriven system. Notice that there are two types of transients. The first, lasting for about 25 time units, is chaotic and begins to disappear once the system notices the time-dependent modulation (once a period of one delay, $\tau = 10$, units elapses). The second transient lasts longer than 50 time units (5 delay periods) and shows low-amplitude oscillation primarily at the driving frequency, $\omega = 3$. After about 100 time units, the system settles into a higher-amplitude periodic oscillation whose main

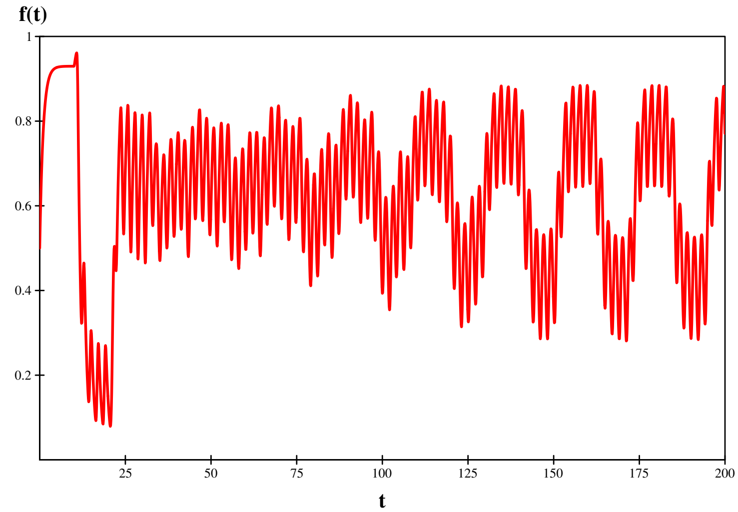

Figure 2.9: *Transient adaptive behavior* of a system of cooperative agents with the payoffs $G_1 = 4 + 7f_1 - 5.333f_1^2$ and $G_2 = 7 - 3f_2$ driven at frequency $\omega = 3$, amplitude $A = 1$, for parameters $\alpha = 1$, $\tau = 10$, and $\sigma = 0.32$. The undriven system behaves chaotically for these parameters.

frequency components are $\omega = 0.2876$ and $\omega = 3$. We will discuss this behavior further in the following section in the context of adaptability.

Lastly, these results also indicate that one straightforward way to control chaos might be to simply incorporate a sinusoidally fluctuating background component into the payoffs of some of the resources. Even low amplitude modulation of the resource utilities results in behavior which is periodic with a sharply peaked power spectrum. Consequently, in cases in which periodic behavior is judged superior to chaotic, periodic modulation of the utility functions is a viable method for inhibiting chaotic behavior. However, as discussed in the next section, a system whose behavior fluctuates periodically may be less adaptable to a changing environment than one which behaves chaotically, and thus could be considered inferior.

2.5 Issues of Adaptability

A natural question to ask, in light of the inevitability of unpredictable changes in the environment, is what type of computational ecosystem is most adaptive. A system which performs well in the absence of change is not necessarily the best choice if the environment is expected to vary unpredictably with time. Observing the altered behavior of ecosystems in the presence of sinusoidal time-dependent resource payoffs may provide insight into the question of adaptability.

If we can define an acceptable quantitative measure of adaptability and apply it to determine the response of the system to a sinusoidal signal, we can then attempt to state which characteristics of an ecosystem are indicative of adaptability.

The measure of adaptability we choose is also the simplest: the performance of the driven system averaged over time, where the performance is defined as the sum over resources of the resource payoff times the fraction of agents using the resource,

$$Performance = \sum_i G_i f_i \ . \tag{2.10}$$

Note that this is also the average payoff received per agent.

The performance of a driven system reflects how well the agents are able to track the driving frequency since it is maximized when the behavior of the system fluctuates at the same frequency as the driving term and with zero phase shift. We can easily find the "optimal" behavior of the system ("optimal" defined circularly by the measure of performance) by calculating the form of $f(t)$ which maximizes the performance given a particular set of resources, payoffs, and agents. Consider the competitive payoffs of Section 2.3. Setting $d(performance)/df_1 = 0$, we find that optimal behavior occurs when

$$\begin{aligned} f_1(t) &= 0.75 + \frac{A}{8}\cos(\omega t) \\ f_2(t) &= 1 - f_1(t). \end{aligned} \tag{2.11}$$

Similarly, we derive the optimal behavior of the system of cooperating agents presented in the previous section:

$$\begin{aligned} f_1(t) &= \frac{1}{4}\left[1 + \sqrt{4 + A\sin(\omega t)}\right] \\ f_2(t) &= 1 - f_1(t). \end{aligned} \tag{2.12}$$

The above equations only hold for *A* small enough such that f_1 and f_2 take on appropriate values constrained to lie in the interval [0,1]. Larger values of *A* effectively correspond to the limit of large amplitude. We will confine ourselves to intermediate amplitudes in our quest for the meaning of adaptability since these are both most interesting and most relevant.

Now that we have a means of measuring performance and a definition of optimal performance, we can turn to specific systems and analyze their responses to sinusoidal signals. Our motivation lies in the possibility that a system in the unstable region of Fig. 2.1 may be more adaptable than one in the stable region. If the ecosystem is in a static environment, then it has been previously shown in (Kephart *et. al.* 1989) that the optimal choice of parameters locates the system on the stability boundary. Specifically, if the system is in an unstable regime, the most straightforward and viable way to simultaneously stabilize the system and maximize its performance is to increase the uncertainty or decrease the reevaluation rate until the system just becomes stable, given the constraints of a minimal intrinsic delay and uncertainty.

(a) **(b)**

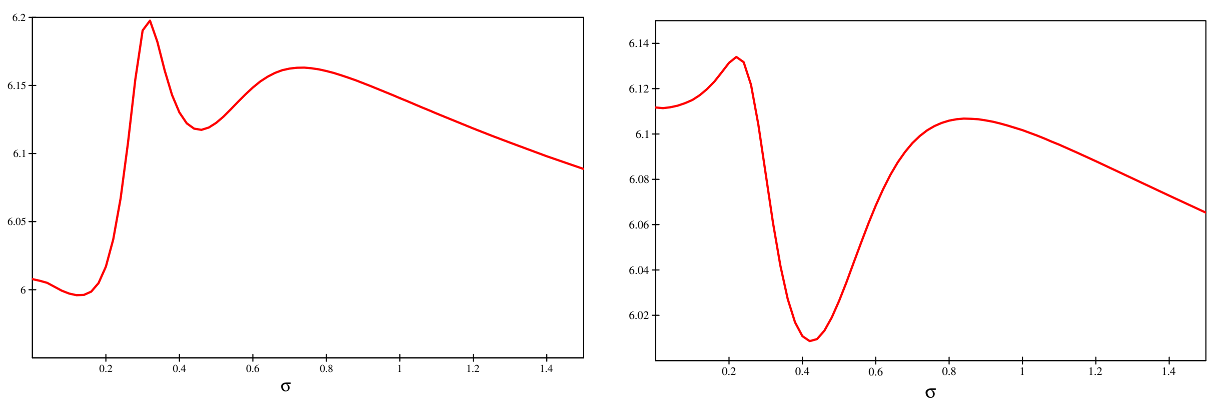

Figure 2.10: *Adaptability.* Performance of a system of cooperative agents (payoffs given by $G_1 = 4 + 7f_1 - 5.333f_1^2$ and $G_2 = 7 - 3f_2$) as a function of σ. In (a) the system is driven at frequency $\omega = 3$, amplitude $A = 1$, for parameters $\alpha = 1$, $\tau = 10$, and $\sigma = 0.32$. The performance is averaged over time $t = 25$ to 75. In (b) the system is driven at frequency $\omega = 4$, amplitude $A = 1$, for parameters $\alpha = 1$, $\tau = 15$, and $\sigma = 0.22$. The performance is averaged over $t = 50$ to 100. In both, the point of optimal performance corresponds to chaotic behavior in the undriven system.

We find that once the system interacts with the outside world through a time-dependent modulation of the relative payoffs, the boundary between the two behavioral regimes is no longer always the best place to be. On the one hand, if the parameters α, τ, and σ are chosen to lie in the oscillatory regime of the undriven system, then increasing the uncertainty σ will indeed lead to increased performance. Agents oscillating at one frequency adapt worse to an external driving term than do their stable counterparts working with increased uncertainty. On the other, if the ecosystem behaves chaotically in the absence of driving, it is possible for the system to perform better than its counterpart on (or above) the stability boundary, but only transiently. Figs. 2.10(a) and (b) depict two examples of cases in which the system responds better to the changing environment (as quantified by higher performance) when starting from a chaotic state than when starting from either an oscillating or stable state, arrived at by varying σ. The points of maximal performance correspond to undriven chaotic behavior. Specifically, the point of highest performance in Fig. 2.10(a) corresponds to the time series shown in Fig. 2.9. If the performance were to be averaged over longer time scales, higher σ would then yield superior performance.

However, a second criterion for adaptability, in addition to good performance, is fast response to a changing environment. Originally, we introduced a sinusoidal modulation in the payoffs as a simple model of external time-variations. Real environments rarely follow such a simple pattern of variation. Thus, the important measure of adaptability is how the system responds to new types of variations. In the case of the periodic modulations we have studied, this measure can be approximated by averaging over short time scales. We are not interested in equilibrium behavior for long times in the periodic case, but rather in the immediate reaction of the system to external influences.

Improved response in the chaotic region of the undriven system is attributed to the wide range of frequencies associated with a chaotic time series. Note, though, that the frequency components of the chaotic behavior observed in the undriven system of cooperative agents, while broad-band, are dominated by a few frequencies, and the enhancement of performance is not as large as it might be for other chaotic systems.

2.6 Summary

In this chapter, we have applied the Huberman-Hogg theory of computational ecosystems to the study of systems with time-modulated resource utilities. We have used linear stability analysis to demonstrate that, in the limits of low amplitude and low frequency modulation, the stable regimes of the driven ecosystem correspond to those of the original one in a fixed environment. In particular, in the stable regime the system oscillates at the periodicity of the payoff utilities. For low amplitudes, the system's response lags the modulation by a phase lag depending on the delay, the reevaluation rate, and the amplitude of the modulation. For low frequency driving, the phase lag is negligible.

We also found an exact solution for the case of high-amplitudes; namely, the system exhibits step-function–like behavior, with the agents first flocking to one resource and then to the other, as their preference switches back and forth in tandem with the sign of the sinusoidal payoff utilities. Since the information used by the agents to calculate the payoffs is delayed, the transition between the two states is also delayed by the same amount.

The high frequency limit was also discussed, and it was argued that the effect of modulating the payoff utilities at frequencies high in comparison with the reevaluation rate was the same as increasing the uncertainty, as long as the number of agents is large enough.

Knowledge of the behavior of the ecosystem in these limits was helpful in interpreting the results obtained by numerical integration of Eq. 2.1 for the more relevant cases of intermediate amplitude and intermediate frequency modulation of the utility functions. We explored the parameter regions, as specified by the values of α, σ, and τ, which corresponded to the unstable regimes of Fig. 2.1 for the original undriven ecosystems studied in previous work.

In the case where the original system exhibits oscillatory behavior, three types of responses ensue in response to modulation of the resource utilities. At low amplitudes, the modulation has no effect on the behavior of the system, and the agents continue to oscillate at their original frequency. At high amplitudes, the system is forced back and forth between the two resources by the external modulation, as described above. However, at intermediate amplitudes (*i.e.*, when the amplitude of the modulation is comparable to the strength of system's natural oscillation), we found a rich set of dynamical behavior, including oscillation at subharmonics of the modulation frequency and aperiodicity. In these last cases, the deterioration in the agents' ability to track the modulation in the payoff utility was most dramatic.

We also observed that external modulation quells the chaotic tendencies of computational ecosystems, and concluded that such environmental effects might be one method to control chaos. We further noted the transient effects of modulating the payoffs of a chaotic system. After an initial adjustment to the external influence lasting for a few delay periods, these transients displayed superior tracking in comparison with the system's long-term equilibrium behavior.

Finally, we considered the question of adaptability by measuring the average performance of the system over short time intervals. By an adaptable system we mean, in this context, one that performs well in response to a changing environment. We found that in some cases periodic modulation of the payoffs changes the parameter values which display the most adaptive behavior. Specifically, chaotic systems are sometimes more adaptive, by our definition, than stable ones.

These results, which were obtained in the context of computational ecosystems, might have some relevance to other distributed forms of organization as well, such as biological ecosystems, economies and social structures. In order for our results to be applicable to such systems, many factors which were excluded from this work will have to be incorporated, such as actual payoffs, many more strategies, discounted future values for the resource utilities, transaction costs, and the dynamics of expectations. Such an effort will be attempted in the following chapter.

3 Expectations and Resource Contention

3.1 Introduction

In most economic and social systems, expectations about the future, along with imperfect knowledge about the present and the past, determine individual choices. Stock trading, currency markets, and the planning of production cycles provide a few examples of such behavior. Social agents form expectations because of their intentional nature, a property which is absent in the objects that are the purview of physics and chemistry. Consequently, a dynamical formulation of social interactions requires a different approach than traditionally taken, one in which the future enters explicitly in the equations describing the evolution of the system.

In this chapter, we extend the theory of computational ecosystems in several ways in order to account for agents that form expectations and have beliefs concerning the future. Most importantly, the agents take into account their expectations about the future when making decisions in the present. As a result, the perceived payoff for using a resource depends on how an agent expects the global behavior of the system to evolve as well as on the resource's current perceived load. We consider two different types of expectations in the course of the chapter. Both are deliberately simplistic since our main concern is not to study how agents might predict the future, but rather to learn how their expectations affect the system's overall dynamical behavior.

Two scenarios of expectation formation will be studied. The first, we call *flat expectations*. In this scenario, each agent assumes that the system remains essentially constant; thus an agent's prediction for the future state of the system is simply the present state. Since the information is delayed, the present state is actually the state of the system some time in the past. The second will go by the name *trendy expectations* or linear extrapolation. In this case, each agent expects the system to follow its present trend. Again, because the agent's information is delayed, the agent's decisions are actually based on the trend perceived at some delayed time in the past.

We emphasize that we consider only the simplest types of expectations. Our concept of flat expectations harks back to the static or "naive" hypotheses of expectation formation tacitly assumed in classical economic theory (Shaw 1984, Carter 1984). While our focus on the effect of expectations on dynamics follows recent research in economics on the role of psychological forces in market movements (Boldrin 1986, Smith 1988, Day 1990), we turn away from the approach to expectations in economics adopted in past decade. This approach postulates that agents have rational expectations; namely, that agents form expectations about the future using knowledge of the model governing the interaction. Problems studied using rational expectations are most tractable when models are linear (Blanchard 1989) or when the agents are assumed to have perfect foresight, see for example (Grandmont 1986b, Azariadis 1981). However, the assumption that agents have accurate knowledge of the underlying model is empirically not well-founded (Shaw 1984, Carter 1984), and the presumption of self-fulfilling prophecies raises difficult problems in expectations coordination (Grandmont 1986b). For these reasons, we eschew rational expectations.

We also extend the computational ecosystem model to include the notion of transaction costs. These are particularly relevant when the cost of moving between resources becomes high compared to the instantaneous payoff. Only when the accumulated utility becomes large enough to justify the transaction cost will agents change strategy. Besides transaction costs and future expectations, we also include discounts as an essential variable in our study. Discounts simply mean that a dollar now is worth more than a dollar in the future. Since an agent could place its money in the equivalent of a bank and let it grow at interest rates, economic rationality dictates that it discount future payoffs with respect to immediate payoff.

The next section describes how we model expectations in the framework of computational ecosystems. In Sections 3.3 and 3.4, analytical and computational results are presented for flat and trendy expectations, respectively. In particular, in Section 3.3 we show how long forecasts tend to destabilize the system and obtain a phase diagram that exhibits a sharp boundary separating stable equilibria from complex dynamics as the horizon length increases. We also demonstrate the existence of an optimal strategy for flat expectations, in the sense that profits are maximized and behavior is stable. Expectations that rise over time are seen to lead to a crash in which the fixed point gives way to wild oscillations. In the case of trendy agents discussed in Section 3.4, we obtain qualitatively similar behavior to that of flat predictions, but with a more pronounced tendency for expectations to destabilize the system.

The more striking behavior is described in Section 3.5. There we show how a diversity of expectations among the agents, coupled to reward mechanisms in which those agents that do well increase at the expense of others, leads to overall dynamics characterized by cycles of almost stable behavior interrupted by sudden crashes. This process is accompanied by an ever-changing diversity in the composition of the system.

3.2 Modeling Expectations

In this section we extend the computational ecosystem model to include expectations about the future and transaction costs. The external environment is now assumed to be fixed in time. Later, in section 3.5, we will further extend the model to include reward mechanisms in which those agents that do well increase at the expense of inferior performers. The general exposition presented below is valid for *N* agents choosing among *R* resources. Subsequently, we specialize to a 2–resource ecosystem.

3.2.1 Formulation

In evaluating the perceived expected payoff for using each resource, agents take into account future anticipated earnings as well as present perceived earnings. The agents discount future earnings expected at a time s units from the present time t at a rate $e^{-s/H}$, reflecting their bias towards the present, where *H* is the agents' horizon length and is constant in time. The instantaneous density-dependent payoff is defined to be $G(\mathbf{f}(t))$. $\mathbf{f}(t)$ is a vector describing the fraction of agents using each resource at time *t* and has scalar components $f_i(t)$ corresponding to the fraction of agents using resource *i*. Now let $\mathbf{f}^*(t+s)$ be a vector describing the expected behavior of the system at time s units from the present. (The components $f_i^*(t+s)$ represent the expected fraction of agents utilizing resource *i* at time *s* units from the present.) This vector function $\mathbf{f}^*(t+s)$ represents the mathematical expression of the way in which agents form expectations about the future. In principle it could depend on the system's entire past history as well as additional information made available to the agents. Under the simplifying assumptions of this chapter, $\mathbf{f}^*(t+s)$ will depend only on the delayed values $\mathbf{f}(t-\tau)$ and $\mathbf{f}'(t-\tau)$, the vector with components $f_i'(t-\tau)$ (the time rate of change of the fraction of agents using resource *i* at the delayed time $t-\tau$).

Each agent expects the instantaneous payoff density for resource *i* at a time *s* units from the present to be $G_i(\mathbf{f}^*(t+s))$. The agent's anticipated payoff, $g_i(\mathbf{f}^*, t)$, at time *t* for the *i*th resource is then the discounted time integral of the expected instantaneous payoff density:

$$g_i(\mathbf{f}^*, t) = \int_0^\infty G_i(\mathbf{f}^*(t+s))\exp\left[-\frac{s}{H}\right] ds. \tag{3.1}$$

At this point, we examine the special case of a system with two resources. The scalar function $f(t)$ is the fraction of agents using resource 1 at time *t* and $1 - f(t)$ is the fraction using resource 2. Similarly $f'(t)$ is the time rate of change of the fraction of agents using resource 1 and $1 - f'(t)$ is the time rate of change of the fraction of agents using resource 2.

3.2.2 Flat expectations

Flat expectations are used to model agents who believe the system to be changing very slowly. In this case, the expected future behavior of the system is simply the past observed behavior, or

$$f^*(t+s) = f(t-\tau) \tag{3.2}$$

for all $s \geq 0$, where τ is the delay. The expected payoff can be evaluated trivially since the instantaneous payoff at all future times is constant:

$$g_i^{flat}(f(t-\tau)) = HG_i(f(t-\tau)). \tag{3.3}$$

Here, the explicit time-dependence has dropped out since, the expected payoff depends on time only implicitly through $f(t-\tau)$.

Alternatively, agents with flat expectations might expect future behavior to depend in a more complicated fashion on the past history of the system. For example, the expected future behavior could be a weighted sum of the resource loads in the past: $f^*(t+t') = \sum\limits_n w_n f(t-n\tau)/\sum\limits_n w_n$, which we refer to as weighted flat expectations. In particular, this framework includes adaptive expectations in which agents subsequently revise their expectations according to the degree of error in their previous expectations. Specializing to the case of flat expectations results in little loss of generality: the equilibrium points remain the same for all types of weighted flat expectations, while the stability portrait of the system (although in practice much more difficult to compute) remains qualitatively similar. This should become clearer later.

3.2.3 Trendy expectations

The second type of expectation reflects the agents' belief that the system will continue to follow the observed trend. As a result, their expectations depend both on the state of the system $f(t-\tau)$ and its trend, or slope, $f'(t-\tau)$, the time rate of change of the state. Agents assume that the system will continue to follow the trend observed at the delayed time. They extrapolate to the present and then to the future (we assume agents are aware of the time lag of their information), allowing their predictions to decay asymptotically to the boundaries at $f=0$ and $f=1$ for the sake of continuity. We can then write for the expected evolution of *f*,

$$f^*(t+s) = \theta - (\theta - f(t-\tau))\exp\left[-\frac{f'(t-\tau)}{\theta - f(t-\tau)}(s+\tau)\right], \tag{3.4}$$

where

$$\theta = \begin{cases} 1, & f'(t-\tau) > 0, \\ 0, & f'(t-\tau) \le 0. \end{cases} \tag{3.5}$$

Specifically, the agents' expectations are such that they match actual behavior at time $t-\tau$, that is $f^*(t-\tau) = f(t-\tau)$, $f^{*\prime}(t-\tau) = f'(t-\tau)$, and for small $\tau+s$, the extrapolation is linear, that is $f^*(t+s) = f(t-\tau) + (s+\tau)\, f'(t-\tau)$. Note that when $f'(t-\tau) = 0$, $f^*(t+s) = f(t-\tau)$ for all future times s, reducing to the case of flat expectations.

The payoff expected by trend-followers depends on the functional form of the instantaneous resource payoffs. Taking the case of quadratic payoffs (*i.e.*, payoffs which have both a cooperative and a competitive component),

$$G_i(f(t)) = a_i + b_i f(t) + c_i f^2(t), \tag{3.6}$$

the agent's total accumulated payoff then becomes

$$g_i^{trend}\big(f(t-\tau), f'(t-\tau)\big) =$$

$$H\Bigg\{[a_i+\theta\,(b_i+c_i)] - \frac{(b_i+2\theta c_i)(\theta-f(t-\tau))^2}{Hf'(t-\tau)+\theta-f(t-\tau)}\exp\left[-\frac{f'(t-\tau)}{\theta-f(t-\tau)}\tau\right] \\ +\frac{c_i(\theta-f(t-\tau))^3}{2Hf'(t-\tau)+\theta-f(t-\tau)}\exp\left[-\frac{2f'(t-\tau)}{\theta-f(t-\tau)}\tau\right]\Bigg\}. \tag{3.7}$$

This equation yields the result that when the slope is zero, *i.e.*, $f'(t-\tau) = 0$, then $g_i^{trend}(f(t-\tau)) = HG_i(f(t-\tau))$, again reducing to the case of flat expectations.

3.2.4 Transaction costs

Transaction costs enter whenever an agent evaluates the difference in payoff between two resources. Since an agent switching resources incurs a transaction cost, the benefit of switching must overcome this cost. Thus, in the absence of uncertainty, an agent will move to another resource j only when

$$g_j(f^*,t) - g_i(f^*,t) > T, \tag{3.8}$$

T being the transaction cost. In general, however, the agents' knowledge will be imperfect. This uncertainty is incorporated into the computational ecosystem model through the parameter σ, which characterizes the average error in the agents' perceptions of the instantaneous resource payoffs G_i. The uncertainty in the expected payoff, g_i, for using resource *i* is $H\sigma$, since no new information is gained by projecting into the future. The probability that an agent will choose to switch between two resources is then given by

$$\rho_{i\to j}(f^*,t) = \frac{1}{2}\left\{1+\operatorname{erf}\left[\frac{g_j(f^*,t)-g_i(f^*,t)-T}{2H\sigma}\right]\right\},\ i\neq j, \tag{3.9}$$

while the probability that it does not switch is simply

$$\rho_{i\to i}(f^*,t) = 1-\rho_{i\to j}(f^*,t) = \frac{1}{2}\left\{1-\operatorname{erf}\left[\frac{g_j(f^*,t)-g_i(f^*,t)-T}{2H\sigma}\right]\right\}. \tag{3.10}$$

For both flat and trendy expectations, $\rho_{i\to j}(f^*,t)$ will have no explicit time dependence since the anticipated payoffs $g_i(f^*,t)$ depend on time only implicitly through the delayed values $f(t-\tau)$ and $f'(t-\tau)$. In the case of zero uncertainty ($\sigma = 0$), $\rho(f^*,t)$ reduces to a step function:

$$\rho_{i\to j}(f^*,t) = \begin{cases} 1, & g_j(f^*,t) - T > g_i(f^*,t), \\ 0, & g_j(f^*,t) - T < g_i(f^*,t). \end{cases} \tag{3.11}$$

Note that transaction costs adds asymmetry to the dynamics of the system since $\rho_{i\to j}(t) \neq 1 - \rho_{j\to i}(t)$.

3.2.5 Dynamics

Finally, the dynamics of a system of identical agents choosing between two resources is described by the dynamical equation

$$\frac{df(t)}{dt} = -\alpha[f(t) - f(t)\rho_{1\to 1}(f^*,t) - (1-f(t))\rho_{2\to 1}(f^*,t)], \tag{3.12}$$

where $f(t)$ is the fraction of agents using resource 1, α is the rate at which agents reevaluate their preferences, $\rho_{i\to j}(f^*,t)$ is given by Eq. 3.9, and $g_i(f^*,t)$, the accumulated expected payoff is given by $g_i^{flat}(f(t-\tau))$ of Eq. 3.3 for the case of flat expectations and by $g_i^{trend}(f(t-\tau), f'(t-\tau))$ of Eq. 3.7 for the case of trend-followers. In a later section when rewards are introduced we will consider a more complex system made up of many different types of agents, but still limited to two different kinds of resources or strategies. At that point we will need to generalize Eq. 3.12.

3.3 Flat Predictions

This section focuses on the scenario of flat expectations. We will use linear stability analysis and computational methods to map out the various behavioral regimes of a computational ecosystem with the cooperative payoffs of the previous chapter. Further, we will numerically establish the optimal value of horizon length for fixed transaction cost, given the chosen performance metric. Finally, an example of the dynamics of a system made up of agents whose expectations change in time is presented.

3.3.1 The stability portrait and its behavioral regimes

Transaction costs will tend to stabilize the system, forcing it into an equilibrium which might not be optimal. This observation raises some questions: How far into the future can agents look given a fixed transaction cost without creating instabilities? How does the maximum stable horizon length scale with the agents' uncertainty and the transaction costs?[3]

For a given system characterized by a set of payoffs, reevaluation rate, delay, and uncertainty, we can map out stability boundaries in the space σ versus H/T (H = horizon length, T = transaction cost). Note from Eqs. 3.3 and 3.9 that the stability boundaries in the case of flat expectations will depend only the ratio $H : T$. This implies that at fixed σ the maximum stable horizon will scale linearly with increasing transaction costs. We expect that the functional form of the boundary has the same qualitative features when considering other two-resource computational ecosystems with flat expectations.

Linear stability analysis of the dynamical equation (Eq. 3.12) gives an implicit solution to the functional form of the stability boundary. We linearize Eq. 3.12 in the neighborhood of a fixed point f_0, letting $\delta(t) = f(t) - f_0$. Then, with $\gamma_{1\to1} = \frac{d\rho_{1\to1}}{df}(f)|_{f=f_0}$ and $\gamma_{2\to1} = \frac{d\rho_{2\to1}}{df}(f)|_{f=f_0}$, we obtain

$$\begin{aligned} \frac{d\delta}{dt} &= \alpha\{\delta(t)[\rho_{1\to1}(f_0) - \rho_{2\to1}(f_0) - 1] \\ &+\delta(t-\tau)f_0[\gamma_{1\to1} - \gamma_{2\to1}] + \delta(t-\tau)\gamma_{2\to1}\}. \end{aligned} \tag{3.13}$$

The behavior of the solutions to this equation can be characterized by making the substitution $\delta(t) = \exp(-\alpha s t)$, giving the following equation in the complex variable s:

$$[f_0(\gamma_{1\to1} - \gamma_{2\to1}) + \gamma_{2\to1}]\mathrm{e}^{-\beta s} + \rho_{1\to1}(f_0) - \rho_{2\to1}(f_0) - s - 1 = 0, \tag{3.14}$$

where $\beta \equiv \alpha\tau$. There are an infinite number of roots to the above exponential equation; the system is stable to perturbations about the fixed point only if the largest real part of all roots is less than the zero (Bellman and Cooke 1963).

[3] In practice, of course, there may be a limit on the longest allowable horizon length (due to lifetime or finite patience of the agent/resource, for example).

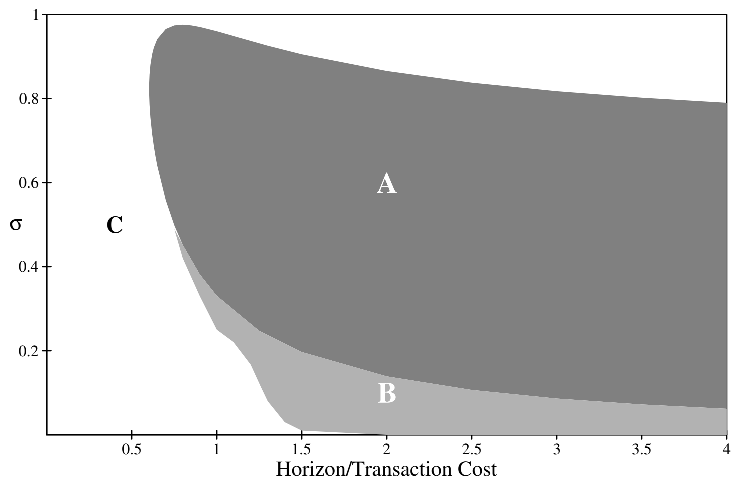

Fig. 3.1: *Stability portrait* in the phase space H/T vs. σ for system with resource payoffs $G_1 = 4 + 7f - 5.333f^2$ and $G_2 = 4 + 3f$, $\alpha = 1$, $\tau = 6$. System is always unstable in region A and always stable in region C. In region B, system either relaxes to a fixed point or is pulled into a limit cycle, depending upon the initial conditions. The boundary to region A was derived analytically, while the lower boundary to region B was found by repeated numerical integration of the dynamical equation.

In the limit $\mathrm{Re}(s) \to 0^-$, an analytical expression can then be derived from Eq. 3.14 which gives the stability boundary $\sigma(H/T)$ implicitly in terms of the accumulated payoffs g_i that enter the ρ functions through Eqs. 3.3 and 3.9:

$$\beta = \left\{ [f_0(\gamma_{1\to 1} - \gamma_{2\to 1}) + \gamma_{2\to 1}]^2 - [1 - \rho_{1\to 1}(f_0) + \rho_{2\to 1}(f_0)]^2 \right\}^{-\frac{1}{2}} * \cos^{-1}\left[\frac{1 - \rho_{1\to 1}(f_0) + \rho_{2\to 1}(f_0)}{f_0(\gamma_{1\to 1} - \gamma_{2\to 1}) + \gamma_{2\to 1}}\right]. \tag{3.15}$$

In order to solve for $\sigma(H/T)$ we also need the fixed point f_0, given by the solution to

$$f_0\rho_{1\to 1}(f_0) + (1 - f_0)\rho_{2\to 1}(f_0) = f_0. \tag{3.16}$$

Solving Eqs. 3.15 and 3.16 simultaneously for σ and f_0 gives the limiting value of the uncertainty at which $\mathrm{Re}(s) \to 0^-$ for fixed H/T and the system becomes unstable. We find that when the uncertainty is large enough, the system is stable for all values of H/T, and that for increasing H/T very small amounts of uncertainty cause instability. The results for the choice of resource payoffs, $G_1 = 4 + 7f - 5.333f^2$, $G_2 = 4 + 3f$, $\alpha = 1$, and $\tau = 6$ are shown in Fig. 3.1. As can be seen, the system is unstable for all choices of initial conditions within region A, delineated by the curve $\sigma(H/T)$ derived above.

This curve delineates the region in phase space for which the system is definitively unstable, but we have yet to investigate the surrounding region. The remaining boundary was computed by repeated numerical integration of Eq. 3.12. In region A, all fixed points of the system are repellors $(\mathrm{Re}(s) > 0)$. Below, in region B, some fixed points are attractors, some are repellors. That there exist multiple fixed points for some pairs of horizon and transaction cost can be seen from the graphical solution of the dynamical equations. For some initial conditions, the system relaxes to a fixed point. For others the system is pulled into a limit cycle. As a result, the behavior of the system depends strongly on the initial conditions. Fig. 3.2 demonstrates this dependence. The same system relaxes to a fixed point for one set of initial conditions but gets pulled into a limit cycle for another. We further add that although the system can be chaotic for large enough values of H/T in region A, within region B the uncertainty is small enough that no chaotic behavior is observed.

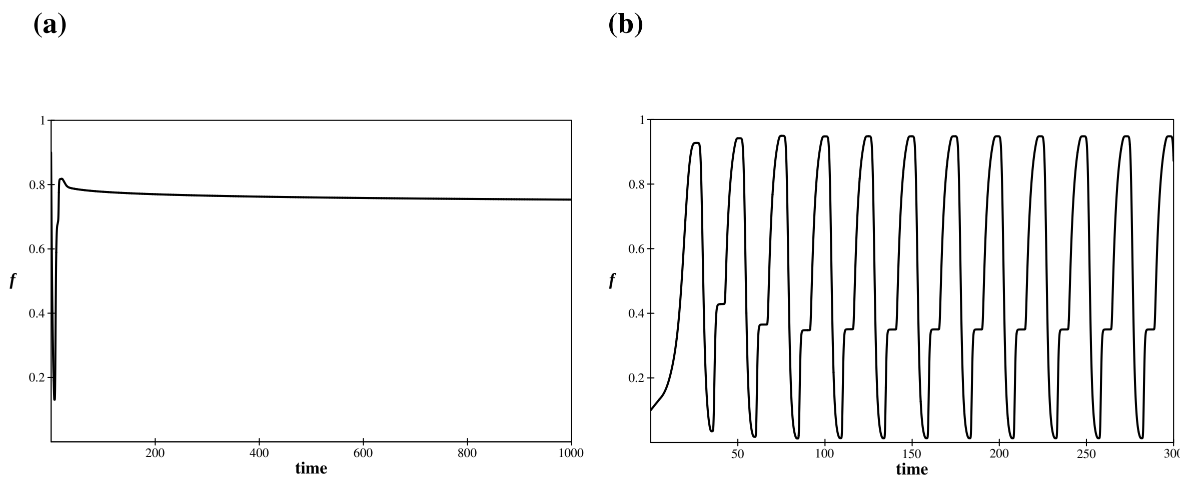

Figure 3.2: *Dependence on initial conditions.* Resource payoffs are given by $G_1 = 4{+}7f{-}5.333f^2$ and $G_2 = 4{+}3f$. $H/T = 1.2333$, $\alpha = 1$, $\tau = 6$ and $\sigma = 0.135135$. System falls into region B of phase diagram of Fig. 3.1, where different initial conditions can lead to different types of behavior. (a) Initial condition $f = 0.9$ — system relaxes very slowly to a fixed point at $f_o = 0.74424$ (after over 8000 time steps); (b) Initial condition $f - 0.1$ — system resource use oscillates. Apparent sharpness of corners is due to the transaction costs.

Region C of Fig. 3.1 gives the values of σ vs. H/T cost for which the system always relaxes to a fixed point. However, it can take a prohibitively long time to do so, depending on the initial conditions — see, for example, Fig. 3.3. In this example, the system is trapped for a long time at the almost-fixed point $f = 0$. The slow rate of escape is caused by the barrier to switching erected by the transaction costs. For vanishing transaction costs, the rate of transitions between resources

$$\alpha(f\rho_{1\to2} + (1-f)\rho_{2\to1}) \tag{3.17}$$

is on the order of the reevaluation rate α. However, for transaction costs such that $T/H \gg |(\mathrm{G}_1(f) - \mathrm{G}_2(f))|$,

$$\rho_{1\to2} \approx \rho_{2\to1} \propto \exp\left(-T^2/4H^2\sigma^2\right), \tag{3.18}$$

Consequently the transition rate is also reduced from α by a factor proportional to $\exp\left(-T^2/4H^2\sigma^2\right)$. In the example shown, $G_1 = G_2$ at $f = 0$ and thus the transaction costs trap the system in a suboptimal state for a length of time roughly exponential in the square of the transaction costs.[4] Also note that if the uncertainty vanishes, the system will be trapped forever in the suboptimal metastable state $f = 0$. This is a manifestation of the persistence of non-optimal strategies as discussed in (Ceccato and Huberman 1989).

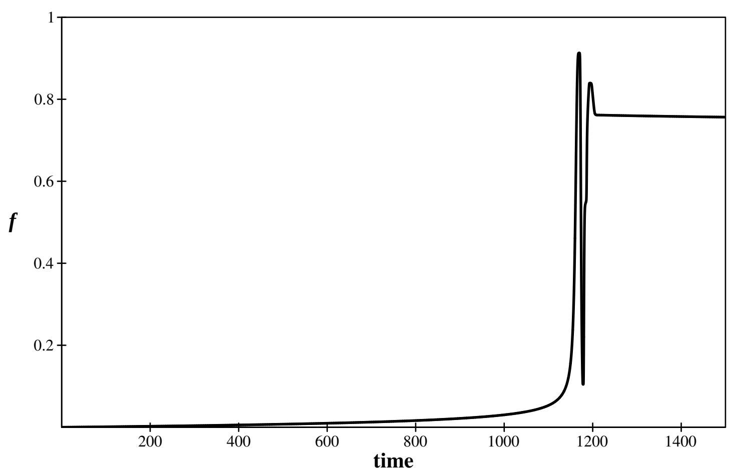

Figure 3.3: *Persistence of non-optimal strategies.* Resource payoffs are given by $G_1 = 4 + 7f - 5.333f^2$ and $G_2 = 4 + 3f$. $H/T = 1.2$, $\alpha = 1$, $\tau = 6$ and $\sigma = 0.13889$. System in region C of *H-T* phase diagram of Fig. 3.1. Initial condition $f = 0$. System takes a very long time to escape the almost fixed point $f = 0$ and overcome the cost barrier erected by the transaction costs.

4 The state $f = 0$ is suboptimal since the value of the instantaneous payoff is far from its maximum value.

As mentioned earlier, the stability boundaries mapped out in Fig. 3.1 depend only on the ratio $H : T$; thus, from this figure we know the linear stability boundaries $H(T)$ for different values of the uncertainty σ. Specifically, when σ is small, the slope of $H(T)$ becomes very steep and only very large horizons destabilize the system at a fixed transaction cost. Consequently, when σ vanishes, infinitesimally small values of T are sufficient for the fixed point to be stable at all values of the horizon length H. This observation can be understood as follows: agents with perfect information will resist any movement away from the fixed point when there are transaction costs. Therefore, the combination of nonzero transaction costs with perfect information effectively erects an insurmountable transition barrier. Note, however, that the fixed point is still unstable when there are no transaction costs ($H/T \to \infty$). Lastly, we point out the anomalous behavior seen for $\sigma \sim 0.9$: as H/T increases the system passes successively through a stable regime, an unstable regime, and then a second unstable regime.

3.3.2 Optimal horizons

For a fixed transaction cost, we have shown that there is a range of possible horizon lengths that lead to stable behavior, defined as relaxation to a fixed point. However, the location of the fixed point varies with horizon length. We now define the optimal horizon length as being the horizon length for which a system yields the highest performance under time and stability constraints. The choice of performance measure is somewhat arbitrary; we choose to equate performance of the system with the average earnings per agent over time in analogy to a market economy. There are two components to the average amount earned per agent at any given time: the average instantaneous payoff received per agent and the mean transaction cost incurred per agent. The average system payoff is simply the sum of the actual payoffs accrued by the entire system per unit time, *i.e.*,

$$\mathrm{Payoff}(t) = f(t)G_1(f(t)) + (1 - f(t))G_2(f(t)). \tag{3.19}$$

The mean transaction cost per unit time at time t is the fraction of agents switching between resources per unit time multiplied by the transaction cost, or

$$\mathrm{Cost}(t) = \alpha T[f(t)\rho_{1\to 2}(f(t-\tau)) + (1 - f(t))\rho_{2\to 1}(f(t-\tau))]. \tag{3.20}$$

Even at a stable fixed point, agents are still switching back and forth between resources, incurring small, but nonzero transaction costs. The amount earned per agent is the payoff per agent minus transaction cost per agent, *i.e.*,

$$\begin{aligned}\mathrm{Earnings}(t) = \mathrm{Payoff}(t) - \mathrm{Cost}(t) = \\ f(t)G_1(f(t)) + (1 - f(t))G_2(t) \\ -\alpha T[f(t)\rho_{1\to 2}(f(t-\tau)) + (1 - f(t))\rho_{2\to 1}(f(t-\tau))].\end{aligned} \tag{3.21}$$

At a fixed point, earnings are accrued at a constant rate. When there are multiple fixed points, the performance of the system depends upon which fixed point is attained. The fixed point which yields the highest performance is defined as the optimal fixed point and for our previous choice of payoffs, is located at $f = 0.75$.

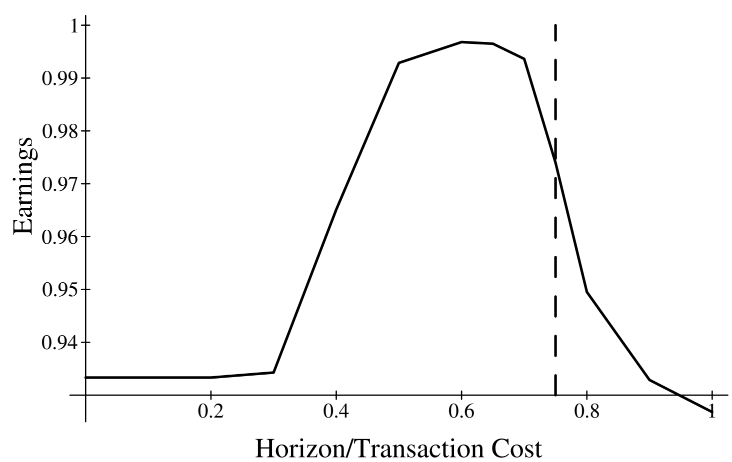

Figure 3.4: *Optimal horizon length.* Normalized mean earnings per agent as a function of horizon length. Payoffs for the two resources are $G_1 = 4 + 7f - 5.333f^2$ and $G_2 = 4 + 3f$. $T = 1$, $\tau = 6$, $\alpha = 1$ and $\sigma = 0.5$. For these parameters the fixed point is stable to left of the dashed line and unstable to the right of it. Earnings are averaged over the first 100 time steps, with the initial condition $f = 0.5$. The optimal horizon length is the one at which the transient earnings are highest: in this example, the optimal horizon is $H = 0.6$.

Determining the optimal horizon length involves trading off the system's performance against its response time. With increasing horizon lengths, stability decreases in the sense that $Re(s) \to 0^-$ and accrued transaction costs increase. Thus, the maximum performance of a system is always greater for smaller horizons. However, the response of the system to excursions away from the fixed point improves as the horizon gets longer (see Eq. 3.18), so that if the performance is averaged over the initial response, the optimal horizon at a fixed level of uncertainty depends strongly on the initial conditions. Consider, for example, a system with the resource payoffs $G_1 = 4 + 7f - 5.333f^2$ and $G_2 = 4 + 3f$. In Fig. 3.4, the normalized average earnings of the system (averaged over the first 100 time steps, with the initial condition $f = 0.5$) is plotted *vs.* the horizon length of the agents' expectations. In this case, the system is stable for horizon lengths less than 0.75, and unstable at horizons greater than this value. The optimal horizon in this case is $H = 0.6$, although performance can be marginally greater at shorter horizons given an ideal choice of initial conditions. At this horizon length, the average earnings are about 99.7% of maximum.

3.3.3 Time–varying expectations

We next consider a society of agents which, unaware of (or indifferent to) the trade-off between stability and response time, plays a dangerous game. Intent on maximizing performance, these agents slowly increase the horizon into the future for which they trust their expectations to remain reliable. As the horizon length of their flat predictions increases, the respective performance of the system also improves, until the optimal horizon length is reached. At this point, since the system is very near the optimal fixed point, increasing the horizon further will not deteriorate performance significantly and greedy agents might unwisely choose to extend the horizons even further. Soon enough the system becomes unstable and begins to oscillate.

Imagine that the horizon length increases linearly with time starting with a zero value horizon. Fig. 3.5 shows the dynamical approach of the system to the optimal fixed point (corresponds to optimal performance) followed by the inevitable crash. In this example, the transaction cost is 1, the uncertainty 0.2, the delay 6 time units, and the reevaluation rate is 1, with the same resource payoffs introduced earlier. The horizon length was made to increase slowly at a rate of 0.005.

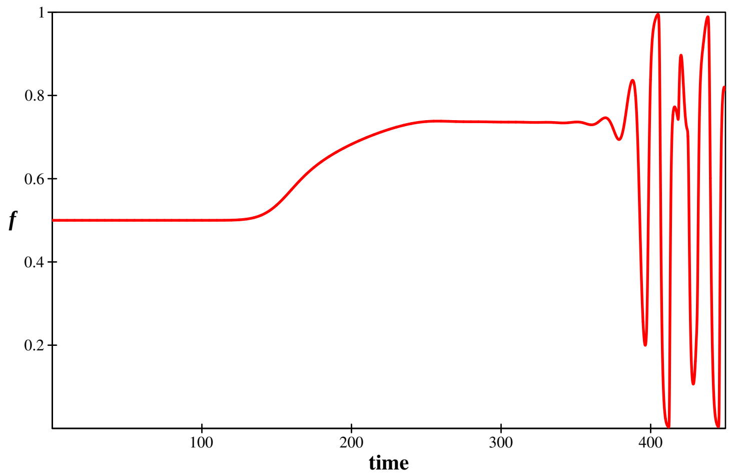

Figure 3.5: *Rising expectations.* Horizon length of agents increases linearly with time from $H = 0$ to $H = 2.25$. Instability first hits at $H \approx 1.7$, but is not noticeable until $H \approx 1.85$. As the horizon increases, the distribution of agents approaches its optimal value reaching its maximum at the horizon length which best counters the trade-off between stability and response time. As the horizon increases further, performance is only slightly compromised, but eventually the horizon becomes large enough that the system destabilizes. Resource payoffs are as in previous figures, $T = 1$, $\tau = 6$, $\alpha = 1$ and $\sigma = 0.2$

Curiously, the point at which the system becomes unstable is not what might have been predicted from examination of Fig. 3.1. Instead of crossing over at $H = 1.5$ into the unstable region A, the system remains stable until $H \approx 1.7$. Furthermore, if we were to next imagine that the oscillating agents started decreasing the horizon length as soon as they observed instability, say when $H = 1.9$ (after 380 time steps) the system would not restabilize (determined here by absence of fluctuations) until $H \approx 0.9$.

What is causing the hysteresis in stability as a function of horizon length? To understand this effect, first note that even with increasing horizon *H*, the unstable fixed point is only marginally unstable. Thus, instabilities grow very slowly when the system is near one of these unstable fixed points. Since the fixed point changes continuously with horizon length, the system remains near the instantaneous fixed point. However, once the system becomes unstable, it takes a long time to restabilize as the fixed point is weakly attractive just below the outside the stability boundary.

Interestingly enough, we also found that a stable system of agents can support a subpopulation of agents with indefinitely increasing expectations. As long as the subpopulation is small, the system may remain stable, but at a new fixed point. However, once the subpopulation becomes large enough, rising expectations once again cause a transition to instability.

3.4 Trends

In the previous section we stipulated that agents believe that system averages remain to remain roughly constant over time and that they persist in this belief despite evidence to the contrary. If the environment in which the system resides varies little with time, a stable system will be better able to profitably use the resources than an oscillating one. Thus, in this case, the agents' flat expectations are counterproductive in that they create instability out of a stable environment.[5]

This motivates us to introduce a new type of agent, the trend-follower. Such an agent not only has access to the delayed (and again, imperfect) state of the system, but can also recall the state of the system back to some point in the past. This means that the agent knows the derivative at the delayed time. We further specify that these agents have short memories, on the order of a few delays, and consequently follow short-term trends as opposed to long-term trends. Examining the dynamical behavior of trend–followers will elucidate the effect of this second type of expectations.

Recalling Eqs. 3.12 and 3.16, we see that the time-dependence of $f(t)$ is now a function of both $f(t-\tau)$ and $df(t-\tau)/dt$, so that the evolution of the state of the system depends on both the delayed value and the delayed derivative. This type of differential equation is called a neutral equation as it contains both delayed and advanced terms (a delayed derivative is equivalent to an advanced term). If we had chosen to give the agents perfect forecasting abilities so that $f^*(t+s) = f(t+s)$, then we would have a very complicated advanced equation, known to always be unstable (Bellman and Cooke 1963). Fortunately, self-fulfilling behavior is so unusual that we need not consider it in this context.

---

[5] Although we showed in the previous chapter that oscillations are not necessarily inferior when the environment itself changes with time.

3.4.1 The revised stability portrait

Once again, linear stability analysis can be used to give an implicit solution to the functional form of the stability boundary. Following the approach taken in the previous section (which is also valid for neutral differential-difference equations), the following characteristic equation for *s* is obtained, valid for all negative values of *s*:

$$\begin{aligned} &\rho_{1\to1}(f_0) - \rho_{2\to1}(f_0) - 1 \\ &+ \exp(-\beta s)[1 + \alpha s(H + T)][f_0(\gamma_{1\to1} - \gamma_{2\to1}) + \gamma_{2\to1}] - s = 0. \end{aligned} \tag{3.22}$$

This differs from the characteristic equation for the case of flat expectations only in the factor of $[1 + \alpha s(H + \tau)]$, which arises from combining terms involving $\gamma = \frac{\partial\rho}{\partial f}|_{f=f_0,f'=0}$ and $\xi = \frac{\partial\rho}{\partial f'}|_{f=f^{\mathrm{o}},f'=0}$. (A straightforward calculation yields that $\xi = (H + \tau)\gamma$ for quadratic resource payoffs.) In the limit $\mathrm{Re}(s) \to 0^-$, an analytic expression for the stability boundary $\sigma(H)$ at fixed transaction cost, *T*, follows and is given by solving the following system along with Eq. 3.16 (which gives f_0 implicitly):

$$\begin{aligned} &[f_0(\gamma_{1\to1} - \gamma_{2\to1}) + \gamma_{2\to}][\cos\beta s_2 + \alpha s_2(H + \tau)\sin\beta s_2] \\ &\quad + (\rho_{1\to1}(f_0) - \rho_{2\to1}(f_0) - 1) \ = 0; \\ &[f_0(\gamma_{1\to1} - \gamma_{2\to1}) + \gamma_{2\to1}][\sin\beta s_2 - \alpha s_2(H + \tau)\cos\beta s_2] + s_2 = 0. \end{aligned} \tag{3.23}$$

The above expressions also depend on the imaginary part of the root, s_2, and we must solve the system for many values of s_2 since the stability boundary is given by the largest unstable region delineated by $\sigma(H)$. A theorem in the theory of neutral differential-difference equations states that the roots of Eq. 3.22 fall asymptotically on a line parallel to the imaginary axis, as $|\mathrm{s}_2| \to \infty$, simplifying the task.

The resulting stability diagram (for fixed transaction cost) is shown in Fig. 3.6 for the same resource payoffs and parameters used to establish the stability portrait for the case of flat expectations. The boundary to region B was again obtained by repeated numerical integration of the dynamical equation. As can be seen, the stability portrait is very similar to that of Fig. 3.1, the phase space diagram for the case of flat expectations. Notice, however, that the unstable regime is much larger for the system with trendy expectations. In region A all fixed points

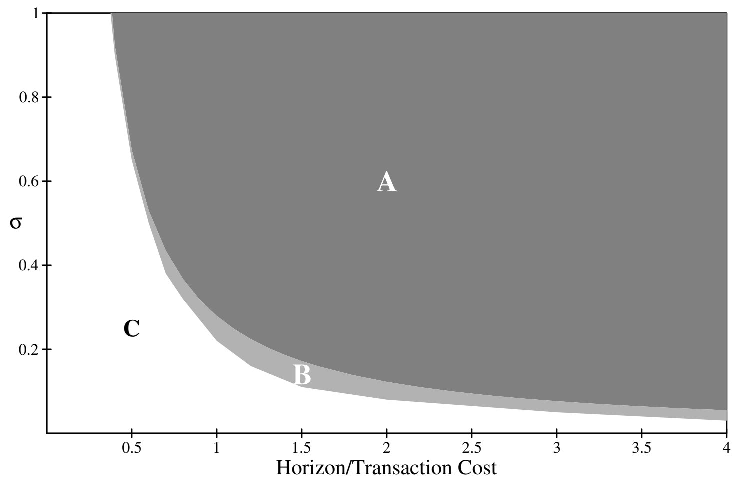

Figure 3.6: *Stability portrait of trend-followers* in σ *vs.* H/T phase space for a system with resource payoffs $G_1 = 4 + 7f - 5.333f^2$ and $G_2 = 4 + 3f$, $T = 1$, $\alpha = 1$ and $\tau = 6$. System is always unstable in region A and always stable in region C. In region B, the system either relaxes to a fixed point or is pulled into a noisy limit cycle, depending upon the initial conditions.

are unstable, while in region C the system is always stable. The upper boundary to region A extends to larger values of uncertainty $(\sigma > 4.5)$. Again, for small values of σ, only very large horizon lengths destabilize the system. In region B, equilibrium behavior depends on the initial conditions: for some initial conditions, the system relaxes to a fixed point; for others, the system is pulled into a limit cycle.

3.4.2 Sensitivity to trends

Closer examination shows that, while the phase diagrams for the two types of expectations are qualitatively similar, a system of trend-followers always becomes unstable at smaller horizon lengths than a collection of agents with flat expectations. This can be understood intuitively as follows. Linear extrapolation is an accurate prediction method only for brief periods following the arrival of information. Consequently, for systems with non-negligible delays in obtaining information, the accuracy of the extrapolation from the past state of the system to the present is already compromised, not to speak of extrapolation to the future. Thus, expectations of trend-followers are often very far off, especially considering that for large gradients, the expectations rapidly saturate to $f = 0$ or $f = 1$, both rather unlikely states of the system. Agents react to a large slope by flocking to the less popular resource, creating a large gradient in the opposite direction. In such a way, trend-followers needlessly cause large fluctuations.

Fig. 3.7(a) exemplifies this type of behavior. The choice of parameters $H = 0.5$, $T = 1$, $\alpha = 1$, $\tau = 6$ and $\sigma = 1.0$ places a collection of trend followers in region A of Fig. 3.6, but in the stable regime (region C) of the *σ–H/T* phase space portrait for the system made up of agents with flat expectations. Thus, a flat expectation system damps down to the fixed point, while the trendy system continues to oscillate.

(a) **(b)**

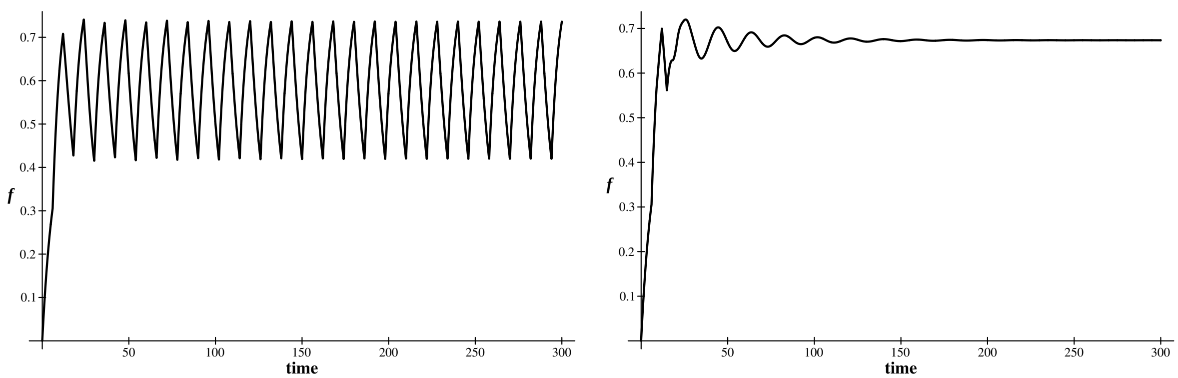

Figure 3.7: *Following trends is destabilizing.* Resource payoffs of Fig. 3.6 with $H = 0.5$, $T = 1$, $\alpha = 1$, $\tau = 6$ and $\sigma = 1.0$. System is in region A of Fig. 3.6. Initially, $f = 0$. (a) demonstrates the destabilizing effect of following short-term trends in a system of agents which base their decisions on expectations of the future. For the same choice of parameters, a system of agents with flat expectations damps out to the fixed point after about 200 time steps. (b) If agents are oblivious to gradients of less than 0.05 change in fractional density per time step, the system converges to the fixed point at $f = 0.674$ after about 150 time steps.

An important point remains to be raised about the way in which we modeled agents which follow trends. Our computational representation of such agents detects gradients up to the numerical precision specified by our experiments, which are limited by the precision of the computer. Real-life agents probably wouldn't notice trends with slopes flatter than about 1% over a few delay periods, as such gradients would not be judged significant. Since small gradients are magnified due to the compensating mechanism described above, slopes that might be ignored in a real system contribute to the instability of the simulated system when the system is near a fixed point. Indeed, increasing the threshold at which agents observe slopes above the numerical precision (for example slopes less than 0.01 are truncated to zero) also increases the stability of the system. Fig. 3.7(b) shows the eventual relaxation to the fixed point of a system of trend-followers blind to gradients in the fractional density of agents larger than 0.05 per time step for the same parameters as in Fig. 3.7(a).

Thus, while following trends may superficially seem like a more intelligent way of formulating expectations than flat expectations, it is inherently destabilizing. In reality, it is often a very unintelligent prediction mechanism, since in a world with finite resources, trends tend to top off or bottom off, a fact of which smart agents should be aware.

3.5 Diversity and Competition

In the previous sections we assumed that all agents have identical methods of forming expectations about the future in a first attempt to gauge the effect of expectations on overall system behavior. In a real system, however, different agents might have different types of expectations. We now incorporate the diversity of expectations by allowing several types of agents with differing expectations to coexist within a system. Specifically, we consider agents with flat expectations, but allow the horizon length to vary among types of agents. As a concrete example, we will study a diverse system made up of agents whose horizon length is inversely proportional to its level of uncertainty. The smaller the agent's uncertainty, the greater its confidence in its ability to extrapolate into a future, and thus the longer its horizon length. The different types of agents are then characterized by their uncertainty and horizon length, where the relation $H\sigma =$ constant is taken to hold across all agent types.

3.5.1 Agent types and rewards

In a closed world with a finite amount of resources, these different strategies are in competition with each other. In order to model the effect of competition, we now introduce a reward mechanism which encourages the spread of high-performing strategies at the expense of the poorer ones. The reward can be thought of in an economic sense as corresponding to monetary benefits, for example, or in a biological sense to reproductive success. The implementation of rewards presented here generalizes the treatment given in (Hogg and Huberman 1991) to include the effect of expectations and transaction costs.

While agents base their decisions on perceived accumulated future payoffs, the system rewards them in proportion to their actual performance or net income. The computational ecosystem can now be described by specifying the fraction of agents, f_{rs}, of a given type s using a given resource r at a particular time t. The fraction of agents using a resource r is then $f_r^{res} = \sum_s f_{rs}$ and the fraction of agents of type s is $f_s^{type} = \sum_r f_{rs}$.

As mentioned previously, the net effect of rewarding performance is to increase the fraction of highly performing agents. If κ is the rate at which performance is rewarded, then Eq. 2.1 is enhanced with an extra term which corresponds to this reward mechanism, giving

$$\begin{aligned}\frac{df_{rs}}{dt} = \alpha[f_{rs}\rho_{r\to r,s}(f_1^{res}(t-\tau)) + f_{r's}\rho_{r'\to r,s}(f_1^{res}(t-\tau)) - f_{rs}] \\ +\kappa(f_r^{res}\eta_s - f_{rs}),\end{aligned} \tag{3.24}$$

with $r \neq r'$. The second term now incorporates the effect of rewards on the population. In this equation

$$\rho_{1\to 2,s}(f_1^{res}(t-\tau)) = \frac{1}{2}\left\{1 + \text{erf}\left[\frac{H_s[G_2(f_1^{res}(t-\tau)) - G_1(f_1^{res}(t-\tau))] - T}{2H_s\sigma_s}\right]\right\} \tag{3.25}$$

is the probability that an agent using resource 1 of type s with horizon length H_s will prefer resource 2 (taking into account transaction costs) and

$$\rho_{1\to 1,s}(f_1^{res}(t-\tau)) = \frac{1}{2}\left\{1 - \text{erf}\left[\frac{H_s[G_2(f_1^{res}(t-\tau)) - G_1(f_1^{res}(t-\tau))] - T}{2H_s\sigma_s}\right]\right\} \tag{3.26}$$

is the probability that it will prefer to stay with resource 1. The function η_s is the probability that agents will be of type s, which we take to be proportional to the earnings achieved by agents of type s. Specifically, we choose

$$\eta_s = \frac{[\text{Earnings}(t)]_s}{\sum_s [\text{Earnings}(t)]_s}\,, \tag{3.27}$$

with

$$\begin{aligned}[\mathrm{Earnings}(t)]_s = [\mathrm{Payoff}(t)]_s - [\mathrm{Cost}(t)]_s = \\ f_{1s}G_1(f_1^{res}(t)) + f_{2s}G_2(f_1^{res}(t)) \\ -\alpha T(f_{1s}\rho_{1\to 2,s}(f_1^{res}(t-\tau)) + f_{2s}\rho_{2\to 1,s}(f_1^{res}(t-\tau))),\end{aligned} \tag{3.28}$$

so that an agent's earnings are linearly proportional to its net earnings.[6]

3.5.2 The crash of rising expectations

How does a diverse system of agents with different horizon lengths evolve when good performance is rewarded at the expense of bad performance? In some cases, competition serves to stabilize behavior to a large extent. For many different populations of agent types, the behavior of the system is stable most of the time with comparatively brief bouts of instability. Consider, for example, a population of agents with horizons ranging from 0 to 2. Initially, 9.1% of the agents are of each different type. The transaction cost is set at $T = 1.5$ with parameters $\alpha = 1$, $\tau = 6$. Each agent type has a different value of uncertainty σ_s such that the smaller the uncertainty, the longer the horizon length; in this example, $H_s\sigma_s = 0.5$. Without the reward mechanism, this mix of agents oscillates between the two resources at a fixed frequency and with large amplitude, due to the influence of the agents with longer horizons. Once the agents are permitted to compete for payment among themselves, the system rapidly stabilizes after about 200 time steps. Subsequently, the system drifts very slowly, as shown in Fig. 3.8(a). As the system drifts, the population of agents gradually redistributes itself among the different types, the ones with longer horizon times garnering a greater share of the system's earnings and consequently attracting more and more of the agents.

This behavior is not surprising, since when the system changes little over time, expectations about the future are significantly more accurate. An agent with a long horizon length correlated to its low level of uncertainty then has a big advantage: the agent can look out far ahead with

[6] It is possible for the total net earnings of the system to be negative. In this case the dynamical equation describing the evolution of f_{rs} becomes

$$\frac{df_{rs}}{dt} = \alpha\left[f_{rs}\rho_{r\to r,s}(f_1^{res}(t-\tau)) + f_{r's}\rho_{r'\to r,s}(f_1^{res}(t-\tau)) - f_{rs}\right] + \kappa(f_{rs} - f_r^{res}\eta_s).$$

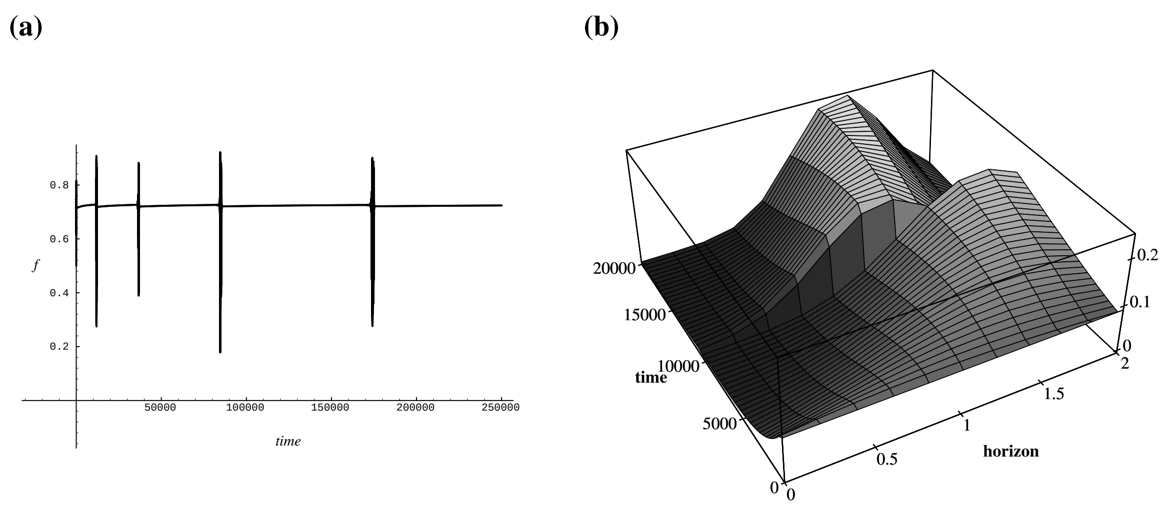

Figure 3.8: *Crash of rising expectations.* (a) Time evolution of the fraction of agents using resource 1 for 250,000 time steps. Competition for rewards among agents with different horizon lengths damps out oscillations, but remains latently unstable. Resource payoffs are $G_1(f_1^{res}) = 4 + 7f_1^{res} - 5.333(f_1^{res})^2$ and $G_2(f_1^{res}) = 4 + 3f_1^{res}$, α=1, τ=6, $H_s\sigma_s = 0.5$, T=1.5, κ=1. Horizon lengths of the 11 different types range uniformly between 0 to 2. Initially, the population of agents is evenly split among the different types. (b) Time evolution of the percentages of agents of different types over the first 20,000 time steps. Note the abrupt decrease of agents with long horizons after the onset of oscillations with the complimentary increase in agents with shorter horizons.

great accuracy. As a result, agents with the longer horizon times distribute themselves among the two resources more optimally, have a greater instantaneous payoff, and a larger share of the system's earnings.

Earlier, we observed that a stable system of agents with identical levels of uncertainty can support a subpopulation with longer horizons as long as that subpopulation remains small. Again, we see here that once a large enough majority of the agents have long expectations, the system becomes unstable. This instability is dramatically depicted in Fig. 3.8(b) where the distribution of agents among different types is seen to change abruptly after the onset of instability. The proportion of agents with long horizons falls abruptly, with types at shorter horizons taking up the slack. The large oscillations of the system render these agents' flat expectations completely inaccurate, despite their smaller uncertainty; however, agents with shorter horizons switch between resources less often, incur less transaction costs and thus receive a greater proportion of the system's earnings. Once again, the reward mechanism serves to damp out the oscillations, and as time passes agents with long horizons regain their advantage.

This crash-stabilize cycle repeats itself over and over, but at longer and longer time intervals, for the transient shown in Fig. 3.8(a). As the cycle repeats, diversity is reduced. Agents with long horizons are weeded out during the crashes while agents with short horizons do poorly during the long periods of quasi-stability. After 250,000 time steps only four types remain in significant numbers: those with horizons of 1.2, 1.4, 1.6, and 1.8, and of those over 80% are of type with horizon length 1.4.

Numerical integration of the system's dynamics over 20 million time steps has been independently carried out by Nils Jensen on a supercomputer[7]. He found that after an initial transient of about 500,000 time steps, the crash-stabilize pattern repeats itself about every 250,000 time steps. This result indicates that system most likely never achieves fixed point behavior by eliminating diversity. If this were to occur, however, it would be on a time scale so large that is inconceivable that a real system would not have meanwhile suffered some destabilizing perturbation.

Analytical results corroborate the experimental finding that no stable mix of agents with different horizon lengths is possible. Fixed point behavior occurs only when one type of agent prevails completely. The following three conditions can be shown to be sufficient for fixed point behavior for a system with uncertainty: (1) the fixed point factorizes, *i.e.*, $f_{rs} = f_r^{res} f_s^{type}$; (2) the different types of agents share a common fixed point in the absence of interactions, that is, $f_1^{res}\rho_{1\to 1,s} + f_2^{res}\rho_{2\to 1,s} - f_1^{res} = 0$ for all s; and (3) the number of transactions per type of agent (again in the absence of interactions) is also independent of type: $\alpha(f_1^{res}\rho_{1\to 2,s} + f_2^{res}\rho_{2\to 1,s}) = \mathrm{const}$. Combining (2) and (3) implies that $\rho_{1\to 2,s}$ is independent of s at a fixed point. This is a contradiction, since $\rho_{1\to 2,s}$ depends on s through H_s. Consequently, no stable mix (of two or more types) exists, although the system can still achieve stability by evolving into a system with only one type of agent.

3.6 Summary

We have shown how the global behavior of a system of locally-acting intentional agents can be analyzed by allowing them to form expectations about the future and act upon them and how

[7] Personal communication.

agents' expectations of the future have a strong effect on the group dynamics. In this spirit, we extended the Huberman-Hogg theory of computational ecosystems to incorporate the effect of individual's expectations about the future. Specifically, we considered two different types of expectations, both of which reflect different beliefs of how systems evolve in time. The first type of agent believes that systems change very little in time. The second type follows short-term trends, trusting that these will continue. All agents discount the future to some extent, at a rate inversely proportional to the horizon length of their expectations. In addition, we also penalized agents for switching between resources by charging transaction costs.

The range of dynamics resulting from this model of computational ecosystems with expectations and transaction costs is very broad. Depending on the type of expectations, the amount of diversity, and the structure of the reward mechanism many different types of behavior can ensue including fixed points, limit cycles, and chaos.

Our results show that, to a large extent, the effect of transaction costs is to stabilize behavior. We demonstrated that transaction costs can also trap the system in suboptimal states for extremely long periods of time, since near a fixed point or an almost fixed point the transition rate between resources drops by a factor exponentially small in the transaction cost. We saw that at small uncertainty, the horizon length of expectations at which the system goes unstable approaches infinity given fixed transaction costs. In addition, for small ratios of the horizon length to transaction costs, the fixed point is stable at all uncertainty levels, and for any choice of horizon length and transaction, it is possible to stabilize the system by increasing the uncertainty. From our computer experiments we also discovered that in some regimes the equilibrium dynamics of the system depends on its initial state: for some initial configurations, the system relaxes to fixed point behavior; for others, it is pulled into a limit cycle. These results hold for agents with either flat or trendy expectations, although the former always do better in the sense that their behavior remains stable for longer horizon lengths.

More significantly, we found that agents with flat expectations attempting to optimize their horizon length face a trade-off between the goals of maximizing performance within a set response time and maintaining stability. Increased horizon length improves the response of the system while jeopardizing stability. Greedy agents hoping to maximize profit might gradually increase their horizon lengths at times when the system appears to be stable, unaware of incipient instabilities caused by nonzero uncertainty.

We also showed how a diversity of expectations among the agents, coupled to the reward mechanisms in which those agents that do well increase at the expense of the others, can lead to overall dynamics characterized by cycles of almost stable behavior interrupted by sudden crashes. This process is accompanied by an ever-changing diversity in the composition of the system. As time progresses the crashes become further apart but remain unpredictable.

PART II

COLLECTIVE ACTION

4 Dynamics of Cooperation with Expectations

In a general social dilemma, a group attempts to obtain a common good in the absence of central authority. Each individual has two choices: either to contribute to the common good, or to shirk and free ride on the work of others. The payoffs are structured so that the incentives in the game mirror those present in social dilemmas. All individuals share equally in the common good, regardless of their actions. However, each person that cooperates increases the amount of the common good by a fixed amount, but receives only a fraction of that amount in return. Since the cost of cooperating is greater than the marginal benefit, the individual defects. Now the dilemma rears its ugly head: each individual faces the same choice; thus all defect and the common good is not produced at all. The individually rational strategy of weighing costs against benefits results in an inferior outcome — no common good is produced.

However, the logic behind the decision to cooperate or not changes when the interaction is ongoing since future expectations influence the rational individual's decision. In particular, individual beliefs concerning the future evolution of the game can play a significant role in each member's decisions.

4.1 Introduction

Collective action problems pose difficult quandaries for communities, be they social, economic, or organizational. Consider a group that wants to attain a common goal, such as cleaner air, better working conditions, or the elimination of poverty. Members of the group may be tempted to benefit from the common good without contributing to its production when the gain for collaborating is less than the cost.

There is a long history of interest in social dilemmas of this type in political science (Schelling 1978, R. Hardin 1982). G. Hardin coined the phrase "the tragedy of the commons" to reflect the fate of the human species if it fails to successfully resolve the social dilemma of limiting population growth (G. Hardin 1968). Furthermore, Olson argued that the logic of collective action implies that only small groups can successfully provide themselves with a common good (Olson 1965). Others, from Smith (1937) to Taylor (1976, 1987), have taken the problem of social dilemmas as central to the justification of the existence of the state. In economics and sociology, the study of social dilemmas sheds light on, for example, the adoption of new technologies (Friedman 1990) and the mobilization of political movements (Oliver and Maxwell 1988).

One may then ask the following question about situations of this kind: if agents make decisions on whether or not to cooperate on the basis of imperfect information about the group activity and incorporate expectations on how their decision will affect other agents, can overall cooperation be sustained for long periods of time? Moreover, how do expectations and group size affect cooperation? These questions are relevant to distributed artificial intelligence as well as to social science. For example, instances of negotiation and cooperative problem solving (Davis and Smith 1983, Gasser and Huhns 1989, Osawa and Tokoro 1991, Zlotkin and Rosenschein 1991) imply the existence of intentional agents. In addition, the emergence of market-like computational systems that have no global controls (Malone *et. al.* 1988, Waldspurger *et. al.* 1992) raises the issue of having computational processes that could free-ride on the work of others (Miller and Drexler 1988a,b).

Several authors have considered the problem of the effect of group size on the possibility for collective action. Their results appear contradictory: some studies have shown that overall cooperation is undermined as the group increases in size (Olson 1965, Bendor and Mookherjee 1987), others, to the contrary, have shown that it is more likely for larger groups (Oliver and Maxwell 1988). However, their conclusions depend strongly on the specific scenario considered, as pointed out by R. Hardin (1982). In particular, collective action problems can be characterized by the degree to which one agent's consumption of the good reduces the amount available to another, called the jointness of supply. In this study, we will consider collective goods whose benefit is divided among the members — one individual's gain is another's loss. Such a divisible good contrasts with a purely public good, such as public radio, in which each individual's benefit is independent of the benefit received by another.

Apart from the question of whether or not a group of a given size can sustain cooperation, there remains the issue of how, if ever, that state is reached. One can imagine an overly large group initially composed of cooperators that eventually evolves into collective defection. Conversely, a large non-cooperative group could break up into smaller sub-groups; the relevant question is then, how long will it take for a sub-group to switch over to global cooperation?

In order to address these issues, we present and study a dynamical model of ongoing collective action among intentional agents whose choices depend not only on the past but also on their expectations as to how their actions will affect those of others. In this model agents act on the basis of information that can be uncertain at times. We show that under these conditions the onset of overall cooperation can take place in a sudden and unexpected way. Likewise, defection can appear out of nowhere in very large, previously cooperating groups. These outbreaks mark the end of long transient metastable states in which defection or cooperation persists in groups that cannot sustain it indefinitely.

Our mathematical treatment of the collective action problem is presented in Section 4.2. There, we state the benefits and costs to the individual associated with the two actions of cooperation and defection, *i.e.* contributing or not to the social good. The problem thus posed is referred to in the literature as the *n*-person prisoner's dilemma (R. Hardin 1971, Taylor 1976, Bendor and Mookherjee 1987). We then allow beliefs and expectations about other individuals' actions in the future to influence each member's perception of which action, cooperation or defection, will benefit it most in the long run.

We recast the interaction as an asynchronous dynamic game in which each individual reconsiders its decision at an average reevaluation rate α, using delayed information of the level of provision of the good. The iterated game is of finite duration, and individuals decide to cooperate or defect by determining which choice maximizes their *expected* share of the good for the remainder of the game. In addition, uncertainty is introduced into the relation between individual effort and group performance to model imperfect information and bounded rationality.

The type of expectations that we consider consist of two components: each individual believes (1) that future aggregate collective behavior is directly influenced by the individual's choices in inverse proportion to the group size; and (2) that the interaction is of finite duration characterized by a horizon length, H. Thus, individuals believe that in the long run their actions encourage similar actions on the part of others in the group. There are various mechanisms by which this might occur: imitation and establishment of conventions or norms are two. In addition, individuals believe that the ability of their actions to encourage like actions increases with the number of individuals who have chosen to contribute to the collective good (as opposed to free riding). In some sense, this reflects a belief that contributing individuals form a core group whose reactions are much more sensitive to fluctuations in the amount of the good produced.

We show that an individually rational strategy of conditional cooperation emerges from the individuals' expectations and beliefs. Individuals cooperate if they perceive the fraction cooperating to be greater than some critical amount and defect otherwise. This strategy of conditional cooperation is reminiscent of the successful tit-for-tat strategy in 2–player prisoner's dilemma in which a player cooperates if and only if its opponent cooperated in the previous turn and defects otherwise (Axelrod 1981). As pointed out by Dawes (1980), however, the n-person case is qualitatively different from the 2–person interaction.

In Section 4.3, we analyze the dynamics of fluctuations away from the equilibria using a thermodynamic-like formalism (Ceccatto and Huberman 1989). Besides confirming that there exists a critical group size beyond which cooperation is not sustainable (Olson 1965, Bendor and Mookherjee 1987), we find additional, intriguing effects. Our results reveal several different dynamical regimes that should be observable as the size of a group changes. In one regime, cooperation is a metastable state: it may persist for very long times even for groups exceeding a critical size, although group behavior eventually decays to overall defection. In another regime, defection is metastable: the system can remain stuck in a non-cooperative state even though its size is well below that guaranteeing long term cooperation. These effects are shown to depend strongly on the degree of uncertainty pervading the system, as well as on the length of the individual's horizons.

Section 4.3 also contains the results of computer simulations which confirm our analytical predictions, in particular the persistence of metastable states. We also present an example of a phenomenon in which the average behavior is not typical. In this example, group behavior flip-flops between the strategies of mutual cooperation and mutual defection, without ever settling into one or the other.

4.2 A Model of Ongoing Collective Action

4.2.1 The economics of free riding

In our model of social dilemmas, each individual can either contribute (cooperate) to the production of the good, or not (defect). While no individual can directly observe the effort of another, each member observes instead the collective output and can deduce overall group participation using knowledge of individual and group production functions. We also introduce an amount of uncertainty into the relation between members' efforts and group performance. There are many possible causes for this uncertainty (Bendor and Mookherjee 1987); for example, a member may try but fail to contribute due to unforeseen obstacles. Alternatively, another type of uncertainty might arise due to individuals with bounded rationality occasionally making suboptimal decisions (Selten 1975, Simon 1962). In any case, we treat here only idiosyncratic disturbances or errors, whose occurrences are purely uncorrelated.

Consequently, we assume that a group member intending to participate does so successfully with probability p and fails with probability $1 - p$, with an effect equivalent to a defection. Similarly, an attempt to defect results in zero contribution with probability q, but results in unintentional cooperation with probability $1 - q$. Then, as all attempts are assumed to be uncorrelated, the number of successfully cooperating members, $\widehat{n}^c$, is a mixture of two binomial random variables with mean $< \widehat{n}^c >= pn^c + (1 - q)(n - n^c)$, where n^c is the number of members attempting to cooperate within a group of size n.

Let k_i denote whether member i intends to cooperate ($k_i = 1$) or defect ($k_i = 0$), and let k_i' denote whether member i is cooperating or defecting in effect. The number of members cooperating is $\widehat{n}^c = \sum_j k_j'$.

The limit p and q equal to 1 corresponds to an error-free world of complete information, while p and q equal to 0.5 reflect the case where the effect of an action is completely divorced from intent. Whenever p and q deviate from 1, the perceived level of cooperation will differ from the actual attempted amount.

In a simple, but general limit, collective benefits increase linearly in the contributions of the members, at a rate b per cooperating member. Each contributing individual bears a personal cost, c. Then the utility at time t for member i is

$$U_i(t) = \frac{b}{n}\widehat{n}^c(t) - ck_i. \tag{4.1}$$

Using its knowledge of the functional form of the utility function,[8] each individual can deduce the number of individuals effectively cooperating at some time t by inverting Eq. 4.1:

$$\widehat{n}^c(t) = \frac{n}{b}(U_i(t) + ck_i). \tag{4.2}$$

Of course, this estimation will differ from the actual number of individuals intending to cooperate in a manner described by the mixture of two binomial distributions. We also define $\widehat{f}^c(t)$, to denote the fraction, $\widehat{n}_c(t)/n$, of individuals effectively cooperating at time t.

When all members contribute successfully, each receives net benefits $(bn/n) - c = b - c$. The production of the collective good becomes an n-person prisoners' dilemma when

$$b > c > \frac{b}{n}. \tag{4.3}$$

Thus, although the good of all is maximized when everyone cooperates $(b - c > 0)$, the dominant strategy in the one-shot game is to defect since additional gain of personal participation is less than the private cost $(b/n - c < 0)$.

4.2.2 Expectations

The logic behind the decision to cooperate or not changes when the interaction is ongoing since future expected utility gains will join present ones in influencing the rational individual's

[8] For a justification of the form of the individual utility function in the context of either divisible goods or pure public goods see (Bendor and Mookherjee 1987).

decision to contribute or not to the collective good. In particular, individual expectations concerning the future evolution of the game can play a significant role in each member's decisions. The importance individuals place on the future depends on how long they expect the interaction to last. If they expect the game to end soon, then rationally, future expected returns should be discounted heavily with respect to known immediate returns. On the other hand, if the interaction is likely to continue for a long time, then members may be wise to discount the future only slightly and make choices that maximize their returns on the long run. Notice that making present choices that depend on the future is rational only if, and to the extent that, a member believes its choices influence the decisions others make.

The time scale of the interaction is set by the rate, α, at which members of the group reexamine their choices. Information about the level of cooperation is deduced from individual utility accrued in the past, as per Eq. 4.2, and can be delayed by an interval τ. Along with expectations about the future, the fraction observed as cooperating, $\widehat{f}^c$, and the reevaluation rate, α, determine in part how individuals believe the level of cooperation will evolve over time.

For simplicity we assume all members of the group share a common rationality in their method of forming expectations. Specifically, all members expect the game to be of finite duration *H*, the horizon length[9]. Thus, future returns expected at a time t' from the present are discounted at a rate $e^{-t'/H}$ with respect to immediate expected returns.

Secondly, each member expects that their choice of action, when reflected in the net benefits received by the others, will influence future levels of cooperation. Since, however, the decision of one individual affects another's return by only $\pm b/n$, each member perceives its influence as decreasing with increasing group size. Furthermore, individual changes in strategy are believed to be most effective in encouraging similar behavior when levels of cooperation are high. We postulate that these two effects compound so that each member expects its decision to cooperate (defect) to encourage an overall growth (decay) in the level of cooperation at a rate that depends on $\widehat{f}^c/n$[10]. Roughly, then, a member expects its cooperative (defecting) action to stimulate an

---

9 The concept of a horizon is formally related to a discount δ, which reflects the perceived probability that the game will continue through the next time step. The two are connected through the relation $\sum_{i=0}^{\infty} \delta^i \,() \to \int_0^{\infty} dt'\, e^{-t'/H} \,()$ which implies $H = \frac{1}{1-\delta}$.

10 Assigning a precise functional form to individual expectations, although somewhat arbitrary, is necessary to study dynamics. Variations on the same theme, however, can be seen to yield similar dynamics, provided that the rate at which similar behavior is

additional $\widehat{f}^c/n$ members to cooperate (defect) during each subsequent time period.

A qualitative formulation of the manner in which cooperation encourages cooperation and defection encourages defection with the appropriate asymptotic properties now follows. Let $\Delta\widehat{f}^c(t+t')$ denote the expected future difference (at time $t+t'$) between the fraction of agents cooperating and the fraction of those defecting. Member *i*'s choice itself causes an instantaneous difference at $t' = 0$ of $\Delta\widehat{f}^c(t, t'=0) = 1/n$. To reflect member *i*'s expectation that during subsequent time steps (of average duration $1/\alpha$), its action will encourage $\widehat{f}^c/n$ additional members per time step to behave likewise, we stipulate that the difference $\Delta\widehat{f}^c(t+t')$ approach 1 asymptotically from its initial value 1/*n* at some given rate. Qualitatively, this rate looks like $e^{-\alpha\widehat{f}^c(t-\tau)t'/n}$, which gives the following expected deviation at time $t+t'$:

$$\Delta\widehat{f}^c\big(t+t'\big) = 1 - (1-1/n)\exp\left(-\frac{\alpha\widehat{f}^c(t-\tau)}{n}t'\right). \tag{4.4}$$

Variations on the precise functional form of the expected deviation $\Delta\widehat{f}^c(t+t')$ would simply cause the deviation to grow faster or slower with increasing t', without yielding significant qualitative changes in the types of dynamical behavior characterizing the interaction.

In summary, member *i* reevaluates its decision whether or not to contribute to the production of the good at an average rate α, using its knowledge of $\widehat{f}^c(t-\tau)$ and following its expectations about the future. Using its prediction of how it expects f^c to evolve in relation to its choice and discounting the future appropriately, member *i* then makes its decision on whether to cooperate or defect.

Putting it all together, member *i* perceives the advantage of cooperating over defecting at time *t* to be the net benefit

$$\begin{aligned}\Delta\widehat{B}_i(t) &= \int_0^\infty dt' e^{-t'/H}\Big\{b\,\Delta\widehat{f}^c\big(t+t'\big) - c\Big\}\\ &= H(b-c) - \frac{Hb(n-1)}{n + H\alpha\widehat{f}^c(t-\tau)}.\end{aligned} \tag{4.5}$$

expected to encourage similar behavior rises monotonically in the ratio $\widehat{f}^c/n$.

Member i cooperates when $\Delta \widehat{B}_i(t) > 0$, defects when $\Delta \widehat{B}_i(t) < 0$, and chooses at random between defection and cooperation when $\Delta \widehat{B}_i(t) = 0$. Its decision is based on the fraction of the group it perceives to have cooperated at a time τ in the past, $\widehat{f}^c(t-\tau)$. These criteria reduce to the following condition for cooperation at time t:

$$f^{crit} \equiv \frac{1}{H\alpha}\left(\frac{nc-b}{b-c}\right) < \widehat{f}^c(t-\tau). \tag{4.6}$$

Since $\widehat{f}^c(t-\tau)$ is a mixture of two binomially distributed variables, Eq. 4.6 provides a full prescription of the stochastic evolution for the interaction. In particular, member i cooperates with probability $P_c(f^c(t-\tau))$ that it perceives cooperation as maximizing its expected future accumulated utility, given the actual attempted level of cooperation $f^c(t-\tau)$. This probability is evaluated in section 4.5.1.

Thus, the individuals engage in conditional cooperation, cooperating when the fraction perceived as cooperating is greater than the critical amount, f^{crit} and defecting when the perceived fraction cooperating is less than f^{crit}. Because of the nature of the individuals' expectations, this behavior, reminiscent of a generalized tit-for-tat strategy, emerges from individually rational choices. In this way, conditional cooperation is a rational strategy. Furthermore, conditional cooperation is a more general type of strategy than might be expected from the explicit specification of the individuals' beliefs. As a result, it holds for a wider range of models than considered here, once one allows for variation in the functional form of f^{crit}.

4.2.3 Average group behavior

A deterministic differential equation describing the stochastic discrete interaction specified above can be derived using the mean-field approximation. This entails assuming that: (1) the size, n, of the group is large; and (2) the average value of a function of some variable is well approximated by the value of the function at the average of that variable. (An introduction to mean field theory can be found in Ashcroft and Mermin (1976), for example.) The model developed below will allow us to determine the equilibrium points and their stability characteristics.

We now specialize to the symmetric case $p = q$ in which an individual is equally likely to effectively defect when intending to cooperate as to cooperate when intending to defect. Later we will comment on the effect of this restriction on the generality of our results. By the Central Limit Theorem, for large *n*, the random variable $\widehat{f}^c$ tends in law to a Gaussian distribution with mean $\langle \widehat{f}^c \rangle = pf^c + (1-p)(1-f^c)$ and variance $\sigma^2 = p(1-p)/n$. Using the distribution of $\widehat{f}^c$, we now want to find the mean probability, ρ_c, that $\widehat{f}^c > f^{crit}$. Using a mean-field approximation, the mean probability that $\widehat{f}^c > f^{crit}$ is approximated by the probability $\rho^c(f^c)$ that $\langle \widehat{f}^c \rangle > f^{crit}$, (*i.e.*, $\langle \rho^c(\widehat{f}^c) \rangle = \rho^c(\langle \widehat{f}^c \rangle)$). Thus, the mean probability that $\widehat{f}^c > f^{crit}$ becomes

$$\rho^c(f^c) = \frac{1}{2}\left\{ 1 + \operatorname{erf}\left[\frac{\langle \widehat{f}^c \rangle - f^{crit}}{\sqrt{2}\sigma} \right] \right\}, \tag{4.7}$$

where $\operatorname{erf}\left(x/\sqrt{2}\sigma\right)$ represents the error function associated with the normal curve. The evolution of the number of agents cooperating in time is then described by the dynamical equation (Huberman and Hogg 1988)

$$\frac{df^c}{dt} = -\alpha[f^c(t) - \rho^c(f^c(t-\tau))], \tag{4.8}$$

where α is the reevaluation rate and τ is the delay parameter, as defined earlier.

The equilibrium points f_0^c of the interaction described by the above equation are obtained by setting the right hand side to zero[11]. In this case they are given by the solutions to

$$\rho^c(f_0^c) = f_0^c. \tag{4.9}$$

Solving the above equation yields the critical sizes beyond which cooperation can no longer be sustained. In the case of perfect certainty ($p = q = 1$), these critical sizes can be expressed in simple analytical form as shown in section 4.5.3. Thus, a group will no longer sustain global cooperation if it exceeds a value n^* given by

$$n^* = H\alpha\left(\frac{b}{c} - 1\right) + \frac{b}{c}. \tag{4.10}$$

[11] Furthermore, linear stability analysis of Eq. 2.1 shows that the stability of the equilibrium points is independent of the value of the delay τ, and of the reevaluation rate α. Thus, the asymptotic behavior of the group interaction does not depend on the delay; this observation is corroborated for the discrete model by numerical computer simulations similar to those presented in Section 4.3. Moreover, the equilibrium points belong to one of two types: stable fixed point attractors or unstable fixed point repellors. Due to the linearity of the condition for cooperation (Eq. 4.6), the dynamical portrait of the continuous model contains no limit cycles or chaotic attractors.

Similarly, cooperation is the only possible global outcome if the group size falls below a second critical size n_{min}:

$$n_{min} = \frac{b}{2c} + \frac{1}{2c}\sqrt{b^2 + 4H\alpha c(b-c)}. \tag{4.11}$$

Notice that these two critical sizes are not equal; in other words, there is a range of sizes between n_{min} and n^* for which either cooperation or defection is a possible outcome, depending on the initial conditions.

An estimate of the possible sizes can be obtained, for example, if one assumes a horizon $H = 50$ (which corresponds to a termination probability $\delta = 0.98$), $\alpha = 1$, $b = 2.5$ and $c = 1$. In this case one obtains $n^* = 77$ and $n_{min} = 10$. Observe that an increase in the horizon length would lead to corresponding increases in the critical sizes.

4.3 The Dynamics of Cooperation

4.3.1 The Ω function formalism

The mean–field approximation gives the average properties of a collection of agents having to choose between cooperation and defection. The asymptotic behavior generated by the dynamics gives only the fixed points of the system; it does not address the question of how fluctuations away from equilibrium state evolve in the presence of uncertainty. These fluctuations are important for two reasons: (1) the time necessary for the system to relax back to equilibrium after small changes in the number of agents cooperating or defecting might be long compared to the time-scale of the collective task to be performed or the measuring time of an outside observer; (2) large enough fluctuations in the number of defecting or collaborating agents could shift the state of the system from cooperating to defecting and *vice-versa*. If that is the case, it becomes important to know how probable these large fluctuations are, and how they evolve in time.

In what follows we will use a formalism specified by Ceccatto and Huberman (1989) that is well suited for studying fluctuations away from the equilibrium behavior of the system.[12] This

[12] A synopsis of their derivation is worked out in Section 4.5.4.

formalism relies on the existence of an optimality function, Ω, that can be constructed from knowledge of the density dependent utilities. The Ω function has the important property that its local minima are the equilibria of the system as well as the most probable configurations of the system. Depending on the complexity of the Ω function, there may be several equilibria, with the overall global minimum being the optimal state of the system.

Specifically, the equilibrium probability distribution $P_e(f^c)$ is given by

$$P_e = C \exp\left[-n\Omega(f^c)\right], \tag{4.12}$$

where the optimality function Ω for our model of ongoing collective action is given by

$$\Omega(f^c) = \int_0^{f^c} df' \left[f' - \rho^c(f')\right] \tag{4.13}$$

in terms of the mean probability $\rho^c(f^c)$ that cooperation is preferred. Thus, the optimal configuration corresponds to the value of f^c at which Ω reaches its global minimum.

Within this formalism it is easy to study the dynamics of fluctuations away from the minima. First, consider the case where there is a single equilibrium (which can be either cooperative or defecting). Fluctuations away from this state relax back exponentially fast to the equilibrium point, with a characteristic time of the order of $1/\alpha$, which is the average evaluation time for the individuals. Alternatively, there may be multiple equilibria, with the optimal state of the system given by the global minimum of the Ω function. A situation in which the system has two equilibria is schematically illustrated in Fig. 4.1.

If the system is initially in an equilibrium which corresponds to the global minimum (*e.g.*, state *A*), fluctuations away from this state will relax back exponentially fast to that state. But if the system is initially trapped in a metastable state (state *B*), the dynamics away from this state is both more complicated and interesting. As was shown in (Ceccatto and Huberman 1989), within a short time scale, fluctuations away from a local minimum relax back to it, but within a long time scale, a giant fluctuation can take place in which a large fraction of the agents switches strategies, pushing the system over the barrier maximum. Once there is the critical mass required for a giant fluctuation, the remaining agents rapidly switch to the new strategy and the system slides into the global equilibrium.

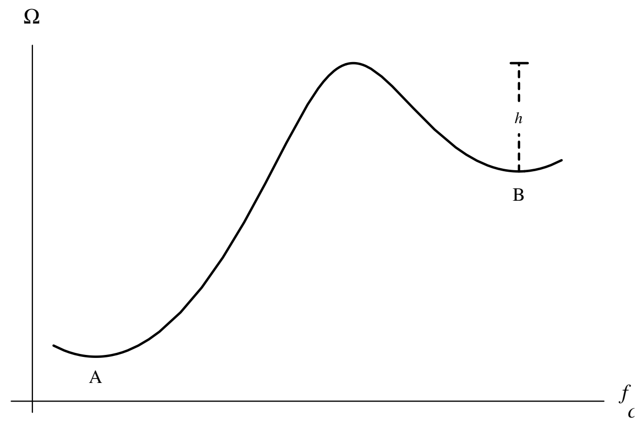

Figure 4.1: *The Ω function.* Schematic sketch of the optimality function Ω *vs.* f^c, the fraction of agents cooperation. The global minimum is at A, local minimum at B. h is the barrier height separating state B from A.

The time scales over which the nucleation of a giant fluctuation occurs is exponential in the number of agents. However, when such transitions take place, they do so very rapidly — the total time it takes for all agents to cross over is logarithmic in the number of agents. Since the logarithm of a large number is very small when compared to its exponential, the theory predicts that nothing much happens for a long time, but when the transition occurs, it does so very rapidly.

The process of escaping from the metastable state depends on the amount of imperfect knowledge that individuals have about the state of the system, in other words, on what other agents are doing. In the absence of imperfect knowledge the system would always stay in the local minimum downhill from the initial conditions, since small excursions away from it by a few agents would reduce their utility. It is only in the case of imperfect knowledge that many individuals can change their behavior. This is because, in evaluating the number of members cooperating, imperfect knowledge can lead to occasional large errors in the individual's estimation of the actual number cooperating.

Determining the time that it takes for the group to crossover to the global minimum is a calculation analogous to particle decay in a bistable potential and has been performed many times (*e.g.*, van Kampen 1981, Suzuki 1977) — the derivation is sketched out in section 4.5.4. The time, t, that it takes for a group of size n to cross over from a metastable Nash equilibrium to the optimal one is given by

$$t = \text{constant } e^{nh/\sigma}, \tag{4.14}$$

with h the height of the barrier as shown in Fig. 4.1 and σ a measure of the imperfectness of information. We should point out, however, that in our model the barrier height itself also depends on *n*, *H*, and *p*, making simple analytical estimates of the crossover time considerably more difficult.

4.3.2 Critical sizes for cooperation

The optimality function reveals a great deal about the dynamics of a system engaged in a collective action problem: it gives the possible equilibria and predicts the long-term stable state, *i.e.*, the state corresponding to the global minimum. It also shows that as the size of the group changes, the relative depths of Ω's minima change, leading to new optimal states.

From here on, we use the continuous approximation to the optimality function given by Eq. 4.13 since it permits clear qualitative description, but employ the exact discrete optimality function given in section 4.5.2 by Eq. 4.19 to derive quantitative results and to conduct computer simulations. Fig. 4.2(a) shows how the optimality function varies as the size of the group, *n*, increases. Eq. 4.6 implies that f^{crit} increases with increasing *n*. As a result, the value of f^{crit} passes from $f^{crit} < 0.5$ to $f^{crit} > 0.5$ as *n* increases. This indicates a transition in the dynamical nature of the interaction: as *n* increases, the interaction switches from having an optimal state of mutual cooperation to having an optimal state of mutual defection. These transition points are obtained numerically for $p < 1$ and derived exactly in 4.5.3 for $p = 1$.

As *p* decreases from 1, numerical analysis of $\Omega(f^c)$ shows that the shape of the optimality function is roughly preserved, but that the barrier height decreases with decreasing *p*. The change in the optimality function as *p* decreases is pictorially depicted in Fig. 4.2(b). The peak drifts away from $f^c = f^{crit}$ (except in the special case $f^{crit} = 0.5$), while the minima may move in from $f_c = 0$ and $f_c = 1$. Eventually, as *p* is decreased further, the barrier height goes to zero at some p_{crit} and only one minimum remains.

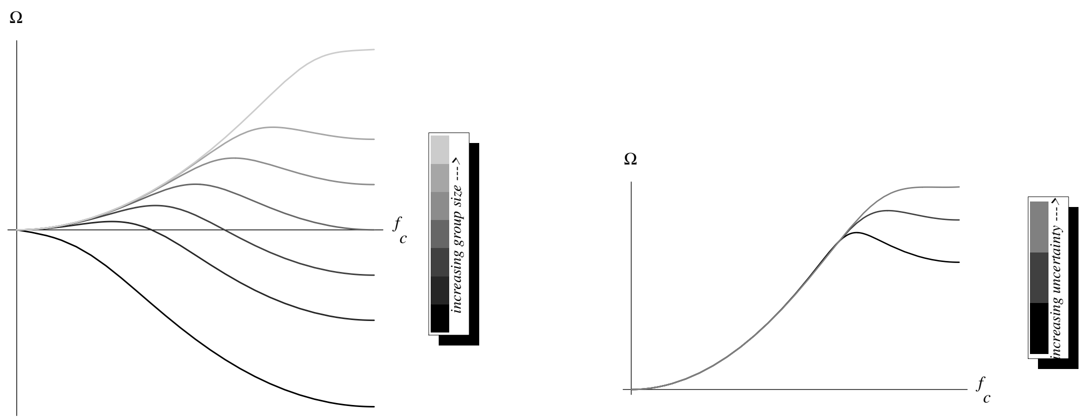

Figure 4.2: *Shifting equilibria.* Optimality function Ω *vs.* f^c, the fraction of agents of cooperating, is shown in (a) for increasing values of group size, *n* and in (b) for increasing uncertainty or decreasing values of *p*. (a) Lighter shades of gray correspond to larger values of *n*. Note the transition as the group size increases from an optimal state (corresponding to the global minimum) of mutual cooperation to an optimal state of mutual defection. Also, the height of the barrier, as measured from the maximum of $\Omega(f_c)$ to the metastable fixed point is greatest when the system has two fixed points of equal optimality. (b) Lighter shades of gray correspond to smaller values of *p*; in the lightest curve, the barrier height has gone to zero and only one equilibrium point is possible — mutual defection.

Until p reaches the critical value p_{crit} at which the barrier height goes to zero for all *n*, the nature of the system's equilibrium points remain similar to the $p = 1$ case. Specifically, for $p < p_{crit}$, three critical values can be obtained: (1) $n_{min}(p)$, the minimum group size below which cooperation is the only fixed point; (2) $\tilde{n}^*(p)$, the critical size below which cooperation is the optimal state; and (3) $n^*(p)$, an upper bound above which cooperation is not sustainable. The values $n_{min}(p)$, $\tilde{n}^*(p)$, and $n^*(p)$ can be determined numerically, demonstrating the emergence of four levels in group size corresponding to the very different resolutions of the collective action problem described in Table 4.1.

Numerical solutions to the discrete optimality function (Eq. 4.19) yield the values $n_{min}(p)$, $\tilde{n}^*(p)$, and $n^*(p)$, along with p_{crit}. In particular, these values were found for the case $H = 50$ (corresponds to discount rate $\delta = 1 - 1/H = 0.98$), $\alpha = 1$, $\tau = 1$, benefit to group of individual cooperation $b = 2.5$ and personal cost of cooperation $c = 1$. The resulting phase diagram delineating the regions of different resolutions to the conflict is shown in Fig. 4.3. At $p = 1$, $n^* = 77$, $\tilde{n}^* = 40$ and $n_{min} = 10$. When $p > p_{crit} = 0.59$, the uncertainty is

$n \leq n_{min}(p)$	one equilibrium point – mostly cooperative
$n_{min}(p) < n \leq \tilde{n}^*(p)$	mostly cooperative optimal state, mostly defecting metastable state
$\tilde{n}^*(p) < n < n^*(p)$	mostly defecting optimal state, mostly cooperative metastable state
$n \geq n^*(p)$	one equilibrium point – mostly defecting

Table 4.1 *Critical sizes for cooperation.*

high enough that all structure in the optimality function is washed out and only one equilibrium point exists for groups of all sizes. For groups of size less than 40 operating with such a high level of uncertainty, the interaction evolves to a mixed equilibrium that is more cooperative than defective. The balance reverses itself when the group size exceeds 40.

The diagram in Fig. 4.3 shows that for the symmetric case $p = q$, $n^*(p)$ is a decreasing function of p while $n_{min}(p)$ is an increasing function of p and $\tilde{n}^*(p)$ is a constant. In general $(q \neq p)$, $n^*(p)$ remains a decreasing function; however, the functional forms of $\tilde{n}^*(p)$ and $n_{min}(p)$ depend on q. For example, $q = 0$ yields decreasing functions $\tilde{n}^*(p)$ and $n_{min}(p)$.

4.3.3 Experimental results

We next present representative results compiled from a number of event-driven Monte Carlo simulations that we used to both test the analytical predictions of the theory and to calibrate the time constants with which cooperation and defection appear in a system of given size. These simulations were conducted in asynchronous fashion, with uncertainty entering through binomially distributed choices by the agents. The choice criterion for each individual agent was given by Eq. 4.6. Statistics were collected for a number of group sizes, horizon lengths and uncertainty levels.

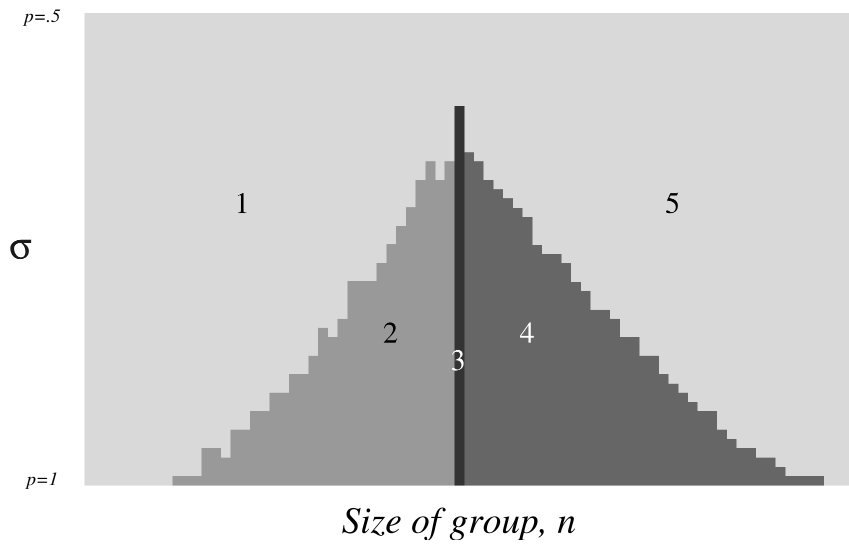

Figure 4.3: *Behavioral phase diagram.* Diagram delineates regions corresponding to different resolutions of collective action problem for parameter values $H = 50$ and $\alpha = 1$. The amount of error, σ, increases vertically, and the size of the group increases horizontally. In region 1, $n \leq n_{min}(p)$, there is one equilibrium point – mostly cooperative; in region 2, $n_{min}(p) < n < \tilde{n}^*(p)$, mostly cooperative is the optimal state, while mostly defecting is a metastable state; in region 3, $n = \tilde{n}^*(p)$,the system is bistable with mostly cooperative and mostly defecting both optimal states; in region 4, $\tilde{n}^*(p) < n < n^*(p)$, mostly defecting is the optimal state, while mostly cooperative is a metastable state; and in region 5, $n^*(p) \leq n$, there is again only one equilibrium point – mostly defecting. There is no sharp boundary between regions 1 and 5 for high levels of uncertainty. Note that region 3 actually has zero width.

As shown in the previous section, our model of the collective action problem, for a large range of group sizes, can yield dynamics with two attractors, one an optimal equilibrium, the other a metastable equilibrium (*e.g.*, regions 2 and 4 of Fig. 4.3). If the group initially finds itself at or near the metastable equilibrium, it may be trapped there for very long times, making it effectively a stable state when the time scale of the interaction is shorter than the crossover time to the optimal configuration. This crossover time, *t*, given in Eq. 4.14, is exponential in the size *n* of the group, in the height of the barrier, *h*, and in the inverse uncertainty, $1/\sigma$. Note that the barrier height is also a function of the uncertainty, σ, and indirectly of *n* and the horizon length *H*, so the functional dependencies for the crossover time are more complex than appears at first sight.

The persistence of non-optimal cooperation Nevertheless, as verified through simulations of groups of individuals engaged in collective action problems, crossover times are exponentially long. As a result, a large group which has an initial tendency to cooperate may remain in a cooperative state for a very long time even though the optimal state is defection. (All groups whose size places them in region 4 of Fig. 4.3 could display this phenomenon, for example). Conversely, small groups whose initial tendency is to defect can persist in non-optimal defection, until a large fluctuation finally takes the group over to cooperation.

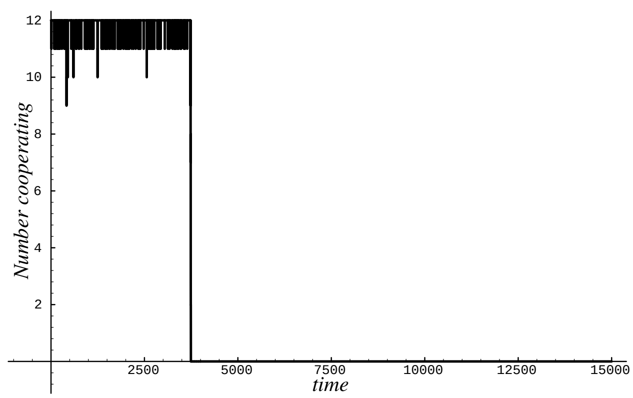

Figure 4.4: *Outbreak of defection.* At $t = 0$, two cooperating groups of size $n = 6$ merge to form a larger, cooperating group of size $n = 12$. All agents have horizon length $H = 9.5$ throughout, with $p = 0.93$, $b = 2.5$, $c = 1$, $\alpha = 1$ and $\tau = 1$. For these parameters, cooperation is the optimal state for a group of size $n = 6$, but for the combined group of size $n = 12$, cooperation is metastable. Indeed, as the figure shows, metastable cooperation persists for almost 4,000 time steps in this example. Uncertainty (p less than one) ensures that eventually a large fluctuation in the perceived number of agents cooperating eventually takes the group over into a state of mutual defection, which is optimal.

As a concrete example, consider two small cooperating groups of size $n = 6$, with horizon length $H = 9.5$, for which the optimal state is cooperation. At $t = 0$, the groups merge to form a larger, cooperating group of size $n = 12$. For the larger group, cooperation is now a metastable state: no one individual will find it to its benefit to defect and the metastable cooperative state can be maintained for very long times, especially if p is close to 1. As shown in Fig. 4.4, in one case mutual cooperation lasts for about 4,000 time steps, until a sudden transition (of duration proportional to the logarithm of the size of the group) to mutual defection occurs, from which the system will almost never recover (the time scale of recovery is many orders of magnitude larger than the crossover time). In this example, $p = 0.93$. If the amount of error increases so that p now equals 0.91 say (thus reducing the height of barrier between cooperation and defection by 21%), the crossover to defection occurs on the order of hundreds of time steps, instead of thousands.

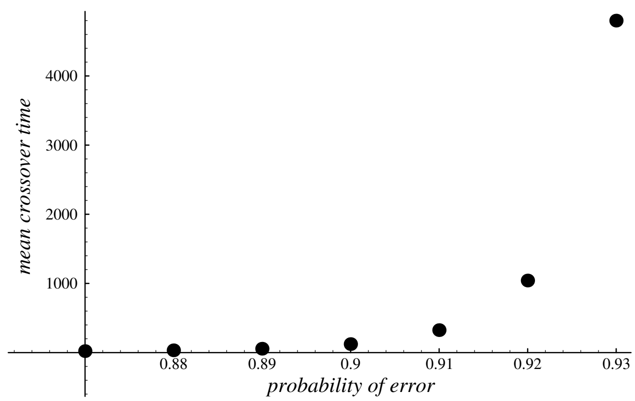

Figure 4.5: *Exponential increase of crossover time* as a function of the probability p that an agent's action is misperceived. The mean crossover time was obtained by averaging over 600 simulations at each value of the uncertainty parameter p. In the all the runs, $n = 12$, $H = 9.5$, $b = 2.5$, $c = 1$, $\alpha = 1$ and $\tau = 1$.

Fig. 4.5 further demonstrates the exponential dependence of the mean crossover time on the probability p that an agent's action is misperceived. For this example, the crossover time from metastable cooperation to overall defection was averaged over 600 simulations at each value of the uncertainty parameter p. Fitting an exponential to the data indicates that the mean crossover time goes as $t \propto e^{f(p)}$, where $f(p)$ is quadratic in p. However, the more relevant relationship is between the mean crossover time and the height of the barrier of the Ω function. The mean crossover time t appears to be given by

$$t = \text{constant } e^{h(n,\sigma)}, \tag{4.15}$$

where the barrier height depends on the group size, n, and the uncertainty, σ. For the example in Fig. 4.5, the barrier height depends quadratically on the parameter, p. The dependence of barrier height on group size is less straightforward, but apart from some outliers, the barrier height also increases quadratically in group size.

Flip-flop strategies We now point out a rather curious phenomena that is observed when the group size is such that the optimality function has two equilibria with nearly the same optimality (that is, those groups whose sizes place them within the narrow interval containing region 3 of the diagram in Fig. 4.3). Thus, the system is bistable with no optimal or preferred equilibrium state. On the whole, then, the group will spend half its time defecting and half its time cooperating, *i.e.*, the group tendency flip-flops between mutual cooperation and mutual defection. This is to be contrasted with a mixed strategy in which half the group cooperates all the time and half the group defects all the time. On the whole, a typical measurement of the level of cooperation would find the group either totally defecting or cooperating. This does not reflect the average behavior, which would correspond to a group made up of half of its agents cooperating and half defecting. In this case, then, typical is not average. These flip-flop strategies are displayed in Fig. 4.6.

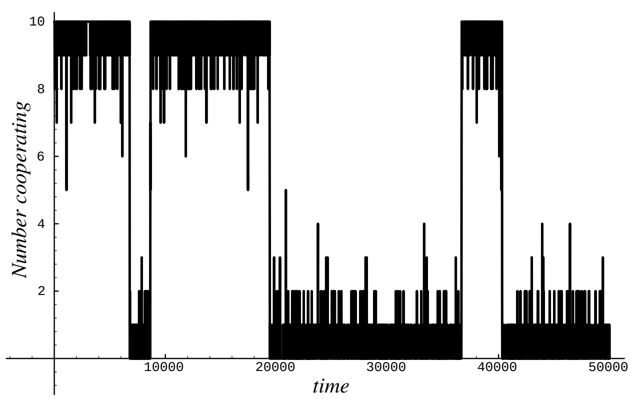

Figure 4.6: *Flip-flops.* The group of size $n = 10$ flip-flops between the two strategies of mutual cooperation and mutual defection. In this example, the horizon $H = 10$, with $p = 0.8$, $b = 2.5$, $c = 1$, $\alpha = 1$ and $\tau = 1$.

4.3.4 Sharp transitions

The phase diagram of Fig. 4.3 showing critical group sizes at various uncertainty levels was computed for a fixed value of the horizon length. As a result, the previous analysis masks the sensitive dependence of the model's long term behavior on the discount parameter $\delta = 1 - 1/H$. (The discount parameter can be thought of as the probability that the interaction continues into the next time step. A discount of 1 corresponds to an infinite horizon.) In fact, the system undergoes a sharp transition at a critical value of the discount parameter, δ_{crit}; that is, small variations in δ near this critical value yield opposite extremes in the long term cooperative behavior of the interaction. The transition is shown in Fig. 4.7 in which the relative probability of mutual cooperation *vs.* mutual defection on the long-term is plotted against the discount δ. The relative probability is calculated using Eqs. 4.12 and 4.19, giving

$$\frac{P_e(f^c = 1)}{P_e(f^c = 0)} = e^{-n[\Omega(f^c=1) - \Omega(f^c=0)]}. \tag{4.16}$$

In the example of Fig. 4.7 for which $p = 0.9$, $n = 50$, then $\delta_{crit} = 0.98$ (corresponding to the horizon length $H = 50$). Note that the vertical axis is in units of 10^{22} showing that the abrupt transition near $\delta = 0.98$ is all or nothing.

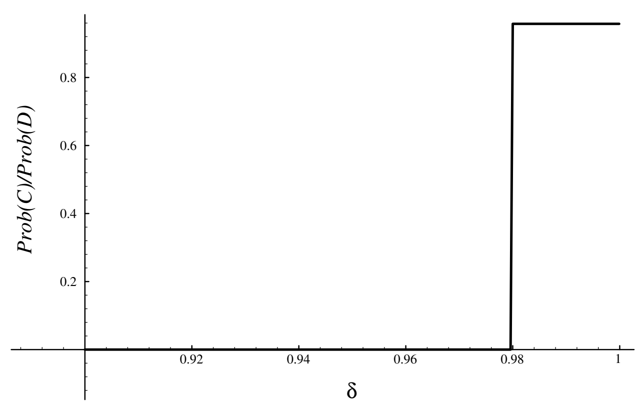

Figure 4.7: *Sharp transition in levels of cooperation.* Relative probability of long-term cooperation *vs.* defection, $P_e(f^c = 1)/P_e(f^c = 0) = e^{-n[\Omega(f^c=1) - \Omega(f^c=0)]}$, is plotted against the discount parameter δ for parameter values $p = 0.9$ and $n = 50$. The vertical axis is in units of 10^{22}. Sharp transition from long-term mutual defection mutual cooperation occurs at $\delta = 0.98$.

The critical value of the discount parameter at which the abrupt transition from long-term defection to long-term cooperation occurs is given in general by solving for the value of δ at which the optimality function describing the interaction is bistable. Thus, solving first for the horizon *H* in the relation $0.5 = \frac{1}{\alpha H_{crit}} \frac{nc-b}{b-c}$ (by setting $f_{crit} = 0.5$, the criterion for metastability), we find

$$\delta_{crit} = 1 - \frac{\alpha}{2} \frac{b-c}{nc-b}. \tag{4.17}$$

Interestingly, while defection is the long-term optimal state for $\delta < \delta_{crit}$, cooperation is still a metastable state. The reverse holds for $\delta > \delta_{crit}$: cooperation is optimal while defection is metastable. As a result, although the transition from long-term defection to long-term cooperation is very abrupt as δ increases through the value δ_{crit}, the time it takes for the transition to actually be observed (for the group interaction to go over to cooperation) is exponentially long!

4.4 Summary

In this chapter we presented and studied a model of ongoing collective action among intentional agents which make choices that depend not only on the past but also on their expectations. This model maps collective action problems onto an iterated *n*-person prisoner's dilemma in an uncertain world. Individuals interact dynamically, making decisions based both on individual preferences and on their expectations as to how their choices will affect other agents in the future. This is a significant departure from previous studies of collective action problems in which the finiteness of the interaction is the only shadow cast by the future.

Using a dynamical formulation of these interactions, we derived the magnitudes of group sizes that are capable of supporting cooperation and established the existence of regimes with multiple equilibria. Extensive computer experiments also allowed us to confirm the analytically derived behavioral diagrams whose distinct regions correspond to different resolutions of the collective action problem.

More significantly, we also studied stochastic fluctuations in the number of individuals cooperating or defecting using a thermodynamic-like formalism. This allowed us to discover the existence of metastable states in which collaboration can persist in groups whose sizes are too big to support it indefinitely. Moreover, the onset of overall defection takes place in a sudden and unexpected way. Likewise, cooperation can appear out of nowhere in small, previously defecting, groups. These outbreaks, which were also confirmed by computer experiments, mark the end of long transient states in which defection or cooperation persists in groups that cannot sustain it indefinitely. Since these discontinuities are not overly sensitive to the analytical form of the expectations, we expect that they might be observed in more complex social settings. An interesting extension would be to examine how diversity within the group affects our results, a direction which we pursue in the next chapter.

4.5 Derivations

4.5.1 Probability of cooperation

To evaluate the probability $P_c(f^c(t-\tau))$ that an individual perceives cooperation as maximizing its expected future accumulated utility, first define $\bar{f}^c(t-\tau)$ to be the fraction of successfully cooperating members at time $t-\tau$ and $1-\bar{f}^d(t-\tau)$ to be the fraction unsuccessfully defecting (*i.e.*, effectively cooperating). This resolves $\widehat{f}^c(t-\tau)$ into its two components: successful cooperators and unsuccessful defectors. Then, applying the criterion of Eq. 4.6, $P_c(f^c(t-\tau))$ is simply the sum of the probability that $\bar{f}^c(t-\tau) > f^{crit}$, the probability that $\bar{f}^c(t-\tau) + \left(1-\bar{f}^d(t-\tau)\right) > f^{crit}$ given $\bar{f}^c(t-\tau) \leq f^{crit}$, and one-half the probability that $\bar{f}^c(t-\tau) + \left(1-\bar{f}^d(t-\tau)\right) = f^{crit}$. Using the Binomial Theorem, we obtain the following expression for $P_c(f^c(t-\tau))$:

$$
\begin{aligned}
P_c(f^c(t-\tau)) = & \sum_{k=n^{crit}+1}^{n^c} \binom{n^c}{k} p^k (1-p)^{n^c-k} \\
& + \sum_{i=1}^{n-n^c} \left[\binom{n^c}{n^{crit}+1-i} p^{n^{crit}+1-i} (1-p)^{n^c-n^{crit}-1+i} \right] * \\
& \left[\sum_{l=0}^{n-n^c-i} \binom{n-n^c}{l} q^l (1-q)^{n-n^c-l} \right] \\
& + \frac{1}{2} \sum_{i=0}^{n-n^c} \binom{n^c}{n^{crit}-i} p^{n^{crit}-i} (1-p)^{n^c-n^{crit}+i} \binom{n-n^c}{i} (1-q)^i q^{n-n^c-i} \ ,
\end{aligned}
\tag{4.18}
$$

where $n^c = nf^c(t-\tau)$ and $n^{crit} = nf^{crit}$.

4.5.2 Discrete Ω function

The exact discrete optimality function for the stochastic model of ongoing collective action is

$$\Omega\left(f^c = \frac{n^c}{n}\right) = \frac{1}{n}\sum_{n_i=0}^{nf^c}\left[\frac{n_i}{n} - P_c\left(\frac{n^c}{n}\right)\right], \tag{4.19}$$

with $P_c(f^c)$ derived in section 4.5.1 and given by Eq. 4.18. Numerical comparison of the continuous and discrete versions of the optimality functions show that the qualitative features of the two are very similar.

4.5.3 Critical sizes and barrier height

For the case of zero error, $p = 1$, exact derivations of the critical sizes n_{min}, $\tilde{n}^*$, and n^* can be obtained using Eq. 4.6, the criterion for cooperation, and an approximate relation for the height of the barrier can be derived using the continuous version of the optimality function, Eq. 4.13.

In the limit of $p = 1$, the critical sizes $n_{min}(p = 1)$, $\tilde{n}^*(p = 1)$ and $n^*(p = 1)$ can be derived analytically. When $f^{crit} > 0.5$, mutual cooperation is a metastable configuration and defection is the optimal state, while when $f^{crit} < 0.5$, the reverse holds. Also, when $f^{crit} = 0.5$, the system is bistable with two optimal states at mutual cooperation and mutual defection. However, zero error means that the barrier of height *h* given by Eq. 4.23, is insurmountable in finite time. As a result, the evolution of the interaction depends completely on the initial actions of the group members. If the initial fraction cooperating is greater than f^{crit}, the group evolves exponentially rapidly to a state of mutual cooperation, while if the fraction cooperating is smaller than f^{crit}, the group is bound for mutual defection.

Using Eq. 4.6, we can transform the above condition $f^{crit} > 0.5$ to obtain the critical size below which cooperation is the optimal state,

$$\tilde{n}^*(p = 1) = \frac{1}{2}H\alpha\left(\frac{b}{c} - 1\right) + \frac{b}{c} \tag{4.20}$$

(rounding down to the nearest integer). An upper bound n^* above which cooperation is not sustainable (defection is the only fixed point) is found by solving for the condition $f^{crit} > 1$:

$$n^*(p=1) = H\alpha\left(\frac{b}{c} - 1\right) + \frac{b}{c} \tag{4.21}$$

(rounding up). So for *n* such that $\tilde{n}^* < n < n^*$, mutual defection is still the optimal state, but mutual cooperation is a metastable configuration. The minimum group size below which cooperation is the only fixed point, n_{min}, is given by the solution to $f_{crit} = 1/n$:

$$n_{min}(p=1) = \frac{b}{2c} + \frac{1}{2c}\sqrt{b^2 + 4H\alpha c(b-c)}, \tag{4.22}$$

rounding down. For $n \leq n_{min}$, the group always evolves to a state of mutual cooperation, and for $n_{min} < n \leq \tilde{n}^*$, mutual cooperation is the optimal state, while mutual defection is metastable.

The height, *h*, of the barrier, is measured from the maximum of $\Omega(f^c)$ to the metastable minimum. For $p = 1$, the height of the barrier (derived using Eq. 4.13) is given by

$$h(p=1) = \Omega\big(f^c = f^{crit}\big) - \Omega(f^c = 1) = \frac{1}{2}\big(f^{crit} - 1\big)^2. \tag{4.23}$$

As *p* decreases from 1, the height of the barrier decreases. This is only an approximate relation since the mean probability $\rho^c(f^c)$ gives only a rough description of $P_c(f^c)$ at $p = 1$.

4.5.4 Ω function formalism

The optimality function formalism of section 4.3.1 is developed in mathematical detail below, replicating the derivation in (Ceccatto and Huberman 1989) supplemented by details from van Kampen (1981). In addition, the expression for the crossover time given in Eq. 4.14 is re-derived.

First, we consider $g(\beta, \mu)$, the derivative of Ω with respect to μ ($g = d\Omega/d\mu$), where the variable $\mu = 2n_c/n - 1$ is the deviation from an even choice among the two strategies of cooperation and defection and β is a measure of the amount of imperfect knowledge ($\beta \propto 1/\sigma$). The fixed points of the dynamics are then determined by the condition

$$g(\beta, \mu) = 0, \tag{4.24}$$

with $g'(\beta, \mu) > 0$. In order to make use of this, we need an explicit form for the dependence of the function $\Omega(\beta,\mu)$ in terms of the difference in payoff between cooperating and defecting, $\Delta G = \langle \widehat{f_c} \rangle - f_{crit}$. Let us denote by $\eta_1(\beta, \mu)$ and $\eta_2(\beta, \mu)$ the probabilities that an agent perceives cooperation as better than defection and *vice-versa*. In terms of ΔG, they can be written as

$$\begin{gathered} \eta_1(\beta, \mu) = D \exp\left[\beta \Delta G\right] \\ \text{and} \\ \eta_2(\beta, \mu) = D \exp\left[-\beta \Delta G\right], \end{gathered} \tag{4.25}$$

where *D* is a small proportionality constant.[13] The connection between the function *g* and payoff difference ΔG is given by

$$g(\beta, \mu) = (1 + \mu)\eta_2(\beta, \mu) - (1 - \mu)\eta_2(\beta, \mu). \tag{4.26}$$

In order to demonstrate that the zeroes of $g(\beta, \mu)$ do indeed generate the stable equilibria, we need to consider the dynamics of the system, which is governed by an equation that determines how the probability of having n_c agents cooperating at time *t*, $P_t(n_c)$, evolves in time. It is given by the master equation

$$P_{t+1}(n_c) = \sum_{m_c = n_c - 1}^{n_c + 1} P_t(m_c | n_c) T(m_c | n_c), \tag{4.27}$$

where the transition matrices, $T(m_c|n_c)$, which determine the rate at which agents switch between cooperation and defection are given in (Huberman and Hogg 1988):

$$\begin{aligned} T_\beta(n_c | n_c + 1) &= \frac{\alpha}{n}(n - n_c)\eta_1(\beta, \mu); \\ T_\beta(n_c + 1 | n_c) &= \frac{\alpha}{n}(n_c + 1)\eta_2(\beta, \mu). \end{aligned} \tag{4.28}$$

The prefactor α/n is the probability per unit time that an agent reevaluates its preferences.

13 This assumption gives the same results as those obtained with the error function in section 4.2.3.

We will now show that the equilibrium probability distribution, $P_e(\mu)$, defined by having $P_{t+1} = P_t = P_e$, is given by

$$P_e(\mu) = C \exp\left[-n\Omega(\beta,\mu)\right], \tag{4.29}$$

with Ω the optimality function defined by $\Omega = \int_{-1}^{\mu} g(\beta,\mu')d\mu'$ and *C* a normalization constant. It can be established that the condition $dP_e(\mu)/d\mu = 0$ is equivalent to Eq. 4.24, with the function $g(\beta,\mu)$ determined by Eq. 4.26. That is, the probability distribution has sharp maxima at the zeroes of $g(\beta,\mu)$. This can be seen from the stationary solution of the master equation, Eq. 4.27, which satisfies the equality

$$P_e(n_c+1)T_\beta(n_c+1|n_c) = P_e(n_c)T_\beta(n_c|n_c+1). \tag{4.30}$$

The solution of Eq. 4.30 is, up to a normalization factor, given by

$$P_e(n_c) = \binom{N}{n_c} \prod_{r=1}^{n_c} \frac{\eta_1(\beta, 2(r-1)/n-1)}{\eta_2(\beta, 2(r-1)/n-1)}, \tag{4.31}$$

which, by taking logarithms and using Stirling's approximation, can be written (to leading order in *n*) in the same form as Eq. 4.29.

Now that we have the equilibrium probability distribution, we can calculate the crossover time. The dynamics of the crossover from a local to a global minimum separated by a local maximum can be modeled as a diffusion process in an external potential. In one dimension this can be written as a special case of the Fokker-Planck equation known as the Smoluchowski equation

$$\frac{\partial P}{\partial t} = \frac{\partial}{\partial \mu}\Omega'(\mu)P + \frac{1}{\beta}\frac{\partial^2 P}{\partial \mu^2} \tag{4.32}$$

where Ω plays the role of a potential, *P* is the probability that a particle is at position μ at time *t*, and β is, as before, the measure of imperfect information. This equation can be solved for the crossover time by using a Laplace transform.

In particular, suppose that the global minimum is at $\mu = a$, the local minimum at $\mu = b$, and the local maximum at $\mu = c$. Then the crossover time from point *b* to the global minimum is related to the probability distribution of the equilibrium state by

$$t_{ab} \propto \int_{b}^{a} \frac{d\mu}{P_e(\mu)}. \tag{4.33}$$

This integral can be approximated when β is small or when the potential barrier is high enough that $P_e(\mu)$ is very small at the instability point, c, by making a parabolic expansion of $P_e(\mu)$ about the instability point. The result yields the expression given for the crossover time in Eq. 4.14:

$$t_{ab} \propto e^{\beta N h}, \tag{4.34}$$

where $h = \Omega(c) - \Omega(b)$ is the difference in the height of the optimality function between the local maximum and the local minimum.

5 Diversity

5.1 Introduction

An essential element contributing to the likelihood of collective action in human societies is that individuals differ from one other. Diversity comes into play in many instances of collective action. Occupation, wealth, age, gender, race, creed all vary from individual to individual and affect each person's valuation of a collective good. Liberty is most valued by those aware of their enslavement; a sustainable ecology and a balanced budget is most desired by those who care about the world our children and grandchildren will inherit.

To the extent that each person's behavior can be approximated by that of an average agent, the analysis of the previous chapter is correct. A first correction to the model involves accounting for diversity within the group. For example, the variety in individual beliefs and expectations can be described by a spread in horizon lengths: some agents look far ahead when planning; others are relatively myopic. Alternatively (or additionally), the benefits and costs of contributing to the collective good may depend on the individual. We showed that individuals in a homogenous group behave as if they are acting out a strategy of conditional cooperation. Thus, the criterion for conditional cooperation in diverse groups can be generalized as follows: whereas one individual might choose to participate in the group effort when the number of cooperators is fairly small, another might wait until there are many cooperators.

How will a diversity of expectations affect the results obtained in the previous chapter? In order to answer this question, we extend the dynamical model of ongoing collective action among intentional agents presented in the previous chapter to account for diversity. Work done by Oliver and Marwell (1988) and Bendor and Mookherjee (1987) indicate that heterogeneous groups should be better able to obtain cooperation. We find that the equilibrium conditions of a homogenous group still hold for a diverse one when the average behavior of the latter is identical to the uniform behavior of the former. However, because the dynamics of reaching the global equilibrium changes, the transitions to cooperation (or defection) speed up and new pathways can unfold.

In Section 5.2, we describe two distinct ways in which diversity can be modeled in the context of collective action. The first type reflects a unimodal distribution of diverse beliefs; the second a multimodal distribution. In Section 5.3, we study the effect of diversity among individuals on the phenomena described in the previous chapter. We find that while diversity does not alter the critical sizes for cooperation, it acts a source of additional uncertainty and shortens the transition time to the long-term stable state of the system. In addition, an analytical form for the increase in uncertainty due to diversity is given. Again, the results of computer simulations confirm our predictions. Lastly, these simulations also show that when there is diversity of beliefs among subgroups contained within a large group, cooperation can be achieved in stages.

5.2 Modeling Diversity

In the first pass at studying collective action problems, we treated all individuals as identical. We now drop that assumption and consider how diversity enters into the model and how behavior changes as a result. As we discuss below, two qualitatively different forms of diversity must be studied. The first type reflects diversity in groups of agents whose differences in beliefs can be captured by a simple spread about some common belief. Thus, the individuals on the whole are similar, but have differences that capture variability in preferences and potentially other additional factors. For example, each group member might be said to have horizon length $H = 10 \pm 2$, instead of $H = 10$ exactly.

On the other hand, the second type of diversity represents differences within a group that cannot be accounted for by a simple variance about an average value. Instead the group acts as the union of several subgroups each best characterized by its own set of beliefs. In this case, one subgroup might have horizon length $H = 2$, another $H = 4$, and yet another $H = 6$.

Of course, there is no line strictly dividing the two types of diversity. Rather, as the amount of environmental noise increases (as parametrized by p), the differences between subgroups are washed out and effectively appear to be simply a large spread about a common belief.

Alternatively, the spread in each of the modes may be large enough that it is not possible to distinguish among them.

A general way of incorporating diversity is to allow the critical fraction, f^{crit}, in the criterion for cooperation (Eq. 4.6) to vary from individual to individual. Thus, each member *i* has a bias b_i, and decides whether or not cooperate based on the revised condition

$$f^{crit} + b_i < \widehat{f}^c(t - \tau). \tag{5.1}$$

The notation we use to denote the biased value of the critical fraction will be $f_i^{crit} \equiv f^{crit} + b_i$.

In addition, it can be easily shown that as long as $b_i \ll f^{crit}$, $\forall i$, then the distribution of biases $\{b_i\}$ can also be used to represent variable horizon lengths or differing benefit-cost ratios b/c. Thus, describing diversity by adding a bias to the criterion for cooperation of Eq. 4.6 is a more general description of diversity than might appear at first glance.

5.3 Analysis and Behavior

5.3.1 Diversity as a form of uncertainty

We examine first the case in which the group's diversity can be modeled by a Gaussian distribution. If we take the biases $\{b_i\}$ to be distributed normally with mean zero, then the critical size beyond which cooperation cannot be maintained remains the same as when there is no diversity. Distributions with non-zero mean would alter the nature of the equilibrium points; so zero mean distributions offer the best means to isolate the effect of adding diversity.

In order to understand the qualitative effect of introducing diversity in this manner, we write the mean probability $\rho_i^c(f^c)$ that member *i* will choose to cooperate as

$$\rho_i^c(f^c) = \frac{1}{2}\left\{1 + erf\left[\frac{\langle\widehat{f}^c\rangle - \left(f^{crit} + b_i\right)}{\sqrt{2}\sigma}\right]\right\}, \tag{5.2}$$

where $\langle \widehat{f^c} \rangle = pf^c + (1-p)(1-f^c)$ and $\sigma = \sqrt{p(1-p)/n}$. The dynamical equation describing the evolution of the system then becomes

$$\frac{df^c}{dt} = -\alpha \left[f^c(t) - \frac{1}{n} \sum_{i=1}^{n} \rho_i^c(f^c(t-\tau)) \right]. \tag{5.3}$$

Letting

$$\tilde{\rho}^c(f^c) = \frac{1}{n} \sum_{i=1}^{n} \rho_i^c(f^c(t-\tau)) \tag{5.4}$$

and linearizing about $\langle \widehat{f^c} \rangle = f^{crit}$, we obtain

$$\tilde{\rho}^c(f^c) = \frac{1}{2} \left\{ 1 + erf \left[\frac{\langle \widehat{f^c} \rangle - f^{crit}}{\sqrt{2}\tilde{\sigma}} \right] \right\}, \tag{5.5}$$

with the renormalized value of uncertainty $\tilde{\sigma}$ given by

$$\tilde{\sigma} \rightarrow \sqrt{\sigma^2 + \sigma'^2}. \tag{5.6}$$

This approximation assumes that the $\{b_i\}$ are distributed normally with mean zero and standard deviation σ'.

Thus, taking this first form of diversity into account simply renormalizes the amount of noise in the system as parametrized by σ in the denominator of the error function. Note, however, that imperfect information also enters in the calculation of the expected value of f^c: $\langle \widehat{f^c} \rangle = pf^c + (1-p)(1-f^c)$. Consequently, adding diversity is not identical to decreasing p away from 1.

From this analysis, it appears that diversity among the agents will shorten the lifetime of the metastable states described in the previous section, while the transitions remain abrupt. Computer experiments verify this prediction, with a sample simulation given in Fig. 5.1. Diversity is modeled as a spread about f^{crit} among the agents; thus, agents differ from one another in their likelihood of cooperating *vs.* defecting. In this example, the individual biases to f^{crit} are distributed normally with standard deviation equal to the amount of imperfect information $\left(\sigma' \equiv \sigma = \sqrt{p(1-p)/n}\right)$. As a result of the diversity among the individuals in the group, the average crossover time becomes about 1,300 time steps, ranging from fewer than 100 time steps to over 2,000. Without diversity, the average crossover time is about 5,000 time steps (see, for example, Fig. 4.4 — note in particular the difference in time scales), ranging from less than 1,000 time steps to over 10,000.

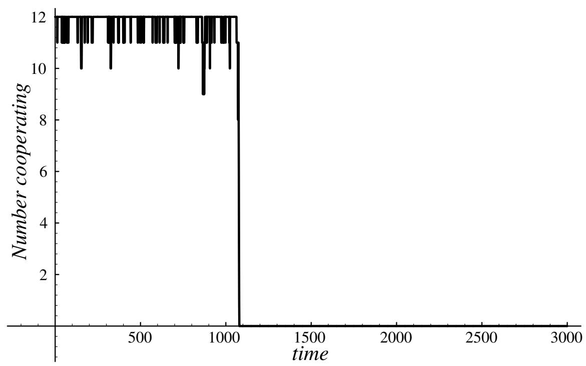

Figure 5.1: *Diversity as noise.* Group of size $n = 12$, with $H = 9.5$, $p = 0.93$, $b = 2.5$, $c = 1$, $\alpha = 1$ and $\tau = 1$ with diversity. Diversity is modeled by adding individual biases to the value of $f^{crit} = 0.6667$. These biases are distributed normally with mean zero and standard deviation $\sigma' \equiv \sigma = \sqrt{p(1-p)/n}$ (the "noise" due to diversity is set equal to the noise due to uncertainty). The moderate amount of diversity is responsible for a transition rate to defection about four times as fast as that for a uniform group (see, for example, Fig. 4.4 — note the difference in time scales).

5.3.2 Stages to cooperation

We next consider the second form of diversity which represents large differences in beliefs between various subgroups within a group. This type of diversity follows a multimodal distribution and does not lend itself to the renormalization of uncertainty performed above. Instead, the summed error function in Eq. 5.3 exhibits a sequence of steps. For example, if in a group of size 12 there is one subgroup biased towards cooperation with $f_1^{crit} = 0.1333$, another with $f_2^{crit} = 0.4667$ and yet another biased towards defection with $f_3^{crit} = 0.8$, then there are three steps in the summed error function of Eq. 5.4. Each step represents the likelihood that an individual from each respective subgroup will find it worthwhile to cooperate. There can now be multiple local minima in which some of the subgroups sustain cooperation while the others do not. As always, the global minimum is the long-term preferred state of the system.

In the example given in the preceding paragraph, overall cooperation can be achieved in stepwise fashion. If the group begins in an uncooperative state, the subgroup with $f_1^{crit} = 0.1333$ is the first to make the transition to cooperation, followed by the subgroup with $f_2^{crit} = 0.4667$ and finally by the subgroup with $f_3^{crit} = 0.8$. These steps can be seen in Fig. 5.2.

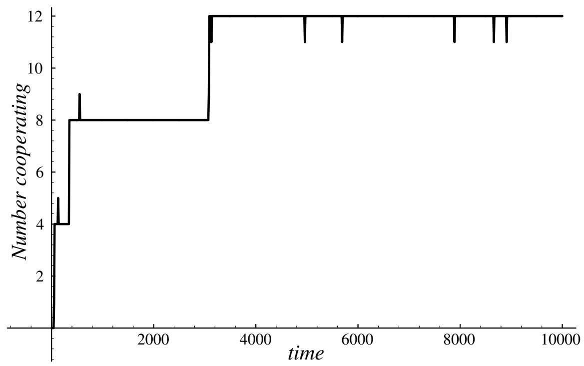

Figure 5.2: *Stages towards cooperation.* A group of size 12 consists of three subgroups. The first is greatly biased towards cooperation with $f_1^{crit} = 0.1333$. The second subgroup has $f_2^{crit} = 0.4667$, while the third is biased towards defection with $f_{crit}^{b_3} = 0.8.$. The figure shows the stepwise transition to mutual cooperation: each subgroup makes the transition to cooperation in turn. Parameter values are $H = 12$, $p = 0.98$, $b = 2.5$, $c = 1$, $\alpha = 1$ and $\tau = 1$.

Compare this example to the uniform case in which all individuals have $f^{crit} = 0.4667$ (the average critical fraction for the diverse group above). This system is now almost bistable, with both cooperation and defection about equally likely to be the long-term stable state, although cooperation is more likely. In this scenario, then, diversity greatly increases the likelihood of attaining a cooperative state. Other scenarios can be envisioned, however, which accomplish the reverse, making defection more likely with diversity than without.

Another interesting consequence of the effect of multimodal diversity is the following. Imagine combining two groups, each composed of 6 individuals. The first group has beliefs such that all members have $f_1^{crit} = 0.3667$. Thus, the long-term stable state for this group is one of cooperation. The second group, on the other hand, has beliefs such that $f_2^{crit} = 0.6667$, and thus tends towards mutual defection. Putting these groups together, then, we have an initial state with 6 of the 12 individuals cooperating. This turns out to be a metastable state of the combined system. Eventually, however, the group undergoes a transition to either complete cooperation or complete defection. Both are equally likely — the two possibilities are illustrated in Fig. 5.3. Thus, either the defecting agents are persuaded to cooperate (although they generally wouldn't when in a group to themselves even though the size of this subgroup is smaller); or the originally cooperating agents lose their incentive to do so and defect.

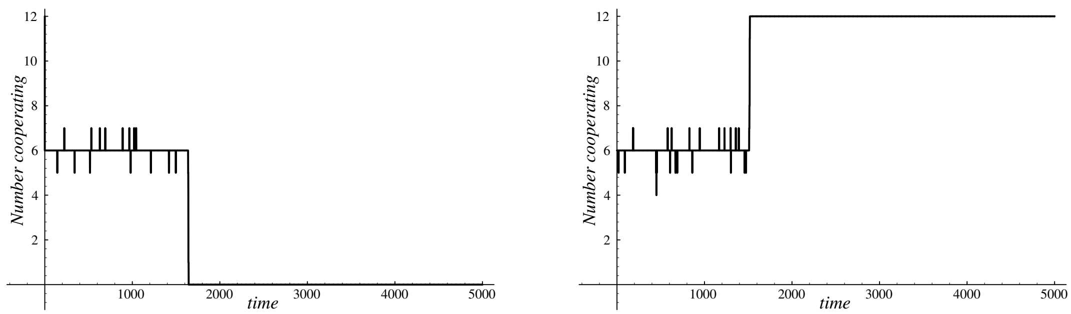

Figure 5.3: *Step down, step up.* Two groups of size 6 are joined together, one with a bias towards cooperation ($f_1^{crit} = 0.3667$), the second with a tendency to defect. ($f_2^{crit} = 0.6333$). Thus, the long-term state for the first group on its own is cooperation, while the long-term state for the second is defection. Combining the two groups with the six agents biased towards cooperation initially cooperating and the second six defecting, the figures show that overall behavior can go either way. The initial state is itself metastable, but fluctuations insure that eventually the group will go over to either mutual cooperation or mutual defection, with both possibilities equally likely. Either the defecting agents are swept over into cooperation or the cooperative agents collapse into defection. Parameter values are $H = 12$, $p = 0.99$, $b = 2.5$, $c = 1$, $\alpha = 1$ and $\tau = 1$.

This example invites comparison to the situation in which all members have $f^{crit} = 0.5666$. Such a group would be nearly bistable, although with a long-term tendency towards defection. A group with half its members initially cooperating would generally evolve rapidly to a state of mutual defection. With diversity as above, on the other hand, the two stable states of defection and cooperation become separated by a metastable state in which only half of the group cooperates while the other half free rides. For this particular choice of diversity, transitions from the metastable state of half cooperation/half defection to either the stable state of overall cooperation or overall defection are equally likely. Thus, diversity in the group has improved its chances of mutual cooperation compared to those of a uniform group with the same average f^{crit}.

5.4 Summary

In this chapter, we relaxed our previous assumption that individuals are identical in their beliefs and preferences. We then studied the effect of diversity on the dynamics of collective action. One type of diversity, represented as a spread about some common average belief among individuals, was shown to act as a source of extra uncertainty, thus shortening the time to an outbreak without affecting its abrupt nature. A second form of diversity, which occurs when a group consists of several subgroups each with its own distinct beliefs was also studied. In this case, we showed that cooperation (or defection) can be achieved by the group in demarcated stages. In general, an additional metastable state may appear for each new subdivision within the group.

The possible implications of this work for the study of cooperation in organizations can be briefly stated. In order to achieve spontaneous cooperation on a global scale, an organization could be structured into small subunits made up of individuals with a well-designed diversity of beliefs. The small size of the units allows for the emergence of sustained cooperation; diversity hastens its appearance. To the extent that these two elements, smallness and diversity, can be preserved in a hierarchical structure, we anticipate that sustained cooperation can be achieved in more complex systems as well.

We further examine the effect of organizational structure on the possibility for cooperation in large groups in the next chapter.

6 Organizational Fluidity

6.1 Introduction

The study of social dilemmas provides insight into a central issue of social behavior: how global cooperation among individuals confronted with conflicting choices can be secured. In the previous two chapters, we showed that cooperative behavior in a social setting can be spontaneously generated, provided that the groups are small and diverse in composition, and that their constituents have long outlooks. Furthermore, the emergence of cooperation takes place in an unexpected fashion, appearing suddenly and unpredictably after a long period of stasis.

In our model of ongoing collective action, intentional agents make choices that depend on their individual preferences, expectations and beliefs as well as upon incomplete knowledge of the past. Because of the nature of individual expectations, a strategy of conditional cooperation was shown to emerge from individually rational choices. According to this strategy, an agent will cooperate when the fraction of the group perceived as cooperating exceeds a critical threshold.

Since the critical threshold grows with group size, beyond a certain size, cooperation is no longer possible in flat, structureless groups. However, groups are often characterized by a distinct social structure that emerges from the pattern of interdependencies among individuals. Clustering in a social structure effectively decreases the size of the group as each individual cares most about the behavior of his/her own cluster. In fact, Bendor and Mookherjee (1987) have recently demonstrated that the level of cooperation depends on the structure of the group, in the context of federated structures with a central monitoring agency. Others who have studied social networks in the context of collective action include Marwell, Oliver and Prahl (1988). For our part, we will show how cooperation is a more likely outcome for hierarchically structured groups even without centralized control. In addition, isolated clusters of cooperation can persist and even trigger a cascade of cooperation throughout the rest of the organization.

The potential for cooperative solutions to social dilemmas increases further if groups allow for structural changes. In a fluid structure, the pattern of interdependencies can vary widely over time, since the sum of small local changes in the structure of a group results in broad restructurings. We find that fluid groups show much higher levels of cooperation over time. By moving within the group structure, individuals cause restructurings that enable cooperation. Computer experiments simulating fluid organizations faced with a social dilemma reveal a myriad of complex behaviors that result from the interplay between various amounts of cooperation and structural changes.

The advantages of fluidity must be balanced against a possible loss in effectiveness for an organization. By effectiveness, we mean how productive a given organization is in obtaining an overall utility over time. The production function for a social good will depend on the tightness or looseness of clustering, and possibly on the stability of the social structure. The form of the production function is not known *a priori*; however, given an optimal level of clustering for a particular good, we illustrate that there is an ideal range of fluidity associated with it.

In the next section, we expand upon the notion of structure in decentralized groups and develop the concept of a hierarchy of interactions. In Section 6.3, we extend our dynamical theory of collective action to include hierarchical groups. In Section 6.4, we present the results of computer experiments which show that, in terms of their ability to sustain high levels of cooperation over time, hierarchical structures are superior to flat groups and fluid hierarchies perform better than fixed ones. In Section 6.5, the trade-off between fluidity and effectiveness is examined.

6.2 The Topology of Organizations

Organizational theorists study group structure in order to elucidate the nature and working of the firm. In firms, there is generally an informal structure that emerges from the pattern of affective ties among participants as well as a formal one imposed from above (Scott 1992). In our study, we will be interested in the former type, the network of activities and interactions that link individuals in a group, commonly known as the sociometric structure.

A school of sociologists and social psychologists also regard the social structure of a group as elemental to understanding group behavior. The approach, called social network analysis, began with Barnes' (1954) and Bott's (1955) first attempts to use the relationship of the linkages in a network to interpret social action, and has gained momentum in the previous decade (Knipscheer and Antonucci 1990). From the micro view of individual interdependencies emerges a global perspective on social structures.

Thus, organizational structure in social groups can spontaneously emerge from the interactions of group members. In contrast with its typical usage in sociology, we use the notion of group in a loose sense, defining it as the community affected by a particular problem or situation. There are a number of different general types of structural topologies: flat, or structureless, hierarchical, matrix, circular, linear, and many others. In this chapter we restrict our attention to the first two.

In addition to its topology, an organizational structure can also be described by the amount of fluidity it exhibits. Fluidity encompasses such features as how flexible the structure is, how readily individuals can modify local structure by changing the strength of their interaction with others.

Consider, as a concrete example, the social problem of limiting air pollution. This is a problem that on one level the whole world faces and must solve collectively. The common good is clean air with a minimum of pollution, not only to make living conditions better today, but, more importantly, to ensure that the world remain hospitable to life in the future. Clean air is a common good because everyone benefits from it regardless of their own efforts to limit pollution.

In this example, the impact of pollution depends partially on relative geographical locations. That is, neglecting prevailing winds and currents, a person is more bothered when her neighbor burns his compost pile than when someone across town does the same. Similarly, the cumulative effect of everyone in town driving their cars to the popular bookstore (instead of bicycling or walking) affects a local resident much more than someone who lives miles away. Of course, some individuals' sphere of influence will range much further than others'. This dilution of impact with distance can be represented as a hierarchy of interactions which reveals itself in the unraveling of layer upon layer: neighborhood, town, county, state, world. The effect of one individual's actions on another depends on how many layers apart they are.

Within this hierarchical structure there is some fluidity. Although people are constrained by their resources and personal ties, they can often choose to move to a new location to escape a neighbor's radio or a textile company's fumes.

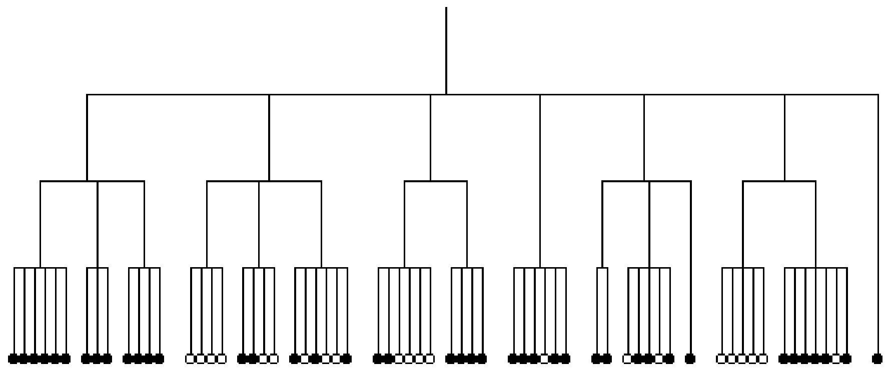

Figure 6.1: *A hierarchical organization can be visualized as an ultrametric tree.* The tree reveals the structure of the whole: each branching represents a subdivision of the higher level. The nodes at the lowest level represent individuals, filled circles mark cooperators, open circles, defectors. One can read directly from the tree the number of layers of the organization separating any two individuals by simply counting nodes backwards up the tree from each individual until a common ancestor is reached. This number determines the distance between two individuals. The larger the distance between them, the less their actions will affect each other.

We now quantify the notion of many levels of clustering within a group. For a group with a hierarchy of connection strengths among the members, the organizational structure can be represented graphically by an ultrametric tree, as in Fig. 6.1. The technique of representing the interdependencies in social networks by hierarchical clustering in tree-like structures dates from the work of Hartigan (1967), and more recently, Burt (1980). The level of interaction between two individuals can be read directly from the tree. Tightly clustered individuals share an ancestor node one level up in the tree. Those more loosely coupled may be connected two levels up. In general, interaction strength is indicated by the number of levels to a common ancestor node. Thus, the tree gives a visual description of the amount and extent of clustering in the group. Note that this interpretation of hierarchy as a pattern of interdependencies is divorced from any notion of rank within the group.

A second more general way to describe the pattern of interactions among individual members of a group is to define the matrix of interdependencies between them. Define **A** to be the matrix of interactions; then a_{ij} is the strength of interaction between individual *i* and *j*. This formalism permits generalization to many other topologies of group structure apart from hierarchical ones.

From this formalism it is easy to determine the extent of influence on any particular individual by the rest of the group. We call the cumulative influence perceived by individual *i* the rescaled quantity

$$\tilde{n}_i = \sum_{j=1}^{n} a_{ij}. \tag{6.1}$$

In a flat group, one in which all influences are equal, **A** would be filled with 1's, and in a group of size *n*, the rescaled size $\tilde{n}_i = n$, for all *i*. In a hierarchical structure, the components, a_{ij}, decrease with the number of levels separating *i* and *j*. One way to scale the strength of interaction with distance in the tree is to let

$$a_{ij} = \frac{1}{a^{d(i,j)}}, \tag{6.2}$$

where $d(i,j)$ is the number of levels separating *i* and *j* ($d(i,j) = 0$ for members of the same cluster) and *a* is the scaling factor. With this scaling, a cluster of size *a* one level distant from individual *i* is equivalent, from *i*'s point of view, to one agent in the same cluster as *i*. The rescaled size then becomes

$$\tilde{n}_i = \sum_{j=1}^{n} \frac{1}{a^{d(i,j)}}. \tag{6.3}$$

Along with social structure comes the notion of fluidity. In a fixed structure the pattern of interactions remains fixed in time. This need not be the case. In the example tying relative geographical positions to influence, the pattern of interactions changes slightly every time someone moves to a new locale, and more significantly when whole groups dissolve and reform. In a corporation, for example, employees may have some leeway to switch departments or even to move between regional branches. More importantly, the sum of small local changes in the structure of a group can result in broad restructurings over time.

The ease with which such moves can be accomplished reveals the amount of fluidity in the group structure. The notion of fluidity actually consists of two elements. One is how easily individuals can move within the structure, the second, how easily they can break away on their own, thereby extending the structure. These moves restructure the group in the sense that they change the pattern of interactions. Breaking away can range from moving to a secluded area to branching off to start a new work group to founding a new company. In many structured groups there will be costs associated with both moving within the structure or breaking away. The higher the costs, the less fluid the structure will be, as individuals must anticipate a clear advantage to themselves before moving or breaking away.

6.3 A Model of Collective Action in Organizations

The model of intentional agents with expectations engaged in ongoing collective action is easily extended to account for structured groups. Basically, the derivation of Section 4.3 holds once the group size and the number cooperating are rescaled, as sketched out below.

Extending the terminology of the previous section, the rescaled number of members cooperating, from the point of view of individual i, is

$$\tilde{n}_i^c = \sum_{j=1}^{n} \frac{k_j'}{a^{d(i,j)}} \tag{6.4}$$

and the rescaled fraction cooperating is

$$\tilde{f}_i^c = \frac{\tilde{n}_i^c}{\tilde{n}_i}. \tag{6.5}$$

Then the one-period utility at time t for member i becomes

$$U_i(t) = \frac{b}{\tilde{n}_i}\tilde{n}_i^c(t) - ck_i. \tag{6.6}$$

In direct analogy to the expectations set forth in Chapter 4, the members of the group expect the evolution of cooperation to qualitatively follow the deviation

$$\Delta\tilde{f}_i^c(t+t') = 1 - (1 - 1/\tilde{n}_i)\exp\left(-\frac{\alpha\tilde{f}_i^c(t-\tau)}{\tilde{n}_i}t'\right). \tag{6.7}$$

Once the individuals' expectations are taken into account, the advantage of cooperating over defecting for member i at time t is the net benefit

$$\Delta \widehat{B}_i(t) = H(b-c) - \frac{Hb(\tilde{n}_i - 1)}{\tilde{n}_i + H\alpha \tilde{f}_i^c(t-\tau)}. \tag{6.8}$$

This criterion reduces to the following condition for cooperation at time t:

$$\tilde{f}_i^{crit} \equiv \frac{1}{H\alpha}\left(\frac{\tilde{n}_i c - b}{b-c}\right) < \tilde{f}_i^c(t-\tau). \tag{6.9}$$

As before, the members of the group engage in conditional cooperation, cooperating when the rescaled fraction perceived as cooperating is greater than the critical amount, $\tilde{f}_i^{crit}$, and defecting when the rescaled fraction perceived cooperating is less than $\tilde{f}_i^{crit}$.

Finally, from Eq. 6.9 we can derive the rescaled size $\tilde{n}^*$ beyond which no individuals will choose to cooperate and consequently cooperation cannot be sustained. This limit is given by $\tilde{f}_i^{crit} = 1$ which yields

$$\tilde{n}^* = H\alpha\left(\frac{b}{c} - 1\right) + \frac{b}{c} \tag{6.10}$$

as the upper limit for the rescaled size beyond which cooperation will not be achieved. The smaller the rescaled size of the group, the larger the potential for sustained cooperation.

6.4 Structures for Cooperation

We previously showed that there is a clear upper limit to the size of a group which can support cooperation. Next we examine how this limit can be stretched by a hierarchically structured group whose members are able to move freely within the group. First, we review the limitations of flat groups and groups with fixed structures. Then we show to what extent these limitations can be overcome by fluid groups. Finally, in the next section we discuss the possible trade-off between fluidity and effectiveness in organizations.

6.4.2 Flat groups

Results of Monte Carlo simulations conducted in asynchronous fashion confirm the theoretical predictions for structureless groups obtained using the stability function formalism. Each individual decides whether to cooperate or defect based on the criterion given in Eq. 4.6 that the perceived fraction cooperating, f^c, must be greater than a critical fraction, f^{crit}. Uncertainty is a factor since these decisions are based on perceived levels of cooperation which differ from the actual attempted amount of cooperation.

As shown earlier there is a critical group size beyond which cooperation will not be sustained. There is also a range of group sizes below the critical size for which the system has two equilibrium points, one with most of the group defecting, the other with most members cooperating.

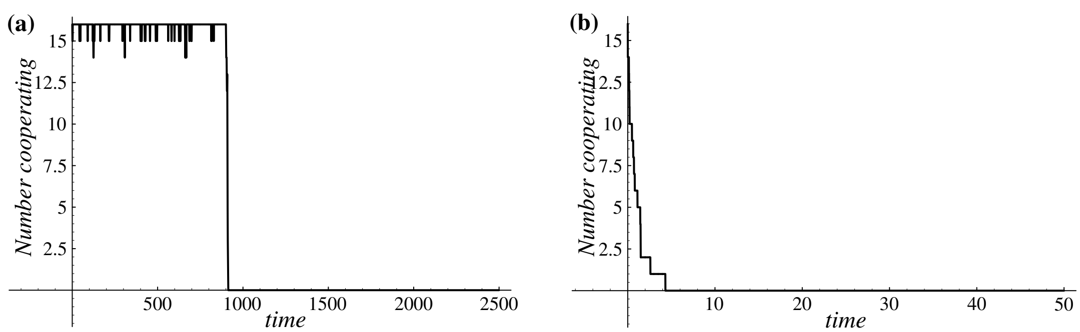

Figure 6.2: *Collapse of cooperation in a structureless group.* The evolution of cooperation is shown here for a structureless group of size $n = 16$. The benefit for cooperating, *b*, is 2.5, the cost, *c*, is 1, the probability *p* that agent cooperates successfully is 0.95, the reevaluation rate α is 1, and the delay τ is also 1. (a) Given these parameters, the horizon length is set at 12.0 to provide an example in which cooperation is the metastable equilibrium and defection, the global equilibrium. Indeed, the group cooperates for a very long time, over 900 time steps. The outbreak of defection is sudden and unpredictable, occurring at widely different times in numerous simulations of the same system. (b) Here the horizon length is set to 8.9; defection becomes the only equilibrium state for the system. As predicted by the theory, the group relaxes exponentially fast to the defecting state.

In Fig. 6.2(a), the evolution of the system is shown for a group of size 16 whose horizon length is such that the critical size for cooperation is less than 16. The system begins in the metastable overall cooperation state and remains there for a very long time. Defection is the global equilibrium, however, and eventually, the transition to defection occurs. In (b), the horizon length is such that the critical size is greater than 16. In this case, the only equilibrium point is overall defection; accordingly, the onset of defection is swift for a group initially cooperating.

6.4.3 Fixed structures

Imagine now that the pattern of interactions among the 16 individuals depicts a group broken down into 4 clusters of 4 agents each. If the amount of interdependency between agents scales down by a factor of $a = 4$ for each level in the hierarchy separating two individuals, then the rescaled size of the group (Eq. 6.3) is $\tilde{n} = 4 + 12/4 = 7$ for all the agents (due to the symmetric structure). If only because the rescaled group size is smaller, cooperation can be sustained under more severe conditions for a hierarchical group than for a flat one.

More interestingly, an enclave of cooperation within the hierarchy can initiate a widespread transition to cooperation within the entire organization. The mechanism that explains this cascade to cooperation is due to the clustering of agents and cannot occur in a flat group. The cascading phenomenon is more vivid in hierarchies with several layers; consider for example a group of 27 agents structured as in Fig. 6.3(a). In this case, the rescaled size $\tilde{n} = 3 + 6/3 + 18/3^2 = 7$ for all the agents. Filled circles at the lowest level represent cooperating agents, open circles are defecting agents. Individuals in different parts of the organization will observe widely differing amounts of cooperation. Those in the leftmost cluster see a fraction of about 3/7 cooperating while agents one level away see 1/7 cooperating and agents two levels away see a fraction 0.33/7 cooperating. So if $1/7 \lesssim f^{crit} \lesssim 3/7$, the agents in the leftmost cluster can sustain cooperation almost indefinitely among the three of them despite the fact that the rest of the organization is defecting. On the other hand, although agents one level away may be defecting initially, for certain parameter choices, cooperation may actually be the long-term stable state for those clusters in the presence of uncertainty. If these agents in clusters one level away start cooperating, they may then trigger the onset of cooperation in the clusters two levels away from the initial cooperators. Fig. 6.3 shows four snapshots in time of a group of 27 agents exhibiting this sequence of behavior over time.

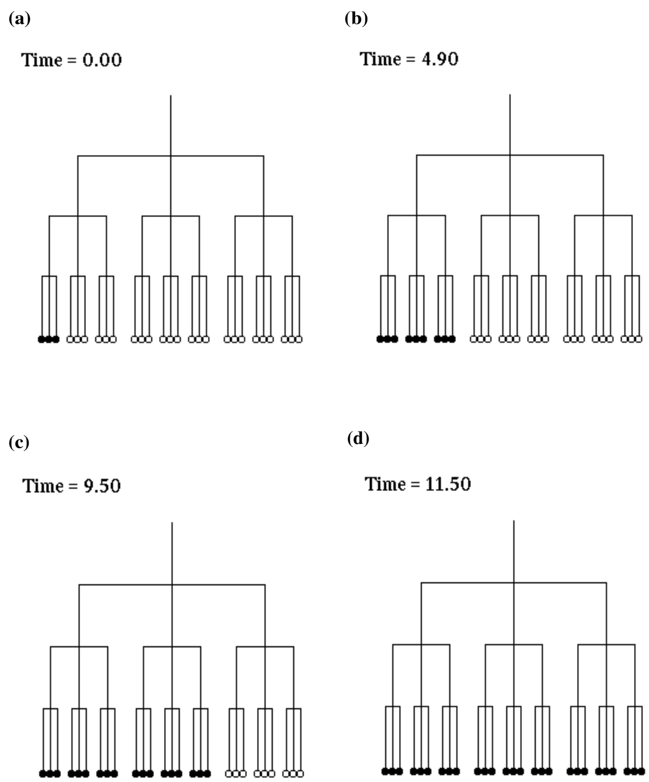

Figure 6.3: *Cascade to cooperation in a fixed hierarchy.* Overall cooperation in a hierarchically structured group can be initiated by the actions of a few agents clustered together. These agents reinforce each other and at the same time can spur agents one level further removed from them to begin cooperating. In turn, this increase in cooperation can spur cooperation in agents even further removed in the structure. In this example, the structure is fixed: the three level hierarchy consists of three large clusters, each subsuming three clusters of three agents each. The parameters in this example were set to $H = 10.0$, $b = 2.5$, $c = 1$, $p = 0.97$, $\alpha = 1$, and $\tau = 0$.

However, as with flat groups, groups with fixed structures easily grow beyond the bounds within which cooperation is sustainable. In the case mentioned earlier of 16 individuals broken down into four clusters, when the rescaled size is larger than $\tilde{n}^*$, the group rapidly evolves into its equilibrium state of overall defection. But even when the rescaled size is less than $\tilde{n}^*$, cooperation will be metastable in the sense that the agents may remain cooperating for a very long time but eventually a transition to overall defection will occur. Before and after snapshots which exemplify this abrupt transition are shown in Fig. 6.4.

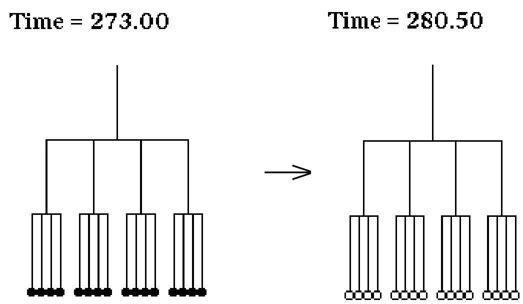

Figure 6.4: *Breakdown of cooperation in a fixed hierarchy.* Here are snapshots taken for a group taken before and after the transition from overall cooperation to overall defection. The times of the snapshots are shown to demonstrate the rapidity of the transition. The group's fixed hierarchical structure is unable to sustain cooperation indefinitely since the rescaled size $\tilde{n} = 7$ is larger than the limiting size for cooperation, in this case, $\tilde{n}^* = 4.42$. The parameters in this example were set to $H = 4.75$, $b = 2.5$, $c = 1$, $p = 0.9$, $\alpha = 1$, and $\tau = 0$.

6.4.4 Fluid structures

In fluid structures, individual agents are able to "move" within the organization. The way in which this might occur depends on what determines the structure, be it geographical location in the world or type of work within a company. The amount of fluidity in an organization is variable. It depends on two factors: (1) the ease with which an individual can switch between two clusters; and (2) how readily agents break away to form clusters of their own. Globally, the sum of individual moves between clusters translates into a mixing and merging of the agents into larger clusters. On the other hand, the local decisions made by individuals to break away and start new clusters expands the structure by decreasing the extent of clustering.

In the simulations, agents regularly reevaluate the situation. Previously, they had only one choice to make: whether to cooperate or defect. In fluid organizations, they must also evaluate how satisfied they are with their location within the structure. However, the agents are assumed to evaluate only one of these two choices at any given decision point.

Individuals make their decision to cooperate or defect according to the long-term benefit they expect to obtain, as before. In order to evaluate their position in the structure, an individual compares the long-term payoff it expects if it stays put with the long-term payoff it expects if it moves to another location, chosen at random. In these calculations, the agent's strategy in response to the social dilemma remains the same, be it cooperation or defection. In order to determine the payoff it expects to obtain by moving, the individual must have access to the world as seen by the individual whose position it is evaluating. That is, if individual *i* wants to evaluate the position of individual *j*, *i* needs to know $\tilde{n}_j$ and $\tilde{n}_j^c$ as well as $\tilde{n}_i$ and $\tilde{n}_i^c$. This additional information is not required by individuals in groups whose structure is fixed. Thus, the validity of this model of fluid organizations is limited by the extent to which this information is available.

In addition, there might be a barrier to moving, either because there is a cost associated with it, or because the agents are risk-averse. So for example, an individual might move only if it perceives the move to increase its expected payoff by a certain percentage, which we shall refer to as the moving barrier.

When evaluating its position, the individual also considers the possibility of breaking away to form a cluster of its own. The agent will do so if it perceives the payoff for either staying put or moving to be small enough that it feels it has nothing to lose by taking a chance and starting its own cluster. The agent can only break away one level at a time, so from a cluster of several agents, it may break away to form a cluster on its own one level distant from its parent cluster. The next time this same individual reevaluates its position, it can then break away an additional layer distant from its parent cluster, if no other agent has come to join it in its new cluster. In this way an agent can move many levels away from its original cluster.

An agent will thus form a new cluster when its current perceived payoff is below a given threshold, called the break away threshold. When the break away threshold is high, agents splinter off most easily; when it is low, they are more reluctant. This is because higher thresholds describe an individual that is more likely to be less satisfied with both its present position and the alternatives and thus will have an increased tendency to break away. We will define these thresholds as a fraction of the maximum possible payoff over time.

Computer experiments implementing this notion of fluid organizations reveal a myriad of complex behavior. Through local moves and break aways, the organization can often restructure itself to recover either from outbreaks of defection or to overcome an initial bias to defection in the group. A series of snapshots over time for a group of size 16 are shown in Fig. 6.5. Initially, the group is divided into four clusters of four defecting agents each — this represents an extreme

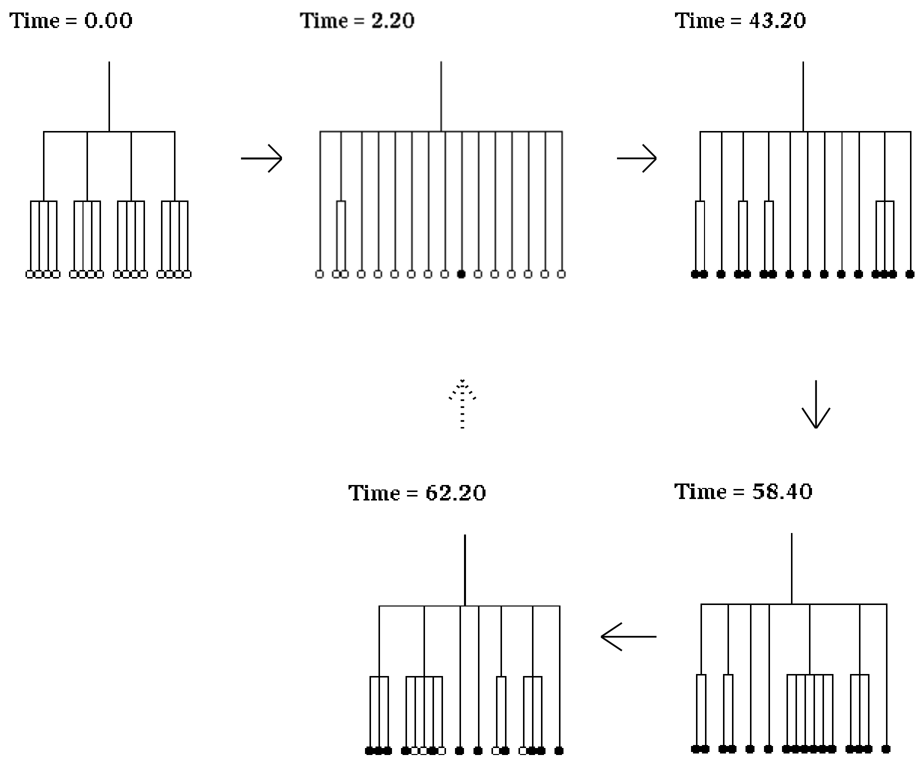

Figure 6.5: *Fluid cooperation.* This figure highlights the ability of fluid groups to recover from overall defection among its members. Initially all members of the group, divided into four clusters of four agents each, are defecting. The next snapshot in time shows that almost all of the agents have broken away on their own. In this dispersed structure, agents are much more likely to switch to a cooperative strategy, and indeed, by the next snapshot shown, all individuals are cooperating. Because of uncertainty, however, agents will occasionally switch between clusters. Eventually, a cluster grows large enough that a transition to defection begins within that cluster. At the same time, individuals will be moving away to escape the defectors. We see these processes happening in the fifth snapshot. At this point, more and more agents will break away on their own, and a similar cycle begins again. The parameters in this example were set to $H = 4.75$, $b = 2.5$, $c = 1$, $p = 0.9$, $\alpha = 1$, and $\tau = 0$, as in the previous figure.

condition and is thus a good test of the group's ability to overcome defection. A later snapshot shows that most of the agents have broken away on their own. This restructuring enables the global switchover to cooperation since the rescaled size is now small enough that cooperation is the global equilibrium for the system, were the structure to remain fixed. Fixed it is not, however, and over time, the now mostly cooperating agents cluster back together again, moving towards clusters where they perceive the amount of cooperation to be highest. Eventually, one or more clusters will grow too large to support cooperation indefinitely, creating the potential for an outbreak of defection. Each outbreak is quelled by a process similar to that responsible for the initial recovery. Cycles of this type have been observed to appear frequently in the lifetime of the simulated organization.

6.5 Effectiveness *vs.* fluidity

In a very fluid organization, individuals break away often, thus founding new clusters, and move between clusters very readily. On the other hand, these actions rarely occur in a group with little fluidity in its structure. In the example of a fluid organization given above, the amount of fluidity of movement was set to an intermediate amount: the moving barrier was set at 15% and the break away threshold was 45% of the maximum possible payoff.

Observations of many simulations at various levels of fluidity and for differing values of horizon length point to the following conclusions. High break away thresholds mean that individuals tolerate little deviation from the maximum available payoff and often break away in search of "greener pastures." Since breaking away is the primary mechanism that allows a group to recover from bouts of defection, a greater tendency to break away is favorable in this sense. However, higher break away thresholds also cause the structure to become very dilute and disconnected. In the extreme case, agents will tend to all splinter off on their own, as in the second snapshot of Fig 6.5. On the other hand, large moving barriers inhibit the clustering of agents and cause the structure to vary little over time. Thus, large moving barriers help stabilize the structure. If the group is cooperating, its structure may then remain relatively fixed over time. Coupled with high break away thresholds, large moving barriers mean that a cooperating system may be frozen into a structure with a very small amount of clustering.

In general, the combination of ease of breaking away and difficulty of moving between clusters enable the highest levels of cooperation. However, this yields the counter-intuitive result of cooperating dis-organizations! That is, agents are all cooperating, but on their own.

The reason this seems counter-intuitive is because, thus far, the effectiveness of the organization has not been considered. By effectiveness, we mean how productive a given organization is in obtaining an overall utility over time. This is revealed by how well an organization achieves its goals. It seems reasonable to assume that there is an optimal degree of clustering for which an organization operates most effectively. The ideal amount of clustering for a particular organization will depend on type of good the group is attempting to provide itself with and how this good is produced. Determining the optimal amount of clustering for a given type of organization is beyond the scope of this thesis; however, we can say something about the range of fluidity that allows an organization to be most effective given an optimal level.

To factor in effectiveness, we must modify the production function of the common good (Eq. 4.1) to reflect increased performance at optimal clustering levels. One way to do this is to posit that the benefit to the group for a cooperative action depends on the amount of clustering. To this end, we introduce a variable

$$\tilde{n}^{clust} = \frac{1}{n}\sum_{i=1}^{n} \tilde{n}_i \tag{6.11}$$

that indicates the average amount of clustering within the structure and postulate that the benefits of cooperation are highest when the amount of clustering is equal to the optimal amount $\tilde{n}^{opt}$ and falls off to either side. Qualitatively, then, the revised individual utility function looks like

$$U_i(t) = b\exp\left[-\left(\tilde{n}^{opt} - \tilde{n}^{clust}\right)^2/n^2\right]\frac{\tilde{n}_i^c(t)}{\tilde{n}_i} - ck_i. \tag{6.12}$$

The subsequent analysis in Section 6.4 still holds by replacing b with

$$b' = b\exp\left[-\left(\tilde{n}^{opt} - \tilde{n}^{clust}\right)^2/n^2\right]. \tag{6.13}$$

Specifically, the collective action problem remains an n-person prisoner's dilemma as long as

$$b' > c > \frac{b'}{\tilde{n}_i}, \tag{6.14}$$

for all possible reconfigurations of the structure.

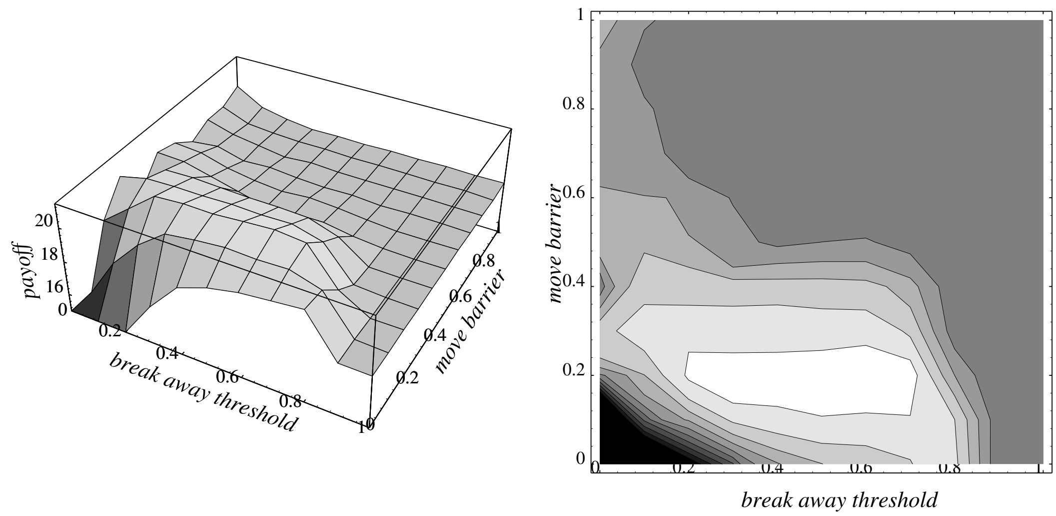

Figure 6.6: *Effectiveness vs. fluidity in changing organizations.* In (a) the total actual payoff to the group, averaged over thousands of time steps, is plotted as a function of the break away threshold and move barrier. High break away thresholds and low move barriers correspond to high amounts of fluidity in the structure of the organization. The contour plot is given in (b) to highlight the gradient of the functional dependence of payoff on fluidity. In this example, the optimal amount of clustering is taken to be $\tilde{n}^{clust} = 10$, a level corresponding to the clustering of the group into two large subgroups of eight agents each. These results show that there is a range in the amount of fluidity at which this organization operates most effectively. The other parameters values in this example are $H = 8$, $b = 2.5$, $c = 1$, $p = 0.95$, $\alpha = 1$, and $\tau = 0$.

The interplay between fluidity and effectiveness is best observed when the critical rescaled size falls inside a certain regime. Within this regime, the critical size is such that high levels of clustering cause the group to be unstable to defection while low levels allow to group to recover. To measure effectiveness, we keep track of the average actual payoff over time for varying amounts of fluidity, given a fixed optimal average amount of clustering. Once again, we consider a group of 16 agents, with a hierarchical structure two levels deep. The ideal amount of clustering is set at $\tilde{n}^{clust} = 10$. Such a high level of clustering occurs when the agents break up into two large clusters. Fig. 6.6 shows that there is indeed a range in the amount of fluidity at which the organization operates most effectively. Systems operating within this range of fluidity may cooperate less over time than those with less fluidity, but compensate by having higher levels of clustering.

6.6 Summary and Outlook

We have shown in this chapter that fluid organizations display higher levels of cooperation than can be attained by groups with either a fixed social structure or lacking one altogether. In fluid organizations, individuals can easily move between groups or even strike out on their own. The sum of many such moves results in a global restructuring of the organization over time.

In an organization faced with the social dilemma of providing itself with a collective good, such ongoing incremental restructurings can provide free riders with the incentive to cooperate. That is, in a group with many free riders, individual moves will cause the structure to disperse into many small clusters. Cooperation is much more likely to emerge spontaneously within these small clusters and then spread to the rest of the organization.

The results of computer experiments bear out these conclusions. In particular, fluid organizations can display long cycles of sustained cooperation interrupted by short bursts of defection. The average level of cooperation sustained over time depends on the amount of fluidity in the organization as well its breadth and extent. A point to consider, however, is that the advantages of fluidity must be balanced against a possible accompanying decrease in the effectiveness of the group.

Although our model assumes hierarchical groups faced with a social dilemma, we expect our results will generalize to less constrained forms of social networks and to more general production problems. Thus, fluidity may play a crucial role in ensuring that a large organization attain high levels of cooperation. In fact, modern firms have begun to foster organizational fluidity: a central theme in recent reorganizations of large corporations has been the creation of pathways that allow project-centered groups to rapidly form and reconfigure as circumstances demand.

The concept of fluidity employed here to study ongoing collective action might be applicable to other complex systems that also have evolving social structures. For example, it has been observed that environmental changes cause the roles and interdependencies among ants in an ant colony to vary over time (Gordon 1991). Some type of fluidity in the ants' roles might be responsible.[14] In addition, the methods used here to explore evolving structures might be applied

[14] Personal communication with Deborah Gordon, March 1993.

to simulations of artificial life, or, more appropriate to theme of cooperation, to negotiation among intelligent agents in distributed systems, or even to neural networks.

Appendix: Evolutionary Games and Computer Simulations

The prisoner's dilemma has long been considered the paradigm for studying the emergence of cooperation among selfish individuals. Because of its importance, it has been studied through computer experiments as well as in the laboratory and by analytical means. However, there are important differences between the way a system composed of many interacting elements is simulated by a digital machine and the manner in which it behaves when studied in real experiments. In some instances, these disparities can be marked enough so as to cast doubt on the implications of cellular automata type simulations for the study of cooperation in social systems. In particular, if such a simulation imposes space-time granularity, then its ability to describe the real world may be compromised. Indeed, we show that the results of digital simulations regarding territoriality and cooperation differ greatly when time is discrete as opposed to continuous.

Over the past decade, the increasing speed and availability of powerful computers has made computer simulations an attractive method to research the behaviors of complex systems. In the physical sciences, simulation techniques have been used to study problems such as critical phenomena, dynamical systems and the large scale structure of the universe. In chemistry and biology, computer simulations provide insight into the mechanics of protein folding and cell metabolism. Similarly, the appearance of behavioral patterns in social settings can be studied using this method, provided that the assumptions of the model capture the underlying interactions.

In a recent paper, Nowak and May (1992) presented a set of intriguing results concerning the evolution of cooperation among players placed on a two-dimensional array and confronted with a prisoner's dilemma, which in recent years has become a metaphor for the evolution of cooperation. By running a number of computer simulations, they showed that when players interact with their neighbors through simple deterministic rules and have no memory of past events, the overall evolution produces striking spatial patterns, in which cooperators and defectors both persist indefinitely. Furthermore, for certain parameter values, they observed that regardless of initial conditions, the frequency of cooperators always reaches the same proportion, raising the interesting issue of the existence of a universal constant governing prisoners' dilemma interactions on a lattice. These results were further elaborated on by Sigmund (1992), who used them as evidence that territoriality favors cooperation among biological organisms and also suggested that similar results would occur in the case of stochastic transition rules.

While it has been known for some time that cellular automata with deterministic rules can generate striking spatio-temporal patterns (Wolfram 1984), their usefulness for studying real world systems is not straightforward. One reason is that the granularity imposed by cellular digital machines on both the spatial and temporal domains can generate behaviors that may not have counterparts in the continuum limit.

In fact, there are important differences between the way a system composed of many interacting elements is simulated by a digital machine and the manner in which it behaves when studied in real experiments. In some instances, these disparities can be marked enough so as to cast doubt on the implications of cellular automata type simulations for the study of cooperation in social systems. These differences, which have been analyzed in detail for simulations of Ising-like magnets and other condensed matter systems (Choi and Huberman 1983, Ceccatto 1986, Gawlinsky *et. al.* 1985), also have important implications for computer studies of evolutionary games. This issue must be addressed if computer simulations are to provide insight into social and biological dynamics.

We begin by analyzing the discrepancies between some digital simulations and real systems and afterwards compare the results of running the same computations presented by Nowak and May in both synchronous and asynchronous fashion. Our results for asynchronous updating demonstrate that defection is the dominant evolutionary outcome if at least one defector is present in the initial configuration. The matrix converges rapidly to steady-state and does not display the widely changing spatial patterns observed in the synchronous case. These results, which do not correspond at all to the behavior found for synchronous updating, cast into doubt the conclusions recently obtained concerning territoriality and the universality of long-term average levels of cooperation.

Before considering the particular computer experiments studied by Nowak and May, it is useful to clarify the potential differences between a cellular automata simulation of a natural process and the dynamics of the same system as found in nature. In a simulation of the type presented by the Nowak and May (1992), the general computation goes through a series of discrete states which are updated by the program at integral values of unit time, according to a set of given instructions. Nothing happens for times shorter than this unit time. The simulations presented in (Nowak and May 1992) are *synchronous*, meaning that all the players are updated in unison at every time step. The resulting global dynamics is mathematically described by a finite difference equation, which if sufficiently nonlinear can generate the complicated patterns and chaotic outcomes that are displayed in their paper.

In natural social systems, however, a global clock that causes all the elements of the system update their state at the same time seldom exists. While clocks and seasonal effects can synchronize metabolic and reproductive processes in biological organisms, in most social settings, players, agents, or organisms act at different and uncorrelated times on the basis of information that may be imperfect and delayed. At about the same time that one player initiates an interaction with its neighbor, another pair may be wrapping up an ongoing game. In the physical realm, the interactions among the individual atoms or molecules making up a magnet or a crystal determines the behavior of the material without there being a global clock which synchronizes their actions. In this *asynchronous* world, the dynamics is usually expressed in the form of differential equations, whose solutions are not always the same as those of their finite difference counterparts

This analysis implies that if a computer simulation is to mimic a real world system with no global clock, it should contain procedures that ensure that the updating of the interacting entities is continuous and asynchronous. This entails choosing an interval of time small enough so that at each step at most one individual entity is chosen at random to interact with its neighbors. During this update, the state of the rest of the system is held constant. This procedure is then repeated throughout the array for one player at a time, in contrast to a synchronous simulation in which all the entities are updated at once.

In order to show the striking differences between synchronous and asynchronous simulations of cooperative games, we conducted a number of computer experiments using the same scenario presented by Nowak and May. Fig. A.1(a) shows the results of a computer experiment with synchronous updating of a system of two dimensional players on a 99×99 square-lattice world with fixed boundary conditions after 217 generations. The initial condition consists of a single defector at the center surrounded by a world of cooperators. As can be verified, the resulting symmetrical pattern is the same as that shown by Nowak and May in Fig. 3(c) of their paper.

The evolution of the same system is very different when the players' strategies are updated asynchronously. In our experiments, we choose the time interval between updates to be small enough so that at most one player updates within that time interval. As a consequence of this asynchronous updating, each player awakens to see a slightly different world than the players acting before or afterwards. The updating is performed as follows. Each player receives a score per unit time that is updated continuously as the matrix of players varies over time. In analogy to the discrete scenario considered in the synchronous case, each player's score per unit time is the sum of its payoffs per unit time. These payoffs are received from interactions with each of its neighbors and itself in a prisoner's dilemma confrontation.

Next, within a microstep, at most one player is replaced by the highest scoring player within its neighborhood. The size of the microstep is chosen so that the average updating time (equivalent to one generation) for the whole array is the same for both the synchronous and asynchronous cases.

Fig. A.1(b) shows the asynchronously updated system at the same point in time as that of Fig. A.1(a), starting from identical initial conditions. Within a hundred generations or so, the array evolves into a fixed state in which all of the players are defecting. In fact, as long as there is at least one defector in the initial state, the matrix always evolves rapidly into a state of overall defection.

Alternatively, one might hypothesize that synchronous simulations of evolutionary games are relevant to real world systems in which there are delays in the transmission of information. Delays would then cause player states to be updated on the basis of neighborhood configurations that correspond to earlier times. Indeed, if the player's score were intended to reflect its fitness function (*i.e.*, its ability to reproduce), one would expect a player's present score (or number of offspring) to depend on past interactions, as well as present ones.

(a)

(b)

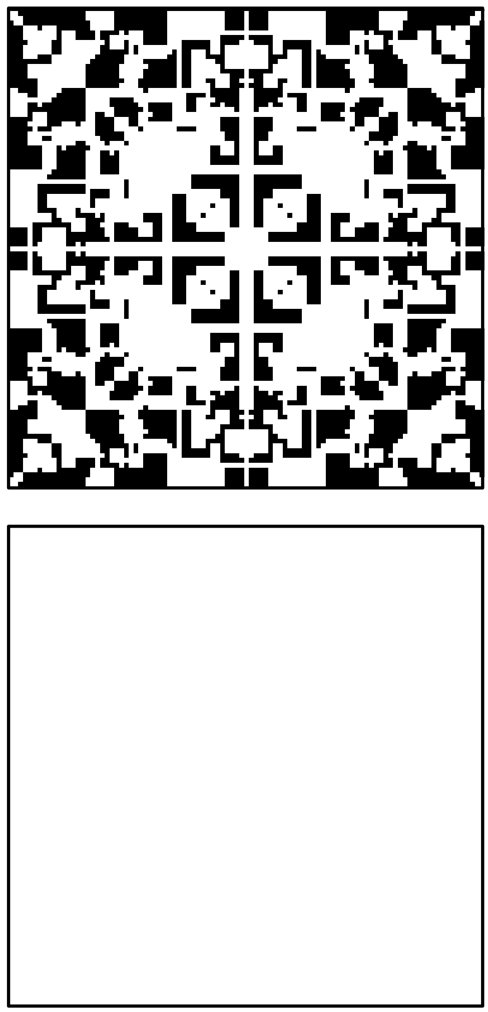

Figure A.1: *Synchronous versus asynchronous updating* of player actions results in differing evolutionary pathways. This simulation starts with a single defector at the center of a 99 x 99 square lattice world of cooperators with fixed boundary conditions. The coding is as follows: black represents a C site and white represents a D site. (a) Player actions are updated synchronously. This snapshot was taken at generation $t = 217$ and corresponds to Fig. 3(b) of Nowak and May's paper. (b) Player actions are updated asynchronously. Within a hundred generations or so, the matrix evolves into a fixed state in which all of the players are defecting.

However, in cases where player interactions depend upon delayed information, a simulation may still possess many complicated dynamical features that will not reduce to the case of synchronous updating. In our experiments, we found that when the players are updated asynchronously based on delayed scores, the evolution of cooperation depends strongly on the size of the delay: the greater the delay the higher the asymptotic level of cooperation. In addition, for very large delays (corresponding to several generations), we observed initial transients that lengthened as delays increased.

These results show that while computer experiments provide a versatile approach for studying complex systems, an understanding of their subtle characteristics is required in order to reach valid conclusions about real world systems. The instance provided by Nowak and May is only one of many cases where the outcomes of synchronous evolutionary games on computers have been invoked to provide insights into the workings of living systems. Other examples are provided by computer experiments used to validate theories about the emergence of order in evolution (Kauffmann 1969), phenotypic novelties through progressive genetic change (Agur and Kerszberg 1987), the molecular origin of life (Tamayo and Hartman 1989) and studies of artificial life forms through computer simulations (Langton 1986). Unless global clocks synchronize mutations and chemical reactions among distal elements of biological structures, the patterns and regularities observed in nature require continuous descriptions (Murray 1989) and asynchronous simulations.

References

E. Abrahams, P. W. Anderson, D. C. Licciardello, and T. V. Ramakrishnan. *Physical Review Letters*, 42:673, 1979.

Dilip Abreu. On the theory of infinitely repeated games with discounting. *Econometrica*, 56(2):383–396, 1988.

Zvia Agur and Michel Kerszberg. The emergence of phenotypic novelties through progressive genetic change. *American Naturalist*, 129(6):862–866, 1987.

Daniel Amit, Hanoch Gutfreund, and Haim Sompolinsky. Spin-glass models of neural networks. *Physical Review A*, 32:1007–1018, 1985.

Kenneth J. Arrow. *The Limits of Organization*. New York, W. W. Norton and Co., Inc., 1974.

N. W. Ashcroft and N. D. Mermin, editors. *Solid State Physics*. Holt, Rinehart, and Winston, New York, 1976.

Robert Axelrod. *The Evolution of Cooperation*. Basic Books, New York, 1984.

Robert Axelrod and William D. Hamilton. The evolution of cooperation. *Science*, 211:1390–1396, 1981.

C. Azariadis. Self-fulfilling prophecies. *Journal of Economic Theory*, 25:380–396, 1981.

Hao Bai-Lin. *Chaos II*. World Scientific, Singapore, 2nd edition, 1990.

Per Bak, Chao Tang, and Kurt Wiesenfeld. Self-organized criticality. *Physical Review A*, 38:364–374, July 1 1988.

J. A. Barnes. Class and committees in a norwegian island parish. *Human Relations*, 7:39–58, 1954.

R. Bellman and K. L. Cooke. *Differential-Difference Equations*. Academic Press, New York, 1963.

Jonathan Bendor and Dilip Mookherjee. Institutional stucture and the logic of ongoing collectiveaction. *American Political Science Review*, 81(1):129–154, 1987.

Olivier Jean Blanchard and Stanley Fischer. *Lectures on Macroeconomics*. The MIT Press, Cambridge, Massachusetts, 1989.

M. Boldrin and L. Montrucchio. The emergence of dynamic complexities in models of optimal growth: The role of impatience. *Journal of Economic Theory*, 40:26–39, 1986.

Elizabeth Bott. Urban families: conjugal roles and social networks. *Human Relations*, 8:345–383, 1955.

S. R. Broadbent and John M. Hammersley. Percolation processes i. crystals and mazes. *Proceedings of the Cambridge Philosophical Society*, 53:629–641, 1957.

R. S. Burt. Models of network structure. *Annual Review of Sociology*, 6:79–141, 1980.

Michael Carter and Rodney Maddock. *Rational Expectations*. MacMillan Publishers, Ltd., London, 1984.

H. A. Ceccatto. Effective discrete-time dynamics in monte carlo simulations. *Physical Review B*, 33:4734–4736, 1986.

H. A. Ceccatto and Bernardo A. Huberman. Persistence of nonoptimal strategies. *Proc. Natl. Acad. Sci. USA*, 86:3443–3446, 1989.

Peter Cheeseman, Bob Kanefsky, and William M. Taylor. Where the really hard problems are. In J. Mylopoulos and R. Reiter, editors, *Proceedings of IJCAI91*, pages 331–337, San Mateo, CA, 1991. Morgan Kaufmann.

M. Y. Choi and Bernardo A. Huberman. Dynamics of nonlinear networks. *Physical Review B*, 28:2547–2554, 1983.

Scott H. Clearwater, Bernardo A. Huberman, and Tad Hogg. Cooperative solution of constraint satisfaction problems. *Science*, 254:1181–1183, November 22, 1991.

Scott H. Clearwater, Bernardo A. Huberman, and Tad Hogg. Cooperative problem solving. In B. Huberman, editor, *Computation: The Micro and the Macro View*, pages 33–70. World Scientific, Singapore, 1992.

Randall Davis and Reid G. Smith. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, 2(3):63–109, 1983.

Robyn M. Dawes. Social dilemmas. *Annual Review of Psychology*, 31:169–193, 1980.

Richard H. Day and Weihong Huang. Bulls, bears and market sheep. *Journal of Economic Behavior and Organization*, 14:299–329, 1990.

D. d'Humieres, M. R. Beasley, B. A. Huberman, and A. Libchaber. Chaotic states and routes to chaos in the forced pendulum. *Phys. Rev. A*, 26:3483–3496, 1982.

David D. Friedman. *Price Theory*. South-Western Publishing Co., Ohio, 1990.

Les Gasser and Michael N. Huhns, editors. *Distributed Artificial Intelligence, volume II*. Morgan Kaufmann Publishers, Inc., 1989.

E. T. Gawlinsky, M. Grant, J. D. Gunton, and K. Kaski. Growth of unstable domains in the two-dimensional ising model. *Physical Review B*, 31:281–288, 1985.

Natalie S. Glance, Tad Hogg, and Bernardo A. Huberman. Computational ecosystems in a changing environment. *International Journal of Modern Physics C*, 2(3):735–753, 1991.

Natalie S. Glance and Bernardo A. Huberman. Dynamics with expectations. *Physics Letters A*, 165:432–440, 1992a.

Natalie S. Glance and Bernardo A. Huberman. Expectations and social outcomes. In S. E. Van der Leeuw, editor, *Dynamical Modelling and the Study of Change in Archaeology*. Edinburgh University Press, Edinburgh, 1992b.

Natalie S. Glance and Bernardo A. Huberman. The outbreak of cooperation. *Journal of Mathematical Sociology*, 17, 1992c, in press.

Natalie S. Glance and Bernardo A. Huberman. Social dilemmas and fluid organizations. In K. Carley and M. Prietula, editors, *Computational Organization Theory*. Lawrence Elbraum Associates, Inc., New Jersey, 1993a, in press.

Natalie S. Glance and Bernardo A. Huberman. Dynamics of social dilemmas. *Scientific American*, 1993b, to appear.

Deborah M. Gordon. Behavioral flexibility and the foraging ecology of seed-eating ants. *The American Naturalist*, 138(2):379–411, 1991.

Jean-Michel Grandmont. Stablility of cycles and expectations. *Journal of Economic Theory*, 40:138–151, 1986.

Jean-Michel Grandmont and Pierre Malgrange. Nonlinear economic dynamics: Introduction. *Journal of Economic Theory*, 40:3–12, 1986.

Mark S. Granovetter. The strength of weak ties. *American Journal of Sociology*, 78:1360–1380, 1973.

Geoffrey Grimmet. *Percolation*. Springer-Verlag, New York, 1989.

Garrett Hardin. The tragedy of the commons. *Science*, 162:1243–1248, 1968.

Russell Hardin. Collective action as an agreeable n-prisoners' dilemma. *Behavioral Science*, 16(5):472–481, 1971.

Russell Hardin. *Collective Action*. Johns Hopkins University Press, Baltimore, 1982.

J. A. Hartigan. Representation of similarity matrices by trees. *Journal of the American Statistical Association*, 62:1140–1158, 1967.

T. Hogg and B.A. Huberman. Artificial intelligence and large scale computation: a physics perspective. *Physics Report*, 156 (5):227–310, 1987.

Tad Hogg and Bernardo A. Huberman. Controlling chaos in distributed systems. *IEEE Trans. on Systems, Man and Cybernetics*, 21(6):1325–1332, 1991.

John H. Holland. *Adaptation in natural and artificial systems*. The University of Michigan Press, Ann Arbor, 1975.

J. J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Science*, 79:2554–2558, 1982.

Bernardo A. Huberman, editor. *The Ecology of Computation*. North-Holland, Amsterdam, 1988.

Bernardo A. Huberman and Natalie S. Glance. Diversity and collective action. In H. Haken and A. Mikhailov, editors, *Interdisciplinary Approaches to Complex Nonlinear Phenomena*. Springer, New York, 1993a.

Bernardo A. Huberman and Natalie S. Glance. Evolutionary games and computer simulations. *Proceedings of the National Academy of Science*, 1993b, to appear.

Bernardo A. Huberman and Tad Hogg. Phase transitions in artificial intelligence systems. *Artificial Intelligence*, 33:155–171, 1987.

Bernardo A. Huberman and Tad Hogg. The behavior of computational ecologies. In B. A. Huberman, editor, *The Ecology of Computation*, pages 77–115. North Holland, Amsterdam, 1988.

Bernardo A. Huberman and W. Wolff. Finite precision and transient behavior. *Phys. Rev. A*, 32:3768–3771, 1985.

Michael D. Intriligator and Eytan Sheshinski. Toward a theory of planning. In W. P. Heller, R. M. Starr, and D. A. Starrett, editors, *Social choice and public decision making: Essays in honor of Kenneth J. Arrow, Volume 1*, pages 135–158. Cambridge University Press, Cambridge, 1986.

E. Ising. *Z. Physik*, 31:253–261, 1925.

Stuart Kauffman. Metabolic stability and epigenesis in randomly constructed genetic nets. *Journal of Theoretical Biology*, 22:437–467, 1969.

J. O. Kephart, T. Hogg, and B. A. Huberman. Dynamics of computational ecosystems. *Physical Review A*, 40:404–421, 1989.

Jeffrey O. Kephart, Bernardo A. Huberman, and Tad Hogg. Can predictive agents prevent chaos? In *Proceedings of CECOIA 2, Second International Conference on Economics and Artificial Intelligence, July 2-6, 1990*.

C. P. M. Knipscheer and Toni C. Antonucci, editors. *Social Network Research: Substantive Issues and Methodological Questions*. Swets and Zeitlinger, Amsterdam, 1990.

Steven E. Koonin. *Computational Physics*. Addison-Wesley, Redwood City, Calif., 1990.

C. J. Krebs. *Ecology*. Harper and Row, New York, 1972.

Christopher G. Langton. Studying artificial life with cellular automata. In D. Farmer, A. Lapedes, N. Packard, and B. Wendroff, editors, *Evolution, Games and Learning*. North Holland, 1986.

Christopher G. Langton. Computation at the edge of chaos: Phase transitions and emergent computation. In S. Forrest, editor, *Emergent Computation: Self-organizing, Collective and Cooperative Phenomena in Natural and Artificial Computing Networks*. MIT Press, 1991.

Samuel Leinhardt, editor. *Social Networks: A Developing Paradigm*. Academic Press, Inc., New York, 1977.

W. A. Little and Gordon L. Shaw. A statistical theory of short and long term memory. *Behavioral Biology*, 14:115–133, 1975.

E. Lumer and B.A. Huberman. Dynamics of resource allocation in distributed systems. Technical Report SSL-P90-00007, Xerox Palo Alto Research Center, 1990.

T.W. Malone, R.E. Fikes, K.R. Grant, and M.T. Howard. Enterprise: A market-like task scheduler for computer environments. In B.A. Huberman, editor, *The Ecology of Computation*, pages 177–206. North-Holland, Amsterdam, 1988.

Gerald Marwell, Pamela E. Oliver, and Ralph Prahl. Social networks and collective action: A theory of the critical mass. iii. *American Journal of Sociology*, 94:502–534, 1988.

Mark S. Miller and K. Eric Drexler. Comparative ecology: A computational perspective. In B. A. Huberman, editor, *The Ecology of Computation*, pages 51–76. North Holland, Amsterdam, 1988a.

Mark S. Miller and K. Eric Drexler. Markets and computation: Agoric open systems. In B. A. Huberman, editor, *The Ecology of Computation*, pages 51–76. North Holland, Amsterdam, 1988b.

J. D. Murray. *Mathematical Biology*. Springer, 1989.

Nils J. Nilsson. *Principles of Artificial Intelligence*. Tiaoga Publishing Co., 1980.

Martin A. Nowak and Robert M. May. Evolutionary games and spatial chaos. *Nature*, 349:826–829, 1992.

Pamela E. Oliver and Gerald Marwell. The paradox of group size in collective action: A theory of the critical mass. ii. *American Sociological Review*, 53:1–8, 1988.

Mancur Olson. *The Logic of Collective Action*. Harvard University Press, Cambridge, 1965.

Ei-Ichi Osawa and Mario Tokoro. Collaborative plan construction for multiagent mutual planning. In *Proceedings of the 3rd European Workshop on Modeling Autonomous Agents in a Multi-Agent World MAAMAW '91*, 1991.

Roy Radner. Dynamic games in organization theory. *Journal of Economic Behavior and Organization*, 16:217–260, 1991.

Eric Rasmusen. *Games and Information*. Cambridge University Press, Great Britain, 1989.

L.E. Reichl. *A Modern Course in Statistical Physics*. Univ. Texas Press, Austin, Texas, 1987.

Thomas C. Schelling. *Micromotives and Macrobehavior*. W. W. Norton and Company, Inc., United States, 1978.

Heinz-Georg Schuster. *Deterministic Chaos*. VCH Verlagsgesellschaft, 1988.

W. Richard Scott. *Organizations: Rational, Natural, and Open Systems*. Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1992.

Reinhardt Selten. Re-examination of the perfectness concept for equilibirium points in extensive games. *International Journal of Game Theory*, 4:25–55, 1975.

G. K. Shaw. *Rational Expectations: An Elementary Exposition*. Wheatsheaf Books, Ltd., Great Britain, 1984.

Karl Sigmund. On prisoners and cells. *Nature*, 359:774, 1992.

Herbert Simon. *The Sciences of the Artificial*. Massachusetts Institute of Technology, Cambridge, 1962.

Steve Smale. The prisoner's dilemma and dynamical systems associated to non-cooperative games. *Econometrica*, 48(7):1617–1634, 1980.

Adam Smith. *The Wealth of Nations*. Random House, New York, 1937.

Vernon Smith, Gerry L. Suchanek, and Arlington W. Williams. Bubbles, crashes and endogenous expectations in experimental spot asset markets. *Econometrica*, 56:1119–1151, 1988.

M. M. Srinivas and A. Beteille. Networks in indian social structure. *Man*, 64:165–168, 1964.

Masuo Suzuki. Scaling theory of transient phenomena near the instability point. *Journal of Statistical Physics*, 16:11–32, 1977.

Pablo Tamayo and Hyman Hartman. Cellular automata, reaction-diffusion systems, and the origin of life. In C. G. Langton, editor, *Artificial Life*. Addison Wesley, Amsterdam, 1989.

Michael Taylor. *Anarchy and Cooperation*. John Wiley and Sons, New York, 1976.

Michael Taylor. *The Possibility of Cooperation*. Cambridge University Press, Cambridge, 1987.

N. G. van Kampen. *Stochastic Processes in Physics and Chemistry*. North-Holland, Amsterdam, 1981.

John von Neumann and Oskar Morgenstern. *The Theory of Games and Economic Behavior*. Wiley, New York, 1944.

C. A. Waldspurger, T. Hogg, B. A. Huberman, J. O. Kephart, and S. Stornetta. Spawn: A distributed computational economy. *IEEE Transactions on Software Engineering*, 18(2):103–117, February 1992.

Colin P. Williams and Tad Hogg. The typicality of phase transitions in search. *Computational Intelligence*, 9(3), 1993.

Stephen Wolfram. Cellular automata as models of complexity. *Nature*, 311:419–424, 1984.

Gilad Zlotkin and Jeffrey S. Rosenschein. Conflict resolution in non-cooperative domain. In *Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI-90)*, pages 100–105, 1990.