

Coevolution to the Edge of Chaos: Coupled Fitness Landscapes, Poised States, and Coevolutionary Avalanches

STUART A. KAUFFMAN† and SONKE JOHNSEN

Department of Biochemistry and Biophysics, School of Medicine, University of Pennsylvania and Sante Fe Institute, Sante Fe, New Mexico, U.S.A.

(Received on 20 April 1990, Accepted in revised form on 8 August 1990)

We introduce a broadened framework to study aspects of coevolution based on the NK class of statistical models of rugged fitness landscapes. In these models the fitness contribution of each of N genes in a genotype depends epistatically on K other genes. Increasing epistatic interactions increases the rugged multi-peaked character of the fitness landscape. **Coevolution is thought of, at the lowest level, as a coupling of landscapes such that adaptive moves by one player deform the landscapes of its immediate partners.** In these models we are able to tune the ruggedness of landscapes, how richly intercoupled any two landscapes are, and how many other players interact with each player. **All these properties profoundly alter the character of the coevolutionary dynamics.** In particular, these parameters govern how readily coevolving ecosystems achieve Nash equilibria, how stable to perturbations such equilibria are, and the sustained mean fitness of coevolving partners. In turn, this raises the possibility that an evolutionary *metadynamics* due to natural selection may sculpt landscapes and their couplings to achieve coevolutionary systems able to coadapt well. The results suggest that sustained fitness is optimized when landscape ruggedness relative to couplings between landscapes is tuned such that Nash equilibria just tenuously form across the ecosystem. **In this poised state, coevolutionary avalanches appear to propagate on all length scales in a power law distribution.** Such avalanches may be related to the distribution of small and large extinction events in the record.

1. Introduction

Our aim in this article is to describe a new class of models with which to investigate some of the problems of coevolution. The class of models is related to the game theoretic models introduced by Maynard Smith & Price (1973) and Maynard Smith (1982). These authors, and many since them, have been concerned primarily (but not exclusively, see Rosenzweig *et al.*, 1987) with intraspecies coevolution, and **the conditions under which coevolving systems attain evolutionary stable strategies, ESS.** Our focus is largely on a general class of models for interspecific coevolution, and the conditions to attain Nash equilibria in such systems.

We consider coevolution occurring among species each of which is itself adapting on a rugged multi-peaked fitness landscape. But the fitness of each genotype, via the phenotype, of each species is affected by the genotypes via the phenotypes of the species with which it is coupled in the ecosystem. Adaptive moves by one

† Address correspondence to: Dr S. A. Kauffman, Department of Biochemistry and Biophysics, University of Pennsylvania School of Medicine, Room 411, Anatomy/Chemistry Bldg., Philadelphia, PA 19104-6059, U.S.A.

coevolutionary partner, therefore, may change the fitness and the *fitness landscapes* of the coevolutionary partners. Anecdotally, development of a sticky tongue by the frog alters the fitness of the fly, and what changes it must now make to increase its fitness; given the frog's sticky tongue, the fly should now develop slippery feet. In this framework, adaptive moves by any partner may *deform* the fitness landscapes of other partners.

To investigate this, we reintroduce the *NK* family of rugged multi peaked fitness landscapes (Kauffman 1989*a, b*, 1990; Kauffman *et al.*, 1989; Kauffman & Weinberger 1989; Weinberger, 1990). In this model, *N* corresponds to the number of genes in a genotype, or traits in an organism, while *K* corresponds to the number of other genes, or traits, which have a bearing on the fitness contribution of each gene or trait. *K*, therefore, corresponds to the richness of epistatic linkages in the system. Tuning *K* from low to high (increasing epistatic linkages) increases the ruggedness of fitness landscapes by increasing the number of fitness peaks, it increases the steepness of the sides of fitness peaks, and *decreases* the typical heights of fitness peaks. The decrease reflects the conflicting constraints which arise when epistatic linkages increase. Thus, the *NK* model provides a tunably rugged family of model fitness landscapes. To study coevolution we couple the fitness landscapes of different interacting species. To do so we suppose that genes or traits in each species make fitness contributions which depend upon *K* other genes or traits within that species itself, but also upon *C* traits in each of the other species with which the species interacts. Therefore, adaptive moves by one species may deform the landscapes of its partners. Altering *C* changes how dramatically adaptive moves by each species deform the landscapes of its partners. The other major parameters in the model are the total number of species which interact, *S*, and the number of these, *S_i*, which may be restricted with which any species, *i*, interacts. One form of the model we investigate ignores population dynamics and focuses upon the evolution of genotypes. A second form includes population dynamics using the familiar Lotka-Volterra logistic equation allowing the evolution of competition or mutualism among the coevolving species.

The framework we consider permits us to study how the ruggedness of fitness landscapes, the richness of coupling among fitness landscapes, the number of species, and the structure of the ecosystem, affect the coevolutionary process. To our knowledge, this is the first model which has allowed these questions to be investigated more or less systematically. We find a number of interesting and apparently novel features of such a process. Among these, the sustained fitness of the coevolutionary partners depends upon all the parameters. Therefore, we are driven to consider the possibility of selective *metadynamics* in evolution in which coevolutionary partners change not only their genotypes in an effort to optimize fitness on a given deforming fitness landscape, but may also change the statistical character of their fitness landscapes by changing its ruggedness, or may change the richness of couplings to other landscapes, and change the number of other partners with whom they coevolve so as to optimize sustained fitness. Rather remarkably, in distributed ecosystems where each partner interacts with only a few other species, this metadynamics appears to lead to a critical "poised" state in which the entire system is tenuously

“frozen” at Nash equilibria. Here each species is maximally fit given the genotypes of its partners, and is characterized by an unchanging, locally optimal genotype and stable population density. But at this poised state avalanches of coevolutionary change unleashed by minor changes in the physical environment, or other exogenous noise, propagate through the system on all length scales, with a power law distribution between sizes of avalanches and numbers of avalanches at each size. Such avalanches typically include fluctuations to lower fitness, hence they are likely to be correlated with increased probabilities of extinction events. Thus, the size distribution of avalanches can be used to try to predict distribution of sizes of extinction events in the record. The observed distribution tends to support a picture of ecosystems which are slightly more rigidly “frozen” than the critical poised state.

In section 2 we reintroduce the *NK* model and extend it to model coevolving species. In section 3 we describe the dynamics of coevolving pairs of species as a function of *K* and *C*. In section 4 we extend the results to distributed ecosystems where each species interacts with only a few of the total. In section 5 we extend the results to include population dynamics via a version of the Lotka-Volterra logistic equation.

2. Coupled *NK* Fitness Landscapes

The *NK* family of rugged fitness landscapes were introduced (Kauffman 1989*a, b*, 1990; Kauffman *et al.*, 1989; Kauffman & Weinberger, 1989) to describe genotype fitness landscapes engendered by arbitrarily complex epistatic couplings. Consider an organism with *N* genes, each with *A* alleles. For simplicity, consider that each has only two alleles, 1 and 0. Alternatively, consider an organism with *N* traits, present or absent. The *fitness contribution* of each gene or trait depends upon itself and epistatically on *K* other traits. Thus, each trait makes a fitness contribution which depends upon the particular combination, one among $2^{(K+1)}$ of the presence or absence of the *K* + 1 traits which bear upon its fitness. The *NK* model is intended to capture the effects of such epistatic coupling. It is similar to physicists' spin-glass models, which assume random energy interactions among spins in a physical material and study the generic properties of the complex energy surfaces in such coupled spin systems (Derrida, 1981; Binder & Young, 1986). The *NK* model is like a spin-glass model in assuming that the epistatic interactions among genes are so complex that their effects on fitness can best be captured by assigning those effects at random. Thus, the *NK* family models the statistical structure of such fitness landscapes by assigning at random to each such combination of “inputs” to the *i*th trait a fitness contribution between 0.0 and 1.0. Then for each organism, the fitness contribution of each of its *N* traits in the context of itself and the *K* which bear upon it are added together. Thereafter, the fitness is normalized through division by *N*. Figure 1 shows a simple organism with three traits, or genes, each of which makes a fitness contribution depending upon itself and the other two. The model assigns a fitness to each of the eight genotypes possible.

A central idea is that of a fitness landscape. Figure 1 shows all $2^3 = 8$ genotypes arranged on a Boolean cube, such that each genotype is a 1-mutant neighbor of the

three others which are accessible by mutating a single gene, or trait, to the opposite allele, or state. Thus, (000) is a 1-mutant neighbor of (001), (010), and (100). Arrangement of the eight genotypes onto the Boolean cube supplies a genotype space, in which each genotype is next to its 1-mutant neighbors. For an N gene, two allele system, the corresponding space is an N -dimensional hypercube. The NK model yields a fitness for each of the eight genotypes, shown in Fig. 1. These fitnesses can be thought of as a heights in a fitness landscapes of the kind initially introduced by Wright (1932).

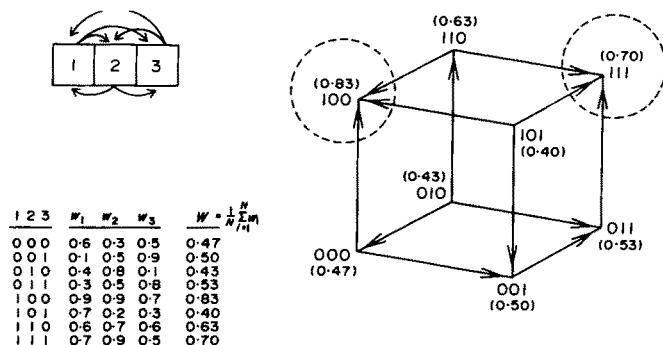


FIG. 1. (a) Assignment of $K = 2$ epistatic "inputs" to each site. (b) Assignment of fitness values to each of the three genes with random values for each of the eight combinations of $K + 1$ alleles bearing on genes 1, 2, and 3. These fitness values then assign a fitness to each of the $2^3 = 8$ possible genotypes as the mean value of the fitness contributions of the three genes, as given in eqn (1). (c) Fitness landscape on the three-dimensional Boolean cube corresponding to the fitness values of the eight genotypes in (b). Note that more than one local optimum exists.

The simplest picture of an adaptive process in a fixed fitness landscape envisions beginning at a genotype and "moving" to a 1-mutant neighbor only if the second genotype is fitter than the first (Maynard Smith, 1970). Then an adaptive walk starts at a genotype and climbs uphill until a genotype which is fitter than all 1-mutant neighbors, a local optimum, is encountered. Figure 1 shows two such local optima. Natural statistical features of landscapes include the number of local optima, the average number of steps to local optima, the number of local optima accessible from a random initial genotype, the number of genotypes which can climb to the global, or any other optima, the number of directions "uphill" from any genotype and how that number dwindles to zero as adaptive walks climb to optima. Other questions ask how these properties change if the adaptive walk can pass via 2-mutant or J -mutant variants, or if it may pass through equally fit or less fit neighbors. Still other questions concern how a population will flow across such a landscape under the drives of mutation, selection, and recombination. Here, as a function of the ruggedness of the landscape, complex population flows occur, with the population able to climb to local peaks and hover in their vicinity at low mutation rates, but melting from those peaks and spreading ever wider reaches of the genotype space as mutation rates increase (Schuster, 1987; Eigen *et al.*, 1988; Fontana & Schuster, 1987; Kauffman, 1989a, 1990).

All these properties depend upon the statistical structure of the underlying fitness landscape. The NK model was introduced because it provides a family of landscapes which range from single peaked and highly correlated to fully random and uncorrelated as K increases from 0 to its maximum, $N-1$. For $K=0$, each gene or each trait makes an independent fitness contribution. The model corresponds to an N locus two allele additive haploid genetic model. Each gene has a favored allele, by chance 1 or 0. Thus, there is an optimal genotype in which each gene has mutated to its optimal allele. For example, without loss of generality, the "1" allele of each gene might contribute more to fitness than the "0" allele. Then for $N=6$, the (111111) genotype is obviously the optimal genotype. Further, any other genotype can "climb" to this optimum by mutating the less fit to the more fit allele of each gene. Hence, there are no optima other than the single and global optimum. On each step uphill, the number of ways uphill decreases by 1. For example, in passing from the least fit genotype, (000000)–(100000), single gene changes from 0–1, hence the number of directions uphill has decreased by 1, from N in this case, to $N-1$. As adaptive walks approach the peak, (111111), the number of ways uphill decrease by 1 at each step, hence they decrease slowly. Since a randomly chosen genotype is half "1" alleles and half "0" alleles, and, without loss of generality we have assumed the optimal genotype to be (111111), on average, the number of steps to the global optimum scales as $N/2$. Further, the fitness of 1-mutant neighbors is nearly the same, since the greatest difference a single gene might make scales as $1/N$, which is small for large N , and reduces towards 0 as the number of genes, N increases. Thus, $K=0$ corresponds to a highly correlated landscape about a single global peak, a "Fujiyama" landscape. Furthermore, the gradient of the fitness slope to the peak is shallow: fitness increases by $1/N$ on each step uphill. Finally, it is easy to show from order statistics that in this model the global optimum has a fitness of $2/3$.

The maximum value of K is $N-1$. For $K=N-1$ the landscape is fully uncorrelated. That is, the fitness of 1-mutant neighbors are random with respect to one another. This can be seen in Fig. 1. Each of the three genes depends upon itself and the other two genes. Thus, changing any gene's allele, say from 0 to 1, changes the combination of allele states affecting each of the three genes to a new combination. For each combination of alleles of the $K+1$ genes which affect each gene, that gene makes a different, randomly assigned fitness contribution. Since the fitness of the new 1-mutant genotype is the mean of N new random variables, it is fully random with respect to the original genotype. Thus, for $K=N-1$, the fitness values of genotypes are fully random. The fitness of each genotype carries no information about the fitness of its 1-mutant neighbors. The landscape is completely *uncorrelated*.

The $K=N-1$ fully random fitness landscape is highly multi-peaked and rugged. It corresponds to the Derrida's (1981) random energy spin-glass model and was analyzed by Kauffman & Levin (1987) and Macken & Perelson (1989). Such landscapes have a very large number of local peaks, on the order of $2^N/(N+1)$. Because there are many local optima, adaptive walks to optima are very short. Indeed, expected walk lengths equal $\ln N$. Thus, a system with $N=10$ genes would typically harbor walks of two or three steps to optima, and a vastly larger system

with $N = 10\,000$ genes would have walks only of about nine steps. Unlike "smooth" $K = 0$ landscapes where, at each step, the number of ways uphill decreases by 1, on $K = N - 1$ random landscapes at each step uphill the number of ways uphill drops by half. A direct implication of the multi-peaked character of such fitness landscapes and the rapid dwindling of remaining directions uphill as walks progress, is that if an initial walk branches along more than one route uphill at each step, as a phylogeny might, then the resulting phylogenetic tree will be "bushy". Initial low fit genotypes might branch in many directions, but the rate of branching must dwindle as fitness increases until high fit genotypes barely manage to find single directions further uphill. In short, something like radiation and stasis is built into the structure of fixed rugged fitness landscapes. Further, due to the rugged multi-peaked character of the landscape, any initial genotype can climb to only a small fraction of the local optima, and only a small fraction can climb to the global optimum. An adapting population tends to become trapped in local regions of the multi-peaked fitness landscape. It follows that selection has a hard time searching across very rugged fitness landscapes. A further critical feature of the NK model is that for $K = N - 1$, the optima are *lower* than for $K = 0$. Indeed, a kind of complexity catastrophe occurs: as N , hence K , increases, the fitness of typical local optima dwindle towards 0.5, the mean fitness in the space. This decrease reflects the *conflicting constraints* among the K genes impinging upon each gene, which makes it increasingly difficult to find allele states of the set of genes which yields high overall fitness.

These results show that as K , which is a measure of the richness of epistatic interaction among the genes or traits, increases, fitness landscapes change from smooth and single peaked to random, and multi-peaked. Thus, as K changes, a family of increasingly rugged multi-peaked landscapes is encountered. In general, as K increases for fixed N , the number of local optima increases, the sides of fitness peaks become steeper, and the heights of typical local optima decrease. As we shall see, these features appear fundamental to coevolutionary behavior when NK landscapes are coupled.

The NK model is a very general model of a family of tunably rugged, correlated fitness landscapes. It, together with rather related spin-glass models, are the first explicit examples where such a family has been examined. It is, at this point, entirely unclear how many such families of rugged correlated landscapes may exist. However, the NK family seems likely to be an important member of such a set of families. Furthermore, the potential biological utility of the NK has recently been enhanced by showing that it is able to account for the statistical features of protein evolution seen in maturation of the immune response (Kauffman & Weinberger, 1989). Because the NK model is both the first, and appears a plausible, model of rugged fitness landscapes which may apply in biological evolution, we extend it here to study coevolutionary processes.

Consider an ecosystem with S species. For simplicity imagine that each species is homogeneous, that is, that all organisms in the species are identical, hence the species currently occupies a particular combination of its N traits. Then an NK landscape can represent the fitness landscape of one "homogeneous" species. This

assumption corresponds to the limit described by Gillespie (1983, 1984), in which mutations arise infrequently in an adapting population compared to the fitness differences between initial and mutant forms. Under these conditions, the population encounters advantageous mutants on a slow time scale and moves as a whole on a short time scale to the new fitter variant. Thus, under these limiting cases, the population can be approximated as homogeneous. With a higher mutation rate, or under other conditions, a species is not homogeneous. It is possible to extend the model to allow the population representing one species to be a cloud distributed over its landscape in which frequency and density dependent coevolution within the species occurs.

In a coevolutionary system we need to represent the fact that both the fitness and the *fitness landscape* of each species is a function of the other species. Thus, in general, it is necessary to *couple* the rugged fitness landscapes for each species, such that an adaptive move by one species "projects" onto the fitness landscapes of other species and alters those fitness landscapes more or less profoundly. Over time, each species jockeys uphill on its own landscape, and thereby deforms those of its ecological neighbors. Any such move by one may increase or decrease the fitness of each neighbor on its own landscape, and alter the uphill walks accessible to that neighbor.

In the context of the *NK* model the natural way to couple landscapes is to assume that each trait in species 1 depends epistatically on K other traits internally, and on C traits in species 2. More generally, in an ecosystem with S species, each trait in species 1 will depend upon K traits internally and on C traits in each of the S_i among the S species with which it interacts. It is also natural to assume symmetry, if species 1 is in the niche of species 2, then 2 is in the niche of species 1. We stress that this assumption does *not* imply a symmetry of effect. We assume that species 1 and 2 must each be in the other's niche, but species 1 may help species 2 while species 2 may harm species 1.

Consider the "frog" as species 1, and the "fly" as species 2. The effects of the frog's N genes upon the frog's fitness is given by its *NK* landscape. Similarly, the effects of the fly's N genes upon its fitness is given by its own *NK* landscape. The C traits in species 1 which affect each trait in species 2 can be thought of as C arrows drawn from each such trait in species 1 to the corresponding trait in species 2. Conversely, the C traits in species 2 which impinge upon each trait in species 1 can be thought of as C arrows drawn from each such trait in species 2 to the corresponding trait in species 1. To represent the effect of the C traits from species 1 which are coupled to each trait in species 2, we expand the fitness tables defining the landscape of species 2 to incorporate the added C traits which couple to each trait in species 2. Hence, for each of the N traits in species 2, the model will assign a random fitness between 0.0 and 1.0 for each combination of the K traits internal to species 2 together with all combinations, present or absent, of the C traits in species 1. In the expanded tables, each trait in species 2 makes a randomly assigned fitness contribution for each of the $2^{(K+1+C)}$ combinations of states of the $K+1$ internal traits in species 2 and the C external traits in species 1. In short landscapes are coupled by expanding the random fitness table for each trait in

species 2 such that it "looks" at its K internal epistatic inputs and also "looks" at the C external epistatic inputs from species 1. Given this, then the fitness landscape of species 2 is a function of the current location of species 1 on species 1s own fitness landscape. Therefore, as species 1 adapts and changes its own traits or genotype, those alterations will change both the fitness of species 2, and also deform 2s fitness landscape. In turn, each trait in species 1 is coupled to C traits in species 2, and the fitness values for each of the N traits in species 1 must be expanded similarly. This couples the two landscapes. Each is in the niche of the other. Moves by each tend to deform the fitness and fitness landscape of the other. As we see below, the resulting coevolutionary process depends critically upon K and C . K governs how multi peaked and rugged landscapes are, C governs how much one partner's landscape buckles when another partner makes an adaptive move. If K and C are properly adjusted relative to one another, coevolutionary "improvement" is easy; if not, landscapes buckle faster than coevolving entities can respond adaptively. All partners do poorly.

In a system of S species, the interactions can be represented by a web of such projections. In so doing, we have at our disposal at least: (1) choice of the number and identity of the traits, C , which couple from one species to each single trait in another species. (2) The number and identity of other species among the S which project onto each species. (3) Finally, we have at our disposal the number of species, S , in the entire ecosystem.

In section 3 we shall assume that each species in an S species system is coupled directly to all other species. This richest coupling is undoubtedly unrealistic. A vast literature studies the hierarchical structure of food webs (e.g. Pimm, 1982). In section 4 we consider "structural" ecosystems where each species interacts with only a few of the other species.

We also assume a second "worst case" condition, namely that in an S species ecosystem there is no similarity between the species, hence the effects of species 1 on 2 and 3 on 2 are randomly assigned. In reality, rabbits and hares probably look much the same to a fox. Similarity of species presumably can be thought of as reducing the number of *effectively different species* with which each species interacts. Finally, we also consider a naive case from an ecological standpoint, namely that each coadapting partner interacts at each moment with all other partners. In section 5 we are more realistic and require that each species interacts with the others in the ecosystem in pair-wise combinations, using plausible population dynamics.

3. Landscape Ruggedness and Couplings Between Landscapes Tune Coevolution

The game theoretic models which have been explored to study coevolution have not as yet been built to take account of the *statistical ruggedness* of each of the coevolving partner's landscapes, the *richness of coupling* of those landscapes, and the implications of those features on the entire coevolutionary process. But surely these are major aspects of the problem. The NK model affords tunably rugged landscapes whose richness of coupling can also be tuned, hence we can study the influences of these factors on coevolution.

Simulations of coevolving systems were carried out under the assumption that each species “acts” *in turn*, in the context of the current state of the other species. In its turn, each species tries a random mutation and “moves” as a whole to that mutant variant if the variant is fitter. If the mutant variant is not fitter, the species does not move. Thus, if the species moves, this at least transiently increases the fitness of the species which has just moved, but may increase or decrease the fitness of its coevolving partners. In addition to this “random” dynamics, we also examine two alternative cases. In the “fitter” dynamics, each species in turn examines all its mutant variants and chooses at random one of the fitter variants if any exist. In the “greedy” dynamics, each species in turn chooses the fittest mutant variant.

3.1. NASH EQUILIBRIA

In general, such a coevolutionary process admits two behaviors. Either the partners keep dancing, or the coupled system attains a steady-state at which the local optimum of each partner is consistent with the local optimum of all the other partners via the “*C*” couplings. Such a steady-state is the analogue of a Nash equilibrium in the current context. We use the word “analogue” for the following reason. A true Nash equilibrium assumes that each agent can, at each moment, choose any one of its possible “actions”. In the present context, this corresponds to each species, in a single moment, altering its current genotype to any of the 2^N possible genotypes. In assuming that each species is able, at each moment, to mutate a single “gene” or trait, we are constraining the range of alternative genotypes, or actions, locally accessible to the species. Thus, the steady-states we will find are similarly constrained. In the remainder of this article we use the term “Nash” equilibria with the understanding that such equilibria are with respect to the mutant search range.

A second caution is required. We consider Nash equilibria, as described. The concept of an evolutionary stable strategy, an ESS, is a further refinement of the concept of a Nash equilibrium, in which the condition of non-invadability by a mutant at an *initial low frequency* in the population is analyzed. In the simplified dynamics used here where the whole population moves in an “instant” to a fitter variant, we do not analyze invadability. Studies with fuller population dynamics, section 5, which do include analysis of invadability, confirm the simpler dynamics.

Simulations were carried out between pairs of coevolving species, or “agents”, each a single organism on an independent *NK* landscape. In addition, simulations were carried out for larger numbers of species, *S*. The first major result is that Nash equilibria are actually encountered. It is not obvious that this should occur, for each species has 2^N genotypes among which it is evolving. An *S* species system has the product of these genotypes in its joint “strategy space”.

Figure 2 shows eight species, each with $N = 24$, $K = 17$, and $C = 1$, coevolving over 2500 generations. Here, then, each site within one species is epistatically affected by 17 other sites within that species and one site in each of the eight other species. Over eight generations, each species in turn tries a random mutation and moves to that new genotype only if it is fitter than the current genotype in the context of the current genotypes of the remaining seven species. At each generation, the fitness of




FIG. 2. Coevolution among eight "species" each governed by an NK landscape. $N = 24$, $K = 13$. Each of the N "traits" in each species is affected by $C = 1$ "trait" in each of the seven other species. System reaches a steady-state about generation 1600. Note, mean fitness in the absence of selection is 0.5. See text.

all species are recalculated in the context of the genotype of each and their couplings. As can be seen, for the first few hundred generations, the mean fitness of the whole set of species increases, rapidly at first, then more slowly. Increasingly long intervals with no change occur, reflecting the fact that as fitness increases, the waiting time to find fitter variants increases for each partner. Sudden bursts of change by many species, however, are instigated by occasional changes by a single partner. By about 1600 generations, however, the entire system stops changing, and in fact remains constant forever thereafter. A Nash equilibrium has been found such that each species is locally fitter than all 1-mutant variants, granted that the others do not change.

3.2. WAITING TIME TO ENCOUNTER NASH EQUILIBRIA

The waiting time to encounter Nash equilibria depends upon N , K , and C . For $K > C$, Nash equilibria are encountered rapidly, for $K < C$ the waiting time to find Nash equilibria becomes very long. In order to examine how N , K , and C bear on the waiting time to encounter a Nash equilibrium, simulations were carried out between two species. For each value of N , K , and C tried, 100 coevolving pairs were released. Over generations, a successively larger fraction will have encountered Nash equilibria and hence stopped evolving. Figure 3 (a) and (b) shows the results, plotting the fraction "still walking", against the generations elapsed, for $C = 1$, Fig. 3(a) and for $C = 8$, Fig. 3(b).

The main point to note is that as K increases relative to C the waiting time to hit a Nash equilibrium *decreases*. Thus, as the ruggedness of landscapes increases, by K increasing, the expected waiting time to find Nash are decreasing. Presumably this reflects the increased number of local optima in NK landscapes as K increases




FIG. 3. (a) Fraction of 100 coevolving pairs of species which have not yet encountered a "Nash" equilibrium, hence still walking as a function of generations elapsed. Curves correspond to different values of K . $N=24$, $C=1$ in all cases. (b) Same as (a) except $C=8$.

for fixed N . Similar studies as N increases, data not shown, for fixed K and C , show that the waiting time to hit Nash equilibria increases, presumably because the density of local optima decreases as N increases. In short, for a pair of species which are coevolving, $K=C$ is a crude dividing line. When K is greater than C , Nash equilibria are found rapidly. When K is less than C , Nash equilibria are still found, but the mean waiting time becomes very long.

3.3. COEVOLUTION WHEN TWO INTERACTING SPECIES HAVE DIFFERENT K VALUES

It is of considerable interest to study the outcome of coevolution in which partners are on landscapes of different ruggedness. Figure 4 reports the results of simulations in which pairs of species have K values of 2, 4, 8, 12, and 16, for three values of

C , 1, 8, and 20. $N=24$ in all cases. These three values of C were chosen to lie below, in the middle, and above the range of K values. For each set of parameter values, 300 pairs of coevolving species was released, and evolved for 250 generations. By that time, Nash equilibrium may or may not have been encountered. In the former case, the fitness of each partner was noted. In the latter case, the average fitness of each partner over the prior 85 generations was tabulated. The left panels in Fig. 4 report the mean fitness of each partner at the Nash equilibria encountered, and the number which reached such Nash equilibria. The middle panels of Fig. 4 show the mean fitness of each partner as it continues to coevolve, over the last 85 generations,




FIG. 4. Results of coevolution among 300 pairs of species for 250 generations for $C=1$, $C=8$, and $C=20$ and K values of 2, 4, 8, 12, 16. $N=24$ in all cases. Left panels, "Yes", correspond to cases where "Nash" equilibrium were encountered. Each "cell" has three numbers. The upper number is the mean fitness of the "Row" player. The middle number is the fitness of the "Column" player. The lowest number is the fraction of the 200 pairs of players which encountered a "Nash" equilibrium. The middle panels, "No", show cases where no "Nash" equilibrium was encountered in 250 generations. The three numbers in each "cell" are the same except that the fitness values correspond to mean fitness over the last 85 generations. At each generation a random mutant was tried; the species moved to that variant only if it was fitter. Right panels "all" show average fitness over generations whether or not Nash equilibria were encountered.

and the number of pairs still coevolving. The right panels show the mean fitness of all partners whether or not Nash equilibria were attained.

Figure 4 shows a number of interesting features:

- (1) For all values of K , as C increases, the fraction of coevolving pairs which encounter Nash equilibria in 250 generations decreases. Conversely, for any fixed value of C , as K increases, the fraction of pairs encountering Nash equilibria increases. This re-exhibits the phenomena of Fig. 3. High K leads to more rugged landscapes and Nash equilibria are encountered more rapidly.
- (2) When C is higher than 1, i.e. $C = 8$, $C = 20$, the fitness at Nash equilibria is higher than the corresponding fitness when the partners are still oscillating.
- (3) As C increases, the fitness of both coevolving partners during the oscillatory phase before encountering Nash equilibria *decreases* for all pairs of coevolving K values. Thus, fitnesses during the preNash oscillatory period for $C = 20$ are substantially lower than for $C = 8$ and still lower than for $C = 1$. Presumably this reflects the fact that, for high values of C , a single move by one partner sharply lowers the expected fitness of the remaining partner.
- (4) When C is high, $C = 20$, note that the high K "players" have higher mean fitness during the oscillatory period before hitting Nash equilibria than do low K "players". More strikingly, in playing against a second player with a fixed K value, the first player would increase mean fitness during the oscillatory period by increasing its own value of K . That is, a $K4$ player does better against a $K2$ player than would a $K2$ player, while a $K8$ player does even better, etc.
- (5) Equally remarkably, when C is high, $C = 20$, a low K player achieves higher mean fitness during the preNash oscillatory period if it plays against a second species of high K . That is, a $K2$ player has higher fitness against a $K4$ player than against a $K2$ player. A $K2$ player fares even better against a $K16$ player. Thus, when C is high, increasing the K value of one partner helps *both* coevolving partners.
- (6) This tendency is reversed when C is low, $C = 1$. Here during the oscillatory period, low K players fare better than high K players.
- (7) At the Nash equilibria encountered, the fitnesses of low K players is clearly higher than that of high K players for each value of C , and indeed seem roughly independent of C .
- (8) Finally, when C is high, overall average fitness is highest when K is high. When C is low, $C = 1$, overall average fitness is highest when K is low. Thus, fitness in coevolving systems would be enhanced were K able to adjust to match C , or more broadly, if K and C were themselves evolvable.

Similar studies were carried out using the "fitter" and the "greedy" dynamics described above. The main results are the same. A major difference arises in the greedy dynamics. Here, for any genotype, there is generically a unique best fit 1-mutant variant in the context of the other species. Thus, if each species "plays" in a deterministic order, each changes to a unique next genotype, and the set of coevolving species can enter a recurrent cycle in the total space of S genotypes. We find that when $K < C$, such periodic attractors arise rather frequently.

The analysis above is based on mutating a single site, gene, or trait, at a time in each coevolving species. We have also examined the sustained fitness of coevolving species as a function of the number of genes or traits which can mutate at each moment. In general, for any value of N , C , and K , an optimal mutation rate exists. Figure 5(a) shows the results of coevolution as the mutation rate, or more accurately, the number of traits randomly mutated in each species, increases to 2, 4, 8, 16, and 24, for coevolving pairs of species on increasingly rugged landscapes, $K = 2, 4, 8, 12, 16$. In all cases each of the $N = 24$ sites in each species is coupled to $C = 1$ site in the other species. As the number of genes mutated simultaneously increases, the number of local optima on a fixed landscape dwindle, hence the probability that the coevolving pair reaches a local Nash equilibrium falls. Thus, the fitness seen in the coevolving pair reflects fitness during the preNash period. Values are averages of the last 85 steps in 250 generations. Figure 5(a) shows that for all values of K , as the number of genes mutated increases, the maintained fitness reaches a maximum for two or four simultaneous mutations, and falls thereafter. The decrease in maintained fitness is greatest for small K values, hence smoother landscapes, and less marked for large K values. Figure 5(b) and (c) shows similar results except that the coupling among the coevolving players, C , is increased to eight and 20. In general, but not uniformly, as C increases the optimal mutation rate decreases. These results suggest that there may typically be a low optimal mutation rate to maintain fitness in coevolutionary processes.

This discussion has ignored the costs of mutational search to a population in terms of genetic load. As mutational search distance increases from genotypes of above average fitness on correlated landscapes, typically a smaller fraction of the mutant genotypes equal or exceed the initial genotype. Hence, load becomes greater as mutation rate increases. This would be expected to further reduce the optimal mutation rate in coevolution.

The results we have seen make a number of intriguing suggestions. All point to the possibility that a coevolving system of species may collectively tune the parameters governing its own coevolution.

First, there may well be selective processes which match " K " to " C " in order to optimize the coevolutionary capacities of the coevolving partners. K "should" increase when it is low relative to C , and decrease when it is high relative to C . As shown in Fig. 4 for the most biologically plausible case in which each species tries random mutations and moves to a variant if it is fitter, if C is high relative to K , any player increases its fitness by increasing its own K value. When C is high, increasing K has two beneficial effects. First, Nash equilibria are encountered more rapidly and are fitter than the prior oscillatory period. Second, fitness during that prior oscillatory period is higher. Thus, it is advantageous to any player to increase K in a high C environment. Perhaps equally remarkably, in the biologically reasonable case of random mutations, Fig. 4(a)–(c) such a move by one species *also helps the second species*. Each has higher preNash fitness and finds Nash equilibria sooner. Conversely, suppose K is high relative to C . Then, as is clear from Fig. 4(a)–(c) and Fig. 3, Nash equilibria are encountered very rapidly. Thus, the fitness in the preNash oscillatory period is of less importance, and the fitness of Nash equilibria are more important. But local optima at Nash equilibria are higher for low K players




FIG. 5. (a) Effects of increasing the mutation rate on mean fitness maintained in coevolving species. Both coevolving species randomly mutate 1, 2, 4, 8, 16, or 24 of the $N=24$ "genes" at each moment. $K=2, 4, 8, 12, 16$. $C=1$. Figure plots average fitness over last third of 250 coevolutionary steps against mutation rate. Note that for all K values, maintained fitness increases and reaches a maximum for two or four mutant search ranges. For all values of K , maintained fitness falls for greater mutation rates. (b) As in (a), except $C=8$. (c) As in (a), except $C=20$.

than high K players. Thus, if K is too high with respect to C , it should be advantageous to decrease K . In short, at this group level of coevolving species, it seems clear that there are reasonable selective advantages to a species as a whole to "tune" K to match C . At that match, given a fixed C , Nash equilibria will be encountered rapidly, and be highly fit, optimizing mean fitness during any preNash periods, minimizing the mean duration of those periods, and maximizing the fitness of the Nash optima attained.

The analysis above is based on advantage to the species as a group in increasing K . In order to avoid "group selection" we seek selective conditions acting on individual members of a single species which might increase K in members of that species, hence in the coupled ecosystem. Within the framework of the NK model, a change in " K " would naturally be envisioned as a mutation which altered the epistatic coupling between traits, or genes, such that a trait or gene now depended on one more, or one fewer epistatic "inputs". That is, we must let " K " itself evolve. In this framework, the natural way to express the consequence of such a mutation which increases K is to expand the fitness table for that trait such that it now "looks at" the new trait as well as the K it initially "looked at". That new epistatic connection, in the context of the current genotype in which the new connection is formed, might increase or decrease the fitness of the current genotype. That is, the new epistatic link may itself alter the fitness of the current genotype for better or worse. Thus, we can envision three ways in which selection on an individual level may allow an increased K value at one genetic locus to spread throughout a population: (1) that new epistatic link, when it forms, causes the genotype to be fitter and is selected, hence it spreads. (2) The new epistatic link is "near neutral" and spreads through the population by random drift. (3) The new link has not only a direct effect on the fitness of the current genotype, but also on the *inclusive fitness of the individual and its progeny* due to the increased fitness of those progeny in the *coevolutionary* process itself, due to increased rapidity of finding Nash equilibria, and higher fitness during the preNash oscillatory period. Presumably, organisms can evolve to alter not only the *structure* of their fitness landscapes, here by altering K , but can alter how badly the landscape is deformed when other partners make adaptive moves. In the current context, evolution can presumably alter C .

These considerations suggest the possibility of a coevolutionary dynamic which optimizes K and C relative to one another in an ecosystem, such that partners maintain high mean fitness.

Next we consider the coevolution of the number of species, S , in the coupled ecosystem. A clear process should tend to limit the number of species. Consider coupled landscapes with fixed K and C couplings. Let the number of species increase, under the assumption that each species is " C " coupled to all other species. For pairs of species, $K = C$ is a rough dividing line separating cases where Nash equilibria are encountered rapidly vs. slowly. In multispecies systems with S species, it appears that when K is greater than $S \times C$, the coevolving partners all encounter a Nash equilibrium rapidly. When K is less than $S \times C$, the coevolving partners do not encounter a Nash equilibrium for a long time. We stress that the exact relation between mean waiting time to encounter Nash equilibrium, and the K , S , and C parameters are unknown. $K = S \times C$ is a rough guide.

Numerical results are shown in Fig. 6(a)–(c) for an increasing number of coupled species. In all cases $N = 24$, $K = 10$, and $C = 2$. In Fig. 6(a) $S = 4$ species are coevolving. Thus, each species senses $C = 2$ inputs to each of its genes from each of the three other species. Each species in turn randomly mutates and tries to find a fitter variant. Over 2000 generations, mean fitness increases and the four species find a Nash equilibrium. In Fig. 6(b) and (c) the number of species increases to eight, and then to 16. Data to 2000 generations are shown for the eight and 16 species cases. No Nash was found for 8000 generations in these cases.

Note the following features in Fig. 6(a)–(c): (1) as the number of species, S , increases, the waiting time to encounter a Nash increases. (2) As S increases, the *mean fitness* of the coevolving partners *decreases*. (3) As S increases, the *fluctuations* in fitness of the coevolving partners increases dramatically. Thus, for the four species case, mean fitness increases rapidly, and achieves fairly high Nash equilibria. For the 16 species case, the entire system fluctuates about a mean fitness slightly above average, 0.5, with dramatic excursions below 0.4.

These results show that as the number of mutually coupled species in the system increases, the mean fitness falls and fluctuations to very low fitness increase. Therefore, if we may assume that fluctuations to low fitness are associated with an increased chance of *extinction* of the unlucky species, these results suggest that if S is too high, the coupled ecosystem will fluctuate dramatically and lead to the extinction and loss of species, thereby *lowering* S . In turn, as S is lowered, the ecosystem behaves less chaotically, mean fitness of all partners improves both during the preNash period and because Nash are encountered more rapidly. The remaining system coevolves “well” despite the fact that landscapes are coupled and deform as each actor moves.

In short, with clear hesitations, and heralded caveats, this framework begins to suggest the possibility that the coevolutionary parameters governing the ruggedness of landscapes, couplings among landscapes and number of coevolving partners might themselves coevolve without group selection, to continuously recreate well-formed ecosystems which are *successfully* able to coadapt. This is no mean feat, for as Fig. 6(c) makes clear, coevolution among coupled species can lead to chaotic fluctuations with no accumulation of improvement. If one wishes a “Red Queen”, here is one to reckon with.

4. Structured Ecosystems and Self Organized Criticality: Adaptation to the Edge of Chaos

Real ecosystems are not totally connected. Typically each species interacts with a subset of the total species, hence the system has some extended web structure. In this section we extend our results to such ecosystems. The supposition that a coevolutionary system can control the ruggedness of coupled landscapes, and its own connectivity by selection itself, and therefore its dynamics, has interesting implications for extended ecosystems. It might be the case that coevolving ecosystems tend toward a state of “self organized criticality” in which parts of the ecosystem are “frozen” at Nash equilibria for long periods such that the species in the frozen component do not change, while other species continue to undergo coevolutionary




FIG. 6. (a) Coevolution among four species. $N=24$, $C=2$, $K=10$. (b) As in (a), but eight species are coevolving. (c) As in (a), but 16 species are coevolving. Note that as the number of species increases, the mean fitness decreases and the variance in fitness increases.

changes. Cascades, or avalanches of changes initiated at local points in the ecosystem web may propagate to various extents throughout the ecosystem. Such avalanches may trigger speciation and extinction events. Furthermore, the endogenous dynamics of the coevolving system under natural selection optimizing sustained fitness for each coevolving partner, may tend toward this characteristic critical poised state in which such avalanches can propagate on a variety of size scales with a power law distribution between sizes of avalanches and their frequencies.

The term, "self organized criticality", was recently coined by physicist Bak (1988), to refer to a quite generic pattern of self organization. Bak asks us to consider a tabletop onto which sand is added at a uniform rate. As the sand piles up on the table, it begins to slide off the edges of the table. Eventually, the system reaches a steady-state at which the mean rate of adding sand equals the mean rate at which sand falls over the edges. At this stage the slopes from the peak to the edges of the table are near the rest angle for sand. Bak asks the following question: If one adds a single grain of sand to the pile at a random location and starts an avalanche, what will the *distribution of avalanche sizes be*? Bak finds a characteristic power law distribution relating the frequencies and sizes of avalanches, with many tiny avalanches and a few large avalanches. He argues that this distribution is characteristic of a wide range of phenomena, including distribution of earthquake sizes, and other examples. The argument requires that the system under investigation attain and maintain a kind of "poised" state which is able to propagate perturbations—avalanches—on all possible length or size scales.

There are at least three ideas, derived from Bak's theory, which seem interesting in the coevolutionary context. Cascades of perturbations, constituted by "packets" of coevolutionary change, with a characteristic relation between size scale and frequency, may propagate through an ecosystem. First, this possibility requires that parts of the ecosystem can be "at rest" while other parts change. Second, the propagating changes are likely to be associated with fluctuations to low fitness which may engender both extinction and speciation events. Extinction events might be expected because of low fitness itself. Speciation events would be expected because, at low fitness, the number of directions of improvement is increased. If the probability of branching speciation is proportional to the number of directions of improved fitness, then low fitness episodes should trigger speciation events. Thus, the propagation of avalanches through the system would be linked to speciation and extinction phenomena. Third, coevolutionary dynamics of linked speciation, extinction, and alterations in coupling among the species in the ecosystem may achieve ecosystems which are "poised" such that avalanches can propagate on a characteristic variety of size scales. We shall see that this poised property is likely to be associated with the existence of nearly melted "frozen" component in ecosystems.

4.1. FLUCTUATING FROZEN COMPONENTS: NASH EQUILIBRIA EXTENDED TO LATTICE ECOSYSTEMS

A first hint that such ideas may apply to ecosystems arises in extending the coevolving *NK* model to structured ecosystems in which each species interacts with

only a few of the other species: *parts of the system may be fixed at Nash equilibria while other parts continue to coevolve*. That is, some species can attain an equilibrium and stop coevolving, hence remain "at rest", while adjacent species in the ecosystem continue to change either transiently or persistently.

To begin to investigate the behavior of structured ecosystems, we carried out simulations on "square" lattice ecosystems in which each interior species interacts with its four neighbors. Corner and edge species interact with two and three neighbors respectively. Model ecosystems have varied from $3 \times 3 = 9$ species to $10 \times 10 = 100$ species. In addition to square ecosystem, which have corners and edges, we have investigated toroidal ecosystems in which the "square" is folded into a cylinder by joining left and right edge species, and then bent into a torus by joining top and bottom species. We have also investigated randomly connected ecosystems with similar general results.

Figure 7 shows 12 successive times in the temporal coevolution of a 10×10 ecosystem. In this study $N = 24$, $K = 10$, $C = 1$. At each time moment, one of the 100 species "plays" and "greedily" chooses the fittest 1-mutant variant if any is fitter than its current genotype. Each species "plays" in turn, and thus 100 "plays" constitutes an ecosystem generation. After each ecosystem generation, hereafter "generation", any species may have changed its genotype, or remained the same. If the species changed, we color it "white", if it remains unchanged, we color it "black". The simulation was run over 200 ecosystem generations. The panels in Fig. 7 represent moments which are four ecosystem generations apart in time, from a period in the middle of the 200 generations. The first question to ask is: Can "frozen" regions of "black" species which are unchanging in their genotypes, and other regions of "white" species which are changing their genotypes coexist in the ecosystem? The salient features to note are these:

- (1) A large fraction of the species are frozen and unchanging over single ecosystem generations.
- (2) Some regions remain frozen over very many ecosystem generations. In Fig. 7, species in the upper left and lower left corner remain frozen over about 48 ecosystem generations.
- (3) One or more "white" unfrozen regions may exist.
- (4) Over time, the location and size of the frozen "black" region waxes and wanes. That is, *a fluctuating frozen component* can exist and extend through some or much of the ecosystem.
- (5) In the simulation carried out here, ultimately, the "frozen" region encompasses the entire ecosystem. That is, the ecosystem comes to rest at a combination of genotypes which are local Nash equilibria for all 100 species. In the absence of exogenous perturbations, the system will remain in this frozen state thereafter.
- (6) In many simulations, particularly using the "greedy" algorithm, the coevolving ecosystem encounters a limit cycle. Typically in these cases, a fraction of the entire ecosystem remains permanently "frozen", while the remainder oscillate through a recurrent set of genotypes.

These results show that one region of an ecosystem can be "frozen" while other regions continue to coevolve. One region persists in a Nash equilibrium something




FIG. 7. Twelve successive time moments four ecosystem generations apart in a $10 \times 10 = 100$ species ecosystem where each species plays with its immediate neighbors. Thus, corner species coevolve with two immediate partners, edge species with 3, and interior species with 4. $N = 24$, $K = 10$, $C = 1$. As time progresses, frozen regions, black, where species are not changing genotype at that generation, emerge, expand and contract over the distributed ecosystem. If all species stop changing, the entire ecosystem is frozen, each species at a local Nash equilibrium.

like an evolutionary stable state, while adjacent regions in the *same ecosystem* persist in “Red Queen” antics.

The existence of frozen components in an ecosystem may bear on, indeed may be fundamental to *evolutionary stasis* for some species despite general changes in ecosystems altering other species. Some organisms may be maximally adapted to a fixed coevolutionary environment in such a fixed component, even though other species continue to undergo coevolutionary change.

4.2. A COEVOLUTIONARY ADAPTIVE PROCESS LEADING TOWARD A SELF ORGANIZED CRITICAL STATE WITH A PERCOLATING FROZEN COMPONENT

The results in the previous sections concerning unstructured ecosystems suggested the possibility that a coevolutionary dynamics might tune the parameters of the coevolving species such that the species coevolved “well”. Here we discuss preliminary results suggesting that these ideas may extend to structured ecosystems. Species may “selfishly” tune their K values such that the coupled system as whole coadapts well. At the optimal state, a frozen component just begins to percolate across and covers the ecosystem “tenuously”.

Figure 8(a) and (b) shows the results of simulation of 5×5 square ecosystems in which the average fitness of the corner species, each coupled to two other species, the “edge” species, each coupled to three other species, and the “interior” species, each coupled to four other species, are accumulated. In the simulations, K is varied from 0–22. For each value of K , 50 ecosystems were analyzed over 200 ecosystem generations each. Figure 8(a) corresponds to the “random” dynamics where a




FIG. 8. (a) Mean sustained fitness in 5×5 ecosystems as K varies from 0-22. $N = 24$, $C = 1$ in all cases. "Corner" species are connected to two others, top curve, edge species are connected to three others, middle curve, interior species are connected to four others, bottom curve. Note sustained fitness increases then decreases as K increases. "Random" dynamics were used, such that each species tried a random 1-mutant variant and "moved" there if that mutant were fitter. (b) As in (a), except the "fitter" dynamics were used. Each species picks at random one of its fitter variants if any exists, at each species generation.

random mutation is tried by each species. Figure 8(b) corresponds to the "fitter" dynamics, where one of the fitter variants is chosen at random by each species at each generation. These figures reveal the following features:

- (1) For all interspecies connectivities, two for the corner species, three for the edge species, and four for the middle species, there is an *optimum value of K* about $K = 8 - 10$. For lower or higher values of K , the average fitness declines.

- (2) Mean fitness *increases*, if connectivity to other species *decreases* from 4-3-2.

If there is an optimal value of K for sustained fitness in coevolving systems, then it might be the case that selective effects might "pull" K values of coevolving partners towards this optimum. This intuition is confirmed in Fig. 9. Here we




FIG. 9. Selection force towards a K_{opt} value of $K = 8-10$. Two "experimental" species located adjacent to the central species in the 5×5 ecosystem were constructed with a different K value than the remainder of the ecosystem. In all cases deviations of the experimental pair, "dot", compared to the rest of the ecosystem, small square, toward K_{opt} increased the sustained fitness of the experimental pair compared to the rest of the ecosystem. In contrast, deviation away from K_{opt} decreased the fitness of the experimental pair compared to the rest of the ecosystem.

investigate the effect on the fitness of two "test" species, located adjacent to the center of the ecosystem, due to increasing or decreasing " K " values, compared to the rest of the ecosystem. To sample fairly the rest of the ecosystem, we monitored the fitness of two "control" species, also adjacent to the center of the square ecosystem. Figure 9 shows that, in all cases, presence of two "test" species with an altered K value had little effect on the fitness of the two control species. More critically, if the rest of the ecosystem had suboptimal K values, namely 0 or 2, or above optimal K values, namely 14 or 23, then deviations of the K values of the test species *to or towards* the K optimal values of $K = 8-10$ *increased the fitness of the test species*. Conversely, if the ecosystem as a whole is at the K optimal values of $K = 8-10$, then deviations of the K values of the test species *away from* the optimal values *decreases* the fitness of the test species. In short, there is a selective force toward the K optimal value of 8-10 which can act on single species, presumably via individual members of that species, and pull each toward the jointly optimal K value.

These results support and extend those discussed in the previous section with respect to completely connected ecosystems. There we found evidence that it was advantageous for a single species to increase or decrease K towards an optimal value relative to C . The results in Fig. 9, are a powerful indication that a kind of selective metadynamics may very well tune K , the ruggedness of landscapes among coevolving species, towards a *joint optimum* where all partners coevolve well.

4.3. COEVOLUTION TO THE EDGE OF CHAOS: OPTIMIZATION OF SUSTAINED FITNESS AT K_{OPT} YIELDS ECOSYSTEMS WITH JUST PERCOLATING FROZEN COMPONENTS

Might optimization of the capacity to coevolve, by optimizing K relative to C , actually selectively attain a posed self-organized critical state? The answer appears

to be "yes" in the current model. Figure 9 shows that selection would be expected to pull coevolving partners to jointly exhibit a near optimal $K = 8 - 10$ value. Figure 10 shows that as K increases the rapidity with which the ecosystem becomes "frozen" at a Nash equilibrium increases as well. Thus, for $K = 8$ or less, no freezing of the entire ecosystem occurs in any of the ecosystems over 200 generations. For $K = 10$ entire ecosystems freeze at Nash equilibria gradually over the 200 generations, while for $K = 14$ or 22 ecosystems freeze rapidly. Thus, the K optimal value for sustained fitness, $K = 8 - 10$, occurs just at that value where freezing begins visibly to occur. Otherwise stated, model ecosystems optimize coevolutionary fitness when frozen components are tenuously extending across the ecosystem, hence when the system is "at the edge of chaos". By this phrase we mean that for values of K larger than K_{opt} , the ecosystem "freezes" into a "solid" state with all partners at Nash equilibria. For values of K smaller than K_{opt} , the ecosystem takes a very long time to achieve a Nash equilibria. The fitnesses of the species fluctuate chaotically during the long preNash period. Hence, the system is chaotic, and can be thought of as being in a kind of "gas" phase. Just at the interface between the solid and gas phase is a kind of "liquid" region, at the edge of chaos (Langton, 1990). Our results suggest that selection, in this model, achieves systems poised at the edge of chaos.

4.4. COEVOLUTIONARY AVALANCHES IN POISED ECOSYSTEMS: A POWER LAW DISTRIBUTION

The "edge of chaos" also corresponds to a poised self organized critical state with respect to coevolutionary avalanches. As seen above, as K decreases from above K_{opt} , the frozen component of all the species at a Nash equilibrium "melt". We investigate next the implications for coevolutionary avalanches, and find that




FIG. 10. Fraction of 5×5 ecosystems which have not yet become frozen in an overall Nash equilibrium plotted against ecosystem generation. Note that for $K = 8$ or less, none of the ecosystems attained a frozen Nash equilibria in the time available. For $K = 10$ or larger, some or most systems freeze at Nash equilibria, and do so more rapidly as K increases.

at the optimized state for sustained fitness, avalanches propagate on all length scales in a power law distribution.

Alterations at one site in an ecosystem can often be expected to cause neighboring species to undergo coadaptive changes. Thus, changes at one point may propagate the various extents throughout the ecosystem. Such avalanches of changes are the analogues to Bak's sand pile avalanches. The propagation of such changes through an ecosystem are clearly of interest. In particular, during such changes, the affected species are likely to fall transiently to lower fitness. If lowered fitness is associated with an increased probability of extinction, then such avalanches might be associated with the propagation of extinction events.

To begin to investigate such avalanches, we modified the ecosystem model to allow each species to be affected not only by its neighbors in the ecosystem, but also by its external "world". The "external world" of each species consists in a binary vector length N . Each site in the species is coupled to W sites in its world. The fitness table of each site in the species is augmented to "look at" the K internal sites, the C sites in each of the species impinging upon that species, and the W sites in that species' world. Thus, alteration in the world of one species deforms the fitness landscape that species is coevolving upon. Typically, such alterations lower the fitness of the species sufficiently so that its current genotype is less fit than one or more of the 1-mutant neighbor genotypes. If so, the species changes, and may then unleash an avalanche of coevolutionary change which propagates through the ecosystem.

The simplest avalanches of changes to visualize are those which are perturbations from the "frozen" state in which all species are at local Nash equilibria. Simulations using the "fitter" move dynamics found the frozen state, then the "world" of a random member of the ecosystem was changed. At the end of each ecosystem generation, each species may have remained the same, or altered its genotype. After the onset of such a change, coevolutionary changes continue until the system returns to a (perhaps new) frozen state with all species at local Nash equilibria. At that point, the "avalanche" has died out.

We used two measures of the "size" of such an avalanche. (1) The first is total number of species which are caused to alter their genotypes. This measures the total number of species in the avalanche which have changed genotype at least once. (2) An alternative measure sums the number of species which have changed at each ecosystem generation from the start of the avalanche until the avalanche stops. Thus, the second measure includes both the number of species which are affected, and the number of ecosystem generations in which each is affected. We denote this measure as (species \times time).

Figure 11(a)-(c) shows the resulting histograms of avalanche sizes for the 5×5 ecosystems discussed above. These figures plot the logarithm of the numbers of avalanches at each size vs. the logarithm of the sizes of avalanches in the species \times time measure of avalanche size. The number of avalanches measured are 228 for $K = 10$, 374 for $K = 14$, and 308 for $K = 20$. Figure 12(a)-(c) shows similar histograms of avalanche sizes for $10 \times 10 = 100$ species square ecosystems, with $N = 24$, and values of K of 12, 14, and 18. The number of avalanches measured are 119 for $K = 12$, 402 for $K = 14$ and 150 for $K = 18$.




FIG. 11. (a) Avalanche size distribution in 5×5 ecosystems. Figure plots the logarithm of the number of avalanches, vs. the logarithm of the size of the avalanche (species \times time), for $N=24$, $C=1$, $K=10$ ecosystems, 228 avalanches are plotted. (b) As in (a), except $K=14$, 374 avalanches are shown. (c) As in (a), except $K=20$, 308 avalanches are shown. Note (a)–(c), as K approaches K_{opt} value of $K=10$, avalanche size distribution appears to approach a power law.

We would expect a relation between the value of K and the sizes of avalanches. Intuitively, the "frozen" state is readily attained and "solid" when K is sufficiently high, but the frozen state becomes more tenuous as K decreases. When the frozen state is very solid, avalanches are not likely to propagate far. When the frozen state is nearly melted, any perturbation is likely to propagate further. This is entirely in accord with the fact that the mean avalanche size increases and so does the variance as K decreases in both the 5×5 and 10×10 ecosystems. It also appears, that as K decreases, the distribution of avalanche sizes approaches a power law. Thus, for the 5×5 ecosystems where $K = 8$ – 10 corresponds to the value which optimizes sustained fitness, and for which ecosystems just begin to freeze, the log-log plot is approaching a straight line. Hence, the distribution appears to be a power law. For larger values of K , the log-log plot is convex, hence it is not a power law. The same features recur in the 10×10 ecosystems Fig. 12(a)–(c) shows that the log-log plot for $K = 18$ and $K = 14$ are clearly convex, hence not a power law, and may be approaching a power law linear relationship by $K = 12$.

Similar results, data not shown, arise with respect to the "species" measure of avalanche sizes.

In sum, the following features are of interest.

- (1) There are more small than large avalanches.
- (2) As K decreases, the *mean and median size of avalanches increases*, and the *variance increases*.
- (3) The distribution of avalanche sizes is clearly not a power law when K is sufficiently large, but appear to be approaching a power law on both measures of avalanche size as K decreases towards a "critical" value at which waiting times to encounter Nash equilibria "diverge".
- (4) The critical value of K corresponds to the K_{opt} value which optimizes sustained fitness in the coevolving ecosystem. Thus, selective forces are expected to pull the ecosystem to the poised self-organized critical state where avalanches exhibit a power law distribution.

4.5. AVALANCHES AND THE DISTRIBUTION OF EXTINCTION EVENTS IN THE RECORD

These results may bear upon the distribution of extinction events in the evolutionary record. Raup (1986) has analyzed the intensity of extinction events at the family level during each of the 79 stages of the entire Phanerozoic. On average, each stage is about 7 million years. Figure 13(a) shows Raup's histogram of the number of extinctions per stage (intensity), graphed against the number of stages exhibiting that intensity. Clearly, there are many more small extinction events than large events. Raup makes the point that the distribution is also clearly continuous. Figure 13(b) replots Raup's data in log-log form. Although the data are obviously too weak to place much weight upon, the log-log plot [Fig. 16(b) is clearly convex], suggesting that the observed distribution is not a power law.

In so far as we wish to take these models seriously, Raup's data look like the $K = 14$ curves for both the 10×10 and 5×5 ecosystems, corresponding to an ecosystem with K large enough that the frozen structure is somewhat firm, rather than extremely tenuous or extremely "solid".




FIG. 12. (a) Propagation of coevolutionary avalanches in 10×10 ecosystems. $N = 24$, $C = 1$, $K = 12$. Figure plots the logarithm of the number of avalanches of a given size, vs. the logarithm of avalanche size, (species \times time). (b) As in (a), except K is 14. (c) As in (a), except K is 18.

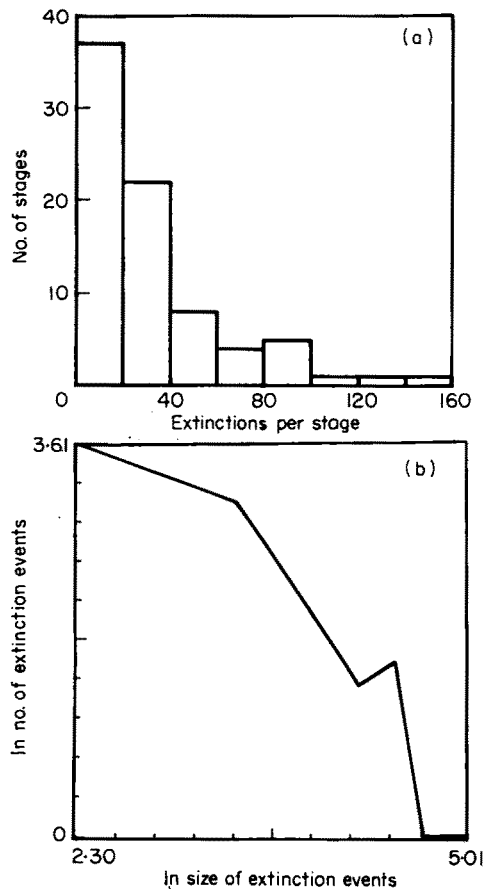


FIG. 13. (a) Raup's data for the sizes of extinction events vs. the number of events at that size. Raup totals the number of extinctions at the family level over each 7 million year interval since the Cambrian. (b) Replot of Raup's data with logarithm of the number of extinction events

What conclusions are warranted by these results? A first general conclusion is the insight that *coevolutionary avalanches propagate* through ecosystems, that such avalanches have characteristic frequency vs. size distributions which *itself changes* depending upon the parameters of the system. In particular the distribution of avalanche sizes depends upon how solid the "frozen" state is. If we tentatively accept Raup's data as weak evidence, the "frozen state" is modestly firm.

A second and critical result is this: *Perturbations of the same initial size can unleash avalanches on a large variety of length scales.* This conclusion is clear and important. In these simulations, the perturbation in each case is a change in the "world" of a single randomly chosen species in the ecosystem. The same "size" perturbation of the external world generates a marked diversity of avalanches, with many small and few large avalanches. If we may tentatively assume that such avalanches can be linked to extinction events, then these results strongly suggest that *uniform* alterations

in the "external world" during evolution can cause a diversity of sizes in extinction events. This possibility stands in contrast to the general hypothesis that small and large extinction events must be associated with small and cataclysmic changes in the external world. Since the external environment has almost certainly undergone changes on a variety of scales, we do not wish to assert that high variability in extinction sizes does not in part reflect a heterogeneity in intensity of causes. But these results place part of the responsibility for the diversity on the dynamics of coupled ecosystems and how "damage" propagates.

5. Inclusion of Density Dependent Population Dynamics

Our analysis above has relied on use of a simplified population dynamics. Each species, at each generation, tries a mutant variant and moves there as a whole if the mutant is better. Each species interacts at each moment with all its immediate neighbors in the ecosystem. The natural extension of this class of models considers a population for each species, subject to population growth characterized by familiar ecological models. Such models are based on "*R*" and "*K*" factors, where *R* reflects the intrinsic growth rate of each population by itself, while *K* reflects "density effects" due to the carrying capacity of the environment which limits each population's increase to a standing abundance, and includes possible competitive, mutualistic and predator-prey interactions among the species (Roughgarden, 1979; May, 1976; Pimm, 1982).

To generalize our *NK* coevolving model, we defined the fitness of each genotype of each species in isolation from all other species to obtain its "*R*" value, and assessed for each species, "*i*", the effect of a pair-wise interaction with a member of each of the other species "*j*" connected to it in the model ecosystem. We define "*R_i*" as the fitness of the species in isolation minus 0.5, the mean fitness in the space of genotypes. This definition allows a species to be an autotroph, *R_i* > 0, or an obligate heterotroph, *R_i* < 0. We define "*a_{i,j}*" as the fitness of species "*i*" when it interacts with species *j*, minus 0.5. This definition allows interactions to be mutualistic, *a_{i,j}* > 0, or competitive, *a_{i,j}* < 0. We assumed that each species had the same carrying capacity, *K*. Then for each species, we utilized the familiar logistic equation:

$$dX_i/dt = X_i[R_i - X_i/K + (\sum_{j=1}^N a_{i,j}X_j)/K].$$

Here in the absence of interactions by other species, each species will attain the stationary population, *R_iK*. Positive or negative interactions by other species, regarded as mutualism and competition, will, in general, alter the population attained by each species. In studies which allowed species to mutate, we first confirmed that a mutant form initially present at low frequency compared to the "wildtype", and which increased in abundance more rapidly than the wildtype, continued to do so until the wildtype population was replaced. Having confirmed this under a number of conditions, we substituted a simpler evolutionary dynamics: if a mutant form

initially increases in abundance more rapidly in the current coevolutionary context than the wildtype, then the whole population "moved" to the mutant form.

Analysis of these population dynamic models appears to confirm the results based on the simpler dynamics discussed above, but allows us to extend our results to look at the actual dynamics of population abundances as well as changes of genotypes. We analyzed 5×5 ecosystems in which each "interior", "edge" or "corner" species interacted with its four, three, or two nearest neighbors, as above, and also 5×5 ecosystems in which each "interior" species interacted with its eight nearest neighbors, adjacent and diagonal, each "edge" species interacted with its five nearest neighbors, and each "corner species interacted with its three nearest neighbors. In addition, we examined the dynamics in 5×5 ecosystems in which each species interacted with all the other species directly. The major results are similar. The major differences, as the numbers of species connected to each species increases, are that ecosystems tend to be more volatile, hence that easy attainment of Nash equilibria requires higher K values. Table 1 summarizes the data for the case in which the interior species interacts with eight species, with respect to mean population attained, average $a_{i,j}$ values, average R_i values, average fraction of couplings from other species which are "mutualistic", and average number of mutations per species accepted over the run. Note that $K = 4$ yields the optimal sustained population, the highest mean $a_{i,j}$ coupling, the highest mean R_i , the highest fraction of mutualistic couplings, and corresponds to a sharp drop in mutations accepted compared to $K = 2$. The results, as the connectivity of interior species increases from 4 to 8 then to 23, data not shown, are that the optimal value of K increases from 2 to 4 and remains at 4 for the completely connected ecosystem.

TABLE 1

Results tabulated for 5×5 ecosystems in which interior species are connected to eight neighbors, edge species are connected to five neighbors, and corner species are connected to three neighbors. K values range from 0-16. Pop is mean sustained population density of a species in the ecosystem. $a_{i,j}$ is mean value of this measure of mutualism among the 25 members of the ecosystem. R_i is the mean self growth rate, autotrophic if positive, of the 25 species in the ecosystem. Fraction mutualistic, the fraction of the couplings into each species which have $a_{i,j}$ values greater than 0, hence it is the fraction of couplings between species which are mutualistic to some degree. Mutations, the number of mutations accepted per species over the simulated coevolutionary period

K	Population	$a_{i,j}$	Fraction mutualistic	R_i	Mutation
0	1593	0.009	0.506	0.158	47.4
2	2103	0.077	0.600	0.205	38.2
4	2264	0.138	0.657	0.210	14.0
8	2064	0.090	0.599	0.196	10.2
16	1276	0.027	0.522	0.158	5.3




FIG. 14. (a)–(c) Onset of frozen components in 5×5 ecosystem, seen by showing the number of mutations which are “accepted” among all 25 species at each ecosystem generation. (a) $K = 0$, (b) $K = 4$, (c) $K = 16$. In these simulations, $N = 24$, $C = 12$.

Figures 14(a)–(c) examines the onset of frozen components in 5×5 species model ecosystems in which each species interacts with eight, five, or three neighbors, by examining the number of mutations which are “accepted” among all 25 species at each ecosystem generation. In these simulations, $C = 12$, and K is 0, 4, and 16 respectively. Note that when K is small relative to C , $K = 0$, the species in the ecosystem continue to mutate persistently throughout the run. Thus, the system never attains a Nash equilibrium. In contrast, when K is much larger than C , $K = 16$, most species stop finding fitter mutants after a few hundred generations, hence the ecosystem is largely frozen and unchanging. Since we here use the “random” rather than “fitter” or “greedy” dynamics, it is not entirely certain that the systems have attained true Nash equilibria.

Figure 15(a)–(c), as well as Table 1, examine the emergence of *mutualism* in these models. Each species grows on its own, hence can be thought of as an autotroph. But in addition, each species, in its interaction with other species, might be helped or harmed by that interaction. The NK model, via the C couplings, permits the possibility that each species can change its genotype such that it is helped by its coevolutionary partners. Table 1 shows that mean $a_{i,j}$ values are greater than 0 in all cases. Thus, mutualism emerges in these systems. Note also in Table 1 that the $a_{i,j}$ values increase then decrease as K increases. Figure 15 shows the details. For $K = 0$ massive fluctuations occur in $a_{i,j}$ values as coevolutionary changes propagate through the unfrozen ecosystem. The fluctuations keep the $a_{i,j}$ values low. Conversely, for the largest value of K , $K = 16$, ecosystems freeze easily, but the fraction of connections which are mutualistic has fallen, and mean $a_{i,j}$ values have decreased from the optimum at $K = 4$. Mutualistic interactions appear harder to attain when K is large and each system harbors more conflicting constraints. Thus, an intermediate value of K , $K = 4$, appears to optimize the ease of forming strong mutualistic interactions.

These figures also give a dramatic view of the periods of quiescence and bursts of change which propagate through these model ecosystems, and for $K = 16$, give clear evidence of the attainment of a Nash or near Nash equilibrium.

Figure 16(a)–(c) shows the actual population dynamics for these conditions. Note that for the highest value of K , $K = 16$, population abundances increase rather smoothly to their carrying capacity. As K decreases, the population behavior becomes more erratic. Genotypic changes show up as discontinuities in the rate of population change. For $K = 0$ it appears that several species actually *decrease* monotonically in abundance during the run. As emphasized in Table 1, for $K = 0$, 2, 4, 8, and 16, the total abundances of all species increases then decreases as K increases, reaching a maximum at $K = 4$. Thus, there is an optimal value of K , as in the simpler dynamics studied above.

5.1. INTERMITTENCY AND PUNCTUATED EQUILIBRIUM

Another general feature of these results, whether on extended or fully connected ecosystems, or including populations or not, is intermittency and bursts of change. The most plausible “move” dynamics which we consider is the “random” dynamics,




FIG. 15. (a)–(c) Evolution of mutualism in 5×5 ecosystem. Figures plot each species average " a_{ij} " coupling to other species over time. Values greater than 0 are mutualistic, those less than 0 are competitive interactions members of the fraction of their couplings which are mutualistic (positive) or competitive (negative). (a)–(c) K values 0, 4, 16.




FIG. 16. (a)–(c) Population dynamics of 5 × 5 ecosystems. (a)–(c) $K = 0, 4, 16$.

where a random mutation is tried, and the species moves if a fitter variant is found. Under this dynamic, even in the absence of exogenous noise, the time history of a coevolving set of species shows periods of quiescence where no species finds a fitter variant, and bursts of change when one species finds a fitter variant and unleashes an avalanche which propagates through the system. Figure 16(a)–(c) demonstrates this abundantly. It is tempting to associate these quiescent periods with stasis, and the bursts with punctuated change.

6. Discussion

The central issues we have investigated in this article are the relations between the *structure* of the fitness landscape of one coevolving partner, and the *couplings* between landscapes such that adaptive moves by one partner deform the landscapes of the adjacent partners, and their effects on the coevolutionary process. Further, we have focused on the possibility that a *selective metadynamics* acting on individual members of individual species, might alter the character of the coevolutionary process such that coevolving partners can coevolve well. With respect to the first set of issues, it is important to stress that our framework is very general. Coevolution is certainly some form of flow by each adapting population on a fitness landscape which deforms due to intraspecies and interspecies couplings. Contemporary discussions of coevolution have developed without yet having a broad theory of the structure of complex fitness landscapes, and therefore without having a body of theory about the relationship between the multi peaked structure of fitness landscapes and their deformability and the consequent character of coevolutionary processes. The *NK* model has been discussed as perhaps the first, but surely not the only, family of rugged multi peaked fitness landscapes. Thus, the *NK* family of landscapes becomes increasingly more rugged as K increases from 0 to N . For $K = 0$, the landscapes have a single Fujiyama peak and smooth sides. For $K = N - 1$, landscapes are fully random. One can conceive of other ways to generate a family of correlated landscapes. For example, one might start with a completely random landscape, and successively “smooth” it by replacing the fitness value at each vertex, or genotype, with the mean of it and its 1-mutant neighbors in the space. Such averaging is an analogue of diffusion. The statistics of such a diffusively smoothed family of landscapes may be similar to or very different from the *NK* family. In general, it seems a critical question whether there are a few basic families of rugged but correlated fitness landscapes, or very many. Whether few or many exist, our aim must be ultimately to characterize the actual statistical structure of fitness landscapes underlying biological evolution at molecular, morphological, and behavioral levels. Thereafter, we need to understand how richly coupled, hence deformable fitness landscapes actually are.

A clear implication of these investigations is that the behavior of a coevolutionary system can be chaotic and disorganized, or highly orderly, depending upon the structures and couplings among the landscapes of the coevolving partners. In the current model, increasing K increases the ruggedness of individual fitness landscapes. Increasing C increases the couplings between landscapes, hence the deforma-

tion to the frog as the fly makes an adaptive move. In completely connected model ecosystems, each species is coupled to all other species. More realistically, each is coupled to only a few other species, S_i . Roughly speaking, if $K > C \times S_i$ for each species, then the ecosystem as a whole reaches a Nash equilibrium rapidly. We referred to this Nash equilibrium, encompassing some or all of the species, as a *frozen component* in the ecosystem. If K is large enough, the ecosystem freezes rapidly and solidly. Conversely, and approximately stated, if $K < C \times S_i$, then the ecosystem does not rapidly reach a Nash equilibrium, and behaves chaotically. Thus, K relative to $C \times S_i$ is a *control parameter*. As this ratio changes, the system passes from a chaotic "gas" like behavior, through a phase where frozen components just form, a kind of "liquid" phase, to a well-frozen "solid" phase.

External perturbations due to changes in the abiotic environment, or internal changes under a random, rather than fitter or greedy evolutionary dynamics, can trigger avalanches of coevolutionary change which propagate through the system. The size distribution of avalanches depends upon how solidly frozen the ecosystem is, and appears to approach a power law distribution in the "liquid" region of the order parameter's value, when frozen components are just arising in the ecosystem. In the solid region, the distribution of the logarithm of the numbers of avalanches against the logarithm of avalanche sizes is convex rather than linear. Since such changes are associated with fluctuations to low fitness it is a real possibility that such avalanches are associated with extinction events. A central conclusion is that small and large extinction events can in principle issue from perturbations of the same small size. Thus, it may prove possible to predict the distribution of sizes of extinction events in the record. The data from the record are convex in a log-log plot, hence tend to support the hypothesis that ecosystems over evolutionary time are somewhat into the solid regime.

Since coevolutionary dynamics depends upon the ruggedness of landscapes and the coupling of landscapes, it is hardly surprising that the sustained average fitness of coevolving partners depends upon these factors as well. This obvious point carries the possibility of a metadynamics in which selection favors individuals whose landscapes and landscape couplings are sculpted to abet coevolutionary success. Regardless of the detailed appropriateness of our NK ecosystem model, therefore, we want to stress that evolution may have itself created the kinds of organisms and couplings among them which favor the capacity of species to actually typically succeed in adapting to one another. In turn, this possibility poses the problem of what the "proper" structure of an ecosystem might be to optimize coevolution. In the context of the NK model, we found evidence that sustained fitness is optimized when the ecosystem is in the liquid regime, "poised at the edge of chaos". Further, we found evidence that selection acting on individual members of individual species might pull each towards the poised state by tuning K relative to $C \times S_i$. Thus, at least the hint of a plausible metadynamics exists, and the hint that the edge of chaos might be an attractor of such an evolutionary metadynamics exists. Nevertheless, the supposition that the poised state where frozen components just form is the natural attractor for a coevolutionary metadynamics must be regarded cautiously. It remains to be seen whether such a state is the one which optimizes sustained

fitness even in the context of the *NK* model, for example if *C* or *S* can evolve as well as *K*, let alone models based on other families of fitness landscapes and their couplings. Thus, we emphasize the distinction between the central idea that a metadynamics might tune the structure of landscapes and their couplings to enhance sustained average fitness, and the more detailed possibility that the edge of chaos is typically the attractor of such a selective metadynamics.

7. Summary

The capacity to coevolve successfully is not trivial, for mere chaotic twitchings of the angry Red Queen may occur. We have introduced an approach to this question based on statistical models of rugged fitness landscapes. Coevolution is thought of, at the lowest level, as a coupling of landscapes, such that adaptive moves by one player deform the landscapes of its immediate partners. In these models we are able to tune the ruggedness of landscapes, how richly intercoupled any two landscapes are, and how many other players interact with each player. We find that all these properties profoundly alter the character of the coevolutionary dynamics.

The results we have reached suggest some tentative conclusions and avenues of investigation. Landscapes need to be of sufficient ruggedness to offset the couplings between landscapes and the number of partners whose moves impinge upon each landscape. Otherwise viewed, epistatic couplings within a species need to be large enough to counterbalance epistatic couplings to the other coevolving partners. We have identified possible selective forces which may tune these parameters such that coevolution is typically successful.

Perhaps most intriguingly, we have found beginning evidence that if each species optimizes its own sustained fitness, such ecosystems might approach a poised self organized critical state, balanced on the edge of chaos. Thus, we must consider a selective "metadynamics" which sculpts the structure of organisms, hence their fitness landscapes, and sculpts their couplings to one another to attain the coevolution of poised, structured ecosystems which harbor nearly melted frozen components and permit propagation of packets of coevolutionary change ringing in new species and ringing out old, with a characteristic distribution between cascade size and frequency. If these views are right, then selection itself achieves ecosystems of a characteristic kind capable of successful coevolution.

This work was partially funded by N.I.H. GM 40186 and O.N.R. N000 14 89 J 1623.

REFERENCES

- BAK, P., TANG, C. & WIESENFELD, K. (1988). Self-organized criticality. *Phys. Rev. A* **38**(1), 364-374.
BINDER, K. & YOUNG, A. P. (1986). Spin glasses: experimental facts, theoretical concepts, and open questions. *Rev. Mod. Phys.* **58**, 801-976.
DERRIDA, B. (1981). Random energy model: an exactly solvable model of disordered systems. *Phys. Rev. B* **24**, 2613.
EIGEN, M., MCCASKILL, J. & SCHUSTER, P. (1988). The molecular quasispecies. *J. Phys. Chem.* **92**, 6881.
FONTANA, W. & SCHUSTER, P. (1987). A computer model of evolutionary optimization. *Biophys. Chem.* **26**, 123.

- GILLESPIE, J. H. (1983). A simple stochastic gene substitution model. *Theor. pop. Biol.* **23**(2), 202–215.
- GILLESPIE, J. H. (1984). Molecular evolution over the mutational landscape. *Evolution* **38**(5), 1116–1129.
- KAUFFMAN, S. A. (1989a). Adaptation on rugged fitness landscapes. In: *Lectures in the Sciences of Complexity. The Santa Fe Institute Series* (Stein, D., ed.) pp. 527–618. New York: Addison Wesley.
- KAUFFMAN, S. A. (1989b). Principles of adaptation in complex systems. In: *Lectures in the Sciences of Complexity. The Santa Fe Institute Series* (Stein, D., ed.) pp. 619–712. New York: Addison Wesley.
- KAUFFMAN, S. A. (1990). *The Origins of Order: Self Organization and Selection in Evolution*. Oxford: Oxford University. In press.
- KAUFFMAN, S. A. & LEVIN, S. (1987). Towards a general theory of adaptive walks on rugged landscapes. *J. theor. Biol.* **128**, 11–45.
- KAUFFMAN, S. A. & WEINBERGER, E. W. (1989). The NK model of rugged fitness landscapes and its application to maturation of the immune response. *J. theor. Biol.* **141**, 211–245.
- KAUFFMAN, S. A., WEINBERGER, E. W. & PERELSON, A. S. (1989). In: *Theoretical Immunology Part 1. Santa Fe Institute in the Science of Complexity* (Perelson, A. S., ed.) pp. 349–382. New York: Addison Wesley.
- LANGTON, C. (1990). Adaptation to the Edge of Chaos. Presented at the Second Artificial Life Conference, Santa Fe. N.M. and forthcoming in *Artificial Life, II*, the Santa Fe Institute Series in the Sciences of Complexity, Addison Wesley, (1990).
- MACKEN, C. A. & PERELSON, A. S. (1989). Protein evolution on rugged landscapes. *Proc. natn. Acad. Sci. U.S.A.* **86**, 6191–6195.
- MAY, R. M. (ed.) (1976). *Theoretical Ecology: Principles and Applications*. Philadelphia: W. B. Saunders.
- MAYNARD SMITH, J. (1970). Natural selection and the notion of a protein space. *Nature, Lond.* **225**, 563–566.
- MAYNARD SMITH, J. (1982). *Evolution and the Theory of Games* Cambridge: Cambridge University Press.
- MAYNARD SMITH, J. & PRICE, G. R. (1973). The logic of animal conflict *Nature, Lond.* **246**, 15–16.
- PIMM, S. L. (1982). *Food Webs*. London: Chapman and Hall.
- RAUP, D. M. (1986). Biological extinction in earth history. *Science* **231**, 1528–1533.
- ROSENZWEIG, M. L., BROWN, J. S. & VINCENT, T. L. (1987). Red queens and ESS: the coevolution of evolutionary rates. *Evol. Ecol.* **1**, 59–94.
- ROUGHGARDEN, J. (1979). *Theory of Population Genetics and Evolutionary Ecology: An Introduction*. New York: Macmillan.
- SCHUSTER, P. (1987). Structure and dynamics of replication–mutation systems. *Physica Scripta* **35**, 402–416.
- WEINBERGER, E. W. (1990). Correlated and uncorrelated fitness landscapes and how to tell the difference. *Biol. Cybernet.* in press.
- WRIGHT, S. (1932). The role of mutation, inbreeding, crossbreeding and selection in evolution. *Proceedings 6th International Congress on Genetics* **1**, 356–366.