Kristian Lindgren
Institute of Physical Resource Theory, Chalmers University of Technology and Göteborg University, S-412 96 Göteborg, Sweden; E-mail: frtkl@fy.chalmers.se

Evolutionary Dynamics in Game-Theoretic Models

A number of evolutionary models based on the iterated Prisoner's Dilemma with noise are discussed. Different aspects of the evolutionary behavior are illustrated: (i) by varying the trickiness of the game (iterated game, mistakes, misunderstandings, choice of payoff matrix), (ii) by introducing spatial dimensions, and (iii) by modifying the strategy space and the representation of strategies. One of the variations involves the finitely iterated game that has a unique Nash equilibrium of only defecting strategies, and it is illustrated that when a spatial dimension is added, evolution usually avoids this state. The finite automaton representation of strategies is also revisited, and one model shows an evolution of a very robust error-correcting strategy for the Prisoner's Dilemma game.

The Economy as an Evolving Complex System II, Eds. Arthur, Durlauf, and Lane
SFI Studies in the Sciences of Complexity, Vol. XXVII, Addison-Wesley, 1997

INTRODUCTION

The term "evolutionary dynamics" often refers to systems that exhibit a time evolution in which the character of the dynamics may change due to internal mechanisms. Such models are, of course, interesting for studying systems in which variation and selection are important components. A growing interest in these models does not only come from evolutionary biology, but also from other scientific disciplines. For example, in economic theory, technological evolution as well as evolution of behavior in markets may call for models with evolutionary mechanisms.

In this chapter, we focus on dynamical systems described by equations of motion that may change in time according to certain rules, which can be interpreted as mutation operations. In models of this type, the variables typically correspond to the number or the mass of certain components (individuals/organisms), or in more mechanistic (or spatial) models the position of certain components.

In the models that we discuss here, the equations of motion for the different variables (individuals) are usually coupled, which means that we have *coevolutionary* systems. The success or failure for a certain type of individual (species) depends on which other individuals are present. In this case, there is not a fixed fitness landscape in which the evolutionary dynamics climbs toward increasing elevation, but a position that at one time is a peak may turn into a valley. This ever changing character of the world determining the evolutionary path allows for complex dynamic phenomena.

Coevolutionary dynamics differ, in this sense, from the common use of the genetic algorithm,[18] in which a fixed goal is used in the fitness function and where there is no interaction between individuals. In the genetic algorithm, the focus is on the final result—what is the best or a good solution? In models of coevolutionary systems, one is usually interested in the transient phenomenon of evolution, which in the case of open-ended evolution never reaches an attractor. There are variations, though, where interactions between individuals have been suggested as a way to improve on the genetic algorithm, see, e.g., Hillis.[16] For a recent review article on genetic algorithms in an evolutionary context, see Mitchell and Forrest.[33]

We make a distinction between evolutionary systems and adaptive systems. The equations of motion in an evolutionary system reflect the basic mechanisms of biological evolution, i.e., inheritance, mutation, and selection. In an adaptive system, other mechanisms are allowed as well, e.g., modifications of strategies based on individual forecasts on the future state of the system. But, increasing the possibilities for individualistic rational behavior does not necessarily improve the outcome for the species to which the individual belongs in the long run. An example of this difference is illustrated and discussed for one of the evolutionary lattice models.

In the next section, we briefly discuss some evolutionary models in the literature. The main part of this chapter is devoted to a discussion of a number of evolutionary models based on the iterated Prisoner's Dilemma game as the interaction between individuals, see, e.g., Axelrod.[3,4] Different aspects will be illustrated:

(i) by varying the trickiness of the game (iterated game, mistakes, misunderstandings, choice of payoff matrix), (ii) by introducing spatial dimensions, and (iii) by modifying the strategy space and the representation of strategies.

The results in this chapter concerning the finite memory strategies have, to a large extent, been reported before; for the mean-field model, see Lindgren,[26] and for the lattice model, see Lindgren and Nordahl;[27] see also the review article by Lindgren and Nordahl.[28] The model based on the finitely iterated game as well as the models using finite automata strategies are new. A detailed analysis of these models will be reported elsewhere.

MODELS OF EVOLUTION

Evolutionary models can be characterized both by the level at which the mechanisms are working and the dimensionality of the system. Usually these characteristics are coupled to some extent, so that low-dimensional models usually reflect mechanisms on a higher system level. An example of that regime is an evolutionary model in which the variables are positions of phenotypic characters in phenotype space, and the dynamics is determined by an ordinary differential equation that uses the distance to fitness maximum.[48] At the other end of the scale, there are models based on "microscopic" components that interact and organize, e.g., in the form of catalytic networks. The work by Eigen,[11] Shuster,[45] Fontana,[15] and many others[6,13,14,21] lies in this region. Some of these models have similarities to spin-glass models in statistical mechanics, see, e.g., the models for evolution of RNA.[1,43,47] Other examples can be found in, e.g., Weisbuch,[51] and one of the most well-known models is the *NK*-model by Kauffman; for a review, see Kauffman.[22] Such spin-glass-like models could also be interpreted as models of interacting individuals or species.

A special class of models is based on instructions floating in a virtual computer memory. Such instructions may eventually organize in a structure that can be interpreted as an organism. Then, the interactions between organisms will not be fixed by the model, but will be a result of the evolution. This is the approach taken by, for example, Rasmussen and coworkers[41] and Ray.[42] Models of this type are obviously high dimensional.

There are many examples of other models, some of them mentioned in Figure 1. The models that will be described in more detail in this paper have been constructed based on mechanisms at the level of individual organisms and their interactions. There is a large number of models of this kind, many of them using various game-theoretic approaches for the interaction between individuals.[4,26,27,30,32,38,46]

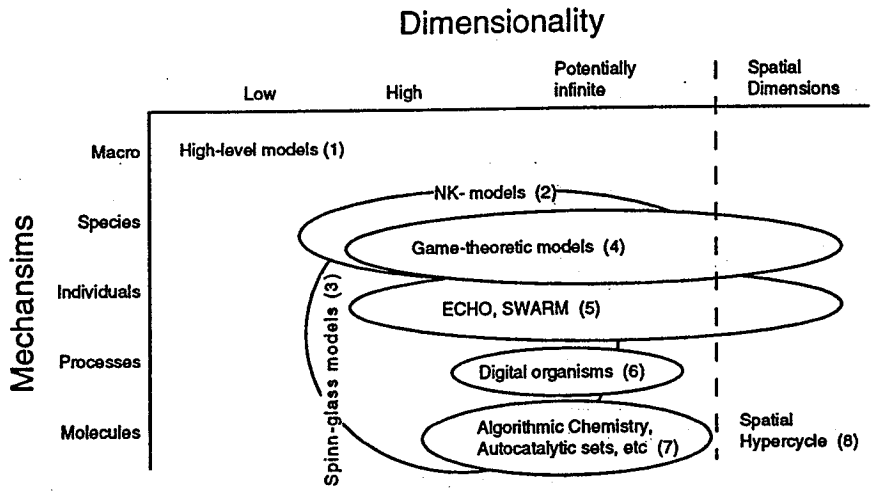

FIGURE 1 Evolutionary models characterized by level and dimensionality. The numbers in the figure refer to various models: (1) Stenseth and Maynard-Smith[48]; (2) Kauffman[22]; (3) Anderson[1] and Stein and Anderson[47]; (4) Matsuo,[30] Axelrod,[4] Miller,[32] Lindgren,[26] Nowak and Sigmund,[38] Stanley et al.,[46] and Lindgren and Nordahl[27]; (5) Holland[18,19] and Langton[24]; (6) Rasmussen et al.[41] and Ray[42]; (7) Eigen,[11] Schuster,[45] Fontana,[15] Kauffman,[21] Farmer et al.,[14] Eigen et al.,[13] and Bagley et al.[6]; (8) Boerlijst and Hogeweg[7,8]; see also Eigen and Schuster.[12]

As a basis for the models presented in the following sections, we use the three necessary mechanisms usually associated with biological evolution:

- inheritance of genetic information
- variation (mutations)
- the feedback effect of natural (or artificial) selection

Furthermore, to get a potential for evolution, the population needs to be provided with a difficult "task" to solve, i.e., complex interactions with environment (including other organisms).

GAME-THEORETIC MODELS AND DYNAMICS

In the choice of interaction model, one has to find a problem that cannot be "solved" too easily in the chosen genetic representation. A solution here means that there is a strategy that is evolutionarily stable (or a set of strategies that together form a

fixed point in the evolutionary dynamics). An evolutionarily stable strategy (ESS) cannot be invaded by any other strategy that initially is present at an arbitrarily small amount.[31]

The Prisoner's Dilemma game, described below, is a suitable model for the interaction between individuals in an evolving population. There are many variations of the game that offer nontrivial problems for the individuals, and still this interaction can be solved analytically given the genetic code for the strategies.

In the Prisoner's Dilemma (PD) game, two players simultaneously choose to cooperate (C) or to defect (D), without knowing the opponent's choice. If they both cooperate they share the highest total payoff and get R points each, but there is a temptation to defect because defection scores $T > R$ against a cooperating opponent, who in that case gets the lowest score S. If both defect they share the lowest total payoff and get P points each. If $T > R > P > S$ and $2R > T + S$, there is a dilemma, since, in an isolated game, rational players choose to defect and then share the lowest total payoff (see Table 1). Cooperative behavior is possible in the iterated game when the same players meet in a series of rounds. In that case, it has been shown that tit-for-tat (TFT), a strategy that mimics the opponent's last action (and starts with C), is capable of establishing cooperation.[3]

Many dynamical systems and evolutionary models have been constructed with the PD game as a model for the interaction between individuals. Axelrod[4] applied a genetic algorithm (GA) to the iterated Prisoner's Dilemma and used a bit-string representation of finite memory strategies. In another GA study, Miller[32] let the bit strings represent finite automata playing the iterated PD game. Miller also introduced noise in the game, but did not find the evolution of the error-correcting strategies like we have seen in the models described below. Stanley et al.[46] also used finite automata to study strategies that could choose their opponents. For recent work along these lines, see Tesfatsion.[50] A third type of strategy representation, suggested by Ikegami,[20] uses a tree structure with varying depth. Strategies could also choose actions with certain probabilities. This has been analyzed by Molander,[34] and population dynamics as well as evolutionary dynamics has been studied for these strategies.[37,38]

TABLE 1 The payoff matrix of the Prisoner's Dilemma game, showing the scores given to player 1 and 2, respectively.

	player 2 cooperates	player 2 defects
player 1 cooperates	(R , R)	(S , T)
player 1 defects	(T , S)	(P , P)

In all models described in this chapter, we have used variations of the Prisoner's Dilemma that can be solved analytically. This reduces the computational power required and allows for long simulations (in terms of generations) of the models.

DYNAMICS BASED ON THE ITERATED PRISONER'S DILEMMA WITH MISTAKES

The iterated game is easily modified to a more complicated one by introducing noise. In this way the game is made more "tricky," and very simple cooperative strategies, like TFT, are not very successful. Noise could either disturb the actions so that the performed action is opposite to the intended one (trembling hand), or the information telling a player about the opponent's action is disturbed (noisy channels).[5] In any case, the complication is sufficient to give rise to interesting evolutionary behavior.

In the first interaction model, there is no noise, but a dilemma is introduced by iterating the game a fixed (known) number of rounds. In that case, the final round is an ordinary Prisoner's Dilemma, and then the next to last round will get a similar status, and so on.

In the second and third models, noise is introduced in the form of mistakes, i.e., the action performed differs from the intended one (given by the strategy), with a certain probability. On the other hand, the problem with the dilemma of the final round is removed, and the game is made infinitely iterated. In the second model, the individuals are equipped with finite memory strategies, while in the third model the more powerful representation of finite automata is used. In all cases the average score per round is calculated, for the finitely iterated game and for the infinitely iterated game, respectively, and used in the dynamics to determine the next generation of the population.

We shall demonstrate various evolutionary phenomena, for all three models by studying the behavior in two completely different "worlds," the mean-field model (all interact with all, i.e., no spatial dimensions), and the cellular automaton (CA) model (local interaction on a lattice with synchronous updating).

In the mean-field model, the dynamics is based on the simple equations,

$$x_k' = x_k + dx_k \left(\sum_i s_{ki} x_i - \sum_{ij} s_{ij} x_i x_j \right), \qquad k = 1, \ldots, N, \tag{1}$$

where x_k is the fraction of the population occupied by strategy k, and s_{ij} is the score for strategy i against j. The prime on the left-hand side denotes the next time step, i.e., we are assuming discrete time. Since the double sum ($\Sigma s_{ij} x_i x_j$) equals the average score in the population, the total population size is conserved. The number of different strategies N may change due to mutations, which are randomly generated after each step of the population dynamics equations (1). We have used continuous variables x_k, with a cut-off value corresponding to a single individual,

below which the strategy is considered extinct and removed from the population. The equations reflect a change of abundancy for a species that is proportional to its own score minus the average score in the population. Each time step is viewed as one generation and, in addition to the equations, mutations are added. This means that generations are nonoverlapping and that reproduction is asexual.

In the CA model, it is assumed that each individual interacts with its four nearest neighbors on a square lattice. The average score for each individual (or lattice site) is calculated, and this determines which individuals will be allowed to reproduce. In each neighborhood (consisting of a cell with its four nearest neighbors), the individual with the highest score reproduces in the middle cell, see Figure 2. (Ties are broken by adding a small random number to each player's score.) The offspring inherits the parent's strategy, possibly altered by mutations to be described below in detail for the different models.

Depending on the choice of payoff matrix, the evolutionary paths may look very different. In the CA model, since the dynamics is determined only by score differences, there are only two independent parameters in the payoff matrix.[27] Therefore, we have chosen to study the parameter region given by $R = 1$, $S = 0$, $1 < T < 2$, and $0 < P < 1$. In the mean-field model, the *value* of the score difference is also important, since the dynamics depends on how fast a score difference changes the composition of the population. Therefore, as a third parameter, we can use the growth constant d in Eqs. (1), while we keep R and S fixed and vary T and P as in the CA model. In the mean-field simulations we have assumed $d = 0.1$. (This results in a growth that is slow enough to view Eqs. (1) as rough approximations to a set of ordinary differential equations, even though this is not the intention.) In all models where noise is present, we have used an error probability of $e = 0.01$.

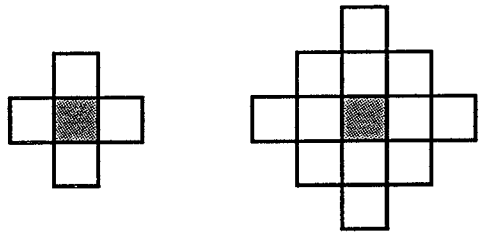

FIGURE 2 (a) In the cellular automaton (CA) model, each individual plays the iterated Prisoner's Dilemma with each of its four nearest neighbors. Simultaneously, for each neighborhood, the player with the highest average score per round reproduces in the "gray" cell. (b) Since the score of the nearest neighbors depends on the strategies present in the next nearest cells, we get a CA with interaction radius 2.

FINITE NUMBER OF ROUNDS

If the iterated game is finite and the number of rounds is known, the dilemma of the single-round PD game remains. Assume that we have two players cooperating, for example, by using the tit-for-tat strategy. When the two players enter the final round of their iterated game, they face the single-round dilemma and profit-maximizing behavior calls for the defect action. But, if we know that our opponent is using a rational strategy, the next to last round will be a dilemma of the same type, and we should defect in that round too. Then the third round from the end will be under consideration, and so on. This implies that the only Nash equilibrium is given by both players defecting all rounds. This does not mean, though, that defecting all the time is the best strategy. Which strategy is the best depends on which other strategies are present in the population. If the population is dominated by, for example, strategies that play TFT for all rounds except the last two when they defect, it is, of course, a bad strategy to start with defection from the beginning. Instead, if one would know the opponent's strategy, one should start with TFT, i.e., cooperate, but switch to always defect (ALLD) the round before the opponent is to defect.

The situation faced by a player entering a population playing the finitely iterated game is an ill-defined one, in the sense discussed by Arthur,[2] where one cannot say which strategy is the best. Like in most real situations, the success or failure of a certain behavior depends on what other individuals or strategies are present.

There are various ways to construct an evolutionary model of the PD game iterated n rounds. Here, we pick a simple example and look at a strategy set containing strategies playing TFT for k rounds and playing ALLD for the remaining $n-k$ rounds. We denote these strategies F_k, with $k = 0, \ldots, n$. The score s_{jk}, for a strategy F_j against F_k, is then given by

$$s_{jk} = \begin{cases} jR + T + (n-j-1)P, & \text{if } j < k\ ; \\ jR + (n-j)P, & \text{if } j = k\ ; \\ kR + S + (n-k-1)P, & \text{if } j > k\ . \end{cases} \tag{2}$$

The population dynamics for a fully connected population (mean-field model) can then be described by Eqs. (1), with mutations at a rate of 10^{-5}, so that a strategy F_i is replaced by a randomly chosen F_j. In Figure 3, the evolutionary path, starting with the strategy F_n, or TFT, shows that cooperative behavior is exploited and that the evolution ends in the ALLD state.[1]

[1]This result also follows from a theorem,[44] saying that an *iteratively strictly dominated strategy* is eliminated in any *aggregate monotonic* selection dynamics. A recent generalization of this result shows that this also holds for a larger class of selection dynamics termed convex-monotone dynamics.[17] In this example, F_n is strictly dominated by F_{n-1}, and F_k is iteratively strictly dominated by F_{k-1}, for $0 < k < n$.

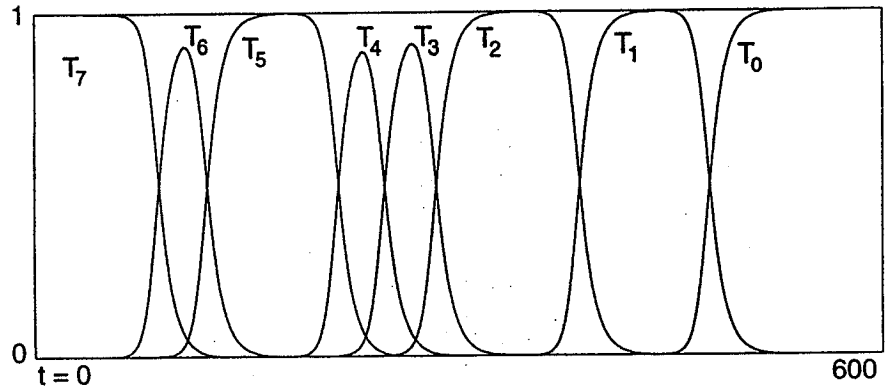

FIGURE 3 The interaction is given by the PD game iterated $n = 7$ rounds. Initially, the population (1000 individuals) consists only of the strategy F_7, or TFT, but due to the dilemma of the fixed number of rounds, the evolution leads to strategies defecting earlier in the iterated game. The final state is then a fixed point, in which all play F_0, or ALLD.

Cooperation cannot be maintained since there will always appear exploiting strategies that grows at the expense of the more cooperative ones. It is well known that the introduction of spatial dimensions may increase the possibility for cooperative behavior.[7,8,27,37,35]

In the CA model, the individuals are put on a 128 × 128 square lattice with periodic boundary conditions. Contrary to what we saw in the mean-field case, cooperation is maintained for a large part of parameter space defining the PD game. This is made possible either by the presence of spatiotemporal patterns, like waves of strategies sweeping over the system, see Figure 4b and 4c, or by formation of stable islands of more cooperative strategies in a background of slightly exploiting strategies, see Figure 4a and 4d.

FINITE MEMORY STRATEGIES

In the infinitely iterated Prisoner's Dilemma with mistakes, a simple type of strategy to consider is the deterministic finite memory strategy. This type of strategy may take into account the actions that have occurred in the game a finite number of rounds backward, and deterministically choose a certain action.

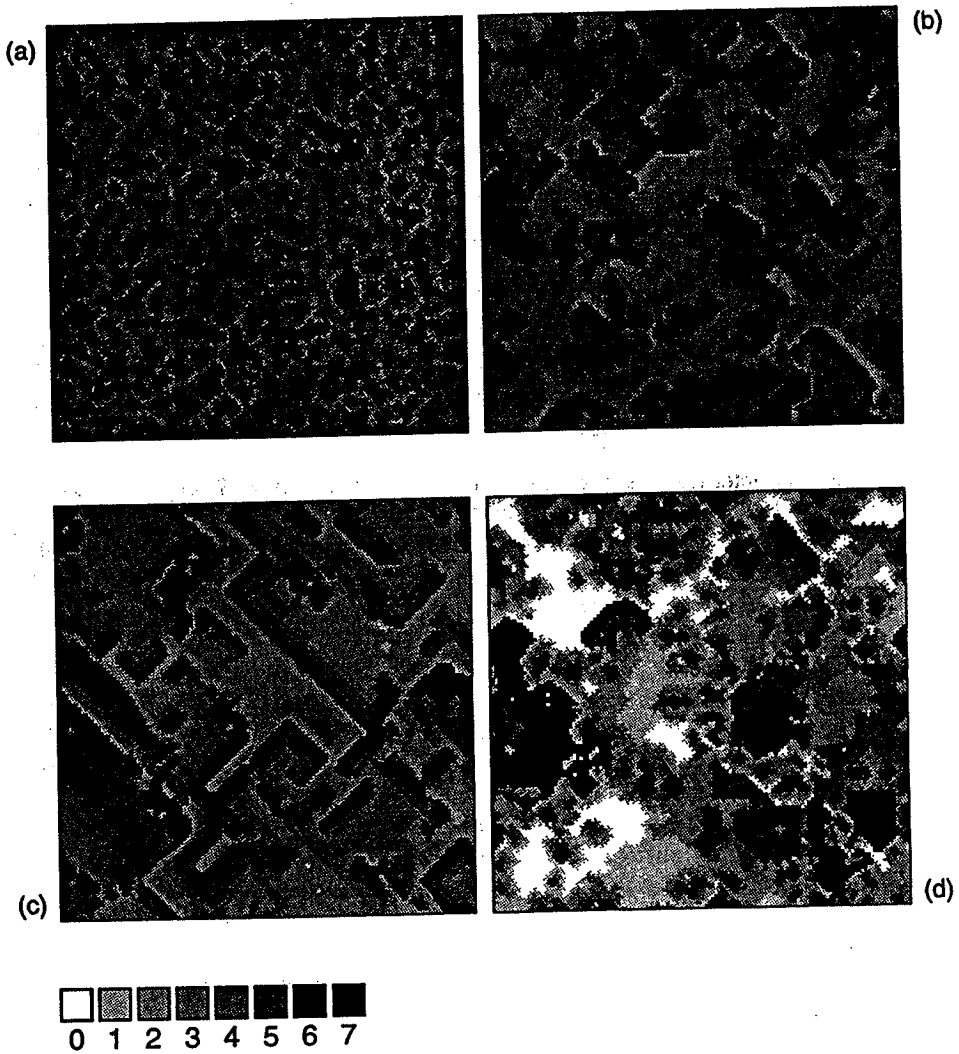

FIGURE 4 The evolution of F_k strategies playing the 7-round iterated Prisoner's Dilemma on a 128 × 128 lattice is shown for four different parameter values ($R = 1$, $S = 0$, T and P are varied). In all cases, we see that the Nash equilibrium of a homogenous ALLD population is avoided. In (a), $T = 1.3$, $P = 0.3$, the state is frozen and dominated by F_7 (TFT), F_6, and F_5. In (b), $T = 1.667$, $P = 0.333$, and in (c), $T = 1.3$, $P = 0.3$, and there are waves running across the system constantly changing the pattern. In (d), $T = 1.7$, $P = 0.7$, and there are frozen patterns that occasionally are modified by waves of TFT. Here, all strategies from F_0 (ALLD) to F_7 (TFT) are present, and they are represented in gray scale from white for ALLD to black for TFT.

A strategy can then be viewed as a look-up table, in which each entry corresponds to each of the possible finite-length histories that can be memorized by the strategy, see Figure 5. The symbols in the output column of the table denote the action to be performed, but there is a small chance (due to noise) that the action is altered.

The game between two finite memory strategies can be described as a stationary stochastic process, and the average score can, therefore, be calculated from a simple system of linear equations. The number of equations necessary equals 2^b, where b is the size of the history needed to determine the next pair of actions. For example, for two memory-1 strategies (as well as memory 2), both actions in the previous round are required and, therefore, the number of equations is 4.

We use three types of mutations. The *point mutation* alters the symbol in the genome. The *gene duplication* attaches a copy of the genome to itself. In this way the memory capacity increases, but it should be noted that this mutation is neutral—the strategy is not changed, but a point mutation is required to make use of the increased capacity. The *split mutation* randomly removes the first or second half of the genome.

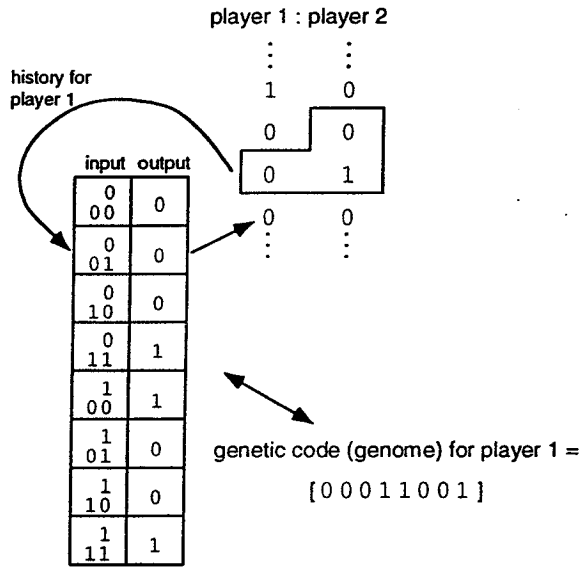

FIGURE 5 The genetic coding of finite memory strategies uses the look-up table. The length of the genome (size of the table) then carries the information on the memory capacity of the strategy, since the table doubles in size when memory increases by 1.

In the simulations, the initial population consists of equal parts of the four possible memory-1 strategies: [00] is ALLD (always defects), [01] is TFT, [10] is ATFT (i.e., anti-tit-for-tat), and [11] is ALLC (always cooperates). Due to mutations, more complex strategies may emerge, and this offers the possibility for evolution of strategies that correct for accidental mistakes.

MEAN-FIELD MODEL

In the mean-field model all interact with all, and the dynamics is given by Eqs (1). The number of individuals is 1000, the point mutation rate is $2 \cdot 10^{-5}$ and the other mutation rates are 10^{-5}.

In Figures 6 and 7, the evolutionary path is shown for two different choices of payoff matrices. The first one, $T = 1.6$ and $P = 0.3$, shows that the system finds a very stable state dominated by the strategy 1001. In the second simulation, at $T = 1.7$ and $P = 0.4$, this strategy does not succeed, but the evolution continues to strategies with longer memory. It turns out that the final state in Figure 7, consists of a large group of strategies having the same mechanism for correcting for mistakes in the iterated PD game. In both cases, the final population is cooperative and cannot be exploited by defecting strategies. For a more detailed discussion on simulations of the finite memory mean-field model, see Lindgren.[26]

Both the strategy 1001 and the final strategies in Figure 7 belong to a class of strategies that can deal with mistakes. These strategies are based on retaliation and synchronization. An accidental defection is followed by a certain number of

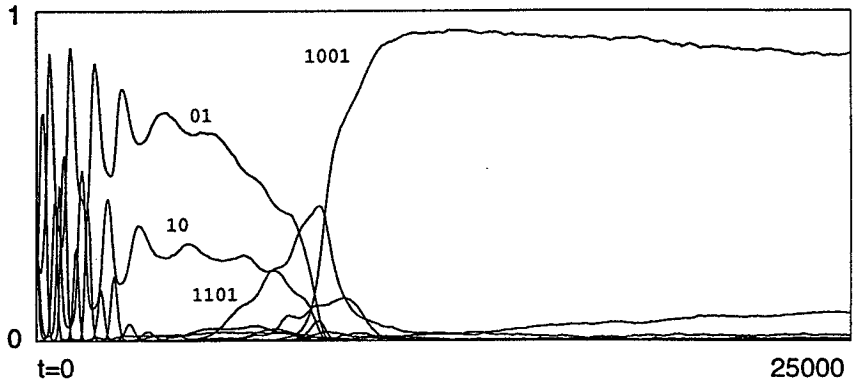

FIGURE 6 The evolution of finite memory strategies in the mean-field model is shown for 25,000 generations, with the parameter values $T = 1.6$, $P = 0.3$. There is a metastable coexistence between 01 (TFT) and 10 (ATFT) before strategies of memory 2 take over. The dominating 1001 strategy is actually evolutionarily stable for these parameter values.

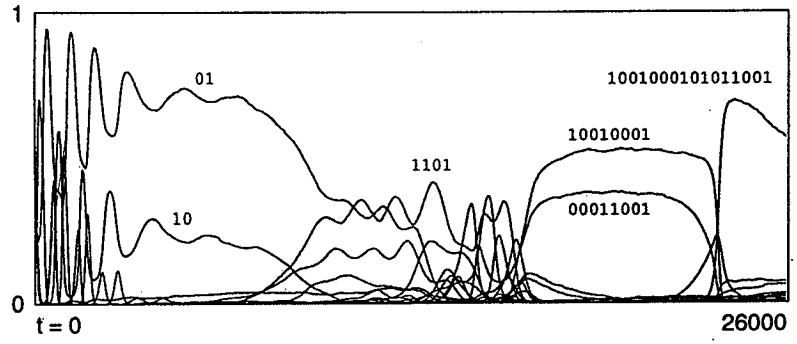

FIGURE 7 The evolution of finite memory strategies in the mean-field model is shown for 26,000 generations, with the parameter values $T = 1.7$, $P = 0.4$. Again, the coexistence between 01 (TFT) and 10 (ATFT) is present, but 1001 does not appear. For these parameter values, longer memories seem to be required to get cooperative and stable (or metastable) strategies. The simulation ends with a strategy that has an unexploitable error-correcting mechanism. The coexistence between the two memory-3 strategies 10010001 and 00011001 is an example of mutualism.[26]

rounds of mutual defection. The rounds of defection both punish possible exploiting strategies and synchronize the strategies before cooperative behavior is established again. The strategies can be described by the simple mechanism:

> After a defection by either of the players, defect until there have been n consecutive rounds of mutual defection, and then start cooperating again.

We shall denote such a strategy by S_n. The first, S_1, often referred to as Simpleton[40] (or Pavlov[38]), is a strategy that cooperates if both players chose the same action in the previous round, and this one is identical to 1001 in our notation. This strategy exhibits the following pattern in case of an accidental defection $\underline{\mathrm{D}}$: $(\ldots, \mathrm{CC}, \mathrm{CC}, \underline{\mathrm{D}}\mathrm{C}, \mathrm{DD}, \mathrm{CC}, \mathrm{CC}, \ldots)$. For the standard parameters used by Axelrod,[3] i.e., $R = 3$, $S = 0$, $T = 5$, and $P = 1$, a population of Simpletons can be invaded, since ALLD (always defects) can exploit it—a single round of mutual defections is not enough to punish an exploiting strategy. The second strategy in this class is S_2, which answers a defection (by either of the players) by defecting twice in a row, and after that returns to cooperation. This behavior leads to the pattern $(\ldots, \mathrm{CC}, \mathrm{CC}, \underline{\mathrm{D}}\mathrm{C}, \mathrm{DD}, \mathrm{DD}, \mathrm{CC}, \mathrm{CC}, \ldots)$, which is the mechanism we find in the population of strategies that dominate at the end of the simulation in Figure 7.

A population of S_n cannot be invaded by a strategy that tries to get back to cooperation faster than does S_n—the invader will only be exploited if that is tried. The only chance for a single invader (or mutant) to succeed is by exploiting the cooperative behavior. Thus, an invader will fail if the number n of defections after

its exploiting defection is large enough to compensate for the temptation score T, i.e., $T + nP < (n+1)R$, if terms of order e are neglected. This means that, in the limit of $e \to 0$, n must be chosen so that

$$n > \frac{T-R}{R-P}. \tag{3}$$

In other words, since $T > R > P$, there is always an error probability e and an integer k, for which the strategy S_k cannot be invaded. Note that $e > 0$ is needed, since it has been shown[9,10,29] that the presence of mistakes are necessary to allow for evolutionary stability. With our choice of $R = 1$, the stability requirement for the strategy S_k, in the limit $e \to 0$, can be written

$$P < 1 + \frac{1}{n} - \frac{T}{n}. \tag{4}$$

These lines are drawn in Figure 8. The two examples discussed above (Figures 6 and 7) are on separate sides of the first line. Parameter choices toward the bottom left hand corner appear to make the situation easier for cooperative strategies, since the temptation to defect is relatively small. In the upper right-hand corner, however, the score difference between cooperating and defecting is smaller, and a strategy that punishes defectors may need a longer memory to be able to avoid exploitation.

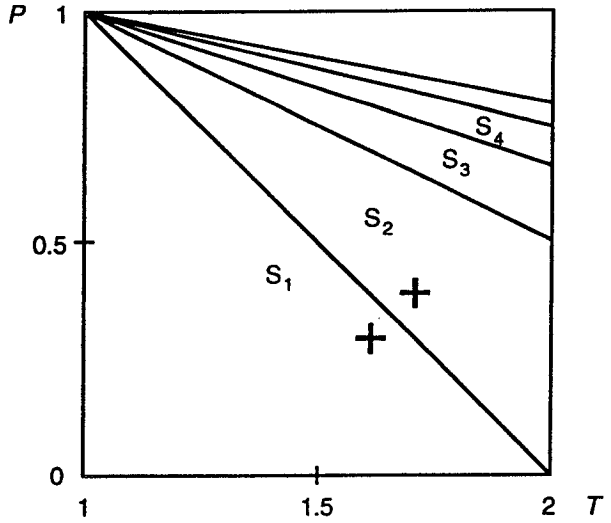

FIGURE 8 The parameter region, $1 < T < 2$ and $0 < P < 1$, is divided into subregions in which the smallest evolutionarily stable strategy of the S_n type is shown, according to the inequality requirement Eq. (4). The two parameter values corresponding to the simulations of Figures 6 and 7 are shown as crosses.

THE CA MODEL

The introduction of spatial dimensions, so that individuals only interact with those in their neighborhood, may affect the dynamics of the system in various ways. It has been demonstrated, for example, in the study by Boerlijst and Hogeweg,[7,8] that parasitic species may be rejected by spiral waves consisting of several species. Other forms of cooperative behavior are also possible, for example regions of cooperative strategies surrounded by exploiting cheaters. The possibility of spatiotemporal structures may allow for global stability where the mean-free model would be unstable. The presence of these various forms of spatiotemoral phenomena may, therefore, also alter the evolutionary path compared with the mean-field case and we may see other strategies evolve.

In Lindgren and Nordahl,[27] we performed a detailed analysis of the dynamic behavior of memory-1 strategies for the parameter region $1 < T < 2$ and $0 < P < 1$. Already at memory 1, there are several spatiotemporal phenomena that may lead to global stability while the system is locally unstable. Spatial coexistence in frozen patterns is also possible. In Figure 9, four examples of spatial patterns from simulations of memory-1 strategies for different parameter choices are shown.

In Figure 10, the strategy abundancy is shown for the first 800 generations of the CA model, with $T = 1.6$ and $P = 0.3$. (In the examples presented here, we have used a 64×64 lattice, and the mutation rates are $p_{\text{point}} = 2 \cdot 10^{-3}$ and $p_{\text{dupl}} = p_{\text{split}} = 10^{-3}$.) It is clear that the strategy 1001 (or S_1) is successful, even though we are close to the line above which it can be exploited, cf. Figure 8. Before 1001 takes over, the system is dominated by 00 and 11, i.e., ALLD and ALLC. This is contrary to what we found in the mean-field case, where 01 (TFT) and 10 (ATFT, does the opposite to TFT) dominate before memory-2 strategies appear. The reason is the possibility to form cooperating regions in the spatial model. Here we get islands of ALLC surrounded by ALLD, similar to the pattern in Figure 9(d). The ALLC individuals at the edges of the islands are certainly exploited by the ALLD individuals and, in the reproduction step, they are removed and replaced by the offspring of the most successful individual in their neighborhood. Now, an ALLC individual that is not in contact with ALLD has the highest score and, therefore, the edge individuals are replaced by an offspring of their own kind. This can be viewed as "kin selection," in which the individual at the edge sacrifices herself since that is beneficial for her own strategy type.

If we would choose a more "rational" procedure for selecting the new strategy, e.g., by taking the strategy in the neighborhood that maximizes the score in the particular cell in the next time step (generation), ALLD would invade the island and wipe out the ALLC strategy from the population. It would be interesting to investigate the difference between "rational" and "natural" selection in this type of model, and to see how the evolutionary dynamics changes.

In Figure 11, the evolutionary path is shown for $T = 1.7$ and $P = 0.4$. Here, 1001 (S_1) should not be stable, according to Eq. (4), but instead, strategies of type

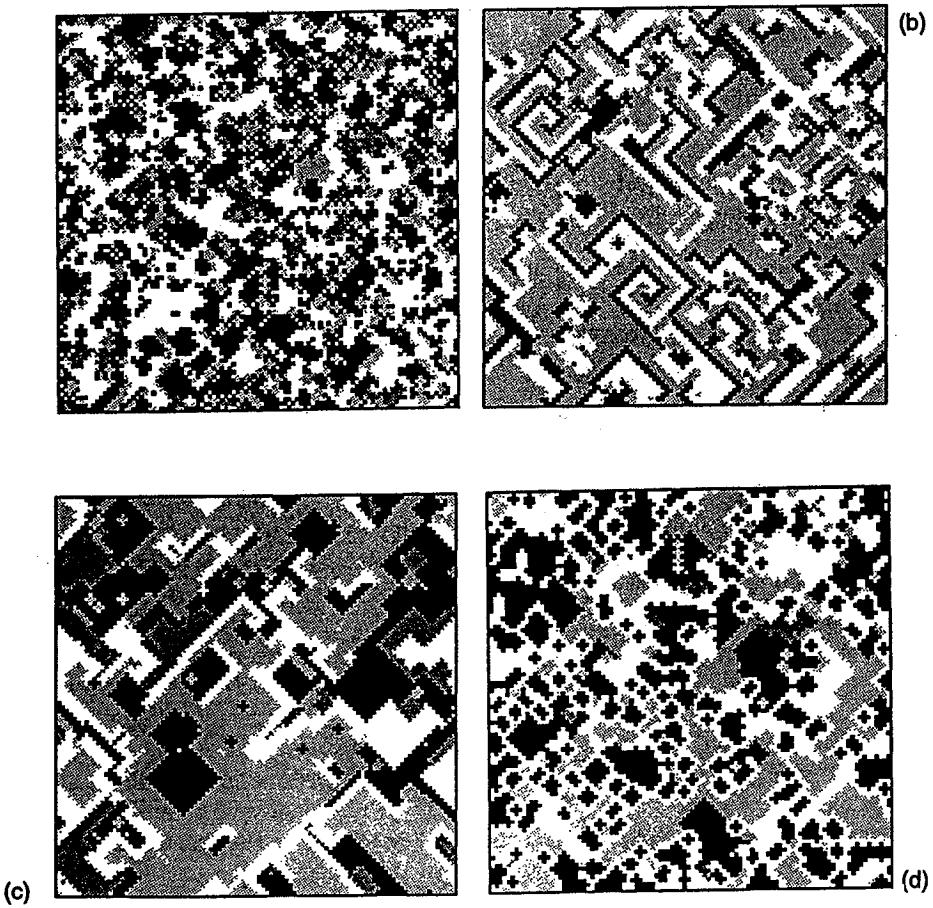

FIGURE 9 Examples of four different types of patterns from simulations of memory-1 strategies on the lattice. The payoffs are (a) $(T, P) = (1.4, 0.05)$, (b) $(T, P) = (1.9, 0.2)$, (c) $(T, P) = (1.9, 0.8)$, (d) $(T, P) = (1.4, 0.5)$. The strategies ALLD, TFT, ATFT, and ALLC are represented in gray-scale from white to black.

S_2 should be able to successfully deal with both mistakes and exploiting strategies. The simulation ends with one of the strategies using the S_2 pattern for error correction, i.e., two rounds of mutual defection after a mistake. There is a large number of strategies in this group, and longer simulations show that in the CA model there is usually a certain memory-5 strategy that takes over the whole population, since even very small score differences may lead to the total dominance of one of the strategies. In the CA model there are also examples of spatial coexistence between strategies of longer memory, that were not observed in the mean-field model. For some examples of these phenomena, see Lindgren and Nordahl.[27]

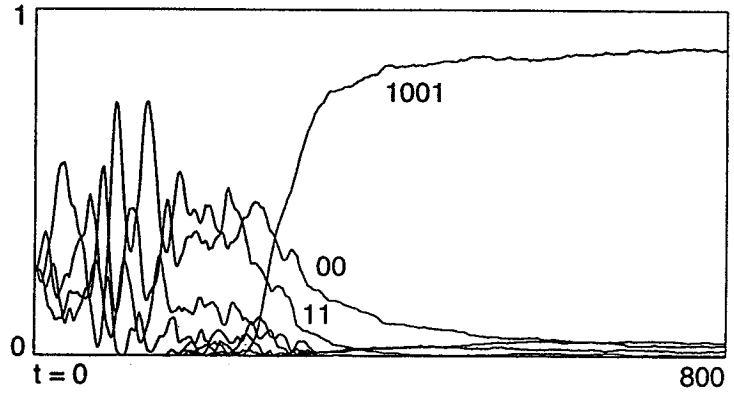

FIGURE 10 The evolution of finite memory strategies in the CA model is shown for 800 generations, with the parameter values $T = 1.6$, $P = 0.3$, cf. Figure 6. There is a metastable coexistence between 00 (ALLD) and 11 (ALLC), which is based on a spatial pattern of ALLC islands surrounded by ALLD, similar to what is seen in Figure 9(d). Still the 1001 strategy dominates at the end.

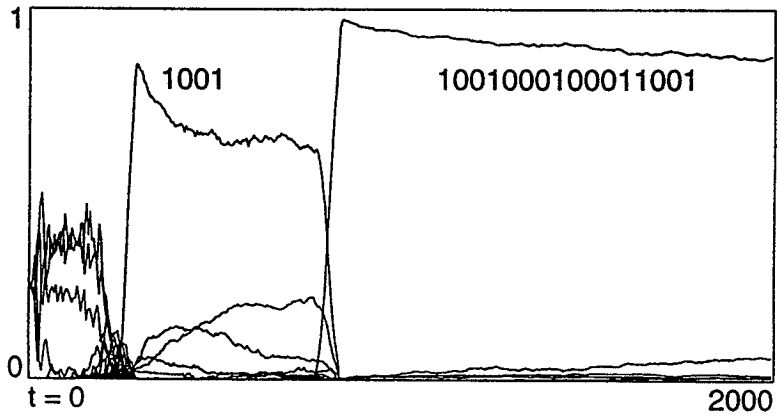

FIGURE 11 The evolution of finite memory strategies in the CA model is shown for 2000 generations, with the parameter values $T = 1.7$, $P = 0.4$, cf. Figure 7. Initially, during the first 200 generations, there is a metastable coexistence between TFT, ALLC, and ALLD in the form of spiral waves, similar to the pattern in Figure 9(b). The spatial dimension allows 1001 to dominate in a metastable state, due to the possibility to form cooperative regions. At the end, the S_2-like strategy from Figure 7 comes to dominate.

Before 1001 dominates the system, in Figure 11, there is a period with three strategies coexisting. This is due to a spatiotemporal pattern of spiral waves involving the strategies TFT, ALLD, and ALLC, similar to Figure 9(b). Again, we see an example of a locally unstable composition of strategies that can be globally stable in the CA model.

FINITE AUTOMATA STRATEGIES

As a final modification of the previous models, we shall look at a different representation of strategies for the game. A natural approach is to use finite automata, and this has been done before by Miller[32] and by Stanley et al.,[46] see also Nowak et al.[39] The novelty in the approach presented here is the combination of an evolvable genome size, the analytically solved iterated Prisoner's Dilemma with mistakes, and the finite automata representation.

The reason why this representation could be interesting to investigate is that there are reasons to believe that more successful strategies may evolve. In fact, there is an extremely simple, cooperative, uninvadeable, and error-correcting strategy in this representation.

This strategy was suggested by Sugden,[49] as a robust strategy that can maintain cooperation in the presence of mistakes. We call the strategy "Fair," since the error-correcting mechanism is built on the idea of a fair distribution of the scores. The Fair strategy can be in one of three possible states: "satisfied" (cooperating), "apologizing" (cooperating), and "angry" (defecting). In the satisfied state Fair cooperates, but if the opponent defects, Fair switches to the angry state and defects until the opponent cooperates, before returning to the satisfied state. If, on the other hand, Fair accidentally defects, the apologizing state is entered and Fair stays in this state until it succeeds in cooperating. Then it returns to the satisfied state again, regardless of the opponent's action. The strategy can be described as a finite automaton, see Figure 12.

This means that two players using the Fair strategy will quickly get back to cooperation after an accidental defection. The player making the mistake enters the apologizing state and the opponent the angry state, so that in the next round they will switch actions making the scores fair. Thus, a sequence of rounds including an accidental defection D looks like $(\ldots, \mathrm{CC}, \underline{\mathrm{D}}\mathrm{C}, \mathrm{CD}, \mathrm{CC}, \mathrm{CC}, \ldots)$. (Only if the apologizing player fails to cooperate, for example by another mistake, they will both stay in the apologizing and angry state, respectively, until the apologizing player succeeds in a C action.)

Before we go on discussing the evolutionary simulations based on finite automata, we shall take a look at the game-theoretic properties of the Fair strategy. The stability properties of the Fair strategies were investigated by Sugden,[49] and

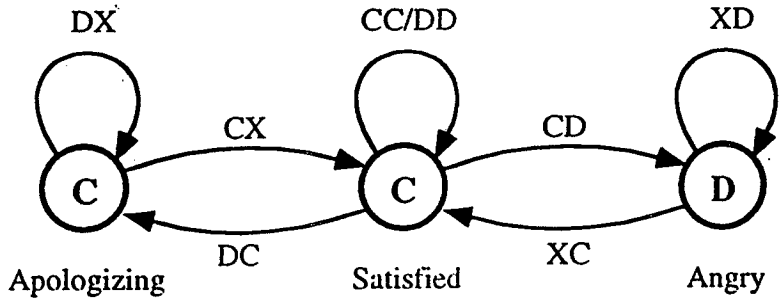

FIGURE 12 The Fair strategy can be represented as a finite automaton with three states. The symbols in the state circles denote the action that is intended, though a mistake may alter the action. After each round of the game, the performed actions are inspected and the strategy enters the state given by the appropriate transition. (On the transition arcs, the left and right symbols correspond to the action performed by the strategy itself and the opponent, respectively, where an X denotes an arbitrary action.) The anti-symmetry assures that when two Fair players meet they will either both be in the satisfied state, or they will be in the Apologizing and the Fair state, respectively, from where they simultaneously return to the satisfied state.

the analysis was extended by Boyd,[9] who labeled the strategy "contrite-tit-for-tat." Wu and Axelrod[52] showed the advantage of the Fair strategy in a population dynamics simulation, starting with a population mainly consisting of the 63 strategies submitted to the second tournament organized by Axelrod.[3]

In the following analysis, we are assuming the infinitely iterated PD with a small probability e for mistakes, i.e., the performed action is not the intended one. In the game between two Fair strategies, the probability is $1/(1+2e)$ to be in the satisfied state and $e/(1+2e)$ for both the apologizing and the angry state.[2] This leads to an average score of $R + 2e(S+T-2R)$, if terms of order e^2 or higher are neglected. (We are assuming small e, and in the following we shall make a first-order analysis, omitting terms of order e^2 or higher.) For the standard parameter values used by Axelrod ($R=3$, $T=5$, $S=0$, and $P=1$) and $e=0.01$, the score is 2.98. This is close to the score that two ALLC strategies get when they meet,[3] $R+e(S+T-2R)$, or, for the same parameters, 2.99.

For a small probability for mistakes (in the limit $e \to 0$), a new player entering a population of Fair players must adopt the same strategy not to be out-competed: If Fair is in the satisfied state, a defect action (by the invader) scoring T must be

[2] A detailed analysis will be presented elsewhere.

[3] Even if ALLC is more cooperative, it cannot invade a population of Fair players for small values of e, since the score received by an invading ALLC, $R+e(2S+T-3R)$, differs by $e(R-T) < 0$ compared to the score for Fair, if terms of order e^2 are neglected.

compensated for by an apology scoring only S, before Fair enters the satisfied state again, or Fair will stay in the angry state giving a score of P. Since $S+T < 2R$ and $P < R$, the invader would have done better cooperating with Fair in the satisfied state. This also implies that when Fair expects an apology, the opponent had better cooperate. In the apologizing state, Fair does not take any notice of the opponent's action and, therefore, the opponent should defect.[4]

REPRESENTATION

Each individual has a strategy represented by a finite automaton like the one in Figure 12. The size of the automaton is not fixed in the evolution, but nodes may be added by mutation increasing the number of internal states. Other mutations may alter the actions associated with the internal states, move transition arcs, and change start node.

The game between two FA strategies is a stochastic process on a finite automaton representing the game between the two strategies. This "game" automaton has, as its internal states, the pairs of possible internal states for the two players, and the transitions are given by the transitions of the two strategies. The average score for the players is then calculated from probability distribution over the internal states of the game automaton.

Note that all finite memory strategies can be represented by FA, but that the FA representation results in a larger evolutionary search space. As in the finite memory models, we start with a population consisting of equal numbers of the memory-1 strategies, ALLD, ALLC, TFT, and ATFT, which can be represented as FA like in Figure 13.

The mutations we allow are the following. The *node mutation* changes the action associated with a certain node. The *transition mutation* changes the destination node of a transition. The *start mutation* changes the start node of the FA. The *growth mutation* adds a node to the FA. This is done in a neutral way by adding a copy of a node that has a transition to itself. After the mutation, this transition connects the original node with a new one that is identical to the original one, see Figure 14. There is also a *delete mutation* that randomly removes a node of the FA. The mutation rate is p_{mut} per node for all mutations, except for the transition mutation that has a rate of $4p_{\text{mut}}$ per node, i.e., p_{mut} per transition. In the mean-field model we have used a rate of $p_{\text{mut}} = 10^{-5}$. In the lattice model, we have used the same mutation rate per individual for all types of mutations, $p_{\text{ind}} = 10^{-3}$.

[4]The Fair strategy can only be invaded if the constants of the payoff matrix are close to the limits, e.g., the punishment score P for mutual defection is not much less than the reward for mutual cooperation R, $R \approx P$, or if the risk for mistake is high. Fixing $S = 0$, one finds that the mistake probability must be of the order $e \approx (1 - P/R)/(4 - 3P/R)$ for ALLD to exploit the Fair strategy. The punishment score P must not be less than $0.99R$, for ALLD to invade with an e of 1%. For the standard parameter values, an error probability of $e \approx 2/9 \approx 22\%$ is required for ALLD to invade.

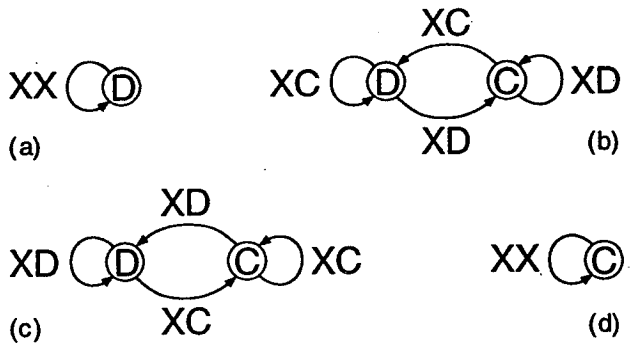

FIGURE 13 The initial state of the finite automata (FA) models consists of equal numbers of four FA that correspond to the memory-1 strategies: (a) ALLD, (b) ATFT, (c) TFT, and (d) ALLC.

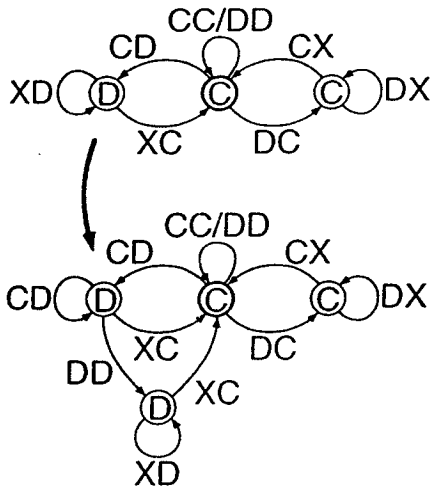

FIGURE 14 The growth mutation can only be applied if the parent FA has a node with a transition to itself. In that case, the mutation creates a copy of the original node with transitions *from* the copied node identical to the original one. The self-transition of the original node is directed to the copy, while the self-transition of the copy remains. This results in a mutation that increases the number of internal states without changing the strategy—the mutation is neutral.

MEAN-FIELD MODEL

Even if we start with the same strategies as in the finite memory representation, the evolutionary paths usually look very different since the evolutionary search space is dramatically increased and the points (in this space) corresponding to the finite memory strategies are not close to each other (in terms of mutations). It turns out that in the mean-field model with finite automata (FA) representation, the ALLD strategy more often takes over the whole system, compared with the finite memory representation. In the simulations shown below, we have chosen a smaller population size ($N = 500$) than in the previous case, which reduces the probability for ALLD to dominate.

In the first illustration of the evolution of FA, Figure 15, the payoff is $T = 1.6$ and $P = 0.3$. After the initial oscillations dominated by TFT, a strategy (not of finite memory type) that only differs from TFT by a transition mutation appears. This strategy is slightly better than TFT since it does not leave the cooperative state if both players defected last round, while in the defecting state it requires a cooperation from the opponent to get back to its cooperating state. The strategies that dominate the picture later on are similar to the finite memory strategy 1001, showing the same pattern for correcting for single mistakes. The slowly growing strategy at which the arrow is pointing in the bottom right-hand corner is the FA representation of 1001.

In the second example, Figure 16, the payoff is $T = 1.7$ and $P = 0.4$, which means that the error-correcting mechanism of 1001-like strategies can be exploited.

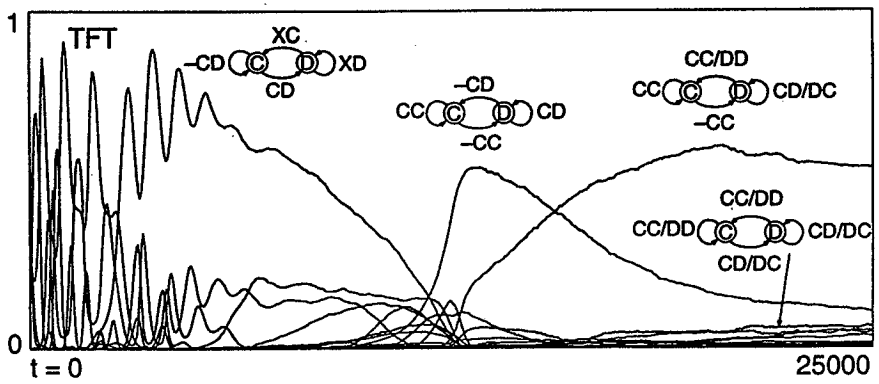

FIGURE 15 The evolution of finite automata (FA) is shown for 25000 generations in the mean-field model, with payoff parameters $(T, P) = (1.6, 0.3)$. Some of the automata that evolve are drawn on top of their abundancy curves, respectively. The FA, drawn at the bottom toward the end of the simulation, is the FA representation of the finite memory strategy 1001. Also, the dominating strategy at the end has an error-correcting mechanism similar to 1001.

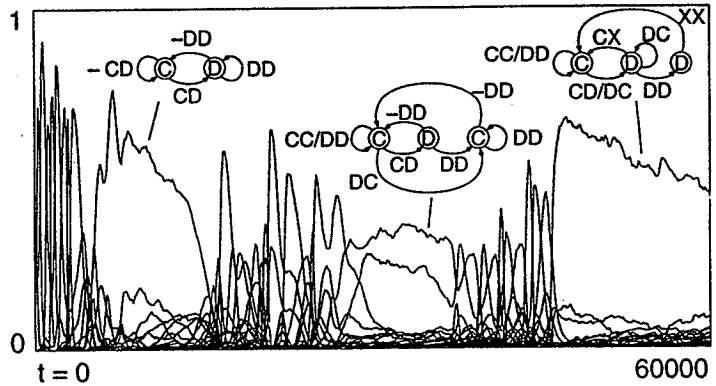

FIGURE 16 The evolution of finite automata (FA) is shown for 60000 generations in the mean-field model, with payoff parameters $(T, P) = (1.7, 0.4)$. Here, strategies corresponding to 1001 (S_1) should not be stable, cf. Figure 8. The automaton shown in the middle is related to the Fair strategy in its mechanism for dealing with accidental defections, but it can be exploited. Instead, an FA that defects twice, after an isolated defect action, dominates at the end similar to the finite memory strategy S_2.

The strategy shown in the middle of the figure, dominating for more than 10000 generations, resembles the Fair strategy, but it is too forgiving, which implies that it will be exploited. The strategy shown at the end, dominating for more than 15000 generations, has an error correction similar to S_2, and it is possible (but not certain) that this simulation would end with a mixture of strategies of this type.

LATTICE MODEL

The spatial dimension again increases the chances for cooperative strategies to succeed in the population. Many simulations end with cooperative strategies, like Fair, dominating the whole lattice for large parts of the parameter region. This is in accordance with the analysis of the Fair strategy above.

In the first example, Figure 17, we have the parameters $(T, P) = (1.7, 0.4)$ on a 64×64 lattice. After the initial oscillations we again see the S_1-like strategy that dominated at the end of the mean-field simulation of Figure 15. This behavior is not stable here, according to Eq. (4), and the simulation continues with a sequence of FA with three and four states, all having an S_2-like mechanism for correcting mistakes. There are many FAs that coincide with the pattern of two mutual defect actions after an isolated mistake, but that differ when it comes to second-order effects, i.e., when there is a mistake within the error-correcting pattern. These second-order effects may be sufficient to give a decisive advantage so that only one of them will dominate.

In the second example, Figure 18, we are in the upper right-hand corner of the payoff matrix parameter region (1.9, 0.9), see Figure 8. For these parameters, not even S_9 will be stable, but, according to Eq. (4), ten rounds of mutual defection are required to punish exploiting strategies. The simulation shows a metastable state dominated by a pair of strategies, of which one is a three-state FA that

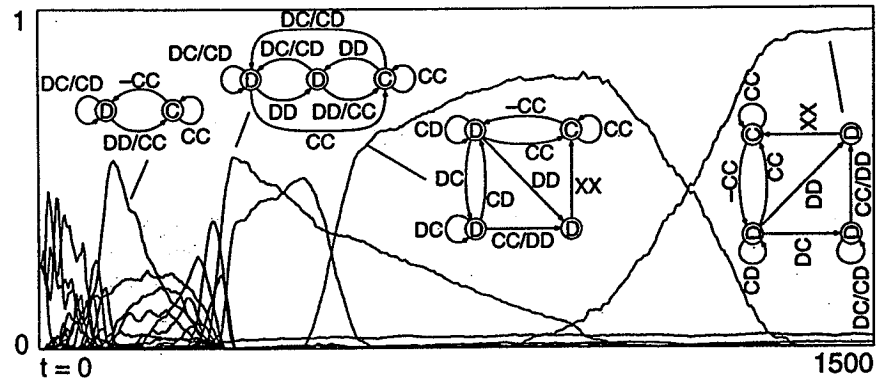

FIGURE 17 The evolution of finite automata for the cellular automaton model is shown for 1500 generations, with the parameter values $(T, P) = (1.7, 0.4)$. The automata of three and four states that evolve all show the S_2-mechanism of the finite memory strategies for error correction.

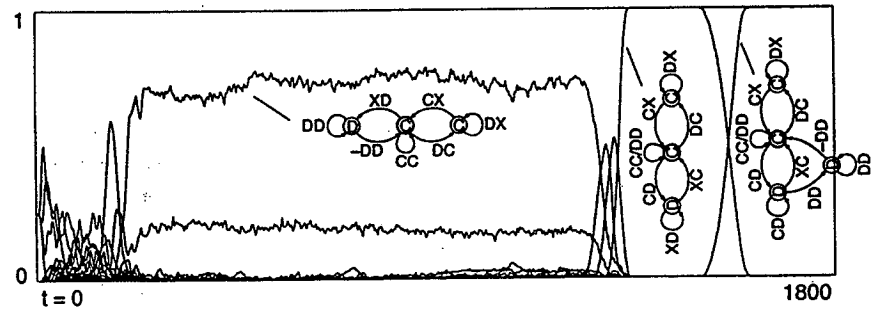

FIGURE 18 The evolution of finite automata for the cellular automaton model is shown for 1800 generations, with the parameter values $(T, P) = (1.9, 0.9)$. This is close to the upper right-hand corner of the payoff parameter region, see Figure 8. The chance that an S_n strategy would evolve here is small since the stability requirement Eq. (4) implies that n would be 10. In the simulation, we see that a predecessor to Fair dominates for a long time until Fair appears, but at the end there is a four-state variation of Fair that manages to take over.

resembles the Fair strategy. The Fair strategy is two transition mutations away, and the intermediate FA can be seen as one of the two narrow peeks before Fair appears, cf. Figure 12. However, there is a four-state variation of Fair that succeeds in out-competing the original Fair. The difference lies in a fourth-order effect, i.e., involving three mistakes in the apology pattern following one mistake. This strategy only differs from Fair by a (neutral) growth mutation and a transition mutation.

Other simulations, from various parts of the parameter region of the payoff matrix in Figure 8, show that Fair is successful and not very sensitive to the parameter choices. In most cases, the simulations end with one strategy dominating the whole lattice, a result that is not general in the CA simulations of the finite memory strategies. From the evolutionary point of view, these simulations may seem less interesting in the sense that the evolutionary transients appear to be short. On the other hand, they demonstrate the important difference between the spatial and the nonspatial models, as well as between different evolutionary search spaces, and that cooperative behavior may not be given the chance unless there is some compartmentalization possible, here in the form of spatial locality.

The examples presented here are, of course, preliminary, and a more detailed study of the evolution of finite automata strategies will be reported elsewhere.

MISUNDERSTANDING VS. MISTAKES

As a possibility for future modifications of these models, one could consider a different kind of noise. Mistakes are a relatively simple form of noise to deal with in games, since both players have the same history of actions to look at. Misunderstandings, on the other hand, leave the two players with different histories, and one cannot be sure whether the distribution of scores is fair or not even for a few rounds. This is of course critical for strategies like Fair, since a misunderstanding may lead them into situations where one player is in a satisfied state while the other one is in an angry state, which may result in sequences of alternating defect and cooperate actions, as in the TFT strategy.

The class of evolutionarily stable strategies (S_n) in the case of mistakes, discussed before, turn out to be error correcting also in the case of misunderstanding, at the same time as there is a possibility of avoiding both exploiting strategies and strategies that are more cooperative.

In the case of *misunderstanding*, two players using, e.g., the S_2 strategy, will show the same pattern as in the case of mistake, (..., CC, CC*, DC, DD, DD, CC, CC, CC, ...), where C* denotes a C action that is perceived as a D action by the opponent. As before, the length of the "quarrel" must be large enough to prevent potential exploiters to defect on purpose, which leads to the inequality requirement of Eq. (3).

But when misunderstandings are involved, there is another possibility to invade a population of S_n strategies, which leads to an additional requirement for the choice of n. If a player ignores a perceived D action from its opponent and still cooperates the score may increase, since the action could have been C, and then the cost of quarrel could have been be avoided, e.g., $(\ldots, \mathrm{CC}, \mathrm{CC}^*, \mathrm{CC}, \ldots)$. On the other hand, if the action was truly D, then the opponent will defect again, and the return to cooperation will be delayed. For example, ignoring a defect action from S_2 leads to the pattern $(\ldots, \mathrm{C}^*\mathrm{C}, \mathrm{C}\underline{\mathrm{D}}, \mathrm{CD}, \mathrm{DD}, \mathrm{DD}, \mathrm{CC}, \mathrm{CC}, \ldots)$, where $\underline{\mathrm{D}}$ is a defect action due to S_2 misunderstanding the previous action C^*. Since the risk for misunderstanding is equal for both players, the two patterns above that can result from ignoring a perceived D action occur with equal probability. The expected score for ignoring instead of answering directly is $\sigma_0 = (n+2)R/2 + (2S + nP)/2$, to compare with the score for S_n, $\sigma_n = (S+T)/2 + nP + R$, for the rounds following a defect action. Thus, for S_n to avoid invasion, the inequality $\sigma_n > \sigma_0$ must be fulfilled. This can be rewritten as a requirement on n, $n < (T-S)/(R-P)$, or $n < (T-R)/(R-P) + 1 + (P-S)/(R-P)$. Since the third term is positive, there is always an integer n satisfying both conditions

$$\frac{T-R}{R-P} < n < \frac{T-R}{R-P} + 1 + \frac{P-S}{R-P}. \tag{5}$$

For the parameters used by Axelrod ($R = 1, S = 0, T = 5, P = 1$), we get $n = 2$, and the score is 2.92. Contrary to Fair, the error correction of these strategies works for both mistake and misunderstanding, but at a cost of a lower average score.

DISCUSSION

By showing examples from various models based on the Prisoner's Dilemma as the interaction between individuals, I have illustrated some aspects of evolutionary models. The differences between the mean-field models and the lattice models are evident in all these cases. The possibility of spatiotemporal phenomena may give rise to a stable coexistence between strategies that would otherwise be outcompeted. These spatiotemporal structures may take the form of spiral waves, irregular waves, spatiotemporal chaos, frozen patchy patterns, and various geometrical configurations.[28] This is very clear in the first model (finite game) where ALLD is the only fixed point in the mean-field model, but where ALLD seldom appears in the simulations of the CA model. Space also allows for increased cooperative behavior by a mechanism that can be described as kin selection. For example, regions of ALLC can be stable in an ALLD environment, since the ALLC individuals at the border are replaced (in the next generation) by more successful ALLC individuals from the interior of the region.

The choice of representation is important, as the difference between the simulations of the finite memory and FA strategies illustrates. Still, the existence of a stable fixed point, e.g., in the form of the Fair strategy, does not necessarily mean that the evolutionary dynamics always finds it. Of course, the regions of attraction of the fixed points are important to characterize the evolutionary search space. This may, however, be a very difficult task since the space of strategies is potentially infinite. The final example also illustrates that the region of attraction may increase dramatically if the spatial dimension is added. The increased probability for the Fair strategy to succeed probably depends on the advantage for cooperative strategies in the spatial world. The Fair strategy may more easily survive in a local situation from where it may then spread.

The models presented here all have in common that they are very simple and can be easily simulated. Despite their simplicity, a number of complex phenomena in the form of evolutionary dynamics and spatiotemporal patterns can be observed. Models similar to these (in terms of simplicity) could be constructed in a more economic context to study economic phenomena like, e.g., trading and price formation. Some of the ingredients of the models presented here should be considered in that work, for example, the advantage of starting small in combination with neutral mutations. (This has been used in another context, where evolving artificial neural networks were selected for recognizing strings of a regular language.[26]) The workshop has inspired us to start constructing economics models along these lines, and we hope to report on these in future contributions.

ACKNOWLEDGMENT

This research is supported by the Swedish Natural Science Research Council. I thank Robert Axelrod for pointing out the reference to Sugden concerning the Fair strategy. I also thank the Santa Fe Institute for supporting several visits during which parts of the work presented here were done.

REFERENCES

1. Anderson, P. W. "Suggested Model for Prebiotic Evolution: The Use of Chaos." *Proc. Natl. Acad. Sci.* **80** (1983): 3386–3390.
2. Arthur, W. B. "Asset Pricing Under Endogenous Expectations in an Artificial Stock Market." In *The Economy as an Evolving Complex System II,* edited by W. B. Arthur, S. N. Durlauf, and D. Lane, 1–30. Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. XXVII. Reading, MA: Addison-Wesley, 1997.
3. Axelrod, R. *The Evolution of Cooperation.* New York: Basic Books, 1984.
4. Axelrod, R. "The Evolution of Strategies in the Iterated Prisoner's Dilemma." In *Genetic Algorithms and Simulated Annealing,* edited by L. Davis, 32–41. Los Altos, CA: Morgan-Kaufmann, 1987.
5. Axelrod, R., and D. Dion. "The Further Evolution of Cooperation." *Science* **242** (1988): 1385–1390.
6. Bagley, R. J., J. D. Farmer, and W. Fontana. "Evolution of a Metabolism." In *Artificial Life II,* edited by C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, 141–158. Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. X. Redwood City, CA: Addison-Wesley, 1991.
7. Boerlijst, M. C., and P. Hogeweg. "Spiral Wave Structure in Prebiotic Evolution: Hypercycles Stable Against Parasites." *Physica D* **48** (1991): 17–28.
8. Boerlijst, M. C., and P. Hogeweg. "Self-Structuring and Selection: Spiral Waves as a Substrate for Prebiotic Evolution." In *Artificial Life II,* edited by C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, 255–276. Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. X. Redwood City, CA: Addison-Wesley, 1991.
9. Boyd, R. "Mistakes Allow Evolutionary Stability in the Repeated Prisoner's Dilemma Game." *J. Theor. Biol.* **136** (1989): 47–56.
10. Boyd, R., and J. P. Lorberbaum. "No Pure Strategy Is Evolutionarily Stable in the Repeated Prisoner's Dilemma Game." *Nature* **327** (1987): 58–59.
11. Eigen, M. "Self-Organization of Matter and the Evolution of Biological Macro-Molecules." *Naturwissenschaften* **58** (1971): 465–523.
12. Eigen, M., and P. Schuster. *The Hypercycle: A Principle of Natural Self-Organization.* Berlin: Springer-Verlag, 1979.
13. Eigen, M., J. McCaskill, and P. Schuster. "Molecular Quasi-Species." *J. Phys. Chem.* **92** (1988): 6881–6891.
14. Farmer, J. D., S. A. Kauffman, and N. H. Packard. "Autocatalytic Replication of Polymers." *Physica* **22D** (1986): 50–67.
15. Fontana, W. "Algorithmic Chemistry." In *Artificial Life II,* edited by C. G. Langton and C. Taylor, 313–324. Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. X. Redwood City, CA: Addison-Wesley, 1991.
16. Hillis, D. "Coevolving Parasites Improve Simulated Evolution as an Optimization Procedure." *Physica D* **42** (1990): 228–334.

17. Hofbauer, J., and J. W. Weibull. "Evolutionary Selection Against Dominated Strategies." Paper presented at The Sixth Conference of the International Joseph A. Schumpeter Society, 1996.
18. Holland, J. H. *Adaptation in Natural and Artificial Ecosystems.* 2nd ed. Cambridge, MA: MIT Press, 1992.
19. Holland, J. H. "Echoing Emergence: Objectives, Rough Definitions, and Speculations for Echo-Class Models." Working Paper # 93-04-023, Santa Fe Institute, Santa Fe, NM, 1993.
20. Ikegami, T. "From Genetic Evolution to Emergence of Game Strategies." *Physica D* **75** (1994): 310–327.
21. Kauffman, S. A. "Autocatalytic Sets of Proteins." *J. Theor. Biol.* **119** (1986): 1–24.
22. Kauffman, S. A. *The Origins of Order: Self-Organization and Selection in Evolution.* New York: Oxford University Press, 1993.
23. Kauffman, S. A., and S. Johnsen. "Coevolution to the Edge of Chaos: Coupled Fitness Landscapes, Poised States, and Coevolutionary Avalanches." In *Artificial Life II,* edited by C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, 325–369. Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. X. Redwood City, CA: Addison-Wesley, 1991.
24. Langton, C. G. Paper presented at The Economy as an Evolving Complex Systems, Santa Fe Institute, Santa Fe, NM, 1995.
25. Lindgren, K. "Evolutionary Phenomena in Simple Dynamics." In *Artificial Life II,* edited by C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, 295–311. Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. X. Redwood City, CA: Addison-Wesley, 1991.
26. Lindgren, K., A. Nilsson, M. G. Nordahl, and I. Råde. "Regular Language Inference Using Evolving Recurrent Neural Networks." In *COGANN-92 International Workshop on Combinations of Genetic Algorithms and Neural Networks,* edited by L. D. Whitley and J. D. Schaffer, 75–86. Los Alamitos, CA: IEEE Computer Society Press, 1992.
27. Lindgren, K., and M. G. Nordahl. "Evolutionary Dynamics of Spatial Games." *Physica D* **75** (1994): 292–309.
28. Lindgren, K., and M. G. Nordahl. "Cooperation and Community Structure in Artificial Ecosystems." *Artificial Life J.* **1** (1994): 15–37.
29. Lorberbaum, J. *J. Theor. Biol.* **168** (1994): 117–130.
30. Matsuo, K. "Ecological Characteristics of Strategic Groups in 'Dilemmatic World.'" In *Proceedings of the IEEE International Conference on Systems and Cybernetics,* 1071–1075, 1985.
31. Maynard-Smith, J. *Evolution and the Theory of Games.* Cambridge, MA: Cambridge University Press, 1982.
32. Miller, J. H. "The Coevolution of Automata in the Repeated Iterated Prisoner's Dilemma." Working Paper No. 89-003, Santa Fe Institute, Santa Fe, NM, 1989.

33. Mitchell, M., and S. Forrest. "Genetic Algorithms and Artificial Life." *Artificial Life J.* **1** (1994): 267-289.
34. Molander, P. "The Optimal Level of Generosity in a Selfish, Uncertain Environment." *J. Conflict Resolution* **29** (1985): 611–618.
35. Nowak, M. A., and R. M. May. "Evolutionary Games and Spatial Chaos." *Nature* **359** (1993): 826–829.
36. Nowak, M. A., and R. M. May. "The Spatial Dilemmas of Evolution." *Intl. J. Bif. & Chaos* **3** (1993): 35–78.
37. Nowak, M. A., and K. Sigmund. "Tit-For-Tat in Heterogenous Populations." *Nature* **355** (1992): 250–253.
38. Nowak, M. A., and K. Sigmund. "A Strategy of Win-Stay, Lose-Shift that Outperforms Tit-For-Tat in the Prisoner's Dilemma Game." *Nature* **364** (1993): 56–58.
39. Nowak, M. A., K. Sigmund, and E. El-Sedy. "Automata, Repeated Games, and Noise." *J. Math. Biol.* **33** (1995): 703–722.
40. Rapoport, A., and A. M. Chammah. *Prisoner's Dilemma.* Ann Arbor: University of Michigan Press, 1965.
41. Rasmussen, S., C. Knudsen, and R. Feldberg. "Dynamics of Programmable Matter." In *Artificial Life II,* edited by C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, 211–254. Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. X. Redwood City, CA: Addison-Wesley, 1991.
42. Ray, T. S. "An Approach to the Synthesis of Life." In *Artificial Life II,* edited by C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, 371-408. Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. X. Redwood City, CA: Addison-Wesley, 1991.
43. Rokhsar, D. S., P. W. Anderson, and D. L. Stein. "Self-Organization in Prebiological Systems: Simulation of a Model for the Origin of Genetic Information." *J. Mol. Evol.* **23** (1986): 110.
44. Samuelson, L., and J. Zhang. "Evolutionary Stability in Assymmetric Games." *J. Econ. Theory* **57** (1992): 363–391.
45. Schuster, P. "The Physical Basis of Molecular Evolution." *Chemica Scripta* **26B** (1986): 27.
46. Stanley, E. A., D. Ashlock, and L. Tesfatsion. "Iterated Prisoner's Dilemma with Choice and Refusal of Partners." In *Artificial Life III,* edited by C. G. Langton, 131–175. Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. XVII. Redwood City, CA: Addison-Wesley, 1993.
47. Stein, D. L., and P. W. Anderson. "A Model for the Origin of Biological Catalysis." *Proc. Natl. Acad. Sci.* **81** (1984): 1751–1753.
48. Stenseth, N. C., and J. Maynard-Smith. "Coevolution in Ecosystems: Red Queen Evolution or Stasis?" *Evolution* **38** (1984): 870–880.
49. Sugden, R. *The Economics of Rights, Co-operation and Welfare.* Oxford: Basil Blackwell, 1986.
50. Tesfatsion, L. "How Economists Can Get ALife." In *The Economy as an Evolving, Complex System II,* edited by W. B. Arthur, S. Durlauf and D. Lane, 533–565. Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. XXVII. Reading, MA: Addison-Wesley, 1997.
51. Weisbuch, G. *C. R. Acad. Sci. III* **298** (1984): 375–378.
52. Wu, J., and R. Axelrod. "How to Cope with Noise in the Iterated Prisoner's Dilemma." *J. Conflict Resolution* **39** (1995): 183–189.