

8

Games on Grids

Martin A. Nowak and Karl Sigmund

8.1 Introduction

The theory of games and the theory of cellular automata seem at first glance to be totally unrelated (despite being created at about the same time – some 50 years ago – and boasting the same father, John von Neumann). In a recent and rather surprising development, the two disciplines have been brought together. The resulting *spatial evolutionary game theory*, which was first used to shed light on the emergence of cooperation, has grown rapidly during the past five years and has proved useful in other biological and economic contexts. In this survey chapter, we concentrate on the game-theoretical aspect, as the other contexts are well covered in other chapters of this volume. Thus, we stress mainly the effects spatial structures have on frequency-dependent selection.

Let us start with an arbitrary 2×2 game (Box 8.1), that is, a game between two players each having two strategies: (1) to cooperate (denoted by C) and (2) to defect (denoted by D). A player using C receives a payoff R (the reward) if the co-player uses C , and S (the sucker's payoff) if the co-player uses D . A player using D obtains the payoff T (the temptation) against a C -player, and P (the punishment) against a D -player. We interpret C as helping in a common enterprise and D as withholding help; therefore, we assume that the payoff R for two C -players is larger than the payoff P for two D -players. Because the evolution of cooperation is only interesting if it is threatened by unilateral defection – that is, if a unilateral D -move is good for the defector and bad for the cooperator – we assume $T > R$ and $R > S$. What is *a priori* less clear is the ranking of P and S . The ranking $P > S$, gives the well-known Prisoner's Dilemma game, where D is the dominant strategy. This game has been used since the early 1970s (Trivers 1971) as a paradigm for the evolution of cooperation (see also Axelrod and Hamilton 1981; Sigmund 1995). However, one can argue that $P < S$ is also an interesting situation: it reflects the dilemma of two players who each

Box 8.1 Payoff matrices for symmetric 2×2 games

	If the co-player plays C	If the co-player plays D
If I play C , I receive	Reward R	Sucker's payoff S
If I play D , I receive	Temptation T	Punishment P

A payoff matrix like the one above describes a Prisoner's Dilemma if its matrix elements obey the following inequalities:

$$T > R > P > S .$$

The Chicken game (also known as the Snowdrift game or the Hawk–Dove game) is defined by

$$T > R > S > P .$$

The following two payoff matrices provide specific examples:

$$\begin{pmatrix} 3 & 0 \\ 4 & 1 \end{pmatrix}$$

Prisoner's Dilemma

$$\begin{pmatrix} 3 & 1 \\ 4 & 0 \end{pmatrix}$$

Chicken game

prefer to play C even if the other plays D . This has been described as the Snowdrift game by Sugden (1986): even if one player refuses to help, the other would be prepared to dig a path through the snowdrift for both, rather than sit tight and freeze. The ranking $P < S$ leads to $T > R > S > P$, and hence to the game called Chicken by classical game theorists, and Hawk–Dove by evolutionary biologists.

In the usual setting of evolutionary game theory, one assumes a well-mixed population of players matched randomly against each other and multiplying at a rate that increases in proportion to their payoff. In the Chicken game, this leads to a mixed population of C - and D -players. In the Prisoner's Dilemma game, the C -players are doomed to extinction. If the probability that players meet for another round (and recognize each other) is sufficiently high, however, then cooperation can be an evolutionarily viable outcome in the resulting *repeated* Prisoner's Dilemma game. In particular, there exist populations using cooperative strategies that cannot be invaded by minorities of players who always defect. This point was made most forcefully in Axelrod's book *The Evolution of Cooperation* (1984).

In Chapter 8 of his book, Axelrod describes an iterated Prisoner's Dilemma game played on a spatial grid. He considers a territory subdivided into square cells, with each cell occupied by one player. Each player plays against his or her four neighbors in a repeated Prisoner's Dilemma game. The total scores are computed. If a player has neighbors who are

more successful, he or she switches to the strategy that obtained the highest score (if there is a tie among the most successful neighbors, one is picked at random). Thus, Axelrod views neighbors as role models whose behavior can be imitated. But one can obviously also interpret the updating differently, as the formation of a new generation, with each cell being taken over by an offspring of the previous owner or of one of the neighbors, depending on who did best in the previous generation – a kind of colonization. Axelrod shows that it is at least as easy for a strategy to protect itself from a takeover in such a territorial structure as it is if the co-players are chosen randomly in the population (i.e., without regard to spatial structure).

Most important, Axelrod shows how a single invader playing the strategy of always defecting (*AD*) can spread in a population of *Tit For Tat* (*TFT*) players (who cooperate in the first round and then do whatever their co-player did in the previous round). This yields fascinating snowflake-like patterns of defectors bypassing islands of cooperators (Figure 8.1).

8.2 One-round Games

In Axelrod's investigations, territoriality was primarily seen as a device for ensuring continuity of the interaction: after defection, the defector cannot run away, but must face an eventual retribution.

Interestingly, cooperators can persist even if there is *no* retribution, that is, even if the game is never repeated. Persistence was first demonstrated by Nowak and May (1992; see also Sigmund 1992; Nowak *et al.* 1995a). Like Axelrod, Nowak and May considered a large lattice with each cell occupied by one player. The players engage in *one* round of the Prisoner's Dilemma game against each of their neighbors. [This could be the four neighbors to the north, south, east, and west, or the eight nearest neighbors corresponding to a chess king's move. Somewhat in the tradition of Axelrod (in whose tournaments each strategy played against itself as well as against other strategies), each player on the lattice could also interact with him- or herself. The main results of the spatial model are independent of these details.] The players are either cooperators or defectors, that is, they play either *C* or *D* in each of the one-shot Prisoner's Dilemma games against their neighbors. Afterward, the next generation is formed: each cell is taken over by a copy of the highest-scoring strategy within the neighborhood (consisting of the previous owner of the site and the nearest neighbors). This simultaneous updating leads to a deterministic transition rule and defines a cellular automaton with the interesting property that the transition rule depends, not only on the states of the nearest neighbors, but

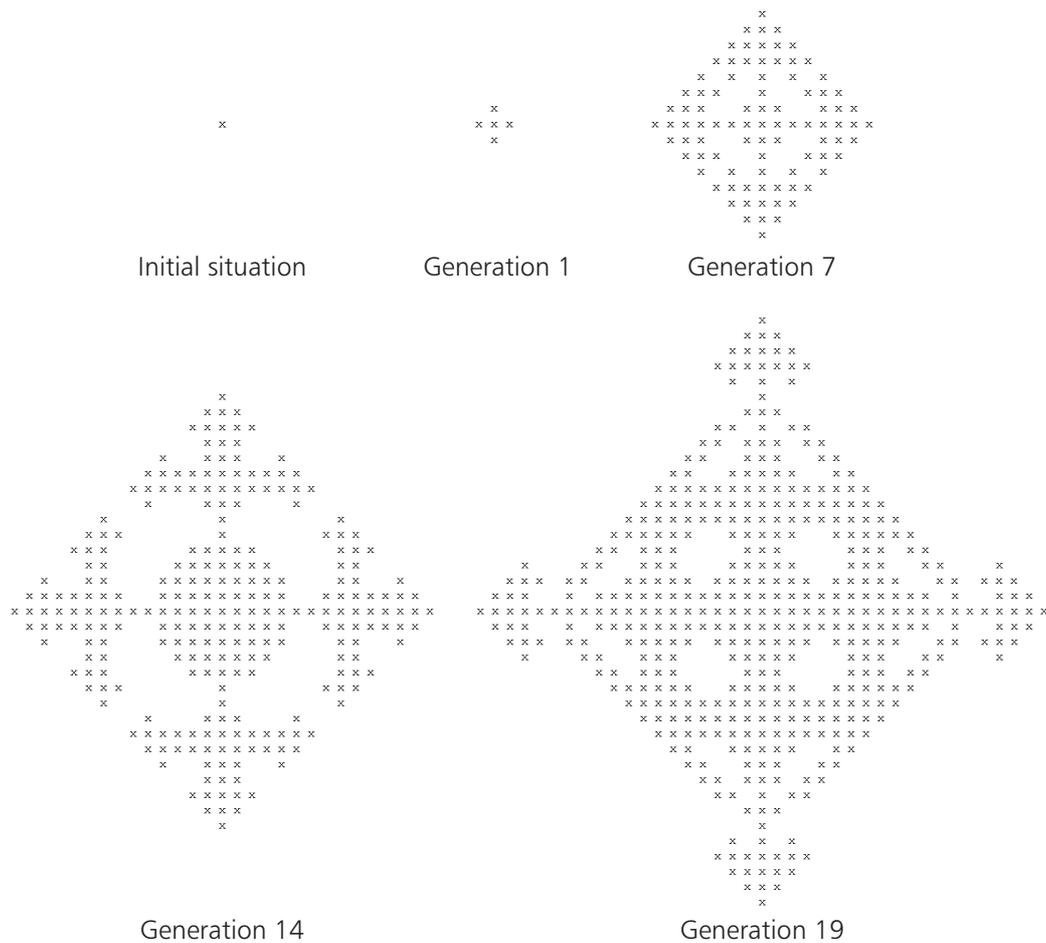


Figure 8.1 Axelrod's snowflakes. Crosses denote defectors, all of whom descend from a single defector who invaded a lattice populated with *Tit For Tat* players. *Source:* Axelrod (1984).

also on those of *their* nearest neighbors (i.e., of a 5×5 square, if the nearest neighbors are the eight cells of a chess king's move).

Let us consider, for instance, a configuration where the left half of the lattice is occupied by cooperators and the right half, by defectors (Figure 8.2). A cooperator sitting at the edge interacts with five cooperators and three defectors, and hence obtains $5R + 3S$ as the payoff, whereas a defector can exploit at most three cooperators and hence obtains the payoff $3T + 5P$. For many parameter values, cooperators earn more than defectors, and hence advance their front by one step to the right per generation. Similarly, if a 2×2 block of cooperators is surrounded by defectors (Figure 8.3), then a cooperator is boosted by the interaction with three other cooperators and obtains $3R + 5S$ points, whereas a defector, as an "outsider," can exploit at most two cooperators and earns $2T + 6P$ points. Again, it is possible for cooperators to spread. Of course, lone defectors will always do well. But by prospering, they surround themselves with defectors and diminish their

C	C	C	D	D	D
C	C	C	D	D	D
C	C	C	D	D	D
C	C	C	D	D	D
C	C	C	D	D	D
C	C	C	D	D	D

Figure 8.2 Phalanx of cooperation. Cooperators sit on the left half of the plane, defectors on the right half. If $5R + 3S > 3T + 5P$, the front of cooperators advances to the right.

own return. For some parameter values, a square of cooperators can expand at the corners but shrink along the sides: this process yields highly intricate patterns of growth, an endlessly milling spatio-temporal chaos with metastatic tentacles flailing at each other.

Nowak and May’s cellular automata can display most of the gadgetry of Conway’s Game of Life – periodic blinkers, for instance, or gliders consisting of teams of cooperators moving through a sea of defectors. If one starts with a symmetric initial condition, the emerging configuration retains this symmetry and can lead to wonderful kaleidoscopic patterns reminiscent of Persian carpets and Andalusian tiles.

The basic message of the Nowak–May model is that, for a substantial subset of the parameter space (i.e., the payoff values R , S , T , and P) and for most initial conditions, defectors and cooperators can coexist forever, either in static irregular patterns or in dynamic patterns with chaotic fluctuations along predictable long-term averages. In stark contrast to the case of well-mixed populations without spatial structure, where cooperators are bound to die out, cooperators frequently persist if the population has a territorial structure. In retrospect, this could have been guessed from Axelrod’s simulations leading to the growing snowflake of invading defectors and persisting islands of cooperators. Indeed, the probability of a further round in the repeated game was chosen by Axelrod to be as low as one-third, so that the average “repeated” game consists of only 1.5 rounds!

Most of the results of Nowak and May (1992, 1993) consist of numerical explorations of the model’s behavior for different payoff values. They consider the case $P = S$, which is the limiting case between the Prisoner’s Dilemma game and the Chicken (or Snowdrift) game, and normalize such that this common value is 0 and the value of R is 1 (such an affine linear normalization, which is standard practice in game theory, obviously does not

D	D	D	D	D	D
D	D	D	D	D	D
D	D	C	C	D	D
D	D	C	C	D	D
D	D	D	D	D	D
D	D	D	D	D	D

Figure 8.3 Island of cooperation. A 2×2 island of cooperators in a sea of defectors.

affect the structure and dynamics of the game). Thus the only remaining parameter is T , the temptation to defect. Starting with a random initial configuration, some values of T lead to an overall structure of defectors interspersed with tiny regions of cooperators (Figure 8.4). For other values of T , the lines of defectors become connected. For still larger values, chaotic regimes appear: the board is covered with defectors, but many small clusters of cooperators persist. These clusters continue to grow as long as they do not come too close to each other. Surprisingly, the average frequency of cooperators converges to about 32%, independent of the initial condition, as long as T has a value between 1.8 and 2.0 (Figure 8.5). Other neighborhood structures and different lattices – for instance, hexagonal lattices – lead to similar results.

A new twist was introduced by Huberman and Glance (1993), who stressed that if the cells on the lattice were updated, not *synchronously* as is usual in cellular automata, but *asynchronously* – picking up an individual cell at random and replacing the player sitting there with the highest-scoring player within the neighborhood – then the evolution leads to the all-out victory of defectors, rather than to the survival of a certain percentage of cooperators. Huberman and Glance claimed that because synchronous updating requires a global clock, it leads to a rather implausible model. They concluded that spatial games offer little hope for explaining the evolution of cooperation in biological communities. A similar argument was made by Mukherji *et al.* (1995).

It quickly became clear, however, that this effect of asynchrony was only valid for a rather limited part of the parameter space. Substantial but distinct regions of the parameter space lead to the persistence of cooperators with both synchronous and asynchronous updating.

In more detail, Nowak *et al.* (1994a, 1994b) added several elements of *stochasticity* to the spatial model. Asynchronous updating of randomly

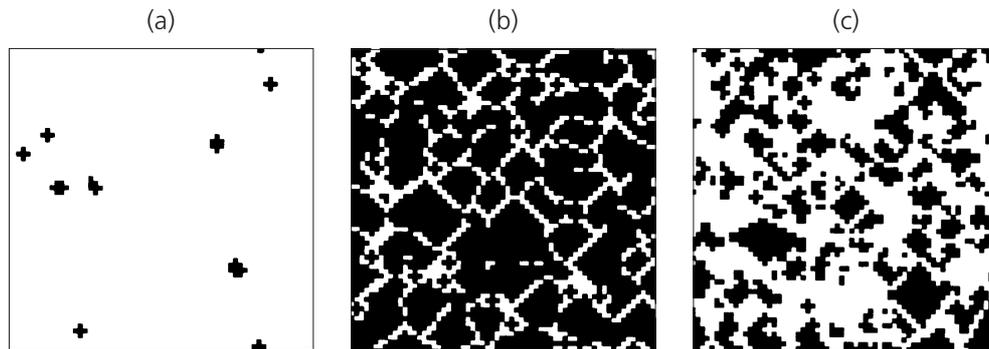


Figure 8.4 The signature of temptation. Different regimes for different values of T : (a) defectors interspersed with small islands of cooperators; (b) lines of defectors on a background of cooperators; (c) spatio-temporal chaos of defectors and cooperators. *Source*: Lindgren and Nordahl (1994).

chosen sites is just one such factor. Another biologically plausible assumption is that the scores obtained by the cells within the neighborhood do not fully determine which player will take over in the next generation (the one with the highest total payoff), but only specify the *probability* that a player will take over. [A somewhat related account of the role of stochasticity upon the updating rule can be found in Mukherji *et al.* (1995).] More generally, if A_i denotes the score of the player at site i , and s_i is 0 if site i is occupied by a defector and 1 if occupied by a cooperator, then the probability that site j will be occupied by a cooperator in the next generation is given by $(\sum A_i^m)^{-1} \sum A_i^m s_i$, where the sum is extended over all neighboring sites i of site j , and where m is a positive real number characterizing the degree of stochasticity of the takeover mechanism. If $m = 1$, the chances of colonizing the site are proportional to the score. The limiting case $m \rightarrow \infty$ gives the original deterministic rule: the site goes with certainty to whoever achieved the highest score. For $m = 0$, on the other hand, we have random drift: the score no longer plays a role, all players in the neighborhood have an equal chance of taking over. Figure 8.6 shows the outcome for the simulations with $0 \leq m \leq \infty$ and $1 < T < 2$ (and, as before, $R = 1$ and $S = P = 0$). It is interesting to compare the effect of synchronous and asynchronous updating in this wide setting of parameter values. As noted by Huberman and Glance (1993), the coexistence of C and D for large values of T holds (for $m \rightarrow \infty$) only in the synchronous case. But with probabilistic updating (i.e., $m = 1$), it actually holds for *more* T values in the asynchronous case. It must be stressed, however, that for $m = 1$ the region where cooperation persists is rather small. The advantage of belonging to a patch of cooperators cannot be efficiently “exported” when $m = 1$, a fact that has also been noted by Wilson *et al.* (1992) in a related model.

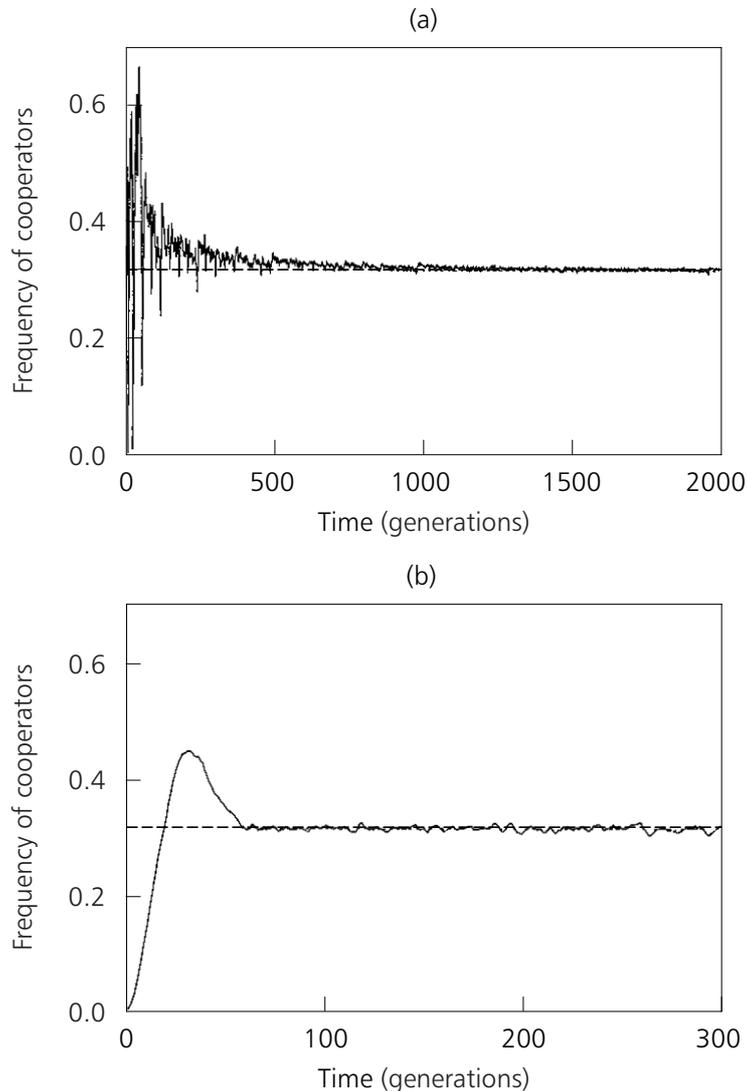


Figure 8.5 Average cooperation. For many initial conditions, the frequency of cooperators converges to $0.3178\dots$ for $1.8 < T < 2.0$ (players interact with their eight neighbors and with themselves); (a) starts with a single defector, (b) starts with a random mixture of cooperators and defectors. *Source:* Nowak and May (1993).

One can randomize the game still further by introducing random neighborhood grids, or random dispersal. These changes lead to no essential alterations of the picture.

Furthermore, Nowak *et al.* (1994b) have established that cooperation can easily be maintained if there exists a certain probability that cells in the neighborhood of individuals with low payoff remain empty. In this case, defectors tend to subvert their own livelihood. Cooperators can then invade the barren patches. This holds even for the unfavorable case $m = 1$, and for very large temptation values T .

Several other results are less intuitive. For certain parameter values it may happen that cooperators vanish despite consistently having a higher

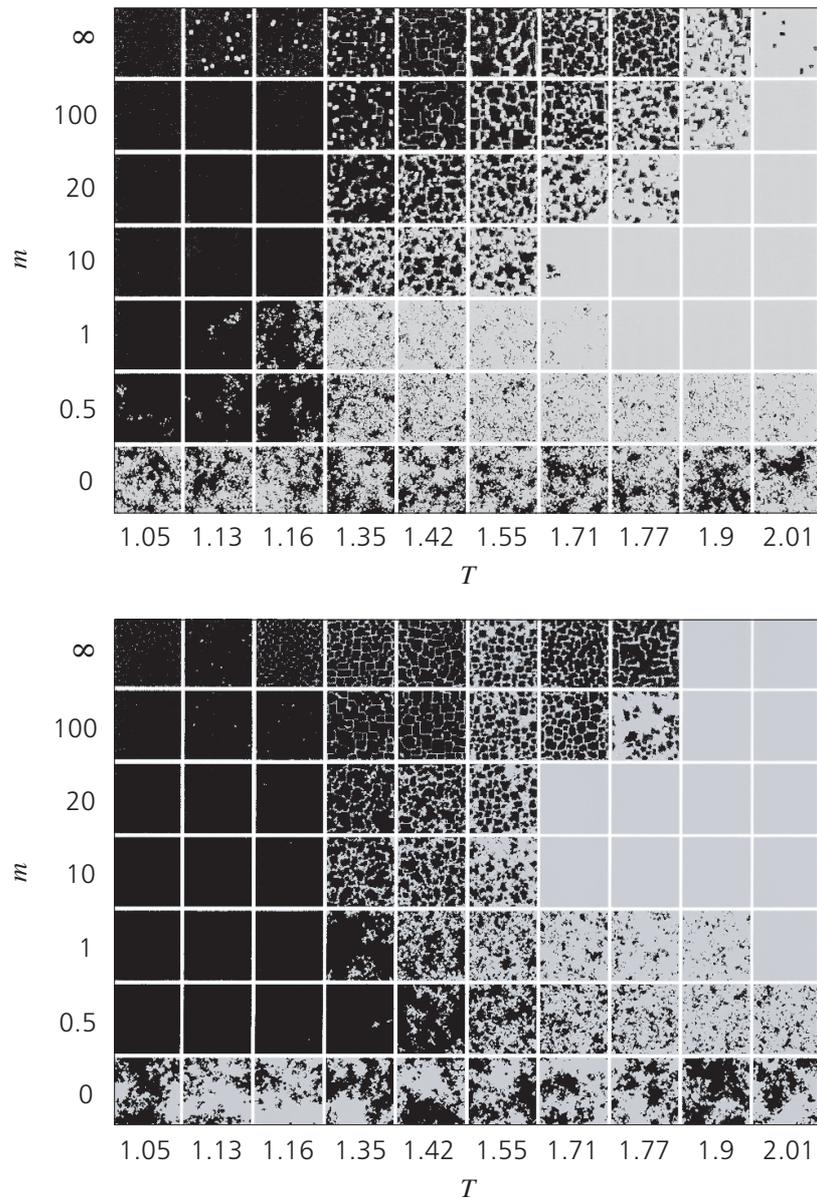


Figure 8.6 Synchronous versus asynchronous updating. Black indicates cooperation; gray, defection. The squares give the outcome of spatial simulations for different values of m and T (see text). The upper figure is based on synchronous updating; the lower figure, on asynchronous updating. *Source:* Nowak *et al.* (1994a).

average payoff than the defectors. The reason is that the few defectors close to cooperators have a higher score and can spread. The low payoff of defectors surrounded by defectors and the high payoff of cooperators surrounded by cooperators do not matter; what counts is the interface where cooperators and defectors are in close proximity.

An even more surprising result of the stochastic game with $m = 1$ is that populations of cooperators can become extinct even though isolated cooperator cells can have more than one offspring on average (i.e., although

their basic reproductive ratio is greater than 1). This extinction occurs when larger clusters of cooperators cannot grow significantly. In this case, random fluctuations may wipe them out in one blow.

Killingback and Doebeli (1996) extended this type of analysis to cover the Chicken game (or Hawk–Dove game). As mentioned earlier, this game has precisely the same structure as the Snowdrift game, but now D is interpreted as the behavioral rule “escalate a conflict until one of the contestants is harmed” and C means “stick to displaying, that is, keep the conflict harmless.” Killingback and Doebeli show that, in general, the long-term proportion of Hawks (or “defectors,” in the context of the Snowdrift game) is smaller than the equilibrium proportion predicted by classical evolutionary game theory. In addition, they observed a type of complex dynamics that is different from the spatio-temporal chaos seen in the spatial Prisoner’s Dilemma of Nowak and May (1992). For a substantial range of parameters, their system organizes itself into a critical state in which its dynamic behavior is governed by long-range spatial and temporal correlations and by power laws (Killingback and Doebeli 1998). An interesting property of such critical systems is the existence of extremely long transients on which the dynamics are initially very complicated but then relax into a simple periodic orbit. They conjectured that suitable spatial extensions of any evolutionary game with a mixed evolutionarily stable strategy will exhibit critical dynamics for appropriate parameters. Killingback and Doebeli (1996) also extended the game to include conditional strategies like *Retaliator* (start by displaying, but escalate the conflict if the opponent escalates) and *Bully* (escalate the conflict, but retreat if the adversary also escalates). The basic result here is that *Retaliator* is much more successful with territoriality than without. Usually, the evolutionary outcome is a population of *Retaliators*, occasionally interspersed with a few *Doves*. This is quite different from the evolutionary outcome without territoriality.

Another one-round game with more than two strategies is the Rock–Scissors–Paper game (a cyclic arrangement of three strategies, each dominating its successor). A spatial version of this game was briefly studied by Nowak *et al.* (1994b), who showed that it may lead to spiral waves with the three strategies endlessly pursuing each other.

Feldman and Nagel (1993) investigated a variant of the spatial Prisoner’s Dilemma where the players are not updated in every round but can accumulate payoff; against this, they must pay a certain fee (the same for all players) to stay in the game. If their savings are eaten up, they are replaced by a wealthier neighbor. The authors observed widespread local coordination on the lattice.

8.3 Repeated Games

The main message so far is that neighborhood structure seems to offer a promising way out of the Prisoner's Dilemma toward the emergence of cooperation. There are many alternative explanations of the prevalence of cooperation, but, arguably, none require less sophistication on the part of the individual agents than those with spatial structure. The latter need no foresight, no memory, and no family structure. Viscosity suffices.

As soon as one considers more highly developed agents who are long-lived enough to interact repeatedly, have enough memory to keep track of their past, and are smart enough to adapt to their co-players' moves, cooperation becomes much easier to sustain, even in the well-mixed case. A retaliatory strategy like *TFT* or *Grim* (the strategy of cooperating until the adversary defects for the first time, and from then on relentlessly defecting) can invade a population of *AD* players provided its initial frequency exceeds a certain low threshold. Further evolution can lead to populations where cooperation is more robustly established and is proof against occasional mistakes. One frequent outcome is *Pavlov*, the strategy that starts with a cooperative move and then cooperates if and only if in the previous round both players used the same move [both playing *C* or both playing *D*, see Nowak and Sigmund (1993)]. *Pavlov* is a strategy based on the win-stay, lose-shift principle: a player repeats his or her previous move if the payoff was large (*T* or *R*) and tries the other move if the payoff was small (*S* or *P*).

It turns out that in models incorporating repeated interactions *and* territoriality, the probability for a cooperative outcome is high indeed. This property was firmly established by the extensive computer simulations of Lindgren and Nordahl (1994). What was less predictable was which strategies would dominate the cooperative regime.

Lindgren and Nordahl (1994) normalized the payoff values to $R = 1$ and $S = 0$, with $1 < T < 2$ and $0 < P < 1$ for the Prisoner's Dilemma. They investigated the infinitely repeated Prisoner's Dilemma game, where the probability for a further round is equal to 1. Like Nowak and Sigmund (1993), they also assumed that players occasionally make mistakes; in their simulations, one move in a hundred was mis-implemented. This means that the initial move is no longer relevant: it only affects the beginning of the interaction, which plays no role in the long run because, in the limit, the overall payoff is now the mean of the payoff per round.

Lindgren and Nordahl (1994) started by investigating strategies with memory spans of 0. This corresponds to unconditional *C* or *D* strategies for the one-shot Prisoner's Dilemma. Using a four-cell neighborhood, they

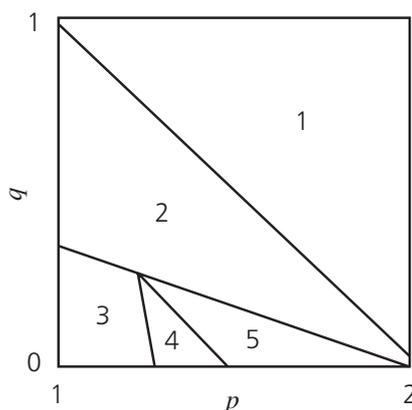


Figure 8.7 Lest we forget. Asymptotic behavior of memory-0 strategies: (1) leads to a homogeneous state of defectors, (2) to small domains of cooperators (3) to percolation networks of defectors in a background of cooperators, (4) to spatio-temporal chaos, and (5) to rigid lines of cooperators; here, $p = \frac{T-S}{R-S}$ and $q = \frac{P-S}{R-S}$. Source: Lindgren and Nordahl (1994).

showed that the evolutionary outcome could be as encapsulated in Figure 8.7. For $T + P > 2$, defectors take over. The remaining parameter values lead to a coexistence of defectors and cooperators either in a stable form (islands of cooperators or percolation networks of defectors) or in the form of spatio-temporal chaos.

Next, Lindgren and Nordahl considered strategies that take into account the *co-player's* previous move – *TFT* is one such strategy, *Anti-TFTat* (which does the opposite of what the opponent did in the previous round and, not surprisingly, never establishes a firm foothold) is another. The remaining two strategies are the unconditional *AD* and *Always Cooperate (AC)*. It turns out that only a tiny corner of the parameter space – far less than one-thousandth of it – leads to the emergence of a homogeneous *TFT* population (indeed, paired *TFT* players suffer severely from errors, as they engage in long runs of alternating defections). A much larger part of the parameter space leads to the dominance of the *AC* strategy. Other parameter regions lead to spatio-temporal chaos or spiral waves, still others to frozen configurations of *AD* and *TFT*.

By enhancing the sophistication of the players so that they can base their next moves on both their co-player's previous move and *their own* previous move (i.e., on the full outcome of the previous round), one obtains 16 strategies that can interact in rather complex ways (see also Nowak *et al.* 1995b). More than half the parameter space leads to a cooperative regime dominated by the *Pavlov* strategy (Figure 8.8). The remaining parameter values lead either to an outcome dominated by a parasitic strategy like *AD* or *Bully* (the strategy of always defecting except after having been punished in the previous round), or to complex behavior with high diversity.

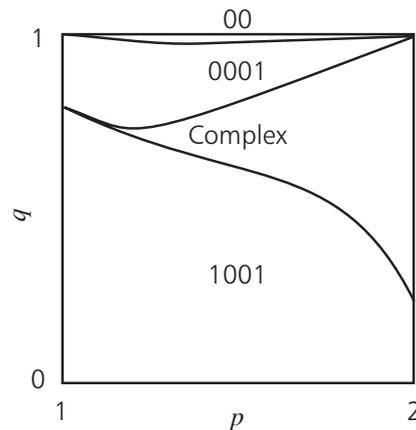


Figure 8.8 *Pavlov's* domain. Asymptotic behavior of strategies based on the outcome of the last round. In the region marked 1001, *Pavlov* prevails. In the regions marked 00 and 0001, a cooperative regime is established; p and q are as in Figure 8.7. *Source*: Lindgren and Nordahl (1994).

One can, of course, consider strategies with even longer memories, based, for instance, on the outcome of the previous *two* rounds. In fact, Lindgren (1991) has developed a genetic algorithm where mutations providing for ever-larger memories can be introduced and eventually selected. If there is no upper bound to the memory length, the evolutionary simulation can lead to extremely diverse patterns consisting of many strategies coexisting in highly entangled spatio-temporal arrangements. Many of the resulting strategies share with *Pavlov* the property that any error leads to a short period of mutual punishment after which, as of a common accord, the players simultaneously return to cooperation. Lindgren and Nordahl (1994) stress that asynchronous updating has little effect on the simulations (although it excludes spiral waves). They argue convincingly that one cannot speak of *the* Prisoner's Dilemma game: for different parameter values one can obtain very different strategic interactions.

At present, there seems little hope for obtaining a full analytic understanding of these phenomena. As with cellular automata in general, most of the results are based on computer simulations and are at best a very crude intuition of some local effects. An interesting attempt at an analytic treatment was made by Nakamaru *et al.* (1997), who considered the interplay of *TFT* with *AD*, as in Axelrod's book (Axelrod 1984). In contrast to Axelrod and the other authors mentioned so far, they assume that the fitness gained by each player (i.e., the total payoff obtained against all neighbors) translates, not into a *transmission rate* (the probability of invading a neighboring cell), but into a *mortality rate*. This assumption yields a variant of asynchronous updating where the probability that a cell is updated depends on the success of the player sitting in that cell, whereas each neighbor has the

same chance of taking over the vacated site. Such a setup offers some scope for spiteful behavior. By withholding cooperation, one decreases the survival rate of a neighbor. This yields a chance of inheriting the neighbor's site – although at a cost to oneself.

Nakamaru *et al.* (1997) considered one-dimensional lattices (where each cell has two neighbors) and two-dimensional lattices (with eight neighbors per cell). The outcome of the spatial simulations is very different from both that of the complete mixing model (i.e., the dynamics without spatial structure) and that of a first, crude approximation called the *mean-field approximation*. But another approximation, the *pair approximation* (Chapters 13 and 18), based on neglecting correlations except among nearest neighbors, frequently leads to analytical results in good agreement with the numerical simulations. In particular, both computer simulation and pair approximation show that *TFT* can invade a population of defectors if the probability of a further round is sufficiently high. Interestingly, the model of Nakamura *et al.* (1997) never leads to the coexistence of cooperators and defectors, which is a frequent outcome in the Nowak–May (1992) model.

Another analytic approach has been investigated by Eshel *et al.* (unpublished). These authors consider players on a one-dimensional lattice. Each player interacts with the k nearest neighbors to the left and to the right. There are two available strategies. After each generation, a player either sticks to his or her own strategy or switches to that of one of the n players to the left or to the right, with probabilities proportional to the neighbors' scores. However, a player can switch strategies only if at least one of the two immediate neighbors uses the alternative strategy. This model is based on an interaction neighborhood and a (possibly different) propagation neighborhood. Eshel *et al.* (unpublished) give conditions on the stability and invasibility of strategies that, for large values of k and n , only depend on the value of k/n . This helps to interpret the concept of population viscosity in terms of inclusive fitness (see Hamilton 1964). In particular, cooperation wins in the Prisoner's Dilemma game if each player observes a large number of players before imitating one.

A similar distinction between one-dimensional propagation neighborhoods and interaction neighborhoods (arranged on the circle) is made by Hoffman and Waring (1996). They studied all deterministic memory-1 strategies for the repeated Prisoner's Dilemma using numerical simulations and found that localization of learning always fosters cooperative behavior, but that the effect of localization of the interaction is not clear-cut and depends on details of how the competing strategies are clustered.

Other papers dealing with the spatial version of the repeated Prisoner's Dilemma include Mar and St. Denis (1994, and unpublished), Grim (1995),

and Kis (unpublished). Mar and St. Denis (unpublished) actually studied a more general variant, where each player can opt for a larger or smaller degree of cooperation. Kirchkamp (unpublished) presents particularly extensive simulations exploring a wide range of learning rules. In particular, he also allows for stochastic timing, not just of the updating, but of the interaction as well.

8.4 Extensions and Related Work

The spatial games of Nowak and May (1992, 1993) emerged at the same time as some models in theoretical ecology and prebiotic evolution, all of which stress the same message: territoriality favors diversity. This point is discussed further in Chapters 9 and 10 in this volume. Durrett and Levin (1994b) provide a general survey of spatial aspects.

Herz (1994) analytically classifies the behavior of 2×2 games in terms of their payoff structure. The spatial versions of some of these games lead to an uphill march in an abstract fitness landscape, whereas other games become bogged down due to a phenomenon called “frustration” in the physics of Ising-type models. The updating rule used by Herz is not evolutionary in the proper sense (neither the transmission rate nor the mortality of a player is affected by the score). However, it is based on a *Pavlov*-like principle that can be interpreted as a learning mechanism: if the player obtains a total score against his or her neighbors that is above a certain threshold, the player uses the same strategy in the following round; if the score is below that threshold, he or she switches to the alternative strategy. The simplicity of this updating rule provides the basis for a theoretical analysis of the resulting cellular automaton.

Another, very different approach to the spatial modeling of the Prisoner’s Dilemma game can be found in the one-dimensional reaction–diffusion models of Hutson and Vickers (1995) and Ferrière and Michod (1995), which lead to traveling waves of *TFT* players swamping the inveterate defectors.

The spatial, iterated Prisoner’s Dilemma has mainly been used to study the effects of neighborhood on the evolution of cooperation between members of the same species. Recently, Doebeli and Knowlton (1998) have suggested an extension in order to study mutualistic interactions between members of different species. In their model, two different species occupy two different, superimposed lattices. Interspecific interactions occur between individuals at corresponding sites in these lattices, and the payoffs from these interactions determine the outcome of competition within each lattice. Doebeli and Knowlton show that in this setting, spatial structure

is even more essential for the evolution of mutualism than in one-lattice games. (In fact, they considered versions of the iterated Prisoner's Dilemma in which investment could vary continuously and always evolved to 0 in the non-spatial games.)

Finally, we mention that lattice models of artificial societies have a long tradition, as can be seen from the work of Schelling (1971) and Sakoda (1971), who used them to simulate migration and segregation behavior based on neighborhood rules. Some recent, related investigations have been done by Hegselmann (1996), for instance, who analyzed a *solidarity game*, which is an extension of the Prisoner's Dilemma, and Epstein and Axtell (1996), who studied highly elaborate artificial societies.

8.5 Concluding Comments

A wealth of computer simulations of artificial populations now exist showing that cooperation can be stably sustained in societies of simple automata. The introduction of spatial structure has shown that cooperation becomes even more likely if interactions are restricted to neighbors. The venerable rule of cooperating with neighbors is certainly not new. But placed alongside other spatial models, it offers wide perspectives for the emergence and stability of cooperative societies. In particular, it shows that even if the interaction between two individuals is not repeated, cooperation can be sustained in the long run.

This positive message of cyber-sociology should not obscure another, darker aspect. In territorial structures, it is easily possible to distinguish between us (me and my neighbors) and them (the foreigners). In fact, most of the challenges experienced by territorial bands of hominids might have been caused by other bands. Cooperation within groups was crucial for dealing with conflicts between groups. It is quite likely that the fiercest of these conflicts were always fights for territory.

References

References in the book in which this chapter is published are integrated in a single list, which appears on pp. 517–552. For the purpose of this reprint, references cited in the chapter have been assembled below.

- Axelrod R & Hamilton WD (1981). The evolution of cooperation. *Science* **211**:1390–1396
- Axelrod R (1984). *The Evolution of Cooperation*. New York, NY, USA: Basic Books
- Doebeli M & Knowlton N (1998). The evolution of interspecific mutualism. *Proceedings of the National Academy of Sciences of the USA*, **95**:8676–8680
- Durrett R & Levin S (1994b). The importance of being discrete (and spatial). *Theoretical Population Biology* **46**:363–394
- Epstein JM & Axtell R (1996). *Growing Artificial Societies*. Boston, MA, USA: The Brookings Institution Press
- Eshel I, Sansone E & Shaked A. Evolutionary dynamics of populations with a local interaction structure. Sonderforschungsbereich Bonn. Unpublished
- Feldman B & Nagel K (1993). Lattice games with strategic takeover. In *Lectures in Complex Systems*, SFI Studies in the Sciences of Complexity, eds. Nadel L & Stein D, Redwood City, CA, USA: Addison-Wesley
- Ferrière R & Michod R (1996). The evolution of cooperation in spatially heterogeneous populations. *The American Naturalist* **147**:692–717
- Hamilton WD (1964). The genetic evolution of social behavior. *Journal of Theoretical Biology* **7**:1–52
- Hegselmann R (1996). Solidarität unter Ungleichen. In *Modelle sozialer Dynamiken*, eds. Hegselmann R & Peitgen HO, pp. 105–128. Vienna, Austria: Hölder-Pichler-Tempski
- Herz AVM (1994). Collective phenomena in spatially extended evolutionary games. *Journal of Theoretical Biology* **169**:65–87
- Hoffman R & Waring N (1996). The simulation of localised interaction and learning in artificial adaptive agents. In *Evolutionary Computing*, ed. Fogarty TC, New York, NY, USA: Springer
- Huberman BA & Glance NS (1993). Evolutionary games and computer simulations. *Proceedings of the National Academy of Sciences of the USA* **90**: 7716–7718
- Hutson V & Vickers GT (1995). The spatial struggle of *Tit-For-Tat* and defect. *Philosophical Transactions of the Royal Society of London B* **348**:393–404
- Killingback T & Doebeli M (1996). Spatial evolutionary game theory: Hawks and Doves revisited. *Proceedings of the Royal Society of London B* **263**:1135–1144
- Killingback T & Doebeli M (1998). Self-organized criticality in spatial evolutionary game theory. *Journal of Theoretical Biology* **191**:335–340
- Kirchkamp O. Spatial evolution of automata in the Prisoner's Dilemma game. Unpublished
- Kis T. Predictable properties of the spatial iterated Prisoner's Dilemma game. Unpublished

- Lindgren K (1991). Evolutionary phenomena in simple dynamics. In *Artificial Life II*, eds. Langton CG *et al.*, pp. 295–312. Redwood City, CA, USA: Addison-Wesley
- Lindgren K & Nordahl MG (1994). Evolutionary dynamics of spatial games. *Physica D* **75**:292–309
- Mar G & St Denis P (1994). Chaos in cooperation: Continuous-valued Prisoner's Dilemmas in infinite-valued logic. *International Journal of Bifurcation and Chaos* **4**:943–958
- Mar G & St Denis P. Chaos in cooperation: Two-dimensional Prisoner's Dilemmas in infinite-valued logic. Unpublished
- Mukherji A, Rajan V & Slagle JR (1995). Robustness of cooperation. *Nature* **379**:125–126
- Nakamaru M, Matsuda H & Iwasa Y (1997). The evolution of cooperation in a lattice-structured population. *Journal of Theoretical Biology* **184**:65–81
- Nowak MA & May RM (1992). Evolutionary games and spatial chaos. *Nature* **359**:826–829
- Nowak MA & May RM (1993). The spatial dilemmas of evolution. *International Journal of Bifurcation and Chaos* **3**:35–78
- Nowak MA & Sigmund K (1993). A strategy of win–stay, lose–shift that outperforms Tit-for-Tat in the Prisoner's Dilemma game. *Nature* **364**:56–58
- Nowak MA, Bonhoeffer S & May RM (1994a). Spatial games and the maintenance of cooperation. *Proceedings of the National Academy of Sciences of the USA* **91**:4877–4881
- Nowak MA, Bonhoeffer S & May RM (1994b). More spatial games. *International Journal of Bifurcation and Chaos* **4**:33–56
- Nowak MA, May RM & Sigmund K (1995a). The arithmetics of mutual help. *Scientific American* **272**:76–81
- Nowak MA, Sigmund K & El-Sedy E (1995b). Automata, repeated games, and noise. *Journal of Mathematical Biology* **33**:703–732
- Sakoda JM (1971). The checkerboard model of social interaction. *Journal of Mathematical Sociology* **1**:119–132
- Schelling TC (1971). Dynamic models of segregation. *Journal of Mathematical Sociology* **1**:143–186
- Sigmund K (1992). On prisoners and cells. *Nature* **359**:774
- Sigmund K (1995). *Games of Life*. Harmondsworth, UK: Penguin
- Sugden R (1986). *The Economics of Co-operation, Rights and Welfare*. New York, NY, USA: Blackwell
- Trivers R (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology* **46**:35–57
- Wilson DS, Pollock GB & Dugatkin LA (1992). Can altruism evolve in purely viscous populations? *Evolutionary Ecology* **6**:331–341