



Evolution of Cooperation in Spatially Structured Populations

KURT BRAUCHLI*, TIMOTHY KILLINGBACK AND MICHAEL DOEBELI

Zoology Institute, University of Basel, Rheinsprung 9, CH-4051 Basel, Switzerland

(Received on 9 April 1999, Accepted in revised form on 13 July 1999)

Using a spatial lattice model of the Iterated Prisoner's Dilemma we studied the evolution of cooperation within the strategy space of all stochastic strategies with a memory of one round. Comparing the spatial model with a randomly mixed model showed that (1) there is more cooperative behaviour in a spatially structured population, (2) PAVLOV and generous variants of it are very successful strategies in the spatial context and (3) in spatially structured populations evolution is much less chaotic than in unstructured populations. In spatially structured populations, generous variants of PAVLOV are found to be very successful strategies in playing the Iterated Prisoner's Dilemma. The main weakness of PAVLOV is that it is exploitable by defective strategies. In a spatial context this disadvantage is much less important than the good error correction of PAVLOV, and especially of generous PAVLOV, because in a spatially structured population successful strategies always build clusters.

© 1999 Academic Press

1. Introduction

The Prisoner's Dilemma is a well-known metaphor for the evolution of cooperation among selfish individuals. It elegantly illustrates the problems of achieving cooperation in any animal society. The Prisoner's Dilemma is a game played by two players who can either cooperate (C) or defect (D). If both players cooperate they are rewarded with R points, if they both defect they get the punishment P . If one player defects while the other cooperates, the defector gets the temptation payoff T , while the other gets the sucker's payoff S (cf. Table 1).

Now, with $T > R > P > S$ and $2R > T + S$ we have an obvious dilemma: if a player plays once against an unknown opponent, it always pays more to defect, whatever the opponent does: if the opponent plays C then $T > R$ and if the opponent plays D then $P > S$. Thus, individuals

that play D will always be favoured by natural selection, even though the average payoff in a population of defectors is less than in a population of cooperators ($P < R$). In terms of evolutionary game theory (cf. Maynard Smith, 1982) D is the unique evolutionary stable strategy (ESS).

Two different types of solutions to this dilemma have been proposed. First, the *Iterated Prisoner's Dilemma* (IPD) was made popular by work of Axelrod (1980, 1981, 1984) and Axelrod & Hamilton (1981). In the IPD there is a probability $w > 0$ that two players will meet again, and the players can remember how they played in previous encounters. This allows more complicated strategies than the simple C or D of a non-iterated Prisoner's Dilemma. Axelrod conducted a computer tournament with 62 strategies submitted by scientists from all around the world (Axelrod, 1980). He concluded from this that successful strategies in the IPD are (1) nice: they never defect first, (2) forgiving: they restore cooperation after an accidental defection

* Author to whom correspondence should be addressed.
E-mail: brauchlik@ubaclu.unibas.ch

TABLE 1
Payoff matrix of the Prisoner's Dilemma game. The game is defined by $P > R > P > S$. In our simulations we use $R = 3$, $T = 5$, $P = 1$ and $S = 0$

	C	D
C	R	S
D	T	P

and (3) retaliatory: they react by playing D if their opponent starts playing D. The most successful strategy in Axelrod's tournament was the very simple TIT FOR TAT (TFT), which starts by playing C and then plays whatever its opponent played in the previous round. Nowak & Sigmund (1993) found another very successful strategy for the IPD, PAVLOV, a "win stay-lose shift" strategy, which will be discussed in detail later.

A second situation where cooperative behaviour can be evolutionary stable are models that introduce some kind of spatial structure, so that within a population all interactions are local: i.e. individuals play against their nearest neighbours and *not* against random opponents or against all other individuals of the population. Nowak & May (1992) have shown that such spatial structure enables the maintenance of cooperation for the simple Prisoner's Dilemma, in contrast to the classical, spatially unstructured Prisoner's Dilemma where defection is always favoured. Similarly, it was found for the Hawk-Dove game that in spatially structured populations there are more cooperative individuals (doves) than in non-structured populations (Killingback & Doebeli, 1996). The main effect of spatial structure is that in structured populations cooperative strategies can build clusters in which the benefits of mutual cooperation can outweigh losses against defectors. Thus, clusters of cooperative strategies can invade into populations of defectors that constitute an ESS in non-spatial populations (Axelrod & Hamilton, 1981; Nowak & May, 1992; Killingback & Doebeli, 1996).

An interesting combination of these two ideas, iterated games and spatial structure, was studied by Lindgren & Nordahl (1994). Not only did they

find that in structured populations there is more cooperation, but also that different strategies are successful in spatially structured populations than in non-structured populations. Lindgren & Nordahl studied strategies which have a memory of 0 to n rounds and always deterministically play either C or D, depending on the outcome of last 0 to n rounds. Through mutations—such as for example gene duplication—the memory length of the strategies could be extended indefinitely. However, one limitation of Lindgren & Nordahl's set of strategy is that no probabilistic strategies are possible. The importance of probabilistic strategies is demonstrated by the work of Nowak & Sigmund (1992) and by Grim's (1995) work on the spatial IPD, who considered stochastic strategies which depend only on the opponent's last move. For example, Grim (1995) found that in spatially structured populations a very generous version of TFT, that responds to defection with a retaliating defection only with a probability of 33%, was the most successful strategy of a fixed set of 121 strategies he considered. Among other things, we will show in this paper that generous variants of PAVLOV are even more successful than generous TFT in spatially structured populations.

We will combine the IPD and spatial structure in a more realistic, stochastic model in which individuals interact only with their neighbours and where new randomly mutated strategies can appear at any time. In contrast to Lindgren & Nordahl (1994), we will examine a different and in a sense more general set of strategies: we will allow all possible stochastic strategies with a memory of one round. We will refer to such strategies as Markov strategies. Instead of playing either C or D with probability 1, these strategies can play C or D—depending on the previous round—with *any* probability between 0 and 1. Markov strategies are defined by the four probabilities (p_R, p_T, p_S, p_P) that an individual plays C after it scored R, T, P, or S, respectively, in the previous round. In populations of such Markov strategies, in which interactions between individuals are of the form of an IPD, we will examine how spatial structure influences the evolution of cooperation and what impact it has on the evolutionary dynamics. By evolutionary dynamics we are here referring to the movement of

the population average in the four-dimensional strategy space (p_R, p_T, p_S, p_P).

In the next section, we will give a detailed overview of the model used. The results of our simulations are divided into three sections. We will first look at spatial structure and the evolution of cooperative behaviour in general, then concentrate on the role of generosity and finally examine the evolutionary dynamics in detail. In the final section, the results are discussed.

2. The Model

The model examined in this paper consists of a square lattice of size $n \times n$ with periodic boundaries, i.e. the marginal cells on one side of the lattice are considered to be the neighbours of the cells on the opposite margin of the lattice. Every cell of the lattice is occupied by an individual with a particular Markov strategy for playing the IPD.

The strategy set is taken to be all stochastic strategies with a memory of one round. This corresponds to the set that Nowak & Sigmund (1993) used in a non-spatial setting. These strategies are given by the four probabilities (p_R, p_T, p_S, p_P) that a player cooperates in round $n + 1$, depending on the outcome of round n : p_R, p_T, p_S, p_P are the conditional probabilities to cooperate after receiving the payoff R, T, S or P , respectively, in the previous round. For instance, the winner of Axelrod's computer tournament, TFT, which always plays C after it scored R or T and D after S or P , corresponds to the strategy given by (1, 1, 0, 0). PAVLOV, another strategy that is very successful in the IPD (Nowak & Sigmund, 1993), is given by (1, 0, 0, 1): PAVLOV plays "win stay-lose shift", it repeats its move if it

scored well (R or T) but changes after it scored badly (S or P).

To make it easier to follow the evolutionary dynamics, we will use the names of the four well-known strategies TFT (1, 1, 0, 0), generous TFT (GTFT) (1, 1, 0.33, 0), PAVLOV (1, 0, 0, 1) and RETALIATOR (1, 0, 0, 0) not only strictly for the single strategy they stand for, but also for a small subset of strategies exhibiting the typical characteristics of the original strategy. Table 2 shows precisely how these subsets of strategies are defined in our context. The values of p_i used to define the subsets of strategies in Table 2 are chosen according to common sense understanding of how the included strategies should behave. Any reasonable alternative choice would not yield qualitatively different results. For example, the whole group of strategies that have values of $p_R > 0.9$, $p_T > 0.8$, $p_S < 0.3$ and $p_P < 0.4$ will be called TFT, because they show the typical characteristics of the original TFT which is strictly only the strategy (1, 1, 0, 0). Thus, if we are talking of PAVLOV, we usually mean the whole group that exhibit typical Pavlovian characteristics. The main difference between PAVLOV and TFT is the value of p_T (cf. Wedekind & Milinski, 1996). While both strategies cooperate with each other, PAVLOV readily exploits unconditional cooperators, but is more heavily exploited by unconditional defectors. In contrast to Wedekind & Milinski (1996), who treated all strategies with high p_R and low p_T as PAVLOV, we divide the group of cooperative strategies with low p_T into PAVLOV with $p_P > 0.6$ and RETALIATOR with $p_P < 0.4$.

To eliminate unrealistic, completely error-free behaviour there is a minimal noise level given

TABLE 2

Interpretation of the p_i -values as strategies. The values of p_i are the probabilities that a player cooperates after it did get a payoff of R, T, S or P respectively in the previous round of an iterated game. See text for the interpretation of the p_i as the strategies named in the left column

Strategy	p_R	p_T	p_S	p_P
Cooperative (COOP)	$p_R > 0.9$			
Retaliator (RET)	$p_R > 0.9$	$p_T < 0.4$	$p_S < 0.4$	$p_P < 0.4$
Tit for Tat (TFT)	$p_R > 0.9$	$p_T > 0.8$	$p_S < 0.3$	$p_P < 0.4$
Generous Tit for Tat (GTFT)	$p_R > 0.9$	$p_T > 0.8$	$p_S > 0.3$	$p_P < 0.4$
Pavlov (PAV)	$p_R > 0.9$	$p_T < 0.4$	$p_S < 0.4$	$p_P > 0.6$

by a small $\varepsilon > 0$ and the values of p_R , p_T , p_S and p_P are restricted to $(\varepsilon, 1 - \varepsilon)$. The number of iterations per game is given by the probability of iteration w . If one round of the simple Prisoner's Dilemma has been played by two individuals, there is a chance w that these two individuals will meet again and the game is iterated for another round. The expected number of iterations for a probability w is $1/(1 - w)$. It is obvious that in such a model a key characteristic of a cooperative strategy is that it will cooperate in round $n + 1$ after mutual cooperation in round n . Cooperative strategies therefore have a value of p_R close to 1.

To model spatial structure, all individuals of the $n \times n$ lattice play an IPD game against their eight nearest neighbours. The payoffs of these eight games are summed and provide the individuals' score or fitness. After every individual has played its eight games, each cell of the lattice is updated according to the general rules of spatial evolutionary game theory (as formulated in Killingback & Doebeli, 1996): the individual on each cell is replaced by an offspring of the highest scoring individual among the former site holder and its eight nearest neighbours. These offspring play the same strategy as their ancestors, except if a mutation occurs which happens at a small mutation rate. If a mutation occurs, the offspring's strategy is not its parent's but a new strategy chosen randomly from the whole strategy space.

In order to identify the effects of spatial pattern formation and self-organized behaviour, the spatially structured populations will be compared to well-mixed populations which are again modelled by a lattice, but on which each individual plays against eight partners that are randomly chosen on the lattice, rather than against its eight nearest neighbours. Likewise the successor of every site is chosen among the site holder and eight random individuals. An alternative possibility for a non-structured population would be to work only with the frequencies of strategies in a population (cf. Nowak & Sigmund, 1993), but then it is not straightforward to determine mutation rates and the initial frequencies of new mutations corresponding to the spatially structured situation.

For populations of a reasonable size this technique will lead to quite gigantic computations if

every single game has to be calculated separately. Fortunately, there are never as many different strategies in a population as there are individuals, because most individuals are offspring of a few successful players of the last generation. For example, in a population of size 50×50 with a mutation rate of 0.002 per individual and generation there are, in one generation, usually from 20 to 50 different strategies in the whole population. It is therefore possible to calculate the payoff of each type of interaction a certain number of times and use the mean as the payoff for all interactions of this type. Specifically, for the populations used for Figs. 1, 2 and 4 we used the following technical simplification. In every generation the payoff for every possible type of game, i.e. for every combination of the present strategies, is calculated five times and the mean value is taken as the expected payoff for every game in which the same combination of strategies interact. So, instead of calculating each of the $n \times n \times 8$ games separately, the payoff for all games in which the same two strategies interact has to be calculated only once (i.e. the average of five repetitions).

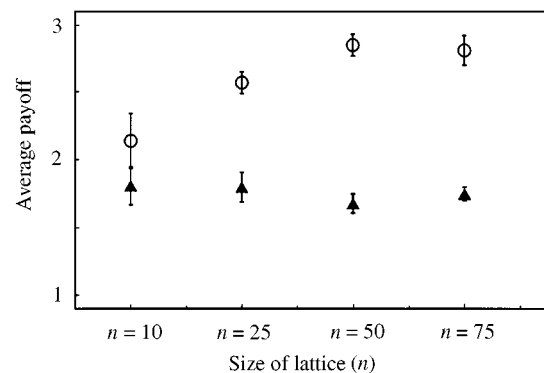


FIG. 1. The average payoff in spatial (○) and non-spatial (▲) populations. One measure of cooperation is the average payoff received by each player. The average payoff of each player is its total payoff divided by the number of iterations and by the number of neighbours. For all population sizes the payoff in the spatially structured populations (○) are higher than in the non-structured control (▲) and in the spatially structured populations the average payoff is higher in larger populations. The figure shows the average payoff from 20 populations which were each run for 10^4 generations with a mutation rate of 0.002. The averages are taken over the whole length of the simulations. All populations were started from one single strategy with $p_i = (0.5, 0.5, 0.5, 0.5)$. (Bars indicate standard deviation between populations.) (▲) Non-spatial; (○) spatial.

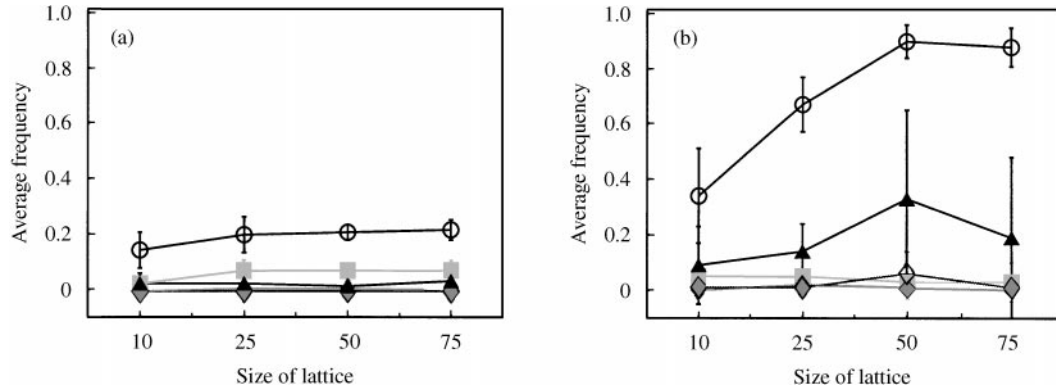


FIG. 2. Strategies present in spatial (a) and non-spatial (b) populations: The average percentage of cooperative strategies ($p_R > 0.9$) in a non-structured population (a) is around 20%. Of the four groups of strategies described in Table 2, RETALIATOR is the most frequent one in the non-spatial populations. In the spatially structured populations (b) the percentage of cooperative strategies is much higher, especially for large populations. In contrast to the non-spatial populations (a), in the structured populations PAVLOV is the most frequent of the strategies of Table 2 (simulations were run as in Fig. 1). (○) COOP; (■) RET; (◆) TFT; (◇) GTFT; (▲) PAV.

For the simulations with small populations (cf. Fig. 3) we used the program with the simplification and also one that calculated every single game. The results are qualitatively identical.

3. More Cooperation in Structured Populations

In a population of cooperative individuals the average payoff per round during an iterated game is close to R because most of the games are cooperative. As Fig. 1 shows, the average payoff per round in spatially structured populations of sufficient size is very close to R . This indicates that in these populations most of the individuals are engaging in cooperative behaviour. Only for small population sizes ($< 50 \times 50$) is the average payoff less than R . On the other hand, the average payoff in the non-structured populations is much smaller than R and not affected by the size of the population.

As mentioned above, one basic characteristic of a cooperative strategy is a value of p_R close to 1, and tracking the p_R -values of all individuals reveals the percentage of cooperative strategies in a population. The uppermost line in Fig. 2(a) and (b) shows the proportion of cooperative strategies (i.e. individuals with a value of $p_R > 0.9$). In the spatially structured populations the average frequency of cooperative strategies during the first 10^4 generations is much higher than in unstructured populations (spatial ≈ 0.9 , non-spatial ≈ 0.25).

A second difference between structured and non-structured populations is the “quality” of the most successful strategies. Of the four categories of strategies described in Table 2, RETALIATOR (RET) is the most frequent in the non-structured populations with an average frequency of almost 10% [Fig. 2(a)]. In the well-mixed populations RET makes up more than a third of all occurring cooperative strategies. RET cooperates only if both players mutually cooperated in the previous round. Therefore, RET is very resistant against exploitation by mutual defectors like ALL D, but at the same time, RET is also very vulnerable to stochastic errors, because it has no means to re-establish cooperation once it accidentally defected. In the spatial settings things look quite different. In the spatially structured populations [Fig. 2(b)] PAVLOV is the most successful of the strategies described in Table 2. PAVLOV is a “win stay-lose shift”-strategy (Nowak & Sigmund, 1993), which starts by playing C and then plays whatever it played in the previous round if it got a good payoff (T or R) or it plays the opposite of its own previous move if it scored badly (P or S). Nowak & Sigmund (1993) state that in a non-spatial setting the strict PAVLOV (1, 0, 0, 1) can be invaded by ALL D but a slightly stochastic variant like (0.999, 0.001, 0.001, 0.995) cannot. As stated above, we use here the term PAVLOV for a whole subset of strategies that have the typical characteristics of PAVLOV: namely, they are (1) cooperative, $p_R > 0.9$

and (2) they play C very often after a mutual defection, $p_P > 0.6$, while they mainly defect otherwise. One very clear advantage of PAVLOV over RET is its ability to return to mutual cooperation with other individuals playing the same strategy, after an occasional mistake leads to a defection (Nowak & Sigmund, 1993). This feature is very important in a stochastic world, where occasional errors occur. Because of this ability PAVLOV will cooperate most of the time when playing against itself, which is particularly important in a spatially structured population, where successful strategies are immediately surrounded by copies of themselves.

Indeed, the fact that successful strategies build clusters due to local reproduction is a crucial aspect of spatial structure. In our extensive simulations, we observed that many cooperative strategies need to be in clusters of a certain minimal size in order to persist over longer periods of time. In a spatial setting, many cooperative strategies are vulnerable to a single defective mutation occurring inside a cooperative cluster. Strategies that try to re-establish cooperation after mutual defections (e.g. PAVLOV) are especially easily exploited by defective mutations, because whenever they try to return to cooperation, they get exploited. Usually, defective mutations cannot invade the cooperative clusters, because as soon as there is more than one defective individual, they defect against each other and so score worse than the cooperators, provided that the cooperators have still enough other cooperators around them with whom they can cooperate. For many cooperative strategies it is therefore very important to build large enough clusters. To illustrate this, we simulated a limited cluster size by computing very small populations ($n = 5$ or 10). These small populations were each started from a single cluster of one strategy and each population was run—with random mutations occurring as described above—until no individual with the initial strategy was left. Figure 3 shows the average persistence (= time until the last offspring of the initial strategy has disappeared) of each of the four types of strategies described in Table 2 for population of size 5×5 and 10×10 . All the strategies were observed for several levels of cooperation, i.e. in six different variants of the value of p_R , ranging from 0.91 to 0.999.

For the non-spatial situation in Fig. 3(a), we see that RET is by far the most persistent strategy in both population sizes. The spatially structured situation [Fig. 3(b)], however, is different. In the smaller populations, RETALIATOR is still the most persistent strategy for all values of p_R . This is particularly interesting, as a retaliator with $p_R < w$, playing against itself, receives a rather poor payoff ($\ll R$), because after an accidental defection of one partner, both partners will very probably defect for the rest of the game (because p_S and p_P are small). In the larger populations of size 10×10 , however, PAVLOV is the most persistent strategy. Especially for relatively low values of p_R , PAVLOV is much better than all other strategies, because PAVLOV's ability to re-establish cooperation after a random defection is particularly important for low values of p_R (≈ 0.9), when accidental defections occur relatively frequently. The point to emphasize in Fig. 3(b) is that in spatially structured populations many of the more forgiving strategies (like PAVLOV and to a certain extent GTFT) that are typically very successful in structured populations (cf. Fig. 2), can only persist against invasion of defectors (or random mutations) if they can build large enough clusters.

As the payoff values chosen for these simulations ($R = 3$, $T = 5$, $S = 0$ and $P = 1$), for which $2R = T + P$, are on the borderline of evolutionary stability for PAVLOV (Boerlijst *et al.*, 1997), we did run other simulations with more rewarding payoffs ($R = 3.5$ or $T = 4$). Although, as one might expect, this generally yields a slight increase in the frequency of cooperative strategies, the results do not qualitatively differ from the results with the classical payoffs presented in this section.

4. The Role of Generosity

Another interesting point is the role of generosity by which we mean the fact that in the structured populations many of the successful strategies that we observed were less severe retaliators than the successful strategies in the unstructured populations. Grim (1995) has found in his simulations (using a simpler set of strategies than we use here) that in spatially structured populations a very generous variant of TIT FOR TAT,

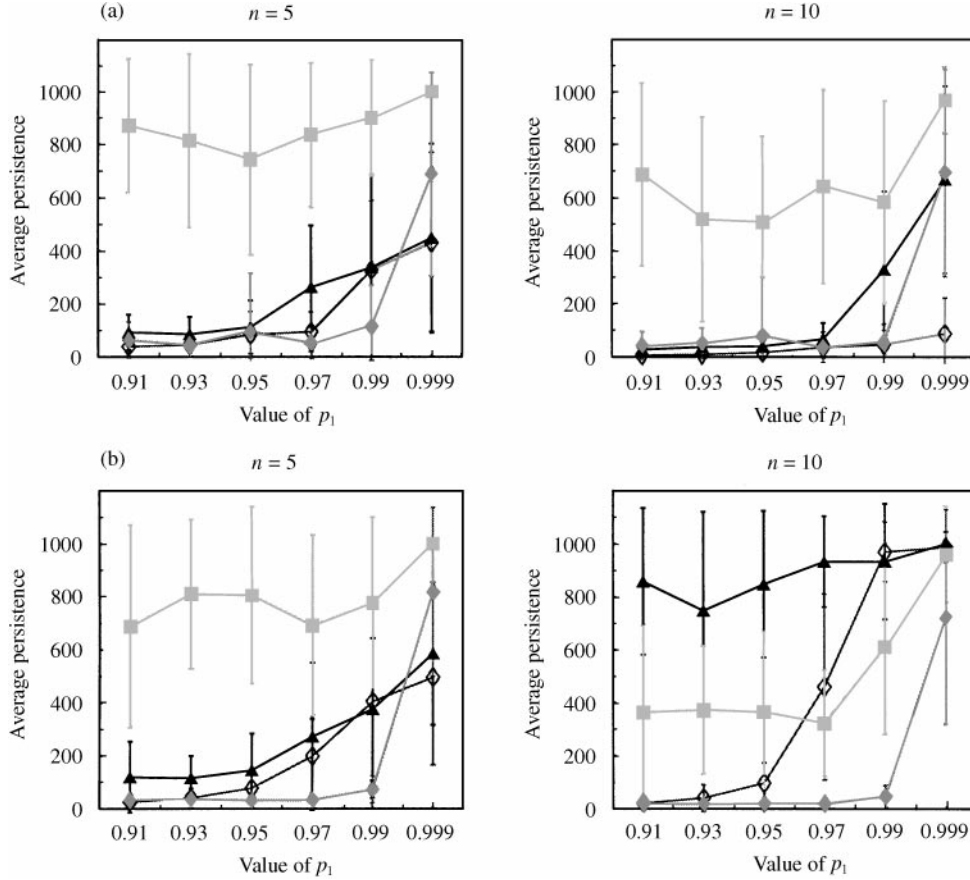


FIG. 3. The persistence of the different types of strategies depends very much on the possible cluster size: (a) non-spatial; (b) spatial. To illustrate this we simulated small cluster size by using very small populations. The populations were started from an artificial situation with all individuals having the same type of strategy. We started populations from all four strategies described in Table 2. For each type of strategy, we used six variants, which differed in the value of p_R , ranging from 0.91 to 0.999: i.e. here RET is $(p_R, 0.01, 0.01, 0.01)$, TFT is $(p_R, 0.99, 0.01, 0.01)$ and so on. To give an example, the leftmost black triangle in (b) ($n = 5$) indicates after how many generations none of the initial (PAVLOV-) strategy, which was in this case (0.91, 0.01, 0.01, 0.99), was left in the population. The parameters used are $w = 0.99$ and $mr = 0.01$, the bars indicate the standard deviation of the ten repetitions per population. (\diamond) GTFT; (\blacktriangle) PAV; (\blacksquare) RET; (\blacklozenge) TFT.

playing D only with a chance of 0.33 after its opponent defected, is a good strategy. In our notation, such a very generous TFT corresponds to (1, 1, 0.66, 0.66). In non-spatial populations the generous variant of TFT which performs best is (1, 1, 0.33, 0.33) (Nowak & Sigmund, 1992). For non-structured populations it is thought that PAVLOV should not be generous (Nowak & Sigmund, 1993), but we shall see that this conclusion does not hold in a spatial context.

To examine the role of generosity we classify the evolving strategies into two groups, one playing TFT-like and the other playing PAVLOV-like strategies. All strategies with values of $p_R > 0.9$, $p_T > 0.7$ and p_S and $p_P < 0.7$ are treated as TFT-like, and strategies with $p_R > 0.9$, p_T and

$p_S < 0.7$ and $p_P > 0.7$ are treated as PAVLOV-like. Additionally, we divided each group into three levels of generosity: generosity is here the value of p_S and p_P for TFT and p_T and p_S for PAVLOV. In the non-generous level, both values are < 0.2 . In the generous level, at least one value is in (0.2, 0.5) while in the very generous level one value is > 0.5 . The results (Fig. 4) show that for the TFT-like strategies we found more or less what could be expected from the literature (Nowak & Sigmund, 1992; Grim, 1995). In the unstructured populations most TFT-like strategies had a generosity between 0.2 and 0.5 but there were also pure TFT and very generous TFT. In the spatially structured populations we found almost no pure TFT but the generous and the very




FIG. 4. The role of generosity: (a) non-spatial; (b) spatial. Classifying the TFT-like and the PAVLOV-like strategies by their level of generosity (cf. text) shows that generous variants are much more successful in spatial populations than in non-spatial populations. This is especially interesting for PAVLOV-like strategies, where generous variants nearly never established in non-spatial populations but dominated the spatial populations. The generosity level on the x-axis (1–3) correspond to the classification of non-generous, generous and very generous used in the text. (The bars indicate averages of 20 simulations as described in Fig. 1.) (□) TFT; (■) PAVLOV.

generous TFT were more frequent. Surprisingly, the difference between the different levels of generosity was much more striking with the PAVLOV-like strategies. In the unstructured populations, the overwhelming majority of the PAVLOV-like strategies were not generous, whereas in the structured populations, the generous and the very generous variants were each at least twice as frequent as the non-generous PAVLOV.

To investigate the stability of the generous variants of PAVLOV and TFT in structured populations we conducted simulations similar to those summarized in Fig. 3, starting populations from a cluster of PAVLOV or TFT with different levels of generosity. This showed very clearly that without spatial structure only very non-generous PAVLOV can survive for more than a few generations, while in structured populations the generous variants did very well. In the spatial model, generous PAVLOV with any value of p_T and p_S with $0 < p_T = p_S < 0.5$ usually persisted for the 1000 generations that we observed. With the TFT-like strategies the difference is not that dramatic. TFT with a generosity < 0.3 usually persisted during the 1000 generations observed in structured as well as in unstructured populations. The more generous TFT, where both p_S and $p_P > 0.3$, only persisted in the structured populations.

5. Dynamics in Structured Populations

Besides the evolution of more cooperative and generous behaviour, spatial structure also has

a profound effect on the evolutionary dynamics of the systems studied. In Fig. 5, the average values of p_i in a population are plotted over 10^4 generations for both a structured and an unstructured population. In the spatially structured population [Fig. 5(a)] the first clear strategy that evolves is TIT FOR TAT. Then, about 400 generations later, a generous PAVLOV-like strategy takes over until it is itself superseded by a “generous” RETALIATOR in generation 2000. From generation 7000 to the end of the 10^4 generations observed, there is again a generous PAVLOV. In the non-structured population [Fig. 5(b)] there is at the beginning a long period without cooperation. From generation 2500 on, some form of RETALIATOR is dominant. Later, at generation 6000, there is a long period of dominance by PAVLOV and at the end of this another “generous” RETALIATOR appears. We also note that unlike the spatially structured case [Fig. 5(a)] cooperative periods are interrupted several times by periods without cooperation. Studying many such simulations shows that once cooperation ($p_R > 0.9$) has evolved in the spatial population, it is very unlikely to be lost again. This certainly cannot be said of the non-spatial population. Overall, in the non-structured populations the evolutionary dynamics show much larger and more persistent fluctuations in strategy space. Thus, evolution is generally much less chaotic in spatially structured populations.

Comparing populations of different sizes reveals more about these dynamics. To characterize the

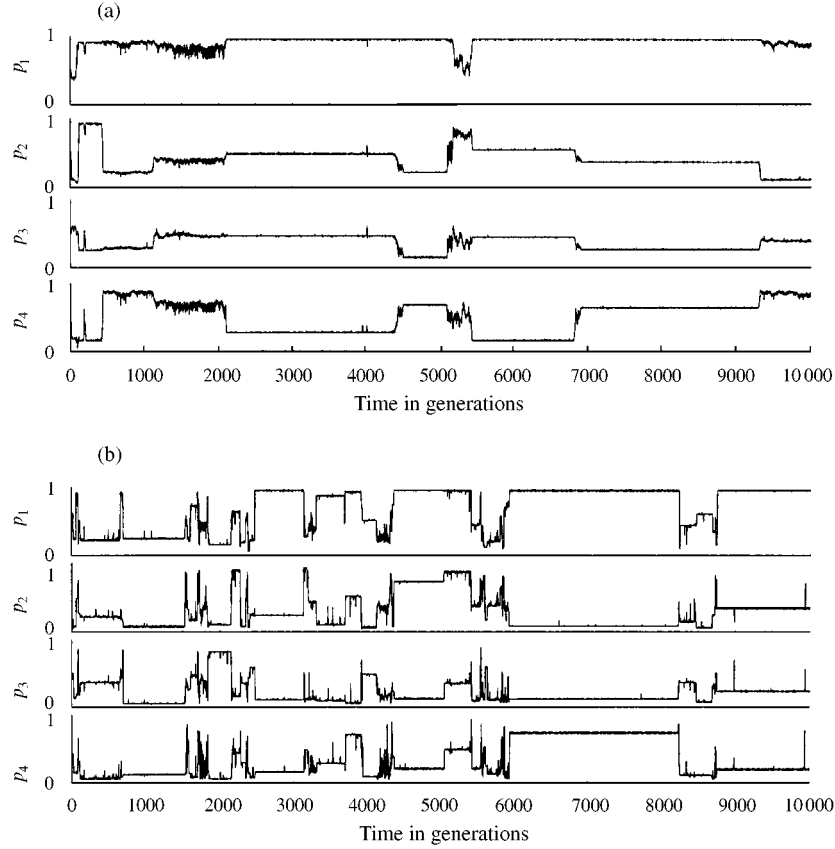


FIG. 5. Time series showing the average value of p_i in a spatial (a) and a non-spatial (b) population. In spatially structured populations the dynamics are typically much less chaotic than in well-mixed populations. Comparing the PAVLOV-like strategies in (a) from generation 7000 on with the PAVLOV in (b) from generation 6000 to 8000 demonstrates the tendency towards more generous strategies in spatially structured populations. For the purpose of illustration, relatively small populations of size 25×25 were used for these plots. While the spatial populations of larger size do not differ much from the one shown here, except that they are even more stable, larger non-structured populations become so chaotic that it is much harder to identify certain strategies. The populations were started from $(0.5, 0.5, 0.5, 0.5)$ with $w = 0.995$ and $mr = 0.002$.

dynamics we chose to look at the length of uninterrupted periods of cooperation. Table 3 shows the average length of periods of cooperation. In contrast to the definition of cooperation in Section 3 ($p_R > 0.9$), an interval of cooperation here means that the population average of p_R is $p_R > 0.5$ during the interval. Of course, $p_R > 0.5$ does not explicitly mean that a population consist of mostly cooperative individuals, but as Fig. 5 and the analysis of the raw data show, the transitions from large to small values of p_R are often relatively sharp, jumping from high to low in a very short time with nearly no intermediate values. The average value of all p_R when $p_R > 0.5$ is around 0.9, and not at 0.75 as one would expect with a uniform random distribution. So if $p_R > 0.5$ most of the individuals are usually

TABLE 3
Average length of uninterrupted periods of cooperation in simulations over 10^4 generations (to reduce the influence of stochastic noise only periods with a length of more than five generations were counted). In the spatially structured populations these periods are much longer, which indicates that cooperative configurations are more stable in a spatially structured context. For large populations, once cooperation was established, it usually lasted for the whole simulation. (Averages from ten simulations with $w = 0.995$ and $mr = 0.002$)

Size	Spatial	Non-spatial
10×50	524	434
25×25	4081	216
50×50	9302	98

cooperators. By using this threshold of $p_R = 0.5$ to distinguish between cooperative and non-cooperative phases, very short stochastic disturbances can be filtered out effectively.

In spatially structured populations, the periods of cooperation are longer if the population size is larger. For a population size of 50×50 or more the length of the cooperative periods are nearly 10^4 generations. This shows that once cooperative strategies have evolved in a structured population, they are usually only superseded by other cooperative strategies, and not by non-cooperative strategies. In the unstructured populations, we see quite the opposite. In large populations, the periods of uninterrupted cooperation are even shorter than in small populations. Two reasons can be given for this: (1) In a large population there are more individuals which means more mutations. Therefore, there is a higher chance that a mutation occurs that can exploit the present strategies. (2) Because interactions are not restricted locally, these new mutants can basically exploit any other individual of the population. In Nowak & Sigmund's (1993) simulations using the same set of strategies as here, but with evolutionary dynamics based only on the frequency of the strategies (no spatial structure), it was found that stable cooperative behaviour established only on much longer time scales than we examined (10^5 – 10^7 generations). Due to the complex calculations of our simulations, it is not possible to observe such long time-scales with our model.

It seems, however, that in the non-spatial situations the class of cooperative strategies that can become established and persist is much more limited than in the spatial situation. We observe that in the spatially structured populations stable cooperative configurations evolve much faster than in non-structured populations. As illustrated in Fig. 5, it usually takes only a few hundred generations until cooperative behaviour has evolved and is established in structured populations. Afterwards, there is much fine tuning, but the cooperative character is not lost thereafter. In non-structured populations, many more generations are needed until a stable cooperative strategy can eventually become established (cf. Nowak & Sigmund, 1993).

Another interesting difference in the evolutionary dynamics is that in the spatially structured

populations there are often combinations of strategies where both strategies can exploit each other, depending on the local configurations. In Fig. 5, this shows up as fast fluctuations of small amplitude of the p_i (e.g. generation 500–2000). Such fluctuating configurations never occur in the non-structured populations and seem very typical of spatial structure. This phenomenon has been found in many other spatial models (Grim, 1995; Killingback & Doebeli, 1996; Nowak & May, 1992) and is extensively described in Lindgren & Nordahl (1994).

6. Discussion

The idea that spatial structure facilitates the evolution of cooperation is already well supported by the work of others (Nowak & May, 1992; Nakamaru *et al.*, 1996; Killingback & Doebeli, 1996; Doebeli & Knowlton, 1998). The presence of spatial structure, however, does not only affect the quantity of cooperative strategies but also their "quality", meaning that in a structured population the successful strategies are not the same as in an unstructured population. For example, in Axelrod's (1984) work, it was found that the most successful strategy in spatial populations was a strategy that ranked only 31st of 62 in the computer tournament. In addition, the evolutionary dynamics of how cooperation emerges and is maintained is very different in spatially structured populations.

One very important consequence of spatial structure is that successful strategies tend to build clusters of similar individuals because of local reproduction. With such clusters some strategies can spread in an environment in which they could not spread without clustering. For example, in a spatial context TIT FOR TAT can invade a population of ALL D (Axelrod, 1984; Nakamaru *et al.*, 1996). Even in the simple Prisoner's Dilemma without iteration C can invade into D under certain conditions (Nowak & May, 1992). In a two-dimensional lattice, this is possible if a C that is at the margin of a cluster of C is scoring higher than the D outside the cluster of C. To illustrate the mechanism favouring cooperation in spatially structured populations, let us consider a cluster of 3×3 C in a sea of D as shown in Fig. 6. The C in the middle




FIG. 6. A cluster of C in a sea of D (cf. text). (■) C; (□) D.

of a margin of the cluster (e.g. no. 1 in Fig. 6) has five C neighbours and three D. It therefore scores $5R + 3S$. A D outside the cluster (e.g. no. 2) can have at most three C neighbours and will therefore score less than or equal to $3T + 5P$. Thus, if $5R + 3S > 3T + 5P$ the D at the border will be replaced by a C in the next generation and the cluster of C grows into the sea of D. We observe that similar mechanisms are operating in our stochastic model. For example, as Table 4 shows, the interaction of PAVLOV vs. RETALIATOR is also a Prisoner's Dilemma, under the conditions of our model. With PAVLOV as the cooperator and RET as the defector, we have $T = 2.98 > R = 2.95 > P = 1.82 > S = 1.72$ (and $2R > T + P$). In particular, we have $5R + 3S = 19.9 > 3T + 5P = 18.0$, which is the sufficient condition in a spatially structured population for clusters of PAVLOV to be able to invade into a sea of RET.

In a population that is not structured, PAVLOV cannot directly invade into a population of Retaliators nor can it invade into a population of ALL D. This does not mean that there is no possibility for PAVLOV to evolve in a non-structured population. Nowak & Sigmund (1993) found that PAVLOV can invade into a population of e.g. ALL D by succeeding other cooperative strategies such as TIT FOR TAT, which in turn *can* invade into populations of less cooperative strategies. In their simulations, which are more or less a mean-field approximation of the spatial simulations in this paper, they found that PAVLOV also dominated the non-spatial situation, but only on much longer time-scales than we used (10^5 – 10^7 generations). In contrast to our model, RETALIATOR played a much less dominant role in Nowak & Sigmund's (1993) simulations. This can be explained mainly by the fact that for computational

TABLE 4
Payoff matrix for a PAVLOV vs. RETALIATOR game. The interaction is a Prisoner's Dilemma where PAV is the cooperator and RET is the defector. Averages from 10^4 rounds per iterated game with $PAV = (0.99, 0.01, 0.01, 0.99)$ and $RET = (0.99, 0.01, 0.01, 0.01)$ and $w = 0.995$, which gives an average of 200 iterations per game

	PAV	RET
PAV	2.95	1.72
RET	2.98	1.82

simplicity, Nowak & Sigmund used infinitely iterated games. The reason why we use finite games, where the probability w that two players will meet again and play another game satisfies $w < 1$, is that under the assumption of infinite iteration certain reasonable strategies receive a much lower payoff than with finite games. For example, RET (cf. Table 2) with a value of $p_R > w$ will cooperate for most of the length of a finite game. In an infinitely iterated game the finite phase of cooperation at the start of the game has no influence compared to the infinite phase of mutual defection that follows as soon as one player accidentally defects. While a Retaliator with $p_R > w$ is a cooperative strategy in finite games, it is a rather non-cooperative strategy in an infinitely iterated game. Thus, with infinitely iterated games a strategy with very limited error correction (e.g. RET) scores very poorly against itself, and therefore plays a much less prominent role overall. With a finite game length, the advantage of error correction is not so big and it is much harder to supersede strategies like RET (in the non-spatial case).

In our simulations we have seen that besides iteration (infinite or not) spatial structure is another strong reason for the evolution of error correction. The importance of an error-correcting mechanism in a spatially structured population is documented in Fig. 3. Provided large enough spatial cluster are possible ($\geq 10 \times 10$), the strategies with good error correction (PAVLOV and GTFT) score best and RET

does worse than in the other situations documented. The raw data even show that in the large spatial populations (10×10), RET was mainly displaced by other cooperative strategies.

Another way of minimizing the losses due to accidental errors is playing more generously. Generous Strategies will defect only with a certain probability in situations where pure, deterministic strategies like TFT or PAVLOV would always play D. Our simulations have shown that in spatially structured populations much more generous strategies evolve. For TFT-like strategies this was also found by Nowak & Sigmund (1992) and by Grim (1995), but in simulations that incorporated only reactive strategies. Reactive strategies depend solely on their opponents' last move and do not take into account what they have played in the last round themselves. Thus, strategies like RET or PAVLOV did not occur in Grim's simulations. While GTFT is also successful in non-spatial populations, we found very few generous PAVLOV in non-spatial populations. One disadvantage of PAVLOV is that it is exploited by very non-cooperative strategies in every second round of an iterated game. Generous variants of PAVLOV are exploited even more. Thus, in a non-spatial setting, PAVLOV is evolutionary stable against defectors only when $T + wP < R + wR$ (Boerlijst *et al.*, 1997). In our simulations with the classical payoff matrix $R = 3$, $T = 5$, $S = 0$ and $P = 1$ this is not the case: as $T + wP > R + wR$ except if $w = 1$, these payoffs are on the borderline of evolutionary stability for PAVLOV. In simulations with payoffs satisfying $T + wP < R + wR$, so that cooperative ESSs are possible in the non-spatial setup (Boerlijst *et al.*, 1997), there is a general tendency towards more cooperative strategies in both spatial and non-spatial setup. Using more rewarding payoffs with $R = 3.5$ or $T = 4$, in the non-spatial case especially the generous variants of PAVLOV and RETALIATOR appear more frequently, whereas nearly no generous PAVLOV appear in the non-spatial set-up with the classical payoffs (cf. Fig. 4). But also in the spatial set-up we observed an even stronger dominance by generous and very generous variants of PAVLOV if R is raised to 3.5 or T lowered to 4. The interesting point, however, is that in the spatial context cooperative strategies evolve

quite frequently even if $T + wP < R + wR$ is not satisfied.

There is also other evidence that PAVLOV is a very successful strategy especially in a spatial context. Lindgren & Nordahl (1994) studied the evolutionary dynamics in spatially structured populations. They represent strategies using a kind of genetic sequence which determines whether a strategy plays C or D according to what was played in the past n moves of a game. These strategies are thus purely deterministic which means for example that no "generous" strategies (in the sense used in this paper) are possible. Lindgren & Nordahl start their simulations with strategies with a memory of $n = 0$. Via genetic mutations, such as gene duplications, the memory of the strategies can increase in the course of evolution. When the memory of the strategies has evolved to a length of 2—they remember their own and their opponents last move, which corresponds to the memory length of the strategies in our paper—Lindgren and Nordahl find that the spatial populations are dominated by a strategy with gene sequence 1001. This 1001 strategy plays C after receiving payoff R or P and plays D after T or S , which is exactly PAVLOV. Only when the memory has evolved to a length of 4 are there other more complicated strategies which supersede 1001. However, the successful memory 4 strategies that evolve are basically more sophisticated versions of PAVLOV which return to cooperation only after a series of two mutual defections or when a mutual defection follows a mutual cooperation. They compare the spatial situation with a mean-field model, where they find that the strategy 0001 dominates much of the initial phase of the populations, when the memory of strategies has length two. In our terminology, this 0001-strategy is RETALIATOR. Interestingly, for the non-spatial set-up these results correspond well with the Wedekind & Milinski (1996) observations of human cooperation in the Prisoner's Dilemma. The most frequent strategy they found examining first year biology students' behaviour in playing the Iterated Prisoner's Dilemma corresponds, in our terminology, best to a moderately generous RETALIATOR ($p_R \approx 0.7$, $p_{T,S,P} \approx 0.2$), although Wedekind & Milinski preferred to call it PAVLOV.

Those strategies which exhibit TIT FOR TAT-like behaviour (e.g. TFT and GTFT), which are classically regarded as particularly important strategies in the Iterated Prisoner's Dilemma (Axelrod & Hamilton, 1981; Nowak & Sigmund, 1992; Grim, 1995), play a surprisingly small role in our investigations. An important disadvantage of TFT-like behaviour is that TFT lacks the ability to exploit unconditional cooperators. TFT for example tolerates the invasion of ALL C by neutral drift. But if there are too many ALL C, then the TFT/ALL C-clusters can be easily invaded by strategies that exploit the ALL C and thereby supersede TFT. This deficiency of TFT-like strategies is more serious in the spatial IPD than in the non-spatial case as many strategies close to ALL C can survive in the spatial game, whereas they would be unable to do so in the non-spatial game. PAVLOV, on the other hand, does not suffer from this deficiency. Once it detects unconditional cooperators by chance, it can exploit them and prevent them from drifting into a PAVLOV-cluster (Nowak & Sigmund, 1993). This ability is particularly valuable in a spatial context.

In conclusion, using a relatively realistic, stochastic model, with a general strategy-space of all stochastic strategies with a memory of one round, we investigated some aspects of the evolution of cooperation. We found—as other work has suggested for models with simpler sets of strategies (Nowak & May, 1992; Grim, 1995; Lindgren & Nordahl, 1994; Killingback & Doebeli, 1996)—that spatial structure greatly influences the evolution of cooperation. By comparing populations with a well-defined spatial structure with populations where interactions and dispersion of individuals are random, we observed the following effects of spatial structure: (1) spatial structure greatly facilitates the evolution of cooperative behaviour; there are more cooperative individuals in a spatially structured population and cooperation evolves faster. (2) In structured populations qualitatively different strategies evolve than in non-structured popula-

tions: in the spatial context, we found a strong tendency towards more generous strategies and generous variants of PAVLOV dominated many of the simulations. In contrast, the non-structured populations were mainly dominated by RETALIATOR. (3) Spatial structure has a strong stabilizing effect on evolutionary dynamics. Clustering, as it occurs in spatially structured populations of sufficient size, is a very strong reason for the evolution of strategies with a good error correction. In our context these are strategies like PAVLOV and generous PAVLOV.

We would like to thank Karl Sigmund for critical comments on the manuscript.

REFERENCES

- AXELROD, R. (1980). Evolution of cooperation. *J. Conflict Resolution* **24**, 3–25.
- AXELROD, R. (1980). More on the evolution of cooperation. *J. Conflict Resolution* **24**, 379–403.
- AXELROD, R. (1984). *The Evolution of Cooperation*. New York: Basic Books Inc.
- AXELROD, R. & HAMILTON, W. D. (1981). The evolution of cooperation. *Science* **211**, 1390–1396.
- BOERLIJST, M. C., NOWAK, M. A. and SIGMUND, K. (1997). The logic of contrition. *J. theor. Biol.* **185**, 281–293.
- DOEBELI, M. & KNOWLTON, N. (1998). The evolution of interspecific mutualisms. *Proc. Nat. Acad. Sci. U.S.A.* **95**, 8676–8680.
- GRIM, P. (1995). The greater generosity of the spatialized prisoner's dilemma. *J. theor. Biol.* **173**, 353–359.
- KILLINGBACK, T. & DOEBELI, M. (1996). Spatial evolutionary game theory: hawks and doves revisited. *Proc. R. Soc. Lond. B.* **263**, 1135–1144.
- LINDGREN, K. & NORDAHL, M. (1994). Evolutionary dynamics of spatial games. *Physica D* **75**.
- MAYNARD SMITH, J. (1982). *Evolution and the Theory of Games*. Cambridge, Cambridge University Press.
- NAKAMARU, M., MATSUDA H. & IWASA, Y. (1996). The evolution of cooperation in a lattice-structured population. *J. theor. Biol.* **184**, 65–81.
- NOWAK, M. A. & MAY, R. M. (1992). Evolutionary games and spatial chaos. *Nature (Lond)* **246**, 15–18.
- NOWAK, M. A. & SIGMUND, K. (1992). Tit for tat in heterogeneous populations. *Nature* **355**, 250–252.
- NOWAK, M. A. & SIGMUND, K. (1993). A strategy of win–stay, lose–shift that outperforms tit-for-tat in the prisoner's dilemma. *Nature* **364**, 56–58.
- WEDEKIND, C. & MILINSKI, M. (1996). Human cooperation in the simultaneous and the alternating prisoner's dilemma: Pavlov versus generous tit-for-tat. *Proc. Nat. Acad. Sci. U.S.A.* **93**, 2686–2689.