Agents with Heterogeneous Strategies Interacting in a Spatial IPD

Frank Schweitzer[1,2*], Robert Mach[1,3], and Heinz Mühlenbein[1]

[1] Fraunhofer Institute for Autonomous Intelligent Systems, Schloss Birlinghoven, D-53757 Sankt Augustin, Germany
[2] Institute of Physics, Humboldt University Berlin, 10099 Berlin, Germany
[3] Institute for Theoretical Physics, Cologne University, D-50923 Köln, Germany

Summary. We use a spatial iterated Prisoner's Dilemma game (IPD) to investigate the spatial-temporal evolution of heterogeneity in agents' strategies. In our model, N agents are spatially distributed on a lattice and each agent is assumed to interact with her 4 local neighbors a number of n_g times during each generation. If the agent has a one-step memory for the last action of each individual neighbor, this results in a total of eight different strategies for the game. After each generation, the agent will be replaced by an offspring that adopts the strategy of her most successful neighbor.

The agents are heterogeneous in that they play different strategies dependent on (i) their past experience, (ii) their local neighborhood. The spatial-temporal distribution of these strategies is investigated by means of computer simulations on a cellular automaton. In particular, we study the incluence of n_g on the dynamics of the global frequencies of the different strategies and the conditions for a stationary (frozen) or non-stationary (dynamic) coexistence of particular strategies on a spatial scale.

1 Introduction

Agent-based models are important for the understanding of microeconomic interaction. A pervasive characteristic is the *heterogeneity* of agent behavior which may result e.g. from interaction with different neighbors (*locality*), or different individual experience in the course of time (*historicity*), or heterogeneous *environmental conditions*, or simply individual *diversity*. While interacting with others, an agent's success or failure often depends on her strategic behavior. To reduce the risk of making the wrong decision, it often seems to be appropriate to copy the most successful strategy, this way adapting to the local environment. This kind of *local imitation* behavior will be used in this paper to explain the spatial evolution of heterogeneity in a multi-agent system.

*corresponding author: schweitzer@ais.fraunhofer.de

We consider a system of N agents spatially distributed on a square lattice, so that each lattice site is occupied by just one agent. Each agent i is characterized by two state variables, (i) her position $\boldsymbol{r}_i$ on the lattice, and (ii) a discrete variable θ_i, describing her possible actions, as specified in Sect. 3. Agents are assumed to directly interact only with their 4 nearest neighbors a number of n_g times. In order to describe the local interaction, we use the so-called *iterated Prisoner's Dilemma* (IPD) game – a paradigmatic example [1, 4, 11] well established in evolutionary game theory with a broad range of applications in politics, economy, and sociology.

In the simple Prisoner's Dilemma (PD) game, each agent i has two options to *act* in a given situation, to *cooperate* (C), or to *defect* (D). Playing with agent j, the outcome of this interaction depends on the action chosen by agent i, i.e. C or D, *without knowing* the action chosen by the other agent participating in a particular game. This outcome is described by a *payoff matrix*, which for the 2-person game, i.e. for the interaction of only two agents, has the following form:

$$\begin{array}{ccc} & \mathbf{C} & \mathbf{D} \\ \mathbf{C} & R,R & S,T \\ \mathbf{D} & T,S & P,P \end{array} \tag{1}$$

In PD games, the payoffs have to fulfill the following two inequalities:

$$T > R > P > S \qquad 2R > S + T \tag{2}$$

The known standard values are $T = 5$, $R = 3$, $P = 1$, $S = 0$. This means in a cooperating environment, a defector will get the highest payoff. From this, the abbreviations for the different payoffs become clear: T means (T)emptation payoff for defecting in a cooperative environment, S means (S)ucker's payoff for cooperating in a defecting environment, R means (R)eward payoff for cooperating in a likewise environment, and P means (P)unishment payoff for defecting in a likewise environment.

In any *one* round (or "one-shot") game, choosing action D is unbeatable, because it rewards the higher payoff for agent i whether the opponent chooses C or D. At the same time, the payoff for *both agents i and j* is maximized when both cooperate. But in a *consecutive* game played many times, both agents, by simply choosing D, would end up earning less than they would earn by collaborating. Thus, the number of games n_g two agents play together becomes important. For $n_g \geq 2$, this is called an *iterated Prisoner's Dilemma* (IPD). It makes sense only if the agents can remember the previous choices of their opponents, i.e. if they have a memory of $n_m \leq n_g - 1$ steps. Then, they are able to develop different *strategies* based on their past experiences with their opponents, which is described in the following.

2 Agent's Strategies

In this paper, we assume only a *one-step* memory of the agent, $n_m = 1$. Based on the known previous choice of her opponent, either C or D, agent i has then the choice between *eight* different strategies. Following a notation introduced by [10], these strategies are coded in a 3-bit binary string $[I_o|I_c\,I_d]$ which always refers to *collaboration.* The first bit represents the *initial* choice of agent i: it is 1 if agent i *collaborates*, and 0 if she defects initially. The two other values refer always to the previous choice of agent j. I_c is set to 1 if agent i chooses to collaborate given that agent j has collaborated before and 0 otherwise. Similarly, I_d is set to 1 if agent i chooses to collaborate given that agent j has defected before and 0 otherwise. Both I_c and I_d can be also interpreted as *probabilities* to choose the respective action given the knowledge of the previous choice of the opponent. But in the *deterministic* case, I_c and I_d are either 0 or 1. Thus, *eight* different strategies ($s = 0, 1, \ldots, 7$) result, which are given in Tab. 1.

s	Strategy	Acronym	Bit String
0	suspicious defect	sD	000
1	suspicious anti-Tit-For-Tat	sATFT	001
2	suspicious Tit-For-Tat	sTFT	010
3	suspicious cooperate	sC	011
4	generous defect	gD	100
5	generous anti-Tit-For-Tat	gATFT	101
6	generous Tit-For-Tat	gTFT	110
7	generous cooperate	gC	111

Table 1. Possible agent's strategies using a one-step memory.

Depending on the agent's *first move*, we can distinguish between two different classes of strategies: (i) *suspicious* ($s = 0, 1, 2, 3$), i.e. the agent *initially defects*, and (ii) *generous* ($s = 4, 5, 6, 7$), i.e. the agent *initially cooperates.* Further, we note that four of the possible strategies do not pay attention to the opponent's previous action, i.e. except for the first move, the agent continues to act in the same way, therefore the strategies sD, sC, gD, gC ($s = 0, 3, 4, 7$) can be also named *rigid* strategies.

Non-trivial strategies are $s = 1, 2, 5, 6$. Strategy $s = 6$, known as (generous) "tit for tat" (gTFT), means that agent i collaborates in the first round and *imitates* her opponent's previous action in every subsequent round. Agents playing strategy gATFT ($s = 5$) also start cooperating but go on to do the opposite of whatever her opponent did in the previous move. sATFT ($s = 1$)

and sTFT ($s = 2$) represent the respective suspicious versions of the above strategies. It is straight forward to see that two agents playing gTFT continue to cooperate in an iterated game, while agents playing gTFT and sTFT alternately cooperate and defect. This illustrates that the *first move* of a strategy can be vital to the outcome of the game and also, that the *heterogeneous behavior* of an agent emerges from the *interaction with different opponents.* Furthermore we note, that the number of interactions n_g is also a crucial parameter in this game, because, if n_g is even, gTFT and sTFT will gain the same, but in case of n_g being odd sTFT will gain more than gTFT.

3 Spatial Interaction

So far, we have explained the interaction of *two agents* with one-step memory. This shall be put now into the perspective of a *spatial game* with *local* interaction among the agents. A spatially extended (non-iterative) PD game was first proposed by Axelrod [1]. Based on these investigations, Nowak and May simulated a spatial PD game on a cellular automaton and found complex spatiotemporal dynamics [8, 9]. A recent mathematical analysis [2] revealed the critical conditions for the *spatial coexistence* of cooperators and defectors with either a majority of cooperators in large spatial domains, or a minority of cooperators in small (non-stationary) clusters.

In the following, we concentrate on the *iterated* PD game, where the number of encounters, n_g, plays an important role. We note that possible extensions of the IPD model have been investigated e.g. by Lindgren and Nordahl [5], who introduced players which sometimes act erroneously, this way allowing a complex evolution of strategies in an unbounded strategy space.

In the spatial game, we have to consider *local configurations* of agents playing different strategies (see Fig.1). As explained in the beginning, each agent i interacts only with her *four nearest neighbors.* Let us define the size of a neighborhood by n (that also includes agent i), then the different neighbors of i are characterized by a second index $j = 1, ..., n-1$. The mutual interaction between them results in a *n-person* game, i.e. $n = 5$ agents interact simultaneously. In this paper, we use the assumption that the 5-person game is decomposed into $(n-1)$ 2-person games, that may occur *independently,* but *simultaneously* [4, 11], a possible investigation of a "true" 5-person PD game is also given in [11].

We further specify the θ_i that characterize the possible actions of each agent as one of the strategies that could be played (Tab. 1), i.e. $\theta_i \in s = \{0, 1, \ldots, 7\}$. The total number of agents playing strategy s in the neighborhood of agent i is given by:

$$k_i^s = \sum_{j=1}^{n-1} \delta_{s\,\theta_{i_j}} \tag{3}$$

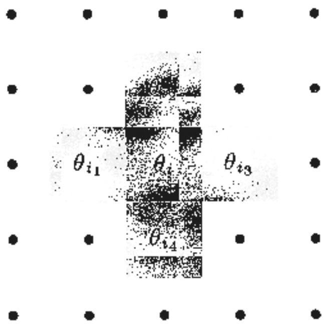

Fig. 1. Local neigbhorhood of agent i. The nearest neighbors are labeled by a second index $j = 1, ..., 4$. Note that $j = 0$ refers to the agent in the center.

where δ_{xy} means the Kronecker delta, which is 1 only for $x = y$ and zero otherwise. The vector $\boldsymbol{k}_i = \{k_i^0, k_i^1, k_i^2, \ldots, k_i^7\}$ then describes the "occupation numbers" of the different strategies in the neighborhood of agent i playing strategy θ_i.

Agent i encounters with each of her four neighbors playing strategy θ_{i_j} in independent 2-person games from which she receives a payoff denoted by $a_{\theta_i \theta_{i_j}}$, which can be calculated with respect to the payoff matrix, eq. (1). The total payoff of an agent i after these independent games is then simply

$$a_i(\theta_i) = \sum_{j=1}^{n-1} a_{\theta_i \theta_{i_j}} = \sum_s a_{\theta_i s} \cdot k_i^s \tag{4}$$

We note again that the payoffs $a_{\theta_i s}$ also strongly depend on the number of encounters, n_g, for which explicit expressions have been derived. They are summarized in a 8×8 payoff matrix not printed here [12].

In order to introduce a *time scale*, we define a *generation* G to be the time in which each agent has interacted with her $n-1$ nearest neighbors n_g times. During each generation, the strategy θ_i of an agent is not changed while she interacts with her neighbors simultaneously. But after a generation is completed, θ_i can be changed based on a comparison of the payoffs received. I.e., payoff a_i is compared to the payoffs a_{i_j} of all neighboring agents, in order to find the maximum payoff within the local neighborhood during that generation, $\max\{a_i, a_{i_j}\}$. If agent i has received the highest payoff, then she will keep her θ_i, i.e. she will *continue* to play her strategy. Otherwise, agent i has the possibility to adopt or to *imitate* the strategy of the most successful agent in her neighborhood. This implies that agent i has the possibility to directly observe the *strategies* of neighboring agents. If there is more than one neighbor with the same highest payoff, agent i randomly chooses one of the strategies of these neighbors.

The maximum payoff in the neighborhood of agent i is given by

$$a_{\max}(i) = \max_{j=1,\ldots,n-1} \{a_i, a_{i_j}\} \tag{5}$$

and

$$j^{\star}(i) = \arg\max_{j=0,\ldots,n-1} a_{i_j} \tag{6}$$

defines the position of the agent that received the highest payoff in the neighborhood. The update rule of the game can then be summarized as follows:

$$\theta_i(G+1) = \begin{cases} \theta_i(G) & \text{if } \; a_i(\theta_i) = a_{\max}(i) \\ \theta_{i_{j^\star(i)[\eta]}}(G) & \text{else} \end{cases} \tag{7}$$

where $j^{\star}(i)[\eta]$ denotes the case, where agent i randomly chooses a strategy of one of the most successful neighbors, η. We note that the evolution of the system described by eq. (7) is deterministic, except that equally sucessful strategies in a local neighborhood need to be chosen randomly. Results for stochastic cellular automata (CA) have been discussed in [3, 7].

The adaptation process leads to an evolution of the spatial distribution of strategies that will be investigated by means of computer simulations on a cellular automaton in the following section.

4 Evolution of Spatial Patterns of Strategies

In order to get a graphic idea of the spatio-temporal evolution of heterogeneity, we have restricted the computer simulation to only *three strategies* instead of eight, namely sD, sATFT and gTFT ($s = 0, 1, 6$). Strategy sD is known to be the winning strategy for a one-shot game, i.e. $n_g = 1$, while gTFT is known to be the most successful strategy for $n_g \geq 4$. Thus when n_g is increased from 1 to 4 a transition from defection to cooperation can be observed in the system. This is of particular interest in this paper, i.e. we mainly investigate $n_g = 2$ and $n_g = 3$ to focus on the transition region.

Agents playing sATFT are added to the initial population, since they behave anti-cyclic, i.e. they defect when the opponent cooperated and vice versa. The simulations are carried out on a 100×100 lattice with periodic boundary conditions, in order to eliminate spatial artifacts at the borders. Initially, all agents are randomly assigned one of the three strategies. Defining the total fraction of agents playing strategy s at generation G as

$$f_s(G) = \frac{1}{N} \sum_{i=1}^{N} \delta_{\theta_i s} \tag{8}$$

$f_0(0) = f_1(0) = f_6(0) = 1/3$ holds for $G = 0$ (see also the first snapshot of Fig. 2[4]).

Because each agent encounters with her 4 nearest neighbors n_g times during one generation, in each generation $(N/2 \times n_g \times 4)$ independent and simultaneous deterministic 2-person games occur. Fig. 2[4] shows snapshots of the spatio-temporal distribution of the three strategies for $n_g = 2$, while Fig. 3[4] shows snapshots with the same setting, but for $n_g = 3$. [4] [5]

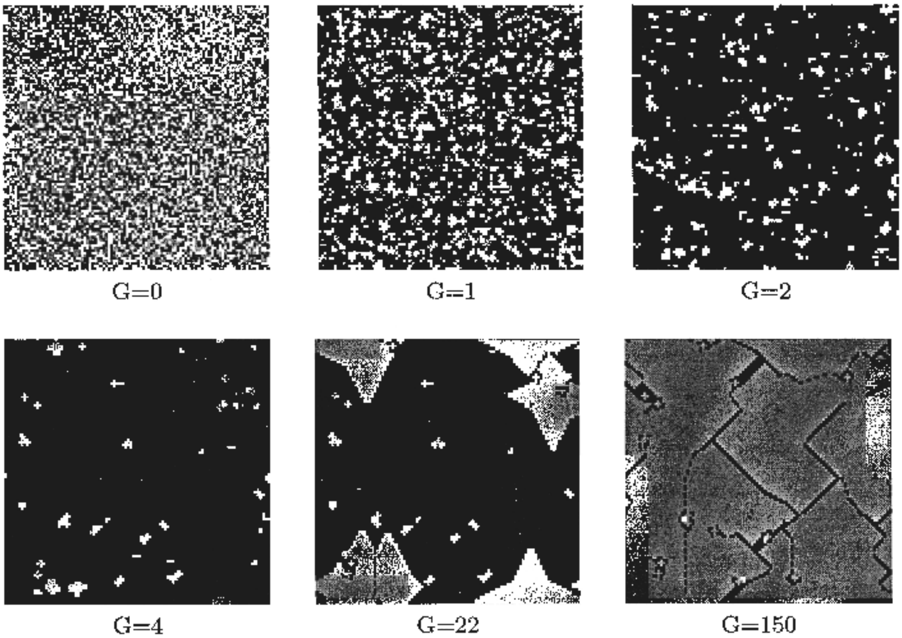

Fig. 2. Spatial-temporal distribution of three strategies: sD (black), sATFT (white), and gTFT (grey). Parameters: grid size: 100×100, payoff values: $T = 5$, $R = 3$, $P = 1$, $S = 0$, $n_g = 2$. [4]

For $n_g = 2$, we see from Fig. 2 that in the very beginning, i.e. in the first four generations, strategy sD grows very fast at the expense of sATFT and especially on gTFT. This can be also confirmed by looking at the global frequencies of each strategy (see left part of Fig. 4). Already for G=4, strategy sD is now the majority of the population – only a few agents playing gTFT and even fewer agents playing sATFT are left in some small clusters. Hence, for the next generation one might assume that the sD will take over the whole population. Interestingly, this is not the case. Instead, the global frequency of sD goes down while the frequency of gTFT starts to increase continuously until it reaches the majority. Only the frequency of sATFT remains at its very low value. On the spatial scale, this evolution is accompanied by a growth of

[4]Note that all spatial snapshots are original color figures, which could not be printed here. Therefore, we encourage the reader to download the color version from `http://www.ais.fraunhofer.de/~frank/p-wehia03.html`

[5]*Computer videos* of these simulations can be found at `http://www.ais.fraunhofer.de/~frank/spatial_game.html`

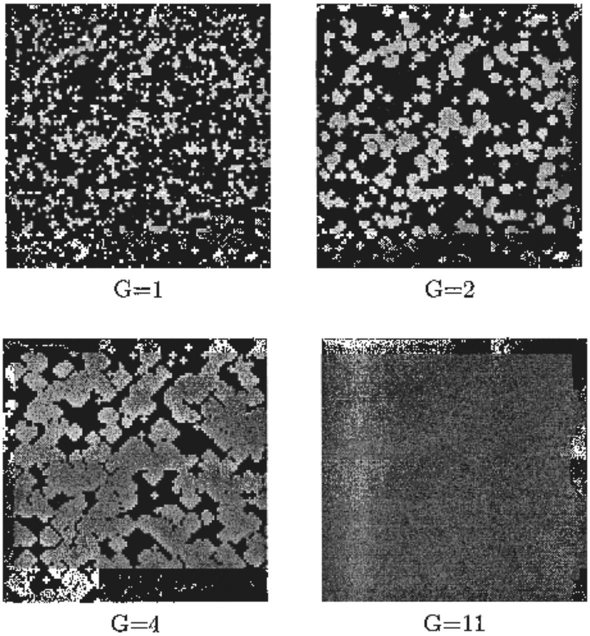

Fig. 3. Spatial-temporal distribution of three strategies: sD (black), sATFT (white), and gTFT (grey) for $n_g = 3$. The comparison with Fig. 2 elucidates the influence of n_g.[4]

domains of gTFT that are finally separated by only thin borders of agents playing sD (cf Fig. 2 for $G = 150$). The reasons for this kind of *takeover dynamics* will be explained later.

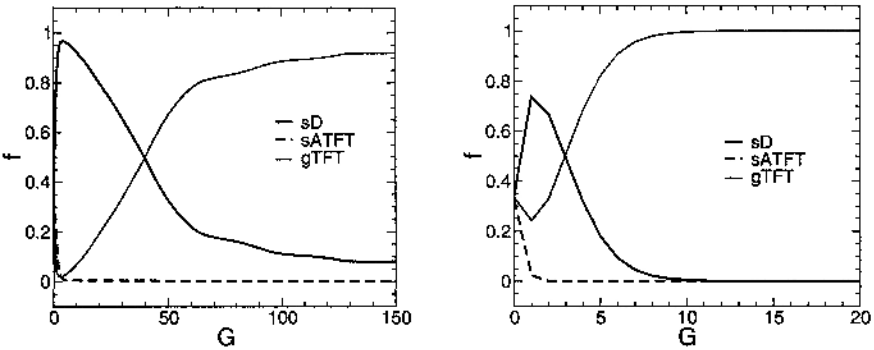

Fig. 4. Global frequencies $f_s(G)$ of the three strategies for $n_g = 2$ (left) and $n_g = 3$ (right). For the spatial distribution, see Fig. 2 and Fig. 3, respectively.

When increasing the number of encounters n_g from 2 to 3, we observe that the takeover of gTFT occurs much faster. Already for $G = 13$, it leads to a situation where *all agents* play gTFT, with no other strategy left. Hence, they will mutually cooperate. The fast takeover is only partly due to the fact that the total number of encounters during one generation has increased. The main reason is that for $n_g = 2$ agents playing sATFT are able to *locally block* the spreading of strategy gTFT, while this is not the case for $n_g = 3$. This is because of the dependence on n_g of both the agent's payoff and the local configuration of players.

To elucidate this, let us for the moment consider a population of only two strategies, namely gTFT and sATFT. For $n_g = 2$ there is only one configuration, shown in Fig. 5, where gTFT can invade sATFT because of the higher payoff. However, as can be seen in Fig. 5(right), after one generation the preconditions for further invasion have vanished, and thus the invasion stops.

The situation is more complicated in the presence of the third strategy sD [6], but it can be summarized, that agents playing sATFT in the early stage are responsible for the delayed takeover by gTFT. The crossover dynamics mentioned in Fig. 4 can be explained by the fact that only agents playing sD are able to invade both, sATFT and gTFT. Thus, their frequency increases strongly in the early stage. Once the number of agents playing sATFT is considerably reduced by sD, gTFT can spread. For $n_g = 3$, this situation is different in that there are more local configurations, where gTFT can invade sATFT. This in turn enhances the takeover of gTFT.

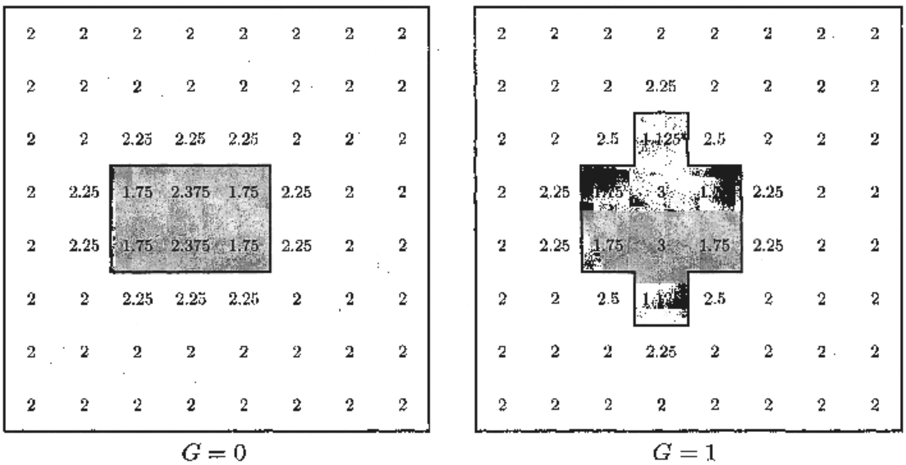

Fig. 5. (left) The smallest configuration of gTFT (grey) that is able to gain ground in sATFT (white) for $n_g = 2$. (right) The result of the invasion after one generation. The numbers show the payoffs of the respective agents within their neighborhood. The values $T = 5, R = 3, P = 1, S = 0$ have been used for this calculation.

In order to demonstrate this *pinning effect* on the global scale – i.e. agents playing sATFT *block* the spreading of agents playing gTFT – we carried out simulations with only two strategies, sATFT and gTFT. Fig. 6 shows spatial snapshots for a $n_g = 2$ game. As expected this time sATFT (white) is the stable majority after only three generations.

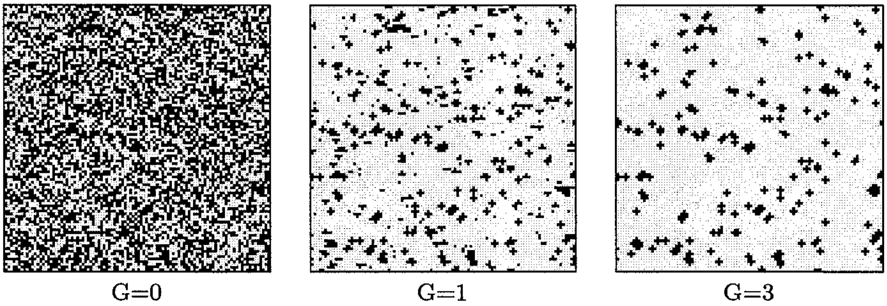

Fig. 6. Spatial snapshots of two strategies, sATFT (white) and gTFT (black). Parameters: 100×100 grid and $n_g = 2$. Initial conditions: random distribution of agents playing sATFT and gTFT.

5 Global Payoff Dynamics

In economics, utility maximizing agents should have an estimate of the probabilities of their expected payoff. In our model, however, the updating of strategies is behavioral instead of rational, i.e. the agent simply imitates the strategy of the best, this way trying to increase her own payoff. It is of interest, whether this simple local imitation strategy would also maximize the global utility. Thus, in the following we investigate the global payoff and the dynamics of the payoffs of the individual strategies.

The average payoff per agent $\bar{a}$ is defined as:

$$\bar{a} = \frac{1}{N}\sum_{i=1}^{N} a_i(\theta_i) = \sum_s f_s(G)\cdot\bar{a}_s\ ; \quad \bar{a}_s = \frac{\sum_i a_i(\theta_i)\delta_{\theta_i s}}{\sum_i \delta_{\theta_i s}} \tag{9}$$

where $f_s(G)$ is the total fraction of agents playing strategy s and $\bar{a}_s$ is the average payoff *per strategy*, shown in Fig. 7 for the different strategies.

We note that the payoffs per strategy for the 2-person games are always fixed dependent on n_g. However, the average payoff per strategy changes in the course of time mainly because the local configurations of agents playing different strategies change. For $n_g = 2$, we have the stable coexistence of all three strategies (cf Fig. 2 and Fig. 4 left), while for $n_g = 3$ only strategy gTFT survives (cf Fig. 3 and Fig. 4 right). Hence, in the latter case we find that the

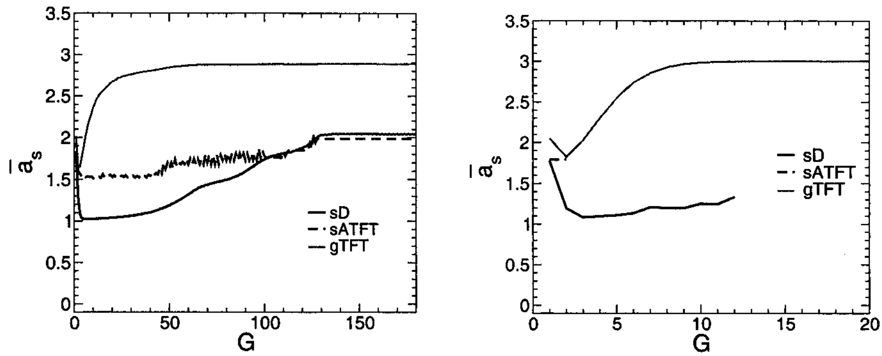

Fig. 7. Average payoff per strategy, $\bar{a}_s$, eq. (9), vs. time for $n_g = 2$ (left) and $n_g = 3$ (right)

average payoff of gTFT reaches a higher value than for $n_g = 2$, while in Fig. 7 the corresponding curves for the other strategies simply end, if one of these strategies vanishes.

The average global payoff is shown in Fig. 8 for different values of n_g. Obviously, the greater n_g, the faster the convergence towards a stationary global value, which is $\bar{a} = 3$ only in the "ideal case" of *complete cooperation.* As we have already noticed, for $n_g = 2$ there is a small number of defecting agents left playing either sD or sATFT, therefore the average global payoff is lower in this case.

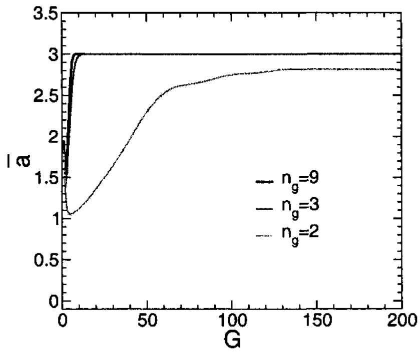

Fig. 8. Average global payoff $\bar{a}$, eq. (9), vs. time for different values of n_g.

6 Evolution of More Diverse Strategy Patterns

So far, we restricted the discussion to only three strategies. But, as explained in Sect. 2, the full strategy space consists of eight possible strategies, Tab. 1 – which in turn may result in much more complex spatial patterns. To elucidate this, in the following we present computer simulations with the full set of strategies.

As before, the initial distribution is chosen randomly, with equal initial frequencies of all eight strategies. Again, we concentrate on $n_g = 2$, and update rule eq. (7) is used. Fig. 9[4] and Fig. 10[4] show snapshots of the spatio-temporal evolution of the strategy distribution. Noteworthy, both runs start with the same setup (except a different randomization) – but, interestingly, the evolution occurs in a rather different manner, as a comparison between the two simulations demonstrates.[6] The evolution of the global frequencies of the respective simulations can be found at Fig. 11.

During the *early stage*, i.e. until $G = 100$ the dynamics of the two runs is rather similar: we observe that all generous strategies (i.e. agents starting with cooperation) die out, except gD (blue) – which later appears to be the global winner in both cases. The critical time window occurs at about $G = 100$, where in the first case the global frequency of gD increases further, until it reaches a *stationary* value of about 0.8, whereas in the second case this increase is turned into a descent, until a global (and *non-stationary*) frequency of about 0.55 is reached.

The different evolution during the critical stage results from the randomness in the update rule for equally best neighbors and from small deviations in the local configurations , where sC (yellow) plays a crucial role (cf Fig. 10 with Fig. 9 at $G = 100$). It seems that sC stabilizes the survival of sD, which can be also observed in Fig. 11. In the first case, both strategies dissapear, while in the second case they *both* survive. The local coexistence of both sC and sD can be seen in the snapshots of Fig. 10.

We emphasize that the coexistence of *four* strategies, i.e. gD, sTFT together with sD and sC, eventually results in a *non-stationary* long-term dynamics, rather different from the first case, where only two stategies, i.e. gD, sTFT, survive. Here, we find a (quasi-) stationary spatial pattern, where the local configuration consecutively flips between two states and thus only tiny changes in the global frequency occur.

The significantly different evolution of the spatial patterns for the same average setup suggests that the dynamics may have at least two different *attractors* that are reached, however, with different probability. From our computer simulations we deduced that the non-stationary pattern shown in Fig.

[6]The observant reader may notice that simulations are not exactly comparable since we have chosen the snapshots at different generations. This was mainly to show the most interesting patterns. But, in order to give an impression about the different evolution, we also provide *computer videos* of the full simulations, which can be found at `http://www.ais.fraunhofer.de/~frank/spatial_game.html`.

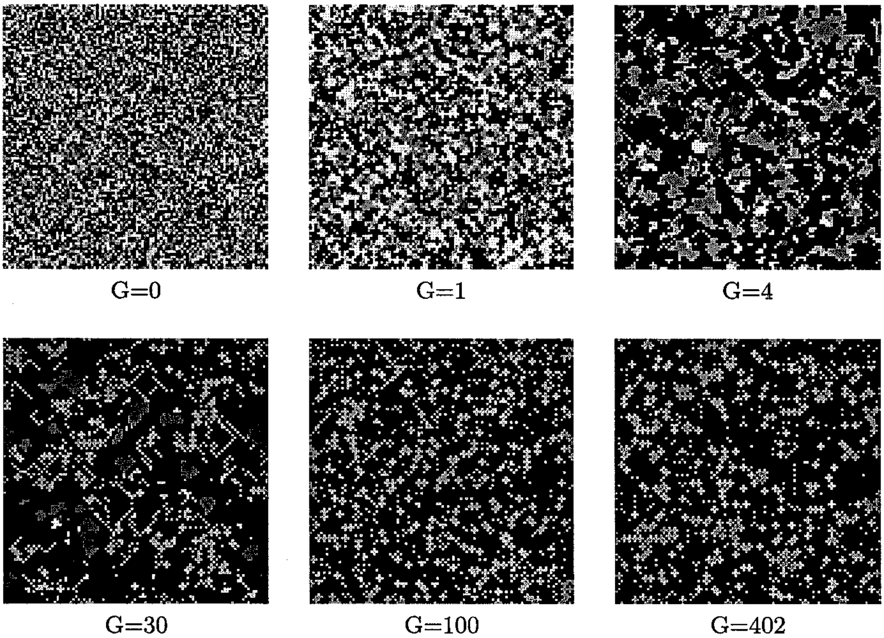

Fig. 9. Spatial-temporal distribution of 8 strategies: sD (black), sATFT (white), sTFT (green), sC (yellow), gD (blue), gATFT (pink), gTFT (red), gC (cyan). For parameters see Fig. 2. [4]

10 is not reached very often, e.g. the attractor size should be relatively small compared to the size of the attractor with the stationary pattern. This shall be investigated in a future paper.

Eventually, we mention an important difference between the games with three and with eight strategies. Whereas in the first case agents playing gTFT are in the majority, in the latter case they completely disappear – which is different from the observations by e.g. [1]. This means that an increase in *heterogeneity* in the agent's strategies, together with consideration of *local interaction*, results in a quite complex (sometimes non-stationary) dynamics in the IPD.

7 Conclusions

In this paper, we have investigated the evolution of *spatial heterogeneity* in agents' strategies. This heterogeneity originates from two different sources: (i) the *local* restriction of agent interaction to nearest neighbors, (ii) the consideration of a one-step memory that allows different responses to individual actions.

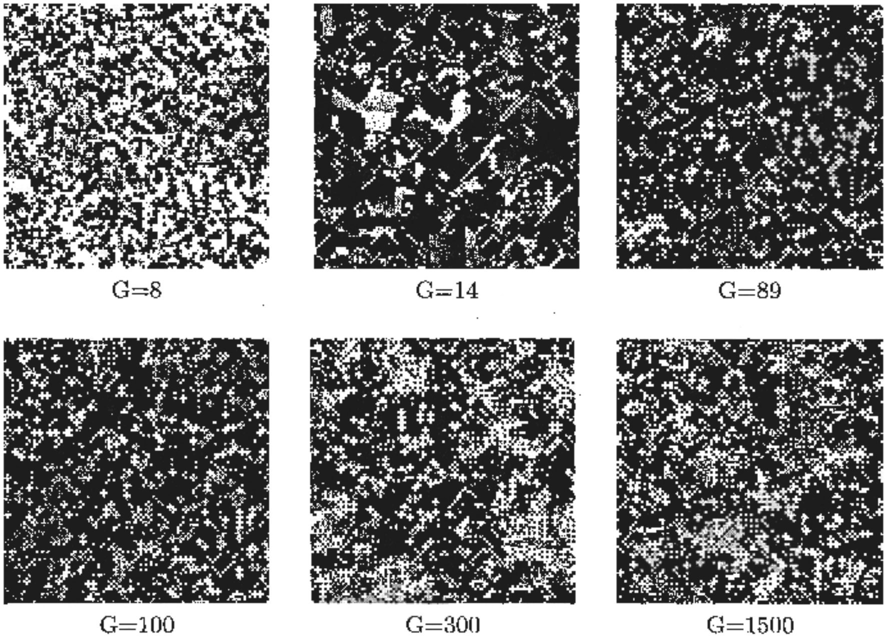

Fig. 10. Spatial-temporal distribution of 8 strategies: sD (black), sATFT (white), sTFT (green), sC (yellow), gD (blue), gATFT (pink), gTFT (red), gC (cyan). Same setup as in Fig. 9, but with a different initial randomization. [4]

As an illustrative example for the heterogeneous interaction, we have investigated the iterated Prisoner's Dilemma (IPD) game, that has been established as a key paradigm for investigating the evolution of cooperation, with many applications in politics, economy, and sociology.

In the first part of the paper, we have investigated the case of three strategies, concentrating mainly on the role of consecutive encounters between any two agents, n_g. We find that a critical value of n_g exists above which no coexistence between agents playing different strategies is observed. Hence, the most successful strategy, i.e. the one with the highest payoff in an iterated game, gTFT, is eventually adopted by all agents. This confirms the findings of [1] also for the spatial case. Below the critical n_g, we find a coexistence between cooperating and defecting agents, where the cooperators are the clear majority (playing gTFT), whereas the defectors play two different strategies, either sD or sATFT. In both cases, we observe that the share of gTFT in the early evolution drastically decreases before it eventually invades the whole agent population. We notice, however, that this picture holds only for a *random initial distribution* of the three strategies. It can be shown [12] that there are always specific initial distributions where gTFT fails to win.

In the second part of the paper, we have investigated the more complex case of eight strategies, with $n_g = 2$ fixed. Different from the case of three

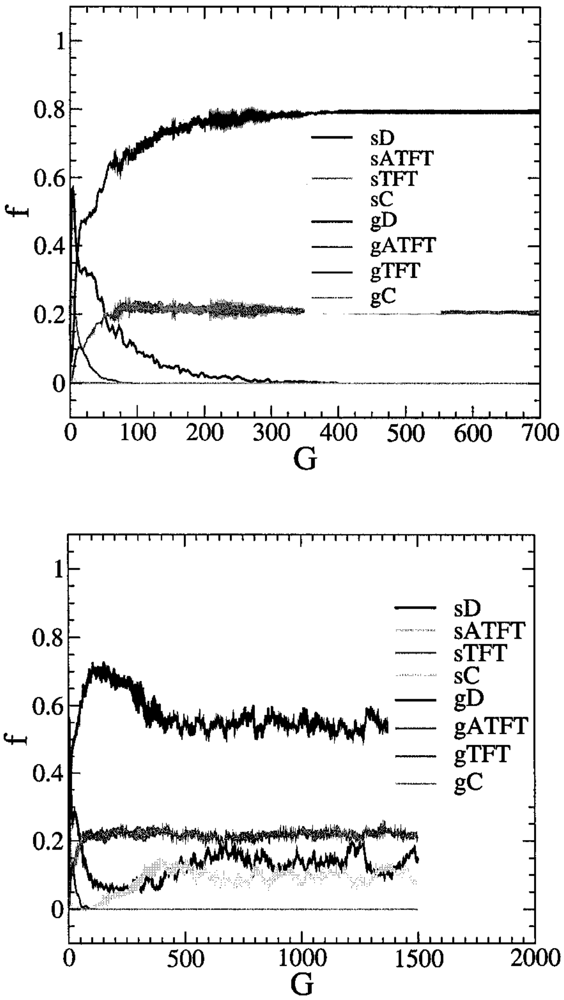

Fig. 11. Frequencies of the simulation shown in Fig. 9 (top) and Fig. 10 (bottom). The surviving strategies are (descending order of relative frequencies) gD, sTFT (top) and gD, sTFT, sD, sC (bottom)[4].

strategies which all coexist for $n_g = 2$, we observe the coexistence of either *two* or *four* strategies. In particular, the popular strategy gTFT does not survive, if initially eight strategies are considered. This is different from the known results in many alternative scenarios, where TFT usually wins. An interesting observation is the non-stationary long-term dynamics of the spatial game in the case of eight strategies (given the standard payoff matrix). Here, random deviations in the initial configuration may lead the global dynamics to different attractors. Thus, *local configurations* are of great influence.

References

1. Axelrod, R. (1984). *The Evolution of Cooperation.* New York: Basic Books.
2. Behera, L.; Schweitzer, F.; Mühlenbein, H. (2004). The Invasion of Cooperation in Heterogeneous Populations, I: Thresholds and Attractor Basins. *Journal of Mathematical Sociology* (to be submitted).
3. Durrett, R. (1999). Stochastic Spatial Models. *SIAM Review* **41(4)**, 677–718.
4. Lindgren, K.; Johansson, J. (2002). Coevolution of strategies in n-person Prisoner's Dilemma. In: J. Crutchfield; P. Schuster (eds.), *Evolutionary Dynamics: Exploring the Interplay of Selction, Accident, Neutrality, and Function*, pp. 341–360.
5. Lindgren, K.; Nordahl, M. G. (1994). Evolutionary dynamics of spatial games. *Physica D* **75**, 292–309.
6. Mach, R.; Schweitzer, F. (2004). Food-Web Representation of a Spatial IPD. (to be submitted).
7. Mühlenbein, H.; Höns, R. (2002). Stochastic analysis of cellular automata with applications to the voter model. *Advances in Complex Systems* **5(2)**, 301–337.
8. Nowak, M. A.; May, R. M. (1992). Evolutionary games and spatial chaos. *Nature* **359**, 826–829.
9. Nowak, M. A.; May, R. M. (1993). The spatial dilemmas of evolution. *International Journal of Bifurcation and Chaos* **3(1)**, 35–78.
10. Nowak, M. A.; Sigmund, K. (1992). Tit for tat in heterogeneous populations. *Nature* **355**, 250–253.
11. Schweitzer, F.; Behera, L.; Mühlenbein, H. (2002). Evolution of cooperation in a spatial prisoner's dilemma. *Advances in Complex Systems* **5(2)**, 269–300.
12. Mach, R.; Schweitzer, F. (2004). Adaptation of Cooperation: Random versus Local Interaction in a Spatial Agent Model. *J. of Artificial Societies and Social Simulation* (submitted).