The Impact of Group Reputation in Multiagent Environments*

Bastian Baranski, Thomas Bartz-Beielstein, Rüdiger Ehlers, Thusinthan Kajendran,
Björn Kosslers, Jörn Mehnen, Tomasz Polaszek, Ralf Reimholz, Jens Schmidt, Karlheinz Schmitt,
Danny Seis, Rafael Slodzinski, Simon Steeg, Nils Wiemann, Marc Zimmermann

Abstract

This paper presents results from extensive simulation studies on the iterated prisoner's dilemma. Two models were implemented: a nongroup model in order to study fundamental principles of cooperation and a model to imitate ethnocentrism. Some extensions of Axelrod's elementary model implemented individual reputation. We furthermore introduced group reputation to provide a more realistic scenario. In an environment with group reputation the behavior of one agent will affect the reputation of the whole group and vice-versa. While kind agents (e. g. those with a cooperative behavior) lose reputation when being in a group, in which defective strategies are more common, agents with defective behavior on the other hand benefit from a group with more cooperative strategies. We demonstrate that group reputation decreases cooperation with the in-group and increases cooperation with the out-group.

1 Introduction

The *iterated prisoner's dilemma* (IPD) is a game for two players, A and B, which is played for k rounds—while k is unknown to the players. In every round, each player has the option to cooperate or defect with the other player in order to get as many rewardas possible. There is a cost for cooperating but also an advantage if the other player cooperates as well(Table 1). In case of mutual cooperation A and B receive a *reward* R (3 points). Player A receives the *temptation* T (5 points) if A defects and B cooperates. Player A receives zero points, the so-called *sucker's payoff* S, if A cooperates while B defects. If both players defect, they receive the *punishment* P (1 point). Thus, the entries in the payoff matrix satisfy the following relationship: $T > R > P > S$. Furthermore the following inequality becomes true: $R > (T + S)/2$ (Axelrod, 1984). For very small and known k values, both players will probably defect, because this is best choice in the single-rounded, or non-iterated *prisoner's dilemma* (PD). This strategy is referred to as *always defect* (ALLD) in the literature. However, playing ALLD results in a lack of the opponent's confidence, making him less liable to cooperate and thus yielding less points than in case of mutual cooperation. The determination of an optimal strategy, i. e. to maximize the gain at minimum cost, is a difficult task and its solution depends on several factors. For example, can the players be sure to receive correct information about the opponent and that their perception of his strategy (cooperative or defective) is adequate or do they have to take into account a certain rate of noise? How often do they play against each other? And can they be sure to play a next round versus the same opponent?

The IPD model can be applied to several real-life situations. Our paper extends the application of models from Axelrod (1997) on agent-based simulations. Axelrod introduced the non-iterated prisoner's dilemma in an agent-based simulation. He analyzed how agents cooperate or defect against each other on a grid. The agent's decision whether to cooperate or to defect influences its *probability to reproduce* (PTR). The decision depends on parameters which are mutated during the game. Two parameters are responsible for the decision, one for the own group and one for every other group, resulting in a friend/foe, or ethnocentric view throughout the game.

Inspired by Axelrod's work our analysis examines effects and mechanisms influencing the evolution of cooperation. Cooperation is essential in many real-life situations—even among rivals. The cooperation of two rivaling companies is a well-known example: they cooperate to produce new goods, because the respective companies would not be able to accomplish this feat on their own. Another example for mutual cooperation aside from this rivalry context is the just-in-time service many suppliers provide, where they deliver their goods in a certain amount to a designated time.

Based thereon, we implemented reputation in Axelrod's elementary model. Axelrod & Hammond (2003) noted: "To be specific, the model shows that in-group favoritism can overcome egoism and dominate a population even in the absence of

---

*Department of Computer Science, Dortmund University, D-44221 Dortmund, Germany (phone: +49-231-755-7705; fax: +49-231-755-7740; The authors of this joint work can be reached via e-mail.

Table 1: Payoff matrix for a two-player IPD; S, P, R, T must satisfy $T > R > P > S$ and $R > (S + T)/2$

	B cooperates	B defects
A cooperates	R R	T S
A defects	S T	P P

[...] reputation [...]" This statement triggered our investigations. The aspect of reputation plays an important role in today's life, because having good reputation can often influence people's behavior significantly. An example from everyday life is the commercial online-marketplace eBay where each user planning to sell or buy items has to get an account with ebay. Each member also has a feedback score, which tells other users how reliable this member was up to now. The further a member's score is below 100%, the more wary other members will be of doing business with him or her. Marler & Evans (1996) and Zahavi & Zahavi (1997) described the influence of renown on human (and also animal) behavior. By altering one's own behavior in the sense of being more altruistic one can increase the chance of being the subject of further altruistic acts by other individuals.

To get reliable empirical data, our team of authors organized a Dortmund IPD tournament in 2005, with more than 50 players participating, and analyzed the various strategies encountered there, more information can be accessed under the given URL (PG 474, 2005). By having the agent remember the decisions in the past we were able to mirror the circumstances in the conventional IPD for the agents in Axelrod's simulation, where the agents were confronted with the circumstances of the PD. The tournament and the simulation software used in this paper can be used for further experiments or to reproduce results presented in this study. The software we designed for this study is available as freeware (PG 474, 2005).

In order to establish a model to examine the working principles and the impact of reputation, we equipped the agents with the knowledge from preceding rounds and the results from simulations based on these modified setups are compared to results from Axelrod's original experiment. Similar to approaches described in Yao & Darwen (2000), Chess (1988), and Fogel (1993), co-evolutionary learning was used in the IPD to determine proper settings for different parameters. We investigated the evolution of cooperation and the effects of reputation thereon.

Axelrod's model was extended in two different ways. In the first model M_1, we removed the affliations to the different groups from the agents, i. e. all agents belonged to the same group, thus played for themself alone because there were no other groups. Furthermore the agents had the ability to remember the behavior of the other agents, including information on how each of the other agents had played so far, not only against a specific agent but against all agents as a whole. This kind of memory enabled the simulation environment to give each agent a reputation, based solely on his behavior during the current game. In turn this reputation influenced the behavior of the other agents towards the respective agent, so a good reputation of an agent, i. e., one who had often cooperated, resulted in a more cooperative behavior of the others towards himself, whereas a bad reputation had the opposite effect.

Our second model M_2, used an enhanced representation of reputation. Here, agents belonged to different groups. Moreover the agents had the ability to differentiate all existing groups on the grid. They had the opportunity to work together with specific other groups. Finally we combined this differentiating view with our concept of reputation to observe the resulting effects.

Our reputation model extends other implementations. Yao & Darwen (2000) used a tag to describe how an agent had behaved so far, but this tag was initialized with 0 and could take on during the game only the values "-1" and "1", while our reputation score could take on every value ranging from 0 to 1 in 0.1 steps. Nowak & Sigmund (1998) used an image score, which was initialized with 0 and increased or decreased by one depending on the decisions the agent has made so far. This is in contrast to our model that stores the percentage of cooperation for each agent. Schenk (1995) used another reputation scheme: he equipped each agent with memory to collect information on how this agent was treated so far. However, the agent's behavior, and not its treatment, is memorized in our model to determine its current reputation. We equipped each agent with a memory of how all other agents played so far.

This paper is organized as follows: The reputation models are introduced in Sect. II. Section III describes results from the experiments, Sect. IV presents an analysis. Before an outlook is given in Sect. VI, conlusions are drawn in Sect. V.

2 Definition

Axelrod (1997) tried to discover important factors influencing the evolution of cooperation and ethnocentrism. In his model, agents were placed on a grid and played against each other. Successful agents receive a higher rate to bring forth offspring.

Table 2: Parameter settings in basic environment. $\mathcal{G}_{\text{size}}$ = side length of the grid, γ = cost to cooperate, PTR = probability to reproduce, δ = amount of PTR gained when opponent cooperated, p_{mut} = mutation rate, p_{die} = risk to die. From *Evolution of Ethnocentric Behavior* Axelrod & Hammond (2003)

$\mathcal{G}_{\text{size}}$	γ	PTR	δ	p_{mut}	p_{die}
50	1.0	12.0	3.0	0.5	10.0

Our goal was to define a model that enables agents to build up a kind of reputation, i. e., a way in which they knew how well or badly their opponent had behaved so far. Knowledge gained by the agents so far was used for this purpose. Each agent has to play several rounds of the PD against several opponents. Our model of reputation used the knowledge gained during this games to create a kind of renown for each agent or, as explained later on, for each group. Furthermore each agent used a parameter which determines the next move. This decision depends on the opponent's reputation. Two variants of this reputation model were implemented. Both variants extend the basic environment from Axelrod (1997) which is explained in Sect. 2.1. Model M_1 is a less complicated extension, in which every agent is reliant on himself only as there are no longer any groups. Model M_2 implements different groups and uses a modified reputation scheme.

2.1 Basic environment

The basic environment bases upon the original experiments (Axelrod, 1997). We use a grid $\mathcal{G}$ on which agents from different groups are randomly placed. This grid is a swap around square with a certain side length, overlapping at the borders in a way that each cell on the left border is a neighbor to the corresponding cell on the right border. In the same way upper and bottom cells are treated. The *neighborhood* of an agent at position p_i is defined as the set of grid points that can be reached from p_i with one vertical or horizontal step.

Each cycle in the simulation consists of four phases: immigration, interaction, reproduction, and death. In the immigration phase an agent with stochastic characteristics is placed randomly on a free cell and every agent receives a certain *probability to reproduce* (PTR). Neighbors (horizontal or vertical) play one round of the Prisoner's Dilemma against each other. Their decision to cooperate or to defect depends on the opponents' reputations.

Let γ denote a constant which models the *costs of cooperation*. Each agent has to pay a certain amount γ from his PTR when cooperating, while paying none when defecting: Consider two agents, a_i and a_j. The PTR value of an agent a_i who cooperates is reduced by γ, if a_j defects. If the opponents a_j cooperates, a_i receives a bonus δ increasing his probability to reproduce. So if both opponents cooperate, each receive a bonus that increases their PTR values.

In the reproduction phase offspring are placed randomly on a directly neighboring cell. Parameters are mutated with a so-called *mutation rate* p_{mut}. The parameters which are exposed to a possible mutation are the group of the agent as well as the parameters concerning the reputation.

In the death phase every agent which was at the beginning of this turn on the board and each immigrated agent (except the offspring) have a fixed *risk of dying* p_{die}. After the dying phase the PTR of each agent is reset to its initial value. To enable a comparison with other models, the model parameters have been chosen as proposed by Axelrod & Hammond (2003)). The experimental setup is summarized in Table 2.

2.2 Nongroup model M_1

Reputation can be interpreted as "the overall quality as seen or judged by people in general." Thus, reputation requires memory. To introduce reputation, agents were equipped with memory. We define reputation according to the number of cooperations this agent has chosen in the PD game so far. In this first model every agent has an individual reputation, meaning every agent's decisions influences his reputation directly. For this purpose different groups were eliminated, so each agent played only for himself, not his group. The following two parameters were introduced in the M_1 model: The *reputation* ρ_{indi} is defined as the relation between the number of cooperations and the number of total moves m an agent a had so far

$$\rho_{\text{indi}}(a) = \text{round}\,(c(a)/m(a))/10. \tag{1}$$

ρ_{indi} ranges from 0.0 to 1.0 in 0.1 steps, where 1.0 means total cooperation, or, said in another way, the agent had cooperated in nearly every move up to now, while 0.0 means this agent had never cooperated during the whole game so far.

The second parameter, the so-called *reputation threshold* τ_{indi}, determines the lowest level of reputation which is necessary to cooperate with another agent. Consider agent a_i with $\rho_{\text{indi}} = 0.3$ and agent a_j with $\tau_{\text{indi}} = 0.6$: a_j does not cooperate with a_i, because $\rho_{\text{indi}}(a_i) < \tau_{\text{indi}}(a_j)$. Agent a_j did not cooperate with a_i because of a_i's bad reputation.

Table 3: Five different configurations for initial ρ_{indi}, ρ_{group} and τ_{indi}

ρ_{indi} ρ_{group}	τ_{indi}	configuration
0	0	κ_0
0	1	κ_1
1	0	κ_2
1	1	κ_3
random $\in [0;1]$	random $\in [0;1]$	κ_4

Note, τ_{indi} is only altered via mutation. These two parameters were randomly initialized for every agent and they also were subject to a possible mutation while passing on to a child. We tried to verify our results by using different initializations, like high expectations (τ_{indi} = 1) and a low own reputation (ρ_{indi} = 0) and vice versa, the results will be presented in the next chapter.

2.3 Multi-group model M_2

The second model analyzes different groups and the ability to differentiate between them. The reputation from model M_1 implemented for each group, meaning that the average number of cooperations of every agent belonging to the group was considered. This parameter is referred to as *groupReputation* ρ_{group}. In this way the individual decision of one agent only contributed partly to the reputation of the whole group. The parameter ρ_{indi} is kept, only the determination of the reputation differed from the first model: In the first model each agent's reputation was calculated on the basis of his behavior only, while in the second model the reputation (of the whole group) was averaged over each reputation of all agents of this group. Let n_j denote the number of agents in the jth group (with $j \in [1;m]$) ρ_{group} can be calculated as

$$\rho_{\text{group}}(j) = \frac{1}{n_j} \sum_{i=1}^{n_j} \rho_{i,j}. \tag{2}$$

The parameter τ_{indi} was extended, so that there was an own τ_{indi} for every agent of every occurring group: $\tau_{i,j} \in [0.0; 1.0]$.

Ethnocentrism was implemented in model M_2 as follows: An *ethnocentric agent* a_{ethno} i is an agent that defected at least once against an agent b of another group j (against an out-group, after Axelrod). At the same time a's $\tau_{\text{indi}}(a_{\cdot,i})$ against this group has to be less than 0.6, i. e.

$$a_{\text{ethno}}(i) = \begin{cases} 1 & \tau_{\text{indi}}(a_{\cdot,i}) \leq 0.6 \text{ and } d \geq 1 \\ 0 & \text{otherwise,} \end{cases}$$

where d denotes the number of defections against outgroups. While playing against other agents from the own group (in-group), agents have to fulfil the following criterion: His τ_{indi} concerning his own group must have been lower or equal to the τ_{group} of his group. This means, the agent cooperated with his own group every time he played against another agent of his own group.

Obviously, M_1 does not model ethnocentrism because it implements no groups. Each simulation contained 2000 generations and we ran each simulation 10 times, using not only the average results over the last 100 generations for our analysis but also the average over all 2000 generations. In addition we also tested model M_2 regarding more groups (6 and 8 groups). Furthermore we examined which effects would take place when the costs (in PTR) for cooperating are lower (0) and higher (2 and 3) in order to validate our results. To enable a comparison of our results to the results from Axelrod & Hammond (2003), we produced the same tables.

3 Results

3.1 Individual reputation in a non-group-environment

In an environment in which all agents belong to the same group, meaning there is no distinction between groups at all, every agent possesses an own reputation ρ_{indi} and a single reputation threshold τ_{indi} against all other agents. ρ_{indi} is built up during the past rounds and its initialization was varied in our experiments. The τ_{indi}, which was also initialized with different values, stayed fix during a run and mutated only if inherited to offsprings. The following parameters were modified: ρ_{indi}, τ_{indi}, and the costs to cooperate γ, which are subtracted from the PTR of the cooperating agent. The different parameter settings for ρ_{indi} and τ_{indi} are shown in Table 3. In addition to the values in this table the costs to cooperate γ were also modified. According to the original experiments in Axelrod & Hammond (2003), γ was set to 0, 1, 2 and 3. Lower values, e. g. 0 or 1, represent a friendly environment, higher values, e. g. 2 or 3, represent more hostile environments. Regarding the amount of help gained, which means a value of 3 to be added to the PTR, when the opponent had cooperated, there would be little or no benefit in cooperating

Table 4: Comparing the percentage of behavior that is cooperation with same type to results from Axelrod & Hammond (2003). $G = 50$, $\gamma = 1$, PTR $= 12.0$, $\delta = 3.0$, $p_{\text{mut}} = 0.5$, $p_{\text{die}} = 10.0$, 1 group, see Table 2 for explanation. $\mathcal{C}_{\text{same}}$ is the absolute value of percentage of behavior that is cooperation with same type over whole game, $\mathcal{C}_{\text{same}}^{\text{last 100}}$ is the same like $\mathcal{C}_{\text{same}}$ but over the last 100 rounds only, $\Delta\mathcal{C}_{\text{same}}$ and $\Delta\mathcal{C}_{\text{same}}^{\text{last 100}}$ are the relative differences to the values from Axelrod & Hammond (2003)

configuration	κ_0	κ_1	κ_2	κ_3	κ_4
$\mathcal{C}_{\text{same}}$	97.41	75.63	97.65	75.44	85.34
$\Delta\mathcal{C}_{\text{same}}$	+20.81	-0.97	**+21.05**	**-1.16**	+8.74
$\mathcal{C}_{\text{same}}^{\text{last 100}}$	97.39	93.07	97.49	94	94.25
$\Delta\mathcal{C}_{\text{same}}^{\text{last 100}}$	+22.09	**+17.77**	+22.19	**+18.7**	+18.95

Table 5: Comparing the percentage of behavior that is cooperation with same type with results from Axelrod & Hammond (2003), same as in table 4 but with cost doubled to $\gamma = 2$

configuration	κ_0	κ_1	κ_2	κ_3	κ_4
$\mathcal{C}_{\text{same}}$	94.64	37.08	95.34	39.49	67.63
$\Delta\mathcal{C}_{\text{same}}$	+71.24	**+13.68**	**+71.94**	+16.09	+44.23
$\mathcal{C}_{\text{same}}^{\text{last 100}}$	94.93	72.37	94.48	74.34	85.97
$\Delta\mathcal{C}_{\text{same}}^{\text{last 100}}$	+80.93	**+58.37**	+80.48	**+60.34**	+71.97

in a hostile environment. Each of the 20 different settings were repeated 10 times. The mean value from these runs were taken to compare our results with the ones presented in Axelrod & Hammond (2003).

We compare the results of selected experiments of our own with the according results from the experiments Axelrod and Hammond conducted. Interesting values were printed in boldface. Table 4 shows different values of the percentage of behavior: Cooperation with the same group depends on the five different configurations for initial ρ_{indi} and τ_{indi} introduced in Table 3. $\mathcal{C}_{\text{same}}$ represents the percentage over all 2000 generations and $\mathcal{C}_{\text{same}}^{\text{last 100}}$ the percentage only over the last 100 generations. For both values the differences $\Delta\mathcal{C}_{\text{same}}$ and $\Delta\mathcal{C}_{\text{same}}^{\text{last 100}}$ to the results from Axelrod & Hammond (2003) are shown. In Table 4 results from the experiment with $\gamma = 1$ are compared to experiment no. 117 from Axelrod & Hammond (2003). In Table 5 results from the experiment with $\gamma = 2$ is compared to experiment no. 119 of Axelrod & Hammond (2003).

The comparison in Table 4 demonstrates how the ratio of cooperation within the same group (in this model there was only one group) varied only insignificantly from the results in Axelrod & Hammond (2003), when unfavorable initial values for ρ_{indi} and τ_{indi} are chosen. Beneficial starting values however boosted the increase of cooperation in a significant way. Even a random initialization of ρ_{indi} and τ_{indi} shows a significant increase in the ratio between the number of cooperations and the number of overall interactions in the whole game.

Table 5 shows the positive effect of reputation on cooperation. In this configuration the costs to cooperate γ were doubled, thus the environment was modeled in a more hostile way. Despite this unfavorable starting-point there still is a larger ratio of cooperation than in Axelrod & Hammond (2003), even in runs, when the initial configuration of ρ_{indi} and τ_{indi} were adversarial for cooperation.

3.2 Group-reputation in a group-environment

In this environment agents also possessed a kind of reputation. Each group had a reputation value ρ_{group}, which was determined by the decisions of each group member. Furthermore each agent possessed a reputation threshold τ_{group} for each group on the grid. The starting values for ρ_{group}, γ, and τ_{group} for each other group (or color) were again varied, see Table 3. In addition the number of groups was varied, too. We conducted experiments with four, six and eight groups. Our results were compared to the results from Axelrod & Hammond (2003).

In Table 6 the ratio $\mathcal{C}_{\text{same}}$ (and $\mathcal{C}_{\text{same}}^{\text{last 100}}$ respectively for the last 100 generations) between the number of cooperations and the number of total interactions between agents of the same group again was compared to the values found in Axelrod & Hammond (2003). In this case results were always lower than the results of the corresponding experiments. This means, that the number of cooperations with agents of the same group decreased in our experiments. The other values depicted in Table 6 are percentage of behavior that is defection with different type $\mathcal{D}_{\text{diff}}$ and $\mathcal{D}_{\text{diff}}^{\text{last 100}}$, which is the same but for the last 100 generations only. Compared to results from experiment no. 104 in Axelrod & Hammond (2003) it is obvious that the amount of defections with agents of a different color decreased massively. Thus there is less defection on the grid.

In Table 7 we compared the results from experiment no. 109 in Axelrod & Hammond (2003) to results from our simulations. The γ value was doubled again. The influence of reputation was high enough to reach a slightly better result than the corresponding experiment in Axelrod & Hammond (2003), even if only in configurations κ_0 and κ_2. These configurations are

Table 6: Comparing the percentage of behavior that is cooperation with same type and of behavior that is defection with different type to results from Axelrod & Hammond (2003). The same settings as in Table 4 were used, but four groups have been simulated. $\mathcal{C}_{\text{same}}$ is the absolute value of percentage of behavior that is cooperation with same type over whole game, $\mathcal{C}_{\text{same}}^{\text{last 100}}$ is similar to $\mathcal{C}_{\text{same}}$ but over the last 100 rounds only, $\mathcal{D}_{\text{diff}}$ is the absolute value of percentage of behavior that is defection with different type over whole game, $\mathcal{D}_{\text{diff}}^{\text{last 100}}$ is the same but over the last 100 rounds only, $\Delta\mathcal{C}_{\text{same}}$ and $\Delta\mathcal{C}_{\text{same}}^{\text{last 100}}$ are the differences relative to the values of Axelrod & Hammond (2003), $\Delta\mathcal{D}_{\text{diff}}$ and $\Delta\mathcal{D}_{\text{diff}}^{\text{last 100}}$, respectively

configuration	κ_0	κ_1	κ_2	κ_3	κ_4
$\mathcal{C}_{\text{same}}$	70.91	49.99	71.49	62.77	73.96
$\Delta\mathcal{C}_{\text{same}}$	-17.39	**-38.31**	**-16.81**	-25.53	-14.34
$\mathcal{C}_{\text{same}}^{\text{last 100}}$	70.35	77.76	70.84	68.05	77.06
$\Delta\mathcal{C}_{\text{same}}^{\text{last 100}}$	-19.45	**-12.04**	**-18.96**	-21.75	-12.74
$\mathcal{D}_{\text{diff}}$	3.61	11.50	0.12	6.68	6.69
$\Delta\mathcal{D}_{\text{diff}}$	-72.39	**-64.50**	**-75.88**	-69.32	-69.31
$\mathcal{D}_{\text{diff}}^{\text{last 100}}$	6.14	7.57	0.50	9.49	8.93
$\Delta\mathcal{D}_{\text{diff}}^{\text{last 100}}$	-77.59	**-76.15**	**-82.23**	-74.24	-74.80

Table 7: Same as in table 6 but with costs doubled to $\gamma = 2$

configuration	κ_0	κ_1	κ_2	κ_3	κ_4
$\mathcal{C}_{\text{same}}$	70.82	2.66	73.29	33.20	58.16
$\Delta\mathcal{C}_{\text{same}}$	+1.32	**-66.84**	**+3.79**	-36.3	-11.34
$\mathcal{C}_{\text{same}}^{\text{last 100}}$	66.42	4.58	73.08	19.35	55.11
$\Delta\mathcal{C}_{\text{same}}^{\text{last 100}}$	-2.98	-64.82	+3.68	-50.05	-14.29
$\mathcal{D}_{\text{diff}}$	5.76	20.86	0.31	13.17	12.44
$\Delta\mathcal{D}_{\text{diff}}$	-78.54	**-63.44**	**-83.99**	-71.13	-71.86
$\mathcal{D}_{\text{diff}}^{\text{last 100}}$	9.75	21.59	0.84	19.51	14.01
$\Delta\mathcal{D}_{\text{diff}}^{\text{last 100}}$	-79.01	**-67.17**	**-87.92**	-69.25	-74.75

beneficial for cooperation. The $\mathcal{D}_{\text{diff}}$ none the less is much lower than the one in Axelrod & Hammond (2003), showing that even under harsh conditions (i. e. $\gamma = 2$) there is less defection due to the factor of cooperation.

Increasing the number of groups has a negative effect on the behavior towards the own group. In Table 8 we compare the model of reputation with eight groups to the corresponding experiment in Axelrod & Hammond (2003). Independently of the quality of the initial configurations of ρ_{group} and τ_{group} the results are inferior regarding the amount of cooperation with the own group. A comparison of different $\mathcal{D}_{\text{diff}}$ values demonstrates the impact of reputation on defection.

3.3 Which factors influence the impact of reputation the most

Our experiments showed, that the above defined influence of reputation towards evolution of cooperation are steady, with external factors like configurations κ, cost γ, or number of groups m only changing the amplitude of differences but not the trend of decreasing ethnocentrism and thus increasing cooperation between the agents. κ_2 is a noteworthy exception concerning the variable γ. As can be seen in Fig. 1, not even γ, which defines the harshness of the environment, and therefore is an indicator for the amount of cooperative behavior, influences the steady high value for κ_2. Figure 1 shows the percentage of cooperation with in-group $\mathcal{C}_{\text{same}}$, it should be annotated, that this result is coherent with all other statistically surveyed values, i. e. $\mathcal{D}_{\text{diff}}$,

Table 8: Same as in table 6 but with number of groups doubled to 8

configuration	κ_0	κ_1	κ_2	κ_3	κ_4
$\mathcal{C}_{\text{same}}$	66.44	55.56	66.68	60.05	69.14
$\Delta\mathcal{C}_{\text{same}}$	-23.76	**-34.64**	**-23.52**	-30.15	-21.06
$\mathcal{C}_{\text{same}}^{\text{last 100}}$	65.71	77.85	64.73	74.96	69.29
$\Delta\mathcal{C}_{\text{same}}^{\text{last 100}}$	-24.79	**-12.65**	-25.77	**-15.54**	-21.21
$\mathcal{D}_{\text{diff}}$	3.88	13.07	0.33	10.76	8.66
$\Delta\mathcal{D}_{\text{diff}}$	-73.41	**-64.22**	**-76.95**	-66.53	-68.63
$\mathcal{D}_{\text{diff}}^{\text{last 100}}$	5.29	8.87	0.87	9.01	8.86
$\Delta\mathcal{D}_{\text{diff}}^{\text{last 100}}$	-79.95	**-76.36**	**-84.37**	-76.22	-76.37

$\mathcal{C}_{\text{same}}^{\text{last 100}}$, and $\mathcal{D}_{\text{diff}}^{\text{last 100}}$, respectively.

4 Analysis

4.1 Analysing model M_1, one-group reputation

First of all we will divide the configurations from Table 3 into two groups: the more positive starting points κ_0 and κ_2 regarding the initial values of ρ_{indi} and of τ_{indi} will henceforth be referred to as χ_{pos}. The worse starting points κ_1 and κ_3 will be referred to as χ_{neg}. This arrangement will be held during the analysis of both models.

One might conclude that under χ_{neg} conditions reputation reduces the number of cooperations compared to Axelrod (where $\Delta\mathcal{C}_{\text{same}}$ is negative), see Table 4. However, regarding only the values taken over the last 100 rounds clearly shows that our model of reputation requires some time to improve the amount of cooperation. The $\mathcal{C}_{\text{same}}^{\text{last 100}}$ value again is higher than in Axelrod's model; meaning that in our model of reputation under the bad starting conditions χ_{neg} there is in the beginning less cooperation compared to Axelrod's simulation—but the increase in cooperation is higher, resulting in a greater amount of cooperation towards the end of the simulation.

Regarding χ_{pos} displays the effect of reputation compared to the results from Axelrod. Here also, especially towards the end of our simulation, the increase in cooperation is obvious and noticeably higher than in Axelrod's simulation. Considering this it is very likely that in longer simulations the amount of cooperation would be higher due to reputation and it's long-term effect.

Table 5 displays the same settings and comparisons as shown in Table 4 with doubled costs γ. Here the effects of reputation are even more obvious. While higher costs represent a more hostile environment in itself, even under bad starting conditions χ_{neg} is a higher amount of cooperation compared to Axelrod's results. Towards the end of the simulation this increase is even higher, thus showing that our model of reputation to a certain extent can compensate several effects which would otherwise decrease the amount of cooperation. This is shown also in a comparison between both Table 4 and Table 5. While the values of $\mathcal{C}_{\text{same}}$ under χ_{neg} in Table 5 each is approximately one half of the corresponding values in Table 4, at the end of the simulation ($\mathcal{C}_{\text{same}}^{\text{last 100}}$) the difference is considerably smaller, proving that there was a compensation during the simulation.

Summarzing, we can state that on the one hand reputation can compensate—at least to a certain extent—the negative factors in our experiment. On the other hand reputation has an overall positive effect on the amount and on the increase of cooperation.

4.2 Analyzing model M_2, multi-group reputation

While comparing Table 6 and 8 one striking fact is obvious: the number of groups does not influence the level of cooperation. The values in Table 8 do not vary much compared to the corresponding values of Table 6. Regarding the level of defection with other groups in Table 6 ($\mathcal{D}_{\text{diff}}$) displays a high amount of cooperation with the out-group's while $\mathcal{C}_{\text{same}}$ displays a somewhat smaller amount of cooperation with the in-group. Logically consistent does the number of out-group's have no or only a small effect on the amount of cooperation overall and on the ratio of cooperation. The reasons for the relatively high amount of defection with the in-group under these conditions will be subject of further simulation studies.

In Table 7 the amount of cooperation under κ_2 is highest ($\mathcal{C}_{\text{same}}$ high, $\mathcal{D}_{\text{diff}}$ small) and only marginally smaller than the corresponding values in Table 6. Even though the costs γ are doubled (Table 7) the positive effect of κ_2 is enough to compensate this negative influence.

Simulating κ_1 leads to a significantly lower amount of cooperation than in every other configuration, especially with the in-group. This effect is caused by the starting values of ρ_{group} and τ_{group}. Starting with $\rho_{\text{group}} = 0$ and $\tau_{\text{group}} = 1$ means that all agents at first did not cooperate with one another. "They believed the worst of others and did nothing to improve their own reputation (or the reputation of their group)" would be a good description of their behavior.

4.3 Analysis using regression trees

To quantify the level of cooperation, we introduced the following measure: The *cooperation level* Y is the percentage of agents with $\rho_{\text{indi}} \geq 0.6$.

$$Y = \frac{1}{mT} \sum_{t=1}^{T} \left\{ \sum_{j=1}^{m} \left(\sum_{i=1}^{n_t} x_{ij} \right) \frac{1}{n_t} \right\}, \tag{3}$$

where

$$x_{i,j} = \begin{cases} 1 & \rho_{i,j} \geq 0,6 \\ 0 & \rho_{i,j} < 0,6. \end{cases} \tag{4}$$

Tools from explorative data analysis and computational statistic were used to screen out important factors of the simulation models. Design plots vary one factor while averaging the response over the other factors. Model M_2 has three important factors: cost, configuration, and color. The design plots from Fig. 2 depict their influence on Y. Therneau & Atkinson (1997) describe regression trees as flexible non-parametric tools for screening variables. The same data as in Fig. 2 have been used to generate the tree in Fig. 3.

When looking closer at the regression trees, one can see that configuration κ_2 causes a high Y-value, not depending on the parameter γ. Even in a hostile environment with $\gamma = 3$, Y stayed over 90%. A possible explanation is rooted in the start values of ρ_{indi} and τ_{indi}. New agents get equipped with a ρ_{indi} of 1 and other agents will cooperate when interacting with a τ_{indi} of 0. It seems probable that most agents and their children, who inherit reputation-level and cooperation-threshold do not live long enough to gain a bad reputation. So even in the configuration κ_2, where an agent would gain a maximum of 3 points to his PTR while paying 3 points every time he cooperates, agents tend to cooperate with their neighbors on the grid, caused by their low τ_{indi} and the high ρ_{indi} of any other agent.

For the configurations κ_1 and κ_3, the value of γ is decisive. The higher γ, the lower the value of Y gets, which was quite predictable. In both configurations agents got a τ_{indi} of 1 when placed on the grid and it obviously takes some time to lower the threshold. The percentage of cooperations in κ_1 is lower than that in κ_3, which is caused by the value of ρ_{indi} (0 in κ_1, 1 in κ_3).

In experiments with configuration κ_0, which works with start values of $\rho_{\text{indi}} = \tau_{\text{indi}} = 0$ and thus is comparable to κ_3 with $\rho_{\text{indi}} = \tau_{\text{indi}} = 1$, the rate of cooperative agents receives a significantly higher value even for high γ values, being not independent from the number of groups on the grid (experiments using κ_0 with $\gamma = 3$ produced results of 10%, 17% and 26% with 4, 6 and 8 groups, while κ_3 results were all about 18%-19% for experiments with a different number of groups). This is even more surprising if one looks at the conditions which lead to the value of Y: in κ_0, new agents do not get summed up into Y because of their low default value of ρ_{indi}. It may take some generations for agents to get a ρ_{indi} over 0,6 (which is essential for a high group reputation), but it's easier for agents in κ_0 to get into cooperative interaction, because it's easier for an agent to gain ρ_{indi} than it's to lower the τ_{indi}-Parameter, which is just influenced by mutation, not by interactive behavior.

Further experiments will be performed to gain knowledge about possible reasons for the high values in κ_0 and interaction with the number of groups.

5 Summary and Conclusions

The conclusion is divided up into several theses and ordered according to the underlying model.

5.1 Individual reputation in a multi-agent non-group environment

5.1.1 Reputation encourages cooperation

In M_1, which uses only one group, the increase in the percentage of cooperation was significant. While the configurations χ_{neg} for ρ_{indi} and τ_{indi} came off even with slightly less cooperation over the whole 2,000 rounds, the cooperative behavior differed only sparsely in the final 100 rounds.

5.1.2 Impact of reputation on cooperation is more obvious when costs increase

With doubled costs for M_1, the impact of reputation on cooperation is higher. Even the χ_{neg} achieves higher percentages over the whole 2,000 rounds. Results from the last 100 rounds again are close to the results from χ_{pos}.

5.1.3 Unfavorable initial values slowly develop to higher cooperation

The initial values for ρ_{indi} and τ_{indi} are decisive for the development of cooperation. Even though in each configuration a steady state will be reached at some time, the amount of rounds to be played by agents with unfavorable initial settings is much higher than with favorable or random settings.

5.2 Group-reputation in a multi-agent environment

5.2.1 Group-reputation encourages cooperation with out-group while discouraging cooperation with in-group

The presence of group-reputation increases cooperative behavior vis-à-vis agents of different colors. Even increasing costs for the altruistic act of cooperating cannot delay this effect, nota bene a really harsh environment (e. g., $\gamma \geq 3$) still reduces cooperation, but for the in-group respectively. However, group-reputation seems to decrease cooperation with the in-group, while the number of groups has no significant influence for this value, the presence of groups principally reduces cooperation with the in-group.

5.2.2 Number of groups does not affect cooperation with in-group

While defection toward the out-group is influenced by the number of groups, there is no such effect regarding the cooperation with the in-group.

5.2.3 Believe in good nature (κ_2) is not affected by a hostile environment, but believe in bad nature decreases cooperative behavior

Configuration κ_2, using the values of $\rho_{\text{indi}} = 1$ and $\tau_{\text{indi}} = 0$, is not affected by a high γ. Even with $\gamma = 3$ the percentage of cooperative behavior Y stays above 90%. Agents believing in the good nature of others tend to cooperate, even if their PTR cannot be increased by cooperating with others. On the other hand, in experiments using the configurations κ_1 and κ_3 (where new agents start with a τ_{indi} of 1, which can be interpreted as a certain skepticism) the cooperative behavior decreases with an increasing γ. The τ_{indi} is only affected by mutation and it obviously takes too long for agents and their children to change their opinion.

5.2.4 The number of groups affects cooperative behavior under certain circumstances (κ_0)

In a hostile environment the Y-value increases with the number of groups on the grid when using configuration κ_0. Possible reasons for this phenomenon will be subject of further studies.

6 Further Work

Further studies will analyze interactions between κ_0 and the number of groups on the grid. An explanation of the phenomenon that group-reputation discourages both cooperation with the in-group and defection with out-groups (model M_2) is of great interest.

Our work can be extended in many ways: experiments with simulated reputation while introducing fading memory and examining the influence of noisy interactions with both models will be topics of future research.

The field of social simulations and agent societies for exploration and understanding of social processes by means of computer simulation are discussed in other communities, e. g. in sociology. For example, Edmonds (2006) presents an evolutionary simulation where the presence of *tags* and an inbuilt specialization in terms of skills result in the development of *symbiotic* sharing within groups of individuals with similar tags. Concepts from sociology can be used as an inspiration for further research.

References

Axelrod, R. (1984). *The Evolution of Cooperation*. New York, USA: Basic Books.

Axelrod, R. (1997). *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton, USA: Princeton University Press.

Axelrod, R. & Hammond, R.-A. (2003). The evolution of ethnocentric behavior. In J.-Q. Hill & J.-E. Leighley (Eds.), *Midwest Political Science Convention* Chicago: Midwest Political Science Association.

Chess, D. M. (1988). Simulating the evolution of behavior: the iterated prisoners' dilemma problem. *Complex systems*, 2(6), 663–670.

Edmonds, B. (2006). The emergence of symbiotic groups resulting from skill-differentiation and tags. *Journal of Artificial Societies and Social Simulation*, 9(1).

Fogel, D. B. (1993). Evolving behaviors in the iterated prisoner's dilemma. *Evolutionary Computation*, 1, 77–97.

Marler, P. & Evans, C. S. (1996). Bird calls: just emotional displays or something more? *Ibis*, 138, 26–33.

Nowak, M. A. & Sigmund, K. (1998). Evolution of indirect reciprocity by image scoring, the dynamics of indirect reciprocity. *Nature*, 393, 573–577.

PG 474 (2005). IPD-Dortmund Homepage. http://ls11-www.cs.uni-dortmund.de/people/pg474/. Cited: 24.1.2005.

Schenk, S. (1995). *Evolution kooperativen Verhaltens*. Wiesbaden, Germany: Gabler.

Therneau, T. M. & Atkinson, E. J. (1997). *An Introduction to Recursive Partitioning Using the RPART Routines*. Technical Report 61, Department of Health Science Research, Mayo Clinic, Rochester NY.

Yao, X. & Darwen, P.-J. (2000). How important is your reputation in a multi-agent environment. *International Journal of Knowledge-Based Intelligent Engineering Systems*, 4(3), 191–200.

Zahavi, A. & Zahavi, A. (1997). *The Handicap Principle : A Missing Piece of Darwin's Puzzle*. New York: Oxford University Press.

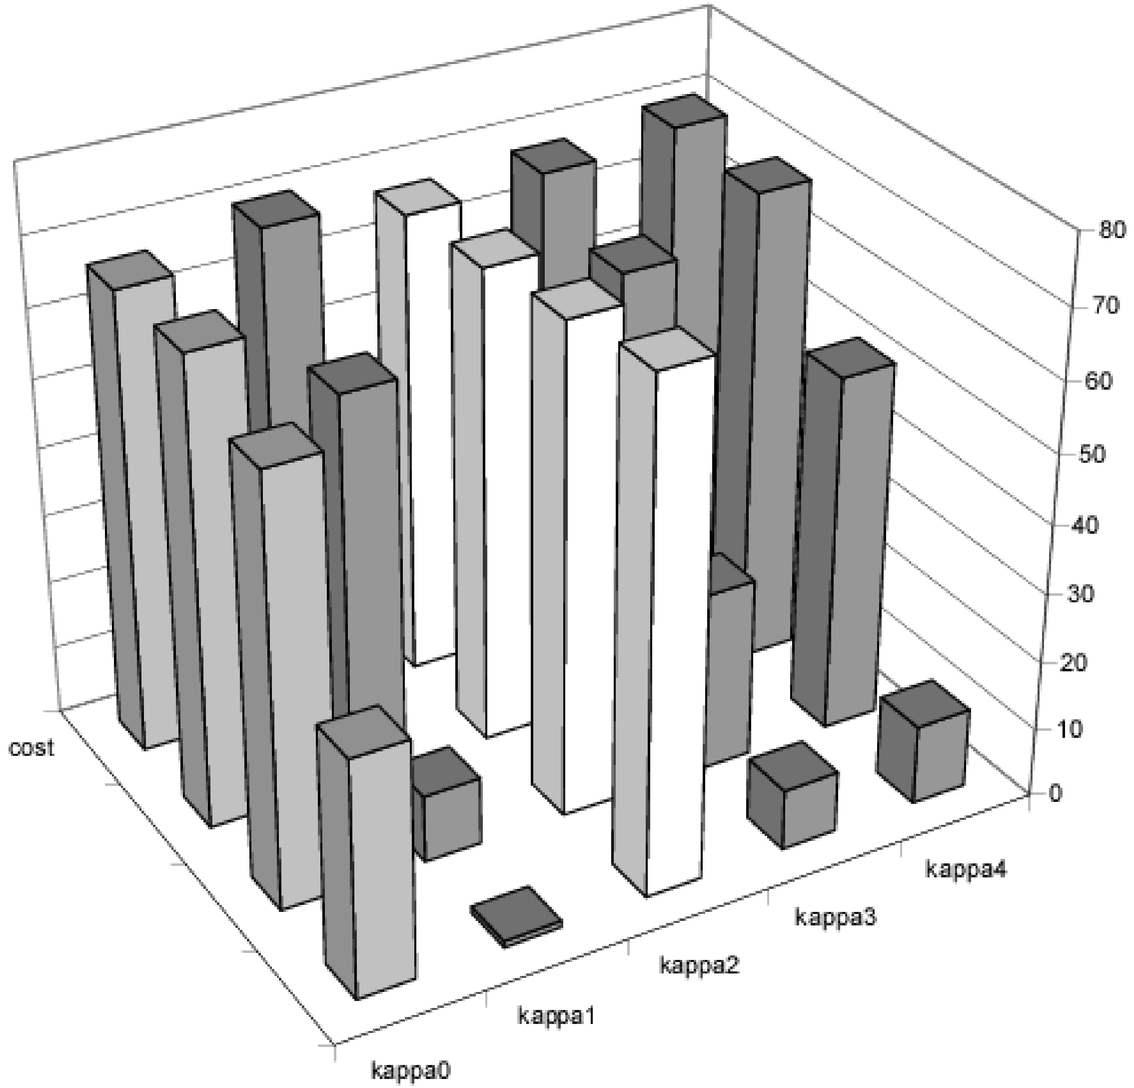

Figure 1: Percentage of cooperation with in-group $\mathcal{C}_{\text{same}}$, cost γ is lowest in rear row increasing to the front row, see Sect. 3.3

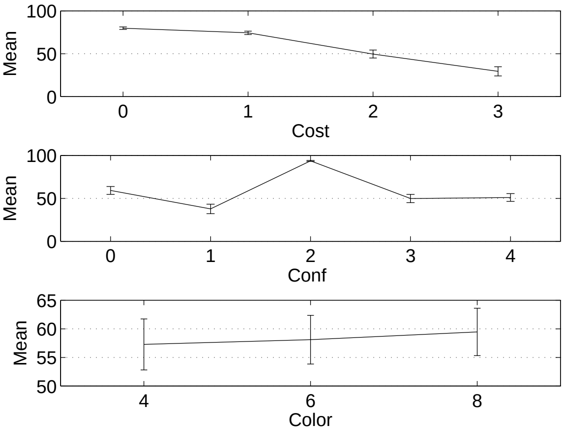

Figure 2: Design plots (with 95% confidence intervals) illustrating the effects of the factors cost, mode, and color on the cooperation level Y. Design plots cannot illustrate interactions between factors, they show only their main effects. Configuration has the largest effect, whereas the choice of the color is negligible

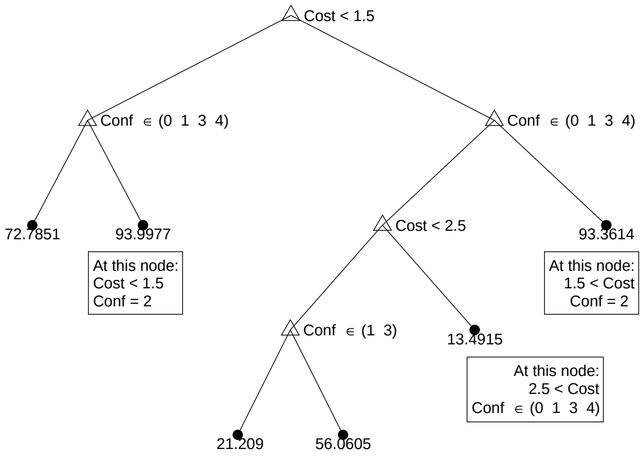

Figure 3: Regression tree to determine significant parameter settings. This tree complements results from the design plot (Fig. 2). The factor at the root node (Cost) has the greatest effect on the response Y. However, high cooperation levels can be obtained if κ_2 is chosen, independently of the costs. This can be seen from the nodes of the third level in the tree