

Article

Ignorance Is Bliss, But for Whom? The Persistent Effect of Good Will on Cooperation

Mike Farjam ^{1,*}, Wladislaw Mill ² and Marian Panganiban ³

¹ Faculty of Social Sciences, Linnaeus University, Växjö 35195, Sweden

² School of Economics, University Jena, Jena 07743, Germany; Wladislaw.Mill@uni-jena.de

³ Max Planck Institute for Research on Collective Goods, Bonn 53113, Germany; mkpanganiban@gmail.com

* Correspondence: Mike.Farjam@lnu.se; Tel.: +49-157-57259445

Academic Editor: Manfred J. Holler

Received: 27 June 2016; Accepted: 17 October 2016; Published: 26 October 2016

Abstract: Who benefits from the ignorance of others? We address this question from the point of view of a policy maker who can induce some ignorance into a system of agents competing for resources. Evolutionary game theory shows that when unconditional cooperators or ignorant agents compete with defectors in two-strategy settings, unconditional cooperators get exploited and are rendered extinct. In contrast, conditional cooperators, by utilizing some kind of reciprocity, are able to survive and sustain cooperation when competing with defectors. We study how cooperation thrives in a three-strategy setting where there are unconditional cooperators, conditional cooperators and defectors. By means of simulation on various kinds of graphs, we show that conditional cooperators benefit from the existence of unconditional cooperators in the majority of cases. However, in worlds that make cooperation hard to evolve, defectors benefit.

Keywords: indirect reciprocity; games on graphs; good will; unconditional cooperation; strategic ignorance

1. Introduction

Why do people cooperate? Economists usually assume that people are actively pursuing their material self-interest, carefully processing and judiciously weighing their options. Obviously, this is not always the case. Some people choose to ignore information on others and cooperate unconditionally. Some governments, for example, self commit to reduce their CO₂-emission regardless of the commitments of others. These decisions would sometimes be pejoratively described as naive or ignorant and subject to the exploitation of others. We show in this paper that unconditional cooperators (UCs) can actually survive in the presence of two other types of agents: defectors (never cooperating) and conditional cooperators (CCs). By simulating interactions between these three types of agents on graphs, we show that a small fraction of UCs helps to sustain and stabilize cooperation in environments where there would be no cooperation without them.

Why is it surprising that UCs survive and even help stabilize cooperation in the population when they are competing for resources with CCs and defectors? In most interactions, cooperation is costly. If using information about others comes with low costs, using information can prevent wasting resources by avoiding cooperation with those that would not reciprocate (e.g., defectors). Furthermore, cooperation with defectors might be interpreted as a signal to others that the agent with whom one has cooperated is considered trustworthy, which lowers the quality of information on which agents like CCs can condition. From an evolutionary point of view, UCs thus seem to favor defectors.

Although ignoring information that comes with no cost seems irrational, experiments show that humans are even willing to pay in order to receive no information about others in strategic interactions [1]. One explanation for this *strategic ignorance* may be that it offers an excuse to

behave egoistically without the psychological costs that come with knowing how one's action affects others [2]. In the context of self-commitment, Carrillo and Mariotti [3] argue that people may prefer no information, fearing the impact that a change of their preconceived ideas could have on their behavior. Additionally to these studies, we show that strategic ignorance can serve not just selfish interests, but can also increase long-run welfare of an individual and other agents. More specifically, by ignoring information about defectors in their interactions, UCs can help promote and stabilize cooperation.

Much research on the interplay between strategies like (un)conditional cooperation and defection has been done by evolutionary game theorists, mostly through the study of games modeling a social dilemma (e.g., the prisoner's dilemma game). Most studies focus on the question of whether a strategy like UC can survive when competing with other strategies, usually not more than 2–3, and whether any cooperation can be sustained on a population level.¹ Although we use a typical model from this area [9,10], our research question is different from the questions usually studied in evolutionary game theory. Within a range of parameters, the model we use is known to sustain cooperation among conditional cooperators and defectors. Our question then is how does the insertion of a fraction of UCs affect (1) the outcomes for UCs, (2) the behavior of agents with other strategies and (3) the performance of the entire system of agents. Hence, our analysis will not focus on the question of whether there is any cooperation in a population of agents, but by how much more cooperation decreases or increases in the population. Furthermore, evolutionary game theorists usually assume that a strategy like UC can enter the population through mutation at any point in time; we induce UCs to the population in a way similar to an intervention of a policy maker or an event changing a fraction of conditional cooperators to unconditional cooperators at once to affect the entire population. Although we frame our study from the perspective of an external decision-maker, agents can nevertheless still explore or mutate. In this light, the work of Han et al. [11] on finite population dynamics bears mentioning. They show cycling among cooperation, defection and a third strategy, which is either (generous) tit-for-tat, win-stay-lose-shift or a generalized version of generous tit-for-tat (intention recognition). Han et al. [12] ask a related question when they study a decision-maker who can interfere in a system to achieve preferred behavioral patterns and aims to optimize the trade-off between the cost of interference and gain from achieving the desired system. Intervention in Han et al.'s study comes in the form of a decision-maker who has fixed resources to reward cooperation, while intervention in our study comes in the form of introducing unconditional cooperators. Phelps et al. [13] also discuss a more general version of the problem by analyzing the optimal design infrastructure of a multi-agent system from the viewpoint of a policy maker.

Unconditional cooperation seems disadvantageous in a three-strategy setting where CCs use a strictly reciprocal/tit-for-tat strategy alongside UCs and defectors. In this setting, defectors might first exploit UCs and benefit from them. Eventually, UCs vanish because their resources are consumed from cooperating with defectors, and defectors receive no further benefits from cooperation, since there are no more agents to cooperate with them. CCs then take over the population since they do not cooperate with defectors and increase their fitness by cooperating with other CCs. The model of Nowak and Sigmund [9,14], which we use, allows for more sophisticated strategies of CCs based on indirect reciprocity.² While CCs using tit-for-tat operate on direct reciprocity, by basing their cooperation on previous interactions with a specific agent, CCs that apply indirect reciprocity base their behavior on the reputation of the agent (i.e., behavior that the other agent has shown in general

¹ For reviews of evolutionary game dynamics, see Nowak and Sigmund [4], Hofbauer and Sigmund [5]. The theoretical background can be found in Maynard Smith [6–8].

² Further fundamental papers using indirect reciprocity are Nowak and Sigmund [9,14], Ghang and Nowak [15], Hoffman et al. [16], Panchanathan and Boyd [17], Ohtsuki et al. [18], Nowak and Sigmund [19], Saavedra et al. [20], Nowak [21]. Closely related to this paper is work on indirect reciprocity like Van Doorn and Taborsky [22], who study reciprocity on social interaction networks; Chen and Liu [23] use indirect reciprocity for cooperation stimulation in cognitive networks; and Zhang et al. [24] are applying this theory for dynamics spectrum access Traag et al. [25] and are observing cooperative clusters by using gossiping as indirect reciprocity.

towards others). Lotem et al. [26] argue that under indirect reciprocity, unconditional altruism or unconditional cooperation can thrive. They interpret unconditional cooperation as a costly signal towards CCs that boosts the UCs' reputation and signals their own trustworthiness to ensure future cooperation with CCs. Furthermore, Ohtsuki and Nowak [10] and Santos et al. [27] find that when agents are interacting according to a network structure, UCs may survive in larger networks when there are at least some small clusters.

Nowak and Sigmund [28] show that when agents' behavior includes a small chance for unintended defection, strategies incorporating what they refer to as forgiveness or goodwill outperform strictly (in)direct reciprocal strategies. In groups where all agents cooperate according to indirect reciprocity, unintended defection causes an agent to immediately get a bad reputation. In this scenario, cooperation breaks down easily, and distrust can spread quickly among agents and cannot be restored. As Nowak and Sigmund show, mild degrees of goodwill rebuild trust and keep up cooperation, while still preventing invasion by defectors. In this light, UCs can inject goodwill or forgiveness in the system by cooperating regardless of the previous actions of players, thereby restoring trust.

Although unconditional cooperation seems to be disadvantageous at first glance, its consequences for an agent and the population turn out to be less clear once we enter worlds with indirect reciprocity. Our question therefore is how a small fraction of UCs set into a population of CCs and defectors affects the evolution of cooperation. In the following sections, we present a formal model based on the work of Nowak and Sigmund (Section 2) and its implementation for simulation in Section 3). In Section 4, we compare worlds with small fractions of UCs with worlds without UCs. We identify in which kinds of environments unconditional cooperation persists and whether it has positive or negative effects on cooperation. Section 5 concludes by elaborating upon the implications of the results in a broader context.

2. Model

2.1. Main Idea

The model we use is based on Nowak and Sigmund [9], Ohtsuki and Nowak [10], which provides a framework that incorporates essential elements needed to study the evolution of cooperation under indirect reciprocity.

At every time step, exactly one random agent is chosen as a potential donor and then one random neighbor of the donor is chosen as a potential recipient. The potential donor only has a binary choice between sharing/cooperating and not sharing/defecting. Under these conditions we get the payoff matrix shown in Table 1. Note that as in Ohtsuki and Nowak [10], $b > c$ and cooperating is thus efficient from the population's point of view, but not from the individual's.

The following subsections discuss other parameters that define the model. Table 2 gives an overview of all of these parameters and a short description of their role in the model.

Table 1. Payoff matrix of agents in the repeated interaction model³.

	Share	Not Share
share	$b - c$	b
not share	b	0

³ The individual interactions of agents are presented in a symmetric way even though each interaction is asymmetric as the recipient can only obtain resources or not. However, each agent will eventually meet the same agent, and hence, the interaction can be seen as symmetric.

Table 2. Summary of all parameters in the model.

Parameter	Meaning
b	resources agents receive from donors
c	resources donors pay to donate
k	number of agents with which an agent can interact
d_i	reputation of agent i
q	probability that d of a neighbor is known
s_i	strategy of agent i
<i>updating</i>	rule that defines updating mechanism
m	games per updating step

2.2. Cooperation on Graphs

In the basic version of their model, Nowak and Sigmund [9] assume a well-mixed population, i.e., interaction between all agents is equally likely. As other more complicated papers, like in Ohtsuki and Nowak [10], we also look at populations that are not well mixed.⁴ The population contains n agents located on a random graph. Every agent has k neighbors, and agents only interact with their neighbors. If $k = n - 1$, the population is well-mixed.

2.3. Indirect Reciprocity

To model indirect reciprocity, Nowak and Sigmund [9] use image scores. These contain information about previous behavior of agents (i.e., the reputation). When an agent cooperates, it receives a “good score”, else a “bad score”. Hence, the image score vector D with elements $d_i \in \{0, 1, 2, 3, \dots\}$ denotes the reputation of all agents i . Agents with the strategy defector will never share their resources irrespective of the reputation of the agent with which they are interacting. The conditional cooperator i shares with player j if $d_j \geq s_i$, meaning that depending on its own cooperating type $s_i \in \{0, 1, 2, \dots, \infty\}$ and the other’s image score, the agent cooperates. Note that conditional cooperators with $s_i = \infty$ are defectors (as for any $j, d_j < \infty$) and conditional cooperators with $s_i = 0$ are unconditional cooperators.

Following Nowak and Sigmund [9], we assume that there are only two possible image scores. When an agent cooperates, then the recipient will receive a “good score”. If the potential donor does not share with the potential recipient, then the potential donor receives a “bad score”. Therefore, we will assume that we have only two possible image scores for an agent: “good” and “bad”, i.e., $d_i \in \{0, 1\}, \forall i \in \{1, n\}$. We include only three possible strategies:

1. Defectors never cooperate with any agent
2. Unconditional cooperators (UCs) cooperate with all agents
3. Conditional cooperators (CCs) cooperate only with an agent i if i is “good”, i.e., $s_i = 1$

As in Nowak [21], each agent knows only a fraction q of the image scores of their neighbors. For the remaining fraction $(1 - q)$, the agent just assumes that the neighbor is “good”.⁵ Nowak derives that a group of indirectly reciprocal agents can promote cooperation if $q > \frac{c}{b}$, i.e., “Indirect reciprocity

⁴ Indeed, there exists a substantial body of literature in evolutionary game theory in spatial settings looking at well-mixed and not well-mixed populations, like in Fu et al. [29], Nowak and May [30], Hauert and Doebeli [31], Nakamaru et al. [32]. Evolutionary game theory on graphs uses research in this area, like Durrett and Levin [33], Ohtsuki et al. [34], Ohtsuki and Nowak [35], Hassell et al. [36]. For an overview of its use in population biology and network structure, see May [37]. It should be noted that many results in evolutionary games are subject to change dependent on the exact parametrization of the model; even more so, when studying graphs [38–40]. Our study should therefore be seen only as a first attempt with regard to the exact influence and interactions of certain parameters.

⁵ This may seem like a rather bold assumption. One might think that the agent who lacks information determines the reputation of the interaction partner, e.g., randomly. However, Nowak and Sigmund [9] show that agent i will outperform agent j if agent i is more optimistic about the reputation of others and that optimistic agents evolve for that reason by themselves. Nevertheless, we also check the robustness of this assumption in our analysis.

can only promote cooperation if the probability q , of knowing someone's reputation exceeds the cost-to-benefit ratio of the altruistic act." [21] (p. 1561).

2.4. Updating Mechanism

During every iteration, one agent is chosen randomly and interacts according to its strategy with a randomly-chosen neighbor. At the end of each generation (m rounds), one agent is chosen randomly to change its strategy by copying the strategy of a fit agent. This process of strategy adaption is called updating, and every updating step can be interpreted as the death of an individual with a certain strategy, while another individual reproduces. The fitness of an agent/its strategy is equivalent to that agent's resources. As the exact dynamics of updating can be implemented in many ways, we consider only three often-used mechanisms [10,41,42].

2.4.1. Birth-Death

Of all agents, one agent ϱ is selected with a probability proportional to its fitness. One randomly-chosen neighbor of ϱ is selected to change its strategy and copies the strategy of ϱ .⁶

2.4.2. Death-Birth

Of all agents, one agent ϱ is selected randomly to copy the strategy of a neighbor, which is selected with a probability proportional to its fitness.⁷

2.4.3. Imitation

Of all agents, one agent ϱ is selected randomly to change its strategy. The new strategy of ϱ will be chosen among the strategies of all neighbors of ϱ and ϱ itself with a probability proportional to everyone's fitness.

3. Simulation

3.1. Overview

Our research questions consider the interplay of three strategies—unconditional cooperation, conditional cooperation and defection—and we study their co-evolution on random graphs with degree k . The multiplicity of strategies and the fact that we are interested in not just well-mixed populations renders the model mathematically close to intractable.⁸ We thus simulate the model with a subset of possible values per parameter to analyze our research question.

Table 3 gives an overview of values with which each parameter in the model was initialized. The initial values of parameters b , k , q , *updating*, *coops* and *conCoops* differed between simulations (discussed in the next subsection). Per possible combination of values per parameter, we ran 110 simulations. This leaves us with 324 combinations of parameter values and a total of 35,640 simulations. Each simulation contains 4000 updating steps, and per updating step, 125 games between agents were played (thus, 500,000 games per simulation). We compared the average ratio of defectors in worlds with 4000 updating steps with long-run simulations of 8000 updating steps (1,000,000 games) and found no meaningful difference. Figure 1 also shows graphically that worlds are not changing after 4000 updating steps. The simulation was programmed in NetLogo [43]. Appendix A contains the pseudo-code of our implementation.

⁶ Note that this updating mechanism is formally identical to pairwise comparison (PC), meaning that two random partners are chosen, and the partner with lower fitness adopts the strategy of the "stronger" partner [10].

⁷ This updating mechanism corresponds to the score-dependent fertility model by Nakamaru et al. [32].

⁸ We do not know of any mathematical solutions for several strategies on a non-well-mixed population.

Table 3. Parameters and their possible values in our simulations. CC, conditional cooperator; UC, unconditional cooperator.

Parameter	Meaning	Value
n	number of agents in the simulation	100
m	games per updating step	125
$iterations$	updating steps per simulation	4000
c	resources donors pay to donate	1
b	resources agents receive from donors	{2, 10}
k	number of agents an agent can interact with	{3, 10, 99}
q	probability that d of a neighbor is known	{0.1, 0.5, 0.9}
$updating$	type of updating mechanism	{BD, DB, IM}
$coops$	fraction of agents that are either CCs or UCs	{0.1, 0.5, 0.9}
$conCoops$	fraction of $coops$ that are CCs	{0.8, 1.0}

3.2. Parameters

The parameters $coops$ and $conCoops$ define the initial distribution of strategies in the population. $coops$ defines the ratio of agents that are cooperators, either conditional (CCs) or unconditional (UCs), and the rest are thus defectors. $conCoops$ defines the initial ratio of cooperators that are CCs. If $conCoops = 1$, the ratio of CCs is equivalent to $coops$, and there are no UCs. If $conCoops = 0.8$, 20% of all cooperators are UCs. The ratio with which each strategy can be found in the population will change throughout the simulation due to the updating of strategies. How the updating mechanism exactly works is defined by the parameter $updating$.

In every interaction between two agents, the donor has to pay c resources (which is one in all simulations) in order for the recipient to get b resources (either two or 10). The number of agents that every agent can interact with (k) is either 3, 10 or 99. We will refer to these k agents as the neighborhood of an agent. If we interpret every agent as a node and each interaction possibility as an edge, we get an undirected, asymmetric graph. Neighbors are assigned randomly, and the resulting graph is thus a random graph. If k is small connectivity and is low, groups are small and scattered, and parts of the graph are more likely to be isolated. The parameter q determines with which probability the donor knows a neighbor. If its value is low, the behavior of CCs becomes very similar to UCs, since (as discussed earlier) CCs assume that the other agent has a good reputation if they do not already know.

3.3. Assumptions

All simulations share the following assumptions. Assumption 1 is used in most theoretical models we encountered (most importantly in Nowak and Sigmund [9]). For the other assumptions, we check the robustness of our results in Section 4.3:

1. All agents start with a “good” reputation.
2. Agents who do not know the reputation of another (which happens with probability of $1 - q$) assume that the reputation of the other is “good”.
3. Only the previous action when the agent was the donor determines its reputation.
4. Donors that do not cooperate with defectors receive a “bad” reputation.
5. Agents do not make mistakes, neither in their perception nor in their actions.
6. There is no mutation in the updating of strategies.

4. Results

4.1. General Results

In this paper, we test how the introduction of a tiny fraction of unconditional cooperators (UCs) affects cooperation in a population where there are also conditional cooperators (CCs) and defectors. We introduce this “tiny fraction” of UCs by replacing 20% of the CCs in the initial population by

UCs. Note that 20% of the CCs may be as little as 2% of the whole population in cases where the starting population of CCs is 10%. Would it lead to more cooperation by encouraging CCs to cooperate more often or will it decrease cooperation by enabling defectors to benefit from the ignorance of UCs? To answer this, we look at the ratio of defectors during the last iteration of a simulation.⁹ A high ratio of defectors implies less cooperation in the population. Since cooperation is the only way for the system to generate resources, the ratio of defectors can be interpreted as the inverse of the efficiency of the population. That is, the more defectors there are in the system, the less resources are generated.

Surprisingly, the simulations in which a tiny fraction of UCs was initially present on average had a lower ratio of defectors, thus were more efficient. On average, worlds with UCs had about 8% less defectors during the final iteration compared to worlds without UCs. A Wilcoxon rank sum test comparing all simulations with and without UCs shows that this difference is clearly significant ($p < 0.01$). This indicates that UCs can actually decrease the fitness of defectors.

To provide a more detailed analysis, we split all simulations into groups. The smallest such group is defined by unique combinations of values of b , k , q , $updating$, $coops$ and $conCoops$. Henceforth, we will refer to these smallest group as worlds. Figure 1 shows that in worlds with $k = 10 \wedge b = 2$, the initial presence of UCs led to more defectors. In 13 out of 162 worlds, the ratio of defectors was higher in worlds with UCs, and in 69 worlds, it was lower in the remaining worlds; the differences in the ratio of defectors was not significant according in a two-sided Wilcoxon rank sum test ($\alpha = 0.05$).¹⁰

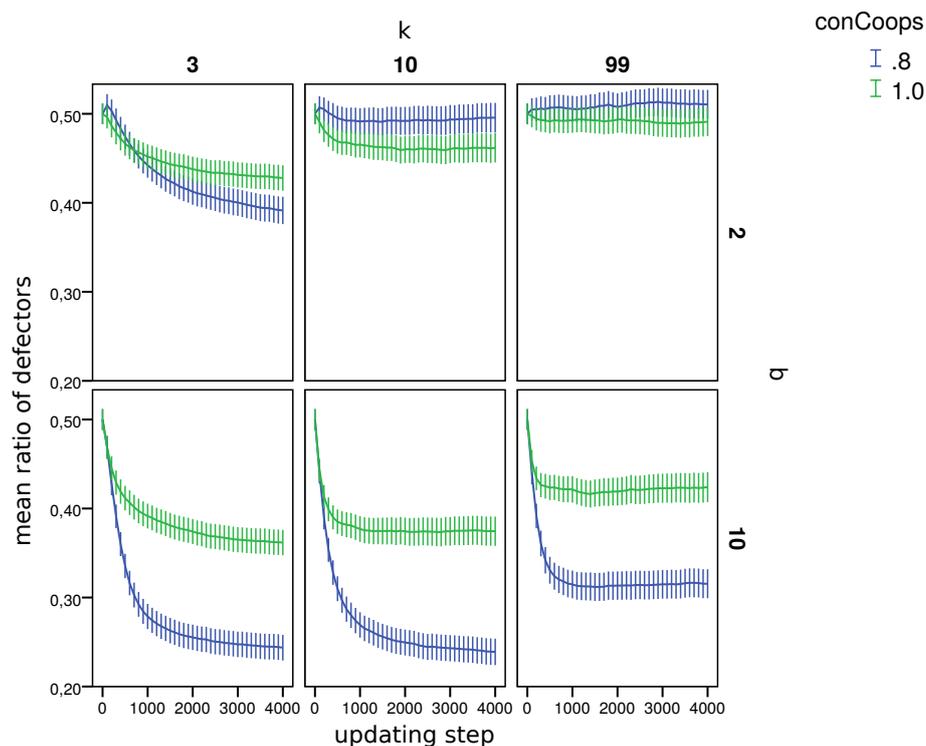


Figure 1. Ratio of defectors in worlds with different values for k and b (bars indicate 95%-confidence intervals).

To identify the parameters that had the largest effect on the ratio of defectors in worlds with and without UCs, we used a step-wise regression builder (based on the AIC). The regression builder

⁹ Similar to other papers analyzing multiple strategies in co-presence [44,45], we also find that in many worlds, the population does not reach a stable distribution of strategies, but oscillates between states with more and less cooperation/defection. Since as we are interested in the general influence of UCs on cooperation, we compare the mean fractions of strategies during the last updating step per world (i.e., 110 simulations).

¹⁰ Appendix B, Table B1 shows the result of the test per world.

used variables b , k , q , $updating$ and $coops$ to predict the difference with respect to the ratio of defectors in the last round in worlds with and without UCs. We also calculated the evolutionary status of CCs. Specifically, we considered for every world if conditional cooperation is an evolutionarily-stable state [46]. Whether a strategy is evolutionarily stable is often used as a main predictor for the survival chances of a strategy [42,47].¹¹ This status of conditional cooperators (*statusCC*) was added as a possible predictor to the step-wise regression builder. An ANOVA comparing whether a model with more predictors performs significantly better in terms of the AIC than a model with less predictors shows that only two predictors (b and $coops$) contribute significantly to the prediction, where b is the most important. The predictive power of *statusCC* thus seems to be negligible.

Figure 2 shows on the left the distribution of the fraction of defectors in all simulations where *conCoops* was one (there were thus no UCs) during the last updating step. One can see that most simulations converged to only containing defectors or only CCs and that simulations where b (benefit of cooperation) is 10 less often converge to a state with a high fraction of defectors. Figure 2 shows on the right the distribution of worlds with regard to the mean fraction of defectors across the 110 simulations per world. While the left density plot shows that most simulations converge to states with only defectors or CCs, the right density plot shows that most worlds allow the population to end up in both states.

Figure 3 also shows the distribution of strategies among all simulations (density plots) and the per world mean fraction of strategies in worlds (points in the simplex) during the last updating step. By comparing the distributions of the fraction of defectors in Figures 2 and 3, one can see that worlds with UCs generally lead to a lower fraction of defectors and that the differences between worlds with $b = 2$ and $b = 10$ became more pronounced.

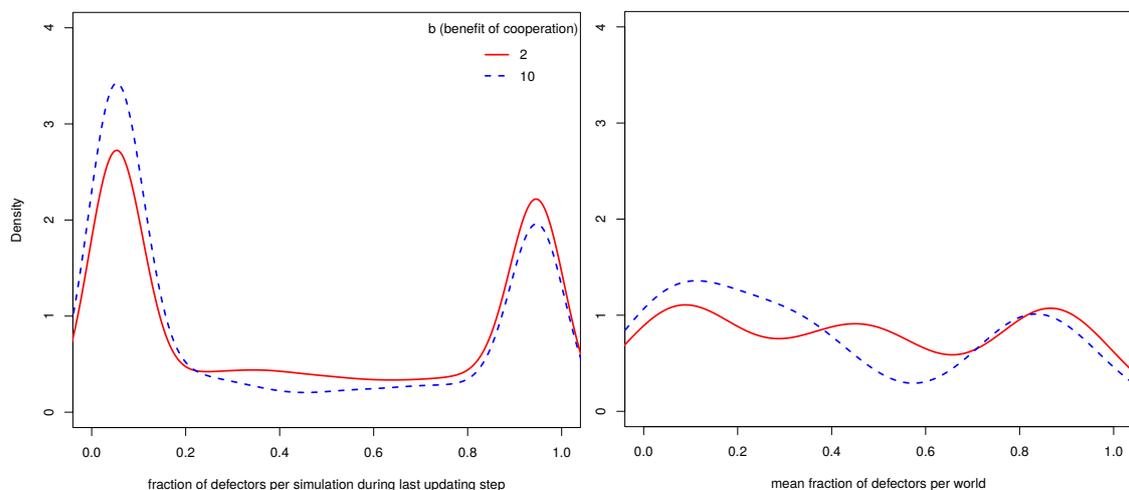


Figure 2. Distribution of the fraction of defectors across all simulations (left) and the per world distribution of the means of the fraction of defectors (right).

¹¹ A meaningful concept of evolutionary stability is also affected by population size [48] and population structure [27,30,49].

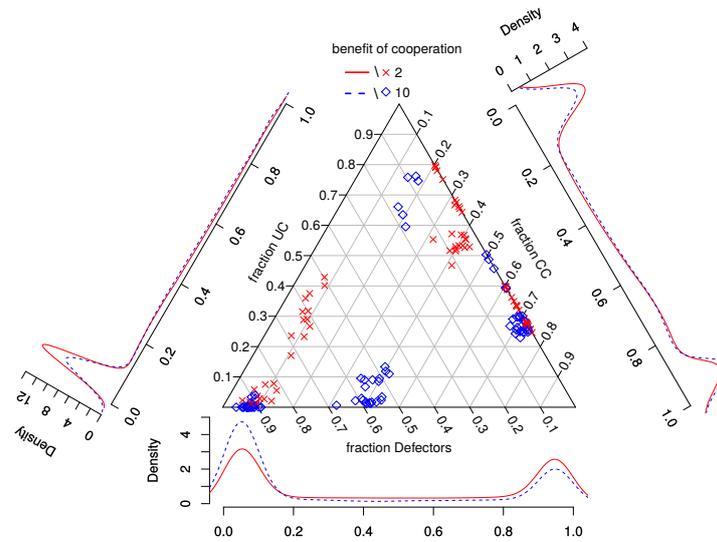


Figure 3. Distribution of the fraction of strategies across all simulations (density plots) and the per world mean fraction of strategies (points in simplex).

4.2. Regression Tree

While the step-wise regression builder indicates which parameters are most important, we discuss in this paragraph how various parameters interact and which of their values make UCs good or bad for defectors. For this purpose, we used the per-world comparison of worlds with *conCoops* = 0.8 and *conCoops* = 1 discussed in Section 4.1 and trained a non-parametric regression tree with the same predictors as those used in the step-wise regression builder to predict the results of the per-world comparison, i.e., whether UCs lead to significantly more or less defectors.¹² Figure 4 shows the resulting tree. Again *b* and *coops* appear the most important predictors since they appear in the top splits of the tree. Below, we discuss the most important splits in the tree.

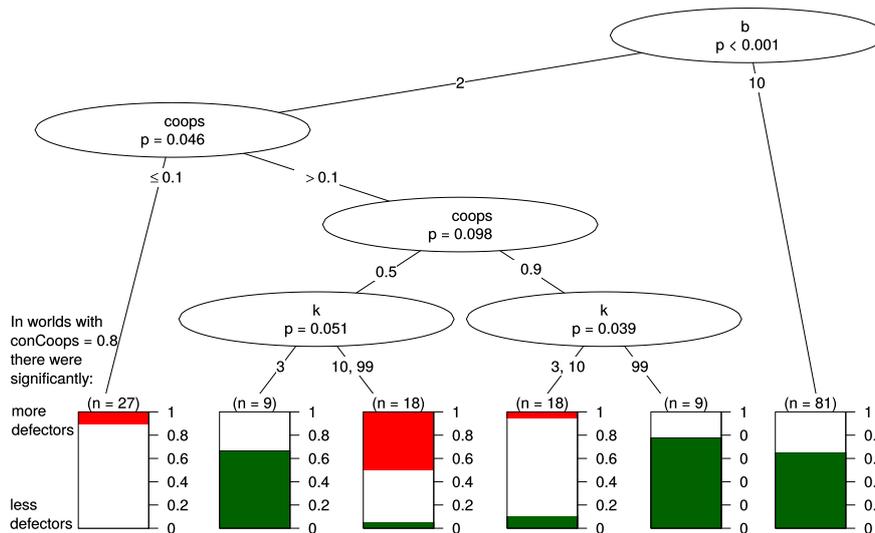


Figure 4. Results of regression tree analysis. The tree predicts whether worlds with *conCoops* = 0.8 lead to more or less defectors than when *conCoops* = 1.

¹² For the generation of the tree we used the R package *ctree* and a minimum information criterion of 0.8.

The first split of the tree is made according to parameter b . In the right branch of the tree where the benefit of cooperation is high ($b = 10$), no other parameters helped with predicting the effect of initial UCs. The effect of UCs was in $\frac{2}{3}$ of all worlds negative for defectors (thus, led to more cooperation), and in the other worlds in this branch, UCs had no significant effect. This effect seems robust as it is not dependent on any of the other parameters, such as the updating mechanism or k .

The left branch of the tree has a more complex structure, and we see that if the benefit of cooperation was rather low ($b = 2$), the effect that UCs had on defectors depends mainly on the value of the parameter $coops$; i.e., in these worlds, it was most important how many agents started being cooperators.

On the very left part of the tree ($b = 2 \wedge coops = 0.1$), there are only very few cooperators, and the benefit of cooperation is rather low. Here, UCs led to more defectors in 10% of all worlds and made no difference in the resulting 90%. We interpret this in the following way: in worlds where the relative benefit of cooperation is small and cooperators are initially in the minority, UCs will, if they make any difference at all, benefit the evolution of defectors. The right branch of the tree, which we discussed in the previous paragraph, contained worlds where the parameter values were set in favor of cooperation ($b = 10$). If the relative benefit of cooperation is high, UCs harm the evolution of defection and benefit cooperators.

The middle part of the tree ($b = 2 \wedge coops \geq 0.5$) classifies worlds where the parameter settings are ambivalent towards cooperation. On the one hand, a large fraction of the population (at least 50%) is initially cooperators. On the other hand, the relative benefit of cooperation is small. Here, only k predicts whether UCs benefit or harm defectors. It seems that when there are only very few defectors ($coops = 0.9$), UCs harm defectors when agents are highly connected ($k = 99$). In contrast, UCs have hardly any effect on defectors when connectedness is low. If there is a considerable amount of defectors ($coops = 0.5$), high connectedness makes the presence of UCs beneficial for defectors, but low connectedness makes UCs harmful to defectors. A possible explanation for this finding is that a low k leads to small isolated clusters of agents. If these clusters do not contain too many defectors or there are only very few connections to other clusters with defectors, cooperators can better withstand invasion by defectors. UCs “fortify” this cluster of cooperation against defection. If there are hardly any defectors, this isolation of cooperators from defectors neither helps nor harms overall cooperation. However, this part of the tree should be interpreted with caution as it seems to depend on a very specific initialization of parameters, and we know that results in evolutionary games are often sensitive to exact parametrization [40].

4.3. Robustness Check

To check the robustness of our results, we changed various assumptions in our model and analyzed how these change the effect of UCs on the final ratio of defectors. Per assumption changed, we reran all worlds another 110 times. Table 4 lists and describes assumptions changed. Appendix C discusses changes and their effect in detail. Figure 5 gives a rough overview on how the change of an assumption changed the influence that UCs had on the ratio of defectors. Red bars indicate the number of worlds where UCs led to more defectors; green bars indicate the ratio of worlds where UCs led to less defectors. Each cell in the figure represents worlds with a unique combination of parameters $coops$ and b . As discussed in the previous section, these turned out to be the most relevant parameters w.r.t. our research question.

Table 4. List of assumptions changed as the robustness check.

Assumption	Description
normal	Simulations with assumption as described in Section 3.3
pessimistic	Unknown reputations are assumed to be bad (good in normal)
memory	Reputations of agents depend on the last 3 actions (1 action in normal)
reproduction	With a 1% chance, offspring will have a random strategy from {UC, CC, defectors} (0% in normal)
perception	With a 5% chance the agent perceives the other agent as good if it is bad and vice versa (0% in normal)
action	With a 5% chance the agent cooperates when he/she wanted to defect and vice versa (0% in normal)

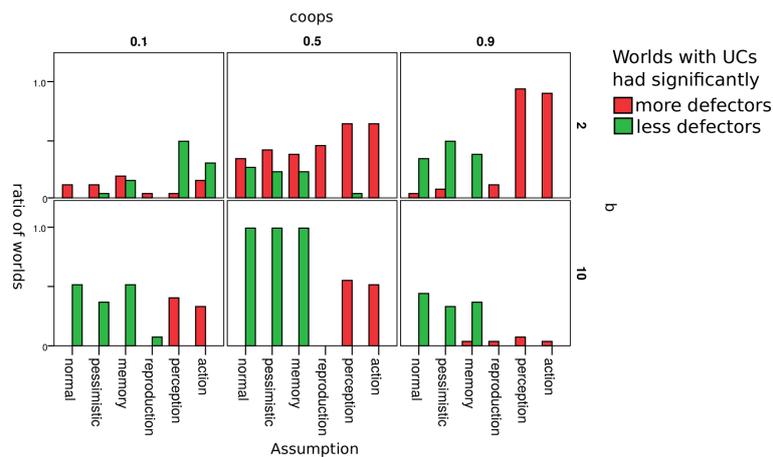


Figure 5. Ratio of worlds where the number of defectors was different in worlds with and without UCs.

Results are hardly affected by whether agents are optimistic or pessimistic w.r.t. agents whose reputation they do not know. Results are also robust w.r.t. to whether an agent’s reputation depends only on the last action of or on the last three actions (assumption memory). As soon as we introduced errors in reproduction, there were hardly any differences between worlds where there were initially UCs and where there were none. We explain this by the fact that UCs can enter populations in worlds with errors in reproduction even if there initially were no UCs. They appear as an error in reproduction. It thus does not seem to matter whether UCs are there from the beginning or enter the population later as mutants. The impact of UCs differs clearly from normal when agents make errors in the perception of the reputation of another agent or in choosing whether to defect or cooperate (action). With these two types of errors, worlds with $conCoops = 0.8$ never lead to less defectors and often to more defectors.

5. Discussion

We tested how the presence of a tiny fraction of unconditional cooperators (UCs) affects cooperation in a population where there are also conditional cooperators (CCs) and defectors. We introduce this “tiny fraction” of UCs by replacing 20% of the CCs in the initial population by UCs. The effect of unconditional cooperation in these populations is two-fold: it allows defectors to benefit from cooperation in their interactions and stops the spread of distrust among CCs. These two mechanisms can, in part, explain our results. Which of the two effects dominates depends on the specific world in which agents interact.

In summary, we find that in worlds where gains from cooperation are high, unconditional cooperation (in small doses) helps sustain cooperation. In worlds where the gains from cooperation are small and where there are initially many defectors, UCs encourage defection. In worlds where

cooperation is only mildly beneficial, but there are relatively few defectors, the effect of UCs on cooperation depends on the social structure, which defines interactions among agents.

If one thinks of unconditional cooperation as a form of strategic ignorance ignoring free information about other agents, this strategy opens an agent to exploitation by defectors and, thus, erodes overall cooperation in the population. We find that this is not always the dominant effect. In our model, we even find that cooperation can be boosted through the presence of a tiny fraction of unconditional cooperators (UCs). Under most of the parameter settings examined in our simulations, conditional cooperators (CCs) benefited more from the presence of UCs, than defectors could exploit the naivete of UCs. However, this was not true among all parameter settings we examined. In about 8% of all cases, defectors benefited more from the presence of UCs.

It is fairly straightforward that unconditional cooperation can be exploited by defectors. However, it is also easy to overlook the positive effects that unconditional cooperators can have on conditional cooperators. We think that the presence of unconditional cooperators has two opposing effects on the evolution of cooperation in a three-strategy setting with UCs, CCs and defectors. On the one hand, UCs benefit defectors as any agent that cooperates with defectors increases defectors' fitness. CCs share less often with defectors since they condition their cooperation. On the other hand, unconditional cooperators rebuild trust among CCs by the following mechanism: defectors immediately receive a bad reputation since they never cooperate. CCs meeting agents with a bad reputation also get a bad reputation because they do not cooperate with that agent. Therefore, one defector can lead to a chain of conditional cooperators getting a bad reputation. This dynamic can lead to a world with very little cooperation even when the beginning ratio of defectors is rather low. In these worlds, defectors would have a higher average fitness as they never "lost" fitness by sharing with others. Unconditional cooperators can stop this spread of bad reputation, since UCs never get a bad reputation because they always share, and CCs hence would always share with UCs and thereby get a good reputation. As a consequence, more CCs will cooperate with each other.

The model we presented in this paper is rather abstract and simplified. Nevertheless, there are phenomena in the real world that share a similar dynamic as depicted in our model. Generally, our findings apply to "missing hero stories", and we will elaborate on this by looking at vaccination and herd immunity in the real world. Herd immunity against a disease can be seen as a public good. It is in the interest of society that a certain number of people are immunized to prevent the spread of diseases. The costs of each individual to get vaccinated are the rare adverse reactions towards vaccination. Obviously, this situation presents a social dilemma, as it is in the interest of everybody that all the others are vaccinated, but each individual is better off by not taking the risk of adverse reactions to vaccines. One can look at individuals fundamentally resisting vaccination as defectors and most other people as conditional vaccinators who will get vaccinated if sufficiently many others are vaccinated. In this context, all people who get vaccinated irrespective of what others do can be seen as unconditional vaccinators. Defectors benefit from people who vaccinate themselves unconditionally and, thus, decrease the chances of an outbreak. In those cases where an outbreak is imminent for the population (this corresponds to worlds in our model where b is high) it is crucial that conditional vaccinators observe others getting vaccinated to get vaccinated themselves. Hence, the benefits of rebuilding trust dominate the disadvantages of avoiding the support of defectors.

Unconditional cooperation on its own should not be deemed as something immediately disadvantageous *per se* for the individual agent or for the population. There is merit in ignoring another agent's history of defection and in acting in good will in spite of another agent's consistent betrayal. Whether unconditional cooperation turns out to be advantageous depends on factors, such as the social structure and the ratio with which other strategies are present in the population, and most importantly, it depends on the gains from cooperating.

Acknowledgments: We would like to thank the Max Planck Society for financial support through the International Max Planck Research School on Adapting Behavior in a Fundamentally Uncertain World. Furthermore, we want to thank the two anonymous referees, which helped us greatly to improve this paper.

Author Contributions: Mike Farjam, Wladislaw and Marian Panganiban conceived of and designed the model and simulation, and Wladislaw Mill was mainly responsible for reviewing the literature. Wladislaw Mill and Mike Farjam analyzed the data. All authors jointly wrote the paper.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Pseudo-Code of Simulation

Pseudo-code: MAIN() IS EXECUTED INITIALLY

```

1: function MAIN()
2:   for 500,000 times do
3:     PLAY()
4:     Once every  $m$  times execute BD(), DB() or IM() depending on updating
5: function PLAY()
6:    $donor = \text{random agent}$ 
7:    $recipient = \text{random neighbor of } donor$ 
8:   if  $donor$  is a defector then
9:      $donor.reputation = 0$ 
10:  if  $donor$  is a conditional cooperator then
11:    if  $recipient.reputation == 1$  then
12:      COOPERATE( $donor, recipient$ )
13:    else
14:       $donor.reputation = 0$ 
15:    if  $donor$  is a unconditional cooperator then
16:      COOPERATE( $donor, recipient$ )
17: function COOPERATE( $donor, recipient$ )
18:    $donor.energy = donor.energy - c$ 
19:    $recipient.energy = recipient.energy + b$ 
20:    $donor.reputation = 1$ 
21: function WEIGHTEDDRAW( $set$ )
22:    $a = \text{one agent from } set \text{ by random weighted (by energy) draw}$ 
23:   return  $a$ 
24: function BD()
25:    $reproducer = \text{WEIGHTEDDRAW}(allAgents)$ 
26:    $replaced = \text{random agent of } reproducer.neighbors$ 
27:    $replaced.strategy = reproducer.strategy$ 
28: function DB()
29:    $replaced = \text{random agent}$ 
30:    $reproducer = \text{WEIGHTEDDRAW}(replaced.neighbors)$ 
31:    $replaced.strategy = reproducer.strategy$ 
32: function IM()
33:    $replaced = \text{random agent}$ 
34:    $reproducer = \text{WEIGHTEDDRAW}(replaced.neighbors \cup replaced)$ 
35:    $replaced.strategy = reproducer.strategy$ 

```

Appendix B. Per World Comparison

The following table shows per the possible combination of the values for parameters k , q , b , $updating$ and $coops$ whether the difference in the ratio of defectors during the last iteration of the simulation was significantly different in worlds with $conCoops = 0.8$ and $conCoops = 1$. A value of one in the column means there were more defectors, zero that there was no significant difference and -1 that there were less defectors.

Table B1. Difference in the ratio of defectors during the last iteration of the simulation.

		3									k	
		BD			DB			IM			Updating	
		0.1	0.5	0.9	0.1	0.5	0.9	0.1	0.5	0.9	conCoop	
0.1	10	-1	-1	-1	-1	-1	0	-1	-1	0		
	2	0	0	1	0	-1	0	0	-1	0		
0.5	10	0	-1	-1	-1	-1	0	0	-1	0		
	2	0	0	0	-1	-1	0	1	-1	0		
0.9	10	0	-1	-1	-1	-1	0	0	-1	0		
	2	0	0	0	-1	-1	0	0	-1	0		
		10									k	
		BD			DB			IM			Updating	
		0.1	0.5	0.9	0.1	0.5	0.9	0.1	0.5	0.9	conCoop	
0.1	10	-1	-1	-1	-1	-1	0	-1	-1	0		
	2	0	1	-1	0	0	0	1	0	0		
0.5	10	0	-1	-1	-1	-1	0	-1	-1	0		
	2	0	0	0	1	1	0	0	1	-1		
0.9	10	0	-1	-1	0	-1	-1	-1	-1	-1		
	2	0	1	-1	0	0	-1	0	1	1		
		99									k	
		BD			DB			IM			Updating	
		0.1	0.5	0.9	0.1	0.5	0.9	0.1	0.5	0.9	conCoop	
0.1	10	0	-1	0	0	-1	-1	0	-1	-1		
	2	0	0	0	0	1	-1	0	1	-1		
0.5	10	0	-1	0	-1	-1	-1	0	-1	0		
	2	0	1	0	0	0	-1	0	0	-1		
0.9	10	0	-1	-1	0	-1	-1	-1	-1	-1		
	2	0	1	-1	0	0	-1	0	1	-1		
q	b											

Appendix C. Detailed Robustness Check

Appendix C.1. Default Assumption of CCs

One of the assumptions is that when the agent does not know the reputation of the recipient, he/she just assumes the recipient to have a good reputation. We label this assumption as optimistic. We could also assume the agents to be pessimistic in the sense that they assume the others to have bad reputation in the case of unknown reputation. As Nowak and Sigmund [19] show, a player using strategy p_1 will out-compete a player using strategy p_2 with $p_1 \geq p_2$, with p being the fraction of unknown reputations considered positively.

We simulated now under the new assumption that unknown reputation is taken as “bad”. The results are very much in line with the previous results, as can be seen in (Figure 5), which supports our findings.

Appendix C.2. Error in Action

Of all those considered errors, we think that the errors in perception and mutation are more important than errors in actions. We state this claim mainly because errors in actions seem rather rare and not obvious. While it seems to make sense that an agent is misinformed about his/her partner’s reputation or perceived his/her partner’s reputation wrongly, it seems less obvious why a mistake should occur in the intended action. For example, defectors cooperate, and unconditional cooperators defect in some instances. That is why we ran two possible error rates in the “errors in actions” stage.

In the first simulation group (324×110 simulations), we used an error rate of 10%, meaning that 10% of the actions intended were not performed and the opposite was implemented. Hence, a defector not wanting to share would in 10% of the cases share. The results in those simulations are surprisingly the opposite of our main results. The results suggest that unconditional cooperation always increases the mean of defectors. A possible reason could be that rebuilding trust is more complicated, because even unconditional cooperators make mistakes.

To see whether this effect is driven by the strong assumption of an error rate of 10%, we implemented an additional simulation group with an error rate of 5%. The new results do not change the outcome. Errors in actions do lead to unconditional cooperators being disadvantageous to the evolution of cooperation.

Appendix C.3. Error in Perception

To test how sensitive our model is to errors of perception, we ran additional simulations wherein each agent perceives the reputation of its partner wrongly in 5% of the cases. This has no influence on defectors or unconditional cooperators because both strategies do not take the reputation of their partner into account. Hence, in 5% of the cases, only conditional cooperators made a mistake by assuming a defector to be good, an unconditional cooperator to be bad or a conditional cooperator to have the opposite of his/her actual reputation.

Unconditional cooperators make conditional cooperators be cooperative by virtue of their good reputation borne out of consistent cooperation. However, if conditional cooperators cannot evaluate the reputation of another agent correctly, the unconditional cooperator does not automatically make conditional cooperators cooperate. Given this, the effect of unconditional cooperation on overall levels of cooperation weakens.

The result, as shown in Figure 5, supports our reasoning. If the world is such that errors of perception occur often, 5%, then the positive effect of unconditional cooperation vanishes, and unconditional cooperators lead to more defectors by supporting them. Hence, unconditional cooperation has a negative effect on the evolution of cooperation.

Appendix C.4. Errors during Reproduction

We decided to implement this error type of reproduction in the following manner. Instead of adapting the strategy following the updating mechanism, the chosen agent adapts the strategy randomly in 5% of the cases. This means that in 5% of the cases, an agent does not take over the strategy of the fittest neighbor (death-birth (DB)), but randomly one of the three possible strategies.

Our reason for implementing this type of error in this way is to be coherent in changing only one aspect of the worlds. The upside of this implementation is that even in worlds without unconditional cooperation, unconditional cooperation can evolve by chance. We decided not to implement errors in reproduction in such a way that in worlds without unconditional cooperation, the agent can randomly take one of the two strategies and in worlds with unconditional cooperation the agent randomly takes one of the three strategies. This type of error manipulation would distort the ratios. In worlds

without unconditional cooperators, the probability that an agent is a defector is 50%. In a world with unconditional cooperators, the probability that an agent is a defector is 33%. Therefore, we would not only change the error in reproduction, but we would change it in two different ways. To have comparable results, we implemented in such a way that we are only manipulating one aspect of our model to give us comparable results.

As it turns out, this upside is driving the effect of the error of reproduction. As in worlds without unconditional cooperation, this strategy can occur by chance, both comparing worlds tend to be similar. Figure 5 supports our argument. We can see, that in both worlds, unconditional cooperation evolves and succeeds in surviving. Hence, we cannot find a statistically-significant effect of unconditional cooperation in most cases, as both worlds tend to be similar.

Appendix C.5. Memory Span

Another possible criticism of our model is that we assume the agents to have a small memory span. So far, it is assumed that each donor takes only the last action of the recipient into account. Obviously, it would be more realistic to assume the agents to be more memorable.

Therefore, we reran all worlds another 110 times with again 4000 updating steps each, with the new assumption of a longer memory span. If donors now take the last three action of the recipient into account, the results do not change much. As Figure 5 indicates, the memory span does not influence the main process of UCs rebuilding trust.

Appendix C.6. Standing

A rather critical assumption in our model is on image scoring. The reputation is reduced if one does not share. Basically, only the last action is considered, regardless of the type of agent with which one was previously interacting. Thus, a conditional cooperator loses his/her reputation if he/she does not share with a defector, as a defector loses reputation by not sharing with an unconditional cooperator. This assumption is rather extreme and was, therefore, discussed much in the literature (see, e.g., Mohtashemi and Mui [50], Brandt and Sigmund [51], Ohtsuki and Iwasa [52], Leimar and Hammerstein [53], Brandt and Sigmund [54]). The idea of losing the reputation only if you do not share with an agent who has good reputation is called standing, as discussed in the work of Sugden [55]. Leimar and Hammerstein [53] show that standing has superior properties, in evolutionary terms, in comparison to image scoring. Brandt and Sigmund [54] and similarly Ohtsuki and Iwasa [52] also support standing's superiority over image scoring in evolutionary processes. The reasoning for this is, according to Brandt and Sigmund [54], that standing can be seen as costless punishment. However, we would argue that costless punishment is an unrealistic assumption and that a conditional cooperator that does not share with a defector and, thus, loses his/her reputation should at least bear the costs of not benefiting from cooperation in the next interaction, because of the change in reputation (for alternative forms of punishment, see [56]). Moreover, we want the agents to use as little cognitive power as possible. Hence, we use image scoring so that the agent only needs to know what happened in the last interaction of his/her partner. The agent does not need to take into account whether he/she defected because his/her partner was a defector or if he/she defected because he/she himself/herself is a defector. The simplicity and minimal cognitive requirements in image scoring were the reasons why we used image scoring in the first place.

Nevertheless, we also ran simulations with standing as the mechanism behind the reputation. However, as standing results in defectors always having bad reputation and all cooperators having good standing, it is already obvious what the effect of unconditional cooperation on the system will be. On the one hand, unconditional cooperation will benefit defectors with whom they share resources with, and on the other hand, rebuild trust and help sustain cooperation when they interact with conditional cooperators. Thus, we expect the effect of unconditional cooperation on overall cooperation to be ambiguous. If standing determines the reputation of an agent, there is no need to rebuild trust because trust does not breakdown in this scenario because conditional cooperators will

always have a good reputation. Hence, unconditional cooperation would only benefit defectors and, therefore, harm cooperation.

As predicted, the results of the simulations supported our reasoning and show that in all worlds, unconditional cooperation had either no effect or a negative effect.

Appendix D. Updating

Updating mechanisms specify stochastic process, which define how the composition of the population changes over time ([10]). The evolutionary game dynamics on regular graphs can be studied by a deterministic replicator equation with transformed payoff matrix. This is also called the replicator equation on graphs ([41]), if the population size is large and if selection is weak. The frequency change over time of each strategy is described by this equation. To illustrate the change of two strategies, the transformed payoff matrix is given by:

$$\begin{array}{c} S_1 \\ S_2 \end{array} \begin{array}{cc} S_1 & S_2 \\ \left(\begin{array}{cc} a_{1,1} & a_{1,2} + H \\ a_{2,1} - H & a_{2,2} \end{array} \right) \end{array} \quad (D1)$$

whereas H depends on the updating rule.

For birth-death updating, we have:

$$H = \frac{a_{1,1} + a_{1,2} - a_{2,1} - a_{2,2}}{k - 2} \quad (D2)$$

For death-birth updating, the term H is described by:

$$H = \frac{(k + 1)a_{1,1} + a_{1,2} - a_{2,1} - (k + 1)a_{2,2}}{(k + 1)(k - 2)} \quad (D3)$$

Additionally, for imitation updating, we obtain:

$$H = \frac{(k + 3)a_{1,1} + 3a_{1,2} - 3a_{2,1} - (k + 3)a_{2,2}}{(k + 3)(k - 2)} \quad (D4)$$

k is the number of neighbors. The term H converges to zero if $k \rightarrow \infty$, which is the definition of a complete graph, meaning that each player is interacting with each possible player.

These calculation are derived in [41] by using pair-approximation ([57,58]) for regular graphs with degree $k \geq 3$.

References

1. Grossman, Z.; van der Weele, J.J. *Self-Image and Strategic Ignorance in Moral Dilemmas*; Technical Report; Available at SSRN 2237496; University of California: Oakland, CA, USA, 2013.
2. Cohen, S. *States of Denial: Knowing About Atrocities and Suffering*; John Wiley & Sons: New York, NY, USA, 2013.
3. Carrillo, J.D.; Mariotti, T. Strategic Ignorance as a Self-Disciplining Device. *Rev. Econ. Stud.* **2000**, *67*, 529–544.
4. Nowak, M.A.; Sigmund, K. Evolutionary Dynamics of Biological Games. *Science* **2004**, *303*, 793–799.
5. Hofbauer, J.; Sigmund, K. Evolutionary game dynamics. *Am. Math. Soc.* **2003**, *40*, 479–519.
6. Maynard Smith, J. *Evolution and the Theory of Games*, 1st ed.; Cambridge University Press: Cambridge, UK, 1982.
7. Maynard Smith, J. The theory of games and the evolution of animal conflicts. *J. Theor. Biol.* **1974**, *47*, 209–221.
8. Maynard Smith, J. Group selection and kin selection. *Nature* **1964**, *201*, 1145–1147.
9. Nowak, M.A.; Sigmund, K. Evolution of indirect reciprocity by image scoring. *Nature* **1998**, *393*, 573–577.
10. Ohtsuki, H.; Nowak, M.A. Direct reciprocity on graphs. *J. Theor. Biol.* **2007**, *247*, 462–470.

11. Han, T.A.; Moniz Pereira, L.; Santos, F.C. Intention recognition promotes the emergence of cooperation. *Adapt. Behav.* **2011**, *19*, 264–279.
12. Han, T.A.; Tran-Thanh, L.; Jennings, N.R. The Cost of Interference in Evolving Multiagent Systems. In Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS '15), Istanbul, Turkey, 4–8 May 2015; International Foundation for Autonomous Agents and Multiagent Systems: Richland, SC, USA, 2015; pp. 1719–1720.
13. Phelps, S.; Mcburney, P.; Parsons, S. Evolutionary Mechanism Design: A Review. *Auton. Agents Multi-Agent Syst.* **2010**, *21*, 237–264.
14. Nowak, M.A.; Sigmund, K. Evolution of indirect reciprocity. *Nature* **2005**, *437*, 1291–1298.
15. Ghang, W.; Nowak, M.A. Indirect reciprocity with optional interactions. *J. Theor. Biol.* **2015**, *365*, 1–11.
16. Hoffman, M.; Yoeli, E.; Nowak, M.A. Cooperate without looking: Why we care what people think and not just what they do. *Proc. Natl. Acad. Sci. USA* **2015**, *112*, 1727–1732.
17. Panchanathan, K.; Boyd, R. A tale of two defectors: The importance of standing for evolution of indirect reciprocity. *J. Theor. Biol.* **2003**, *224*, 115–126.
18. Ohtsuki, H.; Iwasa, Y.; Nowak, M.A. Indirect reciprocity provides only a narrow margin of efficiency for costly punishment. *Nature* **2009**, *457*, 79–82.
19. Nowak, M.A.; Sigmund, K. The Dynamics of Indirect Reciprocity. *J. Theor. Biol.* **1998**, *194*, 561–574.
20. Saavedra, S.; David, S.; Felix, R.T. Cooperation under Indirect Reciprocity and Imitative Trust. *PLoS ONE* **2010**, *5*, e13475.
21. Nowak, M.A. Five Rules for the Evolution of Cooperation. *Science* **2006**, *314*, 1560–1563.
22. Van Doorn, G.S.; Taborsky, M. The evolution of generalized reciprocity on social interaction networks. *Evolution* **2012**, *66*, 651–664.
23. Chen, Y.; Liu, K. Indirect Reciprocity Game Modelling for Cooperation Stimulation in Cognitive Networks. *IEEE Trans. Commun.* **2011**, *59*, 159–168.
24. Zhang, B.; Chen, Y.; Liu, K.J.R. An indirect reciprocity game theoretic framework for dynamic spectrum access. In Proceedings of the 2012 IEEE International Conference on Communications (ICC), Ottawa, ON, Canada, 10–15 June 2012; pp. 1747–1751.
25. Traag, V.; Van Dooren, P.; Nesterov, Y. Indirect reciprocity through gossiping can lead to cooperative clusters. In Proceedings of the 2011 IEEE Symposium on Artificial Life (ALIFE), Paris, France, 11–15 April 2011; pp. 154–161.
26. Lotem, A.; Fishman, M.A.; Stone, L. From reciprocity to unconditional altruism through signalling benefits. *Proc. R. Soc. Lon. B Biol. Sci.* **2003**, *270*, 199–205.
27. Santos, F.C.; Pacheco, J.M.; Lenaerts, T. Evolutionary dynamics of social dilemmas in structured heterogeneous populations. *Proc. Natl. Acad. Sci. USA* **2006**, *103*, 3490–3494.
28. Nowak, M.A.; Sigmund, K. Tit for tat in heterogeneous populations. *Nature* **1992**, *355*, 250–253.
29. Fu, F.; Wang, L.; Nowak, M.A.; Hauert, C. Evolutionary dynamics on graphs: Efficient method for weak selection. *Phys. Rev. E* **2009**, *79*, 046707.
30. Nowak, M.A.; May, R.M. Evolutionary games and spatial chaos. *Nature* **1992**, *359*, 826–829.
31. Hauert, C.; Doebeli, M. Spatial structure often inhibits the evolution of cooperation in the snowdrift game. *Nature* **2004**, *428*, 643–646.
32. Nakamaru, M.; Nogami, H.; Iwasa, Y. Score-dependent Fertility Model for the Evolution of Cooperation in a Lattice. *J. Theor. Biol.* **1998**, *194*, 101–124.
33. Durrett, R.; Levin, S. The Importance of Being Discrete (and Spatial). *Theor. Popul. Biol.* **1994**, *46*, 363–394.
34. Ohtsuki, H.; Nowak, M.A.; Pacheco, J. Breaking the Symmetry between Interaction and Replacement in Evolutionary Dynamics on Graphs. *Phys. Rev. Lett.* **2007**, *98*, 108–106.
35. Ohtsuki, H.; Nowak, M.A. Evolutionary games on cycles. *Proc. R. Soc. B Biol. Sci.* **2006**, *273*, 2249–2256.
36. Hassell, M.P.; Comins, H.N.; May, R.M. Species coexistence and self-organizing spatial dynamics. *Nature* **1994**, *370*, 290–292.
37. May, R.M. Network structure and the biology of populations. *Trends Ecol. Evol.* **2006**, *21*, 394–399.
38. Szabó, G.; Borsos, I. Evolutionary potential games on lattices. *Phys. Rep.* **2016**, *624*, 1–60.
39. Szabó, G.; Fáth, G. Evolutionary games on graphs. *Phys. Rep.* **2007**, *446*, 97–216.
40. Roca, C.P.; Cuesta, J.A.; Sánchez, A. Evolutionary game theory: Temporal and spatial effects beyond replicator dynamics. *Phys. Life Rev.* **2009**, *6*, 208–249.

41. Ohtsuki, H.; Nowak, M.A. The replicator equation on graphs. *J. Theor. Biol.* **2006**, *243*, 86–97.
42. Ohtsuki, H.; Nowak, M.A. Evolutionary stability on graphs. *J. Theor. Biol.* **2008**, *251*, 698–707.
43. Tisue, S.; Wilensky, U. Netlogo: A simple environment for modeling complexity. In Proceedings of the International Conference on Complex Systems, Boston, MA, USA, 16–21 May 2004; pp. 16–21.
44. Imhof, L.A.; Fudenberg, D.; Nowak, M.A. Tit-for-tat or win-stay, lose-shift? *J. Theor. Biol.* **2007**, *247*, 574–580.
45. Imhof, L.A.; Fudenberg, D.; Nowak, M.A. Evolutionary cycles of cooperation and defection. *Proc. Natl. Acad. Sci. USA* **2005**, *102*, 10797–10800.
46. Maynard Smith, J.; Price, G.R. The Logic of Animal Conflict. *Nature* **1973**, *246*, 15–18.
47. Taylor, P.D.; Jonker, L.B. Evolutionary stable strategies and game dynamics. *Math. Biosci.* **1978**, *40*, 145–156.
48. Wild, G.; Taylor, P.D. Fitness and evolutionary stability in game theoretic models of finite populations. *Proc. R. Soc. Lond. Ser. B Biol. Sci.* **2004**, *271*, 2345–2349.
49. Nakamaru, M.; Matsuda, H.; Iwasa, Y. The Evolution of Cooperation in a Lattice-Structured Population. *J. Theor. Biol.* **1997**, *184*, 65–81.
50. Mohtashemi, M.; Mui, L. Evolution of indirect reciprocity by social information: The role of trust and reputation in evolution of altruism. *J. Theor. Biol.* **2003**, *223*, 523–531.
51. Brandt, H.; Sigmund, K. Indirect reciprocity, image scoring, and moral hazard. *Proc. Natl. Acad. Sci. USA* **2005**, *102*, 2666–2670.
52. Ohtsuki, H.; Iwasa, Y. How should we define goodness?—Reputation dynamics in indirect reciprocity. *J. Theor. Biol.* **2004**, *231*, 107–120.
53. Leimar, O.; Hammerstein, P. Evolution of cooperation through indirect reciprocity. *Proc. R. Soc. Lond. B Biol. Sci.* **2001**, *268*, 745–753.
54. Brandt, H.; Sigmund, K. The logic of reprobation: Assessment and action rules for indirect reciprocation. *J. Theor. Biol.* **2004**, *231*, 475–486.
55. Sugden, R. *The Economics of Rights, Co-Operation and Welfare*; Palgrave Macmillan: New York, NY, USA, 1986.
56. Farjam, M.; Faillo, M.; Sprinkhuizen-Kuyper, I.; Haselager, P. Punishment mechanisms and their effect on cooperation: A simulation study. *J. Artif. Soc. Soc. Simul.* **2015**, *18*, 5, doi:10.18564/jasss.2647.
57. Matsuda, H.; Ogita, N.; Sasaki, A.; Sato, K. Statistical Mechanics of Population: The Lattice Lotka-Volterra Model: Invited Papers. *Prog. Theor. Phys.* **1992**, *88*, 1035–1049.
58. Harada, Y.; Ezoe, H.; Iwasa, Y.; Matsuda, H.; Sato, K. Population Persistence and Spatially Limited Social Interaction. *Theor. Popul. Biol.* **1995**, *48*, 65–91.



© 2016 by the authors; licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC-BY) license (<http://creativecommons.org/licenses/by/4.0/>).