Center for the Study of Institutional Diversity

CSID Working Paper Series

#CSID-2013-007

Conditional behavior affects the level of evolved cooperation in public good games

Marco A. Janssen
Arizona State University, USA

Miles Manning
Arizona State University, USA

Oyita Udiani
Arizona State University, USA

August 14, 2013

The Center for the Study of Institutional Diversity resides in the School of Human Evolution and Social Change at Arizona State University. CSID can be found on the internet at: http://csid.asu.edu. CSID can be reached via email at csid@asu.edu.

Conditional behavior affects the level of evolved cooperation in public good games

Marco A. Janssen[a], Miles Manning[b], Oyita Udiani[c]

[a]Center for the Study of Institutional Diversity, Arizona State University, Tempe, AZ 85287-2402, USA;
[b]Arizona State University, USA
[c]Arizona State University, USA

Corresponding author:
Marco A. Janssen
School of Human Evolution and Social Change
Center for the Study of Institutional Diversity
Arizona State University
PO Box 872402
Tempe, AZ 85287-2402, USA
Marco.Janssen@asu.edu

Abstract:
Human societies are unique in the level of cooperation among non-kin. Evolutionary models explaining this behavior typically assume pure strategies of cooperation and defection. Behavioral experiments, however, demonstrate that humans are typically conditional co-operators who have other-regarding preferences. Building on existing models on the evolution of cooperation and costly punishment, we use a utilitarian formulation of agent decision making to explore conditions that support the emergence of cooperative behavior. Our results indicate that cooperation levels are significantly lower for larger groups in contrast to the original pure strategy model. Here, defection behavior not only diminishes the public good, but also affects the expectations of group members leading conditional co-operators to change their strategies. Hence defection has a more damaging effect when decisions are based on expectations and not only pure strategies.

Keywords:
Public Good Games, Group Selection, Other-Regarding Preferences, Conditional Cooperation

Conditional behavior affects the level of evolved cooperation in public good games

Marco A. Janssen
School of Human Evolution and Social Change
Center for the Study of Institutional Diversity
Arizona State University
PO Box 872402
Tempe, AZ 85287-2402
Marco.Janssen@asu.edu

Miles Manning and Oyita Udiani
Mathematical, Computational and Modeling Sciences Center
School of Human Evolution and Social Change
Center for the Study of Institutional Diversity
Arizona State University

Introduction

One of the key puzzles in social and biological sciences is the scale of cooperative behavior in human societies. Unlike other species, humans cooperate in non-repeated interactions with non-kin in large groups. Theories of cooperation have focused on kin selection (Hamilton 1964), direct reciprocity in repeated interactions (Trivers, 1971; Axelrod, 1984), and indirect reciprocity and costly signalling of reputation (Alexander 1987; Nowak and Sigmund 1998; Zahavi, 1977). The structure of reciprocity interactions in spatial networks and group selection have also been proposed as possible explanations (Nowak, 2006).

Most computational and mathematical models on the emergence of cooperation have employed evolutionary game theory with pure or fixed mixed strategists to capture human behavior (Axelrod, 1997; Gintis, 2000; Nowak, 2006). However, experimental studies show that humans make strategic decisions that are contextual and evolve due to a learning process (Camerer, 2003). Participants in behavioral experiments have been found to adjust their level of cooperation in response to an expected level of cooperation by others (e.g., Fischbacher et al., 2001). Thus, decisions participants make in social dilemma situations cannot be explained by selfish rational motivations. Alternative utility functions that have been proposed include weighting the earnings of others relative to individual returns (Fehr and Schmidt, 1999; Charness and Rabin, 2002). As such, participants may not maximize their material wealth when maximizing their utility. Finally, decisions are typically made with partial (or erroneous) information about the intended actions of others. In iterated games, such uncertainty may be mitigated by using prior and current information to inform the expected decisions in future rounds (Camerer, 2003).

The use of punishment has been found to be an important factor explaining cooperation in large human societies. Laboratory experiments on public goods and common pool resources have shown that participants are willing to give up monetary returns to punish non-cooperators (Yamagishi, 1986; Ostrom et al., 1992; Fehr and Gächter 2002). When the option to punish others at a cost to themselves was introduced in experiments, participants utilized this force and the level of cooperation increased. From an evolutionary perspective, it is puzzling why individuals would accept a reduction in payoff to decrease the payoff of other individuals. The possibility of the evolution of altruistic punishment in human societies has been demonstrated by various studies including Boyd et al. (2003).

In this paper we model agent decision making in line with experimental evidence of human behavior in social dilemmas. In contrast to previous models, we assume that tradeoffs between behavioral actions also depend on an expected utility, which includes "other-regarding preferences" (i.e., considerations for the earnings of other group members). Finally, we incorporate an update strategy that allows agents to weight future expectations based on the current level of cooperation in the group.

With these modifications, we explore some conditions that favour the evolution of preferences and expectations that lead to cooperative behavior. In order to investigate this, we will build on the cultural group selection model of Boyd et al. (2003) which was used to evaluate the conditions for the evolution of altruistic punishment. As such we will also test the conditions under which cooperative behavior evolve with and without the option of punishment.

The rest of the paper will be structured as follows: First, we present the model in detail and its underlying assumptions. Next, we present the results of our simulations, and conclude with a discussion including implications for future work. The model code and detailed documentation of the model used in this paper can be found at http://www.openabm.org/model/3887/version/1/view.

Model Description

Our model builds on the work of Boyd et al. (2003). Using a similarly structured cultural group selection framework we will formulate an agent's decision making in line with findings from behavioral economics. Consider a population that is divided into groups of size n. Agents make a decision to contribute or not to the public good. Agents have an initial endowment of 1 unit, and if an agent contributes it will incur a cost c to produce a total benefit b that is shared equally among group members. If an agent does not contribute, it will incur no costs and produce no benefits. If the fraction of contributors in the group is y, the expected payoff x for contributors is $1 + by - c$ and the expected payoff for non-contributors is $1 + by$. Hence the payoff disadvantage of the contributors is a constant c independent of the distribution of types in the population.

An agent decides whether or not to contribute to the public good of her group based on the evaluation of the expected utility of two options. Agents decide to contribute or not in a public good game using the following utility function (U_i):

$$U_i = x_i - \alpha_i * \max(x_i - \bar{x}_{-i}, 0) + \beta_i * \max(\bar{x}_{-i} - x_i, 0) \qquad (1)$$

Here, x_i defines the earning of the focal agent and $\bar{x}_{-i}$ the average earning of all other agents in the group. The parameter α_i defines the strength of aversion to exploiting others (i.e., the level of guilt an agent feels when it earns more than the average). Conversely, the parameter β_i defines the strength of altruistic tendencies (i.e., the pride an agent feels when it contributes more than the average). Together, these parameters, along with expected contribution of the group into the public good, determine whether an agent will cooperate. Note that the decision to contribute or not may depend not only on individual payoffs, but also on the payoffs of others.

In line with Charness and Rabin (2002), we can define the following cases for $\beta_i \leq \alpha_i \leq 1$

Case 1: The players prefer to have their payoffs higher than those of the other players.
If $\beta_i \leq \alpha_i \leq 0$, players are highly competitive.

Case 2: Players prefer the payoffs among all players to be equal. This "inequity aversion" holds when $\beta_i < 0 < \alpha_i \leq 1$ (see Fehr and Schmidt, 1999).

Case 3: The third model approximates a "Social Welfare Consideration" which holds when $0 < \beta_i \leq \alpha_i \leq 1$. The parameter α captures the extent to which a player weighs the average payoffs of the other n-1 agents compared to his own payoff, when his own payoff is higher than the average payoff of the others.

Case 4: If $\alpha_i = \beta_i = 0$, then players only care about their own welfare.

Agents make these calculations based on an expectation of cooperation among group-mates. This "trust" variable (T_c) has an initial value and is updated based on observed levels of cooperation after each round of play.

Now we must define the expected earnings on which the agent decides to contribute or not to the public good. In order for agents to make decisions they need to evaluate the expected utility value of equation (1). Therefore, we need to define the expected earnings, which is based on the expected level of cooperation. For a cooperative agent the expected fraction of cooperation is T_c for the (n-1) other group members and 1 for the agent itself. Therefore we can define the expected earnings for a cooperative agent as

$$x_C^* = 1 + b \cdot \frac{\{T_C \cdot (n-1) + 1\}}{n} - c \tag{2}$$

In a similar way we can define the expected earning of a defecting agent as

$$x_D^* = 1 + b \cdot \frac{\{T_C \cdot (n-1)\}}{n} \tag{3}$$

In order to calculate the utility of agent i, agent i needs to define the expected average earning of the other n-1 agents. We can distinguish the expectations for the case agent i will contribute or not contribute. When agent i will contribute and has the expectation that fraction T_c of the other agents contributes too, the expected earnings of other agents is equal to

$$\overline{x_C^*} = T_c \cdot x_C^* + (1 - T_c) \cdot \left(x_D^* + \frac{b}{n}\right) \tag{4}$$

On the other hand if agent i is not contributing, the benefit of b/n will not be entertained by the other agents, and the expected earnings of other agents is defined as

$$\overline{x_D^*} = T_c \cdot \left(x_C^* - \frac{b}{n}\right) + (1 - T_c) \cdot x_D^* \tag{5}$$

Using equation (1) we can define the expected utility of cooperation and defection in the following way:

$$E[U_C] = x_C^* + \beta_i * \left(\overline{x_C^*} - x_C^*\right) - \alpha_i * \left(x_C^* - \overline{x_C^*}\right) \tag{6}$$

$$E[U_D] = x_D^* + \beta_i * \left(\overline{x_D^*} - x_D^*\right) - \alpha_i * (x_D^* - \overline{x_D^*}) \tag{7}$$

The expected utility values are used to make a decision. The choice with the highest expected utility is chosen. If there is a tie, randomly one of the options is drawn. With a small probability, ε, an error is made and the alternative is selected. If there is a tie one of the options is drawn randomly.

When we allow agents to punish others at a cost to themselves, the cost of punishment is *k* and the penalty of the agent being punished is *p*. This will affect the expected earnings of the agents. For simplicitys' sake, we assume only cooperative agents will use costly punishment. The expected earnings of a cooperative agent who is punishing is therefore defined as:

$$x_{CP}^* = 1 + b \cdot \frac{\{T_c \cdot (n-1) + 1\}}{n} - c - k(1 - T_c) \cdot (n-1) \tag{8}$$

The expected earning of a defecting agent (3) is now updated as follows:

$$x_D^* = 1 + b \cdot \frac{\{T \cdot (n-1)\}}{n} - p \cdot T_p \cdot T_c \cdot (n-1) \tag{9}$$

With T_p the share of cooperative agents who punish defectors.

The expected average earning of other agents can now be updated for the three possible decisions of agent i. If agent i is contributing to the public good, but does not punish, the expected average earnings of the other agents is equal to

$$\overline{x_C^*} = T_c \cdot \{\left(1 - T_p\right) \cdot x_C^* + T_p \cdot x_{CP}^*\} + (1 - T_c) \cdot \left(x_D^* + \frac{b}{n}\right) \tag{10}$$

The expected earning of other agents when agent i is a punisher is

$$\overline{x_{CP}^*} = T_c \cdot \{\left(1 - T_p\right) \cdot x_C^* + T_p \cdot x_{CP}^*\} + (1 - T_c) \cdot \left(x_D^* + \frac{b}{n} - p\right) \tag{11}$$

The expected earning of other agents when agent i is defecting is given by

$$\overline{x_D^*} = T_c \cdot \{\left(1 - T_p\right) \cdot \left(x_C^* - \frac{b}{n}\right) + T_p \cdot \left(x_{CP}^* - \frac{b}{n} - k\right)\} + (1 - T_c) \cdot x_D^* \tag{12}$$

The expected utility of a cooperative agent who uses costly punishment is now defined as

$$E[U_P] = x_{CP}^* - \alpha \cdot \max\left(x_{CP}^* - \overline{x_{CP}^*}, 0\right) + \beta \cdot \max\left(\overline{x_{CP}^*} - x_{CP}^*, 0\right) \tag{13}$$

If an agent has decided to cooperate, the agent also has to make a decision to punish or not. The decision to punish or not is based on the expected benefit utility of using punishment versus no punishment.

The agent evaluates the expected utility of both cooperation and defection and chooses the option with the highest expected utility. Like Boyd et al., we assume a small probability ε the agent will make a mistake. When an agent cooperates she evaluates the expected utility to punish or not, and chooses the option with the highest expected utility.

Every generation agents make decisions, and transfer the expectations about other agents to the next generation. This means that they update the expectations that others cooperate (T_c) and the expectation that other cooperator punish (T_p) based on the observations made during the last generation.

Every generation agents may imitate the behavioral norms (i.e., α, β, T_c and T_p) of a more successful agent. With probability (1-*m*) this agent is from the same group, and with probability m the agent is from another group in the population. This models the individual-level selection of forces that promote payoff-maximizing strategies. This submodel is run with probability m for each agent. Agent i copies the behavioral traits of agent *j* with a probability equal to $\frac{\left(1+(X_j-X_i)\right)}{2}$.

Like Boyd et al. (2003) we assume that group selection occurs through intergroup conflicts. At the end of each generation, groups are randomly paired, and with probability ε the interaction results in one group defeating and replacing the other group. The probability that group *i* defeats group *j* is $\frac{\left(1+(N_{c,j}-N_{c,i})\right)}{2}$, where $N_{c,l}$ is a fraction of contributors in group l. This assumes agents who are more successful in generating the public good for their group, are more likely to be imitated. As a consequence, cooperation is the sole target of the resulting group selection process.

Finally, every generation a mutation process is implemented. New values for α, β, T_c and T_p are drawn from normal distributions with mean equal to the current value and standard deviation μ. Table 1 summarizes all variables used in the model.

Table 1: Parameters as used in the model with the description and default values.

Variable	Description	Value
M	Number of groups	128
N	Population of each group	2, 4, 8, 16, 32, 64 and 128
m	Probability that an agent mimics another	0.01
μ	Mutation	0.01
ex	Probability of conflict	0.015
ε	Probability of making a mistake in decision making	0.02
c	Cost of investment in the public good	0.2
b	Relative benefits from the public good	0.5
p	Cost of being punished	0.8
k	Cost of punishing	0.2
α	Utility reduction for having a higher return than average	$-1 \leq \alpha \leq 1$
β	Utility increase for having a lower return than average	$-1 \leq \beta \leq \alpha$

Analysis

Our main interest was to explore the consequences of the inclusion of agents who make decisions based on expected utility, when that utility includes other-regarding preferences. The model was run fifty times for group sizes 2, 4, 8, 16, 32, 64, and 128 agents. Each model run consisted of 128 groups, and was simulated for 5000 time steps with relevant metrics (e.g., fraction cooperated/punish) calculated at the end of the run as the average of the last 1000 time steps. We used the following default values $c = k = 0.2$; $p = 0.8$; b=0.5; m=0.01; ε=0.015; and σ=0.01 in line with the values used by Boyd et al. (2003). The initial values for α, β, T_c and T_p are zero, representing the selfish rational actors. However, one of the 128 groups consisted initially of cooperative agents, again in line with Boyd et al. (2003) (i.e., $\alpha=\beta=T_c=T_p=1$)

We compared the results of our model with a replication of the original Boyd et al. (2003) model. With both models we found lower levels of cooperation in larger groups (Figure 1). Moreover, a major difference between the models occurs for the simulations with agents who can punish. For small group sizes, the level of cooperation with conditional cooperation is higher than the model with the pure strategies. When group size is 8 (no punish) or 32 (punish) the conditional cooperation model experienced a sharp drop in cooperation. The decline of cooperation in larger groups can be explained by the multitude of consequences of defection. If an agent defects in a group, other agents may reduce the trust in their group members and start defecting too. Thus without a change of the behavioral trait of the agent, the observed behavior is changing. This captures the idea that behavior is not only changing because of imitation, but also due to changes in expectations based on prior experiences.

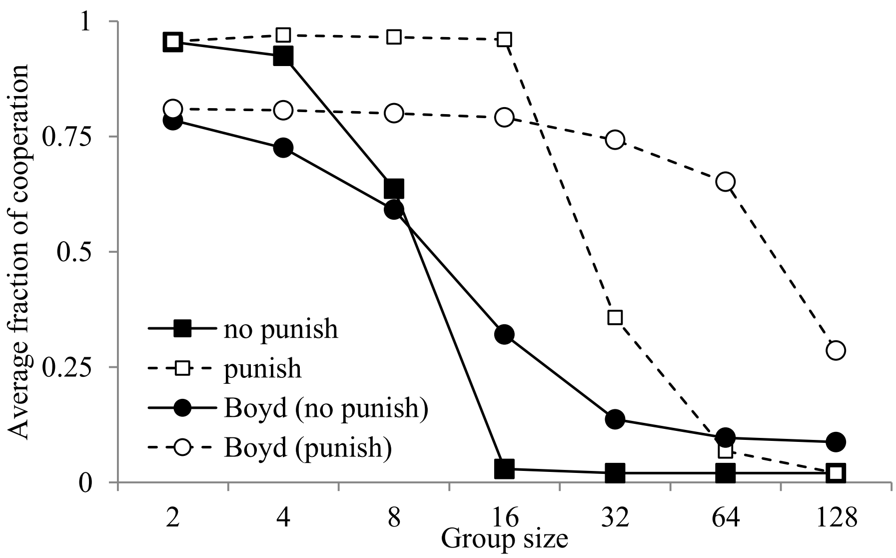

Figure 1: The average level of cooperation for 50 runs for different group sizes n. Results with a replication of the Boyd et al. (2003) model are compared with the model presented in this paper. For both models, two treatments are distinguished: no punishment and punishment.

Figure 2 shows the evolved values of α and β. We observed substantial differences between simulations with and without punishment. The evolved average behavioral type varied with increasing group sizes. For instance, in a group size of 2, agents who valued the earnings of others at a lower level than their own emerged in simulations with (and without) punishment. With a group size of 2 agents cooperate due to direct reciprocity forces. For group sizes 4 and 8 the agents who experience an evolution without punishment evolve norms with greater weight on shame and guilt. Groups who have more individuals with higher values of α and β are more likely to win intergroup conflicts. For group sizes of 16 and higher the penalty of earning more than average disappears for lower levels of cooperation.

When punishment is allowed, there is a higher level of cooperation, but this is not due to a higher level of altruism in the agents. Evolved agents are highly competitive and use punishment to increase their earnings compared to others. Groups cooperate due to the threat of punishment not because of evolved levels of prosocial behavior. The use of punishment in small groups (albeit at a lower level compared to pure strategy model) keeps defection levels low (Figure 3). However, larger groups are more vulnerable to mistakes or defection behaviour, which reduces trust in the cooperative tendency of group members (and conditional likelihood of social sanctions in response to defection) in subsequent time steps. The use of punishment did not appear particularly effective to enforce cooperation.

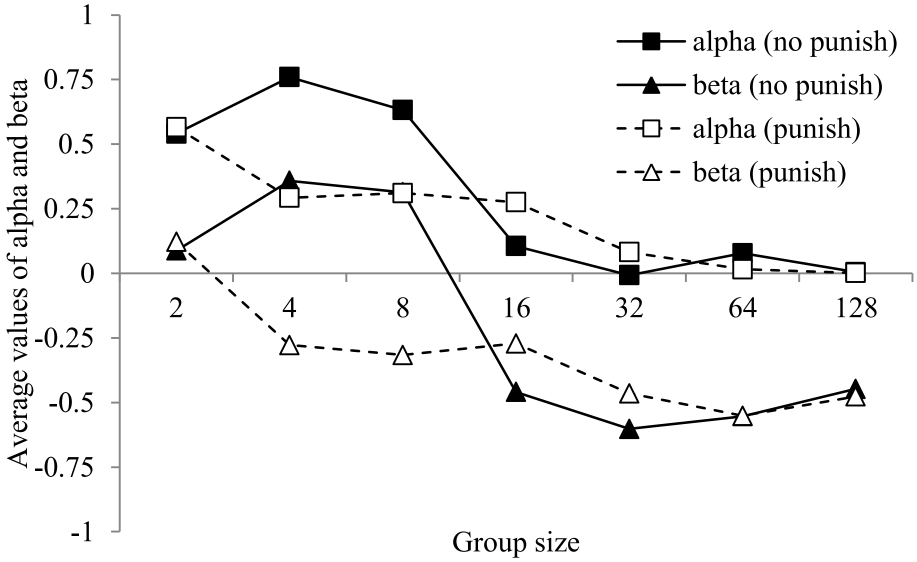

Figure 2: Average level of α and β for two conditions (with and without punishment), for different group sizes.

Figure 3 shows the difference in the degree of punishers in our model versus the results of a replication of Boyd et al. (2003). The Boyd et al. (2003) model experienced a share of punishers between 40% and 50% for groups up to 64 agents. Our model shows significantly lower levels of cooperation in contrast. High levels of cooperation are derived with modest levels of punishers. This can be explained by the punishing effect of conditional cooperation. Agents who do not trust their group members will start defecting. As a consequence, defectors are initially punished by punishers, and subsequently experience a much lower benefit of freeriding since group members have lower trust in the group. Note that punishing agents do not actually punish when there are only cooperative agents in the group.

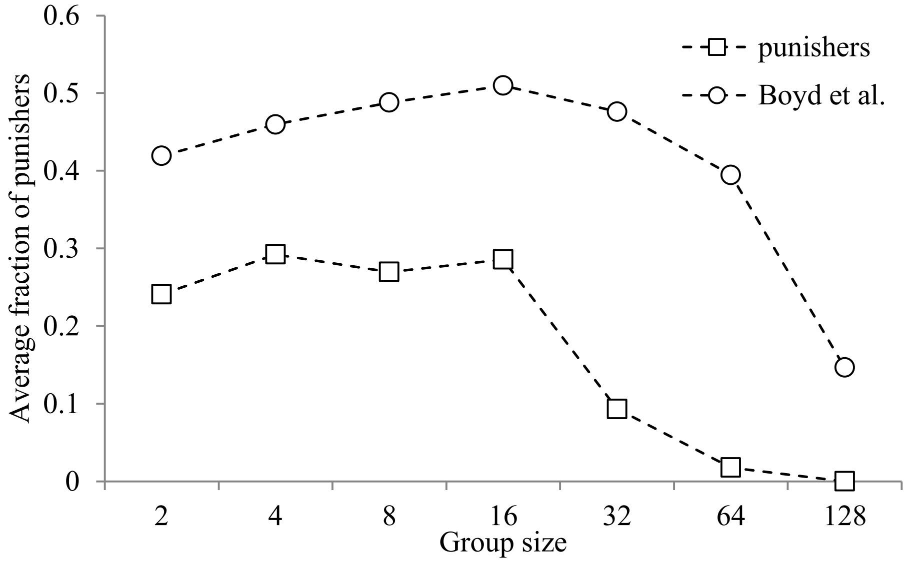

Figure 3: Fraction of punishing behavior within the whole population.

We will now report on some sensitivity analysis. Boyd et al. (2003) showed the sensitivity of the probability of inter group conflict, and the probability of imitating an agent from another group. We hypothesized that higher levels of group competition (a higher value of conflict with other groups) lead to more benefits for cooperative groups and a higher level of cooperative agents. Figures 4 and 5 show that this indeed the case. A higher level of conflict rates leads to cooperative behavior for groups of 16 agents versus 8 agents (no punishment), but the effect for simulations with punishment is modest.

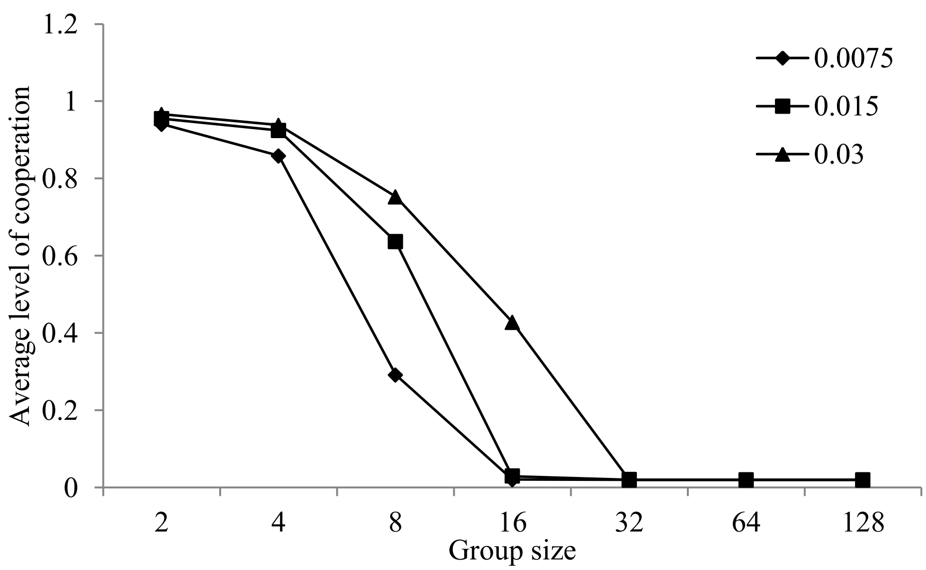

Figure 4. The evolved levels of cooperation with different levels of conflict rates (ex), and when there is no punishment.

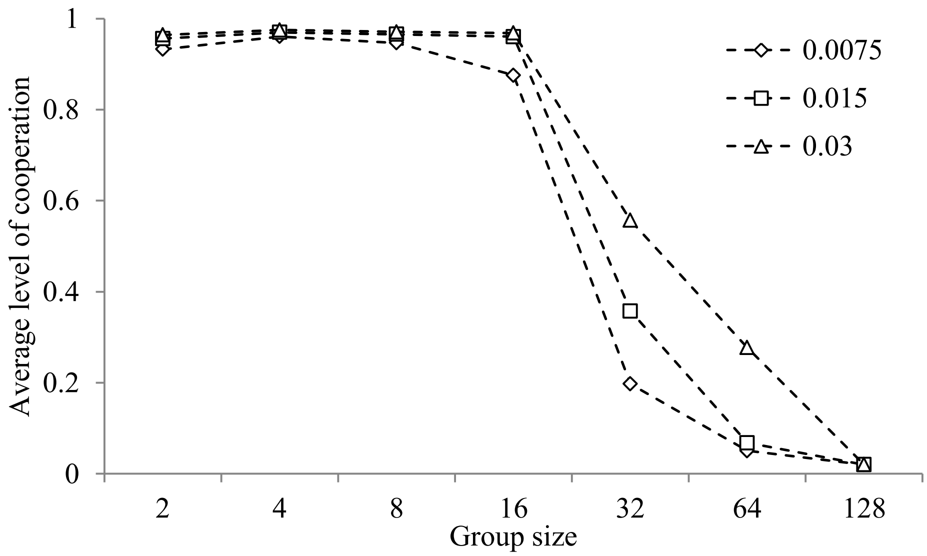

Figure 5. The evolved levels of cooperation with different levels of conflict rates (ex), and when there is punishment.

Each generation an agent considers imitating the behavior of one other agent. With a probability (1-*m*) an agent from her group is drawn and with a probability *m* an agent is drawn from an outside group. Generally speaking, increasing group mixing rates decreased cooperation levels since high-payoff defection strategies spread more quickly. Figures 6 and 7 show that this is indeed the case in our model. When agents cannot punish cooperation is abysmally small even with large mixing rates at a group size of 8. When punishment is allowed, there is some modest level of cooperation for group size 128 even when the mixing rate is very low.

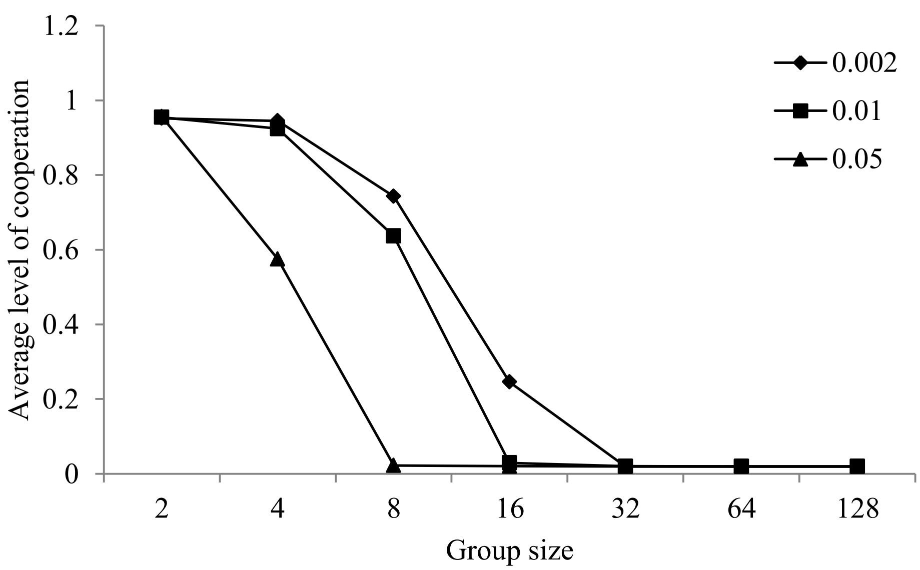

Figure 6. The evolved levels of cooperation with different levels of mixing rates (m), and when there is no punishment.

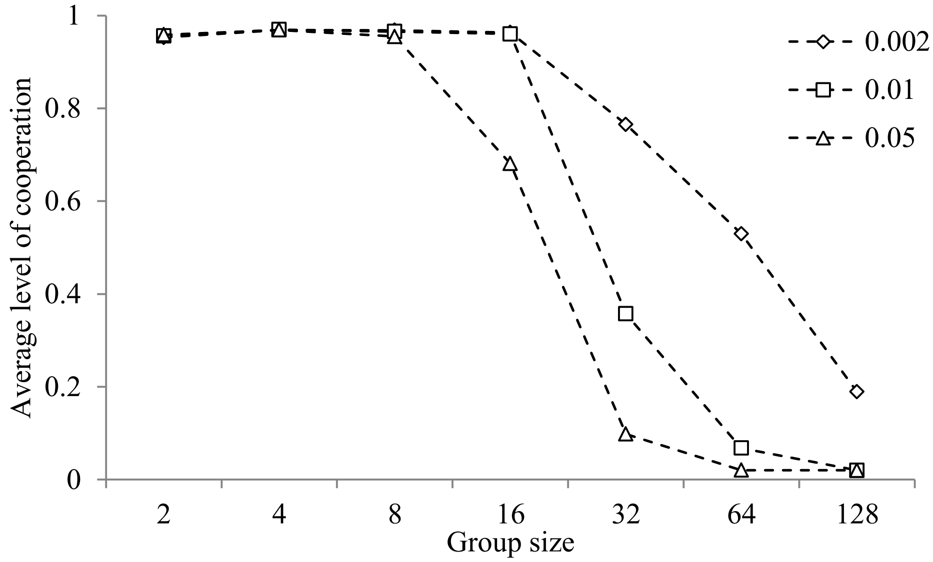

Figure 7. The evolved levels of cooperation with different levels of mixing rates (m), and when there is punishment.

Discussion

Experiments in behavioral economics show that humans make decisions in social dilemmas based on expectations of what others will do. Those experiments also show that humans generally take into account the earnings of others to evaluate the utility of material payoffs (Fischbacher et al., 2001). The inclusion of insights from behavioral economics into a model of cultural group selection of cooperation leads to different levels of cooperation. Cooperation levels are initially relatively higher in small groups (i.e., 8 or fewer agents). The primary driver at these group sizes is a combination of direct reciprocity and altruistic punishment. Interestingly, we also observed a significantly lower fraction of altruistic punishers when punishment was allowed compared to Boyd et al. (2003). Furthermore, we find a comparatively higher level of cooperation even without punishment compared to the results of Boyd et al. (2003). This contrast with the Boyd et al. (2003) model is likely due to the influence of social welfare norms in our agents that promote equitable payoffs within groups.

Cooperation levels drop quite quickly in larger groups (>16, no punishment; > 32, with punishment). This sharp decline of cooperation is caused by the overwhelming presence of second-order free-riders (i.e., non-punishing co-operators) that contribute to the erosion of social welfare norms through the combined effect of decreasing trust in likelihood of group sanctions and increasing chance of erroneous defection among group members. Since decision making is not only based on behavioral types, but also on the expectation of others, we observe a qualitatively different relationship between cooperation levels and group size.

We note here that this relationship is likely more ambiguous if we allow defecting agents the ability to levy sanctions against other defectors. In that case, we may expect greater trust in group sanctions and perhaps some optimal ratio of defecting punishers to altruistic non-punishers that maintains stable levels of cooperation.

We also observe that higher levels of cooperation in simulations with the option of punishment are not caused by higher levels of altruism of the agents. In fact, agents are somewhat competitive and use punishment to increase their earnings relative to others. Hence, introducing the option of punishment allows different types of agent behavior to evolve.

This analysis shows some of the consequences of including more behavioral complexity. How do our results relate with the high levels of cooperation we observe in large human societies? Conditional cooperation is sensitive to the erosion of trust in groups. When trust is lost it is difficult to regain it. This may explain why complex human societies are often characterized by different kinds of rituals, symbols, and social norms that may be especially useful for developing trust relationships. Others have recognized that human societies might be best represented as individual ecologies of games (Bednar and Page, 2007). Experimental research has shown that the alteration of public good games with a reciprocity game leads to a higher level of cooperation compared with playing only public good games (Milinski et al., 2002). Field work on the study of governance of the commons has shown that effective communities use conflict resolution mechanisms as a way to restore trust relationships (Ostrom, 1990).

By formulating agents as conditional strategists in a cultural group selection model, we find that cooperation is sensitive to the expectations of the behavior of others. In real societies, repeated activities are instituted to maintain and develop trust relationships. Future work may focus on the trade-offs in the costs of maintenance of social capital and the evolution of cooperative behavior.

Acknowledgement: We appreciate the support of the National Science Foundation for the grant BCS 0748632

References

Axelrod, R., 1984. The Evolution of Cooperation. New York: Basic Books.

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

Alexander, R. D., 1987. The Biology of Moral Systems. New York: Aldine de Gruyter.

Bednar, J., and S. Page 2007. Can Game(s) Theory Explain Culture? The Emergence of Cultural Behavior Within Multiple Games. *Rationality and Society* 19(1): 65–97.

Boyd, R., Gintis, H., Bowles, S., Richerson, P. J., 2003. The evolution of altruistic punishment, Proceedings of the National Academy of Science of the United States of America 100(6): 3531-3535.

Bowles, S. and H. Gintis (2004) The evolution of strong reciprocity: cooperation in heterogeneous populations, Theoretical Population Biology 65: 17–28.

Camerer, C. (2003) Behavioral game theory: experiments in strategic interaction, Princeton University Press, NJ.

Charness, G., and M. Rabin. 2002. Understanding social preferences with simple tests. *Quarterly Journal of Economics* 117: 817-869.

Fehr, E., and K. Schmidt. 1999. A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics* 114: 817-868.

Fehr, E., Gächter, S. 2002 Altruistic punishment in humans. Nature 415: 137-140.

Fischbacher, U., S. Gächter and E. Fehr (2001) Are people conditionally cooperative? Evidence from a public goods experiment, *Economics Letters* 71(3): 397-404

Gintis, H. 2000. Game theory evolving; a problem-centered introduction to modelling strategic interaction, Princeton University Press, NJ.

Hamilton, W. D., 1964. Genetical evolution of social behavior I and II. Journal of Theoretical Biology 7, 1-52.

Janssen, M.A. (2008) Evolution of cooperation in a one-shot prisoner's dilemma based on recognition of trustworthy and untrustworthy agents, *Journal of Economic Behavior and Organization*, 65: 458-471.
Milinski, M., D. Semmann and H.J. Krambeck 2002. Reputation helps solve the 'tragedy of the commons', *Nature* 415: 424-426
Nowak, M.A. 2006. Five Rules for the evolution of cooperation, Science 314: 1560-1563.
Nowak, M. A., Sigmund, K., 1998. Evolution of indirect reciprocity by image scoring. Nature 393, 573-577.
Ostrom, E. 1990. Governing the commons, Cambridge University Press, MA.
Ostrom, E., Walker, J., Gardner, R., 1992. Covenants with and without a sword: self-governance is possible. American Political Science Review 86(2): 404-417.
Trivers, R., 1971. The evolution of reciprocal altruism. Quarterly Review of Biology 46, 35-57.
Yamagishi, T., 1986. The provision of a sanctioning system as a public good. Journal of Personality and Social Psychology 51(1): 110-116.
Zahavi, A., 1977. The cost of honesty. Journal of Theoretical Biology 67, 603-605.