The Evolution of Other-Regarding Preferences in Networked Populations

Austen Mack-Crane[1] Jeremy Van Cleve[2]

27 August 2012

[1] Brown University, Providence, RI 02912

[2] Santa Fe Institue, Santa Fe, NM 87501

Abstract

We expand upon previous theoretical treatments of cooperation on networks by considering the effects of conditional strategies and continuous strategy space. We model agents as utility-maximizers, with utility functions over their own and their partner's payoffs, weighted by a parameter denoting other-regard. The other-regard parameter evolves over time through learning from nearest neighbors based on payoffs. In simulation experiments, we find that agents using a multiplicative form of the utility function, which allows conditional behavior, support higher levels of cooperation than those endowed with an additive utility function, which does not allow conditionality. We also find a positive effect on cooperation of the continuous strategy space, though only in the presence of the multiplicative utility function. Our results suggest that comparisons of the evolutionary viability of different utility functions can contribute to economic theories of social preferences.

1 Introduction

In many social interactions, cooperative behavior creates greater overall welfare than selfishness. Some of these situations may be loosely modeled by coordination, anti-coordination, or prisoner's dilemma games, which model how individual incentives can lead to inefficient group outcomes and suboptimal long-term individual payoffs. Yet, in human societies, group members can often achieve socially efficient or cooperative outcomes in these scenarios. This suggests two questions: how do the levels of cooperation we see in humans come about in the presence of conflicting individual incentives; and can we gain an understanding of the mechanisms of cooperation that would allow the design of social institutions favorable to it? In this paper, we address the first question by testing social structures and player characteristics that lead groups to resolve the conflict of incentives, while the second question falls in the purview of mechanism design (Fudenberg and Tirole [1991]).

Several mechanisms have been identified which allow the evolution of cooperative behavior in humans and other animals, some of which are based on individual behaviors and others on population structure.

Group selection supports cooperation because cooperative behavior advances group and thus too individual survival; kin selection supports behaviors that help relatives of the focal individual, who share genes; direct reciprocity describes cooperation in repeated interactions, which allow the emergence of conditionally cooperative behaviors; indirect reciprocity occurs when players can respond to others' reputations, an indirect signal of past behavior, and condition cooperation on good reputation; and network reciprocity describes cooperation in a networked population, where individuals interact repeatedly with a small, fixed set of neighbors, which benefits grouped cooperators (Nowak [2006]).

In this study we focus on direct and network reciprocity. Direct reciprocity in repeated games is able to support the evolution of cooperation because conditionally cooperative strategies can earn high payoffs when paired together and can resist invasion by defectors. This concept is introduced by Trivers [1971] and the canonical Tit for Tat strategy is discussed in Axelrod and Hamilton [1981]. Other work has developed a "win-stay lose-shift" strategy that outperforms tit-for-tat under some circumstances (Nowak and Sigmund [1993]).

Previous work studying prisoner's dilemma (PD) games on networks has found that cooperation can survive on simple graphs such as cycles (Eshel *et al.* [1998]), though networks of higher average degree lower cooperation's viability (Ohtsuki *et al.* [2006]). Ohtsuki *et al.* [2006] derive a rule that approximates the condition for the viability of cooperation on lattices: $\frac{b}{c} > k$, or the benefit-cost ratio of cooperation exceeds the average degree in the network. They show additionally that this rule is a necessary but not always sufficient condition on more complex network types. Specific network types, though, fare better than regular ones – scale-free networks allow cooperation for larger ranges of the game parameters than regular networks (Santos *et al.* [2006]; Santos and Pacheco [2005]).

Most of the literature on direct and network reciprocity makes use of the discrete prisoner's dilemma, but several studies consider the continuous version of this game. An example of a prisoner's dilemma payoff matrix, showing payoffs for the row player, reads:

	C	D
C	3	1
D	4	2

(1)

We can think of the prisoner's dilemma as a situation in which players have an endowment and are given the choice to pay a cost in order to transfer a benefit to their partner (Ohtsuki *et al.* [2006]). Thus, we can write the payoffs as:

	C	D
C	e+b-c	e-c
D	e+b	e

(2)

where e designates an endowment, b is the benefit of cooperation, and c is the cost. For the payoffs given above, $e = 2$, $b = 2$, and $c = 1$. To implement the continuous prisoner's dilemma, we write payoff functions as:

$$x_i = e_i - C(I_i) + B(I_j) \tag{3}$$

where x_i is player i's payoff, e_i is again endowment to player i, I_i is a player's cooperative investment, and $C(I_i)$ and $B(I_j)$ are cost and benefit as a function of investment level of the focal individual and its partner, respectively. These cost and benefit functions may be linear, as in Wahl and Nowak [1999], or one or both may be nonlinear, as in Killingback *et al.* [1999]. Continuous prisoner's dilemma games have been studied in structured populations (Killingback *et al.* [1999]; Ifti *et al.* [2004]) and in unstructured populations as repeated games (Wahl and Nowak [1999]; André and Day [2007]; Killingback and Doebeli [2002]), but to our knowledge this type of game has not been explored as a repeated game using a network to model population structure.

In this study we explore these three factors of cooperation - reciprocity, networks, and strategy space - in concert. This enables us to discern quantitatively the effect of reciprocity and strategy space on the evolutionary success of cooperation. We do so in the context of networked populations, the combination of which with reciprocity is relatively little explored. A networked population of reciprocators is realistic if we consider humans having the cognitive complexity to respond to past interactions and living in networks where social ties are important to daily life.

Work in economics has developed an indirect evolutionary approach to the evolution of strategies in social dilemmas and other games, which is well articulated by Bester and Guth [1998]. Instead of applying evolutionary dynamics to game strategies, as is the typical approach of evolutionary game theory, this approach has players maximize preferences and applies evolutionary dynamics to those preferences, selecting for those which lead to higher objective payoffs.

Several studies investigate the evolutionary stability of other-regarding preferences using this indirect approach. Bester and Guth [1998], as well as Heifetz *et al.* [2007], consider the dependence of other-regard on the type of payoffs used in the game. Alger [2010] studies the relationship between assortative matching and other-regard in the population, finding that spite, or negative other-regard, is stable for randomly mixed populations, while increasing assortativity leads to the stability of increasingly positive degrees of other-regard. (While other-regard may be positive or negative, in the remainder of this work we use the term to refer to positive other-regard.)

These studies and most others we have considered in this literature use an additive form of the other-regarding utility function to represent preferences. As we will show later, an additive utility function leads to rather simplistic, unconditional behavior in the prisoner's dilemma. Akçay *et al.* [2009] apply indirect evolutionary dynamics using a multiplicative utility function and find a positive degree of other-regard stable under direct reciprocity. In the present work, we expand upon this tradition by making a direct comparison of behavior stemming from additive and multiplicative utility functions. The use of utility functions can also contribute to our understanding of the evolution of cooperation, because they represent a particular class of behavior that results from rational choice based on a set of preferences, rather than unconstrained and less easily interpretable strategic behavior.

2 Model

In this study, we model agents interacting along the edges of a network, who choose their actions to maximize a utility function. These utility functions are over a player's own and her partner's payoffs, weighted by a parameter denoting other-regard. This parameter evolves across time steps according to a payoff-based imitation rule, and we observe the magnitude of the parameter and the overall level of cooperation in the population once an equilibrium is reached.

We consider additive and multiplicative forms of the utility function. In both equations, α is allowed to vary on $[0, 1]$. The additive form reads:

$$U_i(x_i, x_j) = (1 - \alpha)x_i + \alpha x_j \tag{4}$$

where x_i is player i's payoff and α denotes other-regard. Alternatively, we can use the multiplicative form:

$$U_i(x_i, x_j) = x_i^{(1-\alpha)} x_j^{\alpha} \tag{5}$$

In each time step, agents play the game with each of their neighbors, choosing actions synchronously. This gameplay entails forming an expectation of the partner's next action, and then choosing one's own action to maximize utility contingent on this expectation. In our model, players simply assume that a partner's action will remain the same as in the last round.

When agents maximize an additive utility function, their response to a partner's previous action is either unconditional cooperation or unconditional defection – there is a cutoff for the other-regard parameter above which the agent prefers complete selflessness and below which she favors complete selfishness, regardless of the partner's action. When agents have a multiplicative utility function, however, their response is a linear function (in the case of a continuous strategy space) of their partner's past action. At some moderate value of the other-regard parameter, the agent will exhibit perfect reciprocity, playing exactly the action that their partner played in the last round. At higher values, the slope of the response function increases past one, and the intercept, the response to a partner's zero contribution, becomes positive. For lower values of the parameter, the slope is less than one and the intercept decreases, such that the agent will contribute zero for nonzero contributions by her partner.

Once each player has chosen actions, payoffs are calculated, and then a fixed number of players are picked to update their other-regard values based on our imitation rule. Each player picked considers the players in her neighborhood, including herself, and copies the other-regard value associated with the highest payoff in the neighborhood. We include some amount of error in this copying, such that $\alpha_{new} \sim$Uniform$(\alpha - \epsilon, \alpha + \epsilon)$.

This evolutionary rule is the same as that employed by Nowak and May [1993] and Abramson and Kuperman [2001]. Other possibilities include choosing the type observed to gain highest average payoffs in the neighborhood, as in Eshel *et al.* [1998], or choosing a new type with probability proportional to its fitness, as in Ohtsuki *et al.* [2006]. We interpret our rule as modeling social learning in a networked population.

2.1 Simulation Experiments

To explore this model we employ simulations over a range of parameter values. Several parameters, based on initial explorations of the sensitivities of the model, we have held fixed: for all simulations reported here, the number of nodes is 1000, the error range, ϵ, is fixed at .025, and the initial distribution of the other-regard parameter is at 0, or full selfishness.

Our main parameters of interest are the benefit-cost ratio of cooperation and the average degree in the network. We expect, based on the rule derived in Ohtsuki *et al.* [2006], that these parameters will have strong effects on the evolved level of cooperation. Additionally, in order to explore the significance of reciprocity, we vary the number of times the game is repeated within each timestep. If the game is repeated more than once, players iteratively update their expectations about their partner's action and re-choose their own, until the final round of the game is played and the payoffs from that round are used to apply the evolutionary step. We also vary the number of players picked to evolve their other-regard parameter at each time step - in the condition of few players picked we expect that repitition of the game will have little effect, because the timescales of gameplay and update are sufficiently separate, while in a case of more synchronous update,

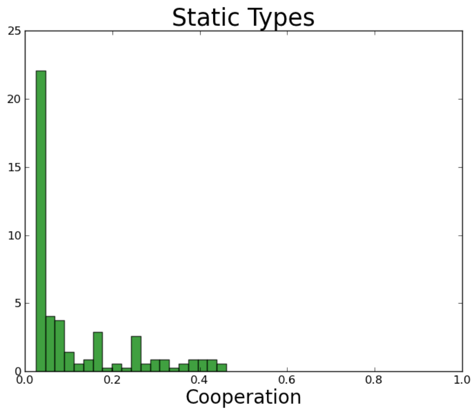

Figure 1: Distribution of equilibrium cooperation levels for 400 simulation runs using static types. Here data are restricted to the runs with $\frac{b}{c} < 4$

where more players are picked, we might expect to find an effect of game repitition on reaching efficient cooperative states.

In addition to running simulations of players with additive and multiplicative utility functions in the discrete and continuous prisoner's dilemma, we also do a set of simulations with the simpler representation of players as universal cooperators or universal defectors. We refer to this model as the static type model. This should help us compare our work with previous results on static type-players in networks, given that other parts of our model such as the evolutionary rule differ from previous studies.

In our simulations, we set the parameters randomly within reasonable ranges, as a convenient way to obtain independent distributions, and we perform at least 400 simulations for each of the discrete PD game, the continuous game (both including additive and multiplicative utility), and the static types model.

In simulations using the discrete prisoner's dilemma, we quantify the level of cooperation in the population by calculating the proportion of the interactions in the population which are cooperative. In the continuous game, we measure cooperation as the average contribution level across all interactions, which should be directly comparable to the former metric. We also measure the mean level of other-regard in the population. Simulations are run for 20000 time steps, which we determined was sufficient for populations to reach equilibrium, and these data are collected as averages over the last 5000 time steps of the simulation.

3 Results

We present three sets of results from our simulations. Figures 1-3 show distributions of equilibrium levels of cooperation across model runs with varied parameters, Figures 4-6 show cooperation levels as a function

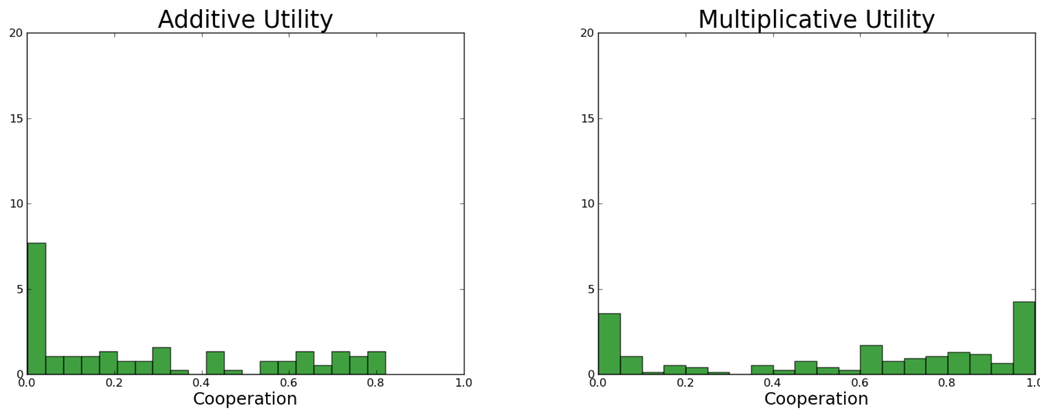

Figure 2: Distribution of equilibrium cooperation for simulations on the discrete prisoner's dilemma with utility types. Data are restricted to the runs with $\frac{b}{c} < 4$

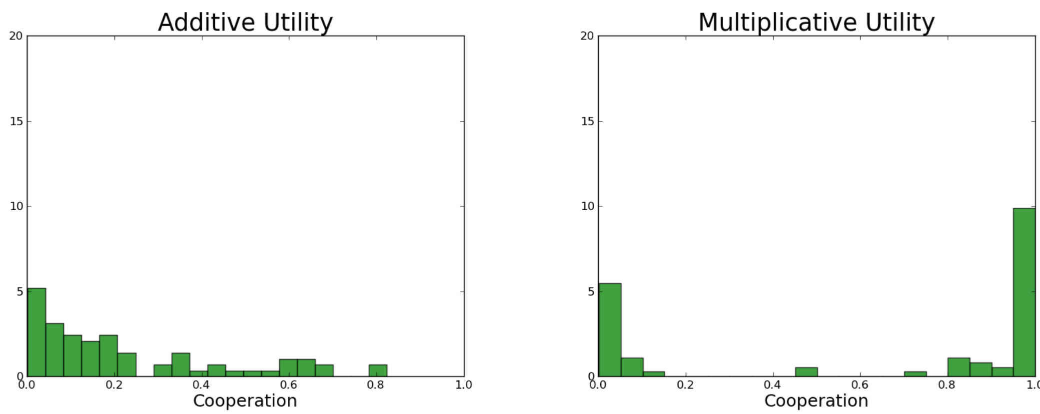

Figure 3: Distribution of equilibrium cooperation levels for 400 simulations of the continuous prisoner's dilemma with utility types. Data are restricted to the runs with $\frac{b}{c} < 4$

of benefit-cost ratio and average degree, and Table 1 shows mean levels of cooperation compared across models. In Figure 1 we see that, for the modest benefit-cost ratios shown, static type players support low levels of cooperation. Figure 2 shows the distribution for simulations of the discrete prisoner's dilemma, with both additive and multiplicative utility forms. Behavior under the multiplicative utility form supports higher levels of cooperation than does the additive utility function. We can also see that the distribution for additive utility in the discrete game shows higher cooperation than the model of the discrete game with static types - in the former, more mass lies above .4 and the cluster around zero is smaller.

In Figure 3 we compare additive and multiplicative utility functions for simulations using the continuous prisoner's dilemma game. We see again that the multiplicative utility function leads to higher levels of cooperation and fewer runs with intermediate levels than does the additive utility. The distribution of

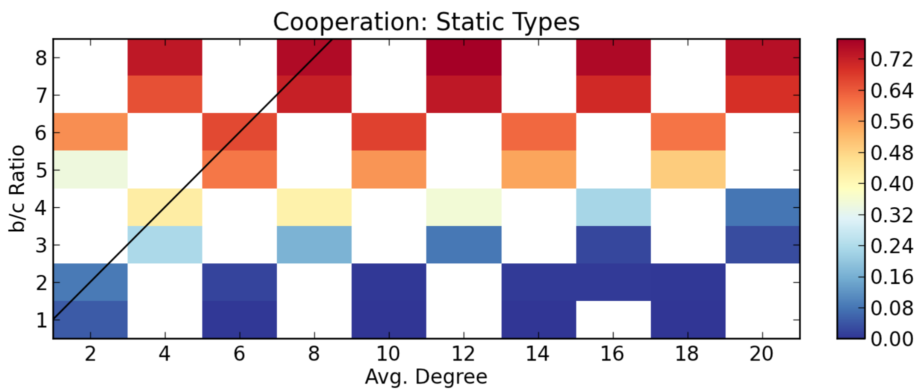

Figure 4: Interaction of the $\frac{b}{c}$ ratio and average degree for static types model. Mean level of cooperation for all runs with a given parameter combination is given by color in the square, with white indicating zero runs for those parameters. The black line indicates the threshold where $\frac{b}{c} = k$.

Model	Mean Cooperation	Standard Error
Static Types	.0703	.00595
Discrete Additive	.206	.018
Discrete Multiplicative	.638	.0203
Continuous Additive	.158	.00686
Continuous Multiplicative	.898	.0276

Table 1: Mean cooperation levels for runs with $\frac{b}{c} = 2$.

cooperation levels under multiplicative utility also shows higher cooperation than the same for the discrete prisoner's dilemma game.

Figures 4 - 6 show the relationship between the effects of $\frac{b}{c}$ ratio and average degree. There is a much more visible dependence of cooperation on the benefit-cost ratio than on the average degree. In some cases, such as the first plot of Figure 5, we can see that cooperation varies with average degree for medium levels of the benefit-cost ratio, but the vertical axis explains the great majority of the variance. The $\frac{b}{c} = k$ threshold, indicated on the plots by the black line, does not predict the patterns of cooperation found here – runs with a high benefit-cost ratio support high levels of cooperation, irrespective of their average degree.

To capture these patterns quantitatively, we compare average levels of cooperation sustained under each of our utility-type and game-type scenarios for certain benefit-cost ratios. Table 1 shows this comparison for $\frac{b}{c} = 2$. This table shows that the static types model leads to the lowest average level of cooperation, while the additive utility model for discrete and continuous prisoner's dilemma has medium levels of cooperation, and the multiplicative utility model has a high level of cooperation, which is higher in the continuous PD than in the discrete game.

We use the Mann-Whitney U-test to compare these models. We test for differences between the static

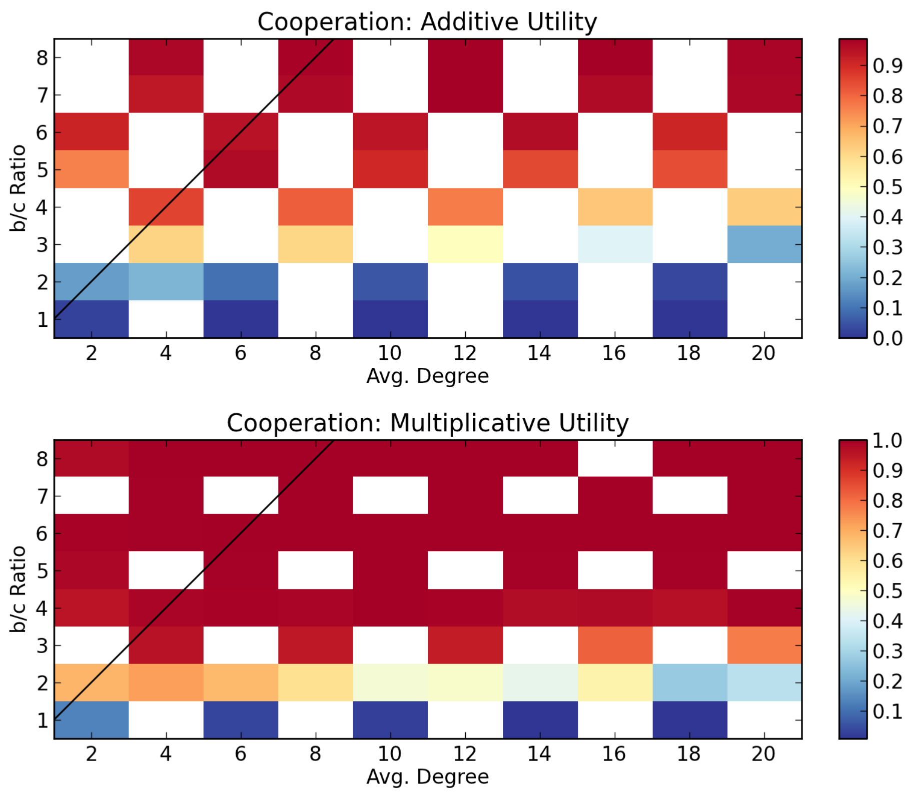

Figure 5: Discrete prisoner's dilemma. Interaction of the $\frac{b}{c}$ and average degree parameters.

types model and the discrete prisoner's dilemma with the additive utility function ($z = 5.498, P < .0001$), between additive and multiplicative utility functions in the discrete prisoner's dilemma ($z = 7.236, P < .0001$) and in the continuous PD ($z = 5.035, P < .0001$), and between discrete and continuous PD with the additive utility function ($z = 1.487, P = .137$) and with the multiplicative utility function ($z = 5.920, P < .0001$). These z-statistics indicate that all pairs are significantly different, at the .01 level with a bonferroni correction, except for the additive utility models with discrete or continuous PD. These results were computed for simulations where the benefit-cost ratio is 2; U-tests run on the data for $\frac{b}{c} = 4$ give identical patterns of significance.

Finally, we find no effect of the repitition of the game within time steps. Comparisons of mean levels of cooperation for the four combinations of high and low repitition and high and low fractions of players updating strategies synchronously show no statistically significant differences.

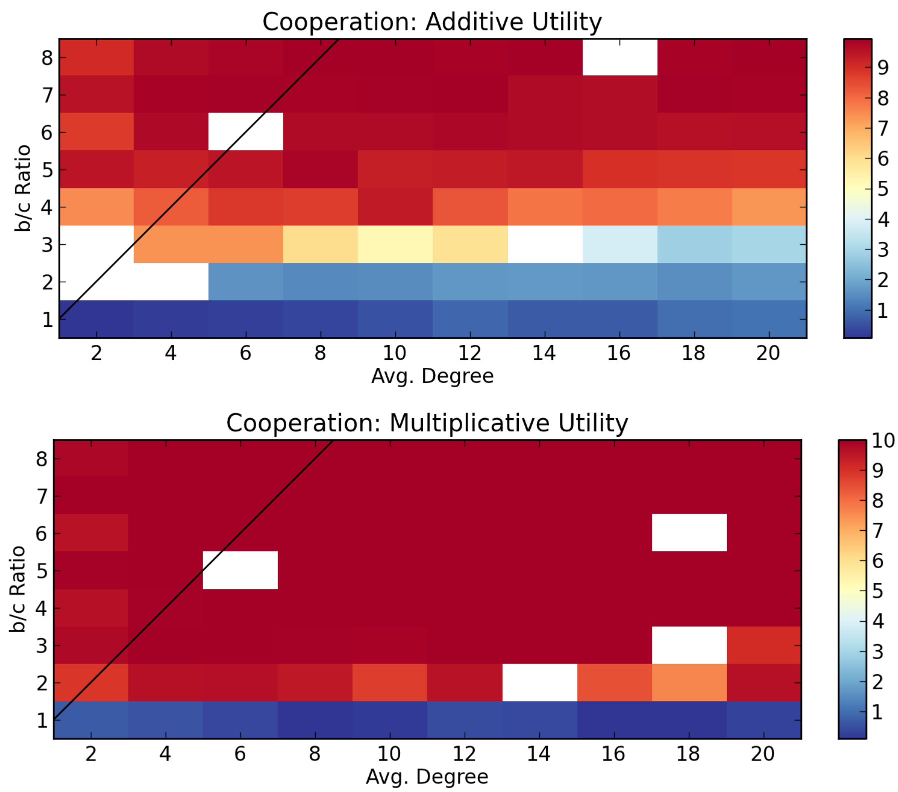

Figure 6: Continuous prisoner's dilemma. Interaction of the $\frac{b}{c}$ and average degree parameters.

4 Discussion

Our results make several statements: first, networked agents maximizing a multiplicative other-regarding utility function, which allows conditional behavior, show higher levels of cooperation than those using an additive utility function, which leads to unconditional behavior. Second, a continuous version of the PD game supports higher cooperation than the discrete game, but only when agents are capable of conditional behavior, with the multiplicative utility function. Last, we find that the static types model shows lower levels of cooperation than the additive utility model.

From the difference between additive and multiplicative utility functions, we can infer that conditional behavior enhances cooperation on networks. Unconditionally cooperative agents, in the additive utility model, can be exploited by defecting neighbors, who earn high payoffs and lower those of the cooperators, which increases the likelihood that a cooperator will "learn" to defect from a successful defector neighbor. In the multiplicative model, conditionally cooperative agents can defect in response to defecting neighbors, mitigating their own payoff loss as well as the defector's gain, which makes cooperation relatively more likely to spread.

By using relatively simple utility functions to define behavior, we limit the behavior space to linear response functions with one degree of freedom. More realistic behavior may be obtainable by allowing the slope and intercept of the response function to vary independently, but that approach does not have the advantage of linking cooperation to the internal representation of other-regard. We use utility functions to allow a straightforward connection to questions about social preferences in the economics literature, while also investigating the determinants of cooperative behavior.

Our second result suggests that a continuous strategy space aids conditionally other-regarding agents in reaching cooperative states. Because our agents form expectations of their partner's actions based on a past action, it is possible for other-regarding agents to be stuck in an equilibrium of defection due simply to the history of their interactions. In a discrete PD game, agents showing a moderate amount of other-regard will simply mirror their partner's last action, which doesn't allow a way out of a state of defection. The continuous strategy space allows more nuanced strategies, so a moderately other-regarding agent can respond to a partner's play with only a slight escalation, which over several rounds can lead to a high level of cooperation, but comes with little risk. This mechanism may allow the continuous strategy space to support cooperation more robustly than the discrete PD.

The third result we find, the difference between cooperation in the static types model and in the additive utility model, is more difficult to explain. Behaviorally, because the additive utility function leads to unconditional cooperation or defection, agents in these models should be indistinguishable. However, they differ in the way that they learn new behaviors – in both models, agents learn from the highest-performing individual in their neighborhood, but in the static types model they mutate with small probability to a random type, while in the utility model they simply introduce some error into the imitation. This difference could mean, for instance, that agents in the utility model mutate to cooperation with higher probability than in the static types model under some circumstances, which could explain the difference in cooperation levels.

In Figures 4-6, we see that the benefit-cost ratio predicts much of the variance in cooperation seen in our simulations, contrary to the finding of Ohtsuki *et al.* [2006], who find the comparison of $\frac{b}{c}$ and average degree important. There may be a few reasons for this discrepancy, which require further investigation. Our simulations, described in Section 2, use an update rule that is somewhat starker than that of Ohtsuki *et al.* [2006], in that ours always replicates the highest payoff viewed, while the latter does so probabilistically. Also, Ohtsuki *et al.* [2006] consider fixation probability as a measure of evolutionary viability, while we simply look at equilibrium levels of cooperation in our simulations. Either of these factors may explain the difference between our static types model and that of Ohtsuki *et al.* [2006].

If the prisoner's dilemma played in a social network is thought to be a realistic model of situations that humans commonly faced in the recent or distant past, then our simulation results indicate that preferences represented by a multiplicative utility function may better represent humans' other-regarding decision-making compared to an additive utility function. A multiplicative utility function supports higher levels of cooperation, which indicates that other-regarding individuals with this preference type have higher fitness than those with additive-type preferences. We suggest that the economic literature on evolving preferences would benefit from exploring the evolutionary stability of other-regard with multiplicative-type preferences. Additionally, an evolutionary modeling perspective could be employed in the broader study of social preferences to suggest which forms for utility functions may have evolved under different historical circumstances. In this vein, an interesting extension to the present work would be to model the evolution of other-regarding preferences in asymmetric social dilemma games, which could privilege those that give weight to fairness such as the model of Rabin [1993].

References

Guillermo Abramson and Marcelo Kuperman. Social games in a social network. *Physical Review E*, 63, 2001.

Erol Akçay, Jeremy Van Cleve, Marcus Feldman, and Joan Roughgarden. A theory for the evolution of other-regard integrating proximate and ultimate perspectives. *Proceedings of the National Academy of Sciences*, 106(45):19061–19066, 2009.

Ingela Alger. Public goods games, altruism, and evolution. *Journal of Public Economic Theory*, 12(4):789–813, 2010.

Jean-Baptiste André and Troy Day. Perfect reciprocity is the only evolutionarily stable strategy in the continuous iterated prisoner's dilemma. *Journal of Theoretical Biology*, 247:11–22, 2007.

Robert Axelrod and William Hamilton. The evolution of cooperation. *Science*, 211:1390–1396, 1981.

Helmut Bester and Werner Guth. Is altruism evolutionarily stable? *Journal of Economic Behavior and Organization*, 34:193–209, 1998.

Ilan Eshel, Larry Samuelson, and Avner Shaked. Altruists, egoists, and hooligans in a local interaction model. *The American Economic Review*, 88(1):157–179, 1998.

Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, Boston, 1991.

Aviad Heifetz, Chris Shannon, and Yossi Spiegel. The dynamic evolution of preferences. *Economic Theory*, 32:251–286, 2007.

Margarita Ifti, Timothy Killingback, and Michael Doebeli. Effects of neighborhood size and connectivity on the spatial continuous prisoner's dilemma. *Journal of Theoretical Biology*, 231:97–106, 2004.

Timothy Killingback and Michael Doebeli. The continuous prisoner's dilemma and the evolution of cooperation through reciprocial altruism with variable investment. *The American Naturalist*, 160(4):421–438, 2002.

Timothy Killingback, Michael Doebeli, and Nancy Knowlton. Variable investment, the continuous prisoner's dilemma, and the origin of cooperation. *Proceedings of the Royal Society B*, 266:1723–1728, 1999.

Martin Nowak and Robert May. The spatial dilemmas of evolution. *International Journal of Bifurcation and Chaos*, 3(1):35–78, 1993.

Martin Nowak and Karl Sigmund. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner's dilemma game. *Nature*, 364:56–58, 1993.

Martin Nowak. Five rules for the evolution of cooperation. *Science*, 314:1560–1563, 2006.

Hisashi Ohtsuki, Christoph Hauert, Erez Lieberman, and Martin Nowak. A simple rule for the evolution of cooperation on graphs and social networks. *Nature*, 441:502–505, 2006.

Matthew Rabin. Incorporating fairness into game theory and economics. *The American Economic Review*, 83(5):1281–1302, 1993.

F. C. Santos and J. M. Pacheco. Scale-free networks provide a unifying framework for the emergence of cooperation. *Physical Review Letters*, 95, 2005.

F. C. Santos, J. M. Pacheco, and Tom Lenaerts. Evolutionary dynamics of social dilemmas in structured heterogeneous populations. *Proceedings of the National Academy of Sciences*, 103(9):3490–3494, 2006.

Robert Trivers. The evolution of reciprocial altruism. *The Quarterly Review of Biology*, 46(1):35–57, 1971.

Lindi Wahl and Martin Nowak. The continuous prisoner's dilemma: I. linear reactive strategies. *Journal of Theoretical Biology*, 200:307–321, 1999.