

N-Player Prisoner's Dilemma in Multiple Groups: A Model of Multilevel Selection

Jeffrey A. Fletcher and Martin Zwick

Systems Science Ph.D. Program
Portland State University
Portland, Oregon 97207-0751
jeff@pdx.edu
zwickm@pdx.edu

Introduction

Simulations of the n-player Prisoner's Dilemma (PD) in populations consisting of multiple groups reveal that Simpson's paradox (1951) can emerge in such game-theoretic situations. In Simpson's paradox, as manifest here, the relative proportion of cooperators can *decrease* in each separate group, while the proportion of cooperators in the total population can nonetheless *increase*, at least transiently. The increase of altruistic behavior exhibited in these simulations is not based on reciprocal altruism (Trivers 1971), as there are no strategies (e.g. Tit-for-Tat) conditional on other players' past actions, nor does it depend on kin selection via inclusive fitness (Hamilton 1964), as there are no genomes. This model is very general in that it can represent both biological and social non-zero sum situations in which utility (fitness) depends upon both individual and group behavior. The two parameters of the PD in this model, which determine the gain in individual utility for defection and the dependence of utility on collective cooperation, are respectively analogous to within-group and between-group selective forces in multilevel selection theory. This work is more fully described in Fletcher and Zwick (2000).

The notion that a system (group) does better when it achieves cooperation among its parts (individuals), often against the self-interest of those parts, goes beyond just biological systems undergoing natural selection. It is applicable to hierarchical systems across a variety of fields. The non-zero sum nature of aggregation is general and optimization by subsystems often results in sub-optimization at a higher level. The PD is often used to model such non-zero sum situations. Like Simpson's paradox, the PD involves an anomaly of composition: individually-rational strategies, when aggregated, give a deficient collective outcome.

As Sober and Wilson (1998) have demonstrated, Simpson's paradox (even if not always identified as such) is important in understanding multilevel selection. These authors show (pp. 18-26) that this paradox can be derived from simple fitness functions for altruists and non-altruists in two populations. These functions amount to an n-player

PD (see Appendix A), although Sober and Wilson do not call attention to this fact. In this paper and in Fletcher and Zwick (2000), we make the connection between the PD and Simpson's paradox explicit. Our main finding is that Simpson's paradox emerges transiently, but for a wide range of PD parameter values, when a minimal group structure is imposed on an n-player PD. This result is produced in a model which involves an implicit competition between two groups and a simple n-player PD in each. The model is based on only two parameters which correlate with the within-group and between-group selection components in multilevel selection theory.

N-Player Prisoner's Dilemma

The n-player PD offers a straightforward way of thinking about the tension between the individual and group levels of selection. In real-world biological and social systems the effects of cooperation or defection are often distributed diffusely to other members of a group, i.e., they do not necessarily arise via pair-wise interactions. When there is a common and finite resource, each individual benefits by using more than its share of that resource, but when all players apply this individual rationality it can lead to collective irrationality. For example, each country that fishes international waters can increase its utility by taking more of the fish in this common resource, but as more and more countries overfish, the common stock is depleted beyond where it can quickly replenish and so in subsequent years all have less. This leads to decreased utility for both countries that overfish (defectors) and those that don't (cooperators).

An n-player PD involving the exploitation of a common resource (e.g., the fisheries example) is also known as the "tragedy of the commons" (Hardin 1968). A simple payoff scheme for such an n-player PD is illustrated by Figure 1. On the horizontal axis is the fraction of individuals cooperating for the common good. On the vertical axis is the average utility to each individual. For convenience, we assume a linear relationship between utility and percent cooperators. The upper line denotes the utility for a

defector (D) while the lower line is the utility for a cooperator (C). The defector’s line dominates the cooperator’s line, i.e., selfish individual behavior always has a higher utility than cooperating no matter what the fraction of cooperators. The resulting dynamic tends to decrease the number of cooperators within a group.

N-Player Prisoner’s Dilemma (Tragedy of the Commons)




Figure 1. Utility lines for defectors and cooperators as a function of the fraction of cooperators for a simple n-player prisoner’s dilemma (PD).

The deficient outcome of the PD here inheres in the fact that the utility to defectors when there is a minimum number of cooperators (point A) is lower than the utility to cooperators when there is a maximum number of cooperators (point B). So even though for a given state of the system an *individual* benefits more by defection than cooperation, still cooperators in a *group* of cooperators get more benefit than defectors in a group of defectors.

The “tragedy” (and what makes this a PD) is that whatever the current state of the system, individual rationality or individual selection favors defection, which tends to drive the system to a (boundary) equilibrium state less beneficial to all (point A). This state is a non-Pareto optimal and irrational collective outcome. To summarize algebraically: $U_D(f_c) > U_C(f_c)$ for all f_c causes f_c to decrease for all f_c , but $U_C(1.0) > U_D(0.0)$. The co-parallel lines used here are the simplest of many cooperator and defector utility curves that can satisfy these PD conditions.

Note that it may be tempting to think of cooperation as selfish rather than altruistic because a group of all cooperators gets more utility per individual than a group of all defectors, but this would be incorrect and misses the crux of the PD. In the 2-player PD, the players would be better off if they both cooperate, but defecting is still the rational *individual* strategy because the prisoners have no way to coordinate their actions and enforce any agreements to cooperate. Cooperating is always disadvantageous no matter what the other player does. So in the absence of

guarantees of cooperation by other players, cooperating is truly altruistic—it lowers one’s individual utility (fitness) while raising the benefit to others. The same situation holds in the n-player game. Given the absence of coordination between players, each player is better off to defect, but benefits others by not doing so in that the system is kept at a state with a higher fraction of cooperators. Of course, this is the dynamic for a single set of players, or for a multi-group system viewed at the intra-group level. As we shall see, at the higher level of organization, i.e. that of the total population which includes all groups, cooperators can thrive, at least for a while, despite their inferior individual utility (fitness).

The Model

In the simplest form of the model there are two groups with no migration between them. These groups initially are the same size and vary only in their fraction of cooperators and defectors. There are no other strategies besides always-cooperate (C) and always-defect (D). We follow the percentage of cooperators in each group and across the whole population. In each group, the n-player PD (see Figure 1) is described by utility functions for cooperation and defection, which are dependent on the fraction of cooperators in each group, i.e.,

$$U_{Ci} = m f_{Ci} + b_C$$

$$U_{Di} = m f_{Ci} + b_D$$

where U_{Ci} and U_{Di} are the utility for cooperators and defector respectively within a group i ; m is the slope of both the defector and cooperator utility lines; f_{Ci} is the fraction of cooperators in group i ; b_C and b_D are the intercepts for the cooperator’s utility line and the defector’s utility line respectively.

There are two parameters in these utility functions. The first is the slope of the utility lines, which for simplicity are linear and parallel. The slope of both lines affects the disparity in utility for groups of different composition and can be thought of as the magnitude of the group-level selective force. At this level groups containing more cooperators have the advantage. The second parameter is the difference in the intercept for the cooperator’s and defector’s utility lines. Because U_C and U_D have the same slope, the difference in intercepts is the vertical displacement between them at all levels of cooperation. For simplicity we always use $b_C = 0$ so the difference is b_D (≥ 0 in our simulations). This disparity in utility for defectors vs. cooperators within a group can be thought of as the magnitude of the individual-level selective force within each group. At this level defectors have the advantage over cooperators.

From the fact that $b_C = 0$, it follows that here the condition for a PD (e.g., point B above point A in Figure 1) is $m > b_D$. In all runs reported in this paper this condition is

satisfied.

In this model, at each timestep the following action is implemented: Within each group the number of cooperators is increased in proportion to the cooperators' utility based on the group's composition. Similarly, the number of defectors is increased in proportion to the defectors' utility based on the group's composition. Specifically, the increase of each strategy within a group equals the number of individuals utilizing the strategy times its utility payoff per individual:

$$N_{Ci}(t+1) = N_{Ci}(t) + N_{Ci}(t) U_{Ci}$$

$$N_{Di}(t+1) = N_{Di}(t) + N_{Di}(t) U_{Di}$$

where N_{Ci} and N_{Di} are the number of cooperators and defectors respectively in group i . To aid in comparisons among runs, the population of each group is proportionally scaled back (preserving the ratio of cooperators and defectors) so that the total population size matches the original total. Scaling does not do anything substantive. For convenience we also define $N_C = N_{C1} + N_{C2}$ and $N_D = N_{D1} + N_{D2}$, and define $N_1 = N_{C1} + N_{D1}$, and define N_2 and N_3 similarly. All of our experiments included here involve two groups with initial conditions of 90 defectors and 10 cooperators in group 1 and 90 cooperators and 10 defectors in group 2, but similar results are obtained with other initial distributions.

Because the utility for defectors is always higher than that for cooperators, in the long run defectors will dominate both in each group and across the whole population. Yet while the percentage of cooperators decreases within each group, the overall percentage of cooperators in the whole population can increase. This effect is transient without mechanisms for reestablishing variation between groups. The effect depends upon initial conditions. Specifically, given that in our model $N_1 = N_2 = N_C = N_D$, the condition for Simpson's paradox to emerge is:

$$m / b_D > N_i^2 / (N_{Ci} - N_{Di})^2$$

where $i = 1$ or 2 (see Fletcher and Zwick, 2000, Appendix A). This equation also implies that m must be greater than b_D , and thus the PD is a necessary (but not sufficient) condition for Simpson's paradox.

Experiments and Results

We explore how combinations of our two parameters affect the magnitude and longevity of the Simpson's paradox effect. Figure 2 shows the beginning of a typical run with two groups where Simpson's paradox is evident. Here the slope of the utility lines are 0.01 and the intercept of the defectors' line is 0.003. Notice that although the percentage of cooperators is decreasing in each group monotonically, the total percentage of cooperators is increasing until timestep 328. Run 1 in Table 1 shows the maximum total percent cooperators reached is 66.6% (initially 50%) for this

same run. The overall increase in percent cooperation, despite decreasing within each group, is due to group 2 (cooperator dominated) expanding, i.e. from 100.0 to 170.9 players by the time the maximum is reached, while group 1 (defector dominated) shrinks from 100.0 to 29.1 players. The percentage of cooperators in group 2 is continuously decreasing (see Figure 2) which after timestep 328 causes the overall percentage of cooperators to also decrease. By timestep 4,000 (not shown) the overall percentage of cooperators is essentially zero ($< 0.01\%$).

Percentage of Cooperators vs. Time
(Slope = 0.01 and Intercept = 0.003)




Figure 2. Percentage of cooperators in group 1, group 2, and total. (Results of Run1 in Table 1.)

Run	Slope	Intercept	Max % C	Time at max
1	0.01	0.003	66.6	328
2	0.05	0.003	85.5	121
3	0.01	0.001	82.4	520
4	1.0	0.3	70.2	6
5	0.0001	0.00003	66.6	32,501
6	0.01	0.008	50.0	0

Table 1. Results of various runs with varied slope and intercept.

Table 1 shows the results of several other runs with varying slope (m) and intercept (b_D), but all with the same initial population sizes. Runs 2 and 3 show the effect of increasing the slope and decreasing the intercept respectively compared to Run 1. Both changes cause the maximum overall cooperation reached to increase, but the time it takes to reach these peaks is quite different. Increasing the slope in Run 2 causes group 2 to dominate sooner and the maximum cooperation is attained at timestep 121 compared to timestep 328 in Run 1. Increasing the slope can be thought of as increasing *between-group* selection.

Decreasing the intercept in Run 3 on the other hand, decreases the advantage defectors have over cooperators *within* each group. This also causes the maximum level of cooperation reached to be higher than in Run 1 (82.4% vs. 66.6%), but time to reach this peak is delayed to timestep 520 as compared to timestep 328 in Run 1. Runs 4 and 5 demonstrate that the Simpson's paradox effect can be seen across a wide variety of slope and intercept values. Run 6 demonstrates that a PD, i.e. $m > b_D$, is not a sufficient condition for producing Simpson's paradox.

Conclusion

The model described in this paper is simpler than models of reciprocal altruism based on the iterated 2-player and n-player PD in that here there are no actions (e.g. Tit-for-Tat) conditioned on past behaviors of other players. It is also more abstract than inclusive fitness models in that there are no genomes. Although one could interpret the C and D strategies as alleles of a single gene and interpret these results in terms of inclusive fitness, this would obfuscate the tension between individual and group-level selective forces that our n-player PD model makes explicit.

Although the n-player PD has been utilized in studies of reciprocal altruism (see Boyd and Richerson 1988, Joshi 1987, Motro 1991), its centrality to multilevel selection has not, to the best of our knowledge, been explicitly acknowledged. An overall increase of altruism (in the n-player PD) merely requires a suitable higher level of organization and appropriate PD parameters values. This is consistent with multilevel selection theory. Increasing the slope (group selection parameter) increases the disparity between group size such that the cooperator-dominated group increases and this accounts for the overall increase in cooperators. Decreasing the intercept (individual selection parameter) causes cooperators within each group to be sustained longer and therefore also contributes to an increase in overall cooperators. This simple model lets us tease out the within-group and between-group components of utility (fitness) and is applicable to both biological and social systems in which there is competition at multiple levels.

Appendix A

Sober and Wilson's fitness functions (1998, p. 20) are:

$$W_A = X - c + [b(np - 1)/(n-1)]$$

$$W_S = X + [b(np)/(n-1)]$$

where W_A and W_S are the fitness of altruists and non-altruists respectively; X is the base fitness; c is the cost to cooperators for providing benefit b (which is distributed to all n group members besides themselves); and p is the fraction of altruists in the group.

In our model we assume the benefit a cooperator produces is distributed evenly to all group members, *including the cooperator*. This is consistent with many forms of altruism that contribute to the common fitness of a group and is not significantly different from the model used by Sober and Wilson, especially for reasonably large n . Using this assumption, Sober and Wilson's equations become:

$$W_A = X - c + bp$$

$$W_S = X + bp$$

It is easy to see that these equations represent the same n-player PD as in our model where, $U_C = W_A$; $U_D = W_S$; $f_C = p$; $m = b$; $b_C = X - c$; and $b_D = X$. Note that the difference in the intercept ($b_D - b_C$) = c is the cost to an individual cooperator, and the benefit produced by a cooperator is m , which corresponds to the group-level selective force.

References

- Boyd, R. and Richerson, P.J. 1988. The Evolution of Reciprocity in Sizeable Groups. *Journal of Theoretical Biology* 132:337-356.
- Fletcher, J.A. and Zwick, M. 2000. Simpson's Paradox can Emerge from the N-Player Prisoner's Dilemma: Implications for the Evolution of Altruistic Behavior. In *Proceedings of The World Congress of the Systems Sciences and ISSS 2000*, Allen, J.K. and Wilby, J.M. eds, Toronto, Canada: International Society for the Systems Sciences. Forthcoming.
- Hamilton, W.D. 1964. The Genetical Evolution of Social Behavior I and II. *Journal of Theoretical Biology* 7:1-52.
- Hardin, G. 1968. The Tragedy of the Commons. *Science* 162:1243-48.
- Joshi, N. V. 1987. Evolution of Cooperation by Reciprocation within Structured Demes. *Journal of Genetics* 66(1):69-84.
- Motro, U. 1991. Cooperation and Defection: Playing the Field and the ESS. *Journal of Theoretical Biology* 151:145-154.
- Simpson, E.H. 1951. The Interpretation of Interaction in Contingency Tables. *Journal of the Royal Statistical Society B* 13:238-241.
- Sober, E. and Wilson, D.S. 1998. *Unto Others, The Evolution and Psychology of Unselfish Behavior*. Cambridge, MA:Harvard University Press.
- Trivers, R.L. 1971. The Evolution of Reciprocal Altruism. *Quarterly Review of Biology* 46:35-57.