Alana Schick
43320027, ISCI 330
Apr. 12, 2007

The Evolution of Cooperation: Putting gTheory to the Test

Evolution by natural selection implies that individuals with a better chance of surviving and reproducing (ie. fitness) increase in frequency. Altruism is a behaviour characterized by acting cooperatively to benefits others. More specifically, weak altruism is behaviour that benefits the others more than the altruist, while strong altruism is behaviour that comes at some cost to the altruist (Campbell, 1983). Since this behaviour tends to increase the fitness of others, one would think that altruism would be removed from populations via natural selection. Yet somehow cooperative behaviour can be observed in all societies and many species. For this reason, the evolution of cooperation is a problem that has been perplexing to scientists of many disciplines for decades. In game theoretic terms, the question becomes; what circumstances must be present for an agent to choose a strategy that is considered irrational? In other words, when would an agent not act to optimize their immediate utility?

There are several intuitive explanations for many situations, however, each explanation requires some special circumstances. These include explanations involving group selection, kin selection and reciprocity. The shortcomings and restrictions of these models will be briefly presented below. Some more recent models, however, have explored a mechanism that depends on similarity between agents. This similarity is apparent through a ‘tag’ that each agent carries, like a trait that is observable by individuals in a population. These models are computer simulations, whereas the trials conducted in this experiment use humans as test subjects. This is achievable through a program called gTheory, which allows the user to specify conditions and test predictions in nature that are based on previous models.

One explanation for the existence of cooperation is group selection, which requires that there is some synergistic effect when cooperative agents interact. In other words, the benefit from the combined action is greater than the sum of each individuals action. A simple example of this involves a pack of wolves hunting a deer. In this case, five wolves hunting together could kill one deer and have a hearty dinner, but if they hunted individually, no deer would be caught and the wolves would starve. The problem with this explanation is that among the group of cooperators, a selfish free-loading agent would have the highest fitness and eventually the cooperators would be replaced by non-cooperators.

A similar explanation is kin selection, which requires that the cost an agent pays to cooperate is less than the benefit gained by the agent's offspring. This makes sense in biological terms of "seed-spreading", where the ultimate goal of every agent is to pass on genetic material to future generations. For example, if a parent could make a sacrifice to ensure the survival of it's offspring, they would gain a higher utility than vice versa. The problem with this explanation is that cooperation decreases rapidly as the degree of relatedness declines (Campbell, 1983). Furthermore, it requires that individuals have the capacity to locate and recognize kin, as well as strategize accordingly.

One last explanation involves reciprocity, either directly, or indirectly. This explanation requires that the cost of cooperating is less than the benefit gained from other individuals observing this cooperation and reciprocating. The main limitation is that agents must have the capacity to remember interactions and their outcomes, which is only realistic for reduced population sizes (Trivers, 1971). Also, there would have to be an infinite number of interactions for this mechanism to work. If an agent knew there was a finite number of interactions, he would reason to defect in the last one, and then every other interaction before that.

A recent study conducted by Riolo et al. (2001) explores the tag-based mechanism for cooperation. They use agent-based computer simulations to show that cooperation is maintained without kin selection or reciprocity. The model consists of agents with two traits; a tag, or number between zero and one, and a tolerance value which specifies the range of tags that are cooperated with. Agents are randomly designated as donors or recipients and then paired up to interact. If the tag of the recipient is within the range of the tag of the donor, the donor pays a cost, or increase in fitness, to the recipient. This model does not require that agents can observe and remember each other's actions, nor does it require continuing interactions between the same agents. The problem with this mechanism for interaction is similar to that of the 'green beard' effect (Roberts and Sherratt, 2002). In order for one green beard to help another green beard, there must be a link between altruistic behaviour and having a green beard. The mechanism in the Riolo et al (2001) model requires that a sufficient tag similarity will lead agents to behave cooperatively. In nature, this could apply to situations where a tag, or trait, does in fact specify behaviour, but falls apart when this link is not present. In other words, if agents could choose whether or not to cooperate against their own tag, they would do better to defect.

Upon realization of this pitfall, Axelrod et al. (2004) developed a model in which tags and strategies are not fixed, but instead allowed to evolve independently. This model consisted of 2000 individuals with similar characteristics to those modeled by Riolo et al. (2001). Each individual can have one of four tags and a specific behaviour, assigned at random. Agents were placed at random throughout a 50x50 lattice and then played one round of prisoner's dilemma with each of their four adjacent neighbors. The placing of agents into a lattice is meant to represent what is referred to as "viscosity" of a population (Mitteldorf and Wilson, 2000). This basically means that agents are restricted to interacting with agents that are physically relatively close to them. When interacting, giving help causes a decrease in the donor's fitness, while receiving help has a greater benefit. Agents then spawn an identical individual with probability specified by their fitness after a complete round of interaction. Individuals are then removed from the population at random to return the population size to 2000. The fixed population size is meant to model the competition for resources. After running this simulation for 2000 periods, the average amount of cooperation was found to be 91% (Axelrod et al., 2004). Upon varying many parameters, including number of tags, mutation rate, lattice size, and the length of the run, at least 68% of the choices were cooperative in all cases, demonstrating a plausible mechanism by which altruism could evolve in nature (Axelrod et al., 2004). The study concludes by stating that heritable tags and spatial restriction, or viscosity, are necessary conditions for altruism to evolve.

Actually testing this mechanism in nature requires more than just a computer simulation. It involves replacing the programmed agents with humans and observing the behaviour under similar conditions to those described above. To test if either tags, viscosity, or both are necessary and sufficient conditions for cooperation, a computer program called gTheory is utilized. This is a program that allows the user to set up a game of specified parameters between agents that connect to the game via the internet. To model the spatial restriction, players are organized into a 4x4 lattice. If there are not enough human players to occupy the spaces, gTheory will have computer players fill in the empty spots, posing as actual players. The specific strategy the computer players adopt can be specified by the user. In this example it was chosen to be the "most successful neighbor" strategy in which the computer player chooses the strategy from the previous round that acquired the highest payoff. Once the lattice has been filled, the players interact with each of their four adjacent neighbors in one game of prisoner's dilemma. This game is used to allow altruism, but not direct reciprocity (Axelrod et al. 2004). The following matrix is used to determine the payoff acquired.

Table 1. Payoff matrix.

	C	D
C	(3,3)	(0,5)
D	(5,0)	(1,1)

In order to model the tag based component, agents are given a choice of whether they are tagged as 'red' or tagged as 'blue'. There is no fitness difference between the two. Note that this is different from the model proposed by Axelrod et al. (2004), where the expression of a tag could not be controlled by the individual, but instead an unavoidable display of an individual's genotype. Here, a strategy consists of a tag and an action to take against both the same and the opposite tag, shown in Table 2. Two groups of students (ISCI 422 and ISCI 330) are subjected to three different games, each with different conditions. One game involves tags and fixed partners, the second involves tags and random partners, while the third involves no tags and random partners. Each game is played for 15 rounds, but it is important that the players do no have this information. If the players

Table 2. Strategy Summary (C=Cooperate, D=Defect).

Strategy	Tag	Action against Red	Action against Blue
A	Red	C	C
B	Red	C	D
C	Red	D	C
D	Red	D	D
E	Blue	C	C
F	Blue	C	D
G	Blue	D	C
H	Blue	D	D

are not under the impression that the game could end at any moment, they would reason to defect in the last round and then the second to last round, all the way to the first round, thus cooperation could not evolve because it would never be present.

Since previous models have shown that both spatial restriction and tags are necessary for cooperation, it is predicted that the games with tags and fixed partners will see an eventual increase in cooperative behaviour towards the end of the 15 rounds. The remaining games are predicted to show little to no cooperation. More specifically, the ISCI 330 tags and fixed partner game will show a more rapid increase in cooperation than the ISCI 422 tags and fixed partner game. This prediction is based on the fact that the game theory class will be quicker to realize the prisoner's dilemma setting they are in and reason that cooperating with their partners will yield a higher payoff than always defecting.

Figure 1 shows the results of the games involving tags and fixed partners. The patterns between groups are surprisingly similar, but do show cooperation to increase as predicted. Also contrary to predicted, ISCI 330 showed a lower frequency of cooperation under these conditions. Both of the games with no tags and random partners showed cooperation to remain close to zero. Interestingly, there was one round in the ISCI – no tags and random partners where the number of cooperators jumped from

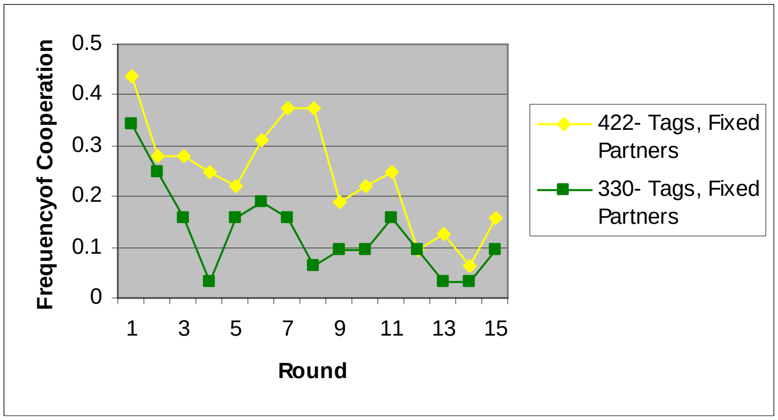

Figure 1. Cooperation in games with tags and fixed partners.

two to eight, then returned to zero (shown in figure 2). Since it is pretty unlikely that half of the players would just randomly cooperate one round then return to defection, there may have been some external influence. For example, someone may have reacted to the amount of defection by a quiet chuckle, provoking others in the room to reconsider their selfish behaviour. Figure 3 shows the results from the games played with tags and random partners. Note that the frequency of cooperation is generally less than the fixed partner case. One further thing to note is that the proportion of cooperation was calculated by the following formula:

A+E+ ½ (B+C+F+G), where the capital letters denote the proportion that each strategy was chosen.

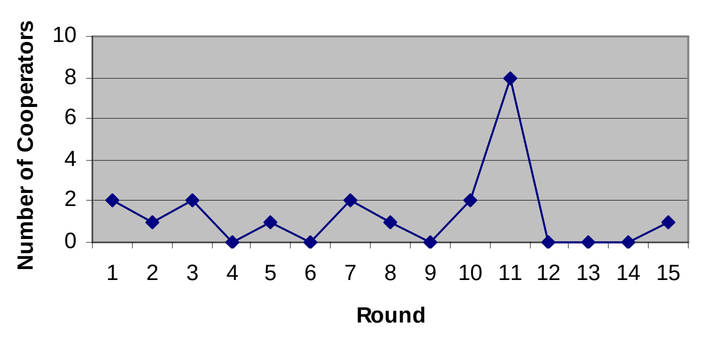

Figure 2. ISCI 330 - No tags, random partners.

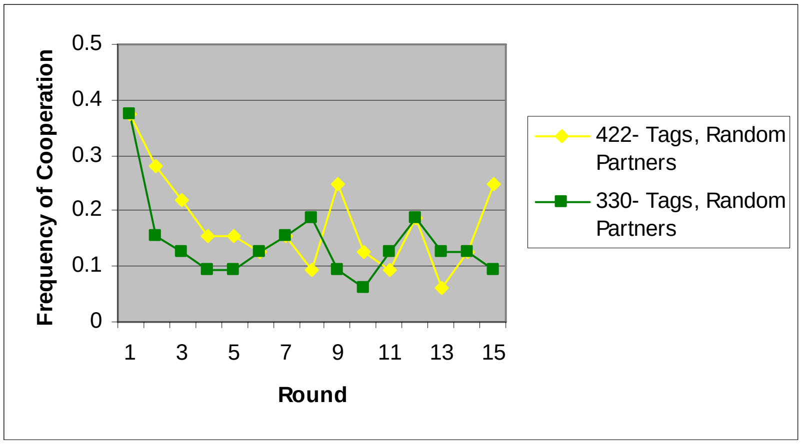

Figure 1. Cooperation in games with tags and random partners.

This is not completely accurate because the proportion of cooperation from a player depends on the tags of that players partners. For example, if an individual suspects that his neighbors might be all tagged as red, he may want to choose to defect against red and cooperate against blue. This would result in zero cooperative actions, but it would be calculated in this method as half cooperation and half defection.

There are a few major problems with the set up of these experiments. First of all, the computer's strategy is specified to play the "most successful neighbor" strategy. This strategy acts as a barrier towards cooperation because the computer will almost always defect. If we begin with all defectors, the computer will defect. In order for cooperation to evolve, one player must pay a cost at first to demonstrate the behaviour in an attempt to spread the strategy. Since the computer player would never pay this cost, and the benefits from cooperation would not be reaped until there were many of them, cooperation would be difficult to spread in the presence of this computer strategy.

Another major problem is the incentive of receiving a prize for achieving the highest payoff. Since rational agents, by definition, only care about their own payoff, introducing a reward for winning is implying that agents must care about the payoffs of other agents. For example, consider two agents playing a repeated prisoner's dilemma game. If an agent wanted to maximize his utility, he might be willing to pay the cost of being the first to cooperate if the other agent followed in suit. If he just wanted to ensure that the other agents payoff was less than or equal to his own, however, he would always defect, regardless of the fact that he could gain a higher payoff by cooperating. In this way, the reward incentive provides yet another barrier to cooperation.

Perhaps if the two barriers discussed above were decreased or removed from the system, a similar experiment would yield much different results. Furthermore, the sample size is too small and the number of trials too few to yield any conclusive results. These experiments, however, did manage to demonstrate the ability of the gTheory program to test game theory predictions with human individuals. The parameters that can be varied and the experiments that can be performed are essentially limitless, proving the value of this program as a tool. Although the experiments conducted here failed to answer the ultimate question of how cooperation evolves, they did show the possibilities for further research. Overall, gTheory provides an excellent interface for deeply exploring game theoretic predictions involving human populations.

References

Axelrod, R., Hammond, R.A., and Grafen, A. 2004. Altruism via kin-selection strategies that rely on arbitrary tags with which they coevolve. *Evolution.* 58(8): 1833-1838.

Campbell, D.T. 1983. Comments on the Sociobiology of Ethics and Moralizing. *Behavioral Science.* 24: 37-45.

Hauber, M.E., and Sherman, P.W. 2000. The armpit effect in hamster kin recognition. *Trends Ecol. Evol.* 15(9): 349-350.

Getty, T. 1987. Dear enemies and the prisoners-dilemma – why should territorial neighbors form defensive coalitions? *Am. Zool.* 27: 327-336.

Mitteldorf, J., and Wilson, D.S. 2000. Population viscosity and the evolution of altruism. *J. Theor. Biol.* 204: 481-496.

Riolo, R.L., Cohen, M.D., and Axelrod, R. 2001. Evolution of cooperation without reciprocity. *Nature.* 414: 441-444.

Roberts, G., and Sherratt, T.N. 2002. Does similarity breed cooperation? *Nature.* 418: 499-500.

Trivers, R.L. 1971. The evolution of reciprocal altruism. *Quarterly Review of Biology*. 46: 35-37.