Effects of misconception on reciprocative agents

Sandip Sen & Anish Biswas
Department of Mathematical & Computer Sciences,
The University of Tulsa
Phone: (918) 631-2985, FAX: (918) 631-3077
e-mail: sandip@kolkata.mcs.utulsa.edu,abiswas@euler.mcs.utulsa.edu

Abstract

In open systems, agents act to serve their self-interests instead of working towards group goals. We investigate the choice of interaction strategies and environmental characteristics that will make the best self-interested actions to be cooperative in nature. In our previous work, we have presented a probabilistic reciprocity mechanism that produces stable, cooperative behavior among a group of self-interested agents. The resultant system was found to exhibit close to optimal throughput with a fair distribution of the workload among the participating agents. In this paper, we test the robustness of this scheme by changing some of the assumptions used before. In particular, we evaluate the performance of reciprocative agents when an agent receiving help from another agent underestimates the cost incurred by the helping agent. This kind of misconceptions can derail cooperative processes leading to disintegration of stable groups. Results from experiments with several kinds of misconceptions, however, testify to the robustness of the reciprocative strategy.

1 Introduction

Researchers involved in the design of intelligent agents that will interact with other agents in an open, distributed system are faced with the challenge of modeling other agents and their behavior [5]. If one can can assume that all agents will be cooperative in nature, efficient mechanisms can be developed to take advantage of mutual cooperation. But, in an open system, assumptions about cooperative agents or system-wide common goals are hard to justify. More often, we will find different agents have different goals and motivations and no real inclination to help another agent achieve its objectives.

We assume agents to be self-motivated in their interactions with other agents, and that interacting agents are uniquely identifiable. An agent may help others in performing tasks. We develop a criteria for an agent to decide to help or not to help another agent when the latter requests for help. The decision criteria should be such that it allows an agent to perform effectively in the long run. This means that to be effective, an agent must be able to adapt its behavior depending on the behavior of other agents in the environment.

In our previous work, we have presented a simple decision mechanism using the principle of *reciprocity*, which suggests that agents help others who have helped them in the past or can help them in the future. By using a multiagent domain where agents can exchange their tasks, we have shown that agents can use the principle of reciprocity to effectively adapt to the environment, optimizing local performance [12]. We developed a probabilistic reciprocity mechanism that was found to be stable against invasion by selfish agents (agents who received help but never helped others).

There were two underlying assumptions in the paper mentioned above which cast some doubts about the possibility of practical application of our proposed strategy:

Accurate cost information: When an agent A incurs a cost c to help agent B and saves the latter an amount s, we say agents have accurate cost information if the values c and s are known to both of them. If agent B has only an estimate, $\tilde{c}$ of c, we say B is underestimating or over-estimating the cost incurred by A if $\tilde{c} < c$ or $\tilde{c} > c$ respectively. Underestimation is a common group malady and can jeopardize agent relationships or destabilize working partnerships. We investigate the effects of over and under-estimation of help costs by reciprocative agents. Such evaluation involves both performance of homogeneous groups, as well as testing the stability of the reciprocative strategy in the presence of selfish agents.

Beneficial cooperation: This assumption states that an agent asks help from another agent to achieve a task only if the cost incurred by the helping agent will be less than the savings obtained by the asking agent. This means that for each instance of cooperation (one agent helping another) there is a net savings in the cost incurred in the system. In real life, if I believe someone else owe me a favor, I can ask that person to complete one of my tasks, even if the cost incurred by that person to complete the task may be more than what it would have cost me to complete that task. This would be an example of what we call a non-beneficial cooperation (from the system point of view, as the recipient of help always benefits). In this paper, we study the performance of our reciprocity mechanism when both beneficial and non-beneficial cooperations are allowed. The latter is expected to erase some of the benefits of beneficial cooperation, but which of the two forms of cooperation dominate is an open question in general.

2 Reciprocity as an adaptive mechanism

The evolution of cooperative behavior among a group of self-interested agents have received considerable attention among researchers in the social sciences and economics community. Researchers in the social sciences have focused on the nature of altruism and the cause for its evolution and sustenance in groups of animals [7, 11, 13]. Mathematical biologist and economists have tried to explain the rationality of altruistic behavior in groups of self-interested agents by proposing various fitness models that analyze the success of altruistic individuals and more importantly the evolution of genetic traits supporting altruistic behavior [4, 8, 9]. Our goal in this paper is not to model altruistic behavior in animals; so we do not address the issues raised in the social science literature on this topic. Our purpose is to propose mechanisms by which cooperation can be encouraged and established in groups of self-interested agents. Other researchers in multiagent systems are beginning to evaluate the effects of mutual help for survival [3]. Most of the work by mathematical biologists or economists on the evolution of altruistic behavior deals with the idealized problem called Prisoner's dilemma [10] or some other repetitive, symmetrical, and identical 'games'. Some objections have already been raised to using such sanitized, abstract games for understanding the evolution of complex phenomena like reciprocal altruism [2].

In a seminal piece of work Robert Axelrod has shown how stable cooperative behavior can arise in self-interested agents when they adopt a reciprocative attitude towards each other [1]. Axelrod shows that a simple, deterministic reciprocal scheme of cooperating with another agent who has cooperated in the previous interaction, is quite robust and efficient in maximizing local utility. Two properties of this *tit-for-tat* strategy deserve special mention: (a) if all agents use this strategy, system performance is optimal, (b) it is stable against invasion by agents using other strategies (i.e., if an agent using another strategy is introduced into a group of tit-for-tat agents, the former cannot obtain greater utility than that obtained by tit-for-tat agents). Though Axelrod's work is interesting and convincing, we have argued [12] that the assumptions used in his work makes the results inapplicable in a number of practical domains. Some of the problematic assumptions are:

Initial decision:
Axelrod assumes that such agents start of cooperating; if agents start off by not cooperating, then the tit-for-tat strategy will never produce cooperative action.

Symmetrical interactions: Axelrod assumes that every interaction is perfectly symmetrical. This implies that if two agents cooperate in any interaction, both incur the same cost and benefit. In real-life interactions, more often than not in any one interaction one agent incurs the cost and the other incurs the benefit.

Repetition of identical scenarios: Axelrod assumes a game-playing framework with continually recurring situations. In real-life, however, more often than not, either the parties involved or the environmental conditions, are different.

Lack of a measure of work: Real life scenarios present differing circumstances which need to be compared based on some common metric. For example, consider a scenario where time is the cost metric of cooperation. Suppose that A helped B by retrieving some information over the internet; this act of cooperation cost A 5 minutes. Now, assume that A asks B to execute some programs that will cost B 2 hours. Should B honor such a request? The simple tit-for-tat mechanism will suggest that B cooperates, but that may not be the best choice. The point is that there is no mechanism for comparing past favors and future expectations in the tit-for-tat strategy. It was not designed for scenarios in which individual cooperation acts benefits one party while the other incurs a cost.

Our proposal, then, is to use a probabilistic, rather than a deterministic reciprocity scheme, with the following properties: (1) allow agents to initiate cooperative relationships (this implies that it should be able to handle asymmetrical interactions), (2) use a mechanism to compare cooperation costs, (3) allow agents to be inclined to help someone with whom it has a favorable balance of help (have received more help than have offered help), (4) be able to flexibly adjust inclination to cooperate based on current work-load (e.g., more inclined to cooperate when less busy, etc.).

3 Probabilistic reciprocity

We assume a multiagent system with N agents. Each agent is assigned to carry out T tasks. The jth task assigned to the ith agent is t_{ij}, and if agent k carried out this task independently of other tasks, the cost incurred is C^k_{ij}. However, if agent k carried out this task together with its own task t_{kl}, the cost incurred for task t_{ij} is C^{kl}_{ij}. Also, the cost incurred by agent k to carry out its own task t_{kl} while carrying out task t_{ij} for agent i is C^{kij}_{kl}. In this paper, we allow an agent to carry out a task for another agent only in conjunction with another of its own tasks.

If an agent, k, can carry out the task of another agent, i, with a lower cost than the cost incurred by the agent who has been assigned that task ($C^i_{ij} > C^{kl}_{ij}$), the first agent can cooperate with the second agent by carrying out this task. If agent k decides to help agent i, then it incurs an extra cost of C^{kl}_{ij} but agent i saves a cost of C^i_{ij}. The obvious question is why should one agent incur any extra cost for another agent. If we consider only one such decision, cooperation makes little sense. If, however, we look at a collection of such decisions, then reciprocal cooperation makes perfect sense.

We now propose a probabilistic decision mechanism that satisfies the set of criteria for choosing when to honor a request for help that we described at the end of the previous section. We will define S_{ik} and W_{ik} as respectively the savings obtained from and extra cost incurred by agent i from agent k over all of their previous exchanges. Also, let $B_{ik} = S_{ik} - W_{ik}$ be the balance of these exchanges ($B_{ik} = -B_{ki}$). We now present the probability that agent k will carry out task t_{ij} for agent i while it is carrying out its task t_{kl}. This probability is calculated as:

$$Pr(i,k,j,l) = \frac{1}{1 + \exp^{\frac{C^{kl}_{ij} - \beta * C^k_{avg} - B_{ki}}{\tau}}}, \qquad (1)$$

where C^k_{avg} is the average cost of tasks performed by agent k, and β and τ are constants. This gives a sigmoidal probability distribution in which the probability of helping increases as the balance increase and is more for less costly tasks. We include the C_{avg} term because while calculating the probability of helping, relative cost should be more important than absolute cost.

We present a sample probability distribution in Figure 1. The constant β can be used to move the probability curve

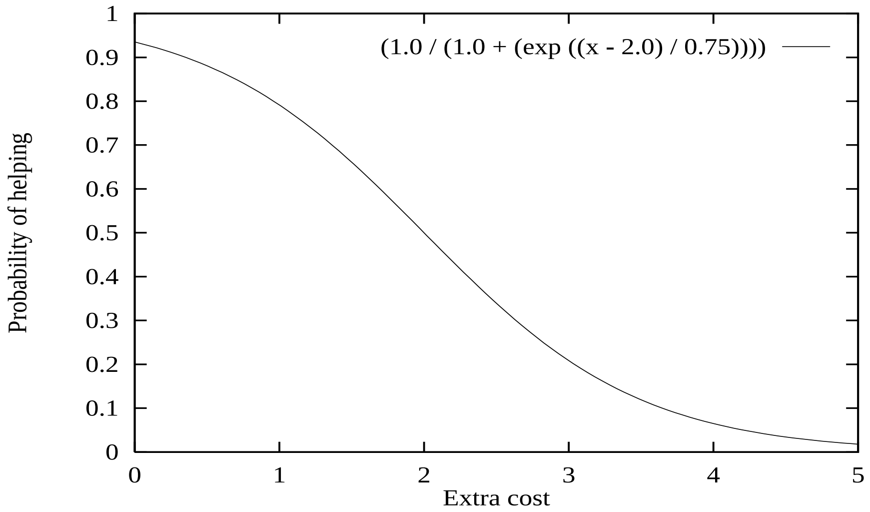

Figure 1: Probability distribution for accepting request for cooperation.

left (more inclined to cooperate) or right (less inclined to cooperate). At the onset of the experiments B_{ki} is 0 for all i and k. At this point there is a 0.5 probability that an agent will help another agent by incurring an extra cost of $\beta * C_{avg}^{k}$. The constant τ can be used to control the steepness of the curve. For a very steep curve approximating a step function, an agent will almost always accept cooperation requests with extra cost less than $\beta * C_{avg}^{k}$, but will rarely accept cooperation requests with an extra cost greater than that value. Similar analyses of the effects of β and τ can be made for any cooperation decision after agents have experienced a number of exchanges. In essence, β and τ can be used to choose a cooperation level [6] for the agents at the onset of the experiments. The level of cooperation or the inclination to help another agent dynamically changes with problem solving experience.

4 Experimental results

In the simple package delivery problem that we have used, we assume N agents, each of which is assigned to deliver T packets. All the packets are located in a centralized depot. The packet destinations are located on one of F different radial fins, and at a distance between 1 and D from the depot. Agents can only move towards or away from the depot following one of the fins; they cannot move directly between fins. On arriving at the depot, an agent is assigned the next packet it is to deliver. At this point, it checks if any other agents are currently located in the depot. If so, it can ask those agents to deliver this packet.

The cost of an agent to deliver one of its packets individually is double the distance of the delivery point from the depot. If it carries another package to help another agent, it incurs one unit of extra cost per unit distance traveled when it is carrying its own packet and this extra packet. In addition, if it is overshooting its own destination to help the other agent, an additional cost measured as double the distance between the destination of its packet and the destination of the other agent's packet is incurred.

We vary the number of agents and the number of packets to be delivered by each agent to evaluate the effects of different environmental conditions. The other parameters are as follows: $F = 4$, $D = 3$, $\tau = 0.75$, and $\beta = 0.5$. Each of our experiments is run on 10 different randomly generated data sets, where a data set consist of an ordered assignment of package deliveries to agents. All the agents are assigned the same number of deliveries. The evaluation metric is the average cost incurred by the agents to complete all the deliveries.

We experimented with the following types of agents:

Philanthropic agents: Agents who will always accept a request for cooperation. Philanthropic agents will produce the best global system performance. To ensure this, we allow only beneficial cooperation between these agents.

Selfish agents: Agents who request for cooperation but never accept a cooperation request.

Reciprocative agents(R): Agents that follow our prescribed strategy of using the balance of cost and savings to stochastically decide whether to accept a given request for cooperation.

Underestimating R agents(UR): R agents who underestimate the help received.

Overestimating R agents(OR): R agents who overestimate the help received.

OUR agents: Agents who at times underestimate and at other times overestimate the help received from others.

Individual agents: Agents who neither receive nor give help to others.

We expect the individual and the philanthropic agents to provide the two extremes of system performance. The individual agents should travel on the average the longest distance to complete their deliveries (because no one is helping them), whereas the philanthropic agents should travel the least. We expect the basic reciprocative agents (referred to as R agents henceforth) behaviors to lie in between. The frequency of occurrence of cooperation possibilities should determine which of the two ends of the spectrum is occupied by the reciprocative agents. We expect that the underestimating reciprocative agents (UR) are going to suffer compared to R agents as they are going to return smaller favors

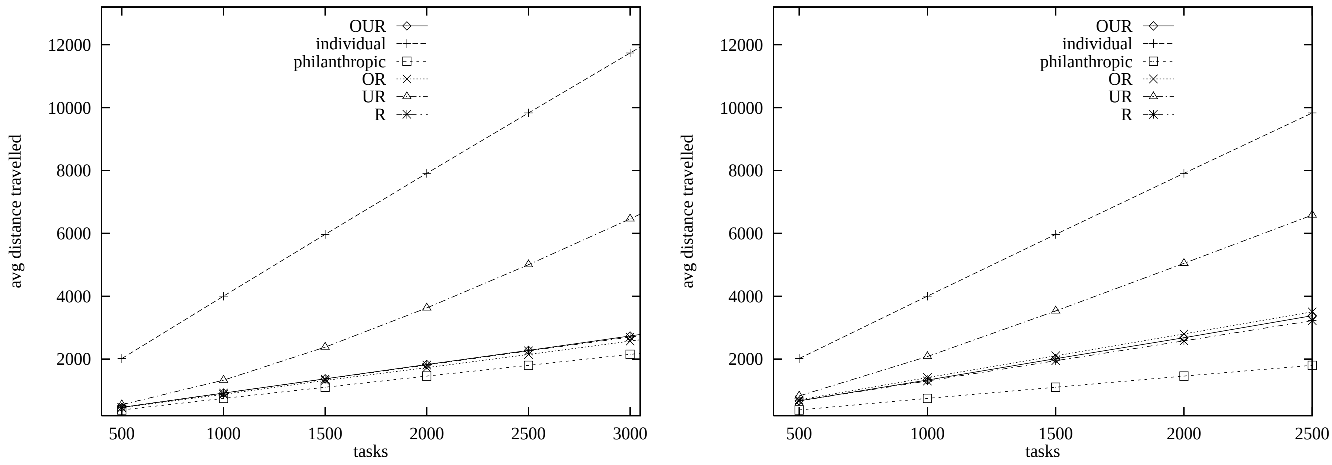

Figure 2: Average distance traversed by each agent to complete all deliveries. Left figure shows the scenario with the restriction of beneficial cooperation and the right one is without the restriction, where R=reciprocative,UR=underestimating reciprocative,OR=overestimating reciprocative and OUR =overestimating and underestimating reciprocative agents

than those received. The overestimating agents (OR) return more favors than received and should perform between the philanthropic and the R agents. We expect the behavior of both over and under-estimating reciprocative (OUR) agents to be close to the reciprocative behavior. Underestimation and overestimation is simulated by subtracting from and adding to, respectively, the actual cost incurred by the helper agent a Gaussian random number generated with mean μ and standard deviation σ.

A critical issue to be investigated empirically whether selfish agents can profit at the expense of reciprocative agents. It would seem that reciprocative agents should perform better because with sufficient interactions they become philanthropic towards each other, a possibility denied to the selfish agents.

For the first set of experiments we chose the number of agents, N, as 100 and varied the number of deliveries per agent from 500 to 3000 in increments of 500. Different experiments were performed on homogeneous sets of individual, R, UR, OR, OUR and philanthropic agents. Results from these set of experiments are presented in Figure 2. As expected, the performance of the individual agents were the worst, and the philanthropic agents were the best. The interesting observation is that the performance of the reciprocative agent is almost identical to that of philanthropic agents. That is, when a reciprocative agent is placed in a group of other reciprocative agents it adapts over time to behave like a philanthropic agent, and this adaptation benefits everybody. This is a significant result because we have been able to show that under proper environmental conditions (frequent interactions with possibilities of cooperation), self-motivated behavior based on reciprocity can produce mutually cooperative behavior that leads to near-optimal system performance. The UR agents perform significantly worse compared to R agents, but are still much better off compared to individual agents. The performance of OR and OUR agents are statistically identical to that of R agents. This is because the performance of R agents are already approximately optimal.

In the second set of experiments we allow non-beneficial cooperation which implies that here agents can ask agents for help irrespective of which fin they plan to travel to. In such cases, cooperation will not always produce a net saving to the system, e.g., agent A delivering a packet for agent B to fin 1 when its own packet has to be delivered to fin 2. When non-beneficial cooperation possibilities are allowed in conjunction with misconception, problems are exacerbated. From Figure 2 we notice that the UR agents are the most affected ones! We also observe that the R, OR, OUR agents perform worse when non-beneficial cooperation is allowed compared to the case where only beneficial cooperation is utilized.

The next set of experiments were designed to evaluate the stability of reciprocative strategies by introducing selfish agents in an otherwise homogeneous group of reciprocative agents. We expected that selfish agents should be able to obtain some help from reciprocative agents. Hence they would perform better than individual agents, but may not be able to match the performance of reciprocative agents. We fixed the number of agents at 100 and the number of deliveries at 1000. We varied the percentage of selfish agents in the group. Results are presented in Figure 3. Our intuitions regarding the relative performance of the agents are corroborated by the figure. The average performance of the group, lies in between the performance of the selfish and reciprocative agents, and moves closer to the performance of the selfish agent as the percentage of the latter is increased. Since reciprocative agents incur extra cost for selfish agents without being reciprocated, their performance is noticeably worse than the baseline performance of the homogeneous philanthropic agents. Moreover we found that the use of reciprocity allows the reciprocative agents to adopt their behavior such that after sufficient number of interactions they learn to reject requests for help from the selfish agents, while at the same time acting "philanthropically" towards other reciprocative agents.

When we evaluate the UR, OR, and OUR agents for stability by including selfish agents in the group, we find that these agents still perform much better than selfish agents (see Figure 3). This is a heartening result. It means that even though misconception will lower system performance compared to the case where agents have accurate cost information, they can still adapt to behave "philanthropically" with similar-minded agents and learn to shun selfish agents

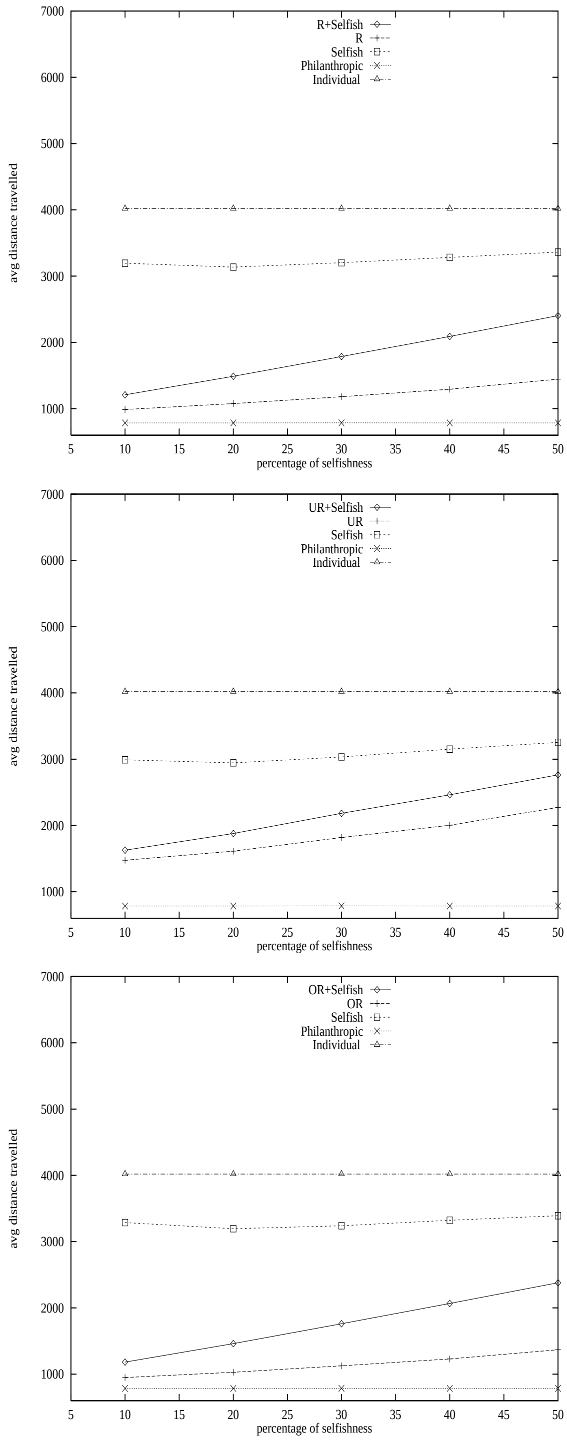

Figure 3: Average distance traversed by each agent to complete all deliveries as the percentage of selfish agent in a group of reciprocative agents is varied. The individual and the philanthropic agent results do not contain selfish agents and are presented for comparison. R=reciprocative, UR=underestimating reciprocative agents

i.e. our reciprocative strategy is stable even in the presence of misconception.

5 Conclusions

In our previous work, we have shown that self-motivated behavior can evolve cooperation among a group of autonomous agents leading to near-optimal local and system performance [12]. In this paper, we have relaxed restrictive assumptions of accurate cost estimation and beneficial cooperation between reciprocative agents. Performance is again close to optimal with both over and under-estimation of the amount of help received by an agent. System performance gracefully degrades, however, when agents underestimate the help they receive or when non-beneficial cooperation is allowed. A very important result is that the reciprocative strategy is still stable against invasion by selfish agents.

Currently, an agent receives help from the first person (from an ordered list) that agrees to help. We plan to study the performance of the mechanism when the agent considers all the offers for help and chooses to take help from the agent with which its got the most negative balance. We also plan to investigate more complex and realistic domains such as distributed monitoring, distributed information gathering, etc. to further evaluate the strengths and limitations of our proposed mechanism. We also want to perform a detailed analysis of how system performance is expected to degrade with the amount and frequency of misconceptions and we plan to present theoretical predictions and experimental verifications from this analysis.

Acknowledgments: This work has been supported, in part, by NSF grants IRI-9410180 and IRI-9702672.

References

[1] Robert Axelrod. *The Evolution of Cooperation*. Basic Books, 1984.

[2] R. Boyd. Is the repeated prisoner's dilemma a good model of reciprocal altruism? *Ethol. Sociobiology*, 9:211–222, 1988.

[3] Amedeo Cesta and Maria Miceli. Help under risky conditions: Robustness of the social attitude and system performance. In *Proceedings of the Second International Conference on Multiagent Systems*, pages 18–25, Menlo Park, CA, 1996. AAAI Press.

[4] L.A. Dugatkin, D.S. Wilson, L. Farrand III, and R.T. Wilkens. Altruism tit for tat and 'outlaw' genes. *Evolutionary Ecology*, 8:431–437, 1994.

[5] Les Gasser. Social conceptions of knowledge and action: DAI foundations and open systems semantics. *Artificial Intelligence*, 47(1-3):107–138, 1991.

[6] Claudia Goldman and Jeffrey S. Rosenschein. Emergent coordination through the use of cooperative state-changing rules. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, pages 408–413, 1994.

[7] D.L. Krebs. Altruism – an examination of the concept and a review of the literature. *Psychological Bulletin*, 73(4):258–302, 1970.

[8] S. Nee. Does hamilton's rule describe the evolution of reciprocal altruism? *Journal of Theoretical Biology*, 141:81–91, 1989.

[9] M.A. Nowak, R.M. May, and K. Sigmund. The arithmetics of mutual help. *Scientific American*, 272(6):76–81, June 1995.

[10] A. Rapoport. Prisoner's dilemma. In J. Eatwell, M. Milgate, and P. Newman, editors, *The New Palgrave: Game Theory*, pages 199–204. Macmillan, London, 1989.

[11] D. Schmitz. Reasons for altruism. *Social Philosophy & Policy*, 10(1):52–68, 1993.

[12] Sandip Sen. Reciprocity: a foundational principle for promoting cooperative behavior among self-interested agents. In *Proceedings of the Second International Conference on Multiagent Systems*, pages 315–321, Menlo Park, CA, 1996. AAAI Press.

[13] R. Trivers. The evolution of reciprocal altruism. *Quarterly Review of Biology*, 46:35–57, 1972.