The evolution and stability of cooperative traits

Sandip Sen
Math & CS Department
University of Tulsa
Tulsa, OK, USA
sandip@kolkata.mcs.utulsa.edu

Partha Sarathi Dutta
Math & CS Department
University of Tulsa
Tulsa, OK, USA
partha@ens.utulsa.edu

ABSTRACT

Recent works in multi-agent systems have identified agent behaviors that can develop and sustain mutually beneficial cooperative relationships with like-minded agents and can resist exploitation from selfish agents. Researchers have proposed the use of a probabilistic reciprocity scheme that uses summary information from past interactions to decide whether or not to honor a request for help from another agent. This behavior has been found to be close to optimal in homogeneous groups and outperform exploiters in mixed groups. A major shortcoming of these experiments, however, is that the composition of the group in term of agent behaviors is fixed. We believe that real-life rational agents, on the contrary, will change their behaviors based on observed performances of different behavioral traits with the goal of maximizing performance. In this paper, we present results from experiments on two distinct domains with population groups whose behavioral composition changes based on the performance of the agents. Based on the experimental results, we identify ecological niches for variants of exploitative selfish agents and robust reciprocative agents.

Keywords

Cooperation, evolution, reciprocity, adaptation and learning

1. INTRODUCTION

With the burgeoning of agent based electronic commerce, recommender systems, personal assistant agents, etc. it is becoming increasingly clear that agent systems must interact with a variety of information sources in an open, heterogeneous environment. One of the key factors for successful agent based systems (ABSs) of the future would be the capability to interact with other ABSs and humans in different social and role contexts and over extended periods of time. Research in societal aspects of agent behaviors has been relatively scarce. Whereas economic models can provide a basis for structuring agent interactions [17], other non-monetary approaches [1, 2, 3, 4, 5, 9, 14] may provide

effective solutions in certain situations. We assume that typical real-world environments abound in *cooperation possibilities*: situations where one agent can help another agent by sharing work such that the helping cost of the helper is less than the cost saving of the helped agent. As agent system designers we can also define rules of interaction to increase the likelihood of cooperation possibilities. We are interested in identifying agent behaviors that allow agents to take advantage of cooperation possibilities in their environments.

The current work is based largely on the work of Sen *et al* [14, 15]. They have presented behaviors that promote cooperation among homogeneous groups and can resist exploitation by malevolent agents in heterogeneous groups. Such behaviors can lead to both improved local performance for individual agents and effective global performance for the entire system.

A restrictive assumption in this line of work has been that agents have fixed behaviors. For example, they have assumed that agents with specified behaviors interact repeatedly over a sustained period of time and their effectiveness is calculated as function of the total cost incurred to complete all assigned tasks. The resultant performance reflects cost incurred for local tasks, cost incurred to help other agents with their tasks, and savings obtained from others when help was received.

A more realistic scenario would be to give an agent the freedom of choosing from one of several of these behaviors and to change its behavior as and when it deems appropriate. An agent may be prompted to adopt a behavior if agents using that behavior is seen to be performing better than others. Such a behavior adoption method leads to an evolutionary process with a dynamically changing composition of agent group behaviors. It is not clear *a priori* if a behavior that produce highest returns would emerge as the dominant behavior in a group where agents change behaviors regularly based on limited-term performance.

Our goal in this paper is to identify the dominant strategies under different environmental conditions including initial population composition and the frequencies with which agents change their behaviors based on observed performance. We experiment with a population of agents with the initial population containing representatives of different behaviors in specified proportions. Each of these agents are then assigned some tasks. The cost of executing a task can be reduced or eliminated if help is obtained from another agent. After all agents have finished processing their assigned tasks, their relative performances are tallied. This

comprises one evaluation period, or *generation*, of the behaviors adopted. The behaviors adopted by the agents in the next evaluation period is determined by a performance-proportionate scheme where the probability with which an agent adopts a strategy increases with the average performance of agents employing that strategy in the most recent evaluation period. Thus, it is likely that more agents are produced with behaviors that generated above-average performance. This generational scheme is semantically equivalent to every agent periodically selecting its behavior based on the current relative performance of the set of available behaviors. This generational approach is akin to work on identifying "evolutionary stable strategy" [6]. The goal of that body of work is to use selection pressure to identify behaviors that can perform dominantly under a variety of group compositions.

The purpose of our current study is two-fold: (a) to evaluate whether conclusions drawn by Sen *et al.* with fixed behavior agents holds up in the more realistic scenarios where agents change their behavior to more successful ones based on experience, (b) to identify the environments for which different behaviors would be dominant in performance. The latter goal includes the sub-goals of understanding the effects of (i) initial group compositions and (ii) behavior evaluation periods or frequency of behavior reconsideration, on the emergence of dominant behavior. Whereas Sen *et al's* study tells us what behaviors are dominant in what group compositions for a fixed number of interactions, it provides little understanding of how the combination of different group compositions and evaluation periods affect the choice of the most effective behavior. Our work identifies such "ecological niches" for different agent behaviors. In particular, it gives us a good idea of the minimum number of tasks needed before reciprocating can be a better choice than exploiting. To demonstrate the feasibility of this approach we experiment with two different types of domains: (a) a physical domain involving delivery of packages, and (b) an information domain requiring processing of information tasks.

In the following, we first briefly discuss the motivation for studying reciprocal behavior in self-interested agents, identify different variants of selfish and cooperative agent behaviors, present our experimental framework to evaluate our conjectures, narrate our experimental results highlighting the non-intuitive results and derive prescriptions for the adoption of different behaviors based on the expected group composition and environmental dynamics, and present summary observations and future research directions.

2. ADAPTATION VIA RECIPROCITY

The social sciences and the economics community have paid considerable attention to the issue of evolution of cooperative behavior among a group of self-interested agents. The social sciences researchers analyze the nature of altruism and the cause for its evolution and sustenance in animal groups [16]. Mathematical biologists and economists evaluate the rationality of altruistic behavior in groups of self-interested agents by proposing fitness models that analyze the success of altruistic individuals and the evolution of altruistic genetic traits [8, 12]. We do not intend to model altruistic behavior in animals or humans and hence do not address the issues raised in the social science or experimental economics literature on this topic [11].

A significant body of work by mathematical biologists or economists on the evolution of altruistic behavior deals with the idealized problem called Prisoner's dilemma [13] or some other repetitive, symmetrical, and identical 'games'. To consider a well-known study in this area, Axelrod demonstrates that a simple, deterministic reciprocal scheme or the *tit-for-tat* strategy is quite robust and efficient in maximizing local utility [2]. Sen criticizes the simple reciprocative strategy is not the most appropriate strategy to use in most real-life situations because most of the underlying assumptions that motivate its use are violated in these situations [14].

The evaluation framework used by Axelrod considers an evolving population composition by allowing propagation of more successful behaviors and elimination of unsuccessful ones. Sen's framework, however, contains a static group composition and relative performance is measured by varying the total number of interactions of this group of agents. Though interesting, we believe it represents an incomplete evaluation of probabilistic reciprocity based behaviors. In this paper, we evaluate the variants of exploitative and reciprocative behaviors suggested by Sen *et al.* [15] in a generational framework as used by Axelrod [2]. This allows us to see what behaviors emerge to be dominant or are evolutionarily stable.

Though Sen *et al.* work with heterogeneous groups, these always have only up to two behavior types. For example they show that reciprocative agents can resist exploitation from selfish agents and outperform them in mixed groups. This, however, does not help us understand if in the presence of naive or philanthropic agents (who always help when asked) the selfish would outperform the reciprocative agents. To understand the dynamics in this more realistic scenarios, we experiment with mixed groups of exploitative, reciprocative and philanthropic agents.

3. PROBABILISTIC RECIPROCITY

We now present our probabilistic reciprocity framework for deciding whether or not to help another agent. Each agent is assigned to carry out T tasks. The jth task assigned to the ith agent is t_{ij} will cost it C_{ij}. However, if agent k carried out this task together with its own task t_{kl}, the cost incurred for task t_{ij} by agent k is C_{ij}^{kl} (no cost is incurred by agent i). If $C_{ij} > C_{ij}^{kl}$, there exists a cooperation possibility as agent k can help agent i save C_{ij} by incurring a cost of only C_{ij}^{kl}.

We define S_{ik} and W_{ik} respectively as the cumulative savings obtained from and extra cost incurred by agent i from agent k over all of their previous exchanges. Also, $B_{ik} = S_{ik} - W_{ik}$ is the balance of these exchanges (note that, in general, $B_{ik} \neq -B_{ki}$).

Sen [14] proposes a probabilistic decision mechanism that satisfies a set of criteria for choosing when to honor a request for help that was described at the end of the previous section. The probability that agent k will carry out task t_{ij} for agent i while it is carrying out its task t_{kl} is given by:

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

where C_{avg}^{k} is the average cost of tasks performed by agent k, and β and τ are constants. This is a sigmoidal probability function (not a probability distribution) where the probability of helping increases as the balance increases and is more for less costly tasks.

Figure 1: Probability distribution for accepting request for cooperation.

A sample probability distribution is presented in Figure 1. β can be set to a low value to move the probability curve left (less inclined to cooperate) or to a high value to move the curve to the right (more inclined to cooperate). Initially, $B_{ki} = 0$ for all i and k. At this point the probability that an agent will help another agent by incurring an extra cost of $\beta * C_{avg}^{k}$ is 0.5. 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. Another way to consider their role is that β and τ can be used to choose a cooperation level [10] for the agents. These are the only two "parameters" in the equation 1 that can be fine tuned to adjust the level of cooperation. The other equation variables determine the actual dynamics of the agent behaviors. The level of cooperation or the inclination to help another agent is dynamically adapted based on past interactions with that agent. Note that the sigmoid is one of several function classes that can be used to represent a probabilistic reciprocity behavior.

4. SET OF AGENT BEHAVIORS

In this paper we experiemnt with the following basic selfish and reciprocative agent types [14]:

Philanthropic agents: Agents that always honor a cooperation request irrespective of past experience.

Selfish agents: Agents who ask for help but never return favors. Selfish agents can thrive on the benevolence of philanthropic agents.

Reciprocative agents: Agents that use the probabilistic reciprocity scheme described above.

The variants on these strategies that we have experimented with are as follows [15]:

Believing reciprocative agents: Agents who use corresponding balances reported by all agents, and not just their own balances, when deciding whether or not to help another agent. In place of using B_{ki} in Equation 1, a believing reciprocative agent k uses $\sum_{j \neq i} B_{ji}$ while calculating the probability of helping agent i. This variant was designed so that reciprocative agents can quickly identify and shun the exploitative selfish agents.

Earned-Trust based reciprocative agents: While evaluating a request for help, these agents consider balances of only those agents with whom they themselves have favorable balances. In place of using B_{ki} in Equation 1, a conservatively trusting reciprocative agent k uses $\sum_{j \neq i \wedge B_{kj} > 0} B_{ji}$ while calculating the probability of helping agent i. This behavior is an augmentation of the believing reciprocative agent and was required to counter false balance reporting by exploitative agents.

Individual lying selfish agents: These agents exploit the fact that believing or trusting reciprocative agents use balances provided by other agents. They reveal false impressions about other helpful agents to ruin their reputation. Whereas this behavior is perhaps difficult to justify for fixed group compositions as used by Sen *et al.* it is all the more reasonable when agents can change behaviors: one can appear to be better than the rest either by doing well itself or by ruining others. When such an agent, j, is asked for its balance with another agent i, it reveals B'_{ji} as

$$
\begin{aligned}
B'_{ji} &= C * (-B_{ji}), \text{ when } B_{ji} > 0 \\
&= B_{ji}, \text{ otherwise,}
\end{aligned}
$$

where C is a positive constant. Hence, the more an agent i helps it, the larger the negative balance an individual selfish agent will report about agent i to other agents.

Collaborative lying selfish agents: These agents not only attempt to tarnish the reputation of other helpful agents, but also collaboratively bolster the reputation of other selfish agents. When such an agent, j is asked for its balance with another agent i, it reveals B'_{ji} given by:

$$
\begin{aligned}
B'_{ji} &= C * (-B_{ji}), \text{ when } B_{ji} > 0 \\
&= \mathcal{P}, \text{ otherwise}
\end{aligned}
$$

where C is a positive constant as above and $\mathcal{P}$ is a large positive constant.

5. EXPERIMENTAL RESULTS

We performed experiments with two domains: a package delivery problem [14] and an information processing domain.

In the package delivery problem N agents are assigned to deliver T packages each. All the packages are picked up for delivery from a central depot. The package destinations are located on one of R 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 package it is to deliver. At this point, it can ask for help in delivering the package it is assigned only from another agent currently at the depot and going to deliver a package on the same fin.

The cost incurred by an agent to deliver one of its packages individually is double the distance of the delivery point

from the depot. If it carries a package to help another agent, it incurs one unit of extra cost per unit distance traveled when it is carrying this extra package up to its own destination. 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 package and the destination of the other agent's package is incurred.

The parameters for the experiments are as follows: $N = 100$, $R = 3$, $D = 3$, $\tau = 0.75$, and $\beta = 2$, $C = 1$, $\mathcal{P} = 10$. Each of our experiments are run on 10 different randomly generated data sets, where a data set consists of an ordered assignment of package deliveries to agents. All the agents are assigned the same number of deliveries with the same total delivery cost. The evaluation metric is the average cost incurred by the agents to deliver all the tasks that they were assigned and those taken from others while honoring help requests.

Agent behaviors are randomly initialized according to some preset distributions. At the end of each evaluation period, i.e., once the agents have delivered all their assigned tasks (perhaps with the help of other agents), their total delivery costs are tallied. Performance of an agent is inversely proportional to the cost it incurs. New agent behavior assignments are made as follows: for each agent i, two agents are selected proportional to their performance, i.e., the probability of selecting the jth individual with performance p_j is $\frac{p_j}{\sum_{k=1}^{N} p_k}$. Then, of these two selected agents, the behavior of the one with higher performance is adopted by agent i[1]. This leads to a propagation of successful behaviors or traits. As a result, if a behavior produces better performance in one evaluation period compared to other behaviors, we are likely to see more individuals adopting that behavior in the next evaluation period. We run this generational process until the population becomes homogeneous or a fixed number of evaluation periods (we have used a limit of 20) is reached[2]. The behavior in the homogeneous population or the majority behavior in the case when evaluation period limit is reached is declared the "winning" or preferred behavior for the corresponding starting conditions (initial behavior distribution and the given number of tasks).

We first describe experiments with two types of behaviors. In these experiments, the initial percentage of selfish agents is varied from 10% to 90%. For a given percentage of selfish agents we ran experiments with different number of tasks per evaluation period. For small number of tasks, the selfish behavior will dominate the population since the agents do not stay 'on the field' interacting for sufficient time and hence reciprocatives fail to identify the exploiters (selfish agents). We started these experiments with 10 tasks assigned to each agent per evaluation period and increased the

Figure 2: **Tasks required before reciprocative agents dominate the population in mixed groups of : Reciprocative and Selfish agents (top), Philanthropic, Reciprocative and Selfish agents (bot).**

number of tasks in increments of 10 until the reciprocative behavior became dominant. The results plotted in Figure 2 (top) shows that as the initial percentage of selfish behaviors is increased, it takes more tasks before the reciprocative strategy evolves and stabilizes to be the dominant strategy. This result has two implications. First it confirms that Sen's conclusion [14] about the reciprocative strategy being dominant do hold in the scenario where agent group composition evolves based on prior performance of different behaviors. Secondly, it gives us a more precise understanding of how many tasks (a measure of evaluation period lengths) are required before reciprocative behavior can dominate exploitative behavior. Another way to interpret the curves is the following: for a given combination of initial selfish percentage and a number of tasks to be accomplished, if the corresponding point lies below the curve then exploitative behavior will prove to be dominant; if the point lies above the curve, reciprocative behavior will be dominant. This interpretation identifies ecological niches where certain behaviors will come to dominate other behaviors.

In the same plot (Figure 2 (top)) we also show the corresponding curve for a mixed group of selfish and believing reciprocative agents. We see that the reputation based believing reciprocity scheme is quickly able to identify exploiters. Hence the corresponding behavior is found to be dominant even when most of the initial group are exploiters and only a few tasks are to delivered.

The next set of experiments are run with mixed groups of philanthropic, exploitative (selfish), and reciprocative agents (see Figure 2 (bot)). We varied the initial percentage of phil-

---

[1]Selection of the best candidate from a set of randomly selected candidates is known as *tournament selection* in the genetic algorithms literature [7]. In our case, the selection pressure is further increased because the candidate set is not chosen randomly but proportionate to the fitness of individuals in the population.

[2]Initially we adopted a purely proportionate selection scheme without the tournament selection component. In such cases, the equilibrium population distribution was mixed, i.e., the dominant population did not completely eliminate the lesser performing behaviors. We believe such mixed equilibrium distributions warrants further detailed investigations. While clearly interesting in its own rights, such an analysis is beyond the scope of the current paper.

Figure 3: Variation in percentage of philanthropic, selfish, and reciprocative agents over two runs.

Figure 4: Tasks required before reciprocative agents dominate the population in mixed groups of : Reciprocative (believing or earned-trust) and Selfish (individual lying) agents (top), Philanthropic, Reciprocative (believing or earned-trust) and Selfish (individual lying) agents (bot).

anthropic agents with the rest of the agents being equally divided between reciprocative and exploitative behaviors. The question was whether the exploitation of the philanthropic agents will allow the selfish to increase its percentage early and enough to dominate the population in the long run. We observed that the selfish exploited the philanthrop ruthlessly to cause its extermination in a few evaluation periods. Whether or not it got to dominate the population depended upon the percentage of the population that adopted the selfish strategy when the philanthrops became extinct. The extermination of the philanthrops created a run-off between the selfish and the reciprocative agents. When a significant percentage of the initial population are philanthrops, the selfish emerge dominant even for very high number of tasks. The plot for the believing reciprocative agents in Figure 2 (bot) shows that the believing reciprocative agents are able to dominate the population even for small number of tasks per evaluation period for all the different initial philanthrop percentages.

Though these results provide summary information over runs, we wanted to understand the dynamics of the population over a number of evaluation periods in a single run. In Figure 3 we plot the variation of group composition over evaluation periods. Both these plots correspond to initial philanthrop percentage of 30%. In the plot of Figure 3 (top), with 650 tasks per agent, the selfish behavior grows to dominate the population. The plot on Figure 3 (bot), with 700 tasks per agent, shows a more interesting phenomena: the selfish behavior initially dominates but finally loses out to the reciprocative behavior. This happens partly because the selfish could not command as high a percentage of the population as in the previous case by the time the philanthropic

agents die out. The reciprocative agents, then could recuperate and eventually dominate the population.

We now summarize observations from experiments with combinations of individual lying (IL) or collaborative lying (CL) selfish agents on one hand and believing (B) or earned-trust (ET) based agents on the other. First, we present comparison of IL-B and IL-ET groups. In Figure 4 (top), the plots for the two compositions are shown with varying initial selfish percentages when there are no philanthrops in the agent population. Figure 4 (bot) shows plots for the same agent compositions with philanthrops introduced in the population. In this later set of plots, for each initial philanthrop percentage, the rest of the population was equally divided among the selfish and reciprocatives. Contrary to our expectations, for smaller selfish and philanthropic percentages, respectively, the B agents are found to require less tasks to dominate the population against lying selfish agents when compared to ET agents. On further analysis of the data, we found that the earned-trust based agents require a nontrivial amount of tasks before they earn the trust from each other and can then mutually identify exploitative agents. The believing agent believes everyone and avoids this initial overhead. This naive believing works when few agents are lying but breaks down completely when a significant portion of the agents are exploitative liars. The lying selfish dominate against the believing reciprocative, but not against the earned-trust based exploitative, for arbitrary number of tasks when they start out in significant majority in the pop-

Figure 5: Minimum tasks before reciprocative agents dominate mixed groups of : Reciprocative (believing or earned-trust) and Selfish (collaborative lying) agents (top), Philanthropic, Reciprocative (believing or earned-trust) and Selfish (collaborative lying) agents (bot).

ulation. Some of the curves are incomplete in the figure because, for the corresponding initial selfish or philanthrop percentages, the reciprocative agents never dominated the population for the largest number of tasks we experimented with, viz. 3000.

Similar observations can be made when comparing CL-B and CL-ET groups (see Figure 5). Actually, the believing reciprocative agents could never dominate the collaborative lying agents when the initial population contained some philanthropic agents.

We ran farther set of experiments with an information processing domain where tasks of several types were assigned to agents. Different agents could achieve tasks of different types with different fixed costs. The agents could hand off a task to another if the latter could do it with less cost than the requesting agents. The principal difference with the previous domain is that in the package delivery domain an agent could not ask for help when it was "on the road". This meant that help-giving behavior influenced the likelihood of interacting with other agents, and indirectly, the likelihood of meeting and receiving help from another agent. The information processing domain decouples the help-giving behavior from the opportunity to ask for help. Results from this domain were qualitatively similar. In Figure 6 we plot the variation of tasks required for reciprocative agents to dominate under different initial percentage of philanthrop agents and in three different combinations of selfish and reciprocative types, viz. CL-B, CL-ET and IL-B. Incomplete
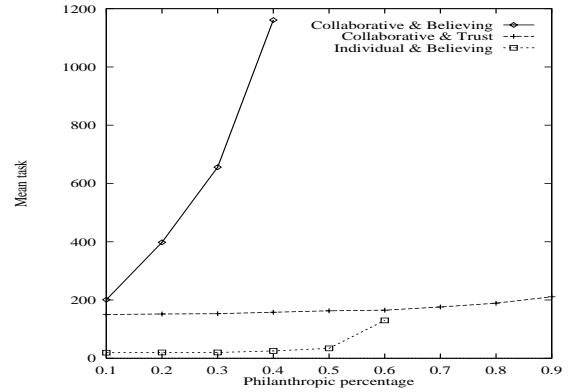

Figure 6: Tasks required for reciprocative agents to dominate in the information processing domain.

curves suggest dominance by selfish for higher philanthrop percentages.

6. CONCLUSIONS AND FUTURE WORK

In this paper we have studied the evolution of variants of selfish and reciprocative behaviors identified by Sen *et al.* [15] in a more realistic scenario where agents change behaviors at regular intervals based on observed performance. Varying initial group composition and the number of tasks to be executed per evaluation period, the dominant agent behaviors that we have identified for different environmental conditions are as follows: (a) collaborative lying for small number of tasks (likely in real world) or with high initial percentage of philanthropic agents (unlikely in real world), and (b) earned-trust based reciprocation for all other scenarios. The results reported in this paper corroborate Sen's prior work. The significance of this work is the demonstration of evolutionary stable strategies that emerge in populations of self-interested agents as the most beneficial strategies under different environmental settings. So, unlike Sen's previous work, the agent behaviors in this paper are adaptive, which is a more realistic and useful model. In addition, these results make available, for the first time, a clear guideline to prescribing agent help-giving behavior when estimates of initial distribution of strategies in the population and the period of interaction in the group, e.g., the number of task deliveries in the package delivery domain, are known.

One future goal is to analytically capture the dynamics of the evolution of agent population. Given a particular group composition and tasks per evaluation period, we plan to analyze and predict the behavioral composition of the group over time. We also plan to investigate alternate schemes for adopting new behaviors, e.g., neighborhood-based rather than global sampling which is more realistic. Another interesting study will be to verify whether groups or coalitions evolve among the agents that exhibit mutually complementary behavior.

7. REFERENCES

[1] A. A. Armstrong and E. H. Durfee. Mixing and memory: Emergent cooperation in an information

marketplace. In *Proceedings of the Third International Conference on Multiagent Systems*, pages 34–41, Los Alamitos, CA, 1998. IEEE Computer Society.

[2] R. Axelrod. *The Evolution of Cooperation*. Basic Books, 1984.

[3] R. Azoulay-Schwartz and S. Kraus. Stable strategies for sharing information among agents. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*, pages 1128–1134, 2001.

[4] C. Castelfranchi, R. Conte, and M. Paolucci. Normative reputation and the costs of compliance. *Journal of Artificial Societies and Social Simulation*, 1(3), 1998.

[5] A. Cesta and M. 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.

[6] T. Day and P. D. Taylor. The evolution of temporal patterns of selfishness, altruism, and group cohesion. *The American Naturalist*, 152(1):102–113, July 1998.

[7] K. Deb and D. Goldberg. A comparative analysis of selection schemes used in genetic algorithms. In G. J. Rawlins, editor, *Foundations of Genetic Algorithms*, pages 69–93, San Mateo, CA, 1991. Morgan Kaufman.

[8] L. Dugatkin, D. Wilson, L. F. III, and R. Wilkens. Altruism tit for tat and 'outlaw' genes. *Evolutionary Ecology*, 8:431–437, 1994.

[9] A. Glass and B. Grosz. Socially conscious decision-making. In *Proceedings of the Fourth International Conference on Autonomous Agents*, pages 217–224, New York, NY, 2000. ACM Press.

[10] C. Goldman and J. 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.

[11] E. Hoffman, K. A. McCabe, and V. L. Smith. Behavioral foundations of reciprocity: experimental economics and evolutionary psychology. *Economic Inquiry*, 36(3):335–352, July 1998.

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

[13] 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.

[14] S. 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.

[15] S. Sen, A. Biswas, and S. Debnath. Believing others: Pros and cons. In *Proceedings of the Fourth International Conference on Multiagent Systems*, pages 279–285, Los Alamitos, CA, 2000. IEEE Computer Society.

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

[17] M. P. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1(1):1–22, 1993.