Believing Others: Pros and Cons

Sandip Sen
Department of Mathematical & Computer Sciences
University of Tulsa
E-mail: sandip@kolkata.mcs.utulsa.edu

Abstract. In open environments there is no central control over agent behaviors. On the contrary, agents in such systems can be assumed to be primarily driven by self interests. Under the assumption that agents remain in the system for significant time periods, or that the agent composition changes only slowly, we have previously presented a prescriptive strategy for promoting and sustaining cooperation among self-interested agents. The adaptive, probabilistic policy we have prescribed promotes reciprocative cooperation that improves both individual and group performance in the long run. In the short run, however, selfish agents could still exploit reciprocative agents. In this paper, we evaluate the hypothesis that the exploitative tendencies of selfish agents can be effectively curbed if reciprocative agents share their "opinions" of other agents. Since the true nature of agents is not known a priori and is learned from experience, believing others can also pose its own hazards. We provide a learned trust-based evaluation function that is shown to resist both individual and concerted deception on the part of selfish agents in a package delivery domain.

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 [7, 9, 10, 19, 20, 34]. 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 role contexts and over extended periods of time. The ABSs of the future will be situated in a social context, playing a variety of roles in different relationships and problem solving situations. Borrowing on the social cliche leveled at humans, we would like to conjecture the following about the agents of the future: *Agents must be social entities.* In particular, we believe 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. Social agents can benefit from such cooperation possibilities by identifying and sustaining mutually beneficial relationships.

Whereas economic models can provide a basis for structuring agent interactions [24, 35], research in multiagent systems involving non-monetary, social reasoning procedures as behavioral strategies for self-interested agents has been relatively scarce. We believe that such societal approaches inspired by non-monetary mechanisms [1, 2, 3, 15, 27] may provide more effective social relationships in certain situations[1]. For example, agents can take advantage of cooperation possibilities by trading helps, where the cost incurred for helping is the time spent in helping the other agent. One strong argument for incurring "time costs" for help rather than "monetary costs" is that time cannot be stored like money. Given a choice, it is preferable to trade non-storable resources than storable resources and individual agents can benefit in the long run by using their unoccupied time to develop mutually beneficial relationships.

But cooperative relationships not only benefit individual agents, but can also enhance the condition of the entire society or the environment. Whereas as individual agent designers, we want to develop strategies which makes our agents profitable, as designers of entire agent systems or infrastructures, we want to maximize the performance of the entire system. For example, as designers of agent systems or infrastructures, we want the entire system to operate smoothly with little or no congestion, high throughput, balanced loads on resources, etc. Can these two different viewpoints be reconciled? Put another way, are there possible worlds where individually rational action leads both to maximizing local utility and improving system-level performance? These two goals cannot be reconciled under all situations. Under a reasonably practical set of assumptions, which apply to a significant set of realistic domains, however, we believe that the desired synthesis can be achieved. For example, as agent designers, we can design interaction protocols, feedback mechanisms, etc. to create abundant cooperation possibilities, which can then be utilized by well-designed agents. Such a scenario can produce high individual as well as system performance. But even though system designers provide cooperation possibilities, agent designers must endow their agents with the capability of identifying and benefiting from such opportunities in the environment.

In this paper, we assume that real-life environments do provide sufficient cooperation possibilities, and focus on the need to design social agents that can form cooperative relationships with other agents to benefit from such cooperation possibilities. But unsuspecting, naive agents, who cooperate with any agent in the environment, can be exploited by malevolent agents who receive but do not return help. We have been interested in agent strategies for interactions with other agents that can promote cooperation in groups while resisting exploitation [5, 27, 28, 29]. Our approach is different from other researchers who have designed effective social laws that can be imposed on agents [8, 30]. In particular, we have studied environments where agents can mutually benefit from sustained interactions. In such environments, appropriately designed agent strategies can lead to both improved local performance for individual agents and effective global behavior for the entire system. These are the desirable features for open systems where self-interested agents are required to share resources.

More specifically, we have developed and analyzed probabilistic reciprocity schemes as strategies to be used by self-interested agents to decide on whether or not to help other agents [27]. The goal of this work has been to identify procedures and environments under which self-interested agents may find it beneficial to help others. By helping we imply incurring some local cost to benefit another agent. We claim that if the group composition changes only slowly, and there is sustained interaction between the agents, probabilistic reciprocity based strategies can be rational, i.e., maximize individual utilities. Probabilistic reciprocity strategies can be considerably more effective than simple deterministic reciprocity schemes like tit-for-tat [2, 13] and avoid major problems associated with the latter [27].

In our experimental package delivery domain, an agent helps another agent by carrying out a task, i.e., by delivering a package, on behalf of the helped agent. But none of the mechanisms presented here are limited to that particular kind of cooperation. For example, the reciprocity approach presented can also be used in domains where cooperation implies the helping agent working together with the helped agent to reduce the latter's workload [26]. Our experiments under a variety of environmental conditions, group composition, work estimate difference, etc. have shown that under prolonged interaction the probabilistic reciprocity strategy produces close to optimal individual and group performance [5, 27, 28, 29]. Additionally, this strategy is stable against selfish intruders, i.e., in the long run, selfish agents perform worse than reciprocative agents in a mixed group.

We now turn to the focus of the current paper. Even though probabilistic reciprocative agents outperform selfish agents in mixed groups, they still waste some efforts in helping out selfish agents. This is because the reciprocative agents have a bias to initiate help to promote cooperative relationships in the future. A selfish agent can then benefit from this initial cooperative advances from each of the reciprocative agents in a mixed group. This is aided by the fact that reciprocative agents do not share their experiences or impressions of the other agents. In other words, there is no "words of mouth" transmission of the reputation or reliability of the agents in the agent group. As a result, the reciprocative strategy was dominant over exploitative strategies only if the agent group was stable, i.e., the agents interacted with a relatively long period of time. Our current work is driven by the need to augment the reciprocative strategy so that it becomes dominant even if the agents interacted with each other for only a limited period of time. We realize that in the extreme cases, where an agent typically interacts with others in the environment for one or a very few number times, the selfish strategy will be the logical choice, as there is not enough time to identify and develop mutually beneficial relationships. In the current work, however, we show that we can significantly reduce the number of interactions an agent should have with others before the reciprocative strategy dominates over exploitative ones.

A hypothesis that follows from the above considerations is the following: *Sharing of experiences about other agents among reciprocative agents will limit the exploitative gains of selfish agents.* Operationalizing this hypothesis, however, requires a closer inspection of the issues at hand. Since it is not clear a priori who is a selfish agent and who is a reciprocative agent (otherwise this whole exercise is moot because accurate identification immediately gives the right strategy to adopt while interacting with others), at the outset it is not possible to limit sharing of experiences only between selfish individuals. When an agent Z decides to use information supplied by an agent X to decide whether or not to help agent Y, then believing X can be advantageous or disadvantageous to Z based on the true nature of X. If X is selfish, it might find it useful to taint Y's reputation, and that of other agents, so that Z will consider X to be a relatively trustworthy agent. As such, we need to augment the reciprocative agents' strategy to believe only the agents who are trustworthy. In this paper, we evaluate the effectiveness of these strategies in mixed groups.

The principal contribution of this paper, therefore, is the development of robust reciprocative strategies, using which self-interested agents can form mutually beneficial relationships with other agents in the environment. These agents are also resistant to exploitation by malevolent agents. Given an environment with sufficient cooperation possibilities, such reciprocative agents can produce desirable individual as well as system level performances.

2. Related Work

The evolution of cooperative behavior among a group of self-interested agents has 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 [21, 25, 33]. 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 [18].

Most of the work by mathematical biologists or economists on the evolution of altruistic behavior deals with the idealized problem called Prisoner's dilemma [23] 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 [6]. In the following we analyze in some detail one of the often-cited work that share the typical assumptions made by economists and mathematical biologists, and then present our own set of suggestions for relaxing the restrictive assumptions made in that work.

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 [2]. The basic assumptions in this work include the following: agents are interested in maximizing individual utilities and are not predisposed to help each other; agents in a group repeatedly interact over an extended period of time; all interactions are identical (they are playing the same "game" again and again, where each game can be represented by payoff matrices which specify the payoff to be received by each agent given their simultaneous choice of actions); agents can individually identify other agents and maintain a history of interactions with other agents; individual agents do not change their behavioral strategy over time; composition of agent groups change infrequently and the changes are minimal (only a few agent leaves and joins a group at a time). Axelrod shows that a simple, deterministic reciprocal scheme of cooperating with another agent who has cooperated in the previous interaction (this strategy, for obvious reasons, is referred to as the *tit-for-tat* strategy), is quite robust and efficient in maximizing local utility.

Though Axelrod's work is interesting and convincing, we believe that the assumptions used in his work makes the results inapplicable in a number of domains of practical interest. In real-life situations, a particular help-giving interaction between two agents often means one agent helps and incurs a cost while the other receives help and obtains a savings in cost or effort. As only one agent decides or acts in each interaction, such interactions are necessarily asymmetrical in nature in contrast to the symmetrical formulation of games like prisoner's dilemma. Another key restrictive feature of Axelrod's experiment with the iterated prisoner's dilemma game is that identical scenarios are repeated. This is not likely in real life as every interaction is different from others. The assumption of repetition of identical scenarios enable Axelrod to work with strategies that do not compare different interactions. In real life, history of interaction will have to capture not only the outcomes, but also the context in which a certain outcome was produced. Also, there has to be a means to compare two different scenarios or two help-giving actions of different magnitude. Comparison of two such different scenarios requires the use of some measure of work or cost involved in help-giving. Such a metric will allow systematic evaluation of different scenarios under different interaction histories.

Based on these observations, we believe that a simple tit-for-tat like deterministic strategy is not adequate for more realistic agent domains [2]. We now identify the desirable features of a behavioral strategy that will be suitable for open environments: a risk attitude that allows the agent to initiate help-giving to a new agent but quickly shun it if requests for help are rejected repeatedly; ability to compare cooperation costs across different scenarios; ability to adjust help-giving behavior based on local work-load.

Over the last few years, multiagent systems researchers have started evaluating non-monetary mechanisms for supporting agent interactions and developing fruitful relationships in a societal context. Our own work [5, 27, 28, 29] have emphasized the use of a probabilistic reciprocity scheme for exchanging help, using which self-interested agents can develop mutually beneficial relationships with other agents while avoiding exploitative agents. Castelfranchi and Falcone have argued the necessity of trust in social interactions between agents with complex mental attitudes [12]. They argue that trust can be based on mental background, and though it necessarily entails risks for delegation and collaboration, considerations of morality and use of reputation can be used to mitigate that risk. We concur with the observation that trust can play a critical role in initiating, nurturing, and supporting collaborative relationships between self-interested agents. Goldman and Rosenschein [16] have suggested an intrinsic *cooperation level* parameter for agents, whereby agents have an innate desire to work to benefit everyone in the environment. This can be viewed as an agent's trust in or attitude towards other agents in the environment. Our reciprocity scheme is a more targeted, one-to-one modeling mechanism, which can be interpreted as the cooperation level of an agent towards another agent depending on their history of interaction.

An example of using morality to promote social relationships, as suggested by Castelfranchi and Falcone, can be found in the SPIRE framework developed by Grosz and collaborators [15, 31]. They examine a framework where agents in a group have to decide whether or not to reneg on commitments to other agents in the group to accept lucrative outside offers. In their model, socially conscious agents combine utility considerations with a *brownie points model*, whereby agents analyze whether their actions will make them “feel good” or be viewed as being a “good guy” by others. Based on their presentation, their agents cannot be termed as self-interested, as these agents evaluate their actions based on not only how that will affect individual payoff but also how the group utility will be affected. In our model, the agents are completely self-interested, and their help-giving behavior is predicated on receiving future help from the other agent to more than compensate for the current help-giving cost.

Cesta, Micelli, and Rizzo evaluate simplistic non-adaptive help-giving, parasitic, and selfish agents in a food-gathering domain [13]. Their results suggest that for very supportive environmental conditions and for a particular choice of evaluation criteria, that suppresses extremes of performance, non-adaptive helping strategies may resist exploitation. But they also report that for a number of environmental settings, selfish and parasitic agent can severely affect the viability of helping agents. A thorough analysis of their work reveals several fundamental weaknesses of agents that do not adapt their behaviors and shun exploiters. In our previous work [27], we have observed similar problems with naive, philanthropic agents, i.e., agents who always help when asked. As these agents get easily exploited by selfish agents, they perform poorly in mixed groups. Hence, we have not included these agents in the current study.

Castelfranchi, Conte, and Paolucci use normative reputation [11] to enhance the performance of agents that comply with social norms. They also experiment with a food-gathering domain, where agents prolong life by finding and consuming food in the environment. Norm-following, *respectful* agents do not attack agents who are consuming food, but aggressors or *cheaters* can snatch food away from those who found it first. Without the use of reputation, cheaters outperform respectful agents. When respectful agents are modified to share their opinion of other agents and to respect the norms only for other agents believed to be respectful, the performance of respectful agent improves to being close to that of cheater agents. There are two basic shortcomings of their approach. The first problem is that a respectful agent believes the opinions provided by another agent, a strategy that is shown to be easily undermined by lying agents in our current work. The second problem is that their deterministic decisions about not following norms for a norm violator is not robust enough to be applied to domains where agents can erroneously or inadvertently violate norms. In domains where an agent cannot help because they are currently occupied or if they fail

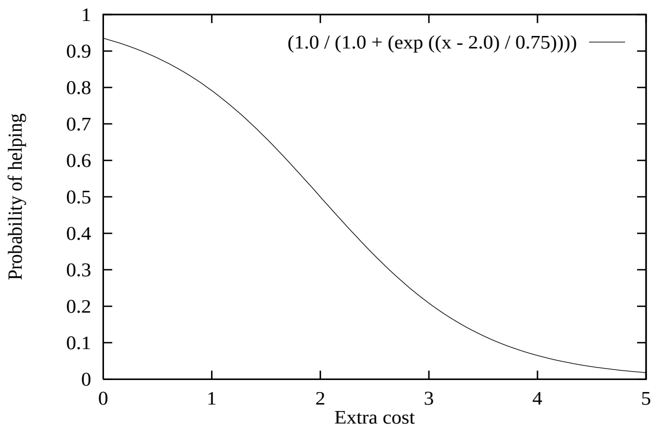

Figure 1. Probability function for accepting request for cooperation.

to complete help-giving behavior because of environmental factors, a decision to shun that agent in the future will be counter-productive and will disrupt the growth and sustenance of collaborative relationships. In a sense, this deterministic reputation mechanism implicitly assumes repetition of identical scenarios and does not compare the costs and benefits of one situation with another. Such simplifications limit the applicability of these strategies in domains where different help-giving situations may result in largely varying costs and savings for the helper and the helped agents. The variants of probabilistic reciprocity strategy, that we evaluate in the current paper, are not susceptible to the above-mentioned problems.

3. Probabilistic reciprocity

Now, we present our probabilistic reciprocity mechanism for deciding whether or not to help an agent who has requested for help. 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 costs it C_{ij}. If agent k carried out this task together with its own task t_{kl}, the cost incurred for task t_{ij} is C_{ij}^{kl}.

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 ($C_{ij} > C_{ij}^{kl}$), 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_{ij}^{kl} but agent i saves a cost of C_{ij}. Since the cost of helping to the helper agent is less than the saving of the helped agent, there exists a cooperation possibility.

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 (note that, in general, $B_{ik} \neq -B_{ki}$). We will see later in Section 5 that in place of the individual balance, an agent can use a combination of balances reported by other agent, i.e., use a social reputation mechanism, to determine if agent i should be helped. 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:

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

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 where the probability of helping increases as the balance increases and is more for less costly tasks[3]. 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 function in Figure 1. The constant β can be used to move the probability curve 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 [17] for the agents. The level of cooperation or the inclination to help another agent is dynamically adapted with problem solving experience. Over time, an agent will adapt to have different cooperation levels for different agents.

We emphasize that we chose the probability function in Equation 1 as it satisfies the desirable features of a behavioral strategy for agent interaction we presented at the end of the last section. In addition, this function provides two well-understood parameters, β and τ, with which agent designers can easily control the degree and nature of cooperation of their agents. Also, a large class of other functions, including step functions and linear functions, can be approximated by this sigmoidal function. We believe that other functional forms can also satisfy a number of our requirements. The choice of this particular function was dictated primarily by the fact that it is well known, well understood, and easy to use.

4. Assumptions

In this section we present several assumptions about the agents and environmental conditions that have motivated the design of the agent strategies and experimental framework that we describe in the following sections:

- We assume that an agent does not change its strategy in the course of an experiment. For example, a selfish agent does not adopt a reciprocative strategy because agents of the latter type are performing well in the environment. A rational agent will be expected to change its strategy based on its observations of the relative performances of agents of different strategies. In the current paper, however, we are primarily interested in evaluating the effectiveness of different strategies and do not address the issue of agents adopting strategies that are perceived to produce higher performance.

- The motivation for the reciprocity work comes from self-interested agents interacting in open environments that abound in cooperation possibilities. In a typical open environment, however, the agents may be using a large number of different strategies. Also, the agents may enter and leave the population at any time resulting in a very volatile agent group. In this paper, we evaluate a restricted class of exploitative and reciprocative strategies in a stable agent population. The particular strategies chosen are representative of their respective classes and are meant to illustrate general properties of such strategies. The variants of reciprocative strategies that we study are designed to reduce the number of interactions necessary between agents to develop mutually beneficial relationships without succumbing to exploitation by malevolent agents. The success of these strategies will allow us to use it in less stable (more volatile) groups. Our goal is to reduce the time period for which a typical agent will have to be a part of the group before reciprocative strategies prove to be more useful than exploitative strategies. While we have made notable progress, it is unlikely that any non-monetary approach will work for a completely open, volatile system where any agent can enter and leave the system at any time, and in particular in those situations where a typical agent interacts with other agents in an environment only once or for very few interactions.

- We assume that helpful agents are also honest, while exploitative agents can be deceitful and can lie to gain undue advantage or hurt other agents. In practical settings, even helpful agents can lie. In those situation, an agent must not only estimate the help-giving nature of another agent, but also the reliability of its opinion about other agents. Techniques similar to action estimation and reinforcement learning [32] can be used to learn such estimates. Our goal in this paper was to focus on the development of trust-based reciprocal relationships, and hence we assumed that agents who are trustworthy help-givers are also trustworthy for their opinions. In domains where such assumptions are expected to be violated, one has to augment the reciprocity mechanism with learning strategies.

- In this paper, we assume that tasks do not require specialization, i.e., any agent can perform a given task. In domains, where agents have differing expertise and tasks require specific expertise, to form fruitful collaborative relationships, agents must learn about the competence level of other agents for different task types [14].

- The *selfish* agents in this paper are not *rational* or self-interested. We have used them to study the effects of disruptive elements in the population. For example, they may not gain significantly by lying, but that behavior can disrupt the formation of beneficial relationships between reciprocative agents.

The communication structure used in this work is simple and is not the focus of our research. When approaching another agent for help, the requesting agent simply states what the task is. The helping agent then decides whether to help the requesting agent by taking over this task. In the process of deciding whether to help or not, the requested agent can ask other agents about their opinion of the help-giving nature of the requesting agent. An agent can be truthful or lying about its opinion of another agent.

We believe that the conclusions drawn in this paper will hold for domains and environments where the following conditions are met:

- The composition of agent group is stable for some amount of time as measured by number of tasks executed or the number of interactions between agents.

- There exists sufficient number of cooperation possibilities (the cost of helping to the helper is less than the saving obtained by the helped agent) with roughly symmetrical possibilities, i.e., sometime one agent can help another and at other times the roles are reversed.

- Agent strategies are fixed for the period under consideration.

In the real-world, we observe reciprocal, mutually beneficial, non-monetary relations develop universally among neighbors, friends, colleagues, etc. who help each other by taking on chores, extra work, etc. without any monetary compensation. Even national and international agencies share information routinely with counterparts as such reciprocal sharing is recognized to be an effective and cost-saving approach to gathering high-quality information which may be difficult or costly to obtain otherwise. We believe that development of reciprocative relationships among diverse groups in widely varying domains cannot be explained by simply attributing this phenomenon to lack of rationality of the participating agencies. One possible benefit of such non-monetary relationships is to eliminate the need for calculating a utility value for each help-giving and help-seeking behavior, thereby reducing the computational load on the agent. In our formulation in this paper, we use time costs for determining help-giving decisions. Time estimates are much easier to form and use and does not require a considerably more complex and uncertain utility calculation.

As a practical example of a an application domain, software agents performing information tasks on behalf of associated users can help each other by sharing information, processing capability, etc. For example personal assistant agents of two users charged with fetching general news information and financial news may decide to split up the work where one gathers general news and the other gathers financial new and then share the retrieved and processed information. This way each agent spends considerably less system resources to provide the required service. We have used a simulation of such a domain to evaluate the basic reciprocity strategy augmented with learning capabilities [14].

5. Agent strategies

There are two types agents that we have used in our previous work on which we will expand in this paper:

Selfish agents: Agents who will request for cooperation but never accept a cooperation request. Selfish agents can benefit in the presence of philanthropic agents (agents who always help when asked) by exploiting their benevolence.

Reciprocative agents: Agents that uses the balance of cost and savings to stochastically decide whether to accept a given request for cooperation.

The augmentations on these strategies are as follows:

Believing reciprocative agents: These are agents who use not only their own balance with another agent, but also the balances as reported by all other agents when deciding whether or not to provide help. More precisely, 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[4].

Learned-Trust based reciprocative agents: These agents also use combined balances, but includes balances of only those agents with whom it has a favorable balance. More precisely, in place of using B_{ki} in Equation 1, a learned-trust based reciprocative agent k uses $\sum_{j \neq i \wedge B_{kj} > 0} B_{ji}$ while calculating the probability of helping agent i[5].

Individual lying selfish agents: These agents are designed to exploit the fact that believing or trusting reciprocative agents use balances provided by other agents. These agents reveal false impressions about other helpful agents to ruin their reputation[6]. More precisely, when such an agent, j, is asked for its balance with another agent i, it reveals B'_{ji} given by:

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

where C is a positive constant. This means that 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 try to spoil the reputation of helping agents, but also collaboratively bolsters the reputation of other selfish agents or agents with whom it has zero balance. More precisely, when such an agent, j is asked for its balance with another agent i, it reveals B'_{ji} given by:

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

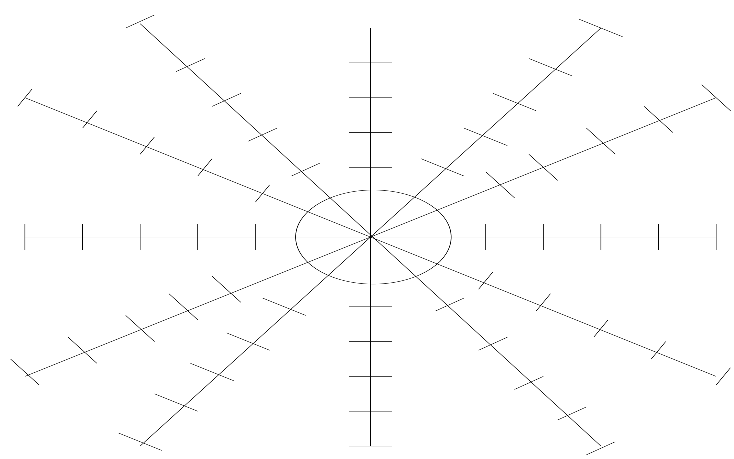

Figure 2. An agent picks up one package from the central depot and delivers it to one of the marked locations on one of the radial fins before going back to the depot to retrieve its next package.

where C is a positive constant as above and $\mathcal{P}$ is a large positive constant. Note that we assume that since the selfish agent never helps anyone, other agents with whom it has 0 balance is to be treated as selfish agents. This means, initially it treats all agents equivalently. Only when the reciprocative agents start helping it does a collaborative lying selfish agent turn against them!

6. Package Delivery Domain

In the simple package delivery problem that we have used for experimentally evaluating strategies, we assume there are 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 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 (see Figure 2). 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 this extra packet. In addition, if the helping agent is going beyond its destination without this packet, then unit extra cost per unit distance is incurred to account for the return journey. If d_1 is the maximum destination distance of all the packets the helping agent is now planning to deliver, and d_2 is the destination distance of the packet of the agent requesting help, then the extra cost is given by:

$$e(d_1, d_2) = \begin{cases} d_2 & \text{if } d_2 \leq d_1 \\ d_2 + (d_2 - d_1) & \text{otherwise} \end{cases}$$

7. Experimental results

In this section, we present experimental results on the package delivery problem with agents using the reciprocity mechanism described in Section 3 to decide whether or not to honor a request for cooperation from another agent (see Figures 3–8). Unless otherwise noted, the parameters for the experiments are as follows: $N = 100$, $T = 500$, $R = 4$, $D = 3$, $\tau = 0.75$, $\beta = 0.5$, $C = 1$, and $\mathcal{P} = 10$. Each of our experiments are 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. We also ensure that all agents are assigned packages such that sum of the destination distances of packages are the same. The evaluation metric is the average time taken by the agents to complete all the deliveries and hence a lower value of this metric is preferred over a larger value. Experiments were run with mixed group of reciprocative and selfish agents. In each figure we present the average performance of the reciprocative agents (Reci), the selfish agents (Self) and the entire group of agents (All).

7.1. Performance in different mixed populations

The first set of experiments we report is from our previous work where reciprocative and selfish agents are evaluated in mixed groups while varying the percentage of selfish agents. In this set of experiments, a $\beta = 0.5$, which means the reciprocative agents are quite cautious about giving help. From the results presented in Figure 3 we see that though the selfish agents are able to exploit the reciprocative agents to reduce their delivery cost (if they had to deliver all of their packets by themselves, their average time taken would be approximately 2000), the reciprocative agents outperform the selfish agents for a wide range of group composition. Only when the percentage of selfish become a large majority ($\geq 80\%$) do the selfish performance dominate the

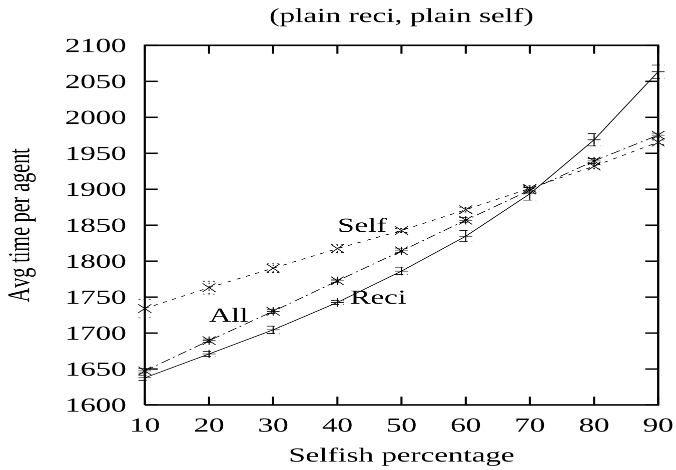

Figure 3. Performance of cautions Reciprocative and Selfish agents in mixed groups (100 agents, 500 tasks/agent, $\beta = 0.5$, $\tau = 0.75$).

performance of the reciprocatives. The performance of both group of agents, and correspondingly the average for all agents, deteriorate with increasing selfish agent percentage in the population. The performance of reciprocative agents decrease as there are fewer reciprocatives to form mutually beneficial partnerships, and also because there are more selfish agents who do not reciprocate help-giving behavior. At about 90% selfish population percentage, the reciprocatives end up doing more work than they would have done if each agent had ignored all other agents and just completed delivery of the packages assigned to it. The performance of selfish agents deteriorate as there are fewer reciprocative agents from whom they can extract help.

Next, we ran a set of experiments where we decreased the cautiousness of the reciprocative agents by increasing β to 2. This meant that the reciprocative agents were willing to incur a larger up-front helping cost to jump-start cooperative relationships. This decrease of cautiousness, however, also meant that selfish agents can exploit reciprocative agents more often and for larger gains. This intuition was verified from the experimental results presented in Figure 4. We observe that the performance of the selfish dominate that of the reciprocatives at all population mixes. From Figures 3 and 4 we see that for selfish percentage of 10%, both the reciprocatives and the selfish perform better with $\beta = 2$ than with $\beta = 0.5$. As the cooperation level, i.e., the willingness of agents to help others, increase, reciprocatives can form more beneficial partnerships with other reciprocatives. The same attitude can also be exploited by the selfish agents for more gains. When

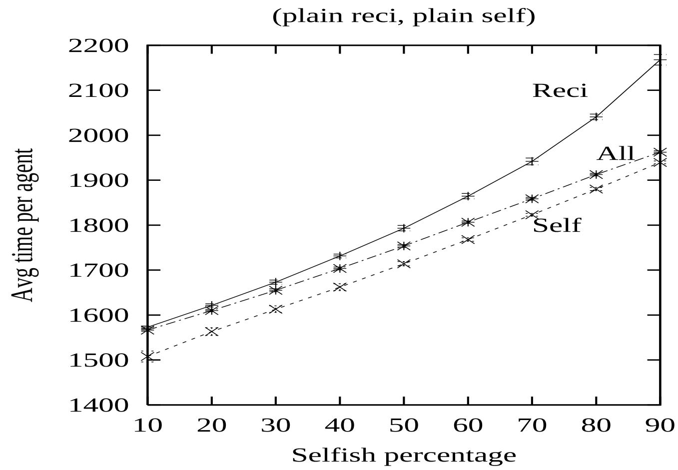

Figure 4. Performance of less cautious Reciprocative and Selfish agents in mixed groups (100 agents, 500 tasks/agent, $\beta = 2$, $\tau = 0.75$).

the percentage of selfish agents in the population is small, the increased loss of reciprocative agents to selfish agents is compensated by increased gains from other reciprocative agents. That is why they perform better with $\beta = 2$ than with $\beta = 0.5$. The corresponding increased gains of the selfish is sufficient to outperform the reciprocatives. As the selfish percentage increases, the performance of both groups deteriorate for reasons mentioned above. When selfish percentage increases beyond 50% the performance gulf between the selfish and the reciprocative rapidly widens.

The above two sets of experiments underlines the need for agent designers to set the β values or cooperation levels appropriately such that mutually beneficial relationships with other reciprocative agents can be nurtured without exposing the reciprocative agents to overt exploitation by selfish agents. The tradeoff can be summarized as follows: with increasing β values, or cooperation level, reciprocatives can more consistently identify other reciprocatives as beneficial partners, but reciprocatives become more vulnerable to increased exploitation by selfish agents.

We hypothesized that a way out of this dilemma would be for reciprocatives to use the opinions of other agents about an agent who has sought help before making the help-giving decision. The underlying motivation is that if the reciprocative agents could share their balances, an agent that refuses to reciprocate help will be identified early by all reciprocative agents. Such early identification will severely limit the exploitative potential of these selfish agents and also enable the

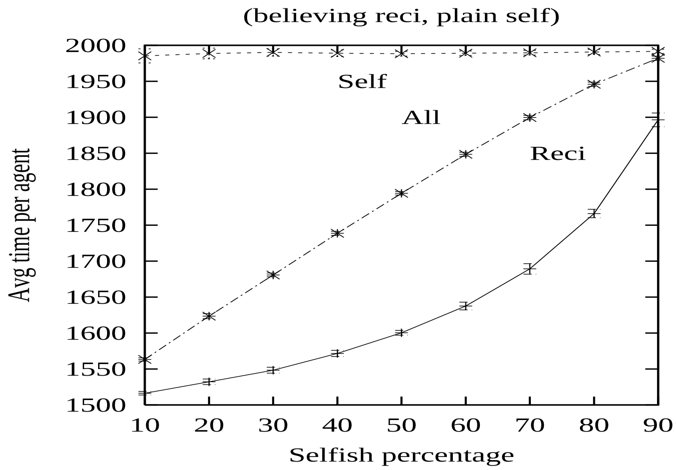

Figure 5. Performance of believing reciprocative and selfish agents in mixed groups (100 agents, 500 tasks/agent, $\beta = 2$, $\tau = 0.75$).

reciprocative agents to perform better by eliminating cost incurred in helping these selfish agents. This line of reasoning led us to designing the believing and learned-trust based reciprocative agents.

In the next set of experiments we evaluated mixed groups of believing reciprocative agents and selfish agents. For this and the following set of experiments, unless otherwise noted, we used $\beta = 2$. We wanted to evaluate the effectiveness of augmented reciprocity mechanisms for producing good performance for a reasonable range of values of β and τ. The goal was to develop mechanisms such that the agent designers do not have to spend significant amount of time on fine tuning these parameter values.

As we see from the results presented in Figure 5, the sharing of balances does indeed severely restrict the exploitative edge of the selfish agents. More importantly, the increase in the proportion of selfish agents in the population does not noticeably increase their capability of exploiting the reciprocative agents. As expected, the early identification of selfish agents also enable the reciprocative agents to improve their performance significantly. At selfish percentage of 10%, the believing reciprocative agents perform much better than the corresponding group of reciprocative agents who do not incorporate opinions of other agents (see Figure 4). The average performance of reciprocative agents suffer with increasing percentage of selfish agents as there are fewer and fewer reciprocative agents with whom mutually beneficial, i.e., cost saving, relationships can be formed.

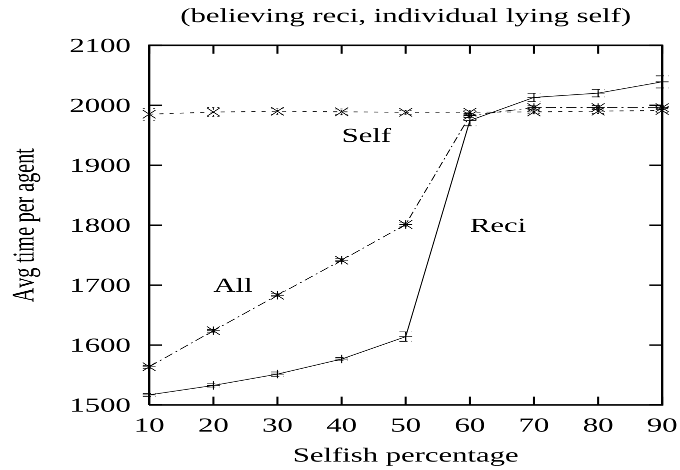

Figure 6. Performance of believing reciprocative and individual lying selfish agents in mixed groups(100 agents, 500 tasks/agent, $\beta = 2$, $\tau = 0.75$).

Even though the believing reciprocity approach appear to be effective from these experiments, it has a serious shortcoming. As it does not know, a priori, which of the other agents are selfish or cooperative, a believing reciprocative agent includes balances from all other agents in its calculations. The selfish agents in the population, therefore, have the opportunity and the incentive to undermine and disrupt this word-of-mouth reputation mechanism by giving false balances about other agents.

In the next set of experiments, we form mixed groups of believing reciprocative agents and individual lying selfish agents. From Figure 6 we observe that when there are few selfish agents, their lying behavior does not noticeably affect the performance of believing reciprocative agents. But as the the percentage of such lying agents increases above a threshold of about 50%, critical mass of negative information surmounts the positive impression created by mutual help between reciprocative agents. At this point the reciprocative agents stop helping each other, and since they do not receive any help from selfish agents, they end up doing all of their work by themselves. With further increase in the percentage of selfish agents, the reciprocatives under-perform the selfish agents as the lying agents are able to extract some help from the reciprocative agents.

A more sinister form of lying occurs when selfish agents collude not only to vilify the reputation of reciprocative agents, but falsely tout themselves to be helpful. The believing reciprocative agent are gullible enough to be swayed by this false group impression which will even

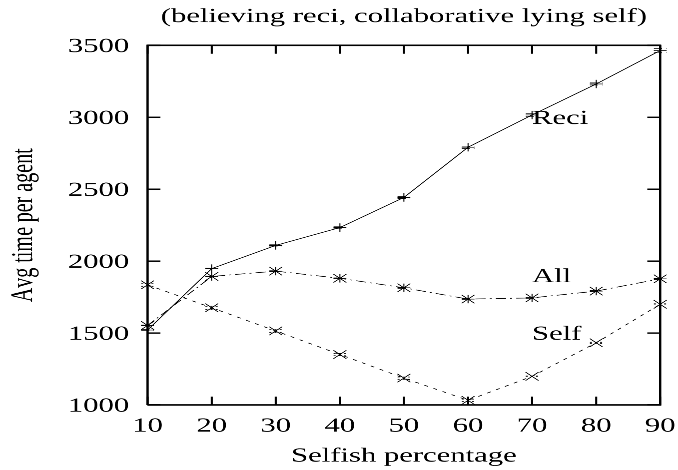

Figure 7. Performance of believing Reciprocative and Collaborative lying Selfish in its calculations agents in mixed groups (100 agents, 500 tasks/agent, $\beta = 2$, $\tau = 0.75$).

override any negative balance it might have individually with those agents. This is actually the other extreme of the effect of word-of-mouth reputation schemes or group balances: instead of correctly identifying "bad guys", now one will incorrectly identify the bad guys as "good guys."

In this set of experiments, we experimented with mixed groups of believing reciprocative agents and collaborative lying selfish agents. From Figure 7 we observe that the collaborative lying agents are able to exploit the reciprocative agents quite effectively and overwhelms them when their percentage in the group is more than 20%. In contrast to the individually lying agents, the collaborative lying agents not only cause poor performance of reciprocative agents, but saves themselves significant delivery costs by receiving help from the reciprocative agents. The reciprocative agents end up doing about 75% more work than if they had just delivered all of their assigned packets on their own, i.e., if they had never explored collaboration possibilities by helping other agents. In contrast, the lying selfish agents can cut their workload by about half when their percentage in the population is about 60%. It is interesting to note that the performance of the selfish agents start deteriorating when selfish percentages increase above 70%. This happens because there are fewer reciprocative agents to receive benefit from. It also means that at less than 50% selfish percentage, there are not enough selfish agents to exploit the reciprocative agents to the extent they are susceptible to exploitation.

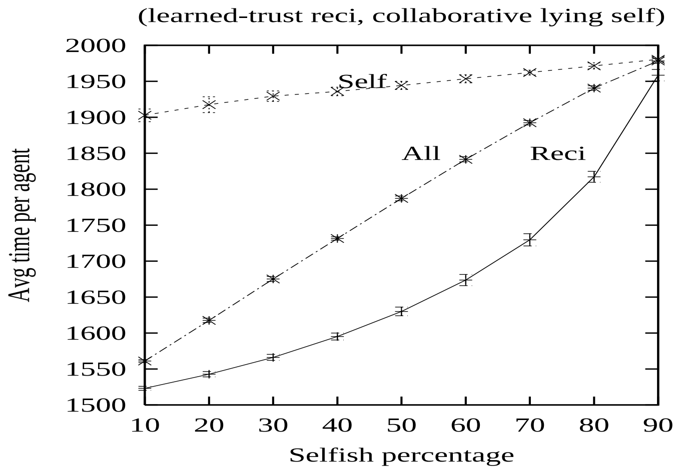

Figure 8. Performance of Learned-trust based Reciprocative and Collaborative lying Selfish agents in mixed groups (100 agents, 500 tasks/agent, $\beta = 2$, $\tau = 0.75$).

It is clear that collaborative lying is a threat which, if not countered, will make the believing reciprocative strategy unusable. One can always revert to using the base reciprocative agent, which does not believe others, and hence is not susceptible to either individual or group lying. But then we have to be happy to concede non-trivial exploitation by even non-lying selfish agents. Our conjecture for a fix to this problem was to alter the believing reciprocative agent strategy to believe only those agents who have proven to be helpful in the past. That is, if someone has consistently been of help, it is reasonable to believe its opinion. On the other hand it is unwise to believe someone who has not reciprocated prior help-giving behaviors. We believed that such a learned-trust based reciprocative agent strategy may withstand both individual and collaborative lying by selfish agents.

In this set of experiments, we evaluated mixed groups of learned-trust based reciprocative and collaborative lying selfish agents. Results presented in Figure 8 show a significant performance improvement for reciprocative agents. The amount of help received by the lying selfish agents is much less than what non-lying selfish agents received from basic reciprocative agents (see Figure 4). An interesting observation is the level of exploitation and hence the performance of selfish and reciprocative agents vary only by a small amount over different group mixes. This set of experiments clearly demonstrated that learned-trust based reciprocative agents can effectively handle lying selfish agents. This variant of reciprocative agents are also able to effectively deal with selfish agents who do not lie. The performance curves of mixed

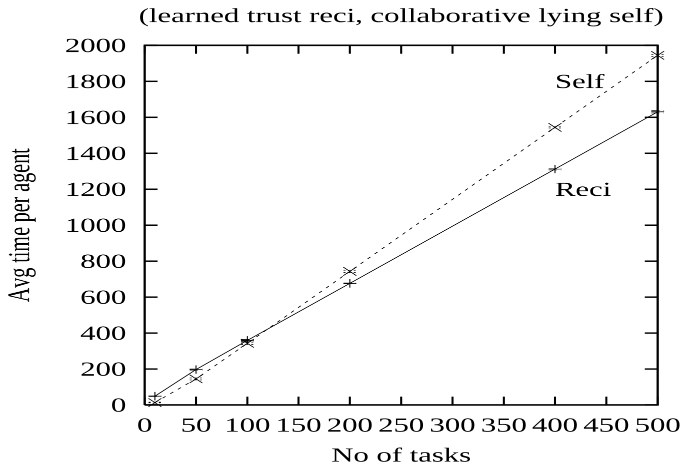

Figure 9. Performance of Learned-trust based Reciprocative and Collaborative lying Selfish agents in mixed groups when varying number of tasks (100 agents, 50% selfish agents, $\beta = 2$, $\tau = 0.75$).

group of learned-trust based reciprocative and individual selfish agents were identical to those for learned-trust based reciprocative and collaborative lying selfish agents. This is because, the learned-trust based reciprocative strategy does not consider the opinion of agents that are not trusted. So, it does not matter if that agent is saying bad things about other reciprocative agents or praising other selfish agents as well.

7.2. Performance sensitivity to parameter values

To further evaluate the robustness of our proposed strategies we run experiments by varying system and agent parameters with mixed groups of learned-trust based reciprocative agents and collaborative lying agents. In all of the following set of experiments, the population was split equally between selfish and reciprocative agents.

In the first set of experiments in this group, we varied the number of tasks assigned, i.e., packages to be delivered, to each agent from 10 to 500. The number of agents used was 100. From the results plotted in Figure 9 we see that when agents have to deliver only a few tasks, the selfish agents can perform marginally better, but as the number of tasks increase beyond 100, the reciprocatives dominate. We have observed that the dominance of reciprocatives happen with even fewer tasks if the selfish agents are not lying. These two results combine to show that learned-trust based reciprocative strategies are dominant when agents interact for much shorter periods compared with basic reciprocity strategy (which takes about 500 tasks per agent to dominate

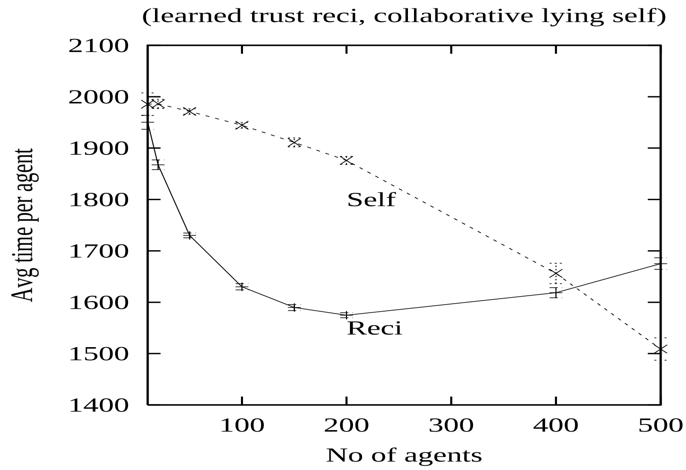

Figure 10. Performance of Learned-trust based Reciprocative and Collaborative lying Selfish agents in mixed groups when varying number of agents (500 tasks per agent, 50% selfish agents, $\beta = 2$, $\tau = 0.75$).

selfish strategies for similar environmental parameters). The performance difference between reciprocative and selfish agents increases with increasing number of tasks. When agents have to deliver only a few tasks, selfish agents can exploit different reciprocative agents before their bad reputation catches on. For the believing reciprocative strategy to be effective reciprocative agents must first identify helpful partners and then use their opinion to shun selfish agents. In this initial period of trust development, reciprocative agents are vulnerable to exploitation by selfish agents. When agents have to deliver only a small number of tasks, this initial period of vulnerability is sufficient for selfish agents to extract enough benefits to outperform reciprocative agents.

In the next set of experiments, we varied the number of agents from 50 to 500, where each agent was assigned the delivery of 500 packages. From Figure 9 we see that the performance of both selfish and reciprocative agents improve as the number of agents increase. This is because as the number of agents increase there are also more agents who are willing to provide help. Initially, the performance of the reciprocative agents improve dramatically, but when the number of agents increases beyond 200, their performance starts to degrade gradually. At this stage the performance of the selfish agents steadily improve and they outperform the reciprocative agents at population size of 500. This happens because with a large population, each reciprocative agent gets exploited a little by a large number of selfish agents. This loss cannot be compensated by gains from the increased number of

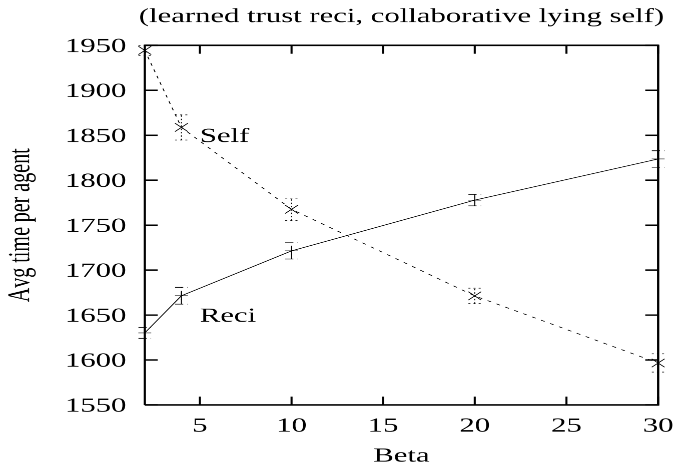

Figure 11. Performance of Learned-trust based Reciprocative and Collaborative lying Selfish agents in mixed groups when varying β (500 tasks per agent, 50% selfish agents, $N = 100$, $\tau = 0.75$).

reciprocative agents in the population unless there are sufficient number of tasks to be delivered. We conjecture that if the number of tasks are increased from 500, the population size at which selfish performance overtakes reciprocative performance will also increase. Moreover, the reciprocative will continue to dominate the selfish if the number of tasks per agent is proportional to the number of agents. The latter situation is not realistic in practice. Hence, in large populations, it is perhaps more effective to be cautious or use a decreased cooperation level, i.e., use a smaller value of β.

In the next set of experiments, we varied the cooperation level or the inclination of reciprocative agents to incur initial cooperation costs to identify mutually helpful partners. We varied β from 0.5 to 30 keeping population size at 100 and number of deliveries at 500 tasks per agent. Results from Figure 11 show that with increasing β value reciprocative performance degrade and selfish performance improve until selfish starts dominating at $\beta > 15$. This is expected because if each reciprocative agent is willing to incur significant helping costs before turning its back on exploiters, the selfish agents can extract enough help to reduce its own workload. The exact β value for performance crossover of selfish and reciprocative agents will depend on the number of agents in the population, the percentage of selfish agents, the number of tasks per agent, etc. But the general nature of the plots presented in Figure 11 will hold for other values of strategy and system parameters.

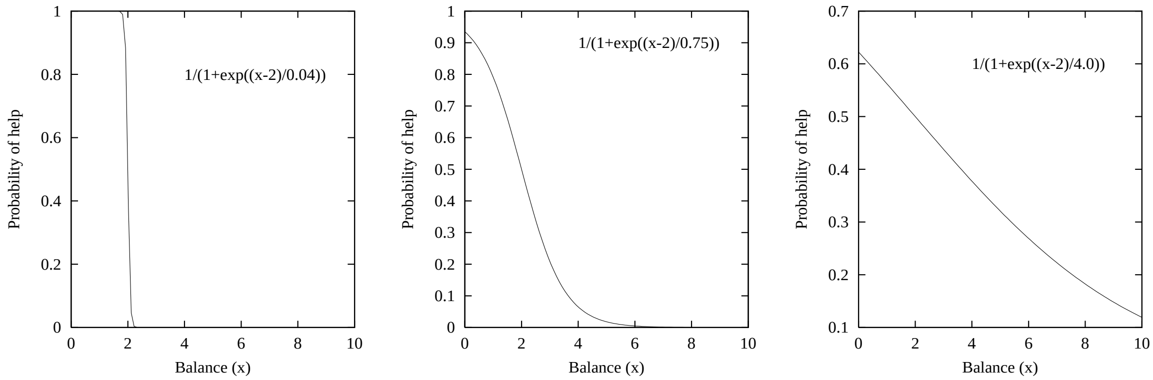

Figure 12. Probability functions for $\tau = 0.04$, 0.75, and 4.

Table I. Performance with varying τ value.

τ	Reci-avg-time	Std-dev	Self-avg-time	Std-dev
0.04	1616.45	5.32509	1968.82	3.71549
0.75	1630.09	6.06221	1944.33	5.6603
4.0	1677.37	7.3255	1858.2	12.1221

In the final set of experiments, we used τ values of 0.04, 0.75, and 4 to denote three different shapes of the probability function used to make help giving decisions (see Equation 1). The three probability functions are depicted in Figure 12. Strategy and system parameters used in this set of experiments were as follows: $N = 100$, $T = 500$, $\beta = 2$, selfish percentage = 50%. From Table I, we see that reciprocative performance dominate selfish performance for all three τ values used. With increasing τ, however, there is a small decrease in performance of reciprocative agents and a corresponding increase in performance of selfish agents.

The above set of experiments demonstrate that the believing reciprocative agents are successful in identifying and benefiting from mutually cooperative relationships and resisting exploitation by lying selfish agents for a wide range of system and strategy parameters. Though we present experimental results under certain values of parameters like R, D, etc., we have observed that the qualitative nature of the plots are consistent across a much larger set of parameter values. We also believe our results are more robust than those derived by Axelrod, and most of the criticisms levied against that body of work [4] does not apply to our case. For example, learned trust based reciprocative agents will be able to collaboratively identify and then shun exploiters if the agents interact for any extended period of time. The deterministic tit-for-tat strategy cannot be used effectively in the package delivery and other real-life task-based situations where the costs for performing different tasks can be significantly different. For example, an exploitative agent can help a tit-for-tat agent on a task with small cost and then extract a large cost from the latter by assigning it a much more complex, time-intensive, or costly task. There are other, related problems with the tit-for-tat strategy. In our previous experiments comparing tit-for-tat strategies, we have found that the performance variations of these agents are significant [27]. This means that one tit-for-tat agent can become envious of other tit-for-tat agents and change strategy. The standard deviation of our probabilistic reciprocity agents are quite small, and hence these strategies are much more stable.

8. Conclusions and Future Work

In this paper, we consider the effects of believing other agents' opinions when deciding to help an agent. Such pooling of opinions is found to effectively restrict the exploitative gains of selfish agents. We then investigate the performance of lying selfish agents, where both individual and group level exploitative schemes may be used. We study the weaknesses of the probabilistic reciprocity based help-giving strategy that uses other's opinions to develop individual and group based exploitative strategies. These schemes are shown to be able to "invade" a homogeneous group of believing reciprocative agents, the latter being particularly susceptible to group exploitation by lying selfish agents. We introduce an experience based trust mechanism for reciprocative agents that is able to successfully withstand invasion by both individual and group level exploitative schemes. The addition of the trust mechanism then restores the stability of the probabilistic reciprocity based strategy. The learned trust based strategy will enable self-interested agents take advantage of cooperation possibilities in the environment by developing stable, mutually beneficial relationships with similar agents without being exposed to exploitation by malevolent agents. As a result both individual and system level performance can be significantly improved.

One of our future goals is to analytically capture the dynamics of the evolution of balance of helps in homogeneous and heterogeneous groups. For example, given a particular group composition and random interactions between members, how do the balances of selfish and reciprocative agents change as a function of time. Difference or differential equation models can be constructed to represent the dynamics of these societies. In addition to identifying the ascendancy of exploitative or cooperative relationships, such models can also allow us to identify the formation of demes or working coalitions based on interaction histories.

We are currently studying the viability of these strategies in an evolutionary setting. In such a scenario, the population would start with some random proportion of different agent types. Based on the relative performances of the different agent types the population would be modified either continually or periodically with "clones" of better performing individuals replacing individuals performing poorly. Such a scenario corresponds to the real-world scenario where self-interested agents are not limited to a fixed strategy but can adopt behaviors that is found to be more effective in practice. It would be instructive to study both the final converged populations given different initial distributions and also to analyze the dynamics of the population as it evolves.

Acknowledgements

This work has been supported in part by an NSF CAREER award IIS-9702672. We would like to acknowledge the programming efforts of Anish Biswas and Sandip Debnath. Special thanks go to Partha Sarathi Dutta for redesigning the code and re-running the experiments as well as generating the plots. The comments of anonymous reviewers were of invaluable help in revising and improving both the content and the presentation of this paper.

Notes

[1] It is often argued that all interactions can be assigned to economic agents. If, in the future, all interactions between any two computational entities on the Internet involved monetary exchanges, then either these agents or their owners have to decide on whether to interact or conserve its monetary allocation for some more important or urgent task that may arrive later. For example, my information gathering agent has to decide between whether to proactively search for information on the net (for which it has to pay) or reactively react to my search requests once it has been allocated $X for the day. This decision making may be difficult to optimize as my requests may vary widely over different days, and I will not take kindly to my agent who cannot process my explicit request because it has already spent its allocation on proactive searches which may have generated useful information but is of less importance to me right now. Neither do I want to micro-manage this monetary allocation to my agent as then the purpose of having an automated assistant is defeated.

[2] There are other, orthogonal criticisms to the generality of the conclusions drawn in Axelrod's work [4, 22].

[3] Note that this function does not represent a probability distribution. In particular $f(x)$ gives the probability that the agent will agree to help when the cost of helping is x. $f(x)$ and $1 - f(x)$ together determine the probability distribution for helping cost x, where the only two options for the agent is to accept or deny the request for help. Also, there does not need to be any correlation between $f(x)$ and $f(y)$ values, where $x \neq y$.

[4] We assume that while k is deciding to help i it finds out the balances that everyone else has with i, but does not ask i itself about it. If k were to ask i about its balance with others, lying agents would be able to easily exploit k.

[5] A key assumption of this strategy is that helpful agents are also truthful, i.e., agents who have reciprocated or provided help in the past are also likely to provide honest estimates of the helping nature of other agents. This assumption, of course, can be violated in practice where helpful agents can lie to prevent its partner from engaging with other agents and thereby being less available for gainful interaction. A solution out of this dilemma is to separately learn the truthfulness and helpful nature of an agent. That extension is beyond the scope of the current paper, where we want to primarily evaluate the feasibility of identifying helpful partners in the presence of deceitful selfish agents.

[6] Another motivation for an agent A to give bad rating for another helpful agent B is to deter other agents from seeking B's help, thereby leaving B available to help A more often.

References

1. Armstrong, A. A. and E. H. Durfee: 1998, 'Mixing and Memory: Emergent Cooperation in an Information Marketplace'. In: *Proceedings of the Third International Conference on Multiagent Systems*. Los Alamitos, CA, pp. 34–41.
2. Axelrod, R.: 1984, *The Evolution of Cooperation*. Basic Books.
3. Azoulay-Schwartz, R. and S. Kraus: 2001, 'Stable Strategies for Sharing Information among Agents'. In: *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*. pp. 1128–1134.
4. Binmore, K.: 1994, *Playing Fair: Game Theory and the Social Contract*. Cambridge: MIT Press.
5. Biswas, A., S. Sen, and S. Debnath: 2000, 'Limiting Deception in Groups of Social Agents'. *Applied Artificial Intelligence, Special Issue on Deception, Fraud, and Trust in Agent Societies* **14**(8), 785–797.
6. Boyd, R.: 1988, 'Is the repeated Prisoner's Dilemma a good model of reciprocal altruism?'. *Ethol. Sociobiology* **9**, 211–222.
7. Bradshaw, J. M.: 1997, *Software Agents*. Menlo Park, CA: AAAI Press/The MIT Press.
8. Briggs, W. and D. Cook: 1995, 'Flexible Social Laws'. In: *Proceedings of the International Joint Conference on Artificial Intelligence*. pp. 688–693.
9. CACM July 1994 issue: 1994, 'Communications of the ACM, July 1994, Volume 37, Number 7'. Special Issue on Intelligent Agents.
10. CACM March 1999 issue: 1999, 'Communications of the ACM, March 1999, Volume 42, Number 3'. Special Issue on Multiagent Systems on the Net and Agents in E-commerce.

11. Castelfranchi, C., R. Conte, and M. Paolucci: 1998, 'Normative reputation and the costs of compliance'. *Journal of Artificial Societies and Social Simulation* **1**(3).
12. Castelfranchi, C. and R. Falcone: 1998, 'Principles of trust for MAS: Cognitive autonomy, social importance, and quantification'. In: *Proceedings of the Third International Conference on Multiagent Systems*. Los Alamitos, CA, pp. 72–79.
13. Cesta, A. and M. Miceli: 1996, 'Help Under Risky Conditions: Robustness of the Social Attitude and System Performance'. In: *Proceedings of the Second International Conference on Multiagent Systems*. Menlo Park, CA, pp. 18–25.
14. Dutta, P. S. and S. Sen: 2001, 'Identifying partners and Sustenance of Stable, Effective Coalitions'. In: *Proceedings of the 5th International Conference on Autonomous Agents*. New York: NY, pp. 23–24.
15. Glass, A. and B. Grosz: 2000, 'Socially Conscious Decision-Making'. In: *Proceedings of the Fourth International Conference on Autonomous Agents*. New York, NY, pp. 217–224.
16. Goldman, C. and J. S. Rosenschein: 1993, 'Emergent Coordination through the Use of Cooperative State-Changing Rules'. In: *Working Papers of the 12th International Workshop on Distributed Artificial Intelligence*. pp. 171–185.
17. Goldman, C. and J. S. Rosenschein: 1994, 'Emergent Coordination through the Use of Cooperative State-Changing Rules'. In: *Proceedings of the Twelfth National Conference on Artificial Intelligence*. pp. 408–413.
18. Hoffman, E., K. A. McCabe, and V. L. Smith: 1998, 'Behavioral foundations of reciprocity: experimental economics and evolutionary psychology'. *Economic Inquiry* **36**(3), 335–352.
19. Huhns, M. N. and M. P. Singh: 1997, *Readings in Agents*. San Francisco, CA: Morgan Kaufmann.
20. Jennings, N., K. Sycara, and M. Wooldridge: 1998, 'A Roadmap of Agent Research and Development'. *International Journal of Autonomous Agents and Multi-Agent Systems* **1**(1), 7–38.
21. Krebs, D.: 1970, 'Altruism – an examination of the concept and a review of the literature'. *Psychological Bulletin* **73**(4), 258–302.
22. Martinez-Coll, J. and J. Hirshleifer: 1991, 'The limits of reciprocity'. *Rationality and Society* **3**, 35–64.
23. Rapoport, A.: 1989, 'Prisoner's dilemma'. In: J. Eatwell, M. Milgate, and P. Newman (eds.): *The New Palgrave: Game Theory*. London: Macmillan, pp. 199–204.
24. Sandholm, T. and V. Lesser: 1997, 'Coalitions among Computationally Bounded Agents'. *Artificial Intelligence* **94**(1), 99–137.
25. Schmitz, D.: 1993, 'Reasons for Altruism'. *Social Philosophy & Policy* **10**(1), 52–68.
26. Sekaran, M. and S. Sen: 1995, 'To help or not to help'. In: *Seventeenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ, pp. 736–741.
27. Sen, S.: 1996, 'Reciprocity: a foundational principle for promoting cooperative behavior among self-interested agents'. In: *Proceedings of the Second International Conference on Multiagent Systems*. Menlo Park, CA, pp. 315–321.
28. Sen, S. and A. Biswas: 1998, 'Effects of misconception on reciprocative agents'. In: *Proceedings of the Second International Conference on Autonomous Agents*. New York: NY, pp. 430–435.

29. Sen, S., A. Biswas, and S. Debnath: 2000, 'Believing Others: Pros and Cons'. In: *Proceedings of the Fourth International Conference on Multiagent Systems*. Los Alamitos, CA, pp. 279–285.
30. Shoham, Y. and M. Tennenholtz: 1992, 'On the Synthesis of Useful Social Laws for Artificial Agent Societies (Preliminary Report)'. In: *Proceedings of the National Conference on Artificial Intelligence*. San Jose, California, pp. 276–281.
31. Sullivan, D. G., B. Grosz, and S. Kraus: 2000, 'Intention Reconciliation by Collaborative Agents'. In: *Proceedings of the Fourth International Conference on Multiagent Systems*. Los Alamitos, CA, pp. 293–300.
32. Sutton, R. S. and A. G. Barto: 1998, *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.
33. Trivers, R.: 1972, 'The Evolution of Reciprocal Altruism'. *Quarterly Review of Biology* **46**, 35–57.
34. Weiß, G.: 1999, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. Cambridge, MA: The MIT Press.
35. Wellman, M. P.: 1993, 'A market-oriented programming environment and its application to distributed multicommodity flow problems'. *Journal of Artificial Intelligence Research* **1**(1), 1–22.