Effect of individual opinions on group interactions

Parijat Prosun Kar, Sandip Sen, & Partha Sarathi Dutta
Mathematical & Computer Sciences Department
University of Tulsa
Tulsa, OK 74104, USA
e-mail: karpa@ens.utulsa.edu, sandip@kolkata.mcs.utulsa.edu

Abstract

We have evaluated the effectiveness of a probabilistic reciprocity scheme for promoting cooperation among self-interested agents. The probabilistic reciprocity mechanism is used to determine whether an agent should cooperate when approached for help by another agent. The situation becomes more complex when a group of agents seek help from another group. The opinion of the members of the helping group about each of the asking group members can be combined to evaluate such a request for help. Exploitative agents would want to be part of groups that receive helps from other groups, but will try to prevent its group from helping other groups. Such agents, revealing false opinion about the reputation of others, can cause unwarranted rejection of help requests from other groups. This leads to global performance degradation in terms of reduced inter-group cooperation, and increased cost for the individual agents. We study the viability of reciprocative agents in randomly formed groups and when groups are formed by agents contracting other helpful agents. Group helping decisions are based on both average and worst combined ratings of group members. A key result from out study is that lying exploitative agents, who provide false opinion about other agents, become ineffective when focused group selection is enabled.

Keywords: reciprocative agents, group, lazy-liar agents

1 Introduction

As agent-based systems continue to be incorporated into novel practical applications, the likelihood of agents interacting in open environments increase [3, 4]. Interaction strategies to promote cooperation in groups of self-interested agents and thwart malevolent behavior of exploitative agents have been an active area of study for social scientists and multiagent researchers [1, 2, 5, 6, 7, 11]. Most of these studies, however, focus on inter-agent interactions rather than interactions between groups of agents [9].

We draw our motivation from a prior work by Sen *et al.* [10] that evaluated the relative performance of different agent behavioral strategies for producing cost savings for the agents. Single agent interactions were considered and two variants of agent behaviors were evaluated, viz., *selfish* and *probabilistic reciprocity*. The selfish agents would ask for help, but never extend help to anyone. *Lying selfish* agents were used who "bad-mouthed" others by reporting inverted ratings of their past interaction histories, i.e., agents who had helped them were reported to have received more help, etc. Reciprocative agents used a probability based mechanism that was based on the balance of past help with other agents to decide whether or not to honor a help request. A variation of reciprocative agents, who based their decision on the opinion of others about the agent asking for help, were also introduced and was found to be more robust in face of the exploitations of the selfish.

In this paper, we extend the inter-agent reciprocity to facilitate interactions among agent groups. We consider a generic situation where tasks are generated and assigned to a group of agents. A task consists of multiple sub-tasks, each requiring a different expertise to be finished. The group that the task is assigned to checks whether its members are able to efficiently complete all the sub-tasks of the assigned task. A group is able to efficiently complete a task only if it has at least one of its members have the expertise needed for each of the sub-tasks. Otherwise, the assigned group, G, asks for help from another group for that task. We assume that each agent is knowledgeable of the expertise of all other agents. This enables the group G, to request help from a group of agents, H, with all the necessary expertise.

The decision to honor the help request depends on the manner in which the members of the selected group H share their opinion about all agents of the asking group. Though the notion of a "contract" is usually associated with monetary transactions, we use the term "contract" to designate the situation when one group helps another to do the task of the latter. The task, in this case, is contracted to the helping group. When a group helps another group, the members of the helping group incur a cost by which they increase their balances with the members of the helped group. The members of the helped group, on the other hand, save a cost and hence, reduce their balances with those of the helping group. The "opinion" that an agent have about another agent is the balance it has with the latter.

In our work, an individual agent can join other agents to form a group and the group acts as a unit where help-giving decisions are based on the opinion of the constituent agent members (more details of alternative group formation mechanisms are presented in Section 2). The motivation behind this conceptualization is to study whether the peculiarities of inter-agent interaction strategies reflect in inter-group interactions where the decision of honoring a help request depends on the shared opinion of the group members [1]. We assume that agents are not statically assigned to groups. Rather, in each time period, new groupings of agents are formed from agents with complementary expertise. So, the same agent can form groups with different agents at different times. An agent, however, accumulates and retains its interaction history over time with other agents in the environment. The characteristics of the individual entities of groups and the way the opinion of different individuals are shared dictate the number of successful contracts and therefore, the performance of individual agents.

We describe two different agent behaviors in this paper. The notion of a selfish agent that was used in [10] is modified in this paper. We introduce *lazy liars* who, while in a group of agents assessing a help request from the asking group, lie about the balances they have with

[1]A parallel to this research is seen in the real world, where organizations pair up and jointly decide to exploit dynamic fleeting opportunities in modern economic markets [8].

the agents in the asking group, thereby reducing the chance of giving help. They are "lazy" because they do not prefer to be in a group that helps another group by taking a contract from the latter. If a contract is accepted, then all agents in the helping group, including the lazy liars, have to perform subtasks that define a contract (details of a contract is given in Section 2). The *reciprocatives* in this case, are those that reveal truthful opinion about the reputation of agents in the asking group.

We use two different strategies to use the collective opinion of the members of the asked group, viz., *average* and *worst*. We present preliminary results to show that such strategies are able to curb the disruptive effects of the lazy liars. Our simulation based study reveals that the reciprocative agents, even in groups of a majority of lazy liars, are able to maintain lower costs on an average. This shows the effectiveness of the probabilistic reciprocity strategy, even in heterogeneous groups, for producing dominant performance. We also show that the *average* strategy for help determination is able to generate more contracts than the *worst* strategy with a non-trivial percentage of lazy liars in the population.

2 Simulation framework

We consider an information processing domain where a set of A agents are considered. Each agent has one of k different expertise. A set of T tasks are generated. Each task t has m ($m \leq k$) sub-tasks, each of which requires one distinct expertise to be completed. Having an expertise in a task type x implies tasks of type x can be completed incurring less cost compared to tasks of other types. We have used two metrics to compute the task costs, *time* to complete and *quality* of performance. An expert in task type x performs tasks of type x with less time and high quality and performs all other task types with higher time and lower quality. Task cost is defined as the ratio of time to quality, hence, an agent incurs low cost in task types for which it is an expert and high cost for all other task types. m is chosen randomly between 1 and k and then, m distinct expertise are assigned to the m subtasks.

For each task, a group G is formed from the set of all agents and is assigned the task. A group is able to complete an assigned task only if it has at least one agent with the required expertise for every sub-task. If G does not have all the necessary expertise to process the assigned task it approaches another group H for carrying out the task on its behalf. In this paper, we experiment with a random (see Section 2.2.1) group selection and a experience-based group formation (see Section 2.2.2) mechanism.

The group H may or may not honor the request for help from group G. For each member of G the members of H provide their *opinion* which is averaged over all members of H. The opinion that agent i gives about agent j is the balance that the agent i has with j, $balance_{i,j}$ [2]. The opinion that a reciprocative reveals about another agent is the true balance it has with the other agent. The opinion $OP_{i,j}$ that a lazy lier agent i reveals about another agent j is defined as follows.

$$OP_{i,j} = \begin{cases} -\gamma * balance_{i,j}, & \text{if } balance_{i,j} < 0 \\ balance_{i,j}, & \text{otherwise} \end{cases} \tag{1}$$

[2] $balance_{i,j}$ is negative if i has received more help from j than it has given to j. When i helps j, $balance_{i,j}$ is increased by the cost that i incurs, whereas $balance_{j,i}$ is reduced by the cost that j would have incurred by doing the task on its own.

This set of equations show that the lazy liars express false opinion about those from whom they had earned more help in the past by increasing the balance γ times. Increasing the balance reduces the probability with which the help request of G is honored. The opinions of all agents in H are combined using two strategies that are described in Section 2.1.

The decision to agree or refuse the help request of G is probabilistic and is based on the following equation. The probability that group H will help group G to do task t is given by

$$Pr(H,G,t) = \frac{1}{1+\exp^{\frac{C_h^t + B_{HG} - \beta}{\tau}}} \tag{2}$$

where C_h^t is the cost to perform task t by group H, which is the sum of the costs of the individual agents in H; B_{HG} is the net *balance* that group H has with G, as computed by one of the strategies. β and τ are the only two constant parameters used in the function, where β is used to set the cost a group is ready to incur to help an unknown group with the hope of initiating a long-term cooperative relationship and τ is used to set the shape of the probability curve. This is a sigmoidal probability function where the probability of helping increases as the balance decreases and is more for less costly tasks.

If H helps G, group members of each of G and H update their balances with the members of the other group. The balances of each agent in H is increased with each of the members in G by $\frac{c_i^{t_i}}{|G|}$, where $c_i^{t_i}$ is the cost incurred by agent i in H to do the i^{th} subtask of t (t_i) and, $|G|$ is the cardinality of G. So it is assumed that i incurred equal cost for each of the agents in G. Also, since there are exactly that many agents in H as there are subtasks in t, each agent i in H does exactly one subtask (i^{th} subtask) of t in which it is an expert.

Now we describe the procedure by which the members of G update their balances with those of H. First the cost savings per member of G is calculated as follows:

$cost \leftarrow 0$
for each subtask $s \in t$ **do**
  $x \leftarrow$ number of agents $\in G$ with expertise required for s
  if $x \geq 1$ **then**
    $cost \leftarrow cost + \frac{\text{expert's cost}}{x}$
  else
    $cost \leftarrow cost + \frac{\text{non-expert's cost}}{|G|}$
  end if
end for

In the above procedure we assume that for a subtask of t for which there are one or more experts in G, that subtask cost is equally shared by the corresponding experts. For a subtask in t for which there are no experts, the cost is divided equally among all members of G (other, perhaps fairer, divisions are possible, but this choice is not going to fundamentally affect the main results of this paper). The balances of each member of G with that of each member of H is decreased by $\frac{cost}{|H|}$, assuming the cost saving was due to all the agents in the helping group.

2.1 Selection Strategies

We have used two simple strategies to combine the individual opinions of the members of H to form the group balance for deciding to help group G. The opinions of the members of H are the balances that they have with the agents in G. The combined opinion of the group H about an agent $j \in G$ is given by,

$$O_{H,j} = \frac{\Sigma_{i \in H} OP_{i,j}}{|H|} \tag{3}$$

Average Opinion Strategy(AOS): In this strategy the balance of H about G is computed as the average of the overall opinions $O_{G,j}$ about each member j of G, as given above. Hence,

$$B_{HG} = \frac{\Sigma_{j \in G} O_{H,j}}{|G|} \tag{4}$$

Worst Opinion Strategy(WOS): In this strategy H tries to punish any selfish agent hiding in G and makes the decision on the basis of the worst or maximum opinion about any member in H. Hence,

$$B_{HG} = max_{j \in G} O_{H,j} \tag{5}$$

2.2 Group formation

We now describe the two distinct group formation methods we have experimented with. The random group formation method is a more centralized approach where the group G is picked by a central arbiter. The group H is also formed by random picks if G cannot process the task efficiently. The other group formation mechanism provides more autonomy to the agents. In this approach, an agent can approach other agents to join the group G, and the requesting agent can accept or refuse that offer. Similarly, the members of group G, jointly decide which other agents should be invited to form the group H. Thus experience plays a major role in determining who is selected for entering a group.

2.2.1 Random group formation

In the random group formation approach, the group that the task is assigned to is selected centrally by randomly selecting m agents among the population. This forms the asking group G. Since the m agents are selected randomly, it is not guaranteed that G will always be able to complete the task assigned to it.

If G is not able to complete its own task, it requires help from another group H. We assume that the members of G have perfect knowledge about the expertise of all agents, using which they select m agents where each agent has a distinct expertise and exactly those required to complete the task t.

2.2.2 Group formation on expected cooperation

In this approach, we assume that an agent, L, takes on the lead role in trying to form group G. L approaches $s-1$ agents $AP_1, AP_2, .., AP_{s-1}$ with which it has the most favorable (negative) balances. Recall that s is the number of subtasks in task t. The idea behind this strategy is to choose those agents in group G from which L has got more help in the past. L expects that if these agents remain in the group G the likelihood of getting help from other groups, if required, will be increased. Each of the invited agent AP_i probabilistically decide to join L using 2. If any one of AP_i declines L's request, L fails to form G. Then a new L is selected and the process is repeated until G is formed or all the agents in the population fail to form G. If G can not be formed by any agent being the leader, then the particular task contracting fails.

$groupFormed \leftarrow false$
repeat
  randomly select $L \in A$
  Sort the agents $A - \{L\}$ in ascending order of $balance_{l,k} \forall k \in \{A - \{L\}\}$
  select first s -1 from sorted list as $AP_1, AP_2, \ldots, AP_{s-1}$ and invite them to join group G
  if All agents $AP_1, AP_2, \ldots, AP_{s-1}$ accept invitation to join G **then**
    $groupFormed \leftarrow true$
  end if
until *groupFormed* or all agents have been considered for leading

If G does not have an expert for any subtask s, members of G seek help from another group H. The selection algorithm for members of H is also experience-based. For each subtask $s \in t$ we consider all the agents who are expert in that type of task and not in G. Members of G probabilistically select one agent out of these experts using their normalized combined balance with those agents. We present the algorithm below:

for each subtask $s \in t$ **do**
  $Hs \leftarrow$ set of agents expert in this subtask and not in G
  for each member $hs \in Hs$ **do**
    $combinedbalance[hs] \leftarrow 0$
    for each member $g \in G$ **do**
      $combinedbalance[hs] \leftarrow combinedbalance[hs] + balance[g][hs]$
    end for
    select one agent from Hs with probability proportional to $combindedbalance[hs]$
  end for
end for

3 Results

In this section we report results from our simulation of the information processing domain. We evaluate the global performance of the system in terms of the total contracts made, as well as average performance of exploitative and reciprocative agents in terms of their average savings. In the simulations, the population of lazy liars is increased from 10 to 90% in steps of 10%. We have used different values of agents and total tasks generated and the results are averaged over 10 runs.

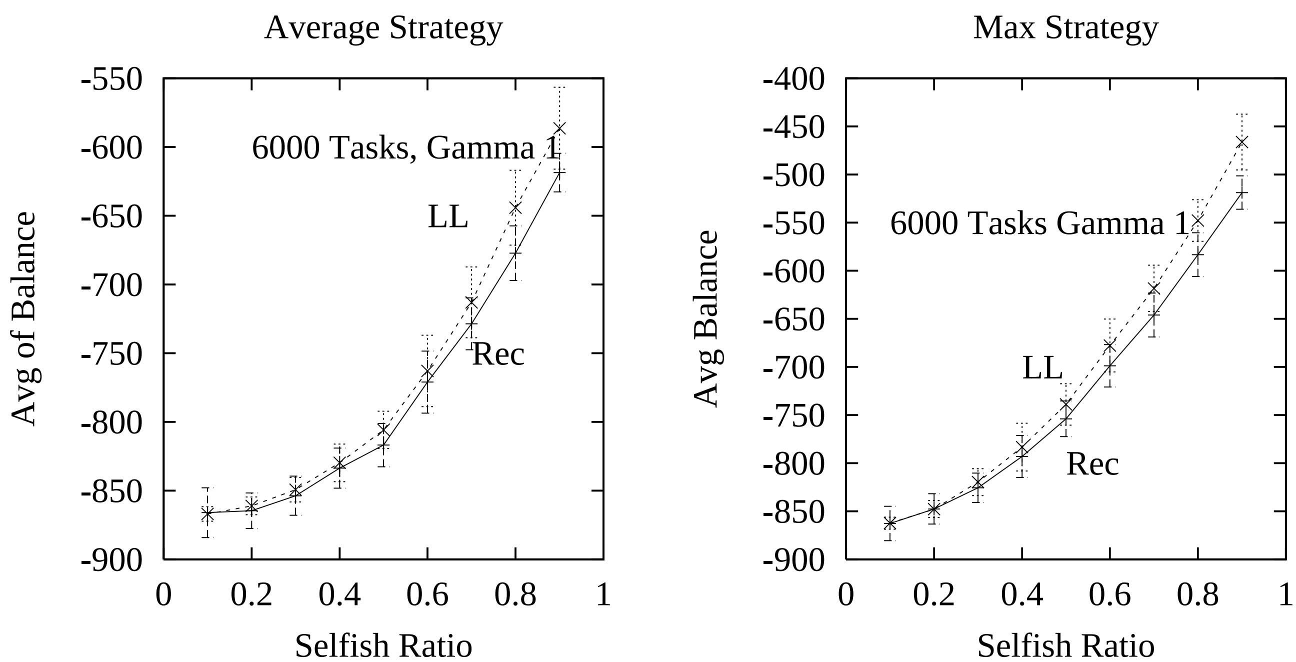

Figure 1: Average balance with average (left) and worst (right) selection strategies ($\gamma = 1$).

3.1 Results with random group formation

In this subsection, we present results when the groups are formed randomly as described in Section 2.2.1. We experimented with a total of 100 agents, where each agent was assigned 6000 tasks with a γ value of 1. Figure 1 shows the variation in the average balance of reciprocative and lazy lier agents for different values of lazy lier percentage in the population. We present plots when using the AOS and the WOS selection strategies. From Figure 1 we see that with an increase in the lazy lier percentage, the average balance of both lazy lier, LL, and the reciprocative agents increase. This means that they are able to save less cost with more lazy liars in the population, which corresponds to a degraded individual performance. For the same percentage of LL agents, and for reasons cited before, the worst strategy leads to poorer performance of both agent types, i.e., generates higher balances. The balance of the reciprocative agents remain similar to the LL agents for most selfish ratio. Because the lying is mild, i.e., $\gamma = 1$, the selfish do not severely impact group decisions when there is a majority of reciprocative in the group. The reciprocative balances are somewhat better (more negative) at very high selfish ratios. This is because the mild lying finally starts influencing group decisions with a majority of selfish agents. The average selection strategy is more beneficial in this context because the reciprocatives earn better balances for the same percentage of LL agents than the worst opinion strategy.

We were interested in investigating whether there exists situations where the LL agents perform better than the reciprocatives. Since the LL agents reveal false reputations about others, we used a more severe lying factor of $\gamma = 40$ in the next set of experiments. As before, we used 100 agents and 6000 tasks per agent and the results are shown in Figure 1.

From Figure 2 we observe that the severe lying has a pronounced negative effect on the

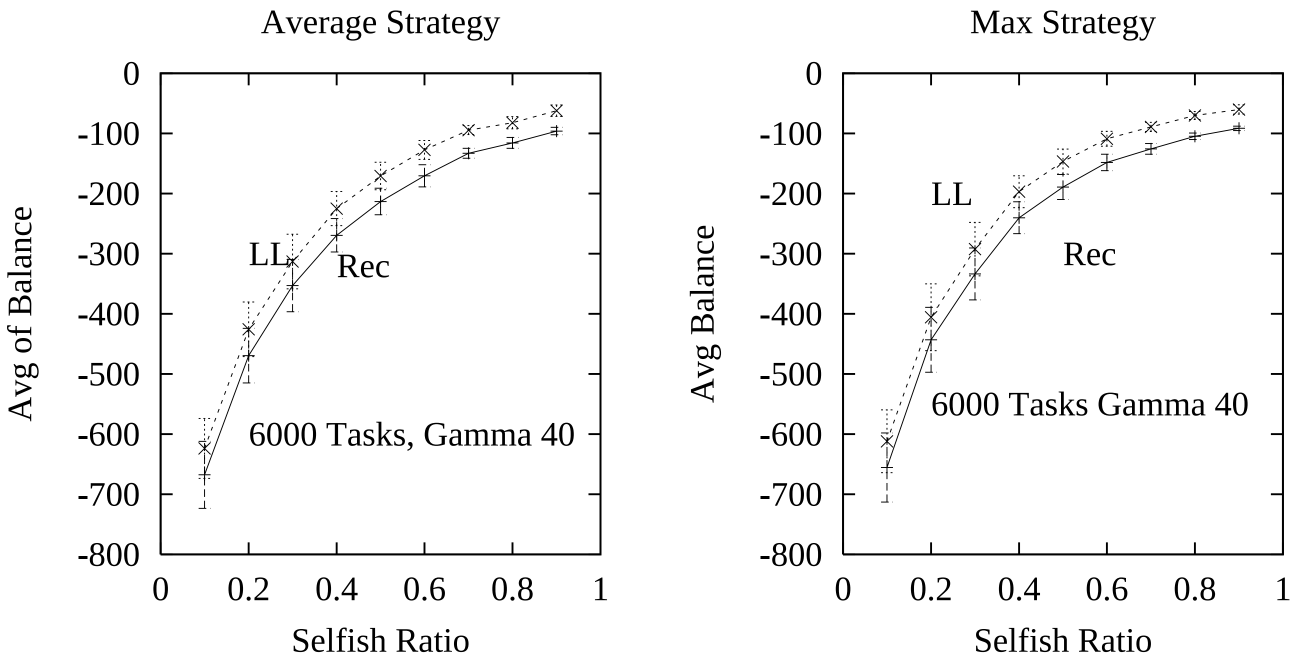

Figure 2: Average balance with average (left) and worst (right) selection strategies ($\gamma = 40$).

group. Thus the shape of the curve changes from the previous experiment set, as even with relatively smally percentage of selfish agents in the population, individual balances worsen quickly. We notice, however, that the reciprocatives perform better than the LL agents by earning more balances.

Results from the above set of experiments give us insight about different situations when the reciprocatives are viable and can earn better balances than the lazy liars. It appears that without a more selective help-seeking or help-giving group formation mechanism, reciprocatives will not be able to significantly outperform lying agents.

3.2 Cooperative group formation

In the following set of experiments we have used group formation based on expectation of cooperation as described in Section 2.2.2. As in the previous set of experiments, we have used 100 agents with 6000 tasks per agent and have varied γ in equation 1 to determine the effect of bad mouthing and exaggeration by the selfish agents on cooperation decisions in a group.

From Figure 3 we can see that for both AOS and WOS decision making strategies the selfish agents are outperformed by reciprocative agents. Thus, and in contrast to Figure 1 even for relatively mild lying agents, the reciprocatives are able to clearly outperform selfish agents. This is made possible by the fact that reciprocative agents both shun joining groups when invited by selfish agents and avoid inviting selfish agents into groups that they are forming. It is interesting to note that the dominance of reciprocatives holds even for high LL percentage in the population.

When γ is increased to 40, i.e., selfish liars severely undermine the reputation of recip-

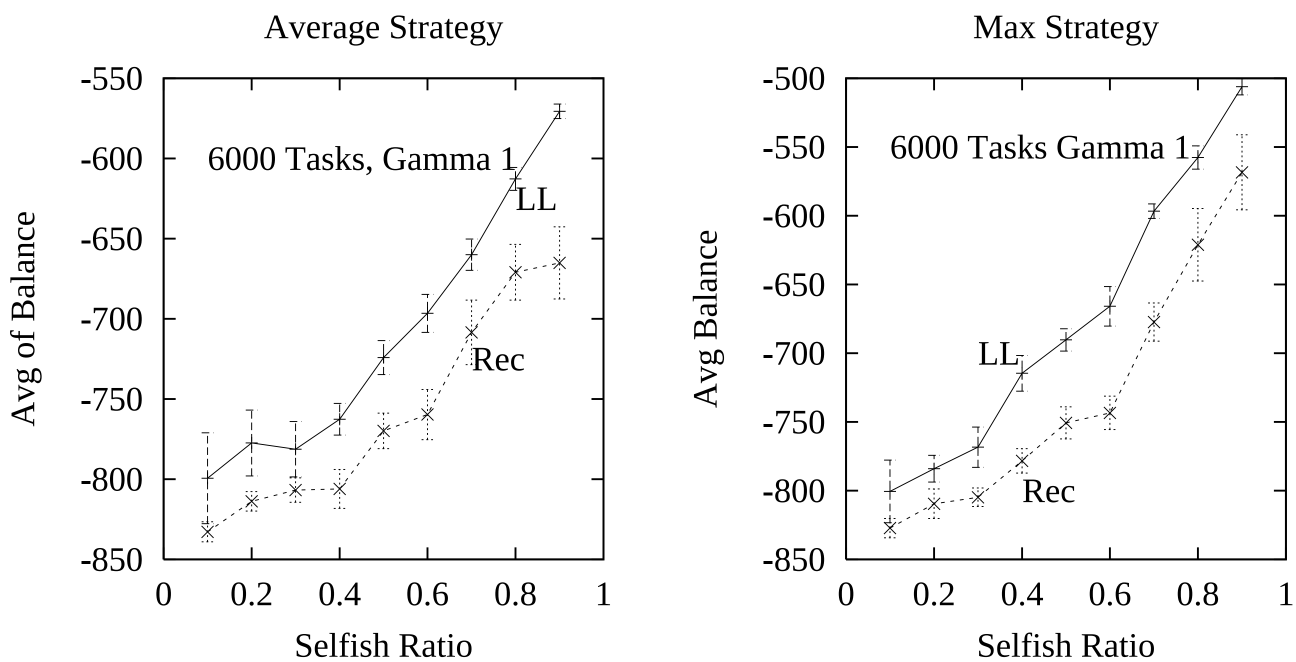

Figure 3: Average balance with average (top) and worst (bot) selection strategies for cooperative group formation ($\gamma = 1$).

rocative agents, the performance difference is accentuated for lower selfish percentages (see Figure 4). This is because the selfish prevents most groups that it becomes a part of from helping. Since it does not help but receives some initial help, they are not invited to the group G later on and hence is not able to save any costs. We have noticed a remarkable difference in the number of groups that include selfish agents, as compared to reciprocative agents, and are still able to contract out their task to other groups. With increase in percentage of selfish population, few groups are able to contract out their tasks, and hence the performance of both the selfish and reciprocative agents deteriorate. At very high percentage of selfish population therefore, few contracts are made, and there is no noticeable difference between the performance of different agent types. The balances of the agents are slightly worse with the WOS selection strategy compared to the AOS selection strategy. This is because of more contract refusals by the more stringent WOS selection criterion.

This set of experiments convincingly demonstrates the effectiveness of the reciprocative strategy in countering the spread of false reputations by the lying selfish agents when selective group formation is enabled.

4 Conclusions and future work

We have investigated the impact of two decision procedures for task exchange between groups, where these decisions are based on the opinions of group members about the individuals in the other group. In particular, we are interested in the number of task exchanges and the relative performance, in terms of cost savings by contracting out tasks, of reciprocative

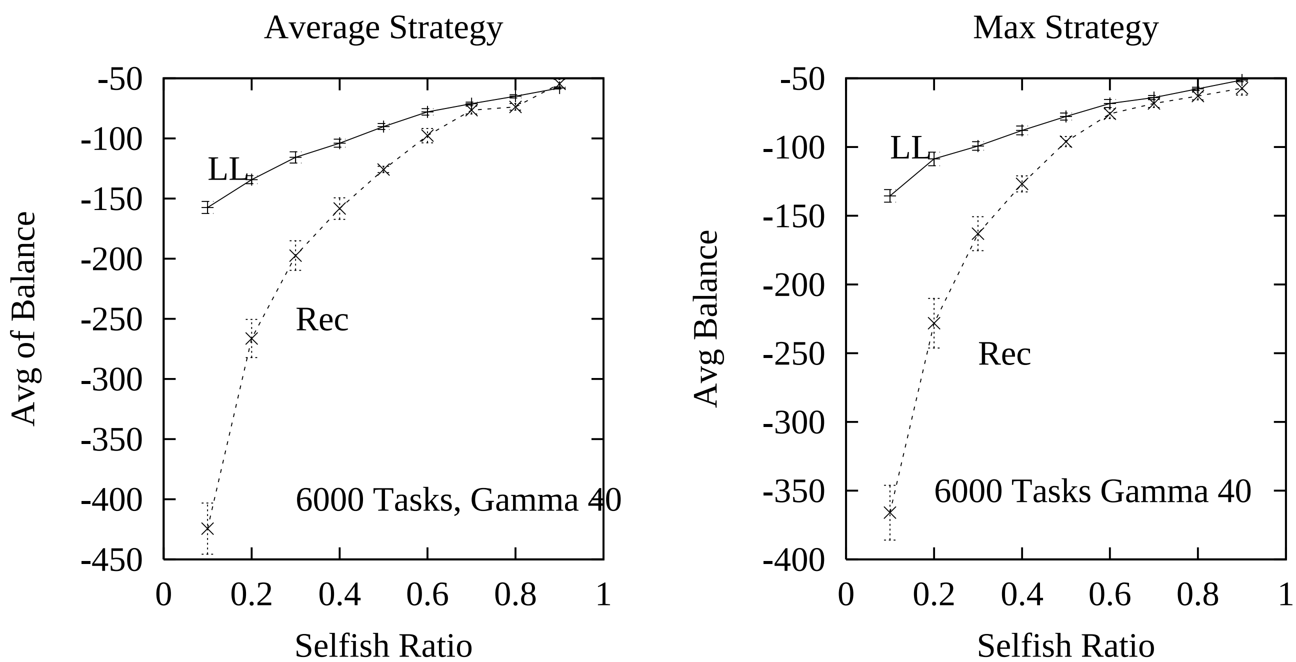

Figure 4: Average balance with average (top) and worst (bot) selection strategies for cooperative group formation ($\gamma = 40$).

and lazy liars as the percentage of the latter is varied in the population.

With random group selection, the selfish liars are seen to perform almost as well as reciprocative agents. Actually, mild lying is not enough to overcome the opinions of truthful reciprocative agents, and hence even selfish agents end up being part of helping groups. In essence, in spite of their lying, they end up having to help other agents. When they *up the ante* by exaggerating their inverted opinions, they bring down the performance of everyone even if they under-perform reciprocative agents.

The reciprocative agents, however, thrive when we allow focused group selection. In this mode, an agent can invite other agents to form a group the group can then decide to form another group to ask for help. Agents being asked to join a group can accept or reject that request based on their own utility perceptions. We find that within this, more realistic, framework, reciprocative agents consistently outperform selfish liars, and even at high percentage of the latter in the population. In particular, when the liars exaggerate their negative opinion, they perform relatively poorly if they are not a large majority in the population.

We plan to experiment with a variable γ by which a selfish agent can adjust their lying pattern to maximally benefit from its environment. It would be interesting to see if, for example, selfish liars can vary their truthfulness to exploit groups differentially.

On the other hand, reciprocative agents can probably use opinion differences between its own rating of another agent in group G and that expressed by other agents in group H to recognize possibly lying agents. This will allow the reciprocatives to detect and avoid the selfish agents quickly, and hence bolster their own performance.

References

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

[2] Anish Biswas, Sandip Sen, and Sandip Debnath. Limiting deception in groups of social agents. *Applied Artificial Intelligence, Special Issue on Deception, Fraud, and Trust in Agent Societies*, 14(8):785–797, 2000.

[3] Jeffrey M. Bradshaw. *Software Agents*. AAAI Press/The MIT Press, Menlo Park, CA, 1997.

[4] Communications of the ACM, march 1999, volume 42, number 3, 1999. Special Issue on Multiagent Systems on the Net and Agents in E-commerce.

[5] Cristiano Castelfranchi, Rosaria Conte, and Mario Paolucci. Normative reputation and the costs of compliance. *Journal of Artificial Societies and Social Simulation*, 1(3), 1998.

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

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

[8] Rebecca A. Henry, Jill Kmet, Erica Desrosiers, and Alyson Landa. Examining the impact of interpersonal cohesiveness on group accuracy interventions: The importance of matching versus buffering. *Journal of Artificial Societies and Social Simulation*, 87(1):25–43, January 2002.

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

[10] Sandip Sen, Anish Biswas, and Sandip 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.

[11] Evelien P.H. Zeggelink, Henk de Vos, and Donald Elsas. Reciprocal altruism and group formation: The degree of segmentation of reciprocal altruists who prefer 'old-help ing-partners'. *Journal of Artificial Societies and Social Simulation*, 3(3):212–15, 2000.