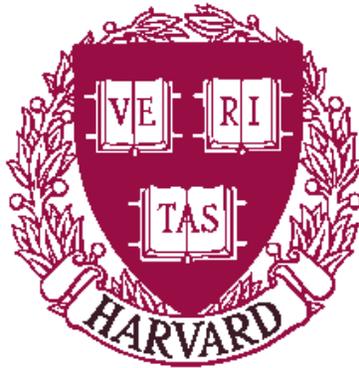


Adapting to Agents' Personalities in Negotiation

Ya'akov Gal
Shavit Talman
Meirav Hadad
and
Sarit Kraus

TR-04-05



Computer Science Group
Harvard University
Cambridge, Massachusetts

Adapting to Agents' Personalities in Negotiation

Ya'akov Gal
Division of Applied Sciences
Harvard University
Cambridge MA 02138

Shavit Talman, Meirav Hadad and Sarit Kraus
Department of Computer Science
Bar Ilan University
Ramat Gan, Israel

Abstract

To establish cooperative relationships, agents must be willing to engage in helpful behavior and to keep their commitments with agents who reciprocate this behavior. However, in uncertain and dynamic environments, it is difficult to identify the degree of helpfulness of other agents. This paper approaches this problem by characterizing agents' helpfulness in terms of cooperation and reliability. An agent chooses an action based on other agents' helpfulness as well as the dependency relationship between the agent and others. This model was evaluated in a negotiation game in which players needed to exchange resources to reach their goals, but did not have information about each other's resources. Results showed that the model allowed agents to identify and to adapt to others' varying degree of helpfulness, even while they are constantly changing their strategy. Moreover, agents who vary their cooperativeness and reliability depending on those traits of others, can outperform agents who do not, as well as increase the social welfare of the group.

1. Introduction

When making decisions, self-interested agents need to reason about the behavior of other agents as well as the effects of their own actions. In particular, when agents depend on each other to achieve their individual goals, agents' success depends on their ability to cooperate with others, i.e. to perform actions that mutually benefit each other. In many of these settings cooperation arises as a result of a process of negotiation, in which agreements are reached and then resources are exchanged. However, deceitful agents can take advantage of agents that are helpful to them and others. As systems become truly open, and consist of multiple agents with no central control of design or behavior, agents must be able to identify those who are cooperative and avoid those who are exploiters.

One option is to use normative models of decision making (e.g. game theory) as a guide for agents' behavior [5]. These prescribe optimal strategies for all agents which take into account the effects of their decisions on each other, and guarantee that no agent is exploited by another. However, in many cases, players who follow these strategies do not help each other, and cooperative behavior, leading to higher utilities for all agents, does not occur. In addition, these models assume all agents deliberate in the same fashion and completely adhere to their assigned strategies, a difficult assumption to make in open environments. Lastly, the complexity of analysis becomes inherently difficult as the number of agents grow or when agents are uncertain about each other's preferences. In short, normative models cannot cope with the diversity that characterizes systems that include many types of agents.

In behavioral economics and social psychology, personality models have been used to explain and predict human behavior across different environments and contexts. In particular, it has been shown that people's personality affects their decision-making strategies in negotiation (e.g., in the ultimatum game) and their level of cooperation [1]. This paper presents a model of negotiation which represents and reasons about agents' personalities in uncertain environments. The model characterizes the personality of others along two dimensions: cooperation (the tendency to propose mutually beneficial exchanges of resources) and reliability (the tendency to fulfill commitments). By including these measures inside agents' utility functions, they can reciprocate helpful behavior rather than maximize their immediate reward. They are also flexible in their ability to adapt to others who change their own behavior. Therefore, they are able to stay clear of unhelpful agents, and realize beneficial exchanges with helpful agents. Agents alternated between updating a model of others and using the model to come up with the best strategy to play.

We tested our model in an open system which included agents that were created by the experiments as well as by other people, who were not required to design their agents in any specific manner. Our experiments varied the com-

plexity of the environment as well as the number of players and the dependency relationship between players. Results showed that agents who reason about the personality were able to identify the cooperation and reliability tendencies of others; they negotiated and reciprocated with helpful agents, while staying clear of agents that were deceitful or unhelpful. In particular, agents who vary their helpfulness depending on the personality tendencies of others, can outperform agents who do not, and include the social welfare of all agents in the system.

Our work extends previous models of social reasoning in negotiation. Hogg and Jennings [4] proposed a model in which agents' utilities were a weighted summation of each other's expected outcomes. By learning these weights from observations, agents could change their measure of helpfulness over time. When all agents in the system were adaptive, high exploration rates led agents to seek out new negotiation opportunities and increased the overall social welfare of the group. Sen and Dutta [6] investigated the effect of agents' helpfulness, and testimonies of others' helpfulness, on the performance of the group. Agents agreed or declined others' requests according to a probability distribution which depended on their past experiences as well as the cost of helping. Agents who weighted the testimony of others based on their reputation were more likely to succeed than those who did not consider reputation of others. They were also able to avoid deceitful agents and to cooperate with helpful agents. Lastly, Zhang et. al [7] explored the tradeoff between selfishness and benevolence in environments where agents are uncertain about the helpful nature of others in the system. They showed that without this knowledge, realizing every opportunity for cooperation is impossible, and that selfish agents do better than helpful agents as the rate of uncertainty in the system grows.

All of these models allow agents to adapt their behavior towards others by changing their measure of helpfulness over time. However, they explicitly include the potential benefit for each agent originating from possible actions of others, which could be computed since all agents in the group were designed by the experimenters. In systems that are truly open, there is no knowledge about any aspect of other agents' design mechanisms, including how they benefit from potential actions of other agents. Our model does not assume or control any shape or form of other's utility function and as we show here, is thus suitable for such systems. In addition, these works examined the effects of agents' helpfulness on the overall performance of the system as well as the level of cooperation between agents. Our focus was to construct an agent which could outperform other agents in open environments, regardless of their helpfulness measure. Lastly, the strategy space of the domains used in these works was limited, as were the number of moves in the games. In the domain we used in this work

there were thousands of potential actions at each phase of the game and the length of each game depended on agents' performance. There was uncertainty not only over agents' resources, but also their measure of helpfulness. In addition, the model was evaluated in repeated game settings which added to its complexity.

2. The Colored Trails Formalism

Our study used the Colored Trails (CT) game designed by Grosz and Kraus [3]. CT is a framework for investigating decision-making processes of agents in contexts in which their outcome depends on each other's actions. By setting the game parameters, it allows to control varying environmental features such as task complexity, the resources available to agents and their capabilities, and the dependency relationship between agents. CT provides a clear analogue between the properties of the game and real world task settings and thus it generalizes to many domains.

CT is played on an $N \times M$ board of colored squares with a set of tiles in colors chosen from the same palette as the squares. One square is designated as the "goal square" and each player has a piece on the board, initially located in one of the non-goal squares. Each player also has a set of colored chips. To move a piece into an adjacent square a player must turn in a chip of the same color as the square. Chips may be exchanged by the players, and the conditions of exchange may be varied to model different decision-making situations.

A player's performance in CT is determined by a scoring function, which is computed when the player is declared "out-of-game". This function may depend on many factors, such as the player's distance from the goal-square, the number of moves made, and the number of tiles the player possesses at the end of the game. In addition, a player's performance in the game may be made to depend on the performance of other players, by including the score of other players in its own scoring function.

For our study, we used a version of CT in which two or four players played on boards of different sizes. Each player had knowledge of the scoring function and full view of the board but cannot see the other player's chips.

The game protocol comprised two phases, a communication phase and a movement phase. During the communication phase, new exchanges could be proposed, pending proposals could be replied to, and chips could be sent from player to player. Agreements reached during the communication phase were not enforced by the game controller, allowing agents to deceive each other. During the movement phase, the game controller automatically advanced each player one square closer along the shortest path to the goal given its resources.

A player was declared “out-of-game” if it reached the goal state or if it stayed dormant for 3 moves, at which point its score, denoted $score_i$, was computed. Each player’s outcome depended solely on its own performance. The scoring rule used was a multi-attribute linear function which incorporated three factors: (1) whether the player reached the goal square; (2) the distance of the player from the goal square; (3) the number of chips the player possessed at the end of the game.

3. Model Construction

We wanted our model to be able to generalize to environments which varied the number of players, the size of the board-game, and the dependency relationship between players. In addition, we wanted the model to be able to play well in systems characterized by uncertainty over others’ resources and utility functions. To this end, our model explicitly represented the helpfulness tendencies of multiple CT players.

We described agents’ personality traits along two dimensions with range $[0, 1)$.

- Cooperation (c) - measured agents’ willingness to share resources with others in the game through initiating and agreeing to proposals.
- Reliability (r) - measured agents’ willingness to keep their commitments in the game through delivering the chips they had agreed to.

We regarded agents that cooperated more than 50% of the time as highly cooperative, and agents that reneged on their commitments less than 20% of the time as highly reliable. More formally, we discretized parameter values into three types.

- low-cooperation: $[0, 0.3)$; low-reliability: $[0, 0.4)$
- medium-cooperation: $[0.3, 0.6)$;
medium-reliability: $[0.4, 0.8)$
- high-cooperation: $[0.6, 1.0)$; high-reliability: $[0.8, 1.0)$

We refer to agents’ personality traits by their cooperation and reliability types. For example, an agent whose behavior exhibits low cooperation and medium reliability will be referred to as a low-cooperation, medium-reliability (or LM) type agent.

According to our model, agent i ’s expected utility of taking action a directed at agent j depended on the following features.

- The personality of agent i , denoted P_i .
- Agent i ’s estimate of the personality of agent j , denoted P_j .
- The expected value of taking action a given the state of the game s , denoted $EV_i(a|s)$.

- The expected cost of future ramifications of taking action a , denoted $EC_i(a)$.

Computing the terms $EV_i(a|s)$ and $EC_i(a)$ is difficult, especially when the chips of the opponent are unknown. They require us to estimate the likelihood of agent i reaching the goal in future phases of the game. In addition, they require modeling the reasoning of opponent j who might itself be modeling agent i . We approximated the probability $P(RG|s)$ of reaching the goal, while at state s , simply as $1 - \frac{\#cn}{M+N}$ where $\#cn$ is the number of chips the agent lacks to get the goal at state s and M and N are the game board dimensions. We can then compute

$$EV_i(a|s) = P(RG|s) \cdot RGwt + score_i - e^{\#np}$$

where $RGwt$ is a constant representing a bonus for reaching the goal, $score_i$ is the score for player i at state s and the term $e^{\#np}$ punishes the player for remaining idle, where $\#np$ equals the number of idle turns the player has had in the game.

To compute the expected ramification cost EC_i , we assumed that selfish actions of a player were punished by others, and that considerate actions were rewarded by others. For any proposal made by agent i towards agent j , we defined a *selfish exchange* of i to be any exchange which is more advantageous to i than to agent j . We defined a *considerate exchange* of agent i to be any exchange that is more advantageous to agent j than to i . We defined a *feasible exchange* of i to be any exchange for which agent i possesses the chips to complete. All exchanges were assumed to be beneficial to the recipient agent.

We estimated $EC_i(a)$ to equal some constant integer t when a is (1) a considerate exchange of agent i ; (2) an agreement of i to a feasible exchange; (3) a transfer of chips to fulfill a promise of exchange. We estimated $EC_i(a)$ to equal $-t$ when a is (1) a selfish exchange for agent i ; (2) an agreement of i to an infeasible exchange; (3) reneging on a promise to send chips.

We constructed a multi-attribute linear utility function associated with weights w_1, \dots, w_4 .

$$u_i(a, j, s) = w_1 \cdot (P_j.c + P_i.c) + w_2 \cdot (P_j.r + P_i.r) + w_3 \cdot EV_i(a|s) + w_4 \cdot EC_i(a)$$

Here, we used object-oriented notation to denote measures of personality traits. For example, $P_j.c$ refers to agent i ’s estimate of the cooperation measure for agent j . The value of the weights depended on agents’ personality traits. Low-cooperation and low-reliability type personalities assigned a higher weight to the expected value EV_i , leading them to adopt behavior that was unhelpful. For example, low-cooperation type agents only proposed selfish exchanges and low reliability-type agents never fulfilled

their commitments to send chips. For medium- and high-cooperation (or reliability) type agents, the weights were tuned empirically.

Given that we have a model that represents the personalities of each opponent, the next step is to show how an agent that used this model behaved in the game. For each opponent j , an optimal action a_j^* maximizes the utility of agent i at state s .

$$a_j^* \in \operatorname{argmax}_a u_i(a, j, s)$$

Note that the equation uses set membership rather than equality to indicate that there may be more than one action that maximize the agent’s utility. An optimal action set for i , denoted A_s is the set of all optimal actions of agent i at state s (It will include at least one action for every opponent). We say two actions a_1 and a_2 contradict if performing a_1 would require at least one chip that is needed to perform a_2 . For each subset of A_s in which no two actions contradict, agent i summed up its utility for performing every action. Agent i chose the subset whose set of actions together yielded maximum utility, and performed every action in that subset at the onset of each communication phase of the game.

Lastly, we show how we updated the personality model after each communication phase k . We computed the reliability measure of agent j at phase k , denoted $P_{j.r}^k$, to be the fraction of times that agent j met its commitments at phase k . The cooperation measure of agent j at phase k , denoted $P_{j.c}^k$, is the fraction of times that agent j proposed exchanges in which it offered chips to others. To update $P_{j.r}$, we computed $\frac{1}{k} \sum_{l=1}^k \delta^{k-l} P_{j.r}^l$, where δ is a discount factor. We updated $P_{j.c}$ in a similar manner.

This model prescribes distinct behaviors for different personality traits. Low-reliability type agents never keep their commitments to others. Medium-reliability type agents are (1) more likely to keep their commitments to medium- and high-reliability type agents than they are to low-reliability type agents; (2) less likely to keep their commitments to medium- and high-cooperation type agents than they are to low-cooperation type agents. High-reliability personalities always keep their commitments, regardless of the personality type of the other.

The effect of an agent’s cooperation level on its behavior, according to the model, is more complex – it varies according to the dependency relationship between players. We say that player i is task-dependent on player j , if player i lacks the chip it needs to reach its goal and depends on some player j , which possesses these chips, to supply them. We say that a player is task-independent if it possesses all the chips it needs to get to the goal. Table 1 presents the behavior that was associated with agents’ level of cooperation given their task dependency type. For example, when it was independent of others, a high-cooperation type agent was benevolent: it accepted any exchange, provided that it

was beneficial to the agent. Medium- and high-cooperation type agents proposed exchanges whose benefit to the recipient agent was correlated with their estimate of its level of cooperation. The higher this estimate, the higher the benefit of the exchange. This was denoted as “cooperation-dependent” behavior within the table. As shown in the table, when agents were task dependent on each other, medium- and high-cooperation type agents were more helpful to cooperative agents than others.

4. Experimental Design and Analysis

We used several classes of agents in our study. Multiple-Personality (MP) and Single-Personality (SP) agents were designed by the experimenters, and explicitly model the cooperation and reliability traits $P_{j.c}$ and $P_{j.r}$ of each opponent j . The personality P_i of a SP agent includes constant cooperation and reliability measures. However, a MP agent adopts different measures of cooperation and reliability for each personality type of its opponents based on a matching scheme we describe later. Both MP and SP agents were adaptive: they changed their behavior as a function of their estimate of others’ measure of helpfulness, given the history of observations. However, the MP agent varied its cooperativeness and reliability traits by adopting a unique personality type for each opponent.

Another class of agents was Peer-Designed (PD) agents. To create a truly open system, the designers of these agents were graduate-level computer-science students who were not given any explicit instructions regarding their decision-making strategies. In particular, the utility function that guided PD agents in their play was not known.

We characterized all agents as “helpful” and “unhelpful” in a preliminary evaluation: Helpful PD agents were those who mainly engaged in *reciprocal-type* exchanges, in which chips were both received and sent by the PDs; unhelpful PD agents were those agents who engaged in (1) *take-type* exchanges in which players received chips, but did not give out chips, or (2) were idle and did not engage in any negotiation with others. Helpful SP agents exhibited medium- and high-cooperation and reliability personality traits, while unhelpful SPs exhibited low-cooperation and reliability personality traits, as described in Section 3.

We expected helpful agents to be able to realize opportunities for exchange with each other more often than unhelpful agents and to exceed them in performance. We also expected that in some cases, unhelpful agents would be able to take advantage of the vulnerability of helpful agents who allow themselves to be exploited. We envisioned that the flexibility of the MP agent would allow it to identify and reciprocate helpful agents more quickly than others, while staying clear of those who would exploit it. As a result, the MP agent would achieve a higher score in the game than

Cooperation Level	Negotiates with Personalities	Types of Exchanges Accepted	types of Exchanges Proposed
Low	none / all	selfish / selfish	none / selfish
Medium	high-reliability / medium- and high reliability	selfish/ cooperation-dependent	cooperation-dependent / cooperation-dependent
High	all / medium- and high- cooperation and reliability	any beneficial / cooperation-dependent	cooperation-dependent / cooperation-dependent

Table 1. Task Independent/Dependent Behavior by Cooperation Level

personality type	personality matched
LL	LL
LM	LM
LH	LM
MM	LM
MH	MM
HM	HM
HH	MM

Table 2. Matching Table for MP agent by Personality Type (cooperation and reliability measure)

all other agents. Our experiments used a single MP agent, 7 SP agents with personality traits LL, LM, LH, MM, HM and HH, and 10 PD agents, all of which achieved the highest score in a preliminary evaluation.

We took an empirical approach to come up with a matching scheme for the MP agent, which assigned a different personality type for each agent type. We ran a series of games in which SP agents played each other using 30 2-player boards, which varied every possible task-dependency combination between two players. Each SP agent played multiple games against each of the other possible SP types. We matched each SP agent’s personality with the personality of the opponent that resulted in the highest average score for the SP agent. We did this separately for single and repeated CT games, and combined the results of both settings to form a matching table, described in Table 4. Based on its estimates of the personality traits of others, the MP agent matches a personality of its own to each agent.

This matching scheme adopted the right “balance” between helpfulness and selfishness. Low-cooperation agents were matched with low-cooperation personalities to avoid getting taken advantage of. Medium and high-reliability agents were matched with medium-reliability personalities, who kept their commitments as long as their opponent kept theirs. Interestingly, some matchings were not intuitive. For example, high-cooperation agents were matched with medium-reliability personalities; this was enough to get high-cooperation agents to behave optimally from the MP agent’s point of view.

We report on the performance of agents in the system by comparing their scores and behavior across environments which varied the dependency relationships between players as well as the number of players. All results are statistically

significant in the 95% confidence interval range unless indicated otherwise.

4.1. Repeated Game Settings

We evaluated the MP agent by playing a series of repeated games with the other agents in the systems. We allowed agents to update their model of others from game to game. We expected the MP agent to score higher than both helpful and unhelpful agents in each of the games, and that the rate of increase in score of the MP agent from game to game to be significantly higher than for other agents who increase their score. Also, we expected that helpful agents would score higher than unhelpful agents in each game, and that helpful agents would improve their performance from game to game while unhelpful agents would decrease their performance. Lastly, we expected that when playing against unhelpful agents, the MP agent would score lower than when playing against helpful agents. This is because one of the characteristics of low-cooperation SPs was that they hardly initiated negotiation. Since we expected the MP agent to avoid them altogether, this would prevent both agents from realizing beneficial opportunities for exchanges in the game.

The evaluation used two types of game boards. One in which all players were task dependent on each other (allDep board), and one in which one player was task independent and the other players were task dependent on it (oneSelf) board. The players in each game consisted of a MP agent, a PD agent and two SP agents. In our experiment we executed 5,040 games, played in 1,080 rounds of three consecutive games each. The board games we used in each round alternated between the (oneSelf, allDep, OneSelf) boards and (allDep, oneSelf, allDep) boards. The players in each game included a MP agent, two SP agents, and one of the PD agents. Each group of four players played all possible task dependency roles, to control for any effect brought about by dependency relationships. Table 3 presents the average score for the MP agent when playing against helpful and unhelpful agents across all games. These results sum over the other players in the game.

As expected, the average score achieved by the MP agent was significantly higher than all other agents, regardless of their measure of helpfulness. Also, the MP agent’s score when playing against helpful agents (170.6) was higher than

	MP agent	PD and SP agents
Helpful	170.6	114.8
Unhelpful	142.5	98.2

Table 3. MP agent Score average (3 repeated games)

Exchange Type	Helpful agents	Unhelpful agents
Reciprocal	60%	25%
Idle	20%	39%

Table 4. Exchange type Percentage of MP agent

its score when playing unhelpful agents (142.5). Helpful agents also benefited from cooperating with the MP agent: Their performance was significantly higher than their unhelpful counterparts (114.8 vs. 98.2).

To show that the MP agent established a cooperative relationship with helpful agents, while staying clear of unhelpful agents, we examined the fraction of reciprocal exchanges the MP engaged in, as well as the fraction of turns it was idle and did not offer any exchange. Results are shown in Table 4 for both MP agents and helpful SP agents. They confirm that the percentage of reciprocal exchanges between an MP agent and helpful agents (60%) were significantly higher than that of the MP agent and unhelpful agents (20%), while the MP agent's percentage of idleness with regard to unhelpful agents (39%) was significantly higher than with regard to helpful agents (20%). This also proves that the MP agent successfully identified agents who were exploitive and kept them from taking advantage of it. In contrast, helpful SP agents were more likely to engage in reciprocal exchanges (65%) and far less likely to remain idle towards others (1%), indicating that they were more vulnerable to exploitation. Note that percentages do not up to 1 because we have left out give- and take-type exchanges from this analysis.

Table 5 shows the average performance for each game. The performance of the PD agent increased from game to game, while the performance of unhelpful PD agents decreased from game to game. This result is supported by the fact that the MP agent avoided interacting with unhelpful agents, as shown in Table 4. Interestingly, there was no significant difference in the performance of helpful PD agents from game to game, indicating that the MP agent was able to capture them and adopt just the right balance of cooperative behavior. However, we were surprised that the average performance of both helpful and unhelpful SP agents increased from game 2 to game 3. We wanted to find out whether this increase signaled a process of adaptation to the MP agent. Would this trend be continued in future games ?

	Game 1	Game 2	Game 3
MP agent	149.87	153.06	154.74
Helpful PDs	115.53	115.92	112.67
Unhelpful PDs	116.85	107.29	102.94
Helpful SPs	117.26	107.93	116.28
Unhelpful SPs	95.33	85.96	91.46

Table 5. Repeated Games Results

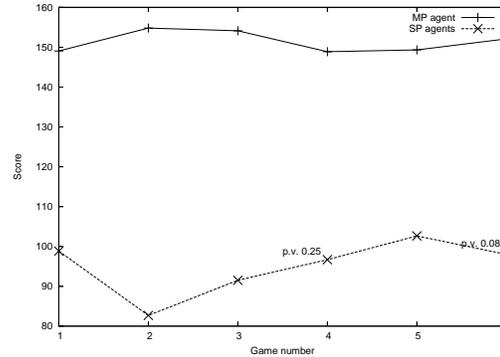


Figure 1. MP agent vs. SP agent Performance (6 repeated games)

To this end, we ran six repeated games played on 4-player board games in which the MP agent played against three other SP agent types. Figure 1 shows average performance of MP vs. SP agents game across the six game series. We labeled the significant differences in score between games whenever the difference was not in the 95% confidence interval. Between game 1 and game 3, the score of the MP agent rose monotonically while the score of the SP agents decreased monotonically. However, the score of the MP agent decreased from game 3 to game 4, while the SP agents increased their score between game 3 and game 5. Although the decrease in the MP agent's score was not statistically significant, we sought to explain it by hypothesizing that unhelpful SP agents who avoided detection prior to game 3, were able to exploit the MP agent in game 4.

This claim was supported by Figure 2 which depicts the difference in score for the MP agent between game 3 and 4 when playing against each SP personality type. As described by the figure, the MP agent suffered a significant decrease in score when playing against unhelpful SP agent types LL and LM, as we expected, but also when playing the helpful agent type HM. To verify that mis-identification is the cause of this behavior, we ran a series of separate games in which a HH SP type agent played against every other SP agent type on each of the boards. We used a HH type agent for this experiment because such a player communicates a lot with other player. Therefore, these results

represent an upper bound on the MP agent’s ability to identify its opponents. We recorded the HH SP agent’s success rate for every personality type of SP agent.

Indeed, the highest error rate in identification (35%) occurred when the HH type agent tried to identify an HM type agent, perhaps because the behavior of this agent is somewhat contradictory. It agrees and proposers considerate exchanges while fulfilling some, but not all, of its commitments. However, we were encouraged by the fact that although some agents escaped identification in earlier rounds, the performance of the MP agent was not affected. It continued to do much better than SP agents, as well as to increase its average score in game 6.

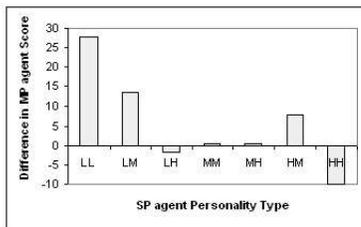


Figure 2. Difference in Score of MP agent between games 3 and 4

4.2. The Influence of Task Dependency

When players need each other’s resources, helpful agents are likely to supply more resources than unhelpful agents. However, they are also taking a chance of getting exploited by players who are deceitful and renege on their commitments. We expected the MP agent to perform better than both helpful and unhelpful agents in any task-dependency role. Also, we hypothesized that when helpful SP and PD agents are task-independent, they will engage in benevolent behavior, and as a result, be taken advantage of by unhelpful players. Therefore we expected their performance to be lower than their unhelpful counterparts. Since all SP agents share the same type of model, we hypothesized that helpful SPs would be able to perform better than unhelpful SPs, because they would identify the exploiters and avoid them.

We analyzed the performance of all agents for each type of task-dependency on the game boards. We presents our results in table 6. Recall that all agents were task dependent in the allDep game board, while one agent was task independent on the oneSelf board, and all other agents were task dependent on it.

Results confirmed the second part of our hypotheses, but not the first: Unhelpful SP and PD agents were more successful than helpful agents when they were task independent,

Agent Type	Task Dependent	task Independent
MP agent	160.32	231.2
Helpful PDs	137.23	221.16
Unhelpful PDs	150.58	253.54
Helpful SPs	113.26	231.93
Unhelpful SPs	93.69	241.2

Table 6. Agent Performance by Task Dependency

and helpful SPs were more successful than unhelpful SPs when they were task dependent. When it was task dependent, the MP agent’s performance was significantly higher than both SP and PD agents’ performance. When it was task-independent, the MP agent performed better than unhelpful SPs, but not better than unhelpful PDs. They were able to take advantage of it. Since the MP agent’s score is higher than that of helpful SPs when they were task-dependent, we inferred that helpful SPs were taken advantage more by unhelpful PDs than the MP agent.

To test this premise, we computed the exchange type ratio that the MP agent engaged in with others in each task dependency role, as described in Figure 3. When it was task dependent, the MP agent engaged in idle behavior much more often when dealing with unhelpful PD agents (42%) than when dealing with unhelpful SPs (21%). This confirmed that the MP agent was able to “cut its losses” when playing with unhelpful PD agents, although they were still able to take advantage of it in this situation. However, the score difference in favor of the MP agent when it was task-dependent was much higher than the score difference in favor of the MP when it was task-independent – therefore on average, the performance of the MP agent was much higher than all other agents in the system, confirming our findings in Section 4.1. Another interesting result was that unhelpful PDs were more successful than helpful PDs when they were task-dependent. We inferred that it was more difficult to predict the behavior of agents for which their was no central design scheme.

4.3. MP agents and Social Welfare

Our hypothesis was that any group of agents would increase its overall social welfare when playing with an MP agent. This is because MP agents engage in helpful exchanges that would not be realized when other agents are playing. To evaluate this claim, we ran a series of 2-player repeated games which included SP and PD type agents, but did not include MP agents, and compared it to the performance of each agent type after including an MP agent in the group. The results are described in Figure 4 and are statistically significant with $p\text{-value} < 0.08$. The performance of helpful and unhelpful agents increased significantly when interacting with the MP agent. As expected,

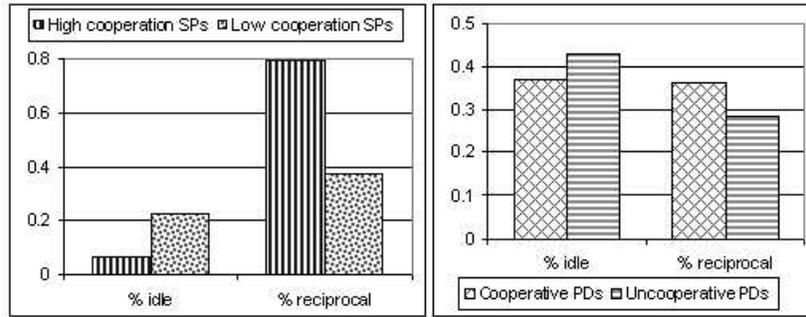


Figure 3. MP Exchange Behavior when task dependent. Left: with SP agents Right: with PD agents

this increase was more profound for helpful SP and PD agents.

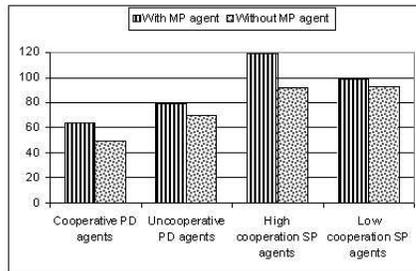


Figure 4. Performance with/without MP agent

5. Conclusions and Future Work

We have presented a model of negotiation which explicitly represents agents' measure of helpfulness personality traits. We evaluated this model in an open system for which there was no central design for the control of agents and in a domain characterized by uncertainty over agents' resources, as well as their cooperative and reliability tendencies. We have shown that agents that adopt a different cooperativeness and reliability measure, depending on who they interact with, could outperform all other agents in the systems. They could identify others personality, and adopt the right balance of behavior towards them more quickly and accurately than other agents. This enabled them to reciprocate helpful behavior while punishing deceitful behavior. Also, they improved the performance of all agents in the system, including unhelpful agents. We showed that when helpful agents are task-independent, they engage in benevolent behavior and are taken advantage by unhelpful agents. However, when they are task-independent, helpful players do better than unhelpful players, who do not realize the full potential of cooperating with others.

One future goal is to use personality models to describe team formation in agent systems. Are agents which exhibit similar personality traits more likely to form teams? How will this affect the performance of the system? Also, we intend to use this model to build computers which interact with people. It has been shown that people's behavior in negotiation is affected by their preferences towards others outcomes, as well as their own [2]. We are interested to see if modeling people's measures of helpfulness can lead to better computer negotiators and explain people's behavioral tendencies outside of the game. We are also planning to tailor the model to deal with higher levels of uncertainty. For instance, by limiting the visibility of the board, it will be difficult for agents to assess their own chances of getting to the goal.

References

- [1] H. Brandsttter and M. Knigstein. Personality influences on ultimatum bargaining decisions.
- [2] Y. Gal, A. Pfeffer, F. Marzo, and B.J Grosz. Learning social preferences in games. In *Proc. 19th National Conference on Artificial Intelligence (AAAI)*, 2004.
- [3] B. Grosz, S. Kraus, S. Talman, B. Stossel, and M. Havlin. The influence of social dependencies on decision-making: Initial investigations with a new game. In *Proceedings of AAMAS-2004*.
- [4] L.M.J. Hogg and N.R. Jennings. Socially intelligent reasoning for autonomous agents. *IEEE Trans on Systems, Man and Cybernetics - Part A* 31(5), pages 381–399, 2001.
- [5] R. Myerson. *Game Theory*. Harvard University Press, 1991.
- [6] S. Sen and P.S. Dutta.
- [7] V. Lesser X. Zhang and T. Wagner. Integrative negotiation in complex organizational agent systems. In *the Proceedings of the First International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2002), research abstract*.