



The Influence of Social Dependencies on Decision-Making: Initial Investigations with a New Game

The Harvard community has made this
article openly available. [Please share](#) how
this access benefits you. Your story matters

Citation	Grosz, Barbara J., Sarti Kraus, Shavit Talman, Boaz Stossel, and Moti Havlin. 2004. The influence of social dependencies on decision-making: Initial investigations with a new game. In Proceedings of the Third International Joint Conference on Autonomous Agents and Multi Agent Systems, (AAMAS-2004), ed. International Joint Conference on Autonomous Agents and Multiagent Systems, and Nick Jennings. 782-789. New York, NY: ACM Press.
Published Version	http://dx.doi.org/10.1109/AAMAS.2004.264
Citable link	http://nrs.harvard.edu/urn-3:HUL.InstRepos:2640578
Terms of Use	This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA

The Influence of Social Dependencies on Decision-Making: Initial Investigations with a New Game*

Barbara J. Grosz
Division of Engineering and Applied Sciences
Harvard University,
Cambridge MA 02138 USA

Sarit Kraus, Shavit Talman,
Boaz Stossel and Moti Havlin
Dept. of Computer Science, Bar-Ilan University
Ramat-Gan 52900 Israel

Abstract

This paper describes a new multi-player computer game, Colored Trails (CT), which may be played by people, computers and heterogeneous groups. CT was designed to enable investigation of properties of decision-making strategies in multi-agent situations of varying complexity. The paper presents the results of an initial series of experiments of CT games in which agents' choices affected not only their own outcomes but also the outcomes of other agents. It compares the behavior of people with that of computer agents deploying a variety of decision-making strategies. The results align with behavioral economics studies in showing that people cooperate when they play and that factors of social dependency influence their levels of cooperation. Preliminary results indicate that people design agents to play strategies closer to game-theory predictions, yielding lower utility. Additional experiments show that such agents perform worse than agents designed to make choices that resemble human cooperative behavior. The paper describes challenges raised by these results for designers of agents, especially agents that need to operate in heterogeneous groups that include people.

1. Introduction

This paper addresses the problem of the design of computer agents that make appropriate decisions in groups comprising both human and computer agents. It investigates settings in which agents' choices affect the outcomes of other agents. In the absence of explicit utility benefits to cooperation, standard economic game theory analyses predict that no cooperation will ensue. A wide-range of results in behavioral economics and psychology contradict these results [7, 8, inter alia]. Prior research in multiagent systems has shown the benefits of cooperativeness to social welfare

[12, 3, 6, 9, inter alia]. We aim to develop models that support the design of self-interested agents that cooperate appropriately with both humans and other agents.

As a first step toward this goal, we undertook experiments based on a new computer game that highlights decision-making in group settings. The game provides a framework for investigating human-decision-making, the effects of different automated decision-making strategies, and comparisons between the two. It also provides a vehicle for examining the ways in which people design computer agents and the performance of different agent designs. The results of these experiments confirm that people cooperate even in the absence of direct utility benefits, that doing so is beneficial and that social dependencies influence behavior.

The remainder of this introduction presents the new game specification and the design desiderata underlying it. Section 2 discusses social dependency factors and the experimental design in which they are explored. Section 3 presents the results of initial experiments of people and computer agents playing the game in different settings; it compares people to computer agents as well as analyzing the effects of different automated strategies. The concluding section discusses implications of the experimental results.

1.1. Testbed for Investigating Decision-Making in Group Contexts

The game Colored Trails (CT) was designed to enable investigation of properties and consequences of decision-making strategies in multi-agent contexts in which agents' choices affect not only their own outcomes but also the outcomes of other agents. It allows for specification of different reward structures, enabling examination of such trade-offs as the importance of the performance of others or the group as a whole to the outcome of an individual and the cost-benefits of collaboration-supporting actions. The game parameters may be set to vary environmental features such as

* The work reported in this paper was supported in part by NSF Grants IIS-0222892 and IIS-0222914. Kraus is also affiliated with UMIACS.

task complexity, availability of and access to task-related information, and the dependencies between agents.

A key determinant of CT design was the goal of providing a vehicle for comparing the decision-making strategies people deploy when they interact with other people with those they deploy when computer systems are members of their groups. We wanted people to be challenged in playing the game and interested in building agents that could play it. The CT architecture allows games to be played by groups comprising people, computer agents, or heterogeneous mixes of people and computers. As a result, CT may also be used to investigate learning and adaptation of computer decision-making strategies in both human and computer-agent settings [4].

As a test-bed environment for decision-making strategies, CT provides several features not present in other multi-agent games and simulation environments [13, 5, 11, *inter alia*]. It highlights the possible influences of inter-agent relationships on decisions, rather than focusing on plan execution, modification or group performance. The complexity of the game may vary across several dimensions including the number of players; the information about the environment available to different players; information about individual agents available publicly to all players, to subgroups, or only privately; the scoring rules; the types of communication possible among agents; and, the negotiation protocol. CT's wider variety of parameters also means it can model more complex scenarios than the games typically used in behavioral economics [2, 10, *inter alia*].

1.2. Colored Trails Game Specification

Colored Trails (CT) is played by two or more players on an $N \times M$ board of colored squares with a set of chips in colors chosen from the same palette as the squares. For each game of CT, one or more squares are designated as goal squares. Each player's piece is located initially in one of the non-goal squares, and each player is given a set of colored chips. The goal squares, distance to the goal and number of chips may vary for different players. A piece may be moved into an adjacent square, but only if the player turns in a chip of the same color as the square. The scoring function and corresponding player-objectives¹ of a CT game may be varied to provide for testing of different kinds of decision-making contexts, but these objectives are all of the general form that certain players or a certain number of players end up in a specified goal square. Chips may be exchanged, and the conditions of exchange varied to model different group dynamics and decision-making situations.

¹ In this paper, "goal" will be used only to refer to those squares on the board designated as goals. Players may have many objectives besides getting their pieces in the goal square. To avoid confusion, the player will use "objectives" rather than "goal" to refer to this larger set.

The scoring function is a parameter of CT game instances and may be set to reflect different possible social policies and utility trade-offs. This function establishes a context in which to investigate the effects of different decision-making mechanisms. For example, by varying the relative weights of individual and group good in the scoring function, collaborative behavior may become more, or less, beneficial.

Two parameters of a CT game may be used to vary the inter-dependence of players. First, the scoring function may stipulate a *reward dependence* by having the scores of a player depend in some way on the scores of other agents. Second, there is a *task dependence* that arises whenever players lack the chips they need to reach their goals and must depend on other players supplying those chips. The term "task dependence" reflects the similarity with agents depending on others for the performance of their tasks.

The CT framework allows agreements between players to be either enforceable or not. If a player's score depends on that player's performance alone, then the combination of unenforceable agreements and a finite horizon leads to a theoretical equilibrium result that no chips will be exchanged. These results apply whether or not players have information about each other's chips and follow from an argument similar to those for the repeated prisoners' dilemma with a finite horizon.

CT satisfies the design desideratum of providing an appropriate abstraction of the general task and decision-making situations faced by computer agents in multi-agent system situations. Play of CT models approximately the performance of actions by a group of agents. Colors correspond to agent capabilities and skills required by tasks; possession of a color chip corresponds to having a skill available for use at a time; not all agents get all colors much as the agents of a group activity have different capabilities and availability. Paths through the board correspond to complex tasks the constituents of which are individual tasks requiring the skills of the corresponding color.

Player objectives may be set to impose a need for players to make task allocation agreements. For instance, if the objectives for a game specify only that a certain number of agents get to a particular goal square without specifying individual goals for each agent, the players need to agree about which agents will head to different goal squares. The game environment may also be set to model different knowledge conditions. For example, varying the amount of the board an agent can "see" corresponds to varying information about task constituents or resource requirements, whereas varying the information players have about each other's chips corresponds to varying information agents have about the capabilities of others.

Various requirements on player-objectives, goal squares and paths correspond to different types of group activities and collaborative tasks. To distinguish collaborative team-

work from settings in which agents act independently, the scoring function may have a significant reward-dependence factor. To model the need for agents to be helpful, players may be task dependent; helpful behavior occurs when a player who has a chip needed by another player gives the chip to that player in some reasonably balanced exchange.

2. Experimental Design: Social Dependency

Our initial investigation using CT examined “helpful behavior choices”, the decisions agents need to make about assisting others in their individual responsibilities, either to make possible another agent’s completion of a task or to improve the quality or decrease the costs of another agent’s task performance. In general task situations, the circumstances under which one agent will decide to help another vary and may depend on such factors as whether the agents are on a team or acting completely independently.

The initial experiments examined the performance of individuals in different reward-dependence conditions (cf. [13]). Although we are ultimately interested in team behavior, teams comprise individuals. To determine team influences on behavior and performance, baseline measures of the performance of individuals in different environments are required, both for people playing one another and for games in which computer agents participate.

To vary reward dependence and thus establish different social environments, a “social dependency factor” ($SDwt$) was included in the scoring function. In particular, for player P_i the experiments used the scoring function,

$$score(P_i) = base(P_i) + SDwt \frac{\sum_{j \in \{1, \dots, N\}, j \neq i} base(P_j)}{(N-1)},$$

where $base(P_i)$ is determined by the performance of P_i alone and N is the number of players. If $SDwt$ is zero, a player’s score is independent of the performance of other players; if it is non-zero, a player’s score combines that player’s individual-performance score and a weighted average of the individual scores of the other players.

The game protocol comprised two phases, a communication phase and a movement phase. Agreements reached during the communication phase were not binding. Players could send chips to other players throughout the communication phase. However, to simulate simultaneous sending of chips, the game controller only delivered chips at the end of the communication phase. To allow comparison of CT play with the no-exchange game theoretic equilibria, chip exchange agreements were not enforceable.

The initial experiments investigated the following two hypotheses: (1) Higher $SDwt$ will lead to an increase in helpful behavior. In particular, when $SDwt$ was higher, we expected agents to give other agents chips more frequently and to ask for fewer chips in exchange. (2) If players can reach the goal without the help of others, they will give

chips to others less frequently and will ask more in exchange. However, players able to reach the goal on their own, will be more helpful if $SDwt$ is higher.

For these experiments, the CT games were played by groups of four players, the board was 6x6, the palette was 5 colors, and there was a single goal square for all players. This setting was chosen to be relatively simple but complex enough both to study the effects of $SDwt$ within group environments and to provide a good set of baselines for subsequent investigations of team behavior. The four-player set-up allows for multiple sources of potential help for players missing chips. It also provides a baseline for comparison in subsequent research with games played by two teams of two players each, which is the smallest possible team size. The 6x6 board size was chosen to restrict the complexity of path-finding for both people and computer agents, and the palette was set to enable interesting chip distributions.

Most of the games were played with players able to see the full board (*board visibility*), but not provided with any information about the chips held by other agents (no *chips visibility*). Full board visibility limits the knowledge acquisition needs of players. They are able to compute the chips they require individually to reach the goal and also the chip needs of other players. The restriction of chip distribution information separates, to some degree, decisions about chip exchanges from scoring information, thereby making helpful behavior distinct from score optimization computations. If players have complete knowledge of the chip distribution as well as the scoring function, they can compute for each possible chip (re)distribution the change in their own scores and the change in all other agents’ scores. Thus, chip-exchange decisions could become simply decisions about relative score improvements. This possibility has three problems. First, there is the computational cost of examining the set of possible chip re-distributions in an attempt to “optimize”. Second, it turns “helpfulness” into a utility function computation rather than separating out a helpfulness “characteristic” factor for examination. Third, full chip visibility corresponds in the task-analogue to full knowledge of other agents’ capabilities, which is an unrealistic assumption. As the remaining sections of this paper reveal, though relatively simple, this setting was complex enough to generate interesting results, lead to a range of systems design choices, and yield varying behavior on the part of both people and computer agents.

Three classes of experiments were performed, one involving 4-player groups of people and the other two involving 4-player groups of computer agents. Subjects for the human player groups were drawn from a population of upperclass and master’s computer science students at Bar Ilan University who were not experts in negotiation strategies nor in economic theories directly relevant to agent design (e.g., game theory, decision theory). Two types of computer agents were deployed: peer-designed agents (PDAs)

and controlled-design agents (CDAs). The PDAs were developed by subjects drawn from the same population as, but distinct from, those who played in the 4-person experiments. One goal of this experiment was to determine whether the subjects would design agents differently depending on whether SDwt was a factor in the score or not. A secondary goal was to determine whether they would design agents to play the way their peers did.

The CDAs were designed to be tunable with respect to the level of cooperativeness of agents as reflected by their willingness to trade chips and the kinds of exchanges they made. Three CDA types were designed, low-cooperative (LC), medium-cooperative (MC), and highly cooperative (HC) agents. The LC agents' strategy was close to the no-exchange game-theory equilibrium strategy. They never gave chips to other agents. HC agents embodied a strategy at the other extreme. They almost always responded to requests or offered chips to other agents. When they were able to reach the goal without help, they would offer 1:1 chip exchanges. If they needed to obtain chips from other agents to reach the goal, they would propose 2:1 deals in which they offered twice as many chips as they requested. They never asked for more chips than they were willing to give. MC agents had a strategy between LC and HC. Like the HC agent, they exchanged chips, but they attempted to obtain more chips than they gave on each exchange. They would propose 1:1 exchanges only if a chip they required to reach the goal was needed for their next move.

The individual performance of a player of CT may be measured according to different criteria, corresponding approximately to different ways of measuring task performance. The scoring rule used for the experiments described in this paper incorporates three factors in the individual-agent score: (1) whether the player reached the goal state (analogous to completing its tasks); (2) the distance of the player from the goal square, if the goal is not reached (closer is analogous to completing more of its tasks); (3) the number of chips the player possessed at the end of the game (related to the cost of performing its tasks). Both the experimental set up for people and the instructions to agent designers made clear that performance of individuals was measured non-competitively; players were to try to maximize their own scores, not to minimize other agents' scores.

In addition to comparing agents' scores, we analyzed the influence of SDwt on the players' exchange rates, distinguishing three kinds of players with respect to exchanges in a game: (1) *reciprocal exchange players*: players who both sent chips to other players and received chips from other players; (2) *take exchange players*: players who received chips from other players, but did not send any chips to others; and (3) *give exchange players*: players who sent chips, but did not get any chips in return. In addition, we categorize as *idle* players who neither sent nor received chips.

The experiments used two different game boards. The

“all dependent” (*AllDep*) board was constructed so that every player was task dependent. To reach the goal, all players needed at least one chip from another player, and each player had some chips it could offer to help other players. The “one self-sufficient” (*OneSelf*) board was constructed such that one player was not task-dependent. This self-sufficient player (the *SelfS* role) was able to reach the goal without help from any other player. The three other players were task-dependent on the *SelfS* player; they needed a chip that only that player could provide.

3. Experimental Results

The experimental results will be presented separately for the games played by people, PDAs, and CDAs. Some results will separate the *AllDep* and *OneSelf* boards, but others will combine performance across these settings. Experiments in which players had chip visibility will be referred to as “full visibility” and those in which players could not see each other's chips will be referred to as “no visibility”. To explore CT's use in investigating a wider range of possibilities, the initial experiments varied a number of features of the experimental setting. As a result, in some cases, the numbers of players for certain settings are not sufficient to provide statistically significant results. We indicate the level of significance when the p-value ≤ 0.05 . A small number of additional especially interesting results are given that though suggestive require more extensive testing to establish significance.

3.1. Experimental Setting and Basic Results

In the experiments in which people played CT, 208 upperclass undergraduate and graduate computer science students at Bar-Ilan University participated in 143 4-player games. On average, each subject participated in 2.75 games. Of these games, a total of 64 games were played with full board visibility and 79 with no visibility. The *AllDep* board games were played in two reward-dependency conditions, SDwt=0 and with SDwt=0.1. The *OneSelf* board games were played with SDwt=0 or SDwt=0.9. Subjects communicated and made moves through CT's GUI using a strictly controlled negotiation language and were not permitted to interact otherwise. They were not told the identities of the subjects with whom they played, nor were they able to see each others' terminals.

The results, analyzed along a number of dimensions in subsequent sections, are summarized in Table 1. The column labelled “# reached goal” gives the average number of people reaching the goal in each game. “Private score” gives the average base score of the agents. “Total Score” is the average total score which includes any influence of other agents' performance.

The PDAs were obtained from 23 one- or two-person teams of upperclass and master's computer science students

		# reached goal	Private score	Total score
No visibility	<i>AllDep</i> $SDwt=0$	2.53	161.03	161.03
	<i>AllDep</i> $SDwt>0$	2.83	177.34	195.08
	<i>OneSelf</i> $SDwt=0$	2.41	141.58	141.58
	<i>OneSelf</i> $SDwt>0$	2.73	158.71	301.55
Full Visibility	<i>AllDep</i> $SDwt=0$	2.44	156.04	156.04
	<i>AllDep</i> $SDwt>0$	3.30	200.04	220.04
	<i>OneSelf</i> $SDwt=0$	2.60	151.87	151.87
	<i>OneSelf</i> $SDwt>0$	2.61	149.87	284.75

Table 1. Results of People Playing CT

		# reached goal	Private score
PDA	<i>AllDep</i>	1.2	94.67
	<i>OneSelf</i>	1.1	70.13
CDA/LC	<i>AllDep</i>	0	35.625
	<i>OneSelf</i>	1	71.625
CDA/MC	<i>AllDep</i>	2	136.75
	<i>OneSelf</i>	1	72.5
CDA/HC	<i>AllDep</i>	4	237
	<i>OneSelf</i>	4	223

Table 2. Homogeneous Groups in games with No Visibility, $SDwt = 0$ Condition: PDAs, CDA/LC: Low Cooperation Controlled Design, CDA/MC: Medium Cooperation, CDA/HC: High cooperation.

at Bar Ilan, peers of (but distinct from) the human players. They were designed only for games with no chips visibility. Eleven agent-design teams were given a scoring rule that included $SDwt$. Twelve teams had no knowledge of possible reward dependence; the game specification they were given did not include $SDwt$. An analysis of the agent-design documents revealed that the scoring rule was seldom used directly in reasoning about exchanges; instead agents were designed to attempt to reach the goal with as many chips as possible. Thus, to our surprise, there were no significant differences between agents designed with and without $SDwt$. As a result, the experiments with PDAs focused on games with $SDwt=0$. Of the 23 teams, 11 implemented agents that could be used in CT experiments.

This agent-design experiment, although preliminary and small in scale, suggests that superficial, implicit mention of reward-dependence in the design specification may not affect design behavior. In contrast, this same incidental mention of $SDwt$ in instructions to people playing the game did engender different behavior as discussed below.

Table 2 summarizes the basic results of games played by

four, identical PDA agents (top section), and those played by four, identical agents of each of the CDA types (remaining sections). Each cell contains an average from the play of at least 40 individual agents. These results are compared with human performance in the next section. We then analyze the results with respect to the social dependency hypotheses and the influence of visibility.

3.2. Analysis and Comparison: PDAs and Humans

We compared the performance in games with no visibility and $SDwt = 0$ of four-player games of homogeneous PDAs with that in games played by people. We expected the performance of these groups to be similar because the PDAs were developed by students drawn from the same population as the human players. Furthermore, we hypothesized that the ability of PDAs to consider a larger set of possible paths to the goal and to send and respond to a large number of messages more easily, would give the PDAs a slight advantage over the human players.

To our surprise, people played significantly better than the PDAs. More players reached the goal and private scores were higher for human players than PDAs, as a comparison of the first and third rows of Table 1 (people) with the top row of Table 2 (PDAs) shows. (Total score cannot be used for comparison, because it differs only in the $SDwt > 0$ settings.) In particular, the average private score for people playing games *AllDep* and *OneSelf* (with no visibility and $SDwt=0$) was significantly higher (t-test, $p \leq 0.05$) than the average private score of the PDAs in these games. Furthermore, the average number of people reaching the goal in these games was significantly higher than the average number of PDAs reaching the goal (χ^2 test, $p \leq 0.001$). CT play on a 6x6 board does not require sophisticated movement strategies, and both human players and PDAs were provided with a path finder procedure that helped them find possible paths to the goal and the chips needed for these possible paths. Thus, these differences in performance cannot be attributed to computational demands on the automated agents.

We explored a number of hypotheses as potential explanations for these results. Superficial possibilities did not hold. PDAs sent significantly more messages in the games than people sent, so lack of communication does not explain the lower performance of PDAs. The hypothesis that the overall number of chips sent by players influences performance was also not supported by analysis of the data. The average number of chips sent by PDAs was similar to the average number of chips sent by human players. However, especially in the no-visibility case, chips that are sent may not be useful for advancing toward the goal. For instance, a player desperate to get a chip to be able to move immediately, may offer a large number of chips in exchange for the one it needs. This kind of exchange increases the overall number of chips exchanged, but may not increase




Figure 1. Chip-exchange Comparison of People and PDAs; y-axis is percentage of players of each exchange type; Left: all games and players; Right: *SelfS* Role *OneSelf* Game.

the number of agents able to reach the goal compared with 1:1 exchanges.

Two analyses were undertaken to examine the hypothesis that people achieved higher scores than PDAs because they were more helpful. First, we examined the percentages across all games of people and of PDAs in each exchange type—reciprocal-, take-, and give-exchange players and idle. Second, the performances of people and PDAs were compared to the different types of CDAs.

The charts in Figure 1 compare human players and PDAs with respect to chip exchange type, both overall (left) and for the *SelfS* role in the *OneSelf* game (right). The percentage of people who were reciprocal- and give-type players was significantly higher (χ^2 test; $p=0.01$) than the percentage of the PDAs in these categories. The percentage of PDAs that were idle (not involved in any chip exchange) was significantly higher (χ^2 , $p \leq 0.001$) than the percentage of people who were idle. These differences are more pronounced for the *SelfS* role in *OneSelf*. This role is one in which the player, from the initial chip distribution at the start of the game, does not need any chips from another player to reach the goal. Any exchanges the player participates in are helpful. Give-type exchanges are evidence of benevolence, since the player in no way takes advantage of the other player’s weaker state. Reciprocal exchanges are also benevolent, because the player, by agreeing to give chips, takes a risk; it may not get any chips in return because agreements are not enforceable and exchanges are simultaneous. No PDA playing the *SelfS* role was of the give-exchange type. Most PDAs playing this role were idle (81%) while few of the human players were (11%). Most people in the *SelfS* role were reciprocal-exchange type (81%).

The results in Table 2 show PDAs perform between LC and MC in *AllDep* games and close to LC in *OneSelf* games. The performance of the human players, as shown in the first and third rows of Table 1, falls between that of MC and HC in both games. Thus, people resemble the more highly co-

operative CDAs, whereas the PDAs resemble the less cooperative ones.

Both analyses support the hypothesis that people’s greater helpfulness led to higher scores. These results resemble research in social psychology on contributions to groups [1, 14].

3.3. The Influence of Reward Dependence

The first hypothesis in Section 2 was that higher *SDwt* would lead to an increase in helpful behavior. We expected agents to give other agents chips more frequently and to ask for fewer chips in exchange when *SDwt* was higher. We also hypothesized that task-independent players would be more helpful if *SDwt* was higher. Our analysis of reward dependence considers only human players, because the PDAs were run only with *SDwt* = 0.

As discussed above, the total number of chips exchanged in a game is not a good indication of helpful behavior. However, overall performance and the percentage of chip exchanges that are reciprocal- and give-type are.

A comparison of the results in Table 1 for human players when *SDwt* = 0 with those when *SDwt* > 0 supports this hypothesis for both the visibility and non-visibility games. The average private score of all games played by people with no visibility in which *SDwt* > 0 was higher than in the games where *SDwt* = 0. Similarly, the average private score of people who played the games with visibility when *SDwt* > 0 was significantly higher than when *SDwt* = 0 (t-test, $p=0.032$).² Considering all the games, with and without visibility the t-test result showed significant of $p=0.057$. In addition, the number of human players who reached the goal in games in which *SDwt* > 0 was significantly higher than for games with *SDwt* = 0 (χ^2 test; $p=0.01$).

² The slight decrease in private score for game *OneSelf* for *SDwt* > 0 from that for *SDwt* = 0 is not significant. Inspection of the game transcripts suggests the relatively poor performance statistics for *SDwt* > 0 in this game setting resulted from small numbers and player errors.




Figure 2. Exchange-type Comparisons for Human Players; y-axis is percentage of players in all games. Left: SDwt; Center: Task Dependence. Right: Visibility;

Reward dependency also influenced the exchange types of the human players. As shown in Figure 2 (left), a higher percentage of players were reciprocal- and give-type when $SDwt > 0$ than when $SDwt = 0$ and a lower percentage were take-type or idle.

3.4. The Influence of Task Dependence

The second hypothesis was that task-independent players, players that do not need other agents' help, would give chips to others less frequently and would ask more in exchange. We thus predicted that such players would be less frequently reciprocal- or give-exchange type players and more often take-type and idle players than would task-dependent players. The analysis of the *Self* role in the *OneSelf* game suggests that task dependency does influence the helpfulness of the players. Figure 2 (middle) shows that reciprocal and idle portions of this hypothesis holds for people. In particular, the percentage of task-dependent players that were reciprocal was significantly larger than the percentage of task independent players (χ^2 test, $p \leq 0.001$). The one surprising result is that the percentage of task-independent players that were give-type was significantly higher than the percentage of task-dependent give-type players (χ^2 test, $p \leq 0.001$). These results suggest that human players are generous and willing to give free chips even when they do not need anything in return. The charts in Figure 3, which give exchange types both people and PDAs for no-chip-visibility games and $SDwt = 0$ are similar.

3.5. The Influence of Visibility

We expected visibility to improve the performance of players and to increase their willingness to help one another. However, as Figure 2 (right) shows, the results only partially support this hypothesis.

The average private score for people in games with chip visibility was higher than their average score in games with no visibility. However, the effect on the players' exchange-type was less consistent. The percentage of players who

were reciprocal-type increased with visibility and the percentage of those who were idle decreased, which supports the hypothesis. However, the percentage of those who were give-types was higher for no-visibility than for visibility. This result may be explained by the fact that with chip visibility, players' needs are known to each other, and as a result, players may be less willing to give out chips freely.

4. Conclusions and Future Work

The experiments presented in this paper demonstrate that the Colored Trails game provides a rich framework for investigating decision-making strategies in multi-agent situations. The experimental results, in particular the superior performances in games played by people and by highly cooperative CDAs, indicate that cooperation is beneficial: it increases the average score and the average number of players reaching the goal. These results suggest that systems designers should build cooperative agents when constructing agents that will engage in group activities with people, not only because it improves performance, but also because people exhibit and expect it.

However, game-theory as well as results of strategic negotiation work suggest a major potential concern: designers of cooperative agents risk their agents being taken advantage of by other, less cooperative agents. To examine this worry, we ran games of 3 MC agents against each of the PDA agents and games of 3 HC agents against each of the PDA agents. The results of these games are reported in Table 3. The values in the table without brackets are those of the PDAs and the values in the brackets are those of the relevant CDA.

We expected the performance of PDAs, which are essentially non-cooperative, to improve when playing with MC and HC agents and to do so at the expense of the CDAs whose performance would degrade significantly. However, the results support only the first part of the hypothesis and a weak version of the second. The performance of the PDAs increases significantly in all games they played with HC agents (t-test, $p=0.06$) and the *AllDep* games with MC




Figure 3. Exchange-types for No Visibility, $SDwt = 0$ Condition. Left: Humans, Right: PDAs

		# reached goal	Private score
With MC	<i>AllDep</i>	1.72 (1.57)	127.25 (138.42)
	<i>OneSelf</i>	1 (1)	74 (69.51)
With HC	<i>AllDep</i>	2 (2.43)	151.05 (193.13)
	<i>OneSelf</i>	3.53 (2.27)	198.19 (168.65)

Table 3. Results of CDA-PDA Play

agents (t-test, $p \leq 0.001$). In the *OneSelf* game when playing with MC agents the average PDA scores did not change.

As expected, the average private score of the HC and MC CDAs decreased when playing with a non-cooperative PDA in the group, as did the number of the agents that reached the goal. However, the decrease was not down to the level of the PDAs or LC agents. HC's average private score changed from 237 to 193.13 in *AllDep* and from 223 to 168.65 in *OneSelf* games; both these scores reflect agents making it to the goal a significant portion of the time. The private score of MC agents slightly increased in *AllDep* games and only slightly decreased in *OneSelf* games. It is also noteworthy that the HC agents had the highest scores among the three type of agents even in this heterogeneous environment.

These results suggest that it is beneficial to design and implement cooperative agents. Together with the main results of games played by people, they suggest that cooperative PDAs will benefit even more from being cooperative in heterogeneous groups. Related CT-based investigations of social reasoning and learning indicate that agents that model and learn the social factors important to people are more successful than agents that adhere to the equilibrium strategy [4].

The initial CT experiments suggest several directions for future research. Game variations other than *AllDep* and *OneSelf* are needed to further explore the conditions under which people are cooperative. Games in which people and computer agents play are needed to understand the ways in

which people will react to different types of agent strategies as well as to formulate better strategies for agents in such settings. In this experiment, the number of student-designed agents was small; additional investigations of agent design are needed to determine whether the contrast between agent design and people's play is idiosyncratic or a general phenomenon.

References

- [1] D. Abrams and M. A. Hogg, editors. *Social Identity and Social Cognition*, chapter 9, pages 197–229. Blackwell Publishers, Oxford, England, 1999.
- [2] C. F. Camerer and E. Fehr. Measuring social norms and preferences using experimental games. In *Foundations of Human Sociality: Experimental and Ethnographic Evidence from 15 Small-scale Societies*. Oxford University Press, 2003.
- [3] U. Endriss, N. Maudet, F. Sadri, and F. Toni. On optimal outcomes of negotiations over resources. In *Proc. of AAMAS-03*, pages 177–184, 2003.
- [4] K. Gal, A. Pfeffer, F. Marzo, and B. Grosz. Learning social preferences in games. In *AAAI-04*, 2004.
- [5] A. Greenwald. The 2002 trading agent competition. *AI Magazine*, 24(1):83–91, 2003.
- [6] N. Griffiths and M. Luck. Coalition formation through motivation and trust. In *Proc. of AAMAS-03*, pages 17–24, 2003.
- [7] J. Kagel and A. Roth, editors. *The handbook of experimental economics*. Princeton University Press, 1995.
- [8] D. Kahneman, P. Slovic, and A. Tversky, editors. *Judgment under Uncertainty: Heuristics and Biases*. Cambridge University Press, 1982.
- [9] S. Kraus, O. Shehory, and G. Taase. Coalition formation with uncertain heterogeneous information. In *Proc. of AAMAS-03*, pages 1–8, 2003.
- [10] H. R. Nalbantian and A. Schotter. Productivity under group incentives: An experimental study. *The American Economic Review*, 87(3):314–341, 1997.
- [11] RoboRescue. www.rescuesystem.org/robocuprescue. 2004.
- [12] S. Sen and A. Biswas. Effects of misconception on reciprocal agents. In *Proc. Agents-98*, pages 9–13, 1998.
- [13] D. G. Sullivan, B. J. Grosz, S. Kraus, and S. Das. The influence of social norms and social consciousness on intention reconciliation. *AIJ*, 142(2):147–177, 2002.
- [14] H. Tajfel. *Social Identity and intergroup behavior*. Cambridge University Press, Cambridge, England, 1982.