The Effect of Context on Decision Making in Colored Trails

A Thesis presented

by

F. Alexander Allain

to

Computer Science

in partial fulfillment of the honors requirements

for the degree of

Bachelor of Arts

Harvard College

Cambridge, Massachusetts

April 4, 2006

Contents

Acknowledgements

I would like to thank Professors Avi Pfeffer, Stuart Shieber and Barbara Grosz for their advice and encouragement, Kobi Gal for his advice, encouragement and aid in running the experiments, Alex Hoffer for her advice and proofreading skills, and Michael Kern and Sam J Lin for their advice regarding the psychology literature.

Chapter 1

Introduction

Decision makers, or *agents*, interact with each other all the time—for instance, two executives negotiating a contract could be thought of as two agents interacting. In the past, most interactions psychologists and economists have concerned themselves with involved humans, but the increased saturation of computers in society has dramatically increased the amount of interaction between humans and non-human agents. Such interactions now take the form primarily of prescribed usage following limited protocols, such as connecting to a web server over the internet. In the future, many interactions are likely to be naturalistic. You may receive technical support not from a call center staffed by humans but by one staffed by computer agents designed both to help you and to avoid costly fees or problems for the company employing the machines.

Work is now being done to address the challenges of designing more sophisticated computer agents capable of interacting in a heterogenous environment of humans and computers [4]. For example, human-computer interfaces may be improved by building agents that adapt their responses to the goals and desires of the human actors—for example, one project had as its goal building agents capable of assisting users of Microsoft Office [5]. In this case, the primary challenge is modeling the human's needs.

More sophisticated interactions may require the computer agent to reason about both the needs of the human and its own “preferences.” For example, computers that understand human decision-making processes can bargain with humans while learning to perform better than human agents [3]. In such scenarios, agents might function as decision support tools for corporate negotiations or even simply to improve AI in strategy games.

In this work, we analyze the assumptions built in to Colored Trails, a game used to study interactions between groups of humans and computers.

1.1 Game Theory

Game theory is the study of how individuals interact when their choices affect the outcome of the interaction. Initial work on decision-making focused on so-called *rational* agents that make choices based strictly on the value of the outcomes to the agents. In practice, however, humans do not always follow the strategy of rational agents [1].

The Ultimatum Game

The ultimatum game provides an example in which game theoretic decision-making does not hold. In the ultimatum game, an allocator is allowed to split a fixed pot of money between himself and a deliberator. The allocator can give the deliberator anywhere from nothing to the entire pot.

If the allocation is accepted, then both participants receive the amount of money offered. If the allocation is declined, both participants receive a default allocation: usually, no money for either participant.

The ultimatum game is a good example of a game where participants do not follow the prescription of game theory. In theory, if the participants will never meet again, a game-theoretic rational deliberator should accept any offer above 0 dollars, so the allocator should offer the minimal amount possible. In practice, offers giving the deliberator less than 20% are accepted only 50% of the time, and most offers give a minimum of 30% [1]. This effect holds even when large sums of money are involved and occurs across nearly all cultures [1].

Behavioral Game Theory

Behavioral game theory attempts to extend game theory to account for other factors that may influence the negotiation, such as whether one participant feels that an offer is unfair. A variety of results suggest different factors, such as a desire for equality on the part of the deliberator, influence decision-making; theorists also provide formulas to compute the best choice of action when these factors must be accounted for [1].

Likewise, computer systems that wish to reason appropriately about humans will almost certainly need to go beyond traditional game theory and focus on factors such as those suggested by behavioral game theory. However, computer agents may wish to learn how humans behave on the fly, using factors from behavioral game theory rather than merely borrowing its results wholesale (for instance, see Gal et. al. [3]) in order to adapt to the specific circumstances in which the interaction occurs or to the specific participants taking part in the interaction.

1.2 Complex Domains and Context

Work on human-computer bargaining cannot restrict itself to simple scenarios—in practice, task domains will vary widely in complexity from relatively simple interactions between few participants with a small number of choices and little interdependence up to markets with large numbers of agents taking part in interdependent interactions. Consequently, it is useful to be able to vary task complexity and features of the scenario being studied [4]. In addition, we would like to be able to create scenarios and contexts that may more closely map to interactions one might expect in the real world.

For instance, when outcomes are the result of probabilistically governed events, participants make different choices depending on how the choices are framed [7, 8]. People are willing to accept less favorable offers when the person making the offer can create the perception of having a reason for the offer [6]. For instance, the allocator might explain that a lowball offer in the ultimatum game is necessary because the money is needed to do laundry. Consequently, when analyzing interactions, in addition to the content of the interaction, the context of the interaction—which could provide either a framing effect or a context in which offers may seem less offensive—may be a significant factor in the offers and responses made.

1.3 Contribution of Work

In this work, we show that participants playing Colored Trails, which was designed to meet the above criteria, behave differently than counterparts who see offers that are of equivalent value but presented outside of the context of the game. We show that both the player making offers and the player responding to offers play differently, and more strategically, when the game context is removed. This suggests that the CT metaphor is a relatively useful way to get at behavior and approximate real world interactions as compared to using normal form tables, which list offers only by the values of each outcome for all participants.

Chapter 2

Background

2.1 Negotiation Games

Negotiations

A *negotiation* is a situation in which some individuals have something that other participants want and there is a structure that enables the participants to make an exchange of resources. A proposed exchange will usually be referred to as an *offer*.

The structure of a negotiation need not be fully specified—negotiations may be essentially free–form—but in many cases, large parts of the structure are prescribed. For instance, the exchangeable resources or the length of negotiation may be limited.

In the negotiations discussed within this paper, only particular individuals, termed *allocators*, are allowed to suggest, or *propose*, offers, and other individuals, termed *deliberators*, are limited to accepting or declining an offer.

Negotiations may proceed in rounds where a single individual may propose an offer and the rest of the participants may accept or decline it. If the offer is declined, then another individual may have the opportunity to propose an offer, or *counter-offer*.

Alternatively, negotiations may be limited to only a single offer that the deliberator may only choose to accept or decline—no counter-offers allowed. The result of a declined offer is the *default outcome*, also known as the non-negotiation alternative (NNA). The default outcome may correspond to a situation in which both sides maintain their original resources or it may result in a penalty that deducts some or all resources from both participants. All experiments reported here followed this style of play.

When the allocator and deliberator will not meet again in the context of a negotiation, the negotiation is called a *one-shot* negotiation.

Value

The extent to which resources are valuable depends on the conditions of each participant, as well as the other resources that the participants possess. We will consider situations in which the value of a resource depends on two factors: the other resources possessed and the state of the participants.

The value of a resource may depend on the other resources possessed because those other resources may enable new uses for the original resource. For instance, an uncut diamond is more valuable if the owner possesses the skills and tools to cut it.

The value of certain resources may also depend on the individual positions or capability of the participants. A tightly capped jar of peanut butter may be worthless to a weak individual incapable of opening it, but it may be especially valuable to a participant interested in selling peanut butter sandwiches.

It is possible for a resource to have a negative value if owning it results in a penalty. For instance, peanut butter may have negative value to a person with a peanut butter allergy. On a frenetic vacation, a heavy tchochke may have negative value. In this work, however, all resources will be assumed to have positive value.

In terms of negotiations, the total value of an outcome will be referred to as the *overall value*—or when it is clear, the *value*—of the exchange.

The increase in value over the default outcome that is conferred by an exchange will be referred to as the *benefit of the exchange*, or simply the *benefit*. Exchanges that have a cost rather than a value may have a *negative benefit*.

For notational convenience, when referring to offers consisting of a benefit to the allocator and a benefit to the deliberator, the offer will be written in the following form: (allocator benefit, deliberator benefit). For instance, "$(100, 0)$" indicates an offer with a benefit of 100 to the allocator and 0 to the deliberator.

Utility

The benefit of an exchange is a knowable outcome that is objectively quantifiable. For two participants with the same relative attributes and in the same condition, the value of an exchange should not differ. While it may depend on measurable criteria of the participants, the benefit does not depend on the desires or wants of the participants.

Utility, on the other hand, refers to the desirability of an outcome for a particular participant. Although in principle the benefit of an exchange may be proportional to the utility of an exchange, this is not necessarily the case. As the benefit increases, the utility may increase at a decreasing rate or may even decrease (for instance, a risk-averse individual may prefer to limit his total holdings to avoid fear of losing large sums).

Social factors may play a role in the utility function. Exchanges of equal benefit to both participants may have different utility to one of the participants; if someone desires fairness, an exchange that gives both participants the same benefit—or, potentially, the same overall value—may be preferential to an exchange that gives one participant a higher value than other. For instance, the exchange $(100, 10)$ may have lower utility for the allocator than the exchange $(100, 100)$.

Aside from concerns about inequity of the benefit or the overall value, participants' utility could be affected by the benefit of the offer in comparison to the benefit of other potential offers. A proposed exchange, $(100, 10)$, may be acceptable to a deliberator when the only alternative is the exchange $(100, 5)$, but it may no longer be acceptable when another offer, $(100, 40)$, is also available.

As with benefit, it is possible for an exchange to have negative utility. In fact, it is possible for an exchange to have a positive benefit but a negative utility. For instance, in the example above, where one participant receives 100 points of benefit, and the other merely 10 points of benefit, the second participant may feel that the 10 points is insufficiently valuable to justify the large difference between the outcomes.

Agents for whom the utility function is equivalent to the value function shall be referred to as *value-maximizing*.

Strategy

Much work has been done on how to make optimal offers in negotiations. Usually, the goal is to provide a strategy for the allocator that will maximize his utility. We take some concepts and terminology from economics and game theory to describe types of offers that follow certain strategies. All terminology will be defined for two player, single-offer allocator-deliberator negotiations.

A *Weak Nash Equilibrium* (WNE) offer is one that maximally benefits the allocator while not hurting the deliberator. If we take the set of all offers with non-negative benefit to the deliberator, then the set of WNE offers consists of those offers in this set for which the benefit to the allocator is maximal. For instance, the offer $(100, 0)$ may be a WNE if no other offer that has non-negative benefit to the responder has a higher benefit to the proposer.

A *Strong Nash Equilibrium* (SNE) offer is one that maximally benefits the allocator while also benefiting the deliberator. If we take the set of all offers with positive benefit to the deliberator, then the set of SNE offers consists of those offers from this set for which the benefit to the allocator is maximal.

The difference between a WNE and an SNE is that a WNE allows a benefit of 0 to the deliberator whereas the SNE requires a positive benefit—in effect, the two classes of offer are drawn from different base sets. As a result, an SNE offer is not necessarily a WNE offer, and vice-versa. For instance, if there are three possible offers, $(100, 10)$, $(110, 0)$ and $(90, 90)$, then the first offer is an SNE and the second offer is a WNE, but the converse is not true. It is, of course, possible for an offer to be both an SNE and a WNE offer. When there is no need for a distinction between the WNE and SNE offers, then the modifier will be dropped and such offers will be referred to simply as *Nash Equilibrium* (NE) offers.

A *Pareto Optimal* (PO) offer is one in which no participant can derive more benefit without another participant's making a sacrifice. There may be more than one PO offer. For instance, if there are three possible offers, $(100, 20)$, $(100, 30)$ and $(50, 50)$, then both the second and third offers are PO. The first offer is not PO because the second offer gives the allocator the same benefit as the first offer while also improving the benefit to the deliberator.

Neither WNE or SNE offers need be PO. For instance, in the previous example, both offers that yield 100 points for the allocator would be possible NE, but only the second offer would be PO. Assuming a value-maximizing deliberator, the NE offer will be accepted regardless of whether the offer is PO. In practice, however, one might expect that allocators would generally prefer PO offers.

Dealing with Non-NE Offers

In actual negotiations, not every offer will be an NE. It is useful to have a notion of how "close" an offer is to an NE offer. The *allocator SNE distance* refers to the difference in benefit to the allocator between an SNE offer and a non-SNE offer. More precisely, it is the allocator's benefit at the SNE subtracted from the allocator's benefit from the offer. For instance, if $(100, 10)$ is an SNE offer but the allocator offers $(90, 90)$, then the allocator SNE distance is $90 - 100 = -10$. Note that if the sign of an allocator's SNE distance is positive, then the benefit to the deliberator must be less than or equal to 0. Otherwise, the offer with a positive allocator SNE distance would have been a SNE offer as the allocator clearly benefits from it as compared to the original SNE offer. On the other hand, if the sign is negative, the deliberator will not necessarily benefit more than at the SNE. It is possible that the allocator has simply made a mistake and chosen a sub-par offer that hurts both players in comparison to the SNE offer.

The *deliberator SNE distance* refers to the difference in benefit to the deliberator between a PO SNE offer and a non-SNE offer. This corresponds to the intuitive notion that the point of making a non-SNE offer is to give the deliberator more benefit than could be gained from an SNE offer. Consequently, it is necessary to consider the best SNE offer available for the deliberator, which is the PO SNE.[1] The deliberator's SNE distance is the deliberator's benefit at the PO SNE subtracted from the deliberator's benefit from the offer. For instance, if $(100, 10)$ is an SNE offer, but the allocator offers $(90, 90)$, then the deliberator SNE distance is $90 - 10 = 80$.

It is important to note that a positive deliberator SNE distance indicates that the allocator *has made a sacrifice*. This is because the deliberator SNE distance is defined

[1]Note that because there may be multiple offers with the same resulting value, it is possible to have multiple PO SNE offers in a negotiation, but they must all have the same benefit to the participants.

in terms of a PO offer: any increase in the deliberator's benefit above that provided by a PO offer must be accompanied by a corresponding decrease in the deliberator's benefit.

The *allocator WNE distance* and the *deliberator WNE distance* are defined in similar fashion.

2.2 Colored Trails

The game Colored Trails (CT) is a board game designed as a metaphor for negotiation exchange scenarios. Gameplay in Colored Trails takes place on a game board consisting of colored tiles. Players start the game possessing a set of colored chips. These chips have colors of some, but not necessarily all, of the tiles on the game board. Each player also starts out on an assigned square and may be given a goal square to reach. Players do not necessarily start on the same square or with the same goal square.

Scoring can be modified to change the payoffs to participants in order to model different types of interactions. For instance, different scoring rules can provide bonuses based on group performance or performance of another participant in the game. As a result, CT allows for modeling both *reward dependence*, situations in which players can be rewarded for the performance of the group, and *task dependence*, situations in which the players must work together to reach a goal [4].

Gameplay proceeds as a set of rounds divided into phases until the last round is reached. Games may have a fixed or a dynamic number of rounds.

In the first phase of each round, an *exchange phase*, the players are allowed to make exchanges of chips that they possess. These exchanges are useful because players may move from square to square on the board only by giving up a chip of the same color as the adjacent square being moved to. During the *move phase* of a round, players may make as many moves as can be "paid" for with chips.

Game Context

In the experiments reported and discussed here, each game consisted of two players on a 4x4 board. Games were one-shot negotiations consisting of an allocator and a deliberator. The default outcome was that both players were left with their original chips. The deliberator was not given the opportunity to counter-offer.

Following the offer and reply, the game was scored as though the players had made the optimal move using the chips gained as a result of the interaction.

Chapter 3

Methodology

3.1 Architecture

At a high level, the architecture of the experiment software was a client-server model. The system was designed to allow simultaneous interactions between an arbitrary number of participants, with the option of assigning pairings based on pairings used in a previous experiment. It also allows the experimenter to control several variables through a web-based interface. The most important variables are the settings for the number of participants expected and for the time limits placed on the phases within each round. The experimenter can also control the type of pairings used–whether from the previous experiment using the game context or assigned by a separate pairing algorithm.

3.2 Breakdown of a User Session

Logging In

The participants interact with the system through a web-based interface. The login system is initiated whenever a user visits the homepage. Upon logging in, the user is assigned a tracking ID. The system tracks the number of users in the session to ensure that the expected number of participants is playing the game. After being assigned a tracking ID, the user is forwarded to a holding area that indicates that the system is waiting for all participants: "When all other participants log in, you will be automatically notified."

Once the expected number of participants logs in, the first round—or an optional tutorial round—automatically begins.

Once a round begins, each pair of participants is conceptually separate from the others. For each pair, the round is divided into phases for making or waiting for an offer, making or waiting for a reply, and viewing the final result of the round. Each pairing of allocator-deliberator occurred only once during the experiment to create a one-shot negotiation scenario. The system can handle repeated games, however.

From here, each round is broken up into a series of phases that continue until the session completes: the proposal phase, the response phase and the result phase.

Proposal Phase

Allocator

Figure 3.1 shows the allocator's screen in the *proposal phase*. The allocator sees the set of all possible offers, as well as the default outcome as a pair of values; the value of the exchange to the allocator is on the left, in bold, and the value of the exchange to the deliberator is on the right. Each pair has a radio button allowing the allocator

Figure 3.1: Screenshot of Allocator Panel Prior to Making Proposal

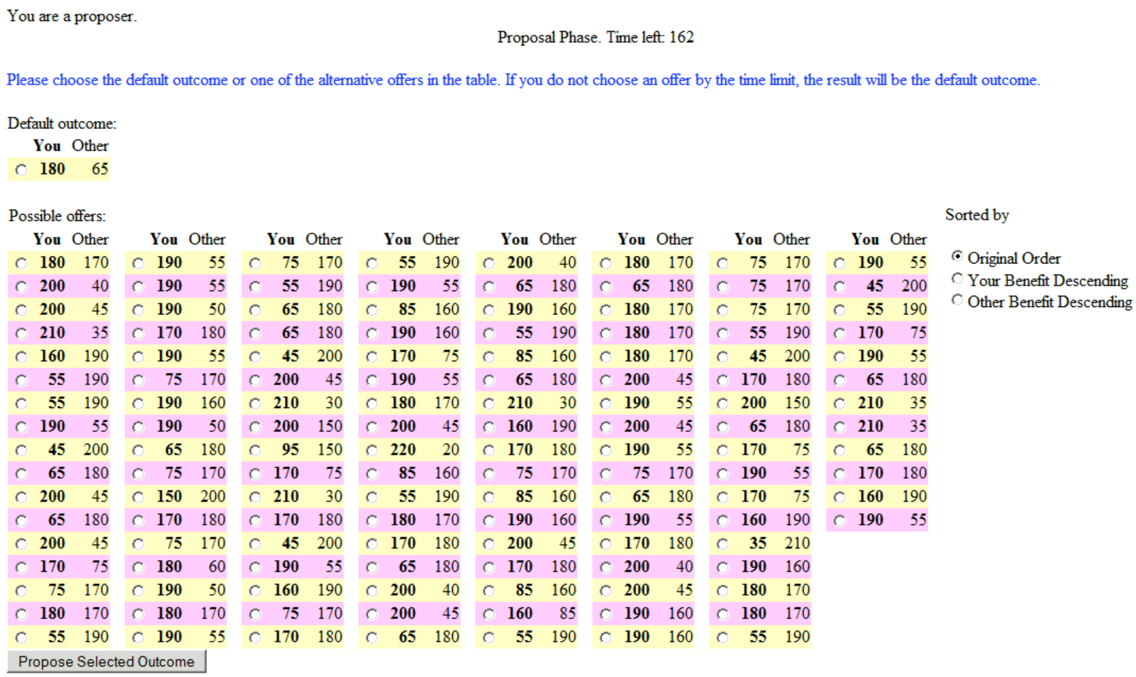

You	Other	You	Other	You	Other	You	Other	You	Other	You	Other	You	Other	You	Other
180	170	190	55	75	170	55	190	200	40	180	170	75	170	190	55
200	40	190	55	55	190	190	55	65	180	65	180	75	170	45	200
200	45	190	50	65	180	85	160	190	160	180	170	75	170	55	190
210	35	170	180	65	180	190	160	55	190	180	170	55	190	170	75
160	190	190	55	45	200	170	75	85	160	180	170	45	200	190	55
55	190	75	170	200	45	190	55	65	180	200	45	170	180	65	180
55	190	190	160	210	30	180	170	210	30	190	55	200	150	210	35
190	55	190	50	200	150	200	45	160	190	200	45	65	180	210	35
45	200	65	180	95	150	220	20	170	180	190	55	170	75	65	180
65	180	75	170	170	75	85	160	75	170	75	170	190	55	170	180
200	45	150	200	210	30	55	190	85	160	65	180	170	75	160	190
65	180	170	180	170	180	180	170	190	160	190	55	160	190	190	55
200	45	75	170	45	200	170	180	200	45	170	180	35	210		
170	75	180	60	190	55	65	180	170	180	200	40	190	160		
75	170	190	50	160	190	200	40	85	160	200	45	180	170		
180	170	180	170	75	170	200	45	160	85	190	160	180	170		
55	190	190	55	170	180	65	180	55	190	190	160	55	190		

to select it.

The offers are initially displayed in an apparently random order, although the ordering is in fact deterministic and repeatable. The allocator can display the offers in descending order by the benefit to the allocator, by the benefit to the deliberator, or display order in the initial state.

The allocator may select the offer before submitting it. The allocator may also choose to take the default outcome, an act that would end the round immediately.

Whenever the allocator elects to sort the table of offers or selects an offer without submitting it, the client sends an event to the server for logging purposes. To maintain a responsive interface, however, the client does not need to wait for the server. Instead, the server sends the client precomputed sorted results as part of the initiation of the round; consequently, the client does not need to wait for a response to the logging event before the participant sees the sorted values.

The allocator's screen indicates that it is the “Proposal Phase” and a countdown ticks off the seconds left in the phase. If the allocator does not make a proposal before the time limit is reached, the result is the default outcome. In the experiment, the proposal phase lasted 180 seconds.

After selecting and submitting an offer, the allocator is taken to a new page displaying the same set of options as before, but with the proposed offer highlighted in a different background color. The options are presented in the original ordering. This starts the "Responder Phase."

Deliberator

For the deliberator, the initial screen is the same set of all possible offers and default outcome that the allocator sees, but with the value of the offer to the deliberator in bold on the left, rather than the right. The deliberator has the same sorting options available as the allocator.

The deliberator is also told that it is the "Proposal Phase" and sees the same countdown as is available to the allocator.

After an offer is made or the time limit is reached, the deliberator is taken to a new page displaying the offer, or, if the offer is the default outcome, the result of the round. If an offer is made, this triggers the "Response Phase" for the deliberator.

Response Phase

Allocator

In the *response phase*, the allocator sees the table and the proposed offer highlighted with a different background color, text indicating that it is the "Response Phase" and a countdown indicating how much longer the deliberator has to make a choice.

Figure 3.2: Screenshot of Deliberator Panel Prior to Making Reply

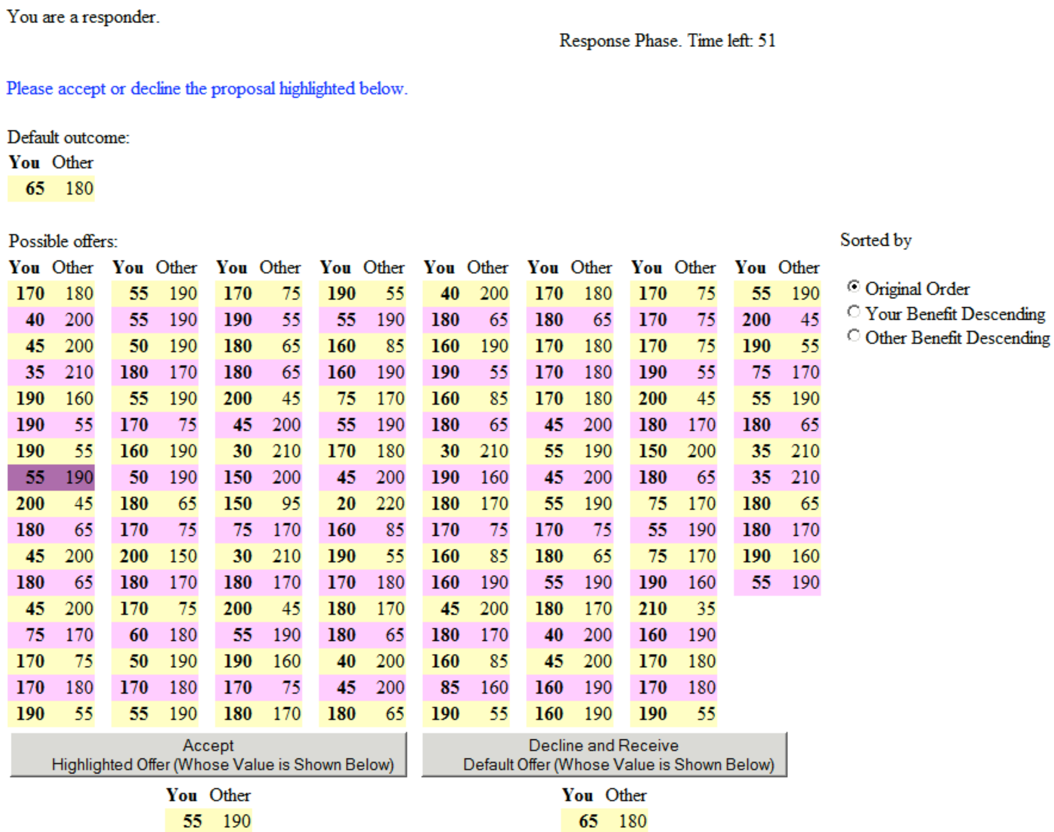

Deliberator

Figure 3.2 shows the deliberator's screen during the response phase. The deliberator is presented with the same table as is shown to the allocator, again with the order of values reversed.

In addition, beneath the table, the deliberator has two buttons, one to accept the proposal, the other to decline the proposal. Beneath each button is shown the outcome corresponding to that choice.

Finally, the response phase is time-limited; if the deliberator does not make a reply within the time limit, then the offer is declined. In the experiment, the response phase lasted for 70 seconds.

Result Phase

Once the deliberator has chosen a response, both participants are forwarded to the *result phase*. Or, if the allocator chooses the default outcome, both participants are immediately forwarded to the result phase, skipping the response phase.

In the result phase, both participants are told the outcome of the round—whether the offer was accepted or declined and the resulting value for each participant.

Figure 3.3: Screenshot of Deliberator in Result Phase

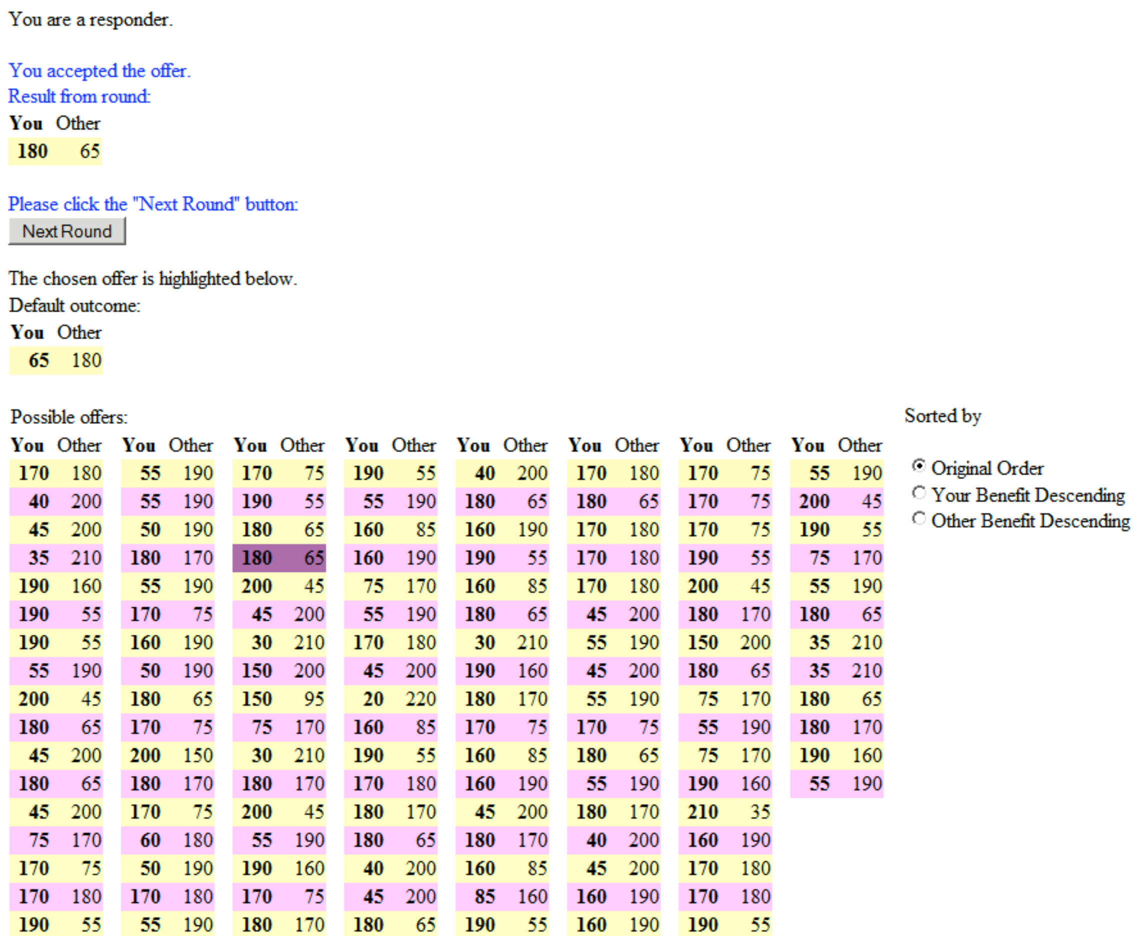

The result phase was not time-limited during the experiment; participants could wait as long as desired before hitting the "Next Round" button.[1]

[1]During the experiment, participants moved quickly through the rounds, but had there been a problem, it would have been useful to be able to pause at the end of a round.

3.3 Technical Details

The Table version of the game was designed to run in a client-server configuration with all pairings and UI generation done on a Linux blade server. Participants used clients running Firefox 1.0.7 with JavaScript enabled. In the experiment reported here, participants used the Macintosh OS X operating system.

Client-Server Interactions

Clients connect to the server over the HTTP protocol. All requests to the server are handled by the Apache web server, version 2. Client connections come in two forms: user-visible connections, the results of which are rendered in the web browser, and asynchronous requests made to the server (using a technology commonly referred to as AJAX). Due to a change to the way Firefox handles JavaScript in Firefox 1.5, Firefox version 1.0.7 or lower is required for the system to function robustly.[2]

Session Tracking

The benefit of the session tracking approach described here is that the user cannot use the back button or the URL navigation bar (hidden by default, but accessible via the keyboard) to cheat in any way—for instance, to attempt to go back to the initial allocator screen and change the offer made.

All user sessions are tracked on the client side using cookies. Whenever a browser visits the initial login page, the server checks the cookie to ensure that it was issued at the start of the current session. If it is not current, the cookie is cleared and the participant using the browser is allowed to log in.

Every page that the user can request contains code that ensures that the user is

[2] In particular, it was not possible to ensure that the back button would be disabled in Firefox 1.5. Changing the system to support Firefox 1.5 would be feasible but was not necessary to run the experiment as we controlled the participants' choice of browsers.

on the correct page. In particular, a JavaScript *onLoad* event takes place that makes an asynchronous request to the server with the URL of the current page and the tracking ID. The server looks up the tracking ID and ensures that the participant is at the correct URL. If the participant is not at the correct URL, then the JavaScript redirects the browser to the proper page.

To support this user-tracking model, all client code must inform the server what page the browser should be at before it redirects the browser to the new page. In addition, all code on the server that redirects a browser must follow the same protocol.

User-Visible Connections

The user-visible connection to the server provides the main game loop. Whenever a client connects to the server, the first step is to ensure that the client is located in the right place. The purpose of this check is twofold. First, it ensures a robust user experience. If the web browser closes—either by accident or through user error—simply reconnecting to server is sufficient to restore the participant to the proper place in the game. (Of course, if a crashed session is not restarted soon, the game may have continued to another phase of the round.)

Timing

All major timing operations take place on the client. Whenever timing is necessary, the client code contacts the server with an asynchronous request for the duration of the period to be timed. Although in certain circumstances this could affect the time limits imposed by perhaps a few hundred milliseconds, in practice this was not a noticeable problem and the server was so close to the clients that the lag time was only a fraction of a second, so other factors, such as server load when requesting a new page, would outweigh any imprecision in the timing code.

3.4 Experiment Conditions

Both conditions were played on a computer implementing some version of the CT system. One condition, the GUI, was designed to imitate the board game, whereas the Table condition showed normal form tables more reminiscent of an economics experiment.

GUI

Figure 3.4: Screenshot of GUI System

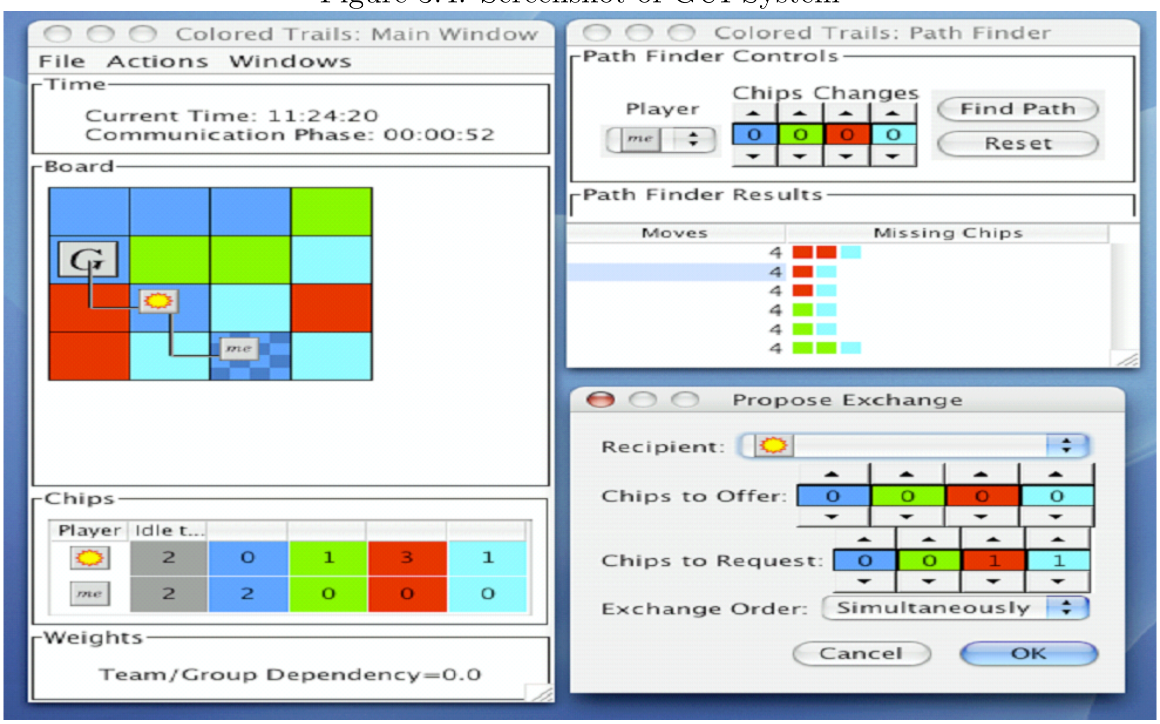

Figure 3.4 shows an example of the GUI system [2]. The GUI version provides a form of CT that matches the description of the game both in terms of possible offers and in terms of information available to players. Players see the board and their chips and have knowledge of the other player's chips. In order to assist players, a pathfinding tool is available that lists possible paths to the goal as well as the chips that are missing along the path route. This tool reduces the complexity of the

game for participants who might otherwise spend much of their time managing the cognitive demands of the task.

Normal Form Tables

A *normal form* table is a table that presents all of the possible outcomes of a negotiation as a list of values for each participant. The purpose of presenting data in normal form tables was to remove elements of the game context. As such, the interactions were not explained using any terminology from CT. Chip exchanges were referred to simply as possible offers, and no mention was made of the underlying game board, the rules of CT, or the scoring function. No information about the other player was available except what that player's outcome from each offer was. For instance, players did not know who was closer to the goal or whether an offer would allow one player to reach the goal.

The sorting functionality is an analog to the decision support tools in the GUI version of the game, providing a means of reducing what might otherwise be a heavy cognitive load in dealing with the clutter of a large number of available offers.

To match the possible exchanges available in the GUI condition and match the cognitive load with respect to the number of available choices, all offers were represented in the list of possible offers. As a result, offers with the same value, which would appear to be duplicate offers but which differed in the underlying chip exchanges, were presented.

3.5 Participants

GUI Condition

Participants for the GUI condition were recruited through the Harvard Gazette. A total of 16 participants participated in this portion of the study.

Table Condition

Participants for the Table condition were recruited from the Harvard and surrounding Cambridge community through house open lists, Craigslist, and text advertisements placed on Facebook.com. Potential participants were told only that they would be participating in a "Harvard-Sponsored Study on Interaction" paying between $18 and $30 and that they must not have participated in previous studies involving the CT interface.

Ultimately, a total of 22 participants were invited to participate in the Table condition. Of these 22, 16 arrived and participated in the study. All participants were given instruction in the game, including a tutorial round. The tutorial round was played with a single game replicated across all eight pairings. In order to allow all participants to see both halves of the interface, deliberators were asked to look over the shoulder of an allocator during the proposal phase, and allocators were asked to look over the shoulder of a deliberator during the response phase.

During the tutorial, the participants were informed that compensation was based only on the value resulting from each interaction.

Limited Interactions

During the experiment, we wished to avoid the possibility of participants' communicating with each other outside of offer proposals made inside the system. Consequently, during both conditions, participants were sequestered in two separate rooms of 8 participants, and participants never interacted with other participants in their own room through the system.

In addition, participants never played the same role twice with another participant, and they were never told who they were participating with except that they would never play the same person twice in the same role.

3.6 Games Used

In both conditions of the experiment, each participant played 12 rounds, for a total of 96 games, in addition to the tutorial round. These games were randomly generated and designed to be *interesting*. Interesting games are those for which at least one player cannot reach the goal without an exchange of chips.

To facilitate comparisons between the GUI and Table conditions, the same games were used in both. In addition, the pairings were the same in each condition. Each participant was assigned a number, $1-16$, and the same pairs of numbers played each other in the GUI condition as in the Table condition. Consequently, participants in the Table condition played the same games as their counterparts would have played in the GUI condition.

Chapter 4

Empirical Evaluation

4.1 Hypotheses

Because of the complexity of comparing offers with radically different chip sets, and because of the game context's promoting a focus on the task, as is usually the case in real life, rather than on maximizing a score (as in the Table condition), we expect that the GUI allocators will be more generous than the Table allocators.

Conversely, we expect that Table allocators will use the raw numbers to facilitate deciding between possible offers to ensure maximal gain, leading to more strategic, selfish play and a higher incidence of NE. It is not clear, a priori, whether this will result in lower gains for the responder. Table allocators may be capable of finding more PO offers that improve the benefit for both participants.

Due to the immediate and obvious complexity of the Table condition and the large number of offers presented to participants, we expect that the decision support tools will be heavily used by allocators in the Table condition. It is not clear whether Table deliberators will use the support tools or rely on making comparisons between offers. If they use the tools, this would imply an interest in checking the relative value of the offer in comparison to other offers and show an attention to social factors in addition to the benefit of the exchange.

Finally, we expect that GUI deliberators will see the allocators' reaching the goal as a reason for an otherwise less beneficial offer and that we will therefore see a correlation with the allocator's reaching the goal and whether the offer was accepted.

4.2 Experimental Results

Benefit of Offers

Table 4.1: Average Benefit of All Offers Proposed

	Allocator	Deliberator
GUI	82.3	47.6
Table	**98.0**	36.0

Table 4.1, which includes the default outcome in games for which the GUI allocator chose it, shows that Table allocators made offers that were on average more selfish, offering less benefit to deliberators (t-test $p < .1$) while significantly increasing their own gains (t-test $p < 0.01$).

Table 4.2: Average Benefit of Accepted Offers

	Allocator	Deliberator
GUI	79.5	56.4
Table	**94.6**	46.7

Although the GUI allocators showed a slight trend toward generosity, Table 4.2 shows that in both conditions, the accepted offers benefit deliberators 9–10 points of benefit over the average offer made and provide 3 fewer points of benefit to the allocator than the average offer. The difference in the deliberator benefit is not significant (t-test $p > .2$), but the difference in benefit to the allocator is significant (t-test $p < .05$).

Within-Game Shifts

Table 4.3: Which Condition Lead to More Favorable to Offers within Individual Games

	% Higher Benefit in Table	% Higher Benefit in GUI	% Same
Allocator	**62**	27	11
Deliberator	28	**60**	12

Table 4.3 shows that within individual games, there was a significant shift toward Table allocators' making offers better for themselves and less beneficial for Table deliberators than the offers made by GUI allocators (chi-square $p < 0.01$). Note that it is possible that within a particular game, both the allocator and deliberator could have fared better in one condition than in another—there is nothing strictly adversarial about the allocator's doing better in the table condition than in the GUI condition. For instance, in a game where the GUI allocator's offer was not PO, then one or both participants could benefit without the other participant's losing value.

Table 4.4: Which Condition Lead to More Favorable to Outcomes within Individual Games

	% Better in Table	% Better in GUI	% Same
Allocator	**43**	21	36
Deliberator	22	**42**	36

Table 4.4 indicates that within individual games, the results of games favored Table allocators significantly more frequently than they favored GUI allocators (chi-

square $p < 0.01$). The higher number of games with the same outcome than the same offer benefit is due to offers that were declined.

In sum, the within-game data suggest that the average increase in benefit to the allocator was due to the significant trend toward making better offers. On the other hand, the trend toward decreased benefit to deliberators in the Table condition can be explained by the same shift. That trend may not have reached the level of significance because the shift in benefit for the deliberator may have been smaller.

Acceptance Rates

Table 4.5: Acceptance Rate

	% Accepted	Total Offers
GUI	**73**	86
Table	**72**	96

Table 4.5 shows that between the conditions, in games in which offers were made, there was no difference in the rate of acceptance, although there was a difference in the total number of offers made as some GUI allocators chose to accept the default outcome.

The equivalent acceptance rates in light of the trend toward decreased average offer benefit for Table deliberators suggest that the Table deliberators may have had a lower benefit threshold for accepting offers.

Figure 4.1 demonstrates that for a lower minimum benefit, Table deliberators were more willing to accept offers. It also shows that GUI deliberators were less likely to accept offers with relatively large benefits than were Table deliberators.

Figure 4.1: Percentage of Offers Accepted Versus Minimum Offer Benefit to Deliberator

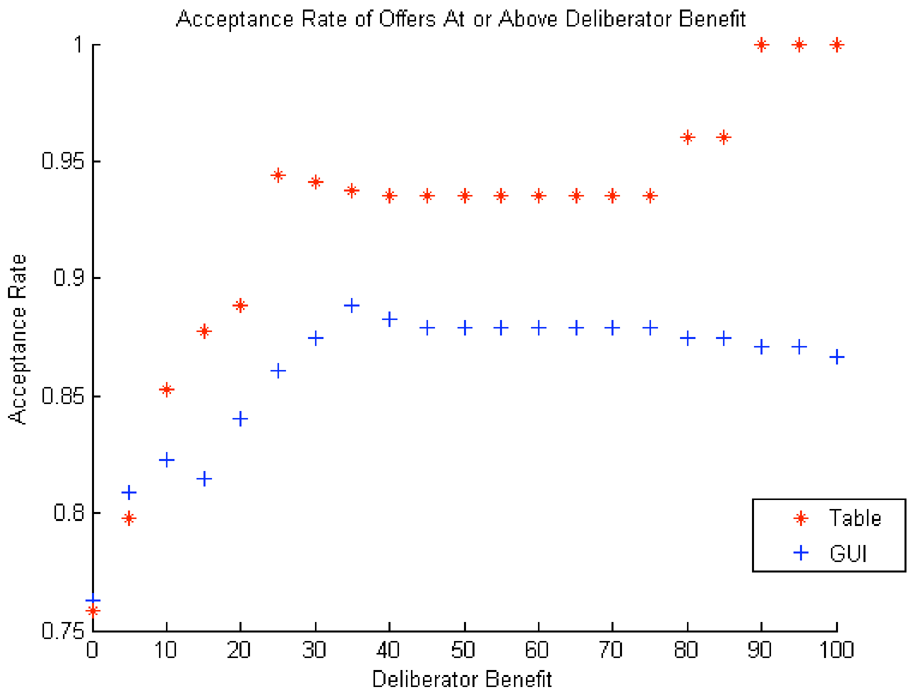

Allocator Strategy

Declining to Make Offers

Table 4.5 also indicates that Table allocators always made offers whereas the GUI allocators declined to make offers in 10 games. In every game used in the experiment, at least one possible offer benefited both allocator and deliberator. As a result, the allocator's taking the default outcome was never the game theoretic value-maximizing choice.

NE and PO Offers

Table 4.6: Percentage of NE and PO Offers in Comparison to All Offers

	% WNE	% SNE	% PO
GUI	20	15	86
Table	**48**	**59**	93

(Note that Table 4.6 includes data only for games where an offer other than the default outcome was chosen.)

Table 4.6 shows that the Table allocators made more game theoretically strategic offers—Table allocators made significantly more WNE and SNE offers (chi-square $p < 0.01$), but not significantly more PO offers (chi-square $p > .05$). The smaller percentage of WNE offers as compared to SNE offers is likely due to an overlap in WNE and SNE offers in some games but not in all games, and to allocators' favoring offers giving some benefit to deliberators in the cases where the WNE and SNE offers differed.

The difference between rates of PO offers is only significant when the 10 games in which no offer was made are accounted for (chi-square $p < 0.01$). Consequently, when GUI allocators chose to make an offer, there was no significant difference in the chance that the offer was PO.

The high percentage of PO offers in both conditions also suggests an explanation for the fact that an increase in the allocator's benefit in the Table condition met with a decrease in the deliberator's benefit.

The higher number of NE offers may account for the higher average Table allocator benefit and lower Table deliberator benefit when compared with the GUI condition.

Table 4.7: Acceptance Rate of NE Offers

	% WNE Accepted	% SNE Accepted
GUI	65	77
Table	70	77

Table 4.7 shows no significant difference in the rate of acceptance of WNE and SNE offers between conditions (chi-square $p > .05$). This suggests that between conditions,

the mere fact that an offer was an NE offer had little effect on participants.

Table 4.8: Average Deliberator Benefit from SNE Offers

	Deliberator Benefit
GUI	21.92
Table	**45.26**

Table 4.8 shows a significant increase in average benefit to deliberators from SNE offers (t-test $p < 0.01$). This implies that GUI allocators, in comparison to Table allocators, were not as strategic in picking *when* to make SNE offers. Because NE offers tend to be thought of as strategic, it can be somewhat disconcerting to think that they can sometimes be more beneficial to deliberators. Nevertheless, in a game where the two offers are $(100, 10)$ and $(110, 90)$, then the second offer is the NE offer in addition to being more beneficial to the deliberator. Not every scenario is a game where one person must make sacrifices for the other to do well, and it appears that Table allocators made NE offers at times when the deliberator would fare well.

NE Distance and Deliberator Benefit

Table 4.9: Average SNE Distance of All Offers

	Allocator	Deliberator
GUI	**-19.24**	**9.19**
Table	**-2.24**	**-1.72**

Table 4.9 demonstrates that for both allocators and deliberators, the difference between the average SNE distance was significant (t-test $p < .05$). In both conditions, the average allocator offer was valued slightly less than the SNE offer, but the

difference between the two conditions is striking. This result is largely because of the smaller number of SNE offers in the GUI condition. GUI deliberators received offers that were on average better than the SNE offer, whereas in the Table condition, the offers were on average slightly less generous than the SNE offer. One explanation is that, due to the high number of SNE offers, only a few outliers were necessary to tilt the average below 0. Taken as a whole, this provides further evidence of the more strategic play evidenced by Table allocators in Table 4.6. It also indicates that the GUI allocators' relative lack of SNE offers was not the result of the allocators' being more demanding than was reasonable (i.e., by asking for more than the SNE, which would have resulted in a positive SNE distance).

Figure 4.2: Deliberator SNE Distances Plotted against Deliberator's SNE Benefit (SNE Offers not Plotted)

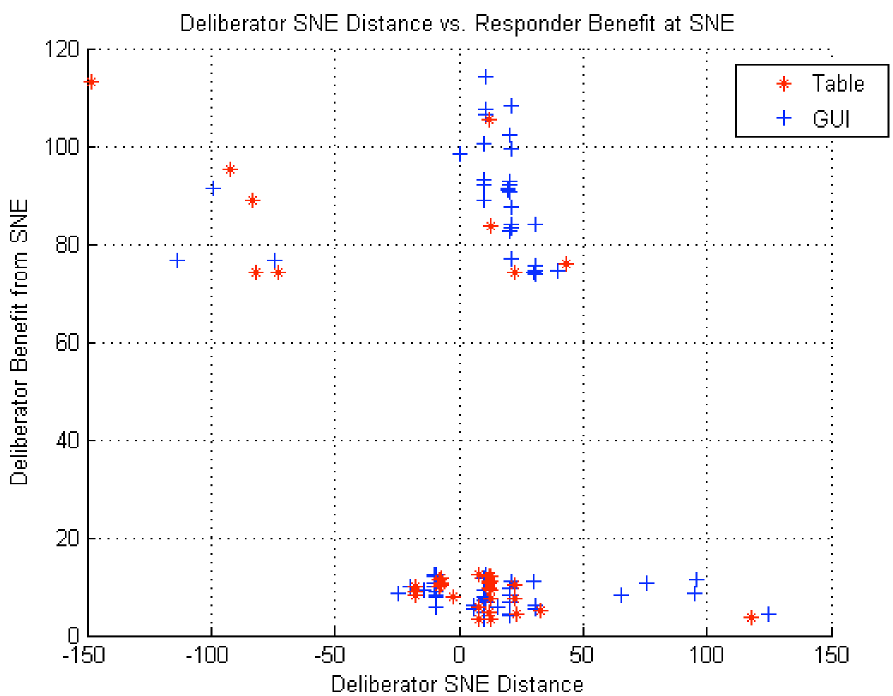

Figure 4.2 shows a striking cluster of games in the GUI condition where the GUI deliberator's SNE benefit was greater than 60 points and the deliberator's SNE distance was less than 50 points (note that in Figure 4.2, SNE offers are omitted).

The number of games from the Table condition in this cluster is significantly lower (chi-square $p < 0.01$). This difference indicates that when the SNE offer was quite good for the deliberator, GUI allocators more frequently increased the benefit of the offer to the deliberator above the SNE benefit than did the Table allocators.

It should be noted that the difference is not due to allocators' making the choice to allow the deliberator to reach the goal square. In particular, most of the offers in the quadrant described above allowed the deliberator to reach the goal square even with the SNE offer.

Table 4.10: Benefit to Deliberator As Compared to SNE in Table

Deliberator SNE Benefit	Overall	Including WNE Offers	Including SNE or Better
5	9.8	9.8	9.8
10	0.5	2.2	4.3

Table 4.10 breaks down Table allocators' offers into several cases: all offers, WNE offers or better for the deliberator, and SNE offers or better for the deliberator. When the deliberator's SNE benefit was 5 points, the allocator gave an average of 9.8 points more than the SNE to the deliberator. On the other hand, when the deliberator's SNE benefit was 10 points, there was a larger spread of offer benefits for the deliberator: on average, the deliberator received almost no benefit, 0.5 points, and even when including WNE offers, the benefit was only 2.2 points; on the other hand, when looking only at the SNE offers or better, the allocator gave 4.3 points to the deliberator. This table will be useful later when we look in more detail at the strategies used by allocators.

Correlations with Acceptance

Table 4.11: Correlations with Offer Acceptance

	GUI	Table
Deliberator Benefit	0.297	**0.418**
Allocator Reaches Goal	-0.07	-0.1754

Table 4.11 shows a larger correlation with deliberator benefit in the Table condition than in the GUI condition.

Surprisingly, Table 4.11 shows that the allocator's reaching the goal was not correlated with acceptance in either condition. Moreover, each of the individual correlations is suspect ($p > .05$).

Use of Decision Support Tools

Table 4.12: Decision Support Tool Use in Proposal Phase

	% Any	% by Allocator Benefit	Deliberator Benefit
Allocator	**100**	**83**	46
Deliberator	75	43	**56**

Table 4.12 shows that every single Table allocator, in every single game, used the decision support tool to sort the offers during the proposal phase. Notably, in different games, different sorting methods were sometimes chosen: Table allocators sorted by their own benefit more often than by the Table deliberator's benefit. Neither sort method was used in every game; the Table allocator sometimes chose one sort and sometimes chose the other. Allocators who focused on a strategy that required deliberators to receive a certain minimal benefit may have found it more useful to sort by benefit to the deliberator rather than to the allocator.

Table 4.12 also shows that deliberators sorted the offers during the proposal phase; it is not clear whether this sorting was done to gain information for later or simply to pass the time while waiting for an offer.

In contrast to their use of the sorting tool during the proposal phase, deliberators rarely used the sorting tools after an offer had been made. In fact, out of the 96 offers made, deliberators sorted the offers only 5 times (a total of 5.2% of the time) during the response phase.

Deliberators using a strategy focusing on the value of the exchange as compared to the value of the default outcome—rather than the value the exchange in relation to alternative exchanges—would not have needed to sort the outcomes.

Self-Reported Strategies

In the Table condition, we asked participants to describe their strategies. Participants indicated strategies similar to playing the NE strategy.

One participant wrote, "simply rank your benefits, and then choose the highest one that would still give the [deliberator] a higher amount in the end." Another wrote, "as the [allocator] I gave selected (sic) the option that gave me the most while still providing some benefit to the [deliberator], as [deliberator] I chose whichever option gave me more."

One participant indicated following a strategy of offering at least 20 points: "I offered the deliberator the highest value for me where they would get 20 more than the default so that they would accept; I accepted anything larger than the default." Here, there appears to be an element of game theory in trying to guess how a deliberator will react.

Another participant indicated a strategy that was close to, but not exactly, one of playing the NE: "I tried to get as much money for myself without letting my opponent gain too much for themselves. If they offered barely more than the default, but increased their percentage too much, I declined the increase."

4.3 Discussion

Results of Hypotheses

Clearly, Table allocators made more offers more beneficial to themselves than GUI allocators while showing a trend of giving somewhat less on average to the deliberators. It is not certain that the GUI allocators were more generous; in support of the hypothesis, the data show a trend toward increased offer values for deliberators, but the results were not quite at the level of statistical significance ($p < .1$ rather than $p < .05$). In both scenarios, allocators seemed to be reasonably adept at finding PO offers, except in the cases when GUI allocators chose the default outcome instead of proposing an offer.

As expected, the Table allocators offered many more NE. This may have come about in part due to the availability of tools that simplified reasoning about these games as participants in the Table condition used their support tools heavily except when choosing whether to accept an offer; this usage pattern implies a strong focus on the relative merits of offers for allocators but not as strong a focus on the relative merits of offers for deliberators. Unfortunately, due to a bug in the GUI system, we cannot make a direct comparison between support tool usage.

Finally, in neither condition was the allocator's reaching the goal square positively correlated with acceptance.[1] This may indicate that the GUI deliberators were not responsive to the allocator's having a good reason for making the offer. This could have been simply because the allocator could not draw attention to the rationale or because participants did not consider it a rationale at all. (Or perhaps they were competitive!)

[1] Participants in the Table condition did not know whether any offer would allow them to reach the goal square. Reaching the goal square is simply a function of whether the underlying offer would allow the participant to reach the goal square.

Game Difficulty

A second explanation for the lower GUI allocator benefit and fewer strategic offers is that certain games may have been more difficult to play in that allocators may have had more trouble finding beneficial exchanges. The fact that in 10 games, no offer was made when making an offer was the game theoretic rational thing to do supports this possibility. The alternative explanation, that the system caused problems for users, is not supported by the fact that not all of the games without offers took place in the early rounds—it was not until the 11th round that allocators stopped accepting the default outcome in the GUI version of the game.

That some games were difficult to evaluate suggests that GUI players may have simply had trouble computing the NE offers. On the other hand, the high percentage of PO offers in the GUI condition indicates that GUI allocators seem to have understood the game because finding PO offers is a non-trivial task that requires understanding the game fairly well. On the other hand, finding a PO outcome requires only manipulating one possible exchange to determine if any improvement can be made, whereas finding a NE offer requires comparing the benefit of an offer with the benefit of a large number of other offers.

A Viable Strategy?

Offers with a benefit of greater than roughly 20 points to the Table deliberator were almost always accepted. This suggests that a simple but powerful strategy is to follow is to offer the PO SNE whenever the SNE benefit to the deliberator is greater than a fixed threshold. Otherwise, make an offer more beneficial to the deliberator than the SNE so that the benefit exceeds the fixed threshold.

Certainly, not every Table interaction followed this strategy: there were a reasonable number of games where the WNE was proposed. But we know that at least one

Table allocator reported following precisely such a strategy for a threshold of 20—the participant wrote as much in the comments about the experiment. Moreover, the data suggest that a similar strategy may have been followed in some of the games for a threshold of 15. If we restrict the data set to games where an offer more generous than the WNE was played—i.e., when the Table allocator ensured that the deliberator gained some benefit—then, for games where the SNE's value to the deliberator was low—either 5 or 10 points (there were no SNEs where the deliberator would have gained 15 or 20 points)—the Table allocators made offers that would have benefitted the allocator sufficiently to reach a threshold of nearly 15 points.

As a result, it might be reasonable to look at the participants in the Table condition as consisting of two populations: those who played the WNE unless the WNE was not PO, and those who played a strategy that required the deliberator to gain some marginal benefit above a small fixed threshold.

In contrast, such a simple strategy categorization does not come close to matching the play of GUI allocators. As indicated in Figure 4.2, GUI allocators more often gave the deliberator more than the SNE even when the deliberator's SNE benefit would have been higher than 60. Moreover, Figure 4.1 shows that even for a large offer benefit, acceptance by a GUI deliberator was not assured.

Chapter 5

Conclusions

The results validate the CT metaphor as a way of gaining insight into a richer style of play approximating real world conditions, as it appears that Table allocators play more strategically and that Table deliberators follow a strategy more closely akin to one predicted by traditional game theory. Nevertheless, their behavior was more complex than simply making NE offers; Table allocators often chose a strategy that met a minimum threshold of benefit. By comparison, however, GUI allocators did not display signs of simply requiring an offer to pass a minimum threshold of benefit to the deliberator, often making non-SNE offers with a moderate positive deliberator SNE distance even when the deliberator's SNE benefit was relatively high.

These results support the claim that CT is a reasonable framework to use when developing computer systems that need to model more complex task domains in a way that more closely approximates real world conditions—where everything is not reduced to numerical outputs, but is instead a situation where resources are controlled and exchanged, and individual offers require effort to evaluate.

It was surprising that past a certain minimum threshold of benefit for the deliberator, offers tend not to be accepted more frequently. This suggests that it is useful for agents to be capable of strategizing in non-linear ways that allow for thresholds.

5.1 Caveats

The results presented here are not the end of the issue. In any situation involving large changes to the user experience, there may be many factors that contribute to the overall decision-making process. For instance, the decision-making tools were used quite frequently in the Table condition. Perhaps using the tools changes how participants reason about the system—focusing less mental capacity the other participant and more on the offers at hand. As we do not have complete data from the GUI condition, it is not possible to make any comparison with which to invalidate this possibility.

While the decision support tools were designed to reduce the cognitive load, it may be that they functioned in different ways in the two games. Finding NE offers in the GUI condition may simply have been a result of the fact that making direct comparisons between offers was not facilitated as much by the GUI pathfinder as it was by the sorting tools in the Table condition. If this was the case, then it would not be surprising that there were more NE offers in the Table condition. A follow-up study asking players to find NE offers in the GUI condition could ascertain whether this is a fundamental difficulty in the GUI condition.

Limitations of the CT Metaphor

Aside from the caveats about these results, we might expect that some interactions would be better modeled by normal form tables than the CT metaphor. Negotiations that involve primarily numerical data—such as in finance—may not be effectively modeled by CT.

Despite the fact that participants play differently in the CT system, we have not shown that they play as they would in real life. In the past, psychologists have found that in certain tests of logical reasoning such as the Wason Selection task, participants perform better when the task is presented with a specific context rather than as an abstract logic puzzle [9]. It may be worthwhile to attempt to set up a more concrete scenario than the CT metaphor provides, as individuals may have domain specific behaviors that are not captured by CT's more abstract interactions. One might expect, for instance, that modeling interactions between social classes within the CT framework would be difficult.

Bibliography

[1] C. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction.* Princeton University Press, 2003.

[2] Y. Gal, B. Grosz, S. Kraus, A. Pfeffer, and S. Shieber. Colored trails: A formalism for investigating decision-making in strategic environments, 2005.

[3] Y. Gal, A. Pfeffer, F. Marzo, and B. Grosz. Learning social preferences in games, 2004.

[4] 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.

[5] E. Horvitz, J. Breese, D. Heckerman, D. Hovel, and K. Rommelse. The lumiere project: Bayesian user modeling for inferring the goals and needs of software users. In *In Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, pages 256–265, Madison, WI, July 1998.

[6] E. Langer, A. Blank, and B. Chanowitz. The mindlessness of ostensibly thoughtful action: The role of "placebic" information in interpersonal interaction. *Journal of Personality and Social Psychology*, 36:635, 1978.

[7] B. McNeil, S. Pauker, H. Sox, and A. Tversky. On the elicitation of preferences for alternative therapies. In *New England Journal of Medicine 306*, pages 1259–1262, 1982.

[8] A. Tversky and D. Kahneman. The framing of decisions and the psychology of choice. *Science*, 211:453–458, January 1981.

[9] P. Wason and D. Shapiro. Natural and contrived experience in a reasoning problem. *Quarterly Journal of Experimental Psychology*, 23:63–71, February 1971.