2008-03-18

Satisficing Theory and Non-Cooperative Games

Matthew S. Nokleby
Brigham Young University - Provo

SATISFICING THEORY AND NON-COOPERATIVE GAMES

by

Matthew Nokleby

A thesis submitted to the faculty of

Brigham Young University

in partial fulfillment of the requirements for the degree of

Master of Science

Department of Electrical and Computer Engineering

Brigham Young University

April 2008

BRIGHAM YOUNG UNIVERSITY

GRADUATE COMMITTEE APPROVAL

of a thesis submitted by

Matthew Nokleby

This thesis has been read by each member of the following graduate committee and by majority vote has been found to be satisfactory.

Date	Wynn C. Stirling, Chair
Date	A. Lee Swindlehurst
Date	Randal W. Beard

BRIGHAM YOUNG UNIVERSITY

As chair of the candidate's graduate committee, I have read the thesis of Matthew Nokleby in its final form and have found that (1) its format, citations, and bibliographical style are consistent and acceptable and fulfill university and department style requirements; (2) its illustrative materials including figures, tables, and charts are in place; and (3) the final manuscript is satisfactory to the graduate committee and is ready for submission to the university library.

____
Date

____
Wynn C. Stirling
Chair, Graduate Committee

Accepted for the Department

____
Michael J. Wirthlin
Graduate Coordinator

Accepted for the College

____
Alan R. Parkinson
Dean, Ira A. Fulton College of
Engineering and Technology

ABSTRACT

SATISFICING THEORY AND NON-COOPERATIVE GAMES

Matthew Nokleby

Department of Electrical and Computer Engineering

Master of Science

Satisficing game theory is an alternative to traditional non-cooperative game theory which offers increased flexibility in modeling players' social interactions. However, satisficing players with conflicting attitudes may implement dysfunctional behaviors, leading to poor performance. In this thesis, we present two attempts to "bridge the gap" between satisficing and non-cooperative game theory. First, we present an evolutionary method by which players adapt their attitudes to increase raw payoff, allowing players to overcome dysfunction. We extend the Nash equilibrium concept to satisficing games, showing that the evolutionary method presented leads the players toward an equilibrium in their attitudes. Second, we introduce the conditional utility functions of satisficing theory into an otherwise traditional non-cooperative framework. While the conditional structure allows increased social flexibility in the players' behaviors, players maximize individual utility in the traditional sense, allowing us to apply the Nash equilibrium. We find that, by adjusting players' attitudes, we may alter the Nash equilibria that result.

ACKNOWLEDGMENTS

There are many people to thank. I am extremely indebted to Professor Wynn Stirling for his assistance throughout my college education. His early faith in me, which led to my involvement in his research, has made a tremendous impact on my academic life. Indeed, the content of this thesis is largely due to the numerous (and often lively) discussions held in his office over the past several years. I also owe a large debt to Professor A. Lee Swindlehurst, whose introductory class in signals and systems persuaded me to pursue graduate education. Professor Swindlehurst funded most of my graduate education, allowing me to apply game-theoretic ideas to his research in signal processing and wireless communications. I am also grateful to the remaining faculty at Brigham Young University—particularly Professors Randal Beard, Travis Oliphant, and Michael Rice—for providing an exceptional educational experience.

This thesis is dedicated to my father, Scott Nokleby, who has always encouraged me towards a lifetime of learning, love, and faith.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Background

Game theory is a mathematically rigorous study of decision making when individuals must interact with others. Individuals are modeled as "players," each of which must choose a strategy to enact. Typically, it is assumed that players are self-interested and therefore play strategies that maximize their payoff. Each player's payoff is quantified by a utility function, which is a function of the strategies played all of the players. This makes maximization a complex problem. Games can be divided into two main categories: in *cooperative* games, players may communicate and form binding agreements prior to play, while in *non-cooperative* games players do not communicate or collaborate.

While game theory is traditionally identified with economics or sociology literature, engineers and computer scientists have increasingly used it in the design of artificial multi-agent systems, such as air-traffic control [1, 2], cooperation in ad-hoc networks [3, 4], and resource allocation in wireless networks [5, 6]. Agents are modeled as players in a game. In non-cooperative games, players choose strategies that maximize individual payoff without regard for the payoffs obtained by others. While a player does consider the strategies played by others, it is only because their actions may affect its utility. Individual maximization, also called individual rationality, leads players to the well-known Nash equilibrium [7], where the players choose strategies such that no single player may increase its utility by changing strategies.

While the maximization of raw payoff is a simple and defensible criterion, this approach has several drawbacks. First, Nash equilibria often lead to highly inefficient results where players each earn low payoff [8]. Second, humans rarely act entirely out of self interest, particularly in social situations. As noted by Sen, "The *purely* economic man is indeed close to being a social moron. Economic theory has much been preoccupied with this rational fool decked in the glory of his one all-purpose preference ordering" [9, p. 15]. However, incorporating social considerations into players' utility functions is a difficult task.

Satisficing game theory [10] has recently been proposed as a means to address social factors in decision-making. It differs from classical game theory in two main ways. First, players' utility functions conform to a probabilistic syntax. Players' utilities are defined *conditionally* on the strategies preferred by other players. We then marginalize the conditional utilities as in probability theory. The "marginal" utility produced intrinsically incorporates the preferences of other players. Second, players posses two utility functions: one to quantify the benefits of a particular strategy and one to quantify its costs. A strategy for which the benefits outweigh the costs is considered "satisficing." The satisficing approach is particularly useful for engineering applications, where taking an action may consume resources such as fuel, power, or computation time. The flexibility afforded by these distinctions naturally incorporates social considerations into players' utilities, allowing for sophisticated behaviors in artificial systems [11, 12].

It should be noted, however, that satisficing game theory represents a significant departure from classical game theory. While the marginal utilities—which are typically used in choosing strategies—are functions of other players' preferences, they are not functions of the strategies played by other players. That is, satisficing players consider the preferences, but not the actions, of other players in making decisions. As a result, satisficing players may play strategies that are difficult to justify given other players' actions. Satisficing theory and classical game theory are two almost entirely separate decision theories with little common ground. The considerable body of knowledge relating to classical game theory therefore simply does not apply to satisficing players.

1.2 Contributions

In this thesis, we attempt to make meaningful connections between satisficing theory and non-cooperative game theory, allowing the social flexibility afforded by satisficing theory while avoiding incoherent behaviors. Further, these connections allow us to apply the theoretical results of traditional non-cooperative game theory.

First, we introduce an element of raw utility maximization into satisficing games by allowing the evolution of players' attitudes. In satisficing games, players' conditional utilities are often parameterized. These parameters, which we refer to as the players' *attitudes*, characterize the dependence of a player's utilities on those of the other players and greatly affect the strategies eventually played. We augment the satisficing game by introducing a classical utility representing the raw payoff to each player. This utility is a function of the strategies played by all of the players. Therefore, we consider the utility of exhibiting particular attitudes by a two-step mapping: attitudes map to the strategies played, which map to raw payoffs.

This leads us to define a non-cooperative game in the players' attitudes rather than in their strategies. We present a method by which players modify their attitudes according to the attitudes of other players and the raw payoff earned. In this method, which is based upon the replicator dynamics from evolutionary game theory, players modify their attitudes to improve individual payoffs. The dynamics leads the players towards a Nash equilibrium in their attitudes, which we call an *attitude equilibrium*. We apply this approach to the Ultimatum game, a simple two-player game.

Second, we introduce the highly useful concept of conditional utility functions into an otherwise standard non-cooperative framework. We retain the conditional utility structure proposed by satisficing theory, but shed the dual-utility concept. As in satisficing theory, players may condition their utilities on the utilities of others and are marginalized to form utility functions that incorporate others' utilities. However, we define the conditional utilities such that the marginal utilities are functions of the strategies chosen by all of the players. Thus, the marginal utilities are functions of other players' preferences *and* actions, rather than only on their actions (as in traditional non-cooperative theory) or their preferences (as in satisficing theory).

After solving for the players' marginal utilities, we can maximize individual utility in the sense of traditional non-cooperative games. We apply traditional solution concepts such as the Nash equilibrium to the proposed framework. We also examine the Ultimatum game under this framework. The results are highly encouraging, suggesting that players still exhibit social behaviors similar to that of satisficing theory. However, the powerful tools of classical game theory still apply.

1.3 Organization

In Chapter 2, we familiarize the reader with the basics of non-cooperative game theory. We define a classical game and formally define the Nash equilibrium. We also discuss extensive-form games, where players make their moves sequentially rather than simultaneously, and define the subgame perfect equilibrium. We also introduce the Ultimatum game, and discuss its difficulties under classical game theory.

In Chapter 3, we introduce satisficing game theory. We present the probabilistic structure of the players' utility functions and formally define the solution concepts that characterize satisficing strategies. We also review a satisficing formulation of the Ultimatum game, and show that players with incompatible attitudes may play incoherent strategies.

Chapter 4 details the attitude dynamics. First, we describe the augmented framework required and define the attitude equilibrium. We summarize the replicator dynamics and extend the dynamics to satisficing games, creating the attitude dynamics, which leads to an attitude equilibrium. Finally, we show the results of applying this method to the satisficing formulation of the Ultimatum game.

In Chapter 5, we apply the conditional utilities to non-cooperative games. After reviewing the theoretical underpinnings of classical utility functions, we show that conditional utilities can be reconciled with the requirements of classical utilities. We propose a conditional-utility model for the Ultimatum game, and show that the subgame perfect equilibrium changes with the players' attitudes.

Finally, Chapter 6 concludes the thesis and suggests areas of further research.

1.4 Notation

We use capital letters (A, B, etc.) to denote sets, and $|A|$ denotes the cardinality of set A. Boldfaced lowercase letters ($\mathbf{x}$, $\mathbf{y}$, etc.) refer to vectors, while boldfaced capitals ($\mathbf{F}$, $\mathbf{G}$, etc.) refer to matrices. Throughout most of this thesis, we do not distinguish between column and row vectors, as the distinction is unimportant. However, when we perform matrix-vector multiplication, vectors are explicitly defined as column vectors. Finally, since mass functions play a dual role as players' utilities, we use expanded notation for clarity. The function $p_X(x)$ represents the marginal probability mass function over random variable X, or $p_X(x) = Pr(X = x)$. The conditional probability mass function $p_{X|Y}(x|y)$ denotes the probability that $X = x$ under the assumption that the random variable Y takes on value y.

1.5 A Final Note

In game-theoretic texts, there is often question over which personal pronouns ("he" or "she") are most appropriate when referring to the players. Since the players typically are not intrinsically endowed with gender, this choice is occasionally motivated by an author's opinion on gender issues. In this thesis, we use the neutral pronoun "it" when referring to players. There are two reasons for this. First, it is the author's opinion that a technical thesis is not the place for a sociological statement, and this choice neatly sidesteps the issue. Second, it underscores the fact that, in the end, the "players" discussed herein are mathematical abstractions rather than sentient beings. While we often anthropomorphize them by referring to their preferences, attitudes, or self-interest, the players do not participate in any true reasoning. We must acknowledge that the mathematical models presented here and elsewhere—useful as they may be—represent at best a crude approximation of the intractably complex process of human decision-making.

Chapter 2

Non-cooperative Game Theory

In this chapter we present several basic concepts and results from classical non-cooperative game theory. We formally define a non-cooperative game and discuss the Nash equilibrium. We then present extensive-form games and the subgame perfect equilibrium. Finally, we present the Ultimatum game and discuss its difficulties under traditional game theory.

2.1 Formalization

We begin by defining a set of n players $I = \{1, 2, \ldots, n\}$. Each player $i \in I$ has a set S_i of *pure strategies*, or the actions which the player may implement. For simplicity, we restrict our attention to games with finite pure-strategy sets. A *pure-strategy profile* $\mathbf{s} = (s_1, s_2, \ldots, s_n)$ is an n-dimensional vector where each element s_i is a member of player i's pure-strategy set S_i. The set of pure-strategy profiles is the *pure-strategy space*, which is the Cartesian product $S = \times_{i=1}^{n} S_i$ of the individual players' pure-strategy sets. Finally, each player i has a scalar *payoff function* $\pi_i(\mathbf{s})$, which quantifies the utility obtained by player i from the implementation of the strategy profile $\mathbf{s} \in S$.

Mixed strategies are probability distributions over players' pure-strategy sets. If we denote $|S_i|$ by m_i, then We can characterize a mixed strategy as an m_i-dimensional vector $\mathbf{x}_i$ with the constraints $\sum_{k \in S_i} x_{ik} = 1$ and $x_{ik} \geq 0$. Each element x_{ik} represents the probability that player i plays pure strategy $k \in S_i$. Therefore, $\mathbf{x}_i$ describes a probability mass function. We assume, as is typical, that players' mixed strategies are probabilistically independent.

Since the probabilities that are the elements of $\mathbf{x}_i$ sum to one, each vector $\mathbf{x}_i \in \mathbb{R}^{m_i}$ represents an element of the m_i-dimensional *unit simplex*, which is defined as

$$\Delta_i = \left\{ \mathbf{x}_i \in \mathbb{R}^{m_i} : \sum_{k=1}^{m_i} x_{ik} = 1, x_{ik} \geq 0 \right\}. \tag{2.1}$$

The unit simplex is a convex set of normalized vectors upon which we impose a geometric interpretation. The *vertices* of the unit simplex are the elementary vectors $\mathbf{e}^1 = (1, 0, \ldots, 0)$, $\mathbf{e}^2 = (0, 1, 0, \cdots, 0)$, etc. These vertices represent the pure strategies, which we may regard as "degenerate" mixed strategies with all of the probability mass on a single element of S_i.

The *interior* of the unit simplex is the set of vectors which have only nonzero entries:

$$\text{int}(\Delta_i) = \{\mathbf{x}_i \in \Delta_i : x_{ik} > 0, \ \forall k\}. \tag{2.2}$$

In game-theoretic terms, vectors on the interior of the mixed-strategy simplex assign non-zero probability to each pure strategy in S_i.

The *mixed-strategy space* Θ is the Cartesian product of the players' mixed-strategy simplexes:

$$\Theta = \times_{i=1}^{n} \Delta_i. \tag{2.3}$$

A *mixed-strategy profile* is a vector of mixed strategies $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n) \in \Theta$ where each element $\mathbf{x}_i, i \in I$ is a mixed-strategy $\mathbf{x}_i \in \Delta_i$. Each entry x_{ik} of the mixed-strategy profile represents the probability with which player i will play pure strategy $k \in S_i$.

With the mixed-strategy profile well-defined, we may discuss players' *expected* utility. Since the players' mixed strategies are independent, the probability that a particular (pure-) strategy profile will occur is the product of the individual probabilities

$$Pr(\mathbf{s}) = \prod_{i=1}^{n} x_{is_i}. \tag{2.4}$$

The expected utility $u_i(\mathbf{x})$ to player i is the sum of the payoffs for each pure strategy multiplied by the probability that they occur:

$$u_i(\mathbf{x}) = \sum_{\mathbf{s} \in S} \pi_i(\mathbf{s}) Pr(\mathbf{s}) \tag{2.5}$$

$$= \sum_{\mathbf{s} \in S} \pi_i(\mathbf{s}) \prod_{i=1}^{n} x_{is_i}. \tag{2.6}$$

2.2 Nash Equilibria

Under traditional game theory, each player seeks to maximize expected utility by selecting an optimal mixed strategy. However, since each other player also attempts to maximize utility, it is difficult to define an optimal mixed strategy. In [7], Nash defines a *non-cooperative* solution. In non-cooperative games, players may not communicate or negotiate prior to play. However, it is typically assumed that players have full knowledge of the each other's payoff functions and that the players' intent to maximize utility is common knowledge.

Definition 2.1 *A* Nash equilibrium *is a mixed-strategy profile* $\mathbf{x}^* \in \Theta$ *such that*

$$u_i(\mathbf{x}_1^*, \ldots, \mathbf{x}_i^*, \ldots, \mathbf{x}_n^*) \geq u_i(\mathbf{x}_1^*, \ldots, \mathbf{x}_i', \ldots, \mathbf{x}_n^*) \tag{2.7}$$

for each $\mathbf{x}_i' \in \Delta_i$ *and for each* $i \in I$*.*

That is, a mixed-strategy profile is a Nash equilibrium if no player—acting alone—can improve its expected utility by choosing a different mixed strategy. If a player chooses a different mixed strategy than $\mathbf{x}_i^*$, and each other player continues to play its part of $\mathbf{x}^*$, then player i's expected utility can only decrease.

Perhaps the most famous example of the Nash equilibrium is the Prisoner's Dilemma [8]. In this scenario, two suspects are separately interrogated by police. They have committed two crimes: a minor offense and a major felony. The authorities have sufficient evidence to convict them of the lesser crime, but need testimony to convict either of them of the felony. The interrogators offer a deal: in return for testifying against the other, a prisoner will be forgiven for the lesser offense. However, even after agreeing to testify, a prisoner can be convicted of the greater crime if the other prisoner also testifies. Here, the set of players is $I = \{1, 2\}$, and each player's pure strategy set is $S_1 = S_2 =$ (Not Confess, Confess). We express the players' payoff functions in the payoff matrix in Table 2.1. Note that the entries in the matrix represent payoff rather than prison time, with higher payoff implying less prison time.

Table 2.1: Payoff matrix for the Prisoner's Dilemma.

	Player 1	
Player 2	Not Confess	Confess
Not Confess	(4, 4)	(5, 1)
Confess	(1, 5)	(2, 2)

From the payoff matrix, we can see that the unique Nash equilibrium is the pure-strategy profile $\mathbf{s} =$ (Confess, Confess), where each player testifies against the other. Even though both players would be better off under $\mathbf{s} =$ (Not Confess, Not Confess), self-interest drives the players towards an inefficient, albeit stable, equilibrium point.

Theorem 2.1 *Every game with a finite pure-strategy space has a Nash equilibrium.*

Proof: Here we reproduce the proof given by Nash in [7]. First, we define a family of continuous functions of $\mathbf{x} \in \Theta$ by

$$\phi_{is_i}(\mathbf{x}) = \max(0, u_i(\mathbf{x}_1, \ldots, \mathbf{x}_{i-1}, \mathbf{e}^{s_i}, \mathbf{x}_{i+1}, \ldots, \mathbf{x}_n) - u_i(\mathbf{x})). \tag{2.8}$$

In other words, ϕ_{is_i} is the increase in expected utility (if any) that player i gets if it changes from its current mixed strategy $\mathbf{x}_i$ to the pure strategy s_i, denoted by the degenerate mixed strategy $\mathbf{e}^{s_i}$.

Using ϕ_{is_i}, we define a mapping that updates the players' mixed-strategy profiles according to the relative payoffs. Define the mapping $T(\mathbf{x}) = \mathbf{x}'$ by

$$x'_{is_i} = \frac{x_{is_i} + \phi_{is_i}(\mathbf{x})}{1 + \sum_{s \in S_i} \phi_{is}(\mathbf{x})}, \ \forall i \in I. \tag{2.9}$$

Since T is a continuous mapping from Θ to itself, and since Θ can always be represented by a unit ball, T is guaranteed to have at least one fixed point by the Brouwer fixed point theorem.

Finally, we show that the fixed points of T correspond to Nash equilibria. If $\mathbf{x}^*$ is a fixed point of T, then $\phi_{is_i}(\mathbf{x}^*)$ must be zero for every i and s_i. Therefore, no player can improve its expected utility by changing to a pure strategy. Since no pure strategy can yield higher payoff, neither can any mixed strategy, making $\mathbf{x}^*$ a Nash equilibrium. Since a fixed point of T must exist, and since any fixed point of T must be a Nash equilibrium, a Nash equilibrium must exist. ■

Theorem 2.1 guarantees that an equilibrium exists in mixed strategies. While many games have pure-strategy equilibria, such equilibria do not exist in general, as shown by the familiar Rock-Paper-Scissors game. In this game, two players ($I = \{1, 2\}$) simultaneously show one of three signs ($S_1 = S_2 = \{\text{Rock, Paper, Scissors}\}$). Rock defeats Scissors, Paper defeats Rock, and Scissors defeats Paper. Otherwise, a tie results. These payoffs are formalized in Table 2.2.

Table 2.2: Payoff matrix for Rock-Paper-Scissors.

	Player 1		
Player 2	Rock	Paper	Scissors
Rock	(0, 0)	(−5, 5)	(5, −5)
Paper	(5, −5)	(0, 0)	(−5, 5)
Scissors	(−5, 5)	(5, −5)	(0, 0)

A careful look at Table 2.2 reveals that there is no pure-strategy equilibrium. Regardless of the pure strategies played, at least one player has incentive to switch to a different pure strategy. However, it is straightforward to show that if both players choose the mixed strategy $\mathbf{x}_1 = \mathbf{x}_2 = (1/3, 1/3, 1/3)$, neither player can improve expected utility by choosing another mixed strategy.

2.2.1 Extensive-Form Games

While every finite game has a Nash equilibrium, it may not be unique. This has motivated a number of equilibrium refinements, perhaps the most notable of which is the *subgame perfect* equilibrium [13]. To discuss the subgame perfect equilibrium, we must first consider games in *extensive form*. In extensive form games we explicitly define the order in which players make their moves and the information available to each player at each stage. We restrict our attention to games with *perfect information*, where each player has complete knowledge of the moves of each previous player. For a more comprehensive study of extensive-form games, see [14].

An extensive-form game will often have significantly different strategic properties. For example, consider a two-player game where player 1 makes its decision first, and player 2 makes its decision after observing player 1's action. In this case, player 2's pure strategy s_2 is specified as a function of the pure strategy played by player 1. That is, player 2 chooses among the functions $s_2 : S_1 \to S_2$ that map from player 1's pure-strategy set to player 2's pure-strategy set. We may think of these strategies as plans for all possible "contingencies": whatever player 1 does, player 2 has a response defined by the function $s_2(s_1)$.

Since player 1 makes its decision first, it may be able to "force" player 2 into making a desired decision. Or, alternatively, player 2 may be able to exploit the fact that player 1 has already made its move.[1]

Consider the game Chicken, where two players drive cars at each other to see who, if anyone, will swerve. The set of players is $I = \{1, 2\}$ and the players' pure strategy sets are $S_1 = S_2 = \{\text{Swerve, Not Swerve}\}$. We present a simple form of the game where each player makes its decision once, without the possibility of changing its mind afterward. We express the payoff functions in Table 2.3. If neither player

[1] In the Paper-Rock-Scissors game, sequential play significantly alters the game. No matter what player 1 does, player 2 can win the game.

Table 2.3: Payoff matrix for Chicken in normal form.

	Player 2	
Player 1	Swerve	Not Swerve
Swerve	$(0,\ 0)$	$(-10,\ 10)$
Not Swerve	$(10,\ -10)$	$(-100,\ -100)$

swerves, the cars collide and both players incur significant negative utility. If both players swerve, no utility is gained. However, if one player swerves and the other does not, the non-swerving player "wins" and garners positive utility while the losing player incurs negative utility. There are two pure-strategy Nash equilibria: $\mathbf{s} =$ (Swerve, Not Swerve) and $\mathbf{s} =$ (Not Swerve, Swerve). Unfortunately, in the normal-form game, it is difficult to say which equilibrium should result from rational players.

In an extensive-form version of Chicken, however, the analysis is more interesting. Extensive-form games are commonly displayed in a game tree, which displays graphically the order in which players move, the players' choices at each stage, and the payoffs to the players. The game tree for the extensive form of Chicken is shown in Figure 2.1. Player 1 moves first, and chooses whether or not to swerve. Player 2, after observing player 1's choice, also chooses whether or not to swerve. The resulting payoffs are the same as those from the payoff matrix in Table 2.3.

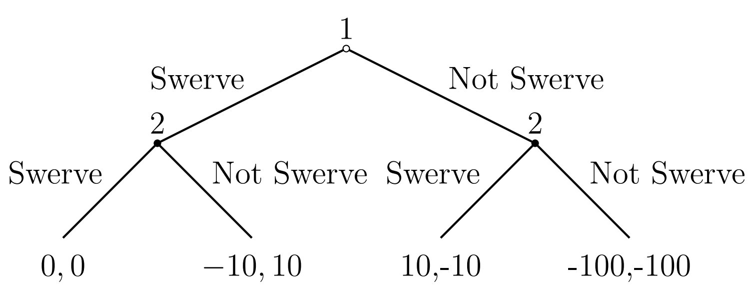

Figure 2.1: Game tree for Chicken in extensive form.

We may also represent the extensive-form game as a normal-form game and display a payoff matrix, as in Table 2.4. That is, rather than viewing the game sequentially where player 2 moves after player 1, we can view it as a game where the players choose strategies simultaneously, except that player 2's strategy is a function of player 1's strategy. That is, player 1 chooses a strategy, and player two chooses a "contingency" strategy that is a function of player 1's choice. Therefore, the distinction between this new payoff matrix and the original one given in Table 2.3 is that player 2's strategies are expressed as functions of player 1's actions, where (x, y) indicates that player 2 plays x if player 1 swerves and plays y if player 1 does not. For convenience, we abbreviate "Swerve" and "Not Swerve" as "S" and "NS," respectively.

Table 2.4: Payoff matrix for Chicken in extensive form.

	Player 2			
Player 1	(S, S)	(S, NS)	(NS, S)	(NS, NS)
Swerve	(0, 0)	(0, 0)	$(-10, 10)$	$(-10, 10)$
Not Swerve	(10, -10)	(-100, -100)	$(10, -10)$	$(-100, -100)$

There are three pure-strategy Nash equilibria, each of which correspond to a pure-strategy equilibrium from Table 2.3. The first two, which result in the same eventual actions, occur when player 1 does not swerve and player 2 plays (S, S) or (NS, S). In the third, player 1 swerves and player 2 plays (NS, NS) and refuses to swerve regardless of player 1's actions.

To choose between these equilibria, we define the subgame perfect equilibrium. In a *subgame*, a history of moves has already been played. In Chicken, for example, there are three subgames. First, there is the trivial subgame where no move has been played, and both players must choose a strategy. There are also two subgames where player 1 has chosen whether or not to swerve. In these subgames, player 2 plays alone. Note that subgames themselves are well-defined games since they have a set of players, a pure-strategy space, and payoff functions. We may therefore consider the Nash equilibria of these subgames, which motivates the subgame perfect equilibrium.

Definition 2.2 *A strategy profile is a* subgame perfect equilibrium *if it is a Nash equilibrium of every possible subgame.*

For Chicken, there is a unique subgame perfect equilibrium. Let's consider each subgame individually. Obviously, each of the three equilibria are still equilibria of the first subgame, since it is equivalent to the full game. Next, consider the subgame where player 1 has chosen to swerve. Since only player 2 chooses an action, the equilibria are simply the actions that maximize player 2's payoff. Here, s_2 = (NS, S) and s_2 = (NS, NS) are equilibria, so s_2 = (S, S) cannot be player 2's part of a subgame perfect equilibrium. Finally, consider the last subgame where player 1 has chosen not to swerve. Here, the equilibria are s_2 = (S, S) and s_2 = (NS, S), which eliminates s_2 = (NS, NS) as player 2's part of a the subgame perfect equilibrium.

Thus, in the unique subgame perfect equilibrium, player 1 chooses not to swerve, and player 2 chooses s_2 = (NS, S) and ends up swerving. In a common abuse of notation, we say that $\mathbf{s}$ = (Not Swerve, Swerve) is the subgame perfect equilibrium even though, strictly speaking, player 2's strategy is a function of player 1's choice. Note that only the "reactive" strategy, where player 2 does the opposite of player 1, survives as an equilibrium of each subgame. In this sense the subgame perfect equilibrium can be said to remove equilibria that result from "non-credible" threats. For example, the strategy (NS, NS) is not a credible threat because player 2, after observing that player 1 chooses not to swerve, will swerve to avoid a collision and maximize utility.

2.3 The Ultimatum Game

The Ultimatum Game has become a common example for illustrating the weakness of classical game theory as a model for human behavior. The game consists of two players ($I = \{1, 2\}$): the proposer (player 1) and the responder (player 2). The proposer and the responder must agree on the division of a dollar. The game is played sequentially: the proposer offers some fraction to the responder, who must decide whether or not to accept it. If the responder accepts the offer, they divide the dollar as proposed. Otherwise, each player receives nothing.

The proposer's strategy space is typically defined as the (uncountable) unit interval $[0, 1]$, which complicates analysis. Thus, we pattern our presentation after Gale et al. [15] and examine a simple quantization which captures the "heart" of the game while simplifying analysis. In this version, the proposer offers either a high or low fraction ($S_1 = \{\text{High, Low}\}$) to the responder, who again may choose to accept or reject ($S_2 = \{\text{Accept, Reject}\}$) the offer. The payoffs are shown in Table 2.5. We use h and l to denote the numerical value of the high and low fraction, respectively, with $h > l$.

Table 2.5: Payoff matrix for the Ultimatum minigame.

	Responder	
Proposer	Accept	Reject
High	$(1-h,\ h)$	$(0,\ 0)$
Low	$(1-l,\ l)$	$(0,\ 0)$

The unique Nash equilibrium (which is also subgame perfect) for this version of the Ultimatum game is $\mathbf{s} = (\text{Low, Accept})$. We can also show that, regardless of the quantization, there is a unique subgame perfect equilibrium in which the proposer offers the smallest possible nonzero fraction and the responder accepts it. However, this strategy is rarely implemented by human decision-makers. Real-life proposers are more likely to give fair offers, and responders often reject unfair offers, even though doing so reduces raw payoff [16, 17].

These findings suggest that players' utility functions are affected by considerations other than raw payoff. Several models have been proposed in recent years to account for this. Fehr and Schmidt [18] propose an "inequity aversion" model in which players exhibit a preference for strategy profiles that limit the difference—positive or negative—between its payoff and the payoff of the other players. If players are sufficiently averse to unfair outcomes, more "realistic" equilibria result. In Chapter 3, we discuss a satisficing model for the Ultimatum game presented by Stirling et. al [19] which also allows for fair outcomes.

Chapter 3

Satisficing Game Theory

3.1 Motivation

In satisficing game theory, players eschew the assumption of individual rationality which is fundamental to non-cooperative game theory. While the simple and reasonable assumption of rationality has given rise to a rich and successful theory, raw maximization may be *too* simple, particularly in describing social situations. As observed by Luce and Raiffa, "general game theory seems to be in part a sociological theory which does not include any sociological assumptions. . . it may be too much to ask that any sociology be derived from the single assumption of individual rationality," [20, p. 196]. Satisficing game theory provides an alternative to the traditional approach. It presents a more elaborate structure which may be more useful in modeling social behaviors. Players may directly concern themselves with the preferences of others and do not explicitly attempt to maximize utility.

To describe this, we alter the structure of the players' utility functions. First, each player possesses *two* utilities: one to characterize the benefits associated with taking an action, and one to characterize the costs. For a satisficing player, an action for which the benefits outweigh the costs is a "good enough" or satisficing decision and may be implemented. Second, the players' utility functions share a common syntax with probability mass functions, allowing probabilistic concepts such as conditioning and independence to be applied to players' preferences—albeit with a significantly different interpretation.

The use of probability mass functions to describe a player's preferences rather than a random phenomenon is an unusual one, and warrants further explanation. Stirling [21] provides a rigorous mathematical justification. It is shown that players' utilities must conform to a probabilistic syntax in order to satisfy several desirable axioms. We present a similar argument in Appendix A. Fortunately, however, the benefits of such a model may also be appreciated intuitively.

For two discrete random phenomena X and Y, where Y is dependent on X, we can express the probabilities for Y by the *conditional* mass function $p_{Y|X}(y|x)$. The conditional mass function gives hypothetical probabilities of Y: what would be the probability that $Y = y$ if we knew that X took on the value x? If we know the probabilities for $X = x$, we can compute the marginal mass function according to basic rules of probability theory: $p_Y(y) = \sum_x p_{Y|X}(y|x)p_X(x)$. The marginal mass function is a convex combination of the different conditional mass functions described by $p_{Y|X}(y|x)$, with the weights of the convex combination given by $p_X(x)$. The marginal probabilities for Y are therefore influenced—but not entirely dictated—by the probabilities of X.

Similarly, players' preferences may depend upon the preferences of others, allowing their utilities (which are called *social utilities*) to be expressed as conditional mass functions. The conditional mass functions allow for hypothetical expressions of utility: what would player 1's utilities be if player 2 unilaterally preferred a particular action? We may compute player 1's marginal utilities—which are the utilities used for decision-making—by summing the conditional utilities over player 2's actual preferences. This structure allows players to consider not simply what actions other players may prefer, but *how strong* the preferences for action are. Their utilities are influenced by others' preferences in a controlled manner which does not require them to discard their own preferences.

3.2 Formalization

As in the classical theory, we first define the set of n players $I = \{1, 2, \cdots, n\}$. We denote each player's pure-strategy set as U_i. A strategy profile is an n-dimensional vector $\mathbf{u} \in U$, where $U = \times_{i=1}^{n} U_i$ is the pure-strategy space. In satisficing game theory, we will concern ourselves only with pure strategies and not consider mixed strategies.

So far, the formalization of satisficing theory is identical to the classical formalization. However, in defining the players' utilities, satisficing theory is significantly different than the classical approach. First, each player possesses two social utilities. To describe these, we first define two "selves" or perspectives from which each player may consider his actions. The selecting self of player i considers actions strictly in terms of their associated benefits, while the rejecting self considers actions only in terms of the costs incurred in implementing them. These selves are described by the *selectability function* $p_{S_i}(u_i)$ and *rejectability function* $p_{R_i}(u_i)$, respectively. Note that, unlike traditional utility functions, social utilities are defined over player i's pure-strategy set U_i rather than over the pure-strategy space U.

Since social utilities are mass functions, they are normalized across the pure-strategy sets and therefore describe the *relative* benefit and cost associated with a pure strategy $u_i \in U$. They also provide players with a formal definition of good enough. A pure strategy is "good enough," or satisficing, if the relative benefit is at least as great as the relative cost. This permits us to define the *individually satisficing set* for player i as

$$\Sigma_i = \left\{ u \in U_i : p_{S_i}(u) \geq q p_{R_i}(u) \right\}, \tag{3.1}$$

where q is the *index of caution*. Typically, $q = 1$, but a player may adjust its definition of "good enough" by changing q. Setting $q \leq 1$ ensures that Σ_i is not empty. We may combine the players' individually satisficing sets by forming the *satisficing rectangle*

$\Re_{12\cdots n}$, which is defined as the Cartesian product

$$\Re_{12\cdots n} = \times_{i=1}^{n} \Sigma_i. \tag{3.2}$$

Any strategy profile $\mathbf{u} \in \Re_{12\cdots n}$ is simultaneously satisficing to each player in terms of its individual preferences.

When specifying the players' conditional utilities, it is convenient to express the relationship between players graphically. In probability theory, relationships between random variables are expressed in Bayesian networks. Similarly, in satisficing theory the relationship between players' utilities are expressed in *praxeic networks*.[1] The praxeic network consists of a directed acyclic graph (DAG), where the nodes are the selecting and rejecting perspectives of each player and the edges are the conditional utility functions. For example, consider the simple two-player community depicted in Figure 3.1. For each player, the rejecting preferences depend on the selecting preferences of the other player, while the selecting preferences are independent. For large communities where the interdependence structure is highly complex, we may employ established methods such as Pearl's Belief Propagation Algorithm [22] to analyze the community.

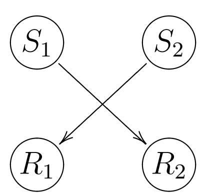

Figure 3.1: A simple praxeic network.

[1]The term *praxeic* is derived from *praxeology*, which refers to the study of human behavior.

In discussing the players' social utilities, we retain the terminology of probability theory. In the community from Figure 3.1, we refer to Player 1's *conditional* rejectability function, denoted $p_{R_1|S_2}(v_1|u_2)$. The conditional mass function expresses hypothetical utility: if Player 2's selecting preferences entirely favored strategy u_2, what would be Player 1's rejectability for v_1? As with probability mass functions, we may compute the *marginal* rejectability by summing over the conditionals, $p_{R_1}(v_1) = \sum_{u_2 \in U_2} p_{R_1|S_2}(v_1|u_2)p_{S_2}(u_2)$. The marginal utilities determine the individually satisficing sets and the satisficing rectangle. If a utility is independent (such as the selectability functions in this example), its marginal is expressed directly, without conditioning.

With the marginal and conditional utilities defined, we define the *interdependence function* $p_{S_1 \cdots S_n R_1 \cdots R_n}(u_1, \cdots, u_n; v_1, \cdots, v_n)$, which is the joint mass function of all players' selecting and rejecting preferences. By the chain rule of probability theory, the interdependence function for this example is $p_{S_1 S_2 R_1 R_2}(u_1, u_2; v_1, v_2) = p_{R_1|S_2}(v_1|u_2)p_{R_2|S_1}(v_2|u_1)p_{S_1}(u_1)p_{S_1}(u_1)$.

Satisficing games are characterized by the triple $(I, U, p_{S_1 \cdots S_n R_1 \cdots R_n})$. From this information, all necessary marginal utilities can be computed and the satisficing rectangle can be determined.

3.3 The Ultimatum Game

In this section we review the satisficing model of the Ultimatum game presented in [19]. It provides both an example of how players' conditional utilities are defined and an illustration of the strengths and weaknesses of the satisficing approach. As in Chapter 2, the set of players is $I = \{1, 2\}$ and the pure-strategy sets are $U_1 = \{\text{High, Low}\}$ and $U_2 = \{\text{Accept, Reject}\}$.

In the satisficing model, the players' behavior is governed by their *attitudes*. The proposer's attitudes are described by its *intemperance index* $\tau \in [0, 1]$. If $\tau = 1$, the proposer is exclusively concerned with maximizing payoff. As τ decreases, it is increasingly willing to compromise. The responder's attitudes are described by its *indignation index* $\delta \in [0, 1]$. If $\delta = 0$, it will accept any fraction offered. As δ increases, the responder becomes increasingly willing to forfeit its share in order to punish the proposer.

This is modeled explicitly by the players' social utilities. In this game, the selectability functions are associated with the benefit—the fraction of the dollar received. The rejectability functions express the risk of losing the entire dollar due to the responder's rejection. Since the proposer acts first, its utilities are specified unconditionally. The proposer's selectability is concerned with benefit, and the desire to keep the larger fraction for is determined by the intemperance index:

$$p_{S_1}(u_1) = \begin{cases} 1 - \tau, & \text{for } u_1 = \text{High} \\ \tau, & \text{for } u_1 = \text{Low} \end{cases}. \tag{3.3}$$

To avoid losing the entire dollar, the proposer considers the responder's indignation index in constructing the rejectability function:

$$p_{R_1}(v_1) = \begin{cases} \tau(1 - \delta), & \text{for } v_1 = \text{High} \\ 1 - \tau(1 - \delta), & \text{for } u_1 = \text{Low} \end{cases}. \tag{3.4}$$

The responder, who plays second, conditions its utility functions on those of the proposer. It wishes to maintain its fraction of the dollar while reserving the right to punish an intemperate proposer. The responder's conditional rejectability is

$$p_{S_2|S_1}(u_2|u_1) = \begin{cases} 1, & \text{for } u_2 = \text{Accept}|u_1 = \text{High} \\ 0, & \text{for } u_2 = \text{Reject}|u_1 = \text{High} \\ 1-\delta, & \text{for } u_2 = \text{Accept}|u_1 = \text{Low} \\ \delta, & \text{for } u_2 = \text{Reject}|u_1 = \text{Low} \end{cases}. \tag{3.5}$$

If the proposer unilaterally favors the high offer ($\tau = 0$), the responder will entirely prefer to accept the offer. However, if the proposer favors the low offer ($\tau = 1$), the responder prefers to reject the offer according to its indignation index δ. The conditional rejectability encodes the same preferences, and is given by

$$p_{R_2|S_1}(v_2|u_1) = \begin{cases} 0, & \text{for } u_2 = \text{Accept}|u_1 = \text{High} \\ 1, & \text{for } u_2 = \text{Reject}|u_1 = \text{High} \\ \delta, & \text{for } u_2 = \text{Accept}|u_1 = \text{Low} \\ 1-\delta, & \text{for } u_2 = \text{Reject}|u_1 = \text{Low} \end{cases}. \tag{3.6}$$

Summing over the conditional mass functions, the responder's marginal utilities are

$$p_{S_2}(u_2) = \begin{cases} 1-\tau\delta, & \text{for } u_2 = \text{Accept} \\ \tau\delta, & \text{for } u_2 = \text{Reject} \end{cases}, \tag{3.7}$$

$$p_{R_2}(v_2) = \begin{cases} \tau\delta, & \text{for } u_2 = \text{Accept} \\ 1-\tau\delta, & \text{for } u_2 = \text{Reject} \end{cases}. \tag{3.8}$$

The interdependence function for the Ultimatum game is constructed according to the chain rule:

$$p_{S_1S_2R_1R_2}(u_1, u_2; v_1, v_2) = p_{S_2|S_1}(u_2|u_1)p_{R_2|S_1}(v_2|u_1)p_{S_1}(u_1)p_{R_1}(v_1). \tag{3.9}$$

With the players' utility functions are defined, we examine their actions according to the satisficing rectangle. Recall that a pure strategy u_i is individually satisficing for player i if $p_{S_i}(u_i) \geq qp_{R_i}(u_i)$. In Figure 3.3 we set $q = 1$ and show the satisficing rectangle as functions of τ and δ. Four possibilities result depending on the players' attitudes. In first region, the proposer is sufficiently greedy and the responder sufficiently conciliatory that $\mathbf{u}$ = (Low, Accept) is the unique member of the satisficing rectangle. If τ is lower and/or δ is higher, the players implement $\mathbf{u}$ = (High, Accept). In the remaining two regions, the responder is sufficiently indignant that the offers are rejected.

Figure 3.3 shows a few interesting properties that are typical of satisficing games. Consider a responder with an indignation index of, say, $\delta = 0.6$. Notice that the responder's actions are not simply (or even primarily) based on the offer proposed. For $\tau = 0.75$, the low offer is accepted. However, if the proposer's intemperance index increases much higher, the proposer refuses. It accepts an low offer from a somewhat moderate proposer and refuses it from an intemperate one. In this highly sophisticated behavior, the responder punishes the proposer not for its actions, but for its *attitudes*. Such social behaviors may be desirable in the synthesis of artificial decision-makers, and are difficult to model under non-cooperative game theory.

However, this framework also allows for undesirable behavior. Consider the region where the responder rejects the high fraction. The intemperance index is sufficiently high that the responder punishes even a moderately intemperate proposer. Such dysfunctional behavior is somewhat common in satisficing games and is

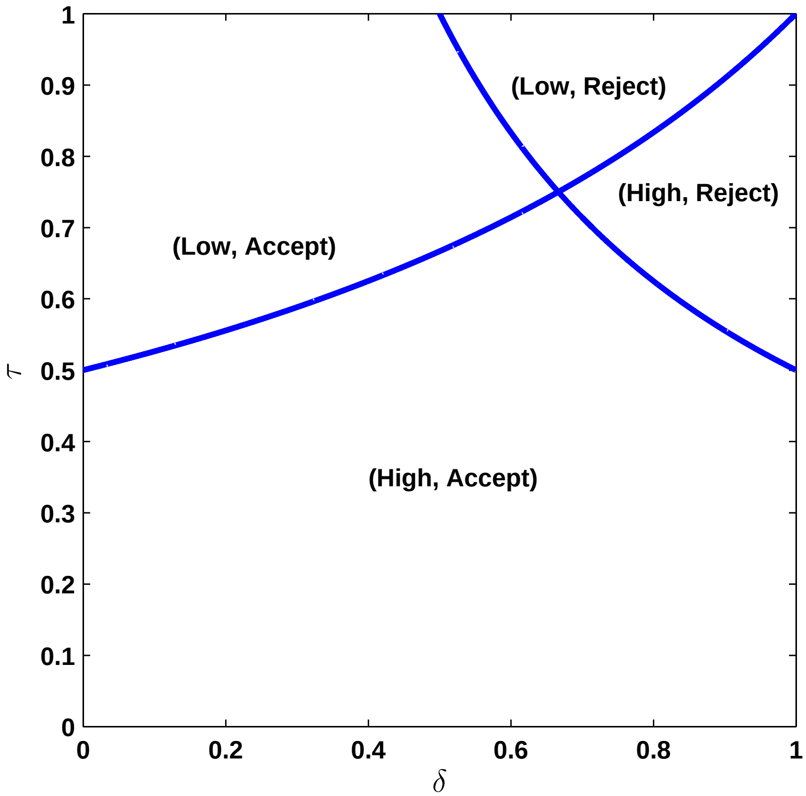

Figure 3.2: (τ, δ) regions for the satisficing rectangles.

a consequence of the structure of the utilities: players' utilities depend on the others' attitudes rather than the strategies they implement. Note that this poor performance is quite distinct from the difficulties under non-cooperative game theory. In the non-cooperative solution, players' narrow focus on payoff allows the proposer to exploit the responder. Here, an overly indignant responder ignores payoff to reject even a high offer.

We hasten to note that dysfunctional behavior is not a failure *per se* of the satisficing model. Dysfunctional societies do exist in practice, and the satisficing model simply suggests that players with incompatible attitudes (in this case a highly indignant responder with a moderately intemperate proposer) may act incoherently.

However, in designing artificial societies, we typically prefer to avoid incoherent behaviors, sociologically justifiable or not.

Chapter 4

Attitude Adaptation in Satisficing Games

We discuss how satisficing players may adjust their attitudes to increase their payoff and overcome dysfunctional behaviors such as those shown in Chapter 3. We augment the satisficing game by introducing a space of attitudes that each player may exhibit and a classical utility defined for each player over the pure strategy space. We define the attitude equilibrium, which is a Nash equilibrium in players' attitudes. After discussing classical replicator dynamics, we show how the dynamics can be used to adapt the attitudes of satisficing players. We then show the results of applying the attitude dynamics to the Ultimatum game.

4.1 Attitude Equilibria

To introduce the attitude equilibrium and the attitude dynamics, we must first embellish the structure of the satisficing game. To do this, we endow each player with a classical (von Neumman-Morgenstern) utility function which is based solely on the pure-strategy profile that the players implement.

Definition 4.1 *An* augmented satisficing game *is a tuple* $(I, U, p_{S_1 \cdots S_n R_1 \cdots R_n}, A, \boldsymbol{\pi}(\mathbf{u}))$. *The first three elements are the set of players, the pure-strategy space, and interdependence function as normal. Additionally, we introduce the product attitude space* $A = \times_{i=1}^{n} A_i$ *containing the attitudes that the players may exhibit and* $\boldsymbol{\pi}(\mathbf{u})$*, a vector payoff function which describes the raw payoff to each player for implementing the pure strategy profile* $\mathbf{u} \in U$*.*

Not all satisficing games may be augmented. To augment a satisficing game, the players' attitudes must be specified as distinct parameters in the social utilities.

Further, we must be able to construct a "raw" payoff function that is separate from the social utilities. Fortunately, the extension is straightforward for the Ultimatum game. The players' attitudes are the intemperance and indignation indices τ and δ, yielding the pure-attitude space $A = [0, 1] \times [0, 1]$. The payoff function $\boldsymbol{\pi}(\mathbf{u})$ is described by the payoff matrix in Table 2.5.

The augmented satisficing game describes a two-step mapping from attitudes to payoffs. The social utilities—determined by the interdependence function—map the players' attitudes to strategy profiles.[1] The payoff function then maps the strategy profile to raw payoffs.

Thus, in an augmented satisficing game, we may evaluate the raw utility of possessing particular attitudes. To simplify notation, we will occasionally refer to $\boldsymbol{\pi}(\mathbf{a})$, the payoff to the players for implementing the strategy profile determined by the attitude profile $\mathbf{a} \in A$. That is, we may think of an augmented satisficing game as a classical non-cooperative game where players' payoffs are determined by the attitudes they adopt, rather than the strategies they implement.

We may also discuss *mixed attitudes* which, by analogy, are probability distributions over the attitudes the players exhibit. Denoting $|U_i|$ by k_i, the mixed attitude of player i is given by a (normalized) k_i-dimensional vector $\mathbf{c}_i$. The discussion of mixed strategies Chapter 2 applies to mixed attitudes. We assume that players' mixed attitudes are probabilistically independent of each other. We define player i's mixed-attitude simplex Δ_i^a and may analogously discuss its vertices and interior. The mixed-attitude space is the Cartesian product $\Theta^a = \times_{i=1}^n \Delta_i^a$, which contains mixed-attitude profiles $\mathbf{c} = (\mathbf{c}_1, \mathbf{c}_2, \ldots, \mathbf{c}_n)$.

[1] Here, we have glossed over the fact that the satisficing rectangle may contain multiple pure-strategy profiles. For simplicity, we will assume that, if necessary, the players employ a tie-breaking mechanism to select a unique pure-strategy profile.

The expected utility $u_i(\mathbf{c})$ to player i when the players exhibit the mixed-attitude profile $\mathbf{c} \in \Theta^a$ is essentially the same as (2.5):

$$u_i(\mathbf{c}) = \sum_{\mathbf{a}\in A} \pi_i(\mathbf{a}) \prod_{i=1}^{n} c_{ia_i}. \quad (4.1)$$

Players may now consider *changing* their attitudes if they result in poor expected utility. This concept provides the motivation for the attitude equilibrium.

Definition 4.2 *An* attitude equilibrium *is a mixed-attitude profile* $\mathbf{c}^* \in \Theta^a$ *such that*

$$u_i(\mathbf{c}_1^*, \ldots, \mathbf{c}_i^*, \ldots, \mathbf{c}_n^*) \geq u_i(\mathbf{c}_1^*, \ldots, \mathbf{c}_i', \ldots, \mathbf{c}_n^*) \quad (4.2)$$

for each $\mathbf{c}_i' \in \Delta_i^a$ *and for each* $i \in I$.

The definition for the attitude equilibrium is almost identical to that of the Nash equilibrium. In fact, we may say that an attitude equilibrium is an equilibrium in players' attitudes, rather than in their strategies. Because of the analogy between the attitude equilibrium and the Nash equilibrium, many theoretical results apply.

Theorem 4.1 *Every game with a finite pure-attitude space has an attitude equilibrium.*

Proof: Since we can view any augmented satisficing game as a classical non-cooperative game where the pure-attitude space takes the place of the pure-strategy space, this result follows directly from Theorem 2.1. ■

As with the Nash equilibrium, the attitude equilibrium may exist only in mixed attitudes.

For the Ultimatum game, even though the attitude spaces are continuous, it is straightforward to show that attitude equilibria exist and that they exist in pure attitudes. In Figure 4.1, the attitude equilibria are the shaded regions. If the players' attitude vector lies in these regions, there is no incentive for either player to alter its attitudes. Consider the shaded portion of the (High, Accept) region. The responder receives maximum payoff, and therefore has no reason to deviate. Similarly, the proposer cannot improve its payoff by changing τ. While the proposer does not earn maximum payoff, changing τ can only drive the responder to reject the offer, resulting in lower payoff. Similarly, in the shaded (Low, Accept) region, altering δ can only result in the offer being rejected. For any other pure attitude profile in A, at least one player stands to increase payoff by changing attitudes.

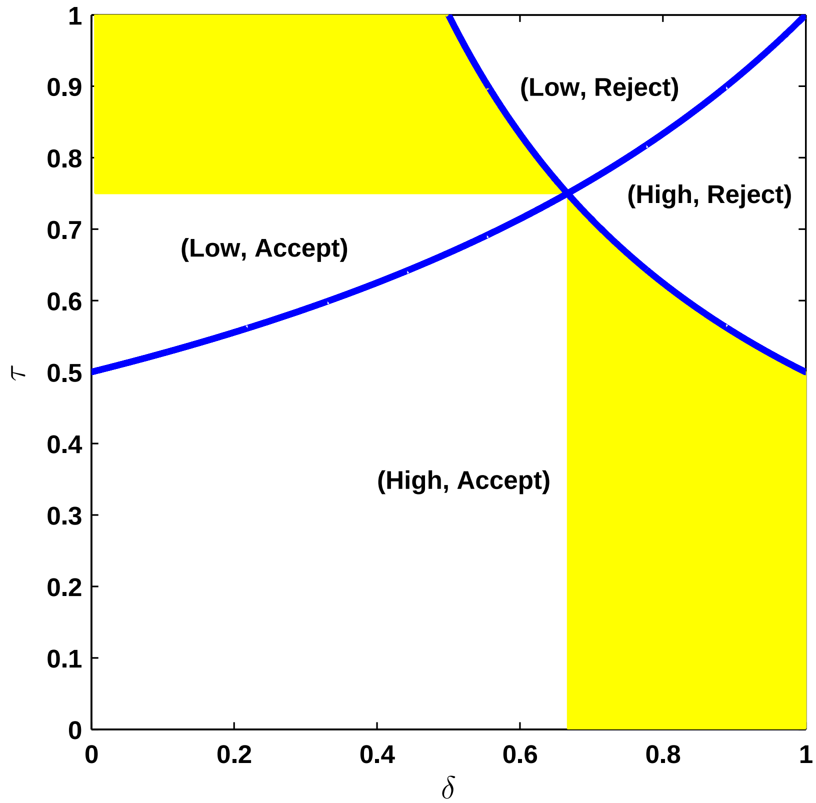

Figure 4.1: Attitude equilibria for the Ultimatum game.

The players' attitude equilibria result in the acceptance of either the high or low offer. By contrast, the traditional Nash equilibrium results in the acceptance of only the low offer. The attitude equilibrium provides a useful juxtaposition of satisficing theory and individual rationality: we retain the social structure which allows the high fraction to be offered, but eliminate the possibility that conflicting attitudes will result in the forfeiture of the entire dollar.

Unfortunately, the attitude equilibrium concept does not tell us which equilibrium these players will adopt. It simply says that if their attitudes lie in an equilibrium region, neither player has incentive to deviate. Therefore, we turn to evolutionary mechanisms to explore which equilibrium will result under different initial conditions.

4.2 Replicator Dynamics

4.2.1 Derivation

In this section we detail the classical replicator dynamics from evolutionary game theory, described in [23]. In its simplest form the replicator dynamics involves a two-player, symmetric[2] game described by a single pure-strategy set $S_1 = S_2$ and payoff function $\pi(s_1, s_2)$. Consider a population of players that are preprogrammed to play some pure strategy $i \in S_1$. That is, regardless of the payoff, each player in the population plays the strategy to which it is programmed.

Let $p_i(t)$ represent the number of players playing strategy i at time t, and let $p(t) = \sum_{i \in S_1} p_i(t)$ represent the size of the population at time t. Finally, define the *population share* for pure strategy i at t as $x_i(t) = p_i(t)/p(t)$. The population share $x_i(t)$ can be interpreted probabilistically: if we draw a member of the population randomly at time t, it will be programmed to play i with probability $x_i(t)$. We define the vector of population shares $\mathbf{x}(t) = (x_1(t), x_2(t), \ldots)$. We may also interpret $\mathbf{x}(t)$ probabilistically as a member of the mixed-strategy simplex Δ.

In continuous time, the members of the population are randomly paired up, where they play the game in question and earn their payoff. In the replicator dynamics, players asexually reproduce according to their payoffs. That is, the number

[2]By "symmetric," we mean that the two players are equivalent in terms of their pure strategy spaces and their payoff functions. In a symmetric game, we can exchange the roles of player 1 and player 2 without effect.

of offspring a player has is equal to its payoff. We assume that each offspring is programmed to the same pure strategy as its parent. If we let the population become infinitely large, we can invoke the law of large numbers, and the evolution of $p_i(t)$ is given by a deterministic differential equation:

$$\dot{p}_i(t) = \left[u(\mathbf{e}^i, \mathbf{x}(t))\right] p_i(t) = p_i(t) \sum_{j \in S_1} \pi(i,j) x_j(t). \tag{4.3}$$

In other words, the rate of increase in $p_i(t)$ is given by the average utility of playing pure strategy i (denoted by the elementary vector $\mathbf{e}^i$) against a member of the population (denoted by the "mixed strategy" $\mathbf{x}(t)$). By the linearity of differentiation, the total population size grows according to

$$\dot{p}(t) = \sum_{i \in S_1} \dot{p}_i(t) = p(t) \sum_{i \in S_1} \sum_{j \in S_1} \pi(i,j) x_i x_j = \left[u(\mathbf{x}(t), \mathbf{x}(t))\right] p(t). \tag{4.4}$$

Of course, as the population size approaches infinity, it becomes more fruitful to examine the evolution of the population shares, which remain finite. We start by rearranging the definition of $x_i(t)$:

$$x_i(t) = \frac{p_i(t)}{p(t)}, \tag{4.5}$$

$$p(t) x_i(t) = p_i(t). \tag{4.6}$$

Next, we differentiate both sides:

$$\frac{d}{dt}\left\{p(t) x_i(t)\right\} = \frac{d}{dt}\left\{p_i(t)\right\}, \tag{4.7}$$

$$\dot{p}(t) x_i(t) + \dot{x}_i(t) p(t) = \dot{p}_i(t), \tag{4.8}$$

$$\dot{x}_i(t) p(t) = \dot{p}_i(t) - \dot{p}(t) x_i(t). \tag{4.9}$$

Finally, we apply (4.3) and (4.4), yielding

$$\dot{x}_i p(t) = \left[u(\mathbf{e}^i, \mathbf{x}(t))\right] p_i(t) - \left[u(\mathbf{x}(t), \mathbf{x}(t))\right] p(t) x_i(t), \tag{4.10}$$

$$\dot{x}_i(t) = \left[u(\mathbf{e}^i, \mathbf{x}(t))\right] x_i(t) - \left[u(\mathbf{x}(t), \mathbf{x}(t))\right] x_i(t), \tag{4.11}$$

$$\dot{x}_i(t) = \left[u(\mathbf{e}^i, \mathbf{x}(t)) - u(\mathbf{x}(t), \mathbf{x}(t))\right] x_i(t). \tag{4.12}$$

A population share's rate of increase depends on its relative expected utility. If a the expected utility of a pure strategy is more successful than average, its population share grows, whereas less successful strategies tend to die out.

Equation (4.12) gives a complete characterization of the evolution of the infinitely large population. The evolution of the population shares is given by system of first-order ordinary differential equations, inducing a unique solution trajectory $\xi(t, \mathbf{x}(0))$ through the initial conditions $\mathbf{x}(0)$. It is shown in [23] that as long as all pure strategies are represented in the initial conditions ($\mathbf{x}(0) \in \text{int}(\Delta)$), any steady-state of the dynamics is a Nash equilibrium.[3] Note that this does not imply that a steady state exists.[4] Rather, if a steady-state exists under well-behaved initial conditions, it must be a Nash equilibrium.

4.2.2 Multipopulation Dynamics

The multipopulation replicator dynamics describes evolution in asymmetric games where players are selected from separate player populations. In the Ultimatum game, for example, there are two player populations: a population of proposers and a population of responders. We therefore have two distinct pure-strategy sets S_1 and S_2 as well as two payoff functions $\pi_1(s_1, s_2)$ and $\pi_2(s_1, s_2)$. In multipopulation

[3] That is, the mixed-strategy profile $\mathbf{x} = (\mathbf{x}^*, \mathbf{x}^*)$, where each player uses $\mathbf{x}^*$ as its mixed strategy, is a Nash equilibrium.

[4] A common example is Rock-Paper-Scissors, described in Chapter 2. Here, the solution to the replicator dynamics "orbits" around the mixed-strategy equilibrium, but does not approach steady-state.

dynamics there are arbitrarily many populations, but we will restrict our attention to the two-population case. Define a vector of population shares for each population: $\mathbf{x}(t)$ describes player population 1, and $\mathbf{y}(t)$ describes player population 2.

At each time t a player is drawn from each population to play the game, earn payoffs, and reproduce. The standard two-player replicator dynamics are given by a system of differential equations similar to (4.12):

$$\dot{x}_i(t) = \left[u_1(\mathbf{e}^i, \mathbf{y}(t)) - u_1(\mathbf{x}(t), \mathbf{y}(t))\right] x_i, \tag{4.13}$$

$$\dot{y}_j(t) = \left[u_2(\mathbf{e}^j, \mathbf{x}(t)) - u_2(\mathbf{y}(t), \mathbf{x}(t))\right] y_j. \tag{4.14}$$

While this system of equations preserves the intuition from (4.12), we have stated them without justification. In [23], Weibull notes that this is just one of many possible multi-population models. This dynamic model glosses over issues peculiar to the multi-population system such as relative population size and the resulting relative velocities with which the populations evolve. However, it is a convenient model because of its simplicity and analogy with the single-population models. Many of the theoretical results from the single-population dynamics hold; namely, the steady state of the dynamics, given well-behaved initial conditions as before, is a Nash equilibrium.

4.3 Attitude Dynamics

The replicator dynamics is traditionally used to describe the evolution of the distributions of large populations. However, it can also be used as a deliberation process [24] where each player updates its mixed strategy according to the mixed strategy of the other player. Here, we interpret $\mathbf{x}(t)$ and $\mathbf{y}(t)$ as the mixed strategies of player 1 and player 2. Each player updates its mixed strategy according to (4.13) and (4.14).

To extend the dynamics to the satisficing case, we operate the replicator dynamics on the players' attitudes rather than on the strategies they implement. Define the vectors $\mathbf{c}(t) \in \Delta_1^a$ and $\mathbf{d}(t) \in \Delta_2^a$, which define the mixed attitudes for each player at time t. We require that both players have finite attitude spaces so that $\mathbf{c}(t)$ and $\mathbf{d}(t)$ are finite-dimensional. The dynamics allows the players to alter the probability with which they will exhibit the attitudes in their attitude spaces. The attitude dynamics is given by (4.13) and (4.14), except that we consider the expected utility of the attitudes rather than the strategy profiles. Thus, as long as all attitudes are represented in the initial conditions, any steady state of the dynamics is an attitude equilibrium.

To apply the attitude dynamics to the Ultimatum game, we first quantize the players' attitude spaces. Each player's attitude space is $A_1 = A_2 = \{a_1, a_2, \ldots, a_{100}\}$, a set of 100 evenly spaced values on the interval $[0, 1]$. Although this provides a finite state space for the attitude dynamics, the high dimensionality and nonlinearity of the system of differential equations makes analysis difficult. However, we can make a few general statements about the results of the attitude dynamics.

We know that, given well-behaved initial conditions, the steady state of the dynamics is an attitude equilibrium in either pure or mixed attitudes. The pure-attitude equilibria are straightforward and have already been shown in Figure 4.1. While the mixed-strategy equilibria are more complicated, we still can extract a few simple and useful facts without overly complicated analysis. First, it is immediate that any attitude profile with all of its probability within one of the equilibrium regions of Figure 4.1 is itself an equilibrium. We cannot rule out, however, the possibility of probability mass located in the non-equilibrium portions of the (Low, Accept) and (High, Accept) regions. Fortunately, an attitude equilibrium cannot have probability mass in either the (Low, Reject) or (High, Reject) regions. If there is probability mass in those regions, each player can improve expected utility by modifying its mixed attitudes until there is no probability of rejecting the fraction. Therefore, the attitude dynamics are guaranteed to eliminate the dysfunctional behavior observed in Chapter 3.

To illustrate the behavior of the attitude dynamics, we study the dynamics of the Ultimatum game under two different scenarios by numerically approximating the solution to the differential equations defined by (4.13) and (4.14). For each player's initial conditions, we use a two-sided exponential distribution similar to the Laplace distribution. Unlike the Laplace distribution, however, the two sides are not symmetric. That is, the initial conditions are given by

$$c_i(0) = \begin{cases} Ce^{\lambda_1(a_i-\mu)}, & \text{for } a_i \leq \mu \\ Ce^{\lambda_2(\mu-a_i)}, & \text{for } a_i > \mu \end{cases}, \tag{4.15}$$

where λ_1 and λ_2 are chosen such that the expected value of the distribution is μ and the variance is an arbitrary σ^2, and C ensures normalization. This choice of distribution provides several benefits in the attitude dynamics. First, we can define an arbitrarily "tight" distribution around the player's desired initial attitudes while still giving nonzero probability to each element in the player's attitude space, ensuring that the steady-state of the dynamics is an attitude equilibrium.

The exponential distribution also encourages players' distributions to "shift" to adjacent values rather than "jump" across the pure-attitude set. Equations (4.13) and (4.14) show that the probabilities grow not only according to their relative utility, but also their current values. Therefore, this distribution ensures that attitudes close to the initial mean can grow more readily than those far away. This allows for a smoother and perhaps more realistic transition in the players' attitudes.

4.3.1 The "Arms Race"

In our simulations we let $l = 0.25$ and $h = 0.75$ be the low and high fractions offered. In this first scenario, we initialize the players' attitudes with means $\mu_1 = \mu_2 = 0.2$ and variances $\sigma_1^2 = \sigma_2^2 = 0.001$. Initially, the players' attitudes almost invariably lead to the offer and acceptance of the high fraction. While such behavior is not necessarily incoherent, their attitudes are not in equilibrium. The dynamics of this scenario provides a useful demonstration of the social Nash equilibrium as well as a highly interesting steady state.

Figure 4.2 shows the initial joint distribution of the player's attitudes. Since the player's mixed attitudes are probabilistically independent, the joint distribution is the product of the marginal distributions. Treating $\mathbf{c}(t)$ and $\mathbf{d}(t)$ as column vectors, we express the joint distribution as the matrix $\mathbf{J}(t) = \mathbf{c}(t)\mathbf{d}^T(t)$. Since the responder initially earns maximal payoff, it has no incentive to shift its attitudes. The proposer, however, can improve its payoff by increasing τ. In Figure 4.3, we see that the proposer shifts its attitudes such that the joint distribution peaks right on the boundary between the (Low, Accept) and (High, Accept) regions of the satisficing rectangle. The proposer has shifted its attitudes just enough move the players into the region where it obtains maximum payoff.

Once the shift is made, however, the responder stands to gain by modifying its preferences. In Figure 4.4, we see the results of an "arms race": the responder slightly increases δ to move the players to the (High, Accept) region, prompting the proposer to increase τ. The players "walk" their attitudes along the High/Low boundary until they intersect the Accept/Reject boundary. At this point (Figure 4.5), neither player can improve payoff by changing attitudes, and the distribution becomes almost entirely focused on the boundary point between the four regions. In this case, the specific behavior is an artifact of the quantization of the attitude spaces, and the players end up in the (High, Accept) region preferred by the responder.

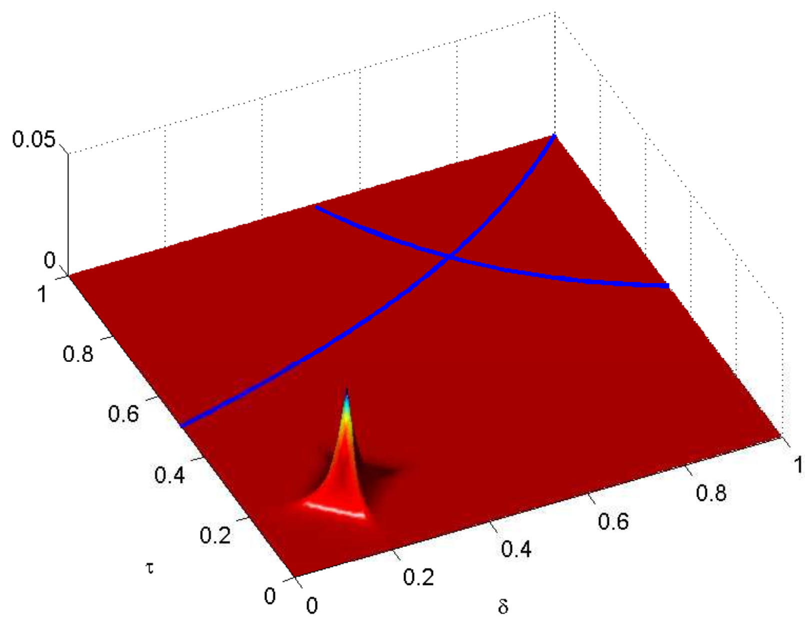

Figure 4.2: "Arms race" joint probability distribution: $t = 0$.

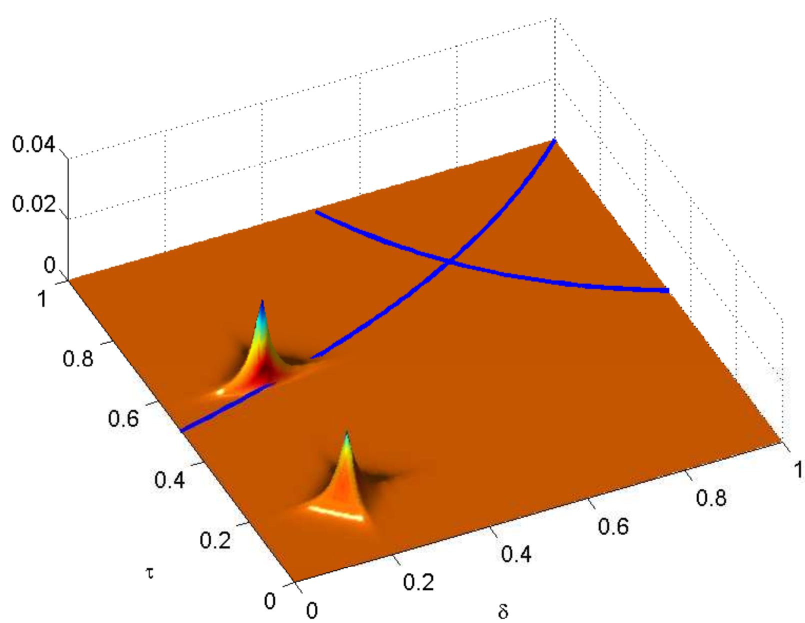

Figure 4.3: "Arms race" joint probability distribution: $t = 35$.

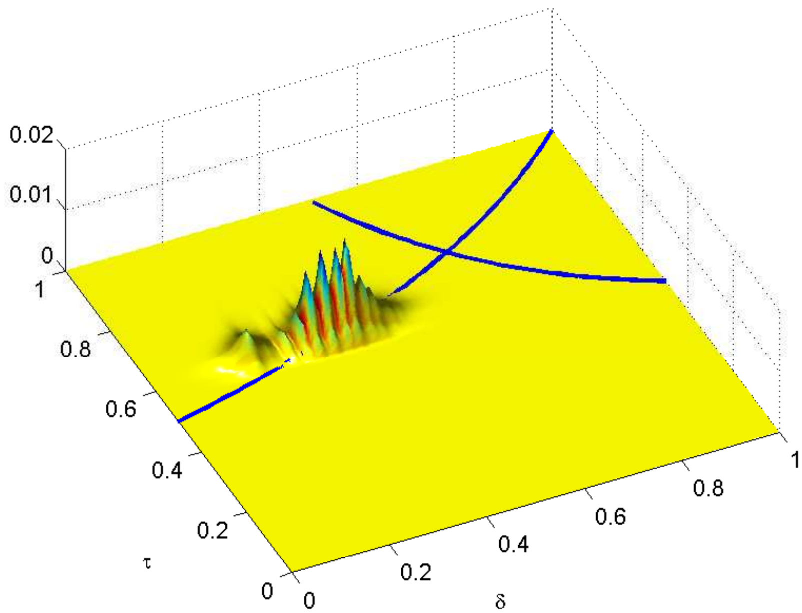

Figure 4.4: "Arms race" joint probability distribution: $t = 60$.

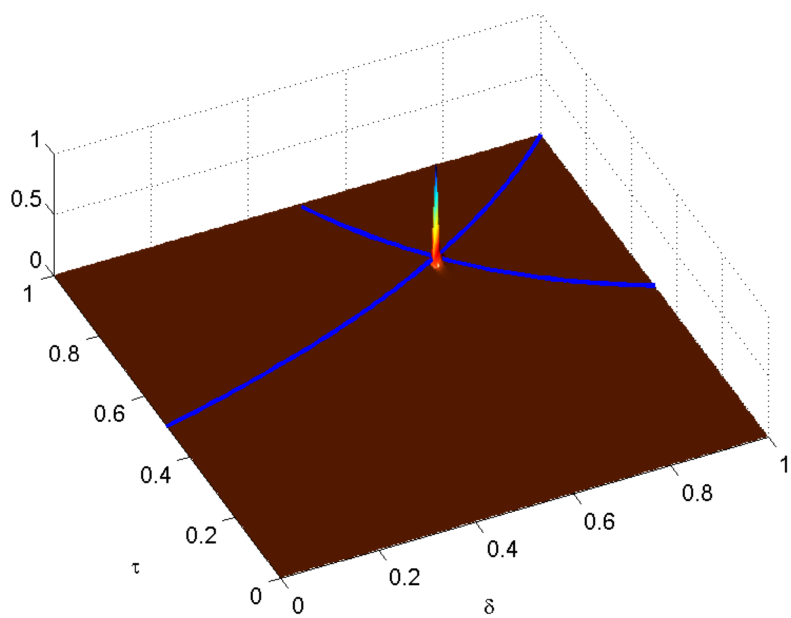

Figure 4.5: "Arms race" joint probability distribution: $t = 125$.

4.3.2 Overcoming Dysfunction

Consider the dysfunctional pair whose initial conditions are characterized by $\mu_1 = 0.8$, $\mu_2 = 0.9$ and $\sigma_1^2 = \sigma_2^2 = 0.001$ (Figure 4.6). Here, the responder rejects the high offer. The dynamics shifts both of the players' attitudes toward an attitude equilibrium, but which one? The proposer, of course, would prefer to end up in the (Low, Accept) region, while the responder prefers (High, Accept). The answer lies with which one adapts most quickly. Given the these initial conditions, the responder is able to shift its attitudes to a different region more quickly than the proposer. The responder begins accepting the low offer, and the proposer no longer has any reason to adjust its attitudes, resulting in (Low, Accept) as the steady state behavior (Figure 4.7). Interestingly, in the Ultimatum game, the player who adapts first ends up in its second-best region of the satisficing rectangle. Here, when the responder adapts most quickly, the players end up choosing (Low, Accept), which is preferred by the proposer. If we were to specify initial conditions in which the proposer were able to adapt more quickly, the players would eventually end up in (High, Accept), which is preferred by the responder. In either case, however, the dysfunctional behavior is overcome and the players choose a strategy profile that results in non-zero payoff to both players.

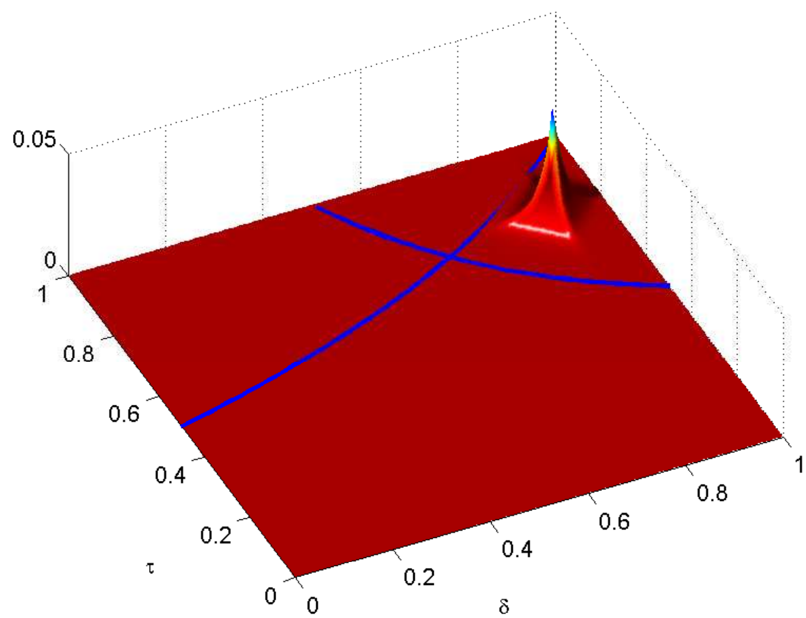

Figure 4.6: "Dysfunctional" joint probability distribution: $t = 0$.

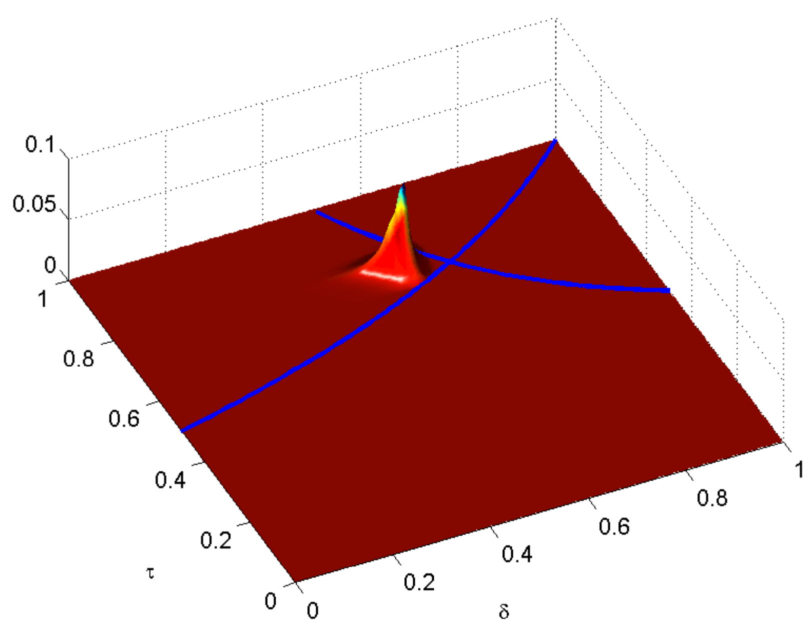

Figure 4.7: "Dysfunctional" joint probability distribution: $t = 100$.

Chapter 5

Social Utilities in Non-cooperative Games

Here we present a more explicit combination of satisficing and non-cooperative game theory. We apply the conditional structure of social utilities to an otherwise standard non-cooperative game-theoretic model. Players possess a single utility, which they may express conditionally, as in satisficing theory. However, each player's utility is defined over the entire pure-strategy space rather than only their own pure-strategy sets. As in satisficing theory, we marginalize the utilities, but now we may apply non-cooperative solutions concepts such as the Nash equilibrium to the marginal utilities. After detailing this framework mathematically, we construct a simple conditional-utility model for the Ultimatum game, showing that by adjusting the players' attitudes we modify the game's subgame perfect equilibrium.

5.1 Von Neumann-Morgenstern Utility Theory

First, we familiarize the reader with the theory of von Neumann-Morgenstern utility functions [25]. As usual, we define a set of players $I = \{1, 2, \ldots, n\}$ who must implement pure-strategy profiles $\mathbf{s}$ in the pure-strategy space $S = \times_{i=1}^{n} S_i$. Let Θ be the mixed-strategy space, or the space of all probability distributions over S. However, instead of defining the utilities directly, we first endow each player with ordinal preferences over the mixed-strategy space.

The ordinal preferences are expressed as a total ordering $\succeq_i$ to player i of the mixed-strategy profiles in Θ. For $\mathbf{x}, \mathbf{y} \in \Theta$ and $i \in I$, let $\mathbf{x} \succeq_i \mathbf{y}$ signify that $\mathbf{x}$ is *at least as preferred to* $\mathbf{y}$ by player i, and let $\mathbf{x} \sim_i \mathbf{y}$ signify that $\mathbf{x}$ is considered *equivalent to* $\mathbf{y}$ by player i. Since the preferences specify a total ordering, the following properties must be satisfied for all $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \Theta$ and $i \in I$:

1. Reflexivity: $\mathbf{x} \succeq_i \mathbf{x}$
2. Antisymmetry: $\mathbf{x} \succeq_i \mathbf{y}$ and $\mathbf{y} \succeq_i \mathbf{x} \Rightarrow \mathbf{x} \sim_i \mathbf{y}$.
3. Transitivity: $\mathbf{x} \succeq_i \mathbf{y}$ and $\mathbf{y} \succeq_i \mathbf{z} \Rightarrow \mathbf{x} \succeq_i \mathbf{z}$.
4. Completeness: Either $\mathbf{x} \succeq_i \mathbf{y}$ or $\mathbf{y} \succeq_i \mathbf{x}$.

With a total ordering across Θ, players can compare any two mixed-strategy profile in terms of their ordinal preferences.

Definition 5.1 *A* von Neumann-Morgenstern utility *is a function $\pi_i(\mathbf{s})$ defined over S whose expected value is consistent with the preference ordering $\succeq_i$:*

$$u_i(\mathbf{x}) \geq u_i(\mathbf{y}) \Leftrightarrow \mathbf{x} \succeq_i \mathbf{y}, \forall \mathbf{x}, \mathbf{y} \in \Theta, \tag{5.1}$$

where $u_i(\mathbf{x})$ denotes the expected value of π_i using the mixed-strategy profile $\mathbf{x}$ as defined in (2.5).

In other words, a probability distribution is preferred to another if and only if the expected von Neumann-Morgenstern utility is greater under the first distribution than it is under the second. In [26], existence conditions and properties for π_i are discussed in detail. For simplicity, we'll assume that $\succeq_i$ is sufficiently well-behaved that π_i exists. To conclude our discussion, we note one important property of the utilities.

Theorem 5.1 *Let π_i be a von Neumann-Morgenstern utility defined over S. Let π_i' be a positive affine transformation of π_i; that is, $\pi_i' = a\pi_i + b$, where a is positive and b is an arbitrary constant. Then π_i' is also a von Neumann-Morgenstern utility over S.*

Proof: This is shown by the linearity of expectation. Given $\mathbf{x} \in \Theta$, We can write the expectation of π_i' as

$$u_i'(\mathbf{x}) = a u_i(\mathbf{x}) + b.$$

Let $\mathbf{y}$ and $\mathbf{z}$ be two mixed-strategy profiles in Θ. Since a is positive and b is simply a constant, $u_i'(\mathbf{y}) \geq u_i'(\mathbf{z})$ if and only if $u_i(\mathbf{y}) \geq u_i(\mathbf{z})$. Since π_i is a von Neumann-Morgenstern utility, this is true if and only if $\mathbf{y} \succeq_i \mathbf{z}$. Thus, $u_i'(\mathbf{y}) \geq u_i'(\mathbf{z})$ if and only if $\mathbf{y} \succeq_i \mathbf{z}$, making π_i' a von Neumann-Morgenstern utility. ■

We can show by counterexample that a transformation that is *not* positive affine will, in general, violate $\succeq_i$. In fact, even a monotonic transformation—where the ordering on pure-strategy profiles is preserved—will, in general, violate the preference ordering $\succeq_i$ unless it is a positive affine transformation.

5.2 Social Utilities

5.2.1 Commitment and Conditional Utilities

As with the von Neumann-Morgenstern utilities, we build the conditional utilities from the "ground up," defining preference orderings from which the utilities are derived. We begin by defining a player's *commitment* to a pure-strategy profile, which was introduced in a slightly different context in [27].

Definition 5.2 *Player i is* committed *to the pure-strategy profile* $\mathbf{s} \in S$ *if*

$$\mathbf{s} \succ_i \mathbf{r}, \forall \mathbf{r} \neq \mathbf{s}$$

and

$$\mathbf{r} \sim_i \mathbf{q}, \forall \mathbf{r}, \mathbf{q} \neq \mathbf{s}.$$

In other words, player i is committed to $\mathbf{s}$ if $\mathbf{s}$ is most the most preferred strategy profile and if every other pure-strategy profile is equally inferior. If player i is committed to $\mathbf{s}_i$, it is the only pure-strategy profile worthy of consideration.

In our model the players' utilities are based upon total preference orderings over Θ. However, the players' preference orderings are conditional upon the commitments of the other players. That is, player i specifies a total preference ordering over the mixed-strategy space *given that* each of the other players are committed to pure-strategy profiles. Let $\mathbf{x} \succeq_{i|\mathbf{s}_1 \cdots \mathbf{s}_{i-1} \mathbf{s}_{i+1} \cdots \mathbf{s}_n} \mathbf{y}$ signify that the mixed-strategy profile $\mathbf{x}$ is at least preferred to $\mathbf{y}$ to player i given that player 1 is committed to $\mathbf{s}_1$, player

2 is committed to $\mathbf{s}_2$, and so on. For convenience, we will use the notation $\mathbf{s}_{-i}$ to refer to the collection of pure-strategy profiles to which the other $n-1$ players commit, allowing us to denote the conditional ordering as $\succeq_{i|\mathbf{s}_{-i}}$. We stress that these preference orderings are conditioned on a hypothetical commitment. The players do not need to be committed to the pure-strategy profiles in $\mathbf{s}_{-i}$ for player i to define its conditional preferences. Instead, if the players *were committed* to the pure-strategy profiles specified, player i's preference ordering over Θ would be $\succeq_{i|\mathbf{s}_{-i}}$. Player i's conditional ordering must be defined for each of the $|S|^{n-1}$ different combinations of pure-strategy profiles to which the other players may commit. These conditional orderings allow us to define conditional utilities:

Definition 5.3 *A* conditional utility *is a function* $\pi_{i|-i}(\mathbf{s}|\mathbf{s}_{-i})$ *defined over* S^n *which satisfies the following properties:*

1. *Consistency:* $u_{i|-i}(\mathbf{x}|\mathbf{s}_{-i}) \geq u_{i|-i}(\mathbf{y}|\mathbf{s}_{-i}) \Leftrightarrow \mathbf{x} \succeq_{i|\mathbf{s}_{-i}} \mathbf{y}, \forall \mathbf{x}, \mathbf{y} \in \Theta$, $\mathbf{s}_{-i} \in S^{n-1}$.
2. *Normalization:* $\sum_{\mathbf{s}\in S} \pi_{i|-i}(\mathbf{s}|\mathbf{s}_{-i}) = 1, \forall \mathbf{s}_{-i} \in S^{n-1}$.
3. *Nonnegativity:* $\pi_{i|-i}(\mathbf{s}|\mathbf{s}_{-i}) \geq 0, \forall \mathbf{s} \in S, \mathbf{s}_{-i} \in S^{n-1}$.
4. *Uniqueness:* $\min_{\mathbf{s}\in S} \pi_{i|-i}(\mathbf{s}|\mathbf{s}_{-i}) = 0$ *unless* $\pi_{i|-i}(\mathbf{s}|\mathbf{s}_{-i}) = \pi_{i|-i}(\mathbf{s}'|\mathbf{s}_{-i}), \forall \mathbf{s}' \in S$.

For each $\mathbf{s}_{-i} \in S^{n-1}$ to which the remaining players commit, $\pi_{i|-i}(\mathbf{s}_i|\mathbf{s}_{-i})$ is a standard von Neumann-Morgenstern utility over the ordering $\succeq_{i|\mathbf{s}_{-i}}$. We additionally impose the normalization and nonnegativity constraints, making the conditional utilities syntactically equivalent to a conditional probability mass function. By forcing the minimum value to zero, we define a unique normalized utility consistent with $\succeq_{i|\mathbf{s}_{-i}}$. Since von Neumann-Morgenstern utilities are unique only up to a positive affine transformation, we simply select the transformation that results in a normalized, nonnegative utility where the minimum is zero.

A player's conditional utility does not necessarily depend on the commitments of all of the remaining players; it may depend on a subset of the remaining players' commitments or upon none at all.

Definition 5.4 *Player i's conditional utility is* preferentially independent *of player j's commitment if*

$$\pi_{i|-ij,j}(\mathbf{s}|\mathbf{s}_{-ij},\mathbf{s}_j) = \pi_{i|-ij,j}(\mathbf{s}|\mathbf{s}_{-ij},\mathbf{s}'_j), \forall \mathbf{s}'_j \in S,$$

or, equivalently

$$\succeq_{i|\mathbf{s}_{-ij},\mathbf{s}_j} = \succeq_{i|\mathbf{s}_{-ij},\mathbf{s}'_j}, \forall \mathbf{s}'_j \in S,$$

where $\mathbf{s}_{-ij} \in S^{n-2}$ *is the collection of pure-strategy profiles to which all the players except for* i *and* j *commit. Player i's conditional preference ordering—and therefore the conditional utility—does not change with the commitment of player* j.

Note that preferential independence[1] is a one-way relationship. That is, player i's independence to player j's commitment does not imply that player j's utilities are independent of player i's commitment. Player j may still condition its utilities on player i's commitment. If player i's utilities are independent of the all of the other players, then we can write its conditional utility and preference ordering as π_i and $\succeq_i$, respectively.

While we regard the players' commitments as hypothetical in terms of the conditional utility functions, a player may actually commit to a pure-strategy profile. If player i commits to the pure-strategy profile $\mathbf{s}^*$, then by the definition of commitment, player i's utility is preferentially independent of the remaining players, and all of player i's utility is placed on $\mathbf{s}^*$, or

$$\pi_i(\mathbf{s}) = \begin{cases} 1, & \text{if } \mathbf{s} = \mathbf{s}^* \\ 0, & \text{otherwise} \end{cases}.$$

5.2.2 Marginal Utilities

The conditional utilities express player i's utility given hypothetical commitments to pure-strategy profiles by the remaining players, prompting an obvious ques-

[1] Preferential independence should not be confused with the familiar notion of "statistical" independence from probability theory. If two players' utilities are statistically independent, then the joint utility is the product of marginal utilities, or $\pi_{ij}(\mathbf{s}_i,\mathbf{s}_j) = \pi_i(\mathbf{s}_i)\pi_j(\mathbf{s}_j)$. While the preferential dependency relationships between players *do* determine whether their utilities are statistically independent, preferential independence does not necessarily imply statistical independence.

tion: what are player i's *actual* utilities? That is, from the conditional utilities, how do we extract a utility $\pi_i(\mathbf{s})$, defined over S, for each player?

The conditional utilities conform to the mathematical syntax of conditional probability mass functions. As in Chapter 3, we may appeal to the argument given in [21], which states that we must use the rules of probability theory in order to satisfy several desirable axioms. From the conditional utility $\pi_{i|-i}$, we can form a utility over S by summing over S^{n-1}:

$$\pi_i(\mathbf{s}_i) = \sum_{\mathbf{s}_{-i} \in S^{n-1}} \pi_{i|-i}(\mathbf{s}_i|\mathbf{s}_{-i})\pi_{-i}(\mathbf{s}_{-i}), \tag{5.2}$$

where $\pi_{-i}(\mathbf{s}_{-i})$ is the *joint* utility of the other players. The specification for $\pi_{-i}(\mathbf{s}_{-i})$ depends on the preferential dependency relationships between the players' utilities, which we will examine in more detail later. However, this joint utility is a mass function and thus is nonnegative and sums to unity.

As with satisficing theory, and by analogy with probability theory, we refer to π_i as player i's *marginal* utility. Equation (5.2) defines π_i as a convex combination of the conditional utilities defined by $\pi_{i|-i}$. Player i specifies its hypothetical utilities by $\pi_{i|-i}$ for each of the possible combinations of pure-strategy profiles in S^{n-1}. We then form a weighted sum of each of those hypothetical utilities, where the weighting is determined by the joint utility π_{-i}.

The joint utilities $\pi_{-i}(\mathbf{s}_{-i})$, which are used to compute each player's marginal utility, depend on the dependency relationships between the players' utilities. These relationships greatly impact whether or not we can extract a unique marginal utility from the conditionals. As we will see, there are dependency relationships for which we cannot solve for unique marginal utilities defined by the conditional utilities. As with the conditional relationships in satisficing theory, we may express the relationships between the utilities graphically.

Definition 5.5 *A* utility network *is a directed (but not necessarily acyclic) graph with n nodes and an arbitrary number of edges. The nodes represent the players' marginal utilities, while the edges represent conditioning between players' utilities.*

Figure 5.1 shows two utility networks that are acyclic. In the utility network shown in Figure 5.1(a), the first three players form their utilities independent of the others' commitments, or $\pi_{i|-i} = \pi_i, i = 1, 2, 3$. Since they are mutally independent, the joint utility of the first three players is simply the product of marginals, or $\pi_{123} = \pi_1\pi_2\pi_3$. Player 4's conditional utility, which depends on each of the remaining players' commitments, is $\pi_{4|123}$. The marginal utility is given by

$$\pi_4(\mathbf{s}_4) = \sum_{\mathbf{s}_1,\mathbf{s}_2,\mathbf{s}_3 \in S} \pi_{4|123}(\mathbf{s}_4|\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3)\pi_{123}(\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3) \quad (5.3)$$

$$= \sum_{\mathbf{s}_1,\mathbf{s}_2,\mathbf{s}_3 \in S} \pi_{4|123}(\mathbf{s}_4|\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3)\pi_1(\mathbf{s}_1)\pi_2(\mathbf{s}_2)\pi_3(\mathbf{s}_3). \quad (5.4)$$

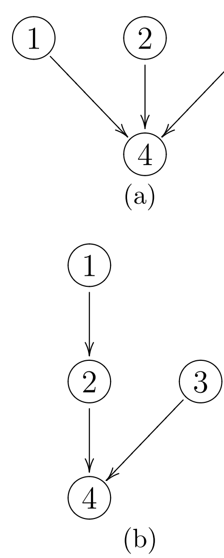

Figure 5.1: Two acyclic utility networks.

In Figure 5.1(b), players 1 and 3 form their utilities independently. Player 2's conditional utility depends on player 1's commitments, and the marginal utility is

$$\pi_2(\mathbf{s}_2) = \sum_{\mathbf{s}_1 \in S} \pi_{2|1}(\mathbf{s}_2|\mathbf{s}_1)\pi_1(\mathbf{s}_1). \tag{5.5}$$

While player 4's conditional utility only explicitly depends on the commitments of players 2 and 3, there is an indirect dependency on player 1's utilities, as seen in the marginal utility:

$$\pi_4(\mathbf{s}_4) = \sum_{\mathbf{s}_2,\mathbf{s}_3 \in S} \pi_{4|23}(\mathbf{s}_4|\mathbf{s}_2,\mathbf{s}_3)\pi_{23}(\mathbf{s}_2,\mathbf{s}_3) \tag{5.6}$$

$$= \sum_{\mathbf{s}_2,\mathbf{s}_3 \in S} \pi_{4|23}(\mathbf{s}_4|\mathbf{s}_2,\mathbf{s}_3)\pi_2(\mathbf{s}_2)\pi_3(\mathbf{s}_3) \tag{5.7}$$

$$= \sum_{\mathbf{s}_2,\mathbf{s}_3 \in S} \pi_{4|23}(\mathbf{s}_4|\mathbf{s}_2,\mathbf{s}_3) \sum_{\mathbf{s}_1 \in S} \pi_{2|1}(\mathbf{s}_2|\mathbf{s}_1)\pi_1(\mathbf{s}_1)\pi_3(\mathbf{s}_3) \tag{5.8}$$

$$= \sum_{\mathbf{s}_1\mathbf{s}_2,\mathbf{s}_3 \in S} \pi_{4|23}(\mathbf{s}_4|\mathbf{s}_2,\mathbf{s}_3)\pi_{2|1}(\mathbf{s}_2|\mathbf{s}_1)\pi_1(\mathbf{s}_1)\pi_3(\mathbf{s}_3). \tag{5.9}$$

In an acyclic graph, it is always possible to express the players' marginal utilities as a function only of the conditional utilities. So, in an acyclic graph, the marginal utilities are always uniquely determined from the conditional utilities.

Figure 5.2 shows a cyclic utility network with two players. In this network, the players' conditional utilities are functions of each other's commitments, meaning that neither of the players expresses its marginal utility directly. To appreciate the difficulty associated with such a network, consider the marginal utility of player 1, which is given by

$$\pi_1(\mathbf{s}_1) = \sum_{\mathbf{s}_2 \in S} \pi_{1|2}(\mathbf{s}_1|\mathbf{s}_2)\pi_2(\mathbf{s}_2). \tag{5.10}$$

However, π_2 is expressed in terms of π_1:

$$\pi_2(\mathbf{s}_2) = \sum_{\mathbf{s}_1 \in S} \pi_{2|1}(\mathbf{s}_2|\mathbf{s}_1)\pi_1(\mathbf{s}_1). \tag{5.11}$$

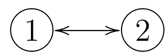

Figure 5.2: Two-player cyclic utility network.

We might wonder whether or not a solution to the coupled equations exists, and whether or not it is unique. Fortunately, the analysis is fairly straightforward if we convert (5.10) and (5.11) to matrix-vector equations. We express the conditional utilities as the $|S| \times |S|$ matrices $\mathbf{\Pi}_{1|2}$ and $\mathbf{\Pi}_{2|1}$. Each column of the matrices represents the conditional utility for a different pure-strategy commitment. The matrices are therefore column stochastic, meaning that each column is nonnegative sums to unity. We express the marginal utilities as $|S|$-dimensional column stochastic vectors $\boldsymbol{\pi}_1$ and $\boldsymbol{\pi}_2$. We can now rewrite (5.10) and (5.11) as

$$\boldsymbol{\pi}_1 = \mathbf{\Pi}_{1|2}\boldsymbol{\pi}_2, \tag{5.12}$$

$$\boldsymbol{\pi}_2 = \mathbf{\Pi}_{2|1}\boldsymbol{\pi}_1. \tag{5.13}$$

Substituting the equations into each other yields

$$\boldsymbol{\pi}_1 = \mathbf{\Pi}_{1|2}\mathbf{\Pi}_{2|1}\boldsymbol{\pi}_1, \tag{5.14}$$

$$\boldsymbol{\pi}_2 = \mathbf{\Pi}_{2|1}\mathbf{\Pi}_{1|2}\boldsymbol{\pi}_2. \tag{5.15}$$

Now we have transformed the problem into an eigenvector problem. If $\lambda = 1$ is an eigenvalue for the matrices $\mathbf{\Pi}_{1|2}\mathbf{\Pi}_{2|1}$ and $\mathbf{\Pi}_{2|1}\mathbf{\Pi}_{1|2}$, then there exists a solution to the equations. If there is a unique linearly independent eigenvector associated with $\lambda = 1$ for these matrices, the solution is unique.

First, we discuss existence. Note that the product of two column stochastic matrices is also column stochastic. From Perron-Frobenius theory [28], we know that every column stochastic matrix has the eigenvalue $\lambda = 1$. Therefore, there must exist marginal utilities that satisfy (5.14) and (5.15). However, there is no guarantee that a unique solution exists. To see this, let $\mathbf{\Pi}_{1|2} = \mathbf{\Pi}_{2|1} = \mathbf{I}$, the identity matrix. Since all of the relevant matrices are the identity, any stochastic vector satisfies (5.14) and (5.15) so long as we choose $\boldsymbol{\pi}_1 = \boldsymbol{\pi}_2$. Here, the conditional utility functions do not uniquely specify the players' marginal utilities.

A closer look at the conditional utilities provides an intuitive explanation. Each player's conditional utility is the identity matrix. Such a conditional utility is an expression of complete indifference: each player has decided that its marginal utility will simply be the other player's utility. Since neither player expresses any meaningful preferences, we cannot solve for unique marginal utilities. In general, when the conditional utilities do not yield unique marginal utilities, we say that the marginal utilities *do not exist*, even though technically there may be an infinity of marginal utilities consistent with the conditional utilities.

Finally, we examine the effect of commitment on the marginal utilities. If every player except for player i commits to a pure-strategy profile, their utilities are preferentially independent, and all of the utility is placed on the pure-strategy profile to which they commit. Letting $\mathbf{s}^*_{-i}$ denote the $|S| - 1$ pure-strategy profiles to which the players commit, the joint utility is

$$\pi_{-i}(\mathbf{s}_{-i}) = \prod_{j \neq i} \pi_j(\mathbf{s}_j^*) \tag{5.16}$$

$$= \begin{cases} 1, & \text{if } \mathbf{s}_{-i} = \mathbf{s}_{-i}^* \\ 0, & \text{otherwise} \end{cases}. \tag{5.17}$$

Therefore, the marginal utility for player i is

$$\pi_i(\mathbf{s}_i) = \sum_{\mathbf{s}_{-i} \in S^{n-1}} \pi_{i|-i}(\mathbf{s}_i|\mathbf{s}_{-i})\pi_{-i}(\mathbf{s}_{-i}) \tag{5.18}$$

$$= \pi_{i|-i}(\mathbf{s}_i|\mathbf{s}_{-i}^*). \tag{5.19}$$

If players commit to pure-strategy profiles, then the resulting marginal utilities are simply the conditional utilities evaluated at the pure-strategy profiles to which the players commit.

5.2.3 Marginal Preferences

We have defined the marginal utility entirely in terms of the conditional and joint utilities rather than in terms of a total preference ordering. That is, we have not specified a "marginal" ordering $\succeq_i$ that is derived directly from $\succeq_{i|\mathbf{s}_{-i}}$. However, when a unique marginal utility π_i exists, it *does* induce a total preference ordering over the mixed-strategy space:

$$\mathbf{x} \succeq_i \mathbf{y} \Leftrightarrow u_i(\mathbf{x}) \geq u_i(\mathbf{y}).$$

Fortunately, we have defined the conditional utilities such that there is a unique conditional utility for every conditional ordering $\succeq_{i|\mathbf{s}_{-i}}$. Therefore, when unique marginal utilities exist, the conditional preferences define unique marginal preference orderings

$\succeq_i$. This gives justification to the fourth property in Definition 5.3, where we required a unique normalization on the conditional utilities. In the next theorem, we show that without a unique normalization, there is not, in general, a unique preference ordering $\succeq_i$ defined by the conditional orderings $\succeq_{i|\mathbf{s}_{-i}}$.

Theorem 5.2 *A unique normalization of the conditional utilities guarantees a unique preference ordering $\succeq_i$ to result from the conditional preferences $\succeq_{i|\mathbf{s}_{-i}}, \forall i \in I$. Moreover, a unique normalization is a necessary condition for the existence of a unique preference ordering in the following sense: if we remove the uniqueness property from the definition of a conditional utility, the conditional preferences no longer, in general, specify unique marginal preference orderings.*

Proof: Sufficiency follows directly from the uniqueness of the conditional utilities. For necessity, note that by removing the requirement of unique normalization, we can apply a positive affine transformation to the conditional utilities as long as it preserves the normalization and nonnegativity constraints. For a conditional utility $\pi_{i|-i}$, such a transformation is of the form $\pi'_{i|-i} = a\pi_{i|-i} + b$ such that

$$-\min_{\mathbf{s}_i \in S} \pi_{i|-i}(\mathbf{s}_i|\mathbf{s}_{-i}) \leq b < |S|^{-1}$$

and

$$a = 1 - b\,|S|\,.$$

The choice of a and b ensure that $\pi'_{i|-i}$ still meets the nonnegativity and normalization criteria of a conditional mass function.

Next, consider a simple example where there are only two players. Player 1's conditional utility depends on the pure-strategy profile to which player 2 (hypothetically) commits, while player 2's conditional utility is independent of player 1. Player 2's conditional utility is therefore equivalent to its marginal utility, or $\pi_{2|1}(\mathbf{s}_2|\mathbf{s}_1) = \pi_2(\mathbf{s}_2)$. Denote the two conditional preference orderings $\succeq_{1|\mathbf{s}_2}$ and $\succeq_2$. Player 1's marginal utility is given by

$$\pi_1(\mathbf{s}_1) = \sum_{\mathbf{s}_2 \in S} \pi_{1|2}(\mathbf{s}_1|\mathbf{s}_2)\pi_2(\mathbf{s}_2). \tag{5.20}$$

Define $\succeq_1$ as the preference ordering induced by the marginal utility π_1.

Now, define π'_2, a positive affine transformation of π_2 that preserves normalization and nonnegativity as discussed. Since it is an affine transformation, π'_2 is still consistent with the preference ordering $\succeq_2$. Player 1's marginal utility using π'_2 is now

$$\pi'_1(\mathbf{s}_1) = \sum_{\mathbf{s}_2 \in S} \pi_{1|2}(\mathbf{s}_1|\mathbf{s}_2)\pi'_2(\mathbf{s}_2) \tag{5.21}$$

$$= \sum_{\mathbf{s}_2 \in S} \pi_{1|2}(\mathbf{s}_1|\mathbf{s}_2)(a\pi_2(\mathbf{s}_2) + b) \tag{5.22}$$

$$= a \sum_{\mathbf{s}_2 \in S} \pi_{1|2}(\mathbf{s}_1|\mathbf{s}_2)\pi_2(\mathbf{s}_2) + b \sum_{\mathbf{s}_2 \in S} \pi_{1|2}(\mathbf{s}_1|\mathbf{s}_2) \tag{5.23}$$

$$= a\pi_1(\mathbf{s}_1) + b \sum_{\mathbf{s}_2 \in S} \pi_{1|2}(\mathbf{s}_1|\mathbf{s}_2). \tag{5.24}$$

Define $\succeq'_1$ as the preference ordering induced by the new marginal utility π'_1. Since in (5.24) we sum over $\mathbf{s}_2$, the second term is not generally constant with respect to $\mathbf{s}_1$. Therefore, π'_1 is not a positive affine transformation of π_1 and is not in general consistent with the original preference ordering $\succeq_1$. Thus, in general, $\succeq_1 \neq \succeq'_1$. Since positive affine transformations in the conditional utilities result in different marginal preference orderings, a unique normalization on the conditionals is required to ensure a unique marginal preference ordering. ■

5.3 The Ultimatum Game

In this section we frame the Ultimatum Game as a non-cooperative game with social utilities. We again employ the simple quantized game presented in [15], where the proposer offers a high or low fraction and the responder accepts or rejects the offer. For convenience, we abbreviate the notation used by denoting the proposer's pure-strategy set as $S_1 = \{H, L\}$ and the responder's as $S_2 = \{A, R\}$. As in previous chapters, we denote the numerical value of the high and low fractions as h and l, respectively, with $h > l$.

We model the players with a cyclic two-player utility network—each player's conditional utility depends on the commitment of the other player. Because of this, we will directly express the players' utilities in the matrix-vector form shown in Section 5.2.2. We assume a sociological model similar to that of the satisficing case presented in [19], where the proposer's only social consideration is fear that the responder might reject a low offer, and the responder's only social consideration is indignation over being offered an unfair fraction.

5.3.1 Conditional Utilities

We begin by defining the proposer's conditional utilities. Since each player chooses between two pure strategies, the pure-strategy space contains four pure-strategy profiles. We express the proposer's conditional utilities with the 4×4 matrix $\mathbf{\Pi}_{1|2}$, which has the form

$$\mathbf{\Pi}_{1|2} = \begin{bmatrix} \pi_{1|2}(H,A|H,A) & \pi_{1|2}(H,A|L,A) & \pi_{1|2}(H,A|H,R) & \pi_{1|2}(H,A|L,R) \\ \pi_{1|2}(L,A|H,A) & \pi_{1|2}(L,A|L,A) & \pi_{1|2}(L,A|H,R) & \pi_{1|2}(L,A|L,R) \\ \pi_{1|2}(H,R|H,A) & \pi_{1|2}(H,R|L,A) & \pi_{1|2}(H,R|H,R) & \pi_{1|2}(H,R|L,R) \\ \pi_{1|2}(L,R|H,A) & \pi_{1|2}(L,R|L,A) & \pi_{1|2}(L,R|H,R) & \pi_{1|2}(L,R|L,R) \end{bmatrix}.$$

The first column describes the proposer's utility given that the responder commits to $\mathbf{s}_2 = (H,A)$, the second column describes the proposer's utility given that the responder is committed to $\mathbf{s}_2 = (L,A)$, and so forth.

Since the proposer's only social consideration is that the responder might reject the low offer, the proposer has no need to modify its utility if the responder commits to a pure-strategy profile in which the offer is accepted. Given these commitments, the proposer's conditional utility is simply the normalized payoffs from the payoff matrix in Table 2.5. This specifies the first two columns of $\mathbf{\Pi}_{1|2}$:

$$\begin{bmatrix} \pi_{1|2}(H,A|H,A) \\ \pi_{1|2}(L,A|H,A) \\ \pi_{1|2}(H,R|(H,A) \\ \pi_{1|2}(L,R|(H,A) \end{bmatrix} = \begin{bmatrix} \pi_{1|2}(H,A|L,A) \\ \pi_{1|2}(L,A|L,A) \\ \pi_{1|2}(H,R|L,A) \\ \pi_{1|2}(L,R|L,A) \end{bmatrix} = \frac{1}{2-l-h} \begin{bmatrix} 1-h \\ 1-l \\ 0 \\ 0 \end{bmatrix}.$$

If, however, the responder commits to a strategy profile in which the offer is rejected, the proposer tempers its desire for payoff by the *temperance index*[2] $0 \leq \tau \leq$

[2]The intuition behind the temperance index is identical to the intemperance index from the satisficing model, except that we reverse the direction: the higher the temperance index, the more

1. If $\tau = 0$, the proposer ignores the responder's preferences and the possibility that it might reject the offer. If $\tau = 1$, the proposer completely defers to the responder's preferences and gives maximal utility to (h, a). In either case, the proposer assigns zero utility to having the offer rejected. This gives us the third and fourth columns of $\Pi_{1|2}$:

$$\begin{bmatrix} \pi_{1|2}(H,A|H,R) \\ \pi_{1|2}(L,A|H,R) \\ \pi_{1|2}(H,R|H,R) \\ \pi_{1|2}(L,R|H,R) \end{bmatrix} = \begin{bmatrix} \pi_{1|2}(H,A|L,R) \\ \pi_{1|2}(L,A|L,R) \\ \pi_{1|2}(H,R|L,R) \\ \pi_{1|2}(L,R|L,R) \end{bmatrix} = \frac{1}{2-l-h} \begin{bmatrix} (1-\tau)(1-h) + \tau(2-l-h) \\ (1-\tau)(1-l) \\ 0 \\ 0 \end{bmatrix}.$$

Now we can write player 1's conditional utilities as a single matrix:

$$\mathbf{\Pi}_{1|2} = \frac{1}{2-l-h} \begin{bmatrix} 1-h & 1-h & (1-\tau)(1-h)+\tau(2-l-h) & (1-\tau)(1-h)+\tau(2-l-h) \\ 1-l & 1-l & (1-\tau)(1-l) & (1-\tau)(1-l) \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

Next we define the responder's conditional utilities. The only social consideration facing the responder is a desire to punish the proposer for his offering an unfair fraction. Therefore, if the proposer commits to a strategy profile in which the high fraction is offered, the responder's utilities come directly from the payoff matrix, giving us columns 1 and 3 of $\mathbf{\Pi}_{2|1}$:

$$\begin{bmatrix} \pi_{2|1}(H,A|H,A) \\ \pi_{2|1}(L,A|H,A) \\ \pi_{2|1}(H,R|H,A) \\ \pi_{2|1}(L,R|H,A) \end{bmatrix} = \begin{bmatrix} \pi_{2|1}(H,A|H,R) \\ \pi_{2|1}(L,A|H,R) \\ \pi_{2|1}(H,R|H,R) \\ \pi_{2|1}(L,R|H,R) \end{bmatrix} = \frac{1}{l+h} \begin{bmatrix} h \\ l \\ 0 \\ 0 \end{bmatrix}.$$

the proposer is willing to accommodate the responder. While this reversal may be slightly confusing, it ensures that if the responder sets $\tau = 0$, its utilities revert to the payoffs from the payoff matrix.

If the proposer commits to a pure-strategy profile in which the low fraction is offered, the responder increasingly prefers to punish the responder by rejecting the offer. We model this with the *indignation index* $0 \leq \delta \leq 1$, which describes the responder's response to being offered the low fraction. As δ approaches 1, the responder is increasingly willing to forfeit the low fraction, and places all of the utility on rejecting the low fraction:

$$\begin{bmatrix} \pi_{2|1}(H,A|L,A) \\ \pi_{2|1}(L,A|L,A) \\ \pi_{2|1}(H,R|L,A) \\ \pi_{2|1}(L,R|L,A) \end{bmatrix} = \begin{bmatrix} \pi_{2|1}(H,A|L,R) \\ \pi_{2|1}(L,A|L,R) \\ \pi_{2|1}(H,R|L,R) \\ \pi_{2|1}(L,R|L,R) \end{bmatrix} = \frac{1}{l+h} \begin{bmatrix} h(1-\delta) \\ l(1-\delta) \\ 0 \\ \delta(l+h) \end{bmatrix}.$$

Combining these columns into the full matrix, we get

$$\mathbf{\Pi}_{2|1} = \frac{1}{l+h} \begin{bmatrix} h & h(1-\delta) & h & h(1-\delta) \\ l & l(1-\delta) & l & l(1-\delta) \\ 0 & 0 & 0 & 0 \\ 0 & \delta(l+h) & 0 & \delta(l+h) \end{bmatrix}.$$

Note that for both players, the conditional utilities reduce to normalized versions of the payoff matrix—and are independent of the commitment of the other player—if social considerations are ignored ($\tau = 0$ or $\delta = 0$).

5.3.2 Marginal Utilities

As in the two-player example from Section 5.2.2, we set up two eigenvector equations to compute the marginal utilities:

$$\boldsymbol{\pi}_1 = \mathbf{\Pi}_{1|2}\mathbf{\Pi}_{2|1}\boldsymbol{\pi}_1,$$

$$\boldsymbol{\pi}_2 = \mathbf{\Pi}_{2|1}\mathbf{\Pi}_{1|2}\boldsymbol{\pi}_2.$$

To find the solution for the marginal $\boldsymbol{\pi}_1$, we consider the eigen-decomposition of $\mathbf{\Pi}_{1|2}\mathbf{\Pi}_{2|1}$. Since the analysis is tedious and not terribly enlightening, we relegate the details to Appendix B. Fortunately, the eigen-decomposition yields a unique, normalized eigenvector for $\lambda = 1$, which is the solution for the marginal $\boldsymbol{\pi}_1$:

$$\boldsymbol{\pi}_1 = \begin{bmatrix} \pi_1(H,A) \\ \pi_1(L,A) \\ \pi_1(H,R) \\ \pi_1(L,R) \end{bmatrix} = \frac{1}{2 - l - h + \tau\delta(1-l)} \begin{bmatrix} (1-h) + \tau\delta(1-l) \\ 1-l \\ 0 \\ 0 \end{bmatrix}.$$

This result is gratifyingly intuitive. If either player ignores social considerations ($\tau = 0$ or $\delta = 0$), player 1's utilities revert to the (normalized) raw payoffs. Otherwise, the relative utility for $\mathbf{s}_1 = (H, A)$ increases according to the product $\tau\delta(1-l)$. Setting $(1-h) + \tau\delta(1-l) = 1-l$, we can show that player 1 is indifferent between the strategy profiles (H, A) and (L, A) when

$$\tau = \frac{h-l}{\delta(1-l)}.$$

When τ is greater than this quantity, player 1 prefers (H, A) to (L, A), even though it results in lower raw payoff.

To find player 2's marginal utilities, we consider the eigen-decomposition of $\Pi_{1|2}\Pi_{2|1}$, which also yields a unique vector for $\lambda = 1$:

$$\boldsymbol{\pi}_2 = \begin{bmatrix} \pi_2(H,A) \\ \pi_2(L,A) \\ \pi_2(H,R) \\ \pi_2(L,R) \end{bmatrix} = \frac{1}{(l+h)(2 - l - h + \tau\delta(1-l))} \begin{bmatrix} h(2-h-l+\delta(1-l)(\tau-1)) \\ l(2-h-l+\delta(1-l)(\tau-1)) \\ 0 \\ \delta(l+h)(1-l) \end{bmatrix}.$$

Although player 2's utility is more complicated, we still can glean an intuition from its form. We see that if $\delta = 0$, the utilities are those from the payoff matrix. Notice, however, that player 2's utilities are *not* simply those from the payoff matrix when $\tau = 0$. Regardless of τ, player 2 assigns nonzero utility to $\mathbf{s}_2 = (L, R)$ as long as player 1 has nonzero utility for (L, A). When

$$\delta = \frac{h(2-h-l)}{(h(2-\tau)+l)(1-l)},$$

player 2 is indifferent between (L, A) and (L, R). When δ exceeds this equality, player 2 prefers (L, R) to (L, A). In fact, it is even possible, for sufficiently high δ, for (L, R) to be the most preferred strategy profile, preferred even to (H, A)! This, however, does not affect the Nash equilibria that result, as we shall see in the next subsection.

5.3.3 Equilibria

To examine the equilibria of the Ultimatum game using social utilities, we take the marginal utilities from $\boldsymbol{\pi}_1$ and $\boldsymbol{\pi}_2$ and arrange them into a game tree. While the exact utilities depend on l, h, δ, and τ, the game will always have the form given in Figure 5.3, where

$$\alpha = \frac{(1-h)+\tau\delta(1-l)}{2-l-h+\tau\delta(1-l)}, \tag{5.25}$$

$$\beta = \frac{h(2-h-l+\delta(1-l)(\tau-1)}{(l+h)(2-l-h+\tau\delta(1-l)}, \tag{5.26}$$

$$\gamma = \frac{l(2-h-l+\delta(1-l)(\tau-1)}{(l+h)(2-l-h+\tau\delta(1-l)}, \tag{5.27}$$

$$\omega = \frac{\delta(l+h)(1-l)}{(l+h)(2-l-h+\tau\delta(1-l)}. \tag{5.28}$$

For convenience, we also display the normal-form representation in the payoff matrix of Table 5.1.

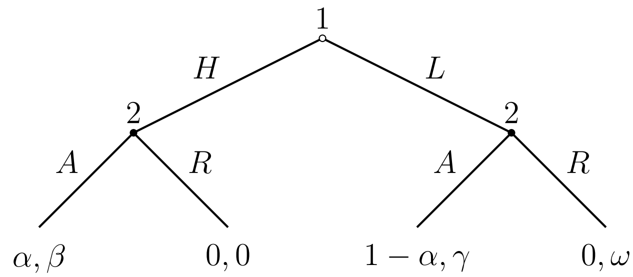

Figure 5.3: Game tree for the Ultimatum game with social utilities.

Table 5.1: Payoff matrix for the Ultimatum game with social utilities.

	Responder			
Proposer	(A, A)	(A, R)	(R, A)	(R, R)
H	(α, β)	(α, β)	$(0, 0)$	$(0, 0)$
L	$(1-\alpha, \gamma)$	$(0, \omega)$	$(1-\alpha, \gamma)$	$(0, \omega)$

Depending on the parameters, there are four possible cases. When $\alpha < 1/2$ and $\omega < \gamma$, there are two[3] pure-strategy equilibria: $\mathbf{s} = (L, A)$ and $\mathbf{s} = (H, A)$. However, $\mathbf{s} = (H, A)$ is an equilibrium only if the responder plays (A, R) and threatens to reject the low fraction. This strategy is not a credible threat since $\omega < \gamma$ and the responder prefers to accept (rather than reject) the low fraction. The unique subgame perfect equilibrium is therefore $\mathbf{s} = (L, A)$.

Next, we consider $\alpha > 1/2$ and $\omega < \gamma$. Now the proposer prefers to offer the high fraction. Again, the Nash equilibria are $\mathbf{s} = (L, A)$ and $\mathbf{s} = (H, A)$. However, in this case the $\mathbf{s} = (L, A)$ occurs only when the responder threatens to reject the high fraction, which is not credible. Therefore, we are left with $\mathbf{s} = (H, A)$ as the subgame perfect equilibrium.

In the third case, $\alpha < 1/2$ and $\omega > \gamma$. Now, the responder prefers to reject the low offer rather than accept it. The two equilibria are $\mathbf{s} = (H, A)$ and $\mathbf{s} = (L, R)$.

[3]Just as in Section 2.2.1, we abuse notation by denoting the eventual actions played out rather than the full strategy of the responder, which is expressed as a function of the proposer's move. In this sense, there are more than two equilibria for these cases, but multiple equilibria result in the same eventual actions.

However, since the $\mathbf{s} = (L, R)$ equilibrium results when the responder threatens to reject the high fraction, only $\mathbf{s} = (H, A)$ is subgame perfect.

The final case, where $\alpha > 1/2$ and $\omega > \gamma$, has nearly the same analysis. Even though the proposer now prefers to offer the high fraction, the equilibria are still $\mathbf{s} = (H, A)$ and $\mathbf{s} = (L, R)$. Again, the $\mathbf{s} = (L, R)$ results only when the responder threatens to reject the high fraction, so only $\mathbf{s} = (H, A)$ is subgame perfect.

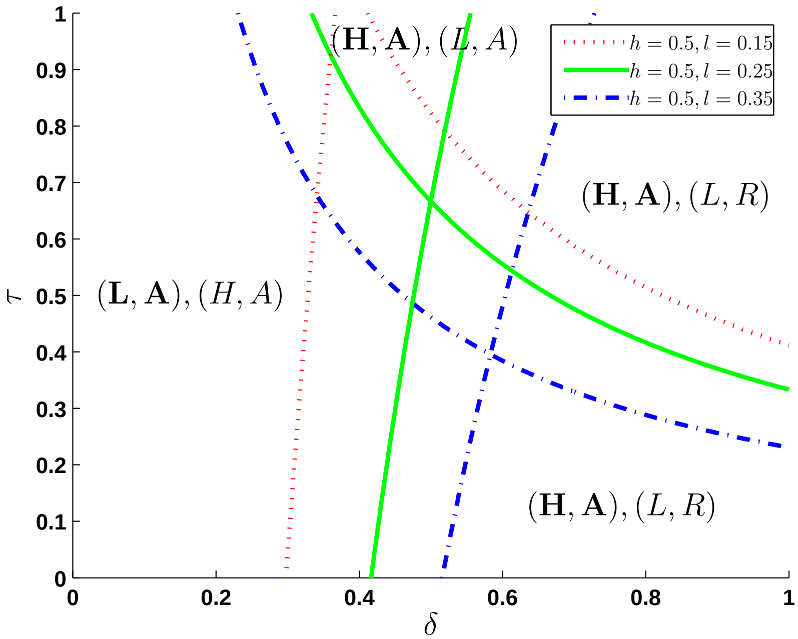

Figure 5.4: Ultimatum game equilibria. Subgame perfect equilibria are bold.

Since the game is played sequentially, the subgame perfect equilibrium depends more on δ than on τ. If the responder's indignation is sufficiently high that (L, R) is preferred to (H, A), then it can credibly threaten to reject the low fraction, and the proposer must offer the high fraction to obtain any payoff. Thus, by changing δ, the responder can force the overall outcome to (H, A). Such reasoning leads us to consider evolutionary games in the attitudes τ and δ similar to the indirect evolutionary dynamics discussed in [29, 30] or the attitude dynamics presented in Chapter 4. However, we do not consider such games in this thesis.

Chapter 6

Conclusion

6.1 Summary

We have sought to combine the merits of satisficing game theory with the more traditional non-cooperative game theory. This effort was motivated by the fact that satisficing players—while socially flexible due to the probabilistic structure of their utilities—may behave incoherently because they consider the preferences of other players, but not the strategies they play. We have presented two distinct methods of accomplishing this task, both of which allow players to consider the actions of other players. We have applied these to the Ultimatum game, which is problematic under both the non-cooperative and satisficing frameworks.

In our first approach, we let the players' attitudes to evolve according to the replicator dynamics of evolutionary game theory. This allows players to adapt their attitudes in order to improve raw payoffs, adding a "layer" of self-interest to a satisficing game. The attitude dynamics leads players toward a Nash equilibrium in their attitudes—rather than in their actions—that we term an *attitude equilibrium.* In the Ultimatum game, attitude equilibria exist where both the high fraction and the low fraction are accepted. However, the solutions where the responder rejects the offer—which were possible under the original satisficing formulation—are eliminated by the attitude dynamics. Incoherent behaviors are eliminated, while social flexibility is preserved.

Our second approach is to introduce conditional utilities directly into non-cooperative games. We define conditional utilities over the entire pure-strategy space so that players' marginal utilities may consider other players' preferences as well as the strategies they choose. We may then apply traditional non-cooperative solution concepts to the marginal utilities such as the Nash equilibrium. We present a simple social utility model for the Ultimatum game, finding that we can change the subgame perfect equilibrium by adjusting the players' attitudes, allowing subgame perfect equilibria where both the high and low fraction are accepted. It is also possible for the rejection of the low fraction to be a Nash equilibrium, but this equilibrium is not subgame perfect.

As with the attitude dynamics, we avoid incoherent behaviors while keeping the flexibility of the conditional utilities. However, it is the author's opinion that this method is the superior approach to combining satisficing and non-cooperative game theory. Instead of introducing a "patch" on the satisficing model by injecting self-interest into of the existing structure, we embed the conditional structure directly into players' utilities. The blend of self-interest and consideration for others' preferences emerges naturally out of the structure of the utilities, and incoherent behaviors are avoided naturally since we can use traditional solution concepts related to the Nash equilibrium.

6.2 Further Research

This thesis leaves open several areas for future work, mostly related to the use of social utilities in non-cooperative games. First is the problem of solving for players' marginal utilities in complicated utility networks. In a two-player network, the cyclical nature of the conditional utilities leads to an eigenvector problem. If a unique eigenvector exists for the eigenvalue $\lambda = 1$, we can solve analytically for the players' marginal utilities. However, for a fully-connected network of three or more players—where each player's utility depends on the commitments of every other player—we cannot solve for the marginal utilities using a simple eigen-decomposition, even if unique utilities exist. It may not even be possible to solve for the marginals in closed form. Therefore, it may be fruitful to explore algorithms by which we may solve iteratively for players' marginal utilities and discover when unique marginals do not exist.

We also have restricted our attention to games where players' strategy sets are finite. One could investigate the consequences of generalizing to the case of uncountable strategy sets. The utilities would then take the form of probability *density* functions rather than mass functions. For acyclic networks, the extension seems straightforward. However, for cyclic networks, even with only two players, continuous strategy sets prevent us from solving for the marginals using a matrix eigendecomposition. However, the conditional utilities may define an infinite-dimensional operator which we may diagonalize to solve for the marginals. Failing this, we may also apply numerical techniques to approximate the players' marginal utilities.

Finally, we previously hinted at extending the conditional structure proposed in Chapter 5 by playing evolutionary games in the player's attitudes, allowing each player to adjust their attitudes in an attempt to force the equilibrium it prefers. Such an approach may provide a method for choosing players' attitudes according to the structure of the conditional utilities.

Bibliography

[1] D. Dugail, E. Feron, and K. Bilimoria, "Stability of intersecting aircraft flows using heading change maneuvers for conflict avoidance," in *American Control Conference*, 2002. 1

[2] S. Resmerita, M. Heymann, and G. Meyer, "A framework for conflict resolution in air traffic management," in *IEEE Conf. on Decision and Control*, 2003, pp. 2035–40. 1

[3] V. Rodoplu and T. Meng, "Core capacity of wireless ad hoc networks," in *Proc. of the 4th Internation Symposium on Wireless Personal Multimedia Communications (WPMC)*, vol. 1, Oct. 2002, pp. 247 –251. 1

[4] V. Srinivasan, P. Nuggehalli, C. F. Chiasserini, and R. R. Rao, "Cooperation in wireless ad hoc networks." INFOCOM 2003, 2003, pp. 808–817. 1

[5] Z. Han, Z. Ju, and K. J. R. Liu, "Fair multiuser channel allocation for OFDMA networks using Nash bargaining solutions and coalitions," *IEEE Trans. Communications*, vol. 53, no. 8, pp. 1366–1376, August 2005. 1

[6] M. F. Demirkol and M. A. Ingram, "Power-controlled capacity for interfering MIMO links," in *Proc. of IEEE Vehicular Technology Conference*, vol. 1, October 2001, pp. 187–191. 1

[7] J. F. Nash, "Non-cooperative games," *Annals of Mathematics*, vol. 54, no. 2, pp. 286–295, September 1951. 1, 9, 10

[8] R. Axelrod, *The Evolution of Cooperation*. New York: Basic Books, 1984. 1, 9

[9] A. K. Sen, "Rational fools: A critique of the behavioral foundations of economic theory," in *Scientific Models and Man*, H. Harris, Ed. Oxford: Clarendon Press, 1979, ch. 1. 2

[10] W. C. Stirling, *Satisficing Games and Decision Making: with applications to engineering and computer science*. Cambridge, UK: Cambridge Univ. Press, 2003. 2

[11] F. R. Johnson, J. C. Hill, J. K. Archibald, R. L. Frost, and W. C. Stirling, "A satisficing approach to free flight," in *IEEE Int. Conf. on Networking, Sensing, and Control*, March 2005, pp. 123–128. 2

[12] M. S. Nokleby and W. C. Stirling, "The stag hunt: A vehicle for evolutionary cooperation," in *Proceedings of The IEEE World Congress on Computational Intelligence*, Vancouver, BC, Canada, July 2006. 2

[13] R. Selten, "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragetragheit," *Zeitschrift für die gasamte Staatswissenschaft*, vol. 12, pp. 301–324, 1965. 12

[14] D. Fundenberg and J. Tirole, *Game Theory*, 2nd ed. Cambrige, MA: The MIT Press, 1991. 12

[15] J. Gale, K. Binmore, and L. Samuelson, "Learning to be imperfect: The ultimatum game," *Games and Economic Behavior*, vol. 8, pp. 56–90, 1995. 16, 57

[16] W. Güth and R. Tietz, "Ultimatum bargaining behavior: A survey and comparison of experimental results," *Journal of Economic Psychology*, vol. 11, pp. 417–449, 1990. 17

[17] K. Binmore, A. Shaked, and J. Sutton, "Testing noncooperative bargaining theory: A preliminary study," *American Economic Review*, vol. 75, pp. 1178–1180, 1985. 17

[18] E. Fehr and K. M. Schmidt, "A theory of fairness, competition, and cooperation," *Quarterly Journal of Economics*, vol. 114, no. 3, pp. 817–868, August 1999. 17

[19] W. C. Stirling, R. L. Frost, and H. L. Miller, "The ultimatum game: A study in social systems control," in *Proceedings of the International Conference on Control*, Glasgow, Scotland, 2006. 17, 23, 57

[20] R. D. Luce and H. Raiffa, *Games and Decisions*. New York: John Wiley, 1957. 19

[21] W. C. Stirling, "Social utility functions, part 1 — theory," *IEEE Transactions on Systems, Man, and Cybernetics (Part C)*, vol. 35, no. 4, pp. 522–532, 2005. 20, 50, 75

[22] J. Pearl, *Probabilistic Reasoning in Intelligent Systems*. San Mateo, Ca: Morgan Kaufmann, 1988. 22

[23] J. W. Weibull, *Evolutionary Game Theory*. Cambridge, MA: MIT Press, 1995. 33, 35, 36

[24] B. Skyrms, *The Dynamics of Rational Deliberation*. Cambridge, MA: Harvard Univ. Press, 1990. 36

[25] J. von Neumann and O. Morgenstern, *The Theory of Games and Economic Behavior*. Princeton, NJ: Princeton Univ. Press, 1944, (2nd ed., 1947). 45

[26] M. H. DeGroot, *Optimal Statistical Decisions.* New York: McGraw Hill, 1970. 46

[27] W. C. Stirling, R. L. Frost, M. S. Nokleby, , and Y. Luo, “Multicriterion decision making with dependent preferences,” in *IEEE Symposium on Computational Intelligence in Multicriteria Decision Making*, Honolulu, HI, 2007. 47, 75

[28] C. D. Meyer, *Matrix Analysis and Applied Linear Algebra.* Philadelphia, PA: Society for Industrial and Applied Mathematics, 2000. 54

[29] W. Güth, “An evolutionary approach to explaining cooperative behavior by reciprocal incentives,” *International Journal of Game Theory*, vol. 24, no. 4, pp. 323–344, 1995. 65

[30] W. Güth and H. Kliemt, “The indirect evolutionary approach,” *Rationality and Society*, vol. 10, no. 3, pp. 377–399, 1998. 65

[31] R. T. Cox, “Probability, frequency, and reasonable expectation,” *Am. Jour. Phys.*, vol. 14, no. 1, pp. 1–13, January 1946. 75

[32] ——, *The Algebra of Probable Inference.* Baltimore, MD: The Johns Hopkins Univ. Press, 1961. 75

[33] J. Y. Halpern, “A counterexample to theorems of Cox and Fine,” *Jour. AI Research*, vol. 10, pp. 67–85, 1999. 75

[34] J. Aczel, *Lectures on Functional Equations and their Applications.* New York: Academic Press, 1966. 77

Appendix A

Justifying the Probabilistic Structure

In this appendix we give an axiomatic justification for the conditional utility structure used throughout this thesis, particularly the way in which we combine the conditional utilities to extract a marginal utility. A similar justification has been given by Stirling in several different contexts (see [21, 27]).

Our axiomatization bears resemblance to that given by Cox in [31, 32]. In these publications, Cox intends to prove that the use of a standard probability measure[1] is a necessary and sufficient condition to satisfy several seemingly natural and intuitive axioms. However, Cox's arguments are somewhat vaguely articulated, and attempts to make Cox's theorem rigorous have been problematic (see [33] for a detailed description of these difficulties).

Our approach is separate from that of Cox. We use a different set of axioms to justify a somewhat weaker result. In our approach, we deal with functions which take point-valued arguments (rather than measures defined on sets as in Cox's treatment), allowing us to avoid the issues raised regarding Cox's result.

In this appendix, we use the conditional-utility notation used in Chapter 5, although the result also applies to the probabilistic structure of satisficing game theory given in Chapter 3. We extend our notation to deal with the utilities of sets of players. For disjoint sets $J, K \subset I$, let π_{JK} denote the joint utility of the players in $J \cup K$, and let $\pi_{J|K}$ denote the joint utility of the players in J conditioned on the commitments of the players in K.

[1] Or, more precisely, the use of a measure isomorphic to a probability measure.

Axiom 1 (Conditioning) *Players may express their utilities conditionally as functions of the (hypothetical) commitments of other players.*

This axiom simply posits the existence of conditional utilities of the form $\pi_{J|K}$.

Axiom 2 (Normalization) *Any utility of the form $\pi_{J|K}$ must be non-negative and sum to unity; that is,*

$$\sum_{\mathbf{s}_j} \pi_{J|K}(\mathbf{s}_j|\mathbf{s}_k) = 1, \forall \mathbf{s}_k \tag{A.1}$$

and

$$\pi_{J|K}(\mathbf{s}_j|\mathbf{s}_k) \geq 0, \forall \mathbf{s}_j, \mathbf{s}_k. \tag{A.2}$$

Normalization is a fairly mild constraint. As we saw in Chapter 5, any von Neumann-Morgenstern utility can be transformed into a normalized utility via a positive affine transformation. Further, normalization limits the interpersonal comparison of utilities that results from the conditional utility structure. Normalized utilities are expressions of dimensionless *relative* benefit rather than utility on a personal—and therefore arbitrary—scale.

Axiom 3 (Endogenous aggregation) *For disjoint subsets $J, K \subset I$, the joint utility π_{JK} is endogenously aggregated if there is some function $F[\cdot,\cdot]$, nondecreasing in both arguments, such that*

$$\pi_{JK} = F[\pi_{J|K}, \pi_K] = F[\pi_{J'|K'}, \pi_{K'}], \tag{A.3}$$

where J' and K' are arbitrary disjoint subsets of I such that $J' \cup K' = J \cup K$.

The endogeny axiom ensures that the joint utilities (and, as we shall see, the resulting marginals) are independent of the "framing" of the problem. For example, if we have two players, where one player's utility is conditioned on the commitments of the other, the joint utility is $\pi_{12} = F[\pi_{1|2}, \pi_2]$. However, endogeny ensures that we can frame the problem differently—where player 2's utility is conditioned on the commitment of player 1—and get the same joint utility $\pi_{12} = F[\pi_{2|1}, \pi_1]$.

Theorem A.1 (The Aggregation Theorem) *The joint utility π_{JK} is endogenously aggregated if and only if $F[\cdot,\cdot]$ is isomorphic to multiplication; that is, there is a bijection f such that, for any endogenous aggregation $F[\cdot,\cdot]$,*

$$f(\pi_{JK}) = f(F[\pi_{J|K}, \pi_K]) = \pi_{J|K}\pi_K. \tag{A.4}$$

Proof: Sufficiency is established by setting $F[x, y] = xy$ and requiring that the utilities be non-negative. $F[\cdot, \cdot]$ is nondecreasing in both arguments, and since multiplication is associative, there exist utilities $\pi_{K|J}$ and π_J such that $\pi_{K|J}\pi_J = \pi_{J|K}\pi_K$.

To show necessity, let J, K, L be mutually disjoint subsets of I. Define the utilities π_{JKL}, $\pi_{J|KL}$, $\pi_{JK|L}$, π_{JK}, $\pi_{J|K}$, and π_J. If $F[\cdot, \cdot]$ is an endogenous aggregation function, it must satisfy

$$\pi_{JKL} = F[\pi_{J|KL}, \pi_{KL}] = F[\pi_{JK|L}, \pi_L]. \tag{A.5}$$

However, we also have

$$\pi_{KL} = F[\pi_{K|L}, \pi_L] \tag{A.6}$$

and

$$\pi_{JK|L} = F[\pi_{J|KL}, \pi_{K|L}]. \tag{A.7}$$

Substituting (A.6) and (A.7) into (A.5) gives us

$$F\left[\pi_{J|KL}, F[\pi_{K|L}, \pi_L]\right] = F\left[F[\pi_{J|KL}, \pi_{K|L}], \pi_L\right]. \tag{A.8}$$

In terms of general arguments, we can rewrite (A.8) as

$$F\left[x, F(y, z)\right] = F\left[F(x, y), z\right], \tag{A.9}$$

which is the well-studied associativity equation. It is shown in [34] that the general solution to (A.9) is isomorphic to multiplication: for any $F[\cdot, \cdot]$ satisfying (A.9), there exists a bijection f such that

$$f[F(x, y)] = f(x)f(y). \tag{A.10}$$

If we choose $F(x, y) = xy$, $f(\cdot)$ is the identity function, giving us $\pi_{JK} = F[\pi_{J|K}, \pi_K] = \pi_{J|K}\pi_K$. Again we require that the utilities be non-negative to make $F(\cdot, \cdot)$ nondecreasing in both arguments. Any other choice of $F(\cdot, \cdot)$ must be isomorphic to such an aggregation. ■

Theorem A.1 gives at least partial justification for the use of $\pi_{JK} = \pi_{J|K}\pi_K$ (called the *chain rule* or the *product rule* in probability theory) in aggregating conditional utilities into a joint utility. The choice of the product rule, combined with the other axioms, necessarily leads to the formation of marginal utilities by summing over the joint utility. For players 1 and 2, summing over the joint utility of both players gives

$$\sum_{\mathbf{s}_1 \in S} \pi_{12}(\mathbf{s}_1, \mathbf{s}_2) = \sum_{\mathbf{s}_1 \in S} \pi_{1|2}(\mathbf{s}_1|\mathbf{s}_2)\pi_2(\mathbf{s}_2) \tag{A.11}$$

$$= \pi_2(\mathbf{s}_2) \sum_{\mathbf{s}_1 \in S} \pi_{1|2}(\mathbf{s}_1|\mathbf{s}_2). \tag{A.12}$$

Since the conditional utility $\pi_{1|2}$ must sum to unity, we get the following expression for the marginal utility $\pi_2(\mathbf{s}_2)$:

$$\pi_2(\mathbf{s}_2) = \sum_{\mathbf{s}_1 \in S} \pi_{12}(\mathbf{s}_1, \mathbf{s}_2). \tag{A.13}$$

Since endogenous aggregation also allows us to partition π_{12} into the product of $\pi_1(\mathbf{s}_1)$ and $\pi_{2|1}(\mathbf{s}_2|\mathbf{s}_1)$, we can also write player 1's marginal utility as

$$\pi_1(\mathbf{s}_1) = \sum_{\mathbf{s}_2 \in S} \pi_{12}(\mathbf{s}_1, \mathbf{s}_2). \tag{A.14}$$

In our treatment, we have required that the joint utilities be formed by an interchangeable aggregation function. If we define $\pi_{12} = \pi_{1|2}\pi_2$, there must also exist utilities $\pi_{2|1}$ and π_1 such that $\pi_{12} = \pi_{2|1}\pi_1$. How are these "alternative" utilities related to the original conditional utility formulation? By setting the two aggregations equal to each other, we see that they are related by the familiar Bayes' rule from probability theory:

$$\pi_{2|1}(\mathbf{s}_2|\mathbf{s}_1)\pi_1(\mathbf{s}_1) = \pi_{1|2}(\mathbf{s}_1|\mathbf{s}_2)\pi_2(\mathbf{s}_2), \tag{A.15}$$

$$\pi_{2|1}(\mathbf{s}_2|\mathbf{s}_1) = \frac{\pi_{1|2}(\mathbf{s}_1|\mathbf{s}_2)\pi_2(\mathbf{s}_2)}{\pi_1(\mathbf{s}_1)}. \tag{A.16}$$

The conditionals $\pi_{2|1}$ and π_1 encode the same dependency information as $\pi_{1|2}$ and π_2, and result in the same joint and marginal utilities. It is interesting to note that we can model the players two different ways—one in which player 1 conditions its utility on player 2's commitments, and one in which player 2 conditions its utility on player 1's commitments—and end up with identical joint and marginal utilities, and therefore identical resulting behaviors.

Appendix B

Marginal Utilities for the Ultimatum Game

First, we solve for the proposer's marginal utilities, which must satisfy the equation $\boldsymbol{\pi}_1 = \mathbf{\Pi}_{1|2}\mathbf{\Pi}_{2|1}\boldsymbol{\pi}_1$. Our first step is to find the product $\mathbf{\Pi}_{1|2}\mathbf{\Pi}_{2|1}$, which is

$$\mathbf{\Pi}_{1|2}\mathbf{\Pi}_{2|1} = \frac{1}{(2-l-h)(l+h)} \begin{bmatrix} 1-h & (1-h)(1-\delta)+\delta(1-h+\tau(1-l)) & 1-h & (1-h)(1-\delta)+\delta(1-h+\tau(1-l)) \\ 1-l & (1-l)(1-\delta)+\delta(1-h+\tau(1-l)) & 1-l & (1-l)(1-\delta)+\delta(1-h+\tau(1-l)) \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}. \tag{B.1}$$

We already know that $\lambda = 1$ is an eigenvalue for this stochastic matrix. To find the corresponding eigenvectors, we must solve the equation $(\mathbf{\Pi}_{1|2}\mathbf{\Pi}_{2|1} - \mathbf{I})\boldsymbol{\pi}_1 = \mathbf{0}$:

$$\begin{bmatrix} 1-h-(2-l-h) & (1-h)(1-\delta)+\delta(1-h+\tau(1-l)) & (1-h) & (1-h)(1-\delta)+\delta(1-h+\tau(1-l)) \\ 1-l & (1-l)(1-\delta)+\delta(1-h+\tau(1-l))-(2-l-h) & (1-l) & (1-l)(1-\delta)+\delta(1-h+\tau(1-l)) \\ 0 & 0 & -(2-l-h) & 0 \\ 0 & 0 & 0 & -(2-l-h) \end{bmatrix} \begin{bmatrix} \pi_{1,1} \\ \pi_{1,2} \\ \pi_{1,3} \\ \pi_{1,4} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}. \tag{B.2}$$

Fortunately, it's clear that $\pi_{1,3} = \pi_{1,4} = 0$, leaving us with the system of equations

$$(1 - h - 2 + l + h)\pi_{1,1} + ((1 - h)(1 - \delta) + \delta(1 - h + \tau(1 - l)))\pi_{1,2} = 0, \tag{B.3}$$

$$(1 - l)\pi_{1,1} + ((1 - l)(1 - \delta) + \delta(1 - l - \tau(1 - l)) - 2 + l + h)\pi_{1,2} = 0. \tag{B.4}$$

These two equations simplify to the single constraint

$$(1-l)\pi_{1,1} = (1-h+\tau\delta(1-l))\pi_{1,2}. \tag{B.5}$$

This equation defines the family of eigenvectors for $\lambda = 1$. Fortunately, there is only one linearly independent eigenvector, so unique marginal utilities exist. The proposer's marginal utility, which is the normalized eigenvector, is

$$\boldsymbol{\pi}_1 = \frac{1}{2-l-h+\tau\delta(1-l)} \begin{bmatrix} 1-h+\tau\delta(1-l) \\ 1-l \\ 0 \\ 0 \end{bmatrix}.$$

To find the responder's marginal utilities, we could solve for the eigenvectors of $\mathbf{\Pi}_{2|1}\mathbf{\Pi}_{1|2}$. However, it is much simpler and more direct to use the fact that $\boldsymbol{\pi}_2 = \mathbf{\Pi}_{2|1}\boldsymbol{\pi}_1$. Since $\boldsymbol{\pi}_1$ is now well-defined, we can evaluate the result:

$$\boldsymbol{\pi}_2 = \frac{1}{l+h} \begin{bmatrix} h & h(1-\delta) & h & h(1-\delta) \\ l & l(1-\delta) & l & l(1-\delta) \\ 0 & 0 & 0 & 0 \\ 0 & \delta(l+h) & 0 & \delta(l+h) \end{bmatrix} \frac{1}{2-l-h+\tau\delta(1-l)} \begin{bmatrix} 1-h+\delta\tau(1-l) \\ 1-l \\ 0 \\ 0 \end{bmatrix} \tag{B.6}$$

$$= \frac{1}{(l+h)(2-l-h+\tau\delta(1-l))} \begin{bmatrix} h(2-h-l+\delta(1-l)(\tau-1)) \\ l(2-h-l+\delta(1-l)(\tau-1)) \\ 0 \\ \delta(l+h)(1-l) \end{bmatrix}. \tag{B.7}$$