Proceedings of the 2007 IEEE Symposium on
Foundations of Computational Intelligence (FOCI 2007)

Attitude Adaptation in Satisficing Games

Matthew S. Nokleby and Wynn C. Stirling
Electrical and Computer Engineering Department, Brigham Young University, Provo, Utah, 84602, USA
Email: {nokleby,wynn}@ee.byu.edu

***Abstract*—Satisficing game theory is an alternative to traditional game theory which offers more flexibility in modeling players in social interactions. Unfortunately, satisficing players with conflicting attitudes may implement dysfunctional behaviors, resulting in poor performance. We present a method based on evolutionary game theory by which players may adapt their attitudes to their circumstances, allowing them to overcome dysfunction. Additionally, we extend the Nash equilibrium concept to satisficing games, showing that the method presented leads the players toward equilibrium in their attitudes. These ideas are applied to the Ultimatum game as a simple example.**

I. Introduction

Game-theoretic models are often used to construct societies of artificial agents. Classical game theory is founded upon rational choice, or *individual rationality*. Individual rationality asserts that players are concerned solely with maximizing individual payoff. Rationality leads players towards *Nash equilibria*, which are strategies such that no single player can improve its payoff by changing strategies. Unfortunately, such self-interested behavior places significant limitations in terms of the players' social interactions. It is typically difficult to engender cooperation and other social behaviors under classical game theory, causing it to come under criticism in recent years [1, 2].

Satisficing game theory [3] is an alternative to classical decision theory that addresses these social shortcomings. Its approach is fundamentally different, relaxing the assumption of individual rationality and instead endowing players with *social rationality*. Social rationality provides a natural mechanism by which players may consider the preferences of others in formulating their utility functions. Players' utilities are expressed conditionally, and may depend on other players' preferences for action.

Satisficing models are often successful in overcoming the social hurdles presented by classical theory. Players can exhibit sophisticated social behaviors such as cooperation, altruism, negotiation, and compromise [4, 5]. However, satisficing theory also presents its own set of challenges. As in real-life social situations, satisficing communities may behave dysfunctionally. When players with incompatible attitudes are grouped together, they may choose incoherent behaviors that lead to a different breed of poor performance.

In this paper, we aim to "bridge the gap" between classical and satisficing game theory by equipping satisficing players with a form of individual rationality. We provide a method whereby players may modify their attitudes according to the game structure and the attitudes of other players. In this method, which is based upon the multipopulation replicator dynamics from evolutionary game theory [6], players modify their attitudes to improve their individual payoffs, allowing dysfunctional societies to adapt. The result is a blend of the two decision theories: players retain the conditional utility structure of satisficing theory while seeking to improve individual payoff. The dynamics lead the society towards a *social Nash equilibrium*, which is essentially a Nash equilibrium in players' attitudes rather than in their actions.

In Section II we familiarize the reader with the basics of satisficing game theory. In Section III we present the Ultimatum game under the classical and satisficing frameworks, which will be used as an example throughout the paper. In Section IV we define the social Nash equilibrium and present the attitude dynamics. We apply these to the satisficing Ultimatum game to explore the dynamics issues from multipopulation replicator dynamics such as initial conditions and varying adaptation rates in Section V. We draw our conclusions in Section VI.

II. Satisficing Game Theory

In satisficing game theory, player eschew individual rationality and instead exhibit social rationality, which comprises two basic tenets: (a) players may consider the preferences of others in constructing their utilities, and (b) instead of seeking for maximal utility, players are content with actions that are "good enough." As mentioned in the introduction, the first tenet is accommodated by introducing conditional utilities (which are called *social utilities*). Social utilities are patterned after probability mass functions, allowing concepts such as conditioning and independence—which are typically used only in the probabilistic sense—to be extended to decision-making problems.

To accommodate the second tenet, each player possesses two social utilities. For Player i, we denote the *selectability* and *rejectability* functions $p_{S_i}(u_i)$ and $p_{R_i}(u_i)$, respectively, where the pure strategies u_i are elements of Player i's strategy space U_i. Each function quantifies the preferences for the strategy u_i from a different perspective. The selectability quantifies them according to their benefits, while the rejectability quantifies them in terms of their costs. Since $p_{S_i}(\cdot)$ and $p_{R_i}(\cdot)$ are mass functions, they are normalized and nonnegative, and thus give measures of the *relative* benefits or costs for implementing a strategy. The two utilities allow a precise definition of "good enough" for Player i. Define the *individually satisficing set* as all strategies for which the relative benefits are at least as great as the relative costs:

$$\Sigma_i = \{u_i \in U_i : p_{S_i}(u_i) \geq q p_{R_i}(u_i)\}, \quad (1)$$

where q is the *index of caution*. Essentially q is a tuning parameter that with which Player i can alter its definition of "good enough." Since $p_{S_i}(\cdot)$ and $p_{R_i}(\cdot)$ are normalized mass functions, setting $q = 1$ ensures that Σ_i contains at least one element.

For multiple players, we typically wish to consider strategy profiles $\mathbf{u} = (u_1, u_2, \ldots u_n)$ that are acceptable to each player. Define the *satisficing rectangle* as

$$\Re_{12\cdots n} = \Sigma_1 \times \Sigma_2 \times \cdots \times \Sigma_n, \tag{2}$$

the Cartesian product of the individual players' individually satisficing sets. Any $\mathbf{u} \in \Re_{12\cdots n}$ is simultaneously acceptable to each player.

In 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 [7, 8]. 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 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 [7] to analyze the community.

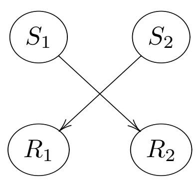

Fig. 1. A simple praxeic network.

In discussing the players' social utilities, we retain the terminology of probability theory. In the community from Figure 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, $p_{R_1} = \sum_{u_2 \in U_2} p_{R_1|S_2}(v_1|u_2) p_{S_2}(u_2)$. 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 may be expressed directly, without conditioning.

With the marginal and conditional utilities defined, we can form 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 $(\mathbf{X}, \mathbf{U}, p_{S_1 \cdots S_n R_1 \cdots R_n})$, where $\mathbf{X}$ is the set of players, $\mathbf{U}$ is the Cartesian product of the players' strategy spaces, and $p_{S_1 \cdots S_n R_1 \cdots R_n}$ is the interdependence function. From this information, all necessary marginal utilities can be computed and solution concepts such as the satisficing rectangle can be determined.

III. The Ultimatum Game

A. Classical Model

The Ultimatum Game has become a common example for illustrating the weakness of classical game theory as a model for human behavior [9, 10]. The game consists of two players: the proposer (Player 1, referred to for convenience as a male) and the responder (Player 2, a female). 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 she does, they divide the dollar as proposed. If not, each player receives nothing.

The proposer's (uncountable) strategy space is the interval $[0, 1]$, which makes analysis difficult. Thus, we will pattern our discussion after Gale et al. [11] and examine the two-option minigame, which captures the "heart" of the game while simplifying analysis. In this minigame, the proposer offers either a high or low fraction (h or l) to the responder, who again may choose to accept or reject (a or r) the offer. The payoffs for this minigame are shown in Table I.[2]

TABLE I
Payoff Matrix for the Ultimatum minigame.

	Responder	
Proposer	a	r
h	$(1-h,\ h)$	$(0,\ 0)$
l	$(1-l,\ l)$	$(0,\ 0)$

As long as $h > l$, the unique Nash equilibrium for is for the proposer to offer the low fraction and for the proposer to accept 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. These results suggest that players' desire to maximize payoffs are tempered by social considerations. Such considerations are difficult to model under classical game theory, which has prompted Stirling et al. to cast the Ultimatum game as a satisficing game [12], where social factors may influence players' decisions. We briefly present their model, which we will use throughout the paper.

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

[2] We use h and l to denote both the strategies of offering the two fractions and the numerical value of each fraction.

B. Satisficing Model

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

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

$$p_{S_1}(u_1) = \begin{cases} 1-\tau, & \text{for } u_1 = h \\ \tau, & \text{for } u_1 = l \end{cases}. \tag{3}$$

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

$$p_{R_1}(v_1) = \begin{cases} \tau(1-\delta), & \text{for } u_1 = h \\ 1-\tau(1-\delta), & \text{for } u_1 = l \end{cases}. \tag{4}$$

The responder, who plays second, conditions her utility functions on those of the proposer. She wishes to maintain her fraction of the dollar, but reserves the right to punish an intemperate proposer. Her conditional rejectability is

$$p_{S_2|S_1}(u_2|u_1) = \begin{cases} 1, & \text{for } u_2 = a|u_1 = h \\ 0, & \text{for } u_2 = r|u_1 = h \\ 1-\delta, & \text{for } u_2 = a|u_1 = l \\ \delta, & \text{for } u_2 = r|u_1 = l \end{cases}. \tag{5}$$

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 ($\tau = 1$), she prefers to reject the offer according to her indignation index δ. Her conditional rejectability essentially encodes the same preferences, and is given by

$$p_{R_2|S_1}(v_2|u_1) = \begin{cases} 0, & \text{for } v_2 = a|u_1 = h \\ 1, & \text{for } v_2 = r|u_1 = h \\ \delta, & \text{for } v_2 = a|u_1 = l \\ 1-\delta, & \text{for } v_2 = r|u_1 = l \end{cases}. \tag{6}$$

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

$$p_{S_2}(u_2) = \begin{cases} 1-\tau\delta, & \text{for } u_2 = a \\ \tau\delta, & \text{for } u_2 = r \end{cases} \tag{7}$$

$$p_{R_2}(v_2) = \begin{cases} \tau\delta, & \text{for } v_2 = a \\ 1-\tau\delta, & \text{for } v_2 = r \end{cases}. \tag{8}$$

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

$$p_{S_1 S_2 R_1 R_2}(u_1, u_2, v_1, v_2) = p_{S_2|S_1}(u_2|u_1) p_{R_2|S_1}(v_2|u_1) \cdot p_{S_1}(u_1) p_{R_1}(v_1). \tag{9}$$

C. The Satisficing Rectangle

Now that the players' utility functions are defined, we can examine their actions according to the satisficing rectangle. Recall that an strategy u_i is individually satisficing for Player i if $p_{S_i}(u_i) \ge q p_{R_i}(u_i)$. In Figure 2 we set $q = 1$ and show the satisficing rectangle as functions of τ and δ. Four possibilities result depending on the players' attitudes. In the (l, a) region, the proposer is sufficiently greedy and the responder sufficiently conciliatory that the low fraction is accepted. In the (h, a), where τ is lower and/or δ is higher, the high fraction is accepted. In regions (h, r) and (l, r), the responder is sufficiently indignant that the offers are rejected.

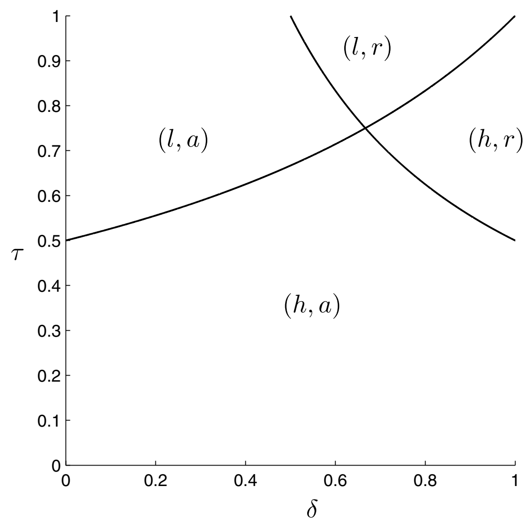

Fig. 2. (τ, δ) regions for the satisficing rectangles.

Figure 2 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 her actions are not simply (or even primarily) based on the offer proposed. For $\tau = 0.75$, she accepts the low offer. However, if the proposer's intemperance index increases much higher, she refuses. She 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 his actions, but for his *attitudes*. Such social behaviors may be desirable in the synthesis of artificial decision-makers, and are extremely difficult to model under classical game theory.

However, this framework also allows for undesirable behavior. Consider the (h, r) region, where the responder rejects the high fraction. Her intemperance is sufficiently high that she punishes even a moderately intemperate proposer. Such dysfunctional behavior is somewhat common in satisficing games and is 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 of the Nash equilibrium. Under classical game theory, 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 satisficing theory simply tells us that players with incompatible attitudes (in this case a highly indignant responder with a moderately intemperate proposer) may act incoherently. Unfortunately, in designing artificial societies, we typically prefer to avoid incoherent behaviors, sociologically justifiable or not. And while there are refinements to the satisficing rectangle that help identify dysfunctional societies, satisficing theory so far has nothing to say about how they might resolve their dysfunction.

IV. Attitude Dynamics

A. Social Nash Equilibria

To introduce the social Nash 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 utility function which is based solely on the strategies that the players implement.

Definition 4.1: An *augmented satisficing game* is a 5-tuple $(\mathbf{X}, \mathbf{U}, p_{S_1 \cdots S_n R_1 \cdots R_n}, \mathbf{A}, \boldsymbol{\pi}(\mathbf{u}))$. The first three elements are the set of players, the product strategy space, and interdependence function as normal. Additionally, we introduce the product attitude space $\mathbf{A} = A_1 \times A_2 \times \cdots \times A_n$ containing the attitudes that the players may adopt, and $\boldsymbol{\pi}(\mathbf{u})$, a vector payoff function which describes the raw payoff to the players for implementing the strategy profile $\mathbf{u} \in \mathbf{U}$.

In order to be able to augment a satisficing game, the players' attitudes must be specified as distinct parameters in the players' social utilities. Further, we must be able to construct a "raw" payoff function that is separate from the social utilities. While this may not be possible for all satisficing games, the extension is straightforward for the Ultimatum game. The players' attitudes are the intemperance and indignation indices τ and δ, yielding a product attitude space of $\mathbf{A} = [0,1] \times [0,1]$. The payoff function $\boldsymbol{\pi}(\mathbf{u})$ is described by the payoff matrix in Table I.

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.[3] 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 $\mathbf{a} \in \mathbf{A}$. That is, we may think of an augmented satisficing game as a classical game where players' payoffs are determined by the attitudes they adopt, rather than the strategies they implement. Players may now consider *changing* their attitudes if they result in poor payoff. This concept provides the motivation for the social Nash equilibrium.

Definition 4.2: An attitude profile $\mathbf{a} \in \mathbf{A}$ is a *social Nash equilibrium* if no single player can improve its payoff by changing attitudes; that is $\pi_i(a_1, \ldots, a_i, \ldots, a_n) \geq \pi_i(a_1, \ldots, a'_i, \ldots, a_n)$ for all $a'_i \in A_i$, $i = 1, 2, \ldots, n$.

In his original paper [13], Nash proves that at least one equilibrium exists for any game with finite strategy spaces. However, this equilibrium may only exist in *mixed strategies*, which are probability distributions over players' strategy spaces. In the case of mixed strategies (as opposed to pure strategies), Nash equilibria are points where no player can improve its *expected* payoff by deviating. By the same argument, a social Nash equilibrium exists for every game with a finite attitude space, but may exist only in a probability distribution over the players' attitude spaces which we will refer to as *mixed attitudes*.

For the Ultimatum game, even though the attitude spaces are infinite, it is straightforward to show that social Nash equilibria exist in "pure" attitudes. In Figure 3, the social Nash equilibria are the shaded regions. If the players' attitude vector lies in these regions, there is no incentive for either player to change attitudes. Consider the shaded region in (h, a). The responder receives maximum payoff, and therefore has no reason to deviate. Similarly, the proposer cannot improve his payoff by changing τ. While the proposer is not earning his maximum payoff, changing τ can only drive the responder to reject the offer, resulting in lower payoff. Similarly, in the (l, a) shaded region, altering δ can only result in the offer being rejected. For any other attitude vector in $\mathbf{A}$, at least one player stands to increase payoff by changing attitudes.

The players' social Nash equilibria result in the acceptance of either the high or low offer. By contrast, the classical Nash equilibrium results in only the low offer being accepted. This new concept provides a useful juxtaposition of social 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 social Nash equilibrium concept does not tell us which equilibrium these players will adopt. It simply says that if their attitudes lie in the equilibrium region, neither player has incentive to deviate. Therefore, we turn to evolutionary mechanisms to explore which equilibrium will result under different conditions.

B. Replicator Dynamics

The multipopulation replicator dynamics describes evolution in asymmetric games where players are selected from separate populations. In the Ultimatum game, for example, there are two populations: a population of proposers and a

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

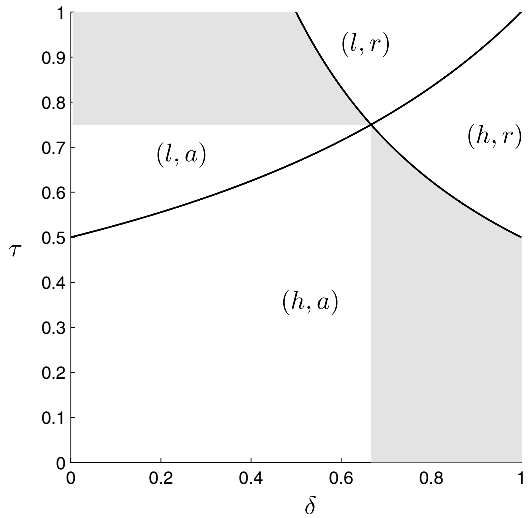

Fig. 3. Social Nash equilibria for the Ultimatum game.

population of responders. Each player is "pre-programmed" to play a particular strategy. In the classical game, the proposers are programmed to offer either the high or low fraction, and the responders are programmed either to accept or reject the offer. Players are randomly selected from the populations, play the game, and earn payoffs according to their payoff functions. They then (asexually) reproduce according to their payoffs: the number of offspring a player has is proportional to its payoffs. A player's offspring always play the same strategy as the parent.

The replicator dynamics examines the ratio of players playing particular strategies. There is in general an arbitrary number of populations, but here we will restrict our attention to the two-population case. Define two normalized state vectors $\mathbf{x} = (x_1, x_2, \ldots, x_n)^T$ and $\mathbf{y} = (y_1, y_2, \ldots, y_m)^T$ for a game where Player 1 and Player 2 have n and m pure strategies to choose from, respectively. Each element x_i or y_j represents the *population shares* of i and j, or the fraction of Population 1 or 2 playing pure strategy i or j. We may equivalently regard $\mathbf{x}$ and $\mathbf{y}$ as probability distributions over the two populations.

The standard two-player replicator dynamics [6] are given by a system of $n + m$ differential equations

$$\dot{x}_i = \big[\pi_1(i, \mathbf{y}) - \pi_1(\mathbf{x}, \mathbf{y})\big] x_i \tag{10}$$

$$\dot{y}_j = \big[\pi_2(j, \mathbf{x}) - \pi_2(\mathbf{y}, \mathbf{x})\big] y_j, \tag{11}$$

where $\pi_1(i, \mathbf{y}) = \sum_k \pi_1(i, k) y_k$ is the expected utility of playing pure strategy i against a member of the population described by $\mathbf{y}$ and $\pi_1(\mathbf{x}, \mathbf{y}) = \sum_l \pi_1(l, \mathbf{y}) x_l$ is the average expected payoff for players in Population 1. If a particular strategy is more successful than average, its population share grows. If a strategy is unsuccessful, the players playing it are overwhelmed by the offspring of more successful players, and its population share diminishes.

The replicator dynamics is typically used to describe the evolution of the distributions of large populations. However, it can also be used as a Bayesian deliberation process [14] where two players may update their mixed strategies according to the mixed strategies of the other player. Here, we interpret $\mathbf{x}$ and $\mathbf{y}$ as the mixed strategies of Player 1 and Player 2. Each player updates its mixed strategy according to (10) and (11). It is shown in [6] that if all pure strategies are represented in the initial conditions $\mathbf{x}(0)$ and $\mathbf{y}(0)$, then any steady state of the dynamics is a Nash equilibrium in the players' strategies.

To extend this to the satisficing case, we operate the deliberation/replicator dynamics on the players' attitudes rather than the strategies they implement. We require that both players have finite attitude spaces so that $\mathbf{x}$ and $\mathbf{y}$ are finite-dimensional. Instead of representing mixed strategies, the state vectors represent "mixed" attitudes, and the dynamics allow the players to alter the probability with which they will exhibit the attitudes in their attitude spaces. The attitude dynamics is exactly as in equations (10) and (11), 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 a social Nash equilibrium.

V. Results

To apply the attitude dynamics to the Ultimatum game, we must quantize the players' attitude spaces. Each player's attitude space is $A = \{a_1, a_2, \ldots, a_{100}\}$, 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.

From the previous section, we know that, given well-behaved initial conditions, the steady state of the dynamics is a social Nash equilibrium in either pure or mixed attitudes. The pure-attitude equilibria are straightforward and have already been shown in Figure 3. While the mixed-strategy equilibria are more complicated, we still can extract a few simple and useful facts without overly complicated analysis. First, it's straightforward to show that any attitude profile with all of its probability within one of the equilibrium regions of Figure 3 is itself an equilibrium. It is also possible to have probability mass located in the non-equilibrium portions of the (l, a) and (h, a) regions. Fortunately, however, a social Nash equilibrium cannot have probability mass in either the (l, r) or (h, r) regions. When there is probability in those regions, both players can improve expected utility by modifying his or her mixed attitudes until there is no probability of rejecting the fraction. Therefore, the attitude dynamics are guaranteed to eliminate the dysfunctional behavior observed in Section III.

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 (10) and (11). 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

$$x_i(0) = \begin{cases} ce^{\lambda_1(a_i-\mu)}, & \text{for } a_i \le \mu \\ ce^{\lambda_2(\mu-a_i)}, & \text{for } a_i > \mu \end{cases}, \tag{12}$$

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. The exponential 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. This is important because nonzero initial conditions are necessary to ensure that the steady-state distribution is a social Nash equilibrium.

Also, the exponential distribution encourages players' distributions to "shift" to adjacent values rather than "jump" across the attitude space. The form of the replicator dynamics equations explains this. Equations (10) and (11) 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.

A. 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 such that $\mu_1 = \mu_2 = 0.2$ and $\sigma_1^2 = \sigma_2^2 = 0.001$. Initially, the players' attitudes almost invariably lead to the offer and acceptance of the high fraction. While such behavior is not necessarily dysfunctional, 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(a) shows the initial joint distribution of the player's attitudes. Since the responder is earning maximal payoff, she has no incentive to shift her attitudes. The proposer, however, can improve his payoff by increasing τ. In Figure 4(b), we see that the proposer shifts his attitudes such that the joint distribution peaks right on the boundary between the (l, a) and (h, a) regions of the satisficing rectangle. He has shifted his attitudes just enough move the players into the region where he gets maximum payoff.

Once the shift is made, however, the responder stands to gain by modifying her preferences. In Figure 4(c), we see the results of an "arms race": the responder slightly increases δ to move the players to the (h, a) region, prompting the proposer to increase τ. The players "walk" their attitudes along the high/low boundary until it intersects the accept/reject boundary. At this point (Figure 4(d)), neither player can improve payoff by changing attitudes, and the distribution becomes almost entirely focused on 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 (h, a) region preferred by the responder.

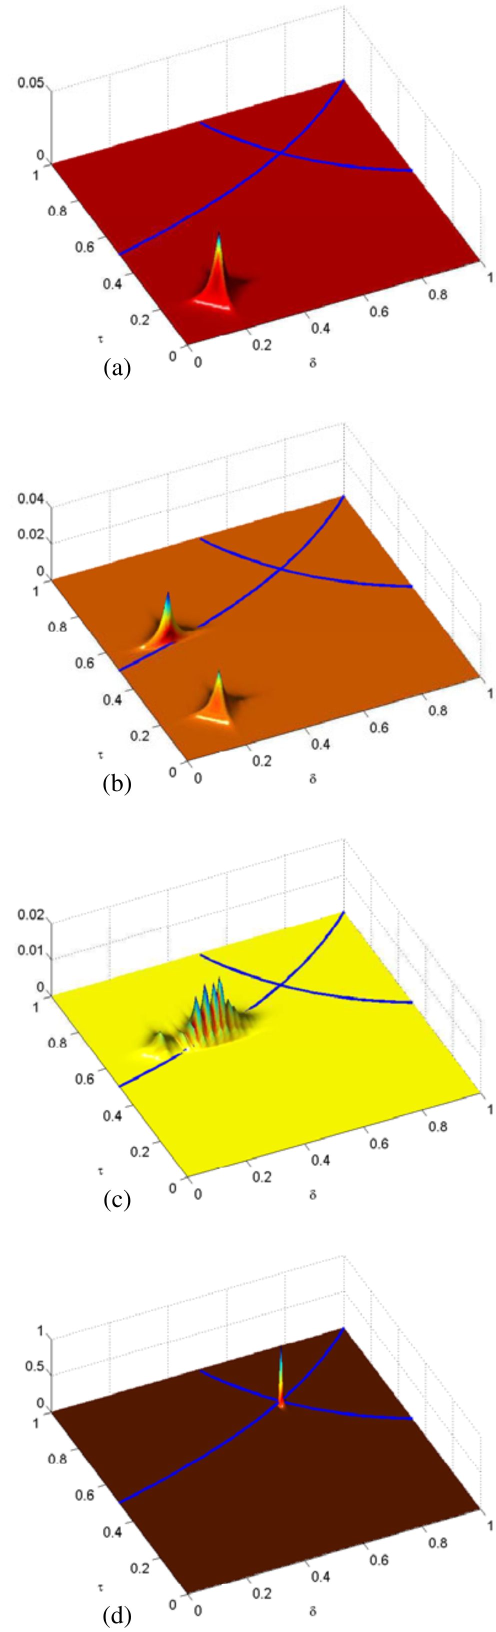

Fig. 4. "Arms race" joint attitude distribution for (a) $t = 0$, (b) $t = 35$, (c) $t = 60$, and (d) $t = 125$.

B. Adaptation Rates

An important consideration in multipopulation replicator dynamics is the relative size of the populations. If, for example, Population 1 is significantly smaller than Population 2, the players from Population 2 are paired less frequently, and the population shares for Population 2 evolve more slowly than for Population 1. While in the attitude dynamics the populations represent single players' attitude distributions, we still may consider the players' adaptation rates. A well-established principle in evolutionary game theory is that if all players' payoffs are multiplied by a constant β, the solution trajectories and steady state behavior remain unchanged; the dynamics simply progresses faster. However, if Player 1's payoffs are multiplied by β_1 and Player 2's payoffs are multiplied by β_2, the players adapt at different rates, and the final distribution may be different.

Consider the dysfunctional pair whose initial conditions are $\mu_1 = 0.8$, $\mu_2 = 0.9$, $\sigma_1^2 = \sigma_2^2 = 0.001$ (Figure 5(a)). Here, the responder rejects the high offer. The dynamics shift both of the players' attitudes toward a social Nash equilibrium, but which one? The proposer, of course, would prefer to end up in the (l, a) region, while the responder prefers (h, a). The answer lies with which one adapts most quickly. If $\beta_1 = \beta_2 = 1$, the responder is able to shift her attitudes more quickly than the proposer. She begins accepting the low offer, and the proposer no longer has any reason to adjust his attitudes, resulting in (l, a) as the steady state behavior (Figure 5(b).

However, if we let $\beta_1 = 5$, the proposer adapts more quickly. He reduces τ until the responder is willing to accept the high offer, leaving the players in the (h, a) equilibrium (Figure 5(c)). In this scenario, we get an interesting result: initially, both players have incentive to adapt to avoid losing the entire dollar. However, the player who adapts more slowly eventually ends up in his or her preferred equilibrium.

VI. Conclusion

In this paper we have presented a solution for overcoming dysfunction in satisficing players. Since players' utilities are functions of others' attitudes, rather than their actions, players with conflicting or incompatible attitudes may enact dysfunctional behaviors. In our solution, we augment the satisficing game with a classical payoff function, allowing us to examine the raw payoff to each player for exhibiting particular attitudes. This approach allows us to incorporate elements of individual rationality (classical game theory) and social rationality (satisficing game theory) into a single framework.

We define a social Nash equilibrium where no player in an augmented satisficing game can improve raw payoff by changing attitudes. We use the standard multipopulation replicator dynamics as a means for players to adjust their attitudes to the game and the other players' attitudes. If the players' initial distributions assign nonzero probability to each attitude in their attitude space, then any steady state of the dynamics is a social Nash equilibrium. The attitude dynamics can be used as a non-heuristic negotiation model for satisficing players: players enter the game with their attitude distributions formed independently. Through the dynamics, the players respond to each other's attitudes, adapting until a social Nash equilibrium results.

Finally, we note that this work does not address all sources of dysfunctional behavior. Players may still act poorly given noisy communication or incomplete knowledge of each other's preferences. However, our solution addresses an issue that is unique to the satisficing approach. Players' conflicting attitudes may be resolved as they adjust their preferences to improve payoff.

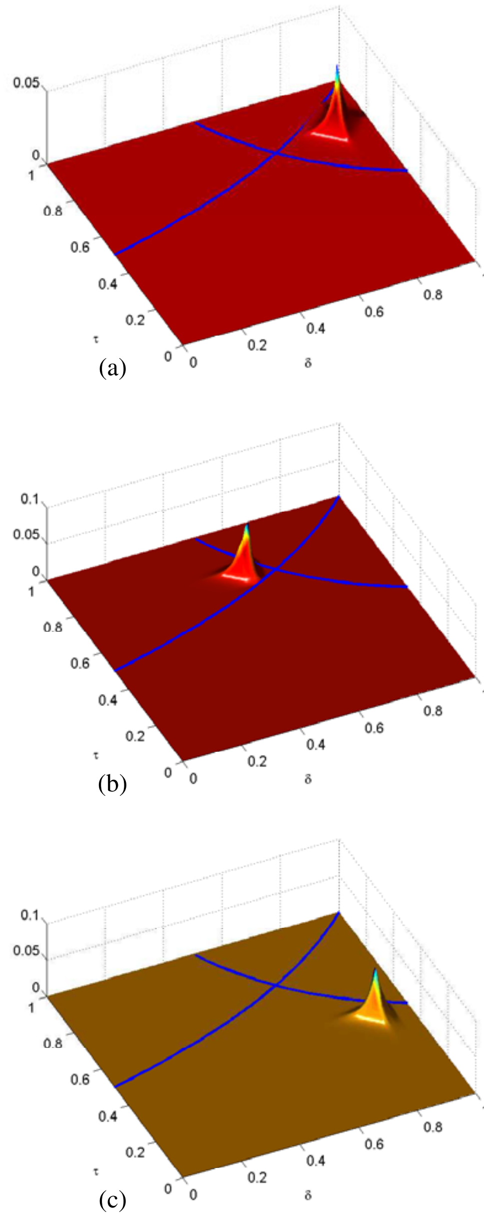

Fig. 5. Adaptation rates: (a) initial distribution, (b) steady state for $\beta_1 = \beta_2 = 1$, (c) steady state for $\beta_2 = 5$.

REFERENCES

[1] E. Sober and D. S. Wilson, *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge, MA: Harvard Univ. Press, 1998.

[2] J. J. Mansbridge, Ed., *Beyond Self-Interest*. Chicago: Univ. of Chicago Press, 1990.

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

[4] W. C. Stirling and R. L. Frost, "Social Utility Functions, Part 2 — Applications," *IEEE Transactions on Systems, Man, and Cybernetics (Part C)*, vol. 35, no. 4, pp. 533–543, 2005.

[5] J. K. Archibald, J. C. Hill, F. R. Johnson, and W. C. Stirling, "Satisficing negotiations," *IEEE Transactions on Systems, Man, and Cybernetics (Part C)*, vol. 36, no. 1, pp. 4–18, 2006.

[6] J. W. Weibull, *Evolutionary Game Theory*. Cambridge, MA: MIT Press, 1995.

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

[8] S. L. Lauritzen, *Graphical Models*. New York: Springer Verlag, 1996.

[9] 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.

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

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

[12] 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.

[13] J. F. Nash, "Non-cooperative games," *Annals of Mathematics*, vol. 54, pp. 289–295, 1951.

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