CONDITIONAL PREFERENCES FOR SOCIAL SYSTEMS

WYNN C. STIRLING and MICHAEL A. GOODRICH

Electrical & Computer Engineering Department, Brigham Young University, Provo, UT, USA 84602
Computer Science Department, Brigham Young University, Provo, UT, USA 84602

Abstract

The design of artificial decision-making systems must be founded on some notion of rationality. Conventional multi-agent decision-making methodologies, such as von Neumann-Morgenstern game theory, are based on the paradigm of individual rationality, which requires decision makers to take the action that is best for themselves, regardless of its effect on other decision makers. Relaxing the demand for the "best possible" decision, however, opens the way to accommodate the preferences of others. Satisficing game theory is a new approach to multi-agent decision making that permits decision makers to adjust their preferences in a controlled way to give consideration to others by permitting *conditional preferences* whereby a decision maker is able to adjust its preferences as a function of the preferences of others.

Keywords

multiple agents, satisficing, game theory, rationality

1 Introduction

Artificial decision systems will be useful to society if they extend the ability of humans to perform complex tasks. For such systems to be accepted and trusted by society, they must perform their tasks in ways that are compatible with the ways that humans would perform them. Although much research has concentrated on such applications, attempts to design systems that are compatible with human behavior have met with limited success.

To illustrate, consider the problem of allocating office space. The potential occupants will each have their own individual needs and preferences which should be respected to the extent possible. The relationships between the occupants maybe complex, however, and behavior models that account only for individual interest may be inadequate. Exclusive self interest fosters competition, exploitation, and even avarice, but those may not be the dominant attributes of the decision makers. Coordinated behavior, however, is difficult to express under a paradigm of individual rationality, since that paradigm does not permit adequate representation of deferential attitudes such as cooperation, unselfishness, and even altruism. We present an alternative notion of rationality that is designed to accommodate both individual and societal interests, and develop a formal decision-making procedure that incorporates this notion of rationality.

2 Individual Rationality

Fundamental rationality requires a decision maker to choose between alternatives in a way that is consistent with its preferences. Preferences arise as a result of the decision maker's evaluations of the expected consequences of the alternatives in the light of all relevant goals. If its evaluations encompass consideration of only its own goals, then the mapping from goals to preferences is straightforward: the decision maker should order its preferences in a way such that acting consistently with them will lead to expected consequences that maximize its own benefit in the sense of achieving its own goals, without regard for the the benefit to others. This is the doctrine of individual rationality, the paradigm upon which conventional decision theory and game theory are based. While this perception may be appropriate for environments of perfect competition or perfect cooperation, the individual rationality hypothesis loses much of its power in more general settings. As expressed by Arrow, when the assumption of perfect competition fails, "the very concept of rationality becomes threatened, because perceptions of others and, in particular, of their rationality become part of one's own rationality [2]." Nevertheless, much of multi-agent decision theory rests upon the foundation of individual rationality. Individual rationality is so ingrained in the decision theoretic culture that its appropriateness is usually unchallenged. But, in the pithy words of Arrow, its use is "ritualistic, not essential [2]."

What is essential, in any formalized approach to multi-agent decision making that claims to be rational (in any sense of the word), is that it be based upon a model of the society that is ecologically balanced in that it is able to accommodate the various relationships that exist between agents and their environment, including other agents. If individual rationality provides that balance, then its use is appropriate. Otherwise, an alternative notion of rationality will be required.

Despite the inability to guarantee consistency between individual and group preferences, there are a number of approaches to group decision making that are seriously considered. Most, however, grow out of the individual rationality paradigm, and are, in effect, extensions of it. One approach is to view the group itself as a higher-level decision making entity, or superplayer, who is able to fashion a group expected utility to be maximized. This "group-Bayesian" approach (as termed by Raiffa [5]), however, fails to distinguish between the notion of group choices and group preferences.

Another approach to group decision making is to invoke the Pareto principle, which is to choose a group decision such that, if any individual were to make a unilateral change to increase its level of satisfaction, the satisfaction level for at least one other member of the group would be decreased. The obvious problem with this approach is that all participants must agree to abide by the Pareto principle. It seems that, for this principle to be viable, some notion of group preference must enter through the back door, if only to agree to be loyal Paretians, not to mention the problem of deciding which Pareto choice to select if there is not a unique one.

As Arrow's impossibility theorem [1] establishes, it is impossible to reconcile the notions of group rationality and individual rationality *and still strictly adhere to the notion of rationality as doing the best thing possible, either for the group or for the individual.* Luce and Raiffa summarize the situation succinctly: "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 [4, p. 196]." Often, the most articulate advocates of a theory are also its most insightful critics. What may help address this problem is a notion of rationality that does not depend, at root, upon the concept of being best. But if we are to abandon individual rationality as the basis for our decision making, we must provide an alternative notion that (a) overcomes the difficulties inherent with individual rationality, but (b) does not conflict with individual rationality.

The appeal of optimization and equilibration is a strongly entrenched attitude that dominates decision theory. There is great comfort in following traditional paths, especially when those paths are founded on such a rich and enduring tradition as individual rationality. But when synthesizing an artificial system, the designer has the opportunity to impose upon the agents a more socially accommodating paradigm. Rather than depending upon the non-cooperative equilibria defined by individual-benefit saddle points, this alternative may lead to the more socially realistic and valuable equilibria of shared interests and acceptable compromises.

3 The Allocation Game

Under a strict interpretation of individual optimization, an agent could not consider sacrificing even a small amount of its own satisfaction, even if doing so would greatly benefit another. If the agents are in a state of perfect competition, such as a zero-sum game, such behavior would clearly be rational, but if an opportunity for cooperation exists it is possible, even important, to dispute the reasonableness of such a stance. If we relax the strict demand for being best, however, we may open the way to accommodate others by giving some deference to their preferences as well as to our own.

Let us replace the demand for the best, and only the best, with a desideratum of being good enough. While the statement "What is best for me and what is best for you is also jointly best for us together" may be nonsense, the statement "What is good enough for me and what is good enough for you is also jointly good enough for us together" may be perfectly sensible, especially when we do not have inflexible notions of what it means to be "good enough." To be useful, we must, of course, precisely define what it means to be good enough, but, if we are able to do so, we may gain some significant benefits from such a stance.

One of the benefits of such a softer stance is to move closer to conforming to Arrow's observation that other agents' rationality is part of one's own rationality. A willingness to be content with a choice that is adequate, if not optimal, opens the way to expanding one's sphere of interest to permit the accommodation of others, thereby paving the way for the development of a notion of rationality that is able to account for the interests of others as well as for one's self interest. The key enabling concept for this to happen is the notion of conditional preferences.

In societies that value cooperation, it is unlikely that the preferences of a given individual will be formed independently of the preferences of others. Knowledge about one agent's preferences may alter another agent's preferences. Such preferences are *conditioned on the preferences of others.* Individual rationality does not accommodate such conditioning. The only type of conditioning supported by individual rationality is for each agent to express its preferences conditioned on the choices of the others but not on their *preferences* about their choices. Each agent then computes its own expected utility as a function of the possible options of all agents, juxtaposes these expected utilities into a payoff array, and searches for an equilibrium. Although the equilibrium itself is governed by the utilities of all agents, the individual expected utilities that define the equilibrium do not consider the preferences of others. We illustrate with the following example.

Example 1 Space Allocation. *Professors X_1 and X_2 are new faculty members who must be assigned offices and laboratories. One of the offices has a window (w), the other has no window ($\overline{w}$). Also, one of the laboratories has air conditioning, (a), while the other laboratory ($\overline{a}$) does not. Although X_1 is mentally and physically healthy, X_2 suffers from claustrophobia and asthma, and would greatly benefit from a win-*

dowed office and an air-conditioned lab. To facilitate the assignment of space, each faculty member provides a ranking, on a scale of 1 to 5, of their preferences for each of the four possible choices. These rankings are displayed in Table 1.

Faculty	w	$\overline{w}$	a	$\overline{a}$
X_1	4	2	3	2
X_2	5	2	4	2

Table 1: Preference Ranking for Space Allocation

These rankings reflect the individual's preferences *acting individually*—they both prefer the superior rooms; hence there is the potential for conflict. This problem seems to fit naturally into the traditional game-theoretic context, where one participant gets to choose the office, and the other gets to choose the laboratory.

Conventional game theory, however, assumes that the participants are committed to seeking their own self interest. But if X_1 has altruistic tendencies, she might be willing to suppress her own egoistic preferences in deference to X_2 and redefine her rankings in Table 1. So doing is a way to trick individual rationality into providing a response that can be interpreted as unselfish. But such an artifice provides only an indirect way to simulate socially useful attributes of cooperation, unselfishness, and altruism under a regime that is more naturally attuned to competition, exploitation, and avarice. Furthermore, if X_1 is to entertain such a move, she must be willing to forego her own preferences regardless of X_2's real motives. It is possible that X_2 is secretly under medication and psychiatric therapy, has effective control of his ailments, and prefers the superior rooms for exactly the same reasons as does X_1 (presumably, for the view and the comfort), and simply wants to manipulate X_2 by exploiting her sympathetic tendencies for his own selfish purposes. Under these circumstances, X_1's altruistic tendencies would be misplaced. Since X_2 cannot ascertain X_1's true motives, she should at least be in a position to control her response as a function of her assessment of X_2's possible motives and attitude.

Suppose, instead of X_1 totally capitulating to X_2's presumed needs, X_1 were willing to accommodate X_2 at least half-way; that is, if X_2 strongly wants w and $\overline{a}$, then she will be happy with $\overline{w}$ and a, and vice versa. But, if X_2 adopts an apparently greedy stance and strongly prefers both of the superior rooms, then she will accommodate that attitude only to the degree α, where $\alpha \in [0, 1]$ is her *index of altruism*. This index is a way for X_1 to temper her generosity and control how much deference she is willing to grant X_2 to accommodate his desires.

Before we present a solution to this game that incorporates conditional values, we first must summarize the theoretical basis for our solution.

4 Satisficing Rationality

An alternative to conventional N-player game theory is a new approach to multi-agent decision making called *satisficing games* [3, 6, 7]. Rather than defining a game in terms of a payoff array of expected utilities, as is done with conventional normal-form games, a satisficing game incorporates the same information that is used to define expected utilities to form two functions, called *selectability* and *rejectability* functions, denoted $p_{S_1 \cdots S_N}$ and $p_{R_1 \cdots R_N}$, respectively. These functions are multivariate mass functions; that is, they are non-negative and normalized (but do not possess the same semantics as is usually attributed to probability mass functions). The selectability function characterizes the degree to which an option leads to success, in the sense of achieving the goal of the decision problem, and the rejectability function characterizes the cost, or degree of resource consumption, that is associated with the option. The two functions are compared for each possible joint outcome, and the set of joint outcomes for which selectability is at least as great as rejectability form the *jointly satisficing set*.

Let $X_1, \ldots, X_N$ be a society of decision makers, and let U_i be the set of options available to X_i, $i = 1, \ldots, N$. The *joint action set* is the product set $\mathbf{U} = U_1, \times \cdots \times U_N$, and denote elements of this set as $\mathbf{u} = \{u_1, \ldots, u_N\}$, where $u_i \in U_i$.

Definition 1 A *satisficing game* is a triple $\{\mathbf{U}, p_{S_1 \cdots S_N}, p_{R_1 \cdots R_N}\}$. The *joint solution* to a satisficing game is the set

$$\mathbf{\Sigma}_q = \{\mathbf{u} \in \mathbf{U}\colon p_{S_1 \cdots S_N}(\mathbf{u}) \geq q p_{R_1 \cdots R_N}(\mathbf{u})\},$$

where q is the *index of caution*, and parameterizes the degree to which the decision maker is willing to accommodate increased costs to achieve success. An equivalent way of viewing this parameter is as an index of boldness, characterizing the degree to which the decision maker is willing to risk rejecting successful options in the interest of conserving resources. Nominally, $q = 1$, which attributes equal weight to success and resource conservation interests. $\mathbf{\Sigma}_q$ is termed the *joint satisficing* set, and elements of $\mathbf{\Sigma}_q$ are *jointly satisficing actions.* □

The jointly satisficing set provides a formal definition of what it means to be good enough for the group; namely, a joint option is good enough it the selectability is greater than or equal to the index of caution times the rejectability.

Definition 2 A decision-making group is *jointly satisficingly rational* if the members of the group choose a vector of options for which joint selectability is greater than or equal to the index of caution times joint rejectability. □

The marginal selectability and rejectability mass functions for each X_i may be obtained by summing over the options of all other participants, yielding:

$$p_{S_i}(u_i) = \sum_{\substack{u_j \in U_j \\ j \neq i}} p_{S_1 \cdots S_N}(u_1, \ldots, u_N) \quad (1)$$

$$p_{R_i}(u_i) = \sum_{\substack{u_j \in U_j \\ j \neq i}} p_{R_1 \cdots R_N}(u_1, \ldots, u_N). \quad (2)$$

Definition 3 The *individually satisficing solutions* to the satisficing game $\{\mathbf{U}, p_{S_1 \cdots S_N}, p_{R_1 \cdots R_N}\}$ are the sets

$$\Sigma_q^i = \{u_i \in U_i \colon p_{S_i}(u_i) \geq q p_{R_i}(u_i)\}. \quad (3)$$

The product of the individually satisficing sets is the *satisficing rectangle*:

$$\mathfrak{R}_q = \Sigma_q^1 \times \cdots \times \Sigma_q^N = \{(u_1, \ldots, u_N) \colon u_i \in \Sigma_q^i\}.$$

□

Definition 4 A decision maker is *individually satisficingly rational* if it chooses an option for which the selectability is greater than or equal to the index of caution times rejectability. □

The individually satisficing sets identify the options that are good enough for the individuals; namely, the options such that the individual selectability is greater than or equal to the index of caution times individual rejectability. It remains, however, to reconcile, if possible, the individual choices with the group choices. To do this, we need to establish the relationship between the jointly satisficing set and the satisficing rectangle.

Definition 5 A *compromise at* q is a jointly satisficing solution such that each element is individually satisficing. We denote this set

$$\mathbf{C}_q = \mathbf{\Sigma}_q \cap \mathfrak{R}_q.$$

A compromise at q exists if $\mathbf{C}_q \neq \emptyset$, otherwise an *impasse at* q occurs. □

The following theorem expresses the relationship between the individual and jointly satisficing sets.

Theorem 1 *If u_i is individually satisficing for X_i, that is, if $u_i \in \Sigma_q^i$, then it must be the ith element of some jointly satisficing vector $\mathbf{u} \in \mathbf{\Sigma}_q$.*

Proof This theorem is proven by establishing the contrapositive, namely, that if u_i is not the ith element of any $\mathbf{u} \in \mathbf{\Sigma}_q$, then $u_i \notin \Sigma_q^i$. Without loss of generality, let $i = 1$. By hypothesis, $p_{S_1 \cdots S_N}(u_1, \mathbf{v}) < q p_{R_1 \cdots R_N}(u_1, \mathbf{v})$ for all $\mathbf{v} \in U_2 \times \cdots \times U_N$, so $p_{S_1}(u_1) = \sum_{\mathbf{v}} p_{S_1 \cdots S_N}(u_1, \mathbf{v}) < q \sum_{\mathbf{v}} p_{R_1 \cdots R_N}(u_1, \mathbf{v}) = q p_{R_1}(u_1)$, hence $u_1 \notin \Sigma_q^1$. □

The content of this theorem is that no one is ever completely frozen out of a deal—every decision maker has a seat at the negotiating table. This is perhaps the weakest condition under which negotiations are possible. Perhaps the most simple way to negotiate is to lower one's standards in a controlled way.

Corollary 1 *There exists an index of caution value $q_0 \in [0, 1]$ such that a compromise exists at q_0.*

The proof of this corollary is trivial and is omitted. If the players are each willing to lower their standards sufficiently by decreasing the level of caution, q, they may eventually reach a compromise soltuion that is both individually and jointly satisficingly rational. The parameter q_0 is a measure of how much they must be willing to compromise to avoid an impasse. Note that willingness to lower one's standards is not total capitulation, since the participants are able to control the degree of compromise by setting a limit on how small of a value of q they can tolerate. Thus, a controlled amount of altruism is possible with this formulation. But, if any player's limit is reached without a mutual agreement being obtained, the game has reached an impasse.

5 Solutions to Allocation Game

5.1 A Game Theoretic Solution

Let us return to the Allocation game. Being game theory enthusiasts, X_1 and X_2 decide to choose offices according to classical game theory. They agree to flip a fair coin, and if it lands heads up, X_1 will choose the office and X_2 will choose the laboratory, with the opposite happening if it lands tails. Table 2 illustrates the payoff matrix for these games. The utilities for each of these games are obtained in accordance with the orderings given in Table 1.

"heads"

	X_2	
X_1	a	$\overline{a}$
w	(6,6)	(7,4)
$\overline{w}$	(4,9)	(5,7)

"tails"

	X_2	
X_1	w	$\overline{w}$
a	(5,7)	(7,4)
$\overline{a}$	(4,9)	(6,6)

Table 2: Payoff matrix for Allocation Game.

These games each have a unique Nash equilibrium. For the "heads" game it is the joint option (w, a); that is, X_1 gets the window office and X_2 gets the air-conditioned laboratory, and for the "tails" game it is (a, w); that is, X_1 gets the air-conditioned laboratory and X_2 gets the window office. In either case, however, the equilibrium solution is for one participant to have the superior office and the other to have the superior laboratory. Neither game, however, permits X_1 to take into consideration X_2's health. Under the von Neumann-Morgenstern setup, the only way that X_1 can accommodate X_2's ailment is to throw the game by redefining her utilities so that X_2 is sure to win both the superior office and the superior laboratory. If X_1 has even the slightest suspicion that X_2 is not as sick as he claims to be, she may not be willing to capitulate, and any altruistic leanings she might have may well be quashed by considerations of her own well being and general feelings of fairness.

5.2 A Satisficing Solution

To obtain a satisficing solution we need to first provide operational definitions for selectability and rejectability, and then define $p_{S_1S_2}$ and $p_{R_1R_2}$. Let us define selectability in terms of the preferences for the rooms, and define rejectability as the condition of dual occupancy (two people in the same room).

Each player has four options, yielding $U_1 = U_2 = \{wa, w\bar{a}, \bar{w}a, \bar{w}\bar{a}\}$. Let us assume that X_1 is willing to adjust her preferences to accommodate X_2 at least to some extent. In other words, X_1 will condition her selectability on X_2's preferences. This conditioning is modeled by means of a conditional selectability function (analogous to a conditional probability). The joint selectability function may be factored according to the chain rule of probability theory as

$$p_{S_1S_2}(u_1, u_2) = p_{S_1|S_2}(u_1|u_2) \cdot p_{S_2}(u_2).$$

We interpret $p_{S_1|S_2}(u_1|u_2)$ as follows. Given that X_2 places all of his selectability mass on a fixed u_2, then this function represents the amount of selectability that X_1 places on the elements $u_1 \in U_1$. There are four conditional selectabilities to compute; namely, $p_{S_1|S_2}(\cdot|wa)$, $p_{S_1|S_2}(\cdot|w\bar{a})$, $p_{S_1|S_2}(\cdot|\bar{w}a)$, and $p_{S_1|S_2}(\cdot|\bar{w}\bar{a})$. According to her principle of meeting X_2 half way, she determines to set all of her conditional selectability mass on the propositions complementary to X_2 having exactly one superior room. This yields the following conditional selectabilities:

$$\mathbf{p}_{S_1|S_2}(\cdot|w\bar{a}) = [0, 0, 1, 0] \quad (4)$$

$$\mathbf{p}_{S_1|S_2}(\cdot|\bar{w}a) = [0, 1, 0, 0]. \quad (5)$$

Now let us deal with the situation where X_2 is greedy, and strongly desires outcome wa for himself. Let $\alpha \in [0, 1]$ be X_1's index of altruism, and set $p_{S_1|S_2}(\bar{w}\bar{a}|wa) = \alpha$, which indicates a willingness to degree α to prefer the two inferior rooms in deference to X_2's preference for the two superior rooms. Furthermore, let us suppose that, if X_2 were to prefer wa, then X_1 would not wish to enter into a dual occupancy arrangement with him, and accordingly set $p_{S_1|S_2}(wa|wa) = 0$. Let us apportion the remaining $1 - \alpha$ of selectability equally among the options $w\bar{a}$ and $\bar{w}a$, yielding

$$\mathbf{p}_{S_1|S_2}(\cdot|wa) = \left[0, \frac{1-\alpha}{2}, \frac{1-\alpha}{2}, \alpha\right].$$

Finally, we complete the specification of the conditional selectabilities by noting that, were X_2 to prefer $\bar{w}\bar{a}$, then there would be no conflict, and X_1 would prefer wa, yielding

$$\mathbf{p}_{S_1|S_2}(\cdot|wa) = [1, 0, 0, 0].$$

The marginal selectability for X_2 is given by the orderings in Table 1, which we simply normalize to obtain

$$\mathbf{p}_{S_2}(\cdot) = \left[\frac{9}{26}, \frac{7}{26}, \frac{6}{26}, \frac{4}{26}\right].$$

Table 3 gives the resulting joint selectability function.

X_1	X_2: wa	$w\bar{a}$	$\bar{w}a$	$\bar{w}\bar{a}$
wa	0	0	0	$\frac{2}{13}$
$w\bar{a}$	$\frac{9(1-\alpha)}{52}$	0	$\frac{3}{13}$	0
$\bar{w}a$	$\frac{9(1-\alpha)}{52}$	$\frac{7}{26}$	0	0
$\bar{w}\bar{a}$	$\frac{9\alpha}{26}$	0	0	0

Table 3: Joint Selectability for Allocation Game.

The marginal selectability for X_1 is obtained by summing across the rows of this matrix, yielding

$$\mathbf{p}_{S_1}(\cdot) = \left[\frac{2}{13}, \frac{3}{13} + \frac{9(1-\alpha)}{52}, \frac{7}{26} + \frac{9(1-\alpha)}{52}, \frac{9\alpha}{26}\right].$$

To form the joint rejectability function, we will ascribe equal rejectability to every condition of dual occupancy, and zero rejectability to every condition of separate occupancy. The resulting joint rejectability function is expressed by the matrix in Table 4.

X_1	X_2: wa	$w\bar{a}$	$\bar{w}a$	$\bar{w}\bar{a}$
wa	$\frac{1}{12}$	$\frac{1}{12}$	$\frac{1}{12}$	0
$w\bar{a}$	$\frac{1}{12}$	$\frac{1}{12}$	0	$\frac{1}{12}$
$\bar{w}a$	$\frac{1}{12}$	0	$\frac{1}{12}$	$\frac{1}{12}$
$\bar{w}\bar{a}$	0	$\frac{1}{12}$	$\frac{1}{12}$	$\frac{1}{12}$

Table 4: Joint Rejectability for Allocation Game.

The marginal rejectability functions are obtained by summing over the rows and columns of the joint rejectability matrix, yielding

$$\mathbf{p}_{R_1}(\cdot) = \mathbf{p}_{R_2}(\cdot) = \left[\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right].$$

With the degree of caution set to unity, the jointly satisficing set is

$$\boldsymbol{\Sigma}_q = \begin{cases} \{(wa, \overline{wa}), (w\overline{a}, \overline{w}a), (\overline{w}a, w\overline{a}), \\ (\overline{wa}, wa)\} & \text{if } \alpha > \frac{14}{27} \\ \{(wa, \overline{wa}), (w\overline{a}, \overline{w}a), (\overline{w}a, w\overline{a}), \\ (\overline{wa}, wa), (w\overline{a}, wa), (\overline{w}a, wa)\} & \text{if } \alpha \leq \frac{14}{27} \end{cases}.$$

We see that all separately-occupied options are jointly satisficing, and if the degree of altruism is sufficiently low, then two dual-occupancy options actually are jointly satisficing.

The individually satisficing sets are

$$\Sigma_q^1 = \begin{cases} \{w\overline{a}, \overline{w}a\} & \text{if } \alpha < \frac{13}{18} \\ \{w\overline{a}, \overline{w}a, \overline{wa}\} & \text{if } \frac{13}{18} \leq \alpha \leq \frac{8}{9} \\ \{\overline{w}a, \overline{wa}\} & \text{if } \alpha > \frac{8}{9} \end{cases} \quad (6)$$

$$\Sigma_q^2 = \{wa, w\overline{a}\}. \quad (7)$$

The possible compromise sets are

$$\mathbf{C}_q = \begin{cases} (\overline{w}a, w\overline{a}) & \text{if } \alpha < \frac{13}{18} \\ \{(\overline{wa}, wa), (\overline{w}a, w\overline{a})\} & \text{if } \frac{13}{18} \leq \alpha \leq 1 \end{cases}.$$

Thus, if X_1's index of altruism is sufficiently high ($\alpha > \frac{13}{18}$), then X_1 is willing to give up both of the superior offices to accommodate X_2. If X_1 is not feeling quite that generous, then she will agree only to a compromise where both players get one superior room. X_1 is not required to capitulate to at least consider accommodating X_2.

This example illustrates the way in which conditioning may enter into a decision problem. It demonstrates that (i) conditioning permits altruism but does not enforce it; (ii) conditioning is compatible with satisficing game theory, but is not easily implemented in either social choice or von Neumann-Morgenstern game theory; (iii) conditioning provides an intuitively reasonable resolution to the conflict; and (iv) conditioning allows the decision maker to formulate its individual preferences in a group context.

As with any multi-agent decision problem, complexity grows with dimensionality, but is mitigated by the sparseness of the relationships. For hierarchical systems, or systems with "Markovian" like relationships such as conditional independence, the complexity will grow essentially linearly with the dimensionality. For more tightly interconnected systems, however, the complexity will grow combinatorically.

6 Discussion

Decision-theoretic researchers have long wrestled with how to deal with Raiffa's deep-seated feeling that "somehow the group entity is more than the totality of its members." Yet, researchers have steadfastly (and justifiably) refused to consider the group entity itself as a decision-making superplayer.

Satisficing game theory offers a way to account for the group entity without the fabrication of a superplayer. This accounting is done through the conditional preference relationships that are expressed through the selectability and rejectability functions due to their mathematical structure as a probability (but not with the usual semantics dealing with randomness or uncertainty). Just as the a joint probability function is more than the totality of the marginals, the selectability and rejectability functions are more than the totality of the individual selectability and rejectability functions. It is only in the case of stochastic independence that a joint distribution can be constructed from the marginal distributions, and it is only in the case of complete interindependence that group welfare can be expressed in terms of the welfare of individuals.

Optimization and equilibration are notions of local interest and global extent, in that they involve only individual interests and require *inter-option* comparisons of all options. Satisficing, on the other hand, is a notion of global interest and local extent, in that it accounts for the interdependencies between all participants and involves only *intra-option* comparisons of each option. Satisficing does not, however, demand the abdication of the individual to the group; rather, it provides an avenue for compromise between individual and group interests.

References

[1] K. J. Arrow. *Social Choice and Individual Values*. John Wiley, New York, 1951. 2nd ed. 1963.

[2] K. J. Arrow. Rationality of self and others. In R. M. Hogarth and M. W. Reder, editors, *Rational Choice*. Univ. of Chicago Press, Chicago, 1986.

[3] M. A. Goodrich, W. C. Stirling, and E. R. Boer. Satisficing revisited. *Minds and Machines*, 10:79–109, February 2000.

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

[5] H. Raiffa. *Decision Analysis*. Addison-Wesley, Reading, MA, 1968.

[6] W. C. Stirling and M. A. Goodrich. Satisficing equilibria: A non-classical approach to games and decisions. In S. Parsons and M. J. Wooldridge, editors, *Workshop on Decision Theoretic and Game Theoretic Agents*, pages 56–70, University College, London, United Kingdom, 5 July, 1999.

[7] W. C. Stirling and M. A. Goodrich. Satisficing games. *Information Sciences*, 114:255–280, March 1999.