Beyond optimization: overcoming the limitations of individual rationality

Wynn C. Stirling
Electrical and Computer Engineering Department
Brigham Young University
Provo, Utah 84602, USA
email: wynn@ee.byu.edu

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.

— *R. D. Luce and H. Raiffa*
Games and Decisions (1957)

Abstract

Von Neumann-Morgenstern game theory is the multi-agent instantiation of individual rationality, and is the standard for decision-making in group settings. Individual rationality, however, requires each player to optimize its own performance, regardless of the effect so doing has on the other players. This feature limits the ability of game theory as a design paradigm for group behavior where coordination is required, since it cannot simultaneously accommodate both group and individual preferences. By replacing the demand for doing the best thing possible for the individual with a mathematically precise notion of being "good enough," satisficing game theory allows both group and individual interests to be simultaneously accommodated.

Introduction

It is a platitude that a decision-maker should make the best choice possible. Typically, this injunction is taken to mean that a decision-maker should optimize, that is, maximize expected utility. Although the exigencies of decision-making under time and computational constraints may require the decision-maker to compromise, resulting in various notions of bounded optimization, the fundamental commitment to optimality usually remains intact. It is almost mandatory that a decision methodology incorporate some instance of optimization, even if only approximately. Otherwise the decision-making procedure is likely to be dismissed as *ad hoc*.

Optimization is founded on the principle that individual interests are fundamental and that social welfare is a function of individual welfare (Bergson, 1938; Samuelson, 1948). This hypothesis leads to the doctrine of *rational choice*, which is that "each of the individual decision-makers behaves as if he or she were solving a constrained maximization problem" (Hogarth and Reder, 1986). This paradigm is the basis of much of the conventional decision theory that is used in economics, the social and behavioral sciences, engineering, and computer science. It relies upon two fundamental premises:

P-1 *Total ordering:* the decision-maker is in possession of a total preference ordering (that is, an ordering that is reflexive, antisymmetric, transitive, and linear) for all of its possible choices under all conditions (in multi-agent settings, this includes knowledge of the total orderings of all other participants).

P-2 *The principle of individual rationality:* a decision-maker should make the best possible decision for itself, that is, it should optimize with respect to its own total preference ordering (in multi-agent settings, this ordering may be influenced by the choices available to the other participants).

Self-interested human behavior is often considered to be an appropriate metaphor in the design of protocols for artificial decision-making systems. With such protocols, it is often taken for granted that each member of a community of decision-makers should try

> ... to maximize its own good without concern for the global good. Such self-interest naturally prevails in negotiations among independent businesses or individuals ... Therefore, the protocols must be designed using a *noncooperative, strategic* perspective: the main question is what social outcomes follow given a protocol which *guarantees that each agent's desired local strategy is best for that agent—and thus the agent will use it.* (Sandholm, 1999, p. 201, 202; emphasis in original).

When artificial decision-makers are designed to function in a non-adversative environment it is not obvious that it is either natural or necessary to restrict attention to noncooperative protocols. decision-makers who are focused on their own self-interest will be driven to compete with decision-makers whose interests might possibly compromise their own. Certainly, conflict cannot be avoided in general, but conflict can just as easily lead to collaboration as to competition.

One of the justifications for adopting self-interest as a paradigm for artificial decision-making systems is that it is a simple and convenient principle upon which to build a mathematically based theory. Self-interest is the Occam's razor of interpersonal interaction and relies only upon the minimal assumption that an individual will put its own interests above everything and everyone else. This simple principle allows the decision-maker to abstract the problem from its context and express it in unambiguous mathematical language. With this language, utilities can be defined and calculus can be employed to facilitate the search for the optimal choice. The quintessential manifestation of this approach to decision-making is von Neumann-Morgenstern game theory (von Neumann and Morgenstern, 1944).

Game theory is built on one basic principle: self-interest—each player must maximize its own expected utility under the constraint that other players will do likewise. Such players will seek an equilibrium; that is, a state such that no individual player can improve its level of satisfaction by making a unilateral change in its strategy. For two-person constant-sum games, this is perhaps the only reasonable, non-vacuous principle—what one player wins, the other loses. Game theory insists, however, that this same principle applies to the general case. Thus, even in situations where there is the opportunity for group, as well as individual interest, only individually rational actions are viable. If a joint (that is, for the group) solution is not individually rational for some decision-maker, that self-interested decision-maker would not be a party to such a joint action.

Coordinated behavior is perhaps the most important (and most difficult) social attribute to synthesize in an artificial decision-making group. Achieve such a design objective, however, will be greatly facilitated if decision-making is based on rationality principles that permit the decision-makers to expand their spheres of interest beyond themselves and give deference to others. As Arrow observed, when the assumption of perfect competition does not apply, "the very concept of [individual] rationality becomes threatened, because perceptions of others and, in particular, of their rationality become part of one's own rationality" (Arrow, 1986). Arrow has put his finger on a critical limitation of individual rationality. As is well known, however, a notion of "group rationality" that requires the group to do the best for itself is not compatible with individual rationality (Luce and Raiffa, 1957). Nevertheless, game theory is often used to characterize situations where coordinated behavior, where the members of a group coordinate their actions to accomplish tasks that pursue the goals of both the group and its members, is of fundamental importance.

In this paper we first review the various notions of group preference that have arisen in the context of game theory, we then present an alternative concept of utility theory and show develop a new class of games, called *satisficing games*. We then show how to formulate these games through the use of conditional preferences and describe how to reconcile individual and group preferences.

Group Preferences

Several attempts have been made to express group preferences in a game-theoretic context. Shubik offers two interpretations of this notion, neither of which game theorists view as entirely satisfactory: "Group preferences may be regarded either as derived from individual preferences by some process of aggregation or as a direct attribute of the group itself" (Shubik, 1982, p. 109). One way to aggregate a group preference from individual preferences is to define a "social-welfare" function that provides a total ordering of the group's strategy profiles. The fundamental issue is whether or not, given arbitrary preference orderings for each individual in a group, there always exists a way of combining these individual preference orderings to generate a consistent preference ordering for the group. In an landmark result, Arrow (Arrow, 1951; Sen, 1979) showed that no social-welfare function exists that satisfies a set of reasonable and desirable properties, each of which is consistent with the notion of self-interested rationality and the retention of individual autonomy.

The Pareto principle provides a concept of social welfare as a direct attribute of the group. A strategy profile is *Pareto optimal* if no single decision-maker, by changing its decision, can increase its level of satisfaction without lowering the satisfaction level of at least one other decision-maker. However, if a Pareto-optimal solution does not provide each each player at least its security level (i.e., the minimum payoff it can be guaranteed, even if all other players conspire against it), it could not be a party to that decision and still be faithful to individual rationality.

To impose a strategy profile, such as a Pareto-optimal solution, on a group would require the existence of a *superplayer*, or, as Raiffa puts it, the "organization incarnate" (Raiffa, 1968), who functions as a higher-level decision-maker. Shubik refers to the practice of ascribing preferences to a group as a subtle "anthropomorphic trap" of making a shaky analogy between individual and group psychology. He argues that, "It may be meaningful... to say that a group 'chooses' or 'decides' something. It is rather less likely to be meaningful to say that the group 'wants' or 'prefers' something" (Shubik, 1982, p. 124). Raiffa, also, rejects the notion of a superplayer, but confesses that he still feels "a bit uncomfortable... somehow the group entity is more than the totality of its members" (Raiffa, 1968, p. 237). Arrow expresses a similar discomfort: "All the writers from Bergson on agree on avoiding the notion of a social good not defined in terms of the values of individuals. But where Bergson seeks to locate social values in welfare judgments by individuals, I prefer to locate them in the actions taken by society through its rules for making social decisions" (Arrow, 1951, p. 106). Evidently, although a satisfactory account of group preferences may be difficult or, perhaps, impossible, to obtain under individual rationality, the desire to accommodate the notion remains.

Perhaps the source of discomfort is that individual rationality by itself does not provide the ecological balance that a group must achieve if it is to accommodate the variety of relationships that can exist between decision-makers and their environment. But achieving such a balance should not require the aggregation of individual interests or the fabrication of a superplayer. While such approaches may be recommended as ways to account for group interests, they may also manifest the limitations of individual rationality.

Of course, one may substitute the interests of others for one's own self-interest, as Sen (1990, p. 19) observed: "It is possible to define a person's interests in such a way that no matter what he does he can be seen to be furthering his own interests in every isolated act of choice ... no matter whether you are a single-minded egoist or a raving altruist or a class-conscious militant, you will appear to be maximizing your own utility in this enchanted world of definitions." Although it is certainly possible to suppress one's preferences in deference to others by redefining one's own utility, doing so is little more than a device for co-opting individual rationality into a form that can be interpreted as unselfish. Such a device only simulates attributes of cooperation, unselfishness, and altruism while maintaining a regime that is competitive, exploitive, and avaricious.

Nevertheless, game theory has been a great success story for economics, political science, and psychology. With these disciplines, however, game theory is used primarily as an *analysis tool* to explain and predict behavior, and there is no *causal* relationship between the performance of the entities being studied and the model used to characterize them. In the engineering context of *synthesis,* however, the goal is to design and build artificial decision-making entities and the models used to characterize behavior are indeed causal. Although von Neumann-Morgenstern game theory has been successfully applied in many disciplines, this success does not imply that self-interest is the only principle that will lead to credible models of behavior, it does not imply the impossibility of accommodating both group and individual interests in some meaningful way, and it does not imply that individual rationality is an appropriate principle upon which to base a theory for the design and synthesis of artificial decision-making entities.

There is an old saying: "If all I have is a hammer, everything looks like a nail." We may paraphrase that sentiment as follows: "If all I know how to do is optimize, every group decision problem looks like a von Neumann-Morgenstern game." If, however, as Luce and Raiffa conjecture, it is indeed too much to ask that a sociology be derived from the single assumption of individual rationality, we may gain some advantage in social situations by considering the use of decision-making tools that are not founded on that single assumption and hence may be better suited for the expression of a sociology. Consider, for example, the following group decision scenario.

Example 1 The Pot-Luck Dinner *Larry, Curly, and Moe are going to have a pot-luck dinner. Larry will bring either soup or salad, Curly will provide the main course, either beef, chicken, or pork, and Moe will furnish the dessert, either lemon custard pie or banana cream pie. The choices are to be made simultaneously and individually following a discussion of their preferences, which discussion yields the following results:*

1. *In terms of meal enjoyment, if Larry were to prefer soup, then Curly would prefer beef to chicken by a factor of two, and would also prefer chicken to pork by the same ratio. However, if Larry were to prefer salad, then Curly would be indifferent regarding the main course.*
2. *If Curly were to reject pork as being too expensive, then Moe would strongly prefer (in terms of meal enjoyment) lemon custard pie and Larry would be indifferent regarding soup or salad. If, however, Curly were to to reject beef as too expensive, then Larry would strongly prefer soup and Moe would be indifferent regarding dessert. Finally, if Curly were to reject chicken as too expensive, then both Larry and Moe would be indifferent with respect to their enjoyment preferences.*

Larry, Curly, and Moe all wish to conserve cost but consider both cost and enjoyment to be equally important. Table 1 indicates the total cost (in stooge dollars) of each of the 12 possible meal combinations (using obvious abbreviations).

	lcst	bcrm
beef (soup/sald)	23/25	27/29
chkn (soup/sald)	22/24	26/28
pork (soup/sald)	20/22	24/26

Table 1: Meal cost structure for the Pot-Luck Dinner.

The decision problem facing the three participants is for each to decide independently what to bring to the meal. Obviously, each participant wants his own preferences honored, but no explicit notion of group preference is provided in the scenario. A distinctive feature of the preference specification for this example is that individual preferences are not even specified by the participants. Rather, the participants express their preferences as functions of other participants' preferences. Thus, they are not confining their interests solely to their own desires, but are taking into consideration the consequences that their possible actions have on others. Such preferences are *conditional.* These interconnections between participants may imply some sort of group preference, but it is not clear what that might be. In fact, if the preferences, either conditional or unconditional (i.e., individual) turn out to be inconsistent, then there may be no harmonious group preference, and the group may be dysfunctional in the sense that meaningful cooperation is not possible. But if they are consistent, then some form of harmonious group preference may emerge from the conditional preferences (and any unconditional preferences, should they be provided). If this is the case, then an important question is how we might elicit a group decision that accommodates an emergent group preference.

To formulate a von Neumann-Morgenstern game-theoretic solution to this decision problem, each participant must identify and quantify payoffs for every conceivable meal configuration that conform to their own preferences as well as give due deference to others. Notice that the unconditional preferences of each of the participants are not specified nor are all of the possible conditional preferences specified. Unfortunately, individual rationality makes it difficult to obviate such requirements, and the lack of a total ordering in the problem statement therefore presents a serious problem for conventional game theory. Without this constraint it is impossible to apply standard solution concepts such as defining equilibria. The desire to apply game theory may motivate decision-makers to manufacture orderings that that are not warranted. Traditional game theory is not an appropriate framework for this problem.

To solve this problem in a way that fully respects the problem statement, we need a solution concept that does not depend upon total orderings. It must, however, accommodate the fact that, even though agents may be primarily concerned with conditional local issues, these concerns can have widespread effects.

Utilities

The chain of logic that supports game theory is as follows: individual rationality leads to optimization, which requires a total ordering of preferences, which in turn motivates the defining of utility functions to characterize these preference orderings. However, as Raiffa observed: "One can argue very persuasively that the weakest link in any chain of argument should not come at the beginning" (Raiffa, 1968, p. 130). Thus, if we are to overcome the limitations imposed by individual rationality, we must forge a new chain. To do so, we must: (1) define a new notion of rationality that accommodates a wider sphere of interest; (2) replace optimization with a criterion that is compatible with non-localized rationality; (3) define preference orderings that accommodate both individual and group interests, and (4) define utility functions that are compatible with these preference orderings.

This paper presents such a chain. To forge it, however, it is more convenient to start at the end and work back to the beginning. Thus, we start by examining the structure of utility functions, which leads to an alternative preference ordering that, in turn, leads to a criterion for decision-making which, finally, defines a new concept of rationality.

Extrinsic Utilities

We concentrate exclusively on finite strategic games of complete information. Such games are defined by a payoff array, the entries of which are the N-tuples of payoffs to the players. Each player defines its payoff as a function of the strategies of all players; that is, the payoff to player i is $\pi_i(s_1, \dots, s_N)$ with $s_j \in U_j$, $j = 1, \dots, N$, where U_j is player j's strategy space. In more compact notation, we write $\pi_i(\mathbf{s})$ where $\mathbf{s} = \{s_1, \dots, s_N\} \in \mathbf{U} = U_1 \times \cdots \times U_N$ is a strategy profile. Individually rational players use such utilities to make comparisons between strategy profiles and thereby form solution concepts such as Nash equilibria to define acceptable strategies. Such comparisons are *inter-strategy* in that they require the comparisons of the attributes of each strategy to the attributes of all other strategies. Utilities that are used for this type of comparisons are *extrinsic*.

The important thing to note about the way these utilities are used is that *it is not until the payoffs are juxtaposed into an array so that the payoffs for all players can be compared that the actual "game" aspects of the situation emerges.* It is the juxtaposition that reveals possibilities for conflict or coordination. These possibilities are not explicitly reflected in the individual payoffs by themselves. In other words, although the individual's payoff is a function of other players' strategies, *it is not a function of other players' preferences.* This structure is completely consistent with exclusive self-interest, where all a player cares about is its personal benefit as a function of its own and other players' strategies, without any regard for the benefit to the others. Under this paradigm, the primary way the preferences of others factor into an individual's decision-making is to constrain behavior so as to limit the amount of damage they can do to oneself.

If the game is not one of pure competition, there may be some benefit to coordinated behavior, whereby players take into consideration the effect of their actions on the welfare of others. One way to account for the interests of others within the von Neumann-Morgenstern framework is to transform the game by introducing new payoffs of the following form (Taylor, 1987): $\pi_i^* = \sum_{j=1}^{N} \alpha_{ij}\pi_j$. By choosing the weights α_{ij}, an altruistic player i may give deference to the preferences of others. This approach, however, requires the imposition of two very strong assumptions: (a) each player must precisely know the other players' numerical payoffs (ordinal rankings are not sufficient), and (b) interpersonal comparisons of utility are implied (the addends must be commensurable). Even if these assumptions apply, choosing the weights α_{ij} requires each player to categorically ascribe a portion of is utility to other players and therefore to defer to some extent to those players in all circumstances. It does not permit a player to choose selectively which of the other player's preferences it will favor or disfavor. To do so would require α_{ij} to be functions of other players' strategies, and the formulation of the game may quickly become intractable.

Perhaps the critical question is not whether it is theoretically possible somehow to account for the interests of others via extrinsic utilities. Rather, the more important question might be: do they offer an adequate platform by which to make rationality of others part of one's own rationality? The lack of a definitive notion of group preference that is consistent with individual rationality would seem to cast doubt on an affirmative answer to the latter question.

Intrinsic Utilities

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.* In contrast to conditioning only on the actions of other participants, conditioning on the preferences of others permits a decision-maker to adjust its preferences to accommodate the preferences of others. It can bestow either deference or disfavor to others according to their preferences as they relate to its own preferences. Since traditional utility theory is a function of participant strategies, rather than participant preferences, it cannot be used to express such relationships.

To address this problem, let us closely examine the way preferences are formed. When defining preferences one often encounters valuations that are in opposition. Any given strategy profile will possess attributes that are beneficial and attributes that are detrimental to each player; such differences in valuation create dichotomies. By separating the favorable and unfavorable attributes of each strategy profile, we may expose the fundamental preference structure. Comparison of the attributes of each profile provides a determination of the benefit that obtains by adopting it relative to the cost. Such dichotomies are ubiquitous. People routinely compare the upside against the downside, the pros versus the cons, the pluses versus the minuses, and they can do this profile by profile *without directly comparing of one profile to another*. In other words, they perform *intra-profile* comparisons. Such comparisons are fundamental, and must be made, even if implicitly, in order to define the utility function which can then be used for inter-profile comparisons (i.e., total preference orderings).

Perhaps, if we start at the headwaters of preference formulation, rather than somewhere downstream, we may be able to characterize these dichotomous relationships more comprehensively and systematically. We thus consider the formation of two utility functions that accommodate dichotomies. A first consideration is that, since the two utility functions are to be compared, they must be expressed in the same units. To avoid arbitrary scalings as well as for reasons that will soon become apparent, it is convenient to define these utilities as mass functions. A function p is a *mass function* if $p(\mathbf{s}) \geq 0$ and $\sum_{\mathbf{s} \in \mathbf{U}}(\mathbf{s}) = 1$. Adopting this convention means that the player has a unit of mass to apportion among the profiles to weight their desirable attributes as well as a unit of mass to apportion to weight their undesirable attributes. These weighting functions thus possess the mathematical properties of probability mass functions, but they *do not possess the same semantics* and do not admit interpretations of belief, propensity, frequency, or any other of the usual probabilistic interpretations. To emphasize the distinction between the mathematical structure and the interpretation of these functions, we will refer to the mass function that characterizes the desirable attributes of the strategy profiles as *selectability*, and we will denote the mass function that characterizes the undesirable attributes as *rejectability*.

When defining the dichotomous utility functions, operational definitions of what is selectable and rejectable about the strategy profiles must be provided. Typically, the attributes of a profile that contribute to a fundamental goal would be associated with selectability and those attributes that inhibit or limit activity would be associated with rejectability. There generally will not be a unique way to frame a given decision problem, but regardless of the way the framing is done, it is essential that the selectability and rejectability attributes not be restatements of the same thing. In general, at least for single-agent decision problems, the selectability of a strategy should be specifiable without taking into consideration its rejectability and vice versa. For multiple-player problems, however, this independence need not apply between players; that is, one player's selectability or rejectability may influence another player's selectability or rejectability.

We are now in a position to complete our chain that links utilities to preference orderings to decision rules to rational behavior. Our procedure involves making intra-profile, rather than inter-profile, comparisons. Utilities used for this purpose are *intrinsic*, meaning that they are used to evaluate a profiles with respect to attributes that it possesses within itself, independently how that profile relates to other profiles. With extrinsic comparisons, the logical notion of rational behavior is to rank-order the strategy profiles and choose one that is optimal (i.e., to equilibrate). With intrinsic comparisons, the logical notion of rational behavior is to choose a profile that is good enough, in the sense that the gains obtained by choosing it outweigh the losses. This defines a new notion of rationality, which we term *satisficing rationality*. This notion is considerably weaker than individual rationality, which asserts that a decision-maker must make the best choice possible. Put in the vernacular, the essence of individual rationality is "only the best will do," while the essence of satisficing rationality is "at least get your money's worth."

Why would a decision-maker choose to do anything other than optimize? For a decision-maker functioning in isolation, there is no incentive, *ceteris paribus*, to eschew an optimal solution. Furthermore, with games of pure competition, there is no incentive to adopt any solution concept that does not maximize the advantage to the player. But with games of mixed motive, the notion of being individually optimal loses much of its force (as Arrow observed). Yet, under the strict paradigm of individual rationality, a player must not modify its choice to is own disadvantage, no matter how slight (unless it also redefines its utility), even if doing so would offer a great advantage to others. However, once such a player starts down the slippery slope of compromise by abandoning individual rationality, there is seemingly no way to control the slide.

Satisficing provides a way to add some friction to the slippery slope. While it does indeed abandon individual rationality, it is not heuristic. Intrinsic utilities are based on exactly the same principles of value that are used to define extrinsic utilities, the valuations are merely applied in a different way. Thus, the satisficing approach is applicable in situations where relationships other than pure competition are relevant.

The justification for using the term "satisficing" is that it is consistent with Simon's original usage of the term—to identify strategies that are good enough by comparing attributes of the strategies to a standard. This usage differs only in the standard used for comparison. Simon's standard is extrinsic; strategies are compared to the "aspiration level" of how good a solution might reasonably be achieved (Simon, 1955; Simon, 1956; Simon, 1990). Satisficing as defined herein, on the other hand, is intrinsic; the comparison is done with respect to the merits of the strategy.

Satisficing Decision-Making

To generate a useful theory of decision-making we must be able to define, in precise mathematical terms, what it means to be good enough, and we must develop a theory of decision-making that is compatible with this notion. An alternative to von Neumann-Morgenstern N-player game theory is a new approach to multiple-agent decision-making called *satisficing games* (Stirling and Goodrich, 1999b; Stirling and Goodrich, 1999a; Goodrich et al., 2000; Stirling, 2002). Two key features of our development are (a) the separation of positive and negative attributes of strategy profiles into selectability and rejectability utility functions and (b) the structure of these utility functions as mass functions. To construct these mass functions, however, we must first define an even more fundamental quantity, which we term the *interdependence mass function*, which accounts for linkages that exist between selectability and rejectability.

An act by any individual player has possible ramifications for the entire group. Some players may be benefited by the act, some may be damaged, and some may be indifferent. Furthermore, although an individual may perform the act in its own benefit or for the benefit of others, the act is usually not implemented free of cost. Resources are expended, or risk is taken, or some other penalty or unpleasant consequence is incurred, perhaps by the individual whose act it is, perhaps by other players, and perhaps by the entire group. Although these undesirable consequences may be defined independently from the benefits (recall the example of choosing an automobile), the measures associated with benefits and costs cannot be specified independently of each other, due to the possibility of interaction (e.g., cost preferences for one player may depend upon style preferences of another player). A critical aspect of modeling the behavior of a group, therefore, is the means of representing the interdependence of both positive and negative consequences of all possible strategy profiles.

Let $X_1, \ldots, X_N$ be a group of decision-makers, and let U_i be the set of strategies available to X_i, $i = 1, \ldots, N$. The *strategy profile set* is the product set $\mathbf{U} = U_1, \times \cdots \times U_N$. Let us denote elements of this set as $\mathbf{s} = \{s_1, \ldots, s_N\}$, where $s_i \in U_i$.

Definition 1 An *interdependence function* for a group $\{X_1, \ldots, X_N\}$, denoted $p_{S_1 \cdots S_N R_1 \cdots R_N} : \mathbf{U} \times \mathbf{U} \to [0, 1]$, is a mass function, that is, it is non-negative and normalized to unity, which encodes all of the positive and negative interrelationships between the members of the group. We will denote this as $p_{S_1 \cdots S_N R_1 \cdots R_N}(\mathbf{s}; \mathbf{r})$, where $\mathbf{s} = (s_1, \ldots, s_N) \in \mathbf{U}$ represents strategy profiles viewed in terms of selectability and $\mathbf{r} = (r_1, \ldots r_N) \in \mathbf{U}$ represents strategy profiles viewed in terms of rejectability. □

The interdependence function provides a complete description of all individual and group relationships in terms of their positive and negative consequences. Let $\mathbf{s}$ and $\mathbf{r}$ be two strategy profiles. $p_{S_1 \cdots S_N R_1 \cdots R_N}(\mathbf{s}; \mathbf{r})$ characterizes the simultaneous disposition of the players with respect to selecting $\mathbf{s}$ and rejecting $\mathbf{r}$. Particularly when $\mathbf{s} = \mathbf{r}$, it may appear contradictory to consider simultaneously rejecting and selecting strategies. It is important to remember, however, that considerations of selection and rejection involve two different criteria. It is no contradiction to consider selecting, in the interest of achieving a goal, a strategy that one would wish to reject for unrelated reasons, nor is it a contradiction to consider rejecting, because of some undesirable consequences, a strategy one would otherwise wish to select. Evaluating such trade-offs is an essential part of decision-making, and the interdependence mass function provides a means of quantifying all issues relevant to this trade-off.

Since it is a mass function, the interdependence function is mathematically similar to a probability mass function, but does not characterize uncertainty or randomness. Nevertheless, it possesses the mathematical structure necessary to characterize notions such as independence and conditioning.

Conditioning

The interdependence function can be rather complex but, fortunately, its structure as a mass function permits its decomposition into constituent parts according to the law of compound probability, or chain rule (Eisen, 1969).[1] Applying the formalism (but not the usual probabilistic semantics) of the chain rule, we may express the interdependence function as a product of conditional selectability and rejectability functions. To illustrate, consider a two-agent satisficing game involving decision-makers X_1 and X_2, with strategy sets U_1 and U_2, respectively. The interdependence function may be factored in several ways, for example, we may write

$$p_{S_1 S_2 R_1 R_2}(s_1, s_2; r_1, r_2) = p_{S_1|S_2 R_1 R_2}(s_1|s_2; r_1, r_2) \cdot p_{S_2|R_1 R_2}(s_2|r_1, r_2) \cdot p_{R_1|R_2}(r_1|r_2) \cdot p_{R_2}(r_2). \quad (1)$$

We interpret this expression as follows. $p_{S_1|S_2 R_1 R_2}(s_1|s_2, r_1, r_2)$ is X_1's conditional selectability of s_1 given that X_2 selects s_2, X_1 rejects r_1, and X_2 reject r_2. Similarly, $p_{S_2|R_1 R_2}(s_2|r_1, r_2)$ is X_2's conditional selectability of s_2, given that X_1 and X_2 reject r_1 and r_2, respectively. Continuing, $p_{R_1|R_2}(r_1|r_2)$ is X_1's conditional rejectability of r_1, given that X_2 rejects r_2. Finally, $p_{R_2}(r_2)$ is X_2's unconditional rejectability of r_2.

Many such factorizations are possible, but the appropriate factorization must be determined by the context of the problem. These conditional mass functions are mathematical instantiations of production rules. For example, we may interpret $p_{R_1|R_2}(r_1|r_2)$ as the rule: If X_2 rejects r_2, then X_1 feels $p_{R_1|R_2}(r_1|r_2)$ strong about rejecting r_1. In this sense, they express *local* behavior, and such behavior is often much easier to express than global behavior. Furthermore, this structure permits irrelevant interrelationships to be eliminated. Typically, there will be some close relationships between some subgroups agents, while other subgroups agents will function essentially independently of each other. For example, suppose that X_1's selectability has nothing to do with X_2's rejectability. Then we may simplify $p_{S_1|S_2 R_1 R_2}(s_1|s_2; r_1, r_2)$ to become $p_{S_1|S_2 R_1}(s_1|s_2; r_2)$.

Such conditioning permits the expression of the interdependence function as a natural consequence of the relevant interdependencies that exist between the participants. Conditioning is the key to the accommodation of the interests

[1] In the probability context, let X and Y be random variables and let x and y possible values for X and Y, respectively. By the law of compound probability, $p_{XY}(x, y) = p_{X|Y}(x|y) p_Y(y)$ expresses the joint probability of the occurrence of the event $(X = x, Y = y)$ as the conditional probability of the event $X = x$ occurring conditioned on the event $Y = y$ occurring, times the probability of $Y = y$ occurring. This relationship may be extended to the general multivariate case by repeated applications, resulting in what is called the *chain rule*.

of others. For example, if X_2 were very desirous of implementing s if X_1 were not to implement r, X_1 could accommodate X_2's preference by setting $p_{R_1|S_2}(r|s)$ to a high value (close to unity). Then, r would be highly rejectable to X_1 if s were highly selectable to X_2. Note, however, that if X_2 should turn out *not* to highly prefer s and so sets $p_{S_2}(s) \approx 0$, then the joint selectability/rejectability of $(s; r)$, namely, $p_{S_2 R_1}(s; r) = p_{R_1|S_2}(r|s) p_{S_2}(s) \approx 0$, so the joint event of X_1 rejecting r and X_2 selecting s has negligible interdependence mass. Thus, X_1 is not penalized for being willing to accommodate X_2 when X_2 does not need or expect that accommodation. By controlling the conditioning values, X_1 is able to achieve a balance between its egoistic interests and its concern for others.

Satisficing Games

From the interdependence function we may derive two functions, called *joint selectability* and *joint rejectability* functions, denoted $p_{S_1 \cdots S_N}$ and $p_{R_1 \cdots R_N}$, respectively, according to the formulas

$$p_{S_1 \cdots S_N}(\mathbf{s}) = \sum_{\mathbf{v} \in \mathbf{U}} p_{S_1 \cdots S_N R_1 \cdots R_N}(\mathbf{s}; \mathbf{v}) \tag{2}$$

$$p_{R_1 \cdots R_N}(\mathbf{s}) = \sum_{\mathbf{v} \in \mathbf{U}} p_{S_1 \cdots S_N R_1 \cdots R_N}(\mathbf{v}; \mathbf{s}) \tag{3}$$

for all $\mathbf{s} \in \mathbf{U}$. These functions are also multivariate mass functions. The two functions are compared for each possible joint outcome, and the set of joint outcomes for which joint selectability is at least as great as joint rejectability form a jointly satisficing strategy profile set.

Definition 2 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

$$\mathbf{\Sigma}_q = \{\mathbf{s} \in \mathbf{U} \colon p_{S_1 \cdots S_N}(\mathbf{s}) \geq q p_{R_1 \cdots R_N}(\mathbf{s})\}, \tag{4}$$

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. Nominally, $q = 1$, which attributes equal weight to success and resource conservation interests. $\mathbf{\Sigma}_q$ is termed the *jointly satisficing set*, and elements of $\mathbf{\Sigma}_q$ are *satisficing strategy profiles*. □

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

Definition 3 A decision-making group is *jointly satisficingly rational* if the members of the group choose a strategy profile for which joint selectability is greater than or equal to the index of caution times joint rejectability. □

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

$$p_{S_i}(s_i) = \sum_{\substack{s_j \in U_j \\ j \neq i}} p_{S_1 \cdots S_N}(s_1, \ldots, s_N) \tag{5}$$

$$p_{R_i}(s_i) = \sum_{\substack{s_j \in U_j \\ j \neq i}} p_{R_1 \cdots R_N}(s_1, \ldots, s_N). \tag{6}$$

Definition 4 The *individual satisficing solutions* to the satisficing game $\{\mathbf{U}, p_{S_1 \cdots S_N}, p_{R_1 \cdots R_N}\}$ are the sets

$$\Sigma_q^i = \{s_i \in U_i \colon p_{S_i}(s_i) \geq q p_{R_i}(s_i)\}. \tag{7}$$

The product of the individual satisficing sets is the *satisficing strategy profile rectangle*:

$$\mathfrak{R}_q = \Sigma_q^1 \times \cdots \times \Sigma_q^N = \{(s_1, \ldots, s_N) \colon s_i \in \Sigma_q^i\}. \tag{8}$$

□

Definition 5 A decision-maker is *individually satisficingly rational* if it chooses a strategy for which the marginal selectability is greater than or equal to the index of caution times the marginal rejectability. □

It is not generally true that the satisficing rectangle will have any close relationship with the jointly satisficing set. What is true, however, is the following theorem:

Theorem 1 *(The negotiation theorem.) If s_i is individually satisficing for X_i, that is, $s_i \in \Sigma_q^i$, then it must be the ith element of some jointly satisficing vector $\mathbf{s} \in \mathbf{\Sigma}_q$.*

Proof We will establish the contrapositive, namely, that if s_i is not the ith element of any $\mathbf{s} \in \mathbf{\Sigma}_q$, then $s_i \notin \Sigma_q^i$. Without loss of generality, let $i = 1$. By hypothesis, $p_{S_1 \cdots S_N}(s_1, \mathbf{v}) < q p_{R_1 \cdots R_N}(s_1, \mathbf{v})$ for all $\mathbf{v} \in U_2 \times \cdots \times U_N$, so $p_{S_1}(s_1) = \sum_{\mathbf{v}} p_{S_1 \cdots S_N}(s_1, \mathbf{v}) < q \sum_{\mathbf{v}} p_{R_1 \cdots R_N}(s_1, \mathbf{v}) = q p_{R_1}(s_1)$, hence $s_1 \notin \Sigma_q^1$. □

Thus, if a strategy is individually satisficing, it is part of a satisficing strategy profile, although it need not be part of all satisficing profiles. The converse, however, is not true: if s_i is the ith element of a jointly satisficing vector, it is not necessarily individually satisficing for X_i. The content of the negotiation theorem is that no one is ever completely frozen out of a deal—every decision-maker has, from its own perspective, a seat at the negotiating table. This is perhaps the weakest condition under which negotiations are possible.

A decision-maker who possessed a modest degree of altruism would be willing to undergo some degree of self-sacrifice in the interest of others. Such a decision-maker would be willing to lower its standards, at least somewhat and in a controlled way, if doing so would be of great benefit to others or to the group in general. The natural way for X_i to express a lowering of its standards is to decrease its index of caution. Nominally, we may set $q = 1$ to reflect equal weighting of the desire for success and the desire to conserve resources. By decreasing q, we lower the standard for success relative to resource consumption and thereby increase the size of the satisficing set. As $q \to 0$ the standard is lowered to nothing, and eventually every strategy is satisficing for X_i. Consequently, if all decision-makers are willing to reduce their standards sufficiently, a compromise can be achieved.

Reconciling Group and Individual Preferences

The satisficing concept induces a simple preference ordering for individuals, namely, we may define the binary relationships "$\succ_i$" and "$\sim_i$" meaning "is better than" and "is equivalent to," respectively, for player i, such that $s_i \succ_i s_i'$ if $s_i \in \Sigma_q^i$ and $s_i' \notin \Sigma_q^i$, and $s_i \sim_i s_i'$ if either $s_i \in \Sigma_q^i$ and $s_i' \in \Sigma_q^i$ or $s_i \notin \Sigma_q^i$ and $s_i' \notin \Sigma_q^i$. If $s_i \in \Sigma_q^i$, then s_i is said to be good enough. This interpretation of Σ_q^i differs from the interpretation of conventional individual rationality primarily in that, in addition to the best strategies it admits all other strategies that also qualify as good enough. An important feature of the satisficing approach, however, is that the individual preference orderings are not specified *a priori*, rather, they emerge *a posteriori* after all of the linkages between the players are accounted for in the interdependence function. In this sense, individual preferences are *emergent*.

Satisficing also induces a preference ordering for the group, namely, if $\mathbf{s} \succ \mathbf{s}'$ if $\mathbf{s} \in \mathbf{\Sigma}_q$ and $\mathbf{s}' \notin \mathbf{\Sigma}_q$, and $\mathbf{s} \sim \mathbf{s}'$ if either $\mathbf{s} \in \mathbf{\Sigma}_q$ and $\mathbf{s}' \in \mathbf{\Sigma}_q$ or $\mathbf{s} \notin \mathbf{\Sigma}_q$ and $\mathbf{s}' \notin \mathbf{\Sigma}_q$. Interpreting this preference ordering, however, is not straightforward; it is not immediately clear what it means to be good enough for the group. It would seem that the notion of group preference should convey the idea of harmonious behavior or at least some weak notion of social welfare, but satisficing game theory is completely neutral with regard to conflictive and coordinative aspects of the game. Both aspects can be accommodated by appropriately structuring the interdependence function. The fact that the interdependence function is able to account for conditional preference dependencies between players provides a coupling that permits them to widen their spheres of interest beyond their own myopic preferences. This widening of preferences does not guarantee that there is some coherent notion of harmony or disharmony. Although such implications are certainly not ruled out, selectability and rejectability do not favor either aspect. They may characterize benevolent or malevolent behavior and they may represent egoistic or altruistic interests. They may result in harmonious behavior or they may result in dysfunctional behavior. With competitive games, conflict can be introduced through conditional selectability and rejectability functions that account for the differences in goals and values of the players—the 'group preference' will then be to oppose one another. On the other hand, constructive coordinated behavior can be introduced through the same procedure, leading to a group preference of cooperation. Thus, as is the case with trying to define group rationality under the optimization regime, the notion of group rationality also appears to be somewhat elusive under the satisficing regime.

However, the apparent elusiveness of a simple interpretation of group rationality is not a weakness of the satisficing approach. On the contrary, it is a strength. Rather than a notion of group preference being defined as an aggregation of individual interests (a bottom-up, or *micro-to-macro,* approach) or imposed by a superplayer (a top-down, or *macro-to-micro,* approach), group preferences are *emergent,* in the sense that they are determined by the totality of the linked preferences (conditional and unconditional) and display themselves only as the links are forged. It is analogous to making a cake. The various ingredients (flour, sugar, water, heat, etc.) influence each other in complex ways, but it is not until they are all combined in proper proportions that a harmonious group notion of "cakeness" emerges.

Thus, just as the satisficing utility functions compare intrinsic, rather than extrinsic, attributes of strategies, the notions of both individual and group preference that emerge from their application are also intrinsic. They develop *within* the group of players as they evaluate their possibilities. Any notions of either group or individual rationality that emerge need not be anticipated or explicitly modeled. Rather than being imposed via either a top-down or bottom-up regime, such preferences may be characterized as inside-out, or *meso-to-micro/macro.* Both individual and group preferences emerge as consequences of local conditional interests that propagate throughout the community from the interdependent local to the interdependent global and from the conditional to the unconditional.

To illustrate the emergence of individual and group preferences, let us now address the Pot-Luck Dinner problem that was introduced in Example 1. To examine this problem from the satisficing point of view, we first need to specify operational definitions for selectability and is a function of six independent variables and may be factored,rejectability. Although there is not a unique way to frame this problem, let us take rejectability as cost of the meal and take selectability as enjoyment of the meal. The interdependence function is a function of six independent variables and may be factored, according to the chain rule, as

$$p_{S_L S_C S_M R_L R_C R_M}(x, y, z; u, v, w) = p_{S_C|S_L S_M R_L R_C R_M}(y|x, z; u, v, w) \cdot p_{S_L S_M R_L R_C R_M}(x, z; u, v, w), \quad (9)$$

where the subscripts L, C, and M correspond to Larry, Curly, and Moe, respectively. The mass function $p_{S_C|S_L S_M R_L R_C R_M}(y|x, z; u, v, w)$ expresses the selectability that Curly places on y, given that Larry selects x and rejects u, that Curly rejects v, and that Moe selects z and rejects w. From the hypothesis of the problem, we realize that, conditioned on Larry's selectability, Curly's selectability is independent of all other considerations, thus we can simplify this conditional selectability to obtain

$$p_{S_C|S_L S_M R_L R_C R_M}(y|x, z; u, v, w) = p_{S_C|S_L}(y|x).$$

Next, we apply the chain rule to the second term on the right hand side of (9), which yields

$$p_{S_L S_M R_L R_C R_M}(x, z; u, v, w) = p_{S_L S_M|R_L R_C R_M}(x, z|u, v, w) \cdot p_{R_L R_C R_M}(u, v, w).$$

But, given Curly's rejectability, the joint selectability for Larry and Moe is independent of all other considerations, so

$$p_{S_L S_M|R_L R_C R_M}(x, z|u, v, w) = p_{S_L S_M|R_C}(x, z|v).$$

By making the appropriate substitutions, (9) becomes

$$p_{S_L S_C S_M R_L R_C R_M}(x, y; z, u, v, w) = p_{S_C|S_L}(y|x) \cdot p_{S_L S_M|R_C}(x, z|v) \cdot p_{R_L R_C R_M}(u, v, w). \quad (10)$$

We desire to obtain Σ_q, the satisficing strategy profiles for the group and Σ_q^L, Σ_q^C, and Σ_q^M, the individually satisficing strategy sets for Larry, Curly, and Moe, respectively. To do so, we must specify each of the components of (10). To compute $p_{S_C|S_L}$, recall that Curly prefers beef to chicken to pork by respective factors of 2 conditioned on Larry preferring soup and that Curly is indifferent, conditioned on Larry preferring salad. We may express these relationships by the conditional selectability functions:

$$\begin{aligned} p_{S_C|S_L}(beef|soup) &= 4/7 & p_{S_C|S_L}(beef|sald) &= 1/3 \\ p_{S_C|S_L}(chkn|soup) &= 2/7 & p_{S_C|S_L}(chkn|sald) &= 1/3 \\ p_{S_C|S_L}(pork|soup) &= 1/7 & p_{S_C|S_L}(pork|sald) &= 1/3. \end{aligned}$$

To compute $p_{S_L S_M|R_C}$, we recall, given that Curly views pork as completely rejectable, Moe views lemon custard pie as highly selectable and Larry is indifferent. Given that Curly views beef as completely rejectable, Larry views soup as selectable, and Moe is indifferent; and given that Curly views chicken as completely rejectable, both Larry and Moe are indifferent. These relationships may be expressed as

$$\begin{aligned} p_{S_L S_M|R_C}(soup, lcst|pork) &= 0.5 \\ p_{S_L S_M|R_C}(soup, bcrm|pork) &= 0.0 \\ p_{S_L S_M|R_C}(sald, lcst|pork) &= 0.5 \\ p_{S_L S_M|R_C}(sald, bcrm|pork) &= 0.0, \end{aligned}$$

$$\begin{aligned} p_{S_L S_M|R_C}(soup, lcst|beef) &= 0.5 \\ p_{S_L S_M|R_C}(soup, bcrm|beef) &= 0.5 \\ p_{S_L S_M|R_C}(sald, lcst|beef) &= 0.0 \\ p_{S_L S_M|R_C}(sald, bcrm|beef) &= 0.0, \end{aligned}$$

and

$$\begin{aligned} p_{S_L S_M|R_C}(soup, lcst|chkn) &= 0.25 \\ p_{S_L S_M|R_C}(soup, bcrm|chkn) &= 0.25 \\ p_{S_L S_M|R_C}(sald, lcst|chkn) &= 0.25 \\ p_{S_L S_M|R_C}(sald, bcrm|chkn) &= 0.25. \end{aligned}$$

Lastly, we need to specify $p_{R_L R_C R_M}$, the joint rejectability function. This is done by normalizing the meal cost values in Table 1 by the total cost of all meals (e.g, $p_{R_L R_C R_M}(soup, beef, lcst) = 23/296$).

With the interdependence function so defined and letting $q = 1$, the jointly satisficing meals are as displayed in Table 2, each of which is good enough for the group, considered as a whole. The individually satisficing items, as obtained by computing the selectability and rejectability marginals, are also provided in Table 2 to be *soup, beef,* and *lemon custard.* Fortunately, this set of choices is also jointly satisficing without lowering the index of caution. Thus, all of the preferences are respected at a reasonable cost and, if pies are thrown, it is only for recreation, not retribution.

Jointly Satisficing		
Meal	$p_{S_L S_C S_M}$	$p_{R_L R_C R_M}$
$\{soup,\ beef,\ lcst\}$	0.237	0.078
$\{soup,\ chkn,\ lcst\}$	0.119	0.074
$\{soup,\ beef,\ bcrm\}$	0.149	0.091
$\{sald,\ pork,\ lcst\}$	0.080	0.074

Individually Satisficing			
Participant	Choice	p_S	p_R
Larry	*soup*	0.676	0.480
Curly	*beef*	0.494	0.351
Moe	*lcst*	0.655	0.459

Table 2: Jointly and individually satisficing choices for the Pot-Luck Dinner.

With the Pot-Luck Dinner, we see that, although total orderings for neither individuals nor the group are specified, we can use the *a priori* partial preference orderings from the problem statement to generate emergent, or *a posteriori*, group and individual orderings. *A posteriori* individual orderings also emerge from this exercise: Larry prefers soup to salad, Moe prefers lemon custard pie to banana cream pie, and Curly prefers beef to either chicken or pork. Note, however, that Curly is not required to impose a total ordering on his preferences (chicken versus pork)—this approach does not force the generation of unwarranted preference relationships. We see that a group-wide preference of avoiding conflict emerges, since the individually satisficing strategies are also jointly satisficing. This group desideratum was *not* specified *a priori*.

Conclusion

Satisficing game theory offers a way for the interests of the group and of the individuals to emerge through the conditional preference relationships that are expressed via the interdependence function due to its mathematical structure as a probability (but not with the usual semantics dealing with randomness or uncertainty). Just as the joint probability function is more than the totality of the marginals, the interdependence function is 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 lack of social concerns that group welfare can be expressed in terms of the welfare of individuals.

Optimization is a strongly entrenched procedure and dominates conventional decision-making methodologies. There is great comfort in following traditional paths, especially when those paths are founded on such a rich and enduring tradition as rational choice affords. But when synthesizing an artificial system, the designer must employ a more socially accommodating paradigm. The approach described in this paper seamlessly accounts for group and individual interests. Order can emerge through the local interactions that occur between agents who share common interests and who are willing to give deference to each other. Rather than depending upon the non-cooperative equilibria defined by individual rationality, this alternative may lead to the more socially realistic and valuable equilibrium of shared interests and acceptable compromises.

References

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

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

Bergson, A. (1938). A reformulation of certain aspects of welfare economics. *Quarterly Journal of Economics*, 52:310–334.

Eisen, M. (1969). *Introduction to Mathematical Probability Theory*. Prentice-Hall, Englewood Cliffs, NJ.

Goodrich, M. A., Stirling, W. C., and Boer, E. R. (2000). Satisficing revisited. *Minds and Machines*, 10:79–109.

Hogarth, R. M. and Reder, M. W., editors (1986). *Rational Choice*. Univ. Chicago Press, Chicago.

Luce, R. D. and Raiffa, H. (1957). *Games and Decisions*. John Wiley, New York.

Raiffa, H. (1968). *Decision Analysis*. Addison-Wesley, Reading, MA.

Samuelson, P. A. (1948). *Foundations of Economic Analysis*. Harvard University Press, Cambridge, MA.

Sandholm, T. W. (1999). Distributed rational decision making. In Weiss, G., editor, *Multiagent Systems*, chapter 5, pages 201–258. MIT Press, Cambridge, MA.

Sen, A. K. (1979). *Collective Choice and Social Welfare*. North-Holland, Amsterdam.

Shubik, M. (1982). *Game Theory in the Social Sciences*. MIT Press, Cambridge, MA.

Simon, H. A. (1955). A behavioral model of rational choice. *Quart. J. Econ.*, 59:99–118.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2):129–138.

Simon, H. A. (1990). Invariants of human behavior. *Annu. Rev. Psychol.*, 41:1–19.

Stirling, W. C. (2002). *Satisficing Games and Decision Making: with applications to engineering and computer science*. Cambridge University Press, Cambridge. Forthcoming.

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

Stirling, W. C. and Goodrich, M. A. (1999b). Satisficing games. *Information Sciences*, 114:255–280.

Taylor, M. (1987). *The Possibility of Cooperation*. Cambridge Univ. Press, Cambridge.

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