Autonomous Agents and Multi-Agent Systems, 5, 305–328, 2002

Satisficing Equilibria: A Non-Classical Theory of Games and Decisions

W. C. STIRLING wynn@ee.byu.edu
Electrical and Computer Engineering Department, Brigham Young University, Provo, UT 84602, USA

M. A. GOODRICH mike@cs.byu.edu
Computer Science Department, Brigham Young University, Provo, UT 84602, USA

D. J. PACKARD dennis_packard@byu.edu
Philosophy Department, Brigham Young University, Provo, UT 84602, USA

Abstract. Satisficing, or being "good enough," is the fundamental obligation of rational decision makers. We cannot rationally choose an option, even when we do not know of anything better, unless we know it is good enough. Unfortunately, we are not often in the position of knowing that there could be no better option, and hence that the option must be good enough. A complete search through all logical possibilities is often impractical, particularly in multi-agent contexts, due to excessive computational difficulty, modeling complexity, and uncertainty. It can be equally impractical, if it is even possible, to determine the cost of the additional required search to find an option that is good enough. In a departure from the traditional notion of satisficing as a species of bounded rationality, satisficing is here redefined in terms of a notion of intrinsic rationality. Epistemic utility theory serves as the philosophical foundation of a new praxeological decision-making paradigm of satisficing equilibria that is applicable to both single- and multiple-agent scenarios. All interagent relationships are modeled by an interdependence function that explicitly accommodates both self and group interest, from which multilateral and unilateral selectability and rejectability mass functions can be derived and compared via the praxeic likelihood ratio test.

Keywords: decision theory, game theory, satisficing, equilibria

1. Introduction

When Simon introduced the notion of "bounded rationality," he appropriated the term *satisficing* to describe a process of constructing expectations, or aspiration levels, of how good a solution might reasonably be achieved, and halting search when these expectations are met [25, 26, 27]. Subsequently, other notions of bounded rationality were generated by relaxing one or more assumptions of standard optimization approaches. Procedures such as augmenting the utility function with computational costs are species of constrained optimization (see, for example, [23, 10], and yield optimal solutions according to modified criteria. Regardless of the details of how a boundedly rational decision is obtained, however, the ultimate reason for adopting it is that it represents a compromise between performance and cost that is, by definition if for no other reason, "good enough."

The notion of being good enough is an underlying issue in all decision problems. While it is tacitly understood that a truly optimal solution to any well-framed problem is certainly good enough, the issue becomes problematic when other-than-optimal procedures are implemented. A characteristic of boundedly rational approaches *à la* Simon is that the aspiration levels are arbitrary; if set too low, performance may be sacrificed needlessly, and if set too high, the solution set may be empty. Although aspiration levels at least superficially establish minimum requirements, the approach provides no guidance as to how these levels are to be specified, and relies instead upon experience-derived expectations. But it is difficult to establish a good and practically attainable aspiration level without first exploring the limits of what is possible, that is, without first identifying optimal solutions—the very procedure satisficing is designed to circumvent. For simple, low-dimensional problems, specifying the aspiration levels may be noncontroversial. But, particularly with multi-agent systems, interdependence between decision makers becomes complex, and aspiration levels become conditional (what is satisfactory for me may depend upon what is satisfactory for you). The current state of affairs regarding aspiration levels does not appear to address completely the problem of specifying minimum requirements in multi-agent contexts. Perhaps what is needed is a notion of "minimum requirements" that does not depend upon aspiration levels.

A characteristic of boundedly rational approaches *à la* constrained optimization is that the modification to strict optimality accounts for criteria (such as computational costs) that are independent of actual performance. It is also possible, even for single-agent decision problems, that slightly modifying one or more assumptions of a bounded optimality-based solution methodology can lead to solutions that would not otherwise have been "good enough." This issue becomes even more delicate for multi-agent systems, since the interconnections between agents can be extremely sensitive, and slightly modifying an assumption can have a substantial effect on the outcome.

In this paper we propose a new definition for satisficing, or being good enough, that applies to both single- and multi-agent situations. As we define it, satisficing does not rely on notions of bounded rationality either *à la* aspiration levels or *à la* constrained optimization. Instead, our definition hinges on a new notion of rationality that is constructive, in that it leads to a practical framework for decision making. In Section 2 we motivate and formalize this definition and establish a new concept for equilibrium. In Section 3 we develop the mathematical structure of our decision framework, in Section 4 we extend it to the multi-agent case, and in Section 5 we apply our theory to the Prisoner's Dilemma game. Finally, in Section 6 we summarize our results, discuss some practical applications, suggest some research directions, and finish with reflections on the relationship between our engineering need to act and our scientific need to know.

2. A comparative paradigm

If the notion of "good enough" is to be made precise in a way analogous to, but distinct from, the way "best" is made precise, we must rely on a notion of rationality that does not depend upon optimization. At the same time, we must not retreat into *ad hocism*, that is, decision making based on notions of desirability or convenience without any definitive measures of quality. Heuristic procedures lack a means of self criticism; there is no way for the decision maker to evaluate its own performance. It may foster good, even very good, outcomes, but we cannot know if they are good enough.

2.1. An example

Strict optimality and pure heuristics represent two extreme views of decision making. We illustrate this situation by means of an example. One of the most widely studied of all games is the Prisoner's Dilemma. This game involves two agents, X_1 and X_2, who have been charged with a serious crime, arrested, and incarcerated in a way that precludes any communication between them. The prosecution has evidence sufficient only to convict them of a lesser crime with a short jail sentence. To get at least one conviction on the more serious crime, the prosecution entices each prisoner to give evidence against the other. Each prisoner has two options: either to confess (C), or to stonewall (S). Confession yields dropped charges if the other stonewalls; stonewalling yields the maximum sentence if the other confesses. If both stonewall, both receive short sentences; if both confess, both receive intermediate-length sentences.

The most widely recognized optimal solution is obtained via von Neumann-Morgenstern (VNM) game theory in terms of a payoff matrix, which is a table of juxtaposed individual expected utility functions. This matrix is illustrated in ordinal form in Table 1. The most widely accepted notion of what constitutes a viable solution to games such is this is the concept of equilibria. An *equilibrium* for a single-stage game such as Prisoner's Dilemma is an array of options (one for each player), or joint option, such that each player's individual option is acceptable to it according to some criterion. Three of the most widely used equilibrium concepts are dominance, Nash equilibria, and Pareto optimality (see, for example, [22] for an expanded discussion of game-theoretic equilibria). A joint option is a *dominant equilibrium* if each individual option is best for the corresponding player, no matter what options the other players choose. A joint option is a *Nash equilibrium* if, were any single agent to change its decision, it would reduce its level of satisfaction. A joint option is Pareto optimal if no single agent, by changing its decision, can

Table 1. Payoff matrix in ordinal form for traditional Prisoner's Dilemma game

	X_2	
X_1	S	C
S	(3, 3)	(1, 4)
C	(4, 1)	(2, 2)

Key: 4 = best; 3 = next best; 2 = next worst; 1 = worst.

increase its level of satisfaction without lowering the satisfaction level of at least one other agent.

By inspection, playing (C, C) is a dominant equilibrium. Furthermore, (C, C) is also the unique Nash equilibrium. Unfortunately, both of these equilibria are inferior to playing the "cooperative," Pareto-optimal equilibrium, (S, S), as they result in next-worst, rather than next-best, play. Hence, the dilemma.

One of the main characteristics of the VNM approach is that it abstracts the game from its context, or story line—all relevant information is captured by the expected utility functions. These expected utilities represent individual preferences as functions of joint actions; they do not represent joint preferences. It is only when the two utility functions are juxtaposed in the payoff matrix, as in Table 1, that the "game" emerges and strategies can be devised. Under this view, the players are assumed to be absolutely certain that self-interest is the only issue. The VNM approach does not countenance mixed motives. Accounting for any dispositions for coordinated behavior would change the payoff matrix and, hence, the game.[1] There is no room for equivocation. This is a powerful, but necessary, assumption under the VNM approach to this game.

This game has also been analyzed in the context of repeated play. Of course, the standard VNM approach must return the Nash solution, no mater how many times the game is played, but other approaches have revealed quite different behavior. A very interesting approach was introduced by Rapoport, who proposed a simple *tit-for-tat* rule of repeated play: start by stonewalling, thereafter play what the other player chose in the previous round. This purely heuristic rule won the Axelrod Tournament [1]. One of its main characteristics is that it does not abstract the game from the story—the actions are taken completely within the game context.

The VNM and Rapoport approaches to the Prisoner's Dilemma represent two extremes. With the VNM approach, all relevant knowledge is assumed to be compactly and precisely encoded into mathematical expressions, self interest is the only issue, and expected utility maximization is the operative solution concept. With the Rapoport approach, the players' dispositions enter into the decision, the solution concept is an *ad hoc* rule to be followed, and there are no attempts to rank-order options or maximize performance.[2]

2.2. *Intrinsic rationality*

The need for an alternative to the extremes of strict optimality on the one hand and pure heuristics on the other was recognized by Kreps, who observed that

> ... the real accomplishment will come in finding an interesting middle ground between hyperrational behavior and too much dependence on *ad hoc* notions of similarity and strategic expectations. When and if such a middle ground is found, then we may have useful theories for dealing with situations in which the rules are somewhat ambiguous [12, p. 184].

In matters of grammar, there are three degrees of comparison. The highest, or superlative, degree, is founded on the notion of being "best," which requires rank-ordering preferences for the consequences associated with the options. The evaluation of an option under the superlative paradigm requires comparing preferences for its consequences with the preferences for the consequences of all other possible options. Such evaluations are extrinsic, since they depend upon qualities that are external to the option under evaluation. In the literature, decision makers who operate under the superlative paradigm are said to be *substantively rational* [8]. To reflect the relativity of this rationality notion, we will refer to substantive rationality also as *extrinsic rationality.*

The lowest, or positive, degree of comparison is founded on the notion of being "good," which requires no explicit preference orderings or comparisons. Decision makers who operate under the positive paradigm are often characterized as being *procedurally rational* [8]. Procedures are developed by experts, and derive their validity from that authority. Substantive rationality tells us where to go, but not how to get there; procedural rationality tells us how to get there, but not where to go. Substantive rationality is viewed in terms of the outcomes it produces; procedural rationality is viewed in terms of the methods it employs.

There is a logical gap between the *superlative* paradigm and the *positive* paradigm. This gap is filled by the *comparative* paradigm of being "better." Literature involving the superlative paradigm, particularly that of statistical decision theory, game theory, optimal control theory, and operations research, is overwhelmingly vast, reflecting many decades of serious research and development. The positive paradigm, manifest in the form of heuristics, rule-based decision systems, and various other *ad hoc* techniques, has also been well-represented in the computer science, social science, and engineering literatures. A formally stated comparative paradigm, however, has not yet been well represented in the literature as a basis for a viable decision-making concept for general application.

As a first step toward the development of a comparative paradigm, we must formalize a notion of rationality that does not depend upon rank-ordering, yet is amenable to self policing. Rank-ordering is an extrinsic exercise involving inter-option comparisons; that is, comparing a given attribute of one option to the same attribute of another option or to a fixed standard. This is not, however, the only way to make comparisons. We may also consider making intra-option, or intrinsic, comparisons; that is, comparisons between different attributes of a given option.

The approach we advocate is very simple, and well-precedes any formal theories regarding its use: a common way people evaluate personal and business propositions is to compare potential gains with potential losses. Forming dichotomies provides a convenient way of accounting for the desirable and undesirable properties associated with options under consideration, and conforms with common praxeological[3] behavior. We will say that a decision is *intrinsically rational* if the gains achieved by making it equal or exceed the loss incurred, provided the gains and losses can be expressed in commensurable units. This is perhaps the most primitive form of self policing—it is local, rather than global.

As a formalized means of decision making, the concept appears in at least two very different contexts: economics and epistemology—the former intensely practical and concrete, the latter intensely theoretical and abstract. Economists implemented cost-benefit analysis in the 1930's. The usual procedure is to express all costs and benefits in monetary units, and to sanction a proposition if "the benefits to whomsoever they accrue are in excess of the estimated costs" [17]. Cost-benefit analysis is a useful way to reduce a complex problem to a simpler, more manageable one. One of its chief virtues is its fundamental simplicity.

Epistemic logic, or the classification of propositions on the basis of knowledge and belief regarding their content, has also made use of a comparative paradigm. The pioneering works of Levi [13] have lead to a distinctive school of thought regarding the evolution of knowledge corpora. Unlike the conventional epistemological doctrine of expanding one's knowledge corpus by adding information that has been justified as unequivocally true, Levi proposes the more modest goal of avoiding error. This theory has been detailed elsewhere (see [5, 31, 34]), and we provide here only a brief summary. The gist of it is that, given the task of determining which, if any, of a set of propositions should be retained in an agent's knowledge corpus, the agent should evaluate each proposition on the basis of two distinct criteria—first, the credal, or subjective, probability of it being true, and second, the informational value[4] of rejecting it, that is, the degree to which discarding the proposition results in a corpus with the kind of information that is demanded by the question. Thus, for a proposition to be admissible, it must be both believable and informative—all implausible or uninformative options should be rejected. Levi proposes a very simple test, based on the mathematical structure of probability theory, to perform such evaluations. He constructs an *expected epistemic utility function*, defined as the difference between credal probability and a certain constant (called the index of boldness) times another probability function, termed the informational-value-of-rejection probability. The set of options that maximizes this difference is the admissible set—all others are rejected. A key feature of this procedure is that the admissible set is generally not a singleton.

2.3. *Satisficing equilibria*

Following the tradition in economics and epistemology, we define a *satisficing option as any option that is intrinsically rational.* We retain the "satisficing" terminology because we are driven by exactly the same issue that motivated Simon's original usage of the term—to find options that are good enough by directly comparing attributes of options. We differ only in the standards used for comparison. Whereas Simon's approach is extrinsic, and compares attributes to fixed aspiration levels, our approach is intrinsic, and compares the positive attributes to the negative attributes of each option.

Under the superlative paradigm, once the utilities are in place, all actual decision making ceases—expected utility maximization is simply a matter of searching. Under the comparative paradigm, the utilities are used to provide rankings of attributes for each option. Thus, instead of making one global decision with respect to the entire collection of options, the comparative paradigm requires a separate local decision to be made for each option. Each option is thus subjected to individual scrutiny, and a decision is made either to retain it as an admissible choice or to reject it.

The comparative paradigm enables the formation of a precise definition of satisficing equilibria without the need to invoke any notions of optimality. We will say that an option is in a state of *satisficing equilibrium* if

S-1 The benefits derived from adopting it at least compensate for the costs incurred.
S-2 No other option provides more benefit without also costing more, or costs less without also providing less benefit.

Condition S-1 provides a weak notion of adequacy, and is a valid praxeological concept. Condition S-2 applies the domination principle to the cost-benefit framework to eliminate options that needlessly either sacrifice performance or incur expense.

In general, the set of satisficing equilibrium options will not be a singleton, and further refinement will be required before action can be taken. Each element of this set, however, enjoys the property that it results in a tradeoff that favors benefits over costs—there are no bad choices. The final choice depends upon the disposition of the decision maker. If the budget is tight, then the results may be commensurately low, but if cost is no object, then performance need not be sacrificed.[5] Whereas the essence of the superlative paradigm is "Nothing but the very best will do," and the essence of the positive paradigm is "It has always worked before," the essence of the comparative paradigm is "You get what you pay for." In this sense, satisficing equilibria provide an operational definition of being "good enough."

By weakening the principles of rationality we hope better to accommodate both individual and group preferences. Our philosophy is consistent with Levi's dictum that

> ... principles of coherent or consistent choice, belief, desire, etc. will have to be weak enough to accommodate a wide spectrum of potential changes in point of view. We may not be able to avoid some fixed principles, but they should be as weak as we can make them while still accommodating the demand for a systematic account [14, p. 24].

In particular, our goal is to develop a satisficing notion that is equally applicable to both group and individual points of view.

We do not view intrinsic rationality, and satisficing equilibrium in particular, as a replacement for substantive rationality, or for conventional game theory in general. We do, however, invite consideration of intrinsic rationality for situations where complexity and uncertainty make it difficult or inappropriate to apply classical superlative approaches. We view it as an additional tool, not as a competitor, and certainly not as a panacea. In the final analysis, all decision methods are subjective; they are simply tools to be used with judgment and skill.

3. Praxeic utility

Our notion of satisficing requires us to obtain dichotomous mathematical descriptions of cost and benefit, and then to use the resulting cost-benefit pair to identify members of the satisficing equilibrium set. Our approach is to adapt the mathematics of epistemic utility theory, originally developed for epistemological decision making (committing to beliefs), to the praxeological domain (taking action). Our interest in this theory is motivated by its striking compatibility with the comparative paradigm, since it permits costs and benefits to be characterized via probability measures that can be evaluated by the expected epistemic utility test discussed in Section 2.

Epistemic utility theory offers a number of significant features. First, it addresses an important issue of insuring comparable units by ascribing a unit of mass to both belief and informational value; second, it permits the evaluation of non-singleton sets of options; and third, it permits the extension of such evaluations to the multivariate case (that is, situations involving more than one decision maker or multiple attribute decision problems for a single agent).

Epistemic utility theory captures the essence of Popper's imperative: "Yet we must also stress that *truth is not the only aim of science.* We want more than truth: what we look for is *interesting truth*" [19, p. 229]. Although truth and informational value are natural semantic notions for cognitive decision making, they are not always natural for practical decision making. To apply the ideas of epistemic utility theory to practical decision making we must formulate praxeological analogs to the epistemological notions of truth and informational value.

A natural analog for *truth* is *success*, in the sense of achieving the fundamental goal of taking action. To formulate an analog for informational value, we observe that, just as the management of a finite amount of relevant information is important when inquiring after truth in the epistemological context, taking effective action requires the management of a finite amount of resource, such as wealth, materials, energy, safety, or other assets, in the praxeological context. Thus, an apt praxeological analog to *informational value of rejection* is *conservational value of rejection.* We thereby rephrase Popper's injunction to become: *we want more than success—what we look for is efficient success.* (Note that this is a softer goal than the superlative imperative of minimizing cost or maximizing payoff.) Thus, we change the context of the decision problem from one of acquiring information while avoiding error to one of conserving resources while avoiding failure. With the context shift from the epistemological issue of belief to the praxeological issue of action, we refer to the resulting utility function as *praxeic utility*, rather than epistemic utility.

Our development mathematically parallels Levi's original development of epistemic utility theory. We will refer to the degree of resource consumption as *rejectability*, and require rejectability to be expressed in terms of a function that conforms with the axioms of probability. We use this new terminology to emphasize the semantic distinction of using the mathematics of probability in a non-conventional way. Thus, for a finite action space, U, rejectability is expressed in terms of a mass function, $p_R\colon U \to [0, 1]$ such that $p_R(u) \geq 0$ for all $u \in U$ and $\sum_{u \in U} p_R(u) = 1$. Inefficient options (those with high resource consumption) should be highly rejectable; that is, if considerations of success are ignored, one should be prone to reject options that result in large costs, high energy consumption, exposure to hazard, etc. Normalizing p_R to be a mass function, which we will term the *rejectability mass function*, insures that the agent will have a unit of resource consumption to apportion among the elements of U. We may view p_R as the inutility of consuming resources. If $u \in U$ is rejected, then the agent conserves $1 - p_R(u)$ worth of its unit of resources which is therefore available to be applied to other options. We will usually assume that $p_R(u) > 0$ for all $u \in U$ (that is, there are no completely cost- or risk-free options).

Failure is avoided if successful options are not rejected, but efficiency, as well as success, must also be considered. Our approach to evaluating candidate sets of options for retention is to define the utility of not rejecting them in the interest of both success and resource conservation, and to retain the set that maximizes this utility. Suppose that implementing $u \in U$ would lead to success, and let $A \subset U$ be a set we are considering for retention. The utility of not rejecting A in the interest of avoiding failure is the *indicator function*,

$$I_A(u) = \begin{cases} 1 & \text{if } u \in A \\ 0 & \text{otherwise.} \end{cases} \quad [1]$$

This utility function is very even-handed; it is concerned only with the value of not rejecting A when u is successful—it does not count one option to be more desirable than another. Matters of desirability are deferred to the conservation-determining function, p_R, and not with accounting for the seriousness of one failure relative to another. We define the *praxeic utility* of not rejecting A when u is successful as the convex combination of the utility of avoiding failure and the utility of conserving resources:

$$\phi(A, u) = \alpha I_A(u) + (1-\alpha)\left(1 - \sum_{v \in A} p_R(v)\right), \quad [2]$$

where $\alpha \in [0, 1]$ is chosen to reflect the agent's personal weighting of these two desiderata—setting $\alpha = \frac{1}{2}$ associates equal concern for avoiding failure and conserving resources.

Clearly, praxeic utility is maximized when $A = \{u\}$ (the singleton set). Unfortunately, it generally is not known precisely which u will lead to success (or that only one u will do so), so we cannot simply reject $U \backslash \{u\}$, the complement of $\{u\}$. We may, however, possess information regarding the degree of success support possessed by each u. Let $p_S\colon U \to [0, 1]$ be a mass function that evaluates each option with respect to the degree to which it accomplishes the objective of the decision problem, independently of how much resource is consumed by implementing it. We will refer to the degree of success support as *selectability*, and we will term p_S the *selectability mass function.*

In the tradition of VNM theory, we may then calculate expected praxeic utility for any set $A \subset U$ by weighting the utility by the degree of success support associated with each u and summing over all $u \in U$. The *expected praxeic utility* is then

$$\begin{aligned} \bar{\phi}(A) &= \sum_{u \in U}\left[\alpha I_A(u) + (1-\alpha)\left(1 - \sum_{v \in A} p_R(v)\right)\right]p_S(u) \\ &= \alpha \sum_{u \in U} I_A(u) p_S(u) - (1-\alpha)\sum_{v \in A} p_R(v) \sum_{u \in U} p_S(u) + (1-\alpha)\sum_{u \in U} p_S(u) \\ &= \alpha \sum_{v \in A} p_S(v) - (1-\alpha)\sum_{v \in A} p_R(v) + (1-\alpha), \end{aligned} \quad [3]$$

Since $\sum_{u\in U} p_S(u) = 1$. Because VNM utility functions are invariant with respect to positive affine transformations, dividing by α and ignoring the constant term yields a more convenient but equivalent form for expected praxeic utility:

$$\bar{\varphi}(A) = \sum_{u\in A}[p_S(u) - bp_R(u)], \quad [4]$$

where $b = \frac{1-\alpha}{\alpha}$.

We may obtain the largest set of options for which the selectability is greater than or equal to b times the rejectability by choosing the set that maximizes[6] expected praxeic utility, resulting in the *satisficing set*

$$\Sigma_b = \arg\max_{A\subset U} \bar{\varphi}(A) = \{u \in U\colon p_S(u) \geq bp_R(u)\} = \left\{u \in U\colon \frac{p_S(u)}{p_R(u)} \geq b\right\}. \quad [5]$$

Σ_b is the set of all options for which the benefits outweigh the costs, as scaled by b. The parameter, b, is the *boldness* of the decision problem, and parameterizes the degree to which the agent is willing to risk rejecting successful options in the interest of conserving resources. Setting $b > 1$ attributes more weight to resource conservation than to success. Setting $b \leq 1$ ensures that $\Sigma_b \neq \varnothing$, since otherwise $p_S(u) < bp_R(u)$ for all $u \in U$, which would imply $1 = \sum_{u\in U} p_S(u) < b\sum_{u\in U} p_R(u) = b$, a contradiction. We will refer to (5) as the Praxeic Likelihood Ratio Test (PLRT).

The set Σ_b contains all options such that the benefits of adoption outweigh the cost and therefore is consistent with condition S-1, but it may include options that needlessly sacrifice benefit or incur cost. We may eliminate such options as follows. For every $u \in U$ let

$$\begin{aligned} B_S(u) &= \{v \in U\colon p_R(v) < p_R(u) \text{ and } p_S(v) \geq p_S(u)\} \\ B_R(u) &= \{v \in U\colon p_R(v) \leq p_R(u) \text{ and } p_S(v) > p_S(u)\}, \end{aligned} \quad [6]$$

and define the set of options that are *strictly better* than u:

$$B(u) = B_S(u) \cup B_R(u); \quad [7]$$

that is, $B(u)$ consists of all possible options that are either less rejectable and not less selectable than u, or are not more rejectable and more selectable than u. If $B(u) = \varnothing$, then no options can be preferred to u on the basis of both selectability and rejectability. The set of equilibrium options is consistent with condition S-2, and is defined as

$$\mathscr{E} = \{u \in U\colon B(u) = \varnothing\}. \quad [8]$$

The set of *satisficing equilibrium* options is the intersection of the satisficing and equilibrium sets:

$$\mathfrak{S}_b = \mathscr{E} \cap \Sigma_b. \quad [9]$$

4. Multi-agent systems

4.1. Background

Perhaps the most intensely studied multi-agent decision formalism is game theory. Game theory is built on one basic principle: individual self-interest. Under the VNM view, when a player is faced with uncertainty, the only justifiable course of action is to choose an action that maximizes expected utility, conditioned on the expected actions of other players. This is called the *principle of individual rationality*. For two-person constant-sum games, individual self-interest is the only possible 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) is not individually rational for some agent, that exclusively self-interested agent would not be a party to such a joint action. This is the crux of the Prisoner's Dilemma.

This limitation of game theory has been observed by Luce and Raiffa:

> ... general game theory seems to be in part a sociological theory which does not include any sociological assumptions ... it may be too much to ask that any sociology be derived from the single assumption of individual rationality [15, p. 196].

With conventional game theory, attitudes such as cooperation and conflict are not built into the expected utility functions. Expected utilities are functions only of the actions, or strategies, of the player, and not of their preferences. It is not until the expected utilities of all players are juxtaposed that the "game" aspect of the decision problem emerges.

Social choice theory has also built largely on the foundation of individual rationality. For example, under Harsanyi's formulation, a social-welfare function is a positive linear combination of individual expected utilities, and each player proceeds under the substantively rational paradigm by maximizing its social-welfare function [7, 9]. Each individual expected utility in this combination serves to map group actions to individual preferences, then the weighted combination of these individual preferences is used to define the preference of society. Thus, this social-welfare function requires a sequence of two couplings: the first provides a group-action to individual-preference coupling, and the second provides an individual-preference to group-preference coupling.

Perhaps the most fundamental type of coupling, however, is that of individual preferences to individual preferences. Suppose decision maker X_i becomes aware of decision maker X_j (perhaps not exclusively). The most fundamental question X_i can consider, in this regard, is "How will X_j's preferences affect my preferences?" VNM utility, however, requires X_i to consider instead the question "How will X_j's actions affect my preferences?" The answers to this latter question are what comprise the individual utility functions, which in turn comprise Harsanyi's social-welfare function. Unfortunately, this function does not necessarily reflect the most fundamental relationships that exist between decision makers, and does not provide a notion of group rationality that is distinct from individual rationality.

4.2. Interagent modeling

We suggest that there should be explicit linkages between different individual preferences and between group preference and individual preferences. These linkages, however, may be quite different from the type of preference orderings that are the substance of conventional utility theory. An act by any single member X_i of a multi-agent system, $\{X_1, \dots, X_N\}$, has possible ramifications throughout the entire community. Some of the agents may be benefited, some may be damaged, and some may be indifferent. Furthermore, although the single agent may perform the act for its own benefit, or for the benefit of other agents, or for the benefit of the entire system, 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 single agent itself, perhaps by other agents, and perhaps by the entire community. Although these undesirable consequences may be defined independently from the benefits, the measures associated with benefits and costs cannot be specified independently of each other, due to the possibility of interaction. A critical aspect of modeling the behavior of such a collection, therefore, is the means of representing the interdependence of both positive and negative consequences of all possible multipartite options that could be undertaken.

This representation must account for group interest as well as for self interest. One way to accommodate group interest is for agent preferences to be sensitive to the preferences, as well as the choices, of other agents. To develop the means of representing these preferences, we require some definitions.

Definition 1 A *mixture*[7] is any subset of agents considered in terms of their interaction with each other, exclusively of possible interactions with other agents not in the subset.

A *selectability mixture*, denoted $\mathcal{S} = S_{i_1} \dots S_{i_k}$, is a mixture consisting of agents $X_{i_1} \dots X_{i_k}$ being considered from the point of view of success. The *joint selectability mixture* is the selectability mixture consisting of all agents in the system, denoted $\mathbf{S} = S_1 \dots S_N$. A *myopic* selectability mixture is a mixture of the form $\mathcal{S} = S_i$; that is, when an agent views success as though it were functioning with complete disregard for all other agents.

A *rejectability mixture*, denoted $\mathcal{R} = R_{j_1} \dots R_{j_\ell}$, is a mixture consisting of agents $X_{j_1} \dots X_{j_\ell}$ being considered from the point of view of resource consumption. The *joint rejectability mixture* is the rejectability mixture consisting of all agents in the system, denoted $\mathbf{R} = R_1 \dots R_N$. A *myopic* rejectability mixture is a mixture of the form $\mathcal{R} = R_i$.

An *intermixture* is the concatenation of a selectability mixture and a rejectability mixture, and is denoted $\mathcal{S}\mathcal{R} = S_{i_1} \dots S_{i_k} R_{j_1} \dots R_{j_\ell}$. The *joint intermixture* is the concatenation of the joint selectability and joint rejectability mixtures, and is denoted $\mathbf{SR} = S_1 \dots S_N R_1 \dots R_N$. A *myopic* intermixture is a mixture of the form $\mathcal{S}\mathcal{R} = S_i R_i$.

Definition 2 Given arbitrary action spaces U_i, $i = 1, \dots, N$, the *product action space*, denoted $\mathbf{U} = U_1 \times \cdots \times U_N$ is the set of all N-tuples $\mathbf{u} = (u_1, \dots, u_N)$

where $u_i \in U_i$. The *selectability action* space associated with a selectability mixture $\mathscr{S} = S_{i_1} \dots S_{i_k}$ is the product space $\mathbf{U}_{\mathscr{S}} = U_{i_1} \times \cdots \times U_{i_k}$. The *rejectability action space* associated with a rejectability mixture $\mathscr{R} = R_{j_1} \dots R_{j_\ell}$ is the product space $\mathbf{U}_{\mathscr{R}} = U_{j_1} \times \cdots \times U_{j_\ell}$. The *interaction space* associated with an intermixture $\mathscr{S}\mathscr{R} = S_{i_1} \cdots S_{i_k} R_{j_1} \cdots R_{j_\ell}$ is the product space $\mathbf{U}_{\mathscr{S}\mathscr{R}} = U_{i_1} \times \cdots \times U_{i_k} \times U_{j_1} \times \cdots \times U_{j_\ell}$. The joint interaction space is $\mathbf{U}_{\mathbf{SR}} = \mathbf{U} \times \mathbf{U}$.

Definition 3 A *selectability mass function* (smf) for the mixture $\mathscr{S} = \{S_{i_1} \dots S_{i_k}\}$ is a mass function denoted $p_{\mathscr{S}} = p_{S_{i_1} \dots S_{i_k}}: \mathbf{U}_{\mathscr{S}} \to [0, 1]$. The joint smf is an smf for $\mathbf{S}$, denoted $p_{\mathbf{S}}$.

A *rejectability mass function* (rmf) for the mixture $\mathscr{R} = \{R_{j_1}, \dots, R_{j_\ell}\}$ is a mass function denoted $p_{\mathscr{R}} = p_{R_{j_1} \dots R_{j_\ell}}: \mathbf{U}_{\mathscr{R}} \to [0, 1]$. The joint rmf is a rmf for $\mathbf{R}$, denoted $p_{\mathbf{R}}$.

An *interdependence mass function* (imf) for the intermixture $\mathscr{S}\mathscr{R} = \{S_{i_1} \dots S_{i_k} R_{j_1} \dots R_{j_\ell}\}$ is a mass function denoted $p_{\mathscr{S}\mathscr{R}} = p_{S_{i_1} \dots S_{i_k} R_{j_1} \dots R_{j_\ell}}: \mathbf{U}_{\mathscr{S}} \times \mathbf{U}_{\mathscr{R}} \to [0, 1]$. The joint imf is an imf for $\mathbf{SR}$, denoted $p_{\mathbf{SR}}$.

Let $\mathbf{v} \in \mathbf{U}_{\mathscr{S}}$ and $\mathbf{w} \in \mathbf{U}_{\mathscr{R}}$ be two joint options. Then $p_{\mathscr{S}\mathscr{R}}(\mathbf{v}, \mathbf{w})$ is a representation of the success support associated with $\mathbf{v}$ and the resource consumption associated with $\mathbf{w}$ when the two joint options are viewed simultaneously. In other words, $p_{\mathscr{S}\mathscr{R}}(\mathbf{v}, \mathbf{w})$ is the mass associated with adopting $\mathbf{v}$ in the interest of success and rejecting $\mathbf{w}$ in the interest of conserving resources.

The joint imf provides a complete description of the individual and interagent relationships in terms of their positive and negative consequences, but specifying the imf can be complex, especially for high-dimensional systems. Fortunately, the structure of the imf makes it possible to generate global behavior (the joint imf) from interdependence mass functions for intermixtures. To develop this theory we require some additional definitions.

Definition 4 Given an intermixture $\mathscr{S}\mathscr{R} = S_{i_1} \dots S_{i_k} R_{j_1} \dots R_{j_\ell}$, a subintermixture of $\mathscr{S}\mathscr{R}$ is an intermixture formed by concatenating subsets of $\mathscr{S}$ and $\mathscr{R}$: $\mathscr{S}_1\mathscr{R}_1 = S_{i_{p_1}} \dots S_{i_{p_q}} R_{j_{r_1}} \dots R_{j_{r_s}}$, where $\{i_{p_1}, \dots, i_{p_q}\} \subset \{i_1, \dots, i_k\}$ and $\{j_{r_1}, \dots, j_{r_s}\} \subset \{j_1, \dots, j_\ell\}$. We shall use the notation $\mathscr{S}_1\mathscr{R}_1 \subset \mathscr{S}\mathscr{R}$ to indicate that $\mathscr{S}_1\mathscr{R}_1$ is a subintermixture of $\mathscr{S}\mathscr{R}$.

The *$\mathscr{S}\mathscr{R}$-complementary subintermixture* associated with a subintermixture $\mathscr{S}_1\mathscr{R}_1$ of an intermixture $\mathscr{S}\mathscr{R}$, denoted $\mathscr{S}\mathscr{R} \backslash \mathscr{S}_1\mathscr{R}_1$, is an intermixture created by concatenating the selectability and rejectability mixtures formed by the relative compliments of $\mathscr{S}_1$ and $\mathscr{R}_1$. Clearly, $\mathscr{S}\mathscr{R} \backslash \mathscr{S}_1\mathscr{R}_1 \subset \mathscr{S}\mathscr{R}$. We will say that $\mathscr{S}\mathscr{R}$ is the union of $\mathscr{S}\mathscr{R} \backslash \mathscr{S}_1\mathscr{R}_1$ and $\mathscr{S}_1\mathscr{R}_1$, denoted $\mathscr{S}\mathscr{R} = \mathscr{S}\mathscr{R} \backslash \mathscr{S}_1\mathscr{R}_1 \cup \mathscr{S}_1\mathscr{R}_1$.

Definition 5 Let $\mathscr{S}\mathscr{R}$ be an intermixture with subintermixture $\mathscr{S}_1\mathscr{R}_1$. A *conditional interdependence mass function*, denoted $p_{\mathscr{S}\mathscr{R} \backslash \mathscr{S}_1\mathscr{R}_1 | \mathscr{S}_1\mathscr{R}_1}$, is a mapping of $(\mathbf{U}_{\mathscr{S}\mathscr{R} \backslash \mathscr{S}_1\mathscr{R}_1} \times \mathbf{U}_{\mathscr{S}_1\mathscr{R}_1})$ into $[0, 1]$ such that, for every $\mathbf{v} \in \mathbf{U}_{\mathscr{S}_1\mathscr{R}_1}$, $p_{\mathscr{S}\mathscr{R} \backslash \mathscr{S}_1\mathscr{R}_1 | \mathscr{S}_1\mathscr{R}_1}(\cdot | \mathbf{v})$ is a mass function on $\mathbf{U}_{\mathscr{S}\mathscr{R} \backslash \mathscr{S}_1\mathscr{R}_1}$.

We require all conditional interdependence mass functions to be consistent with interdependence mass functions, that is, For $\mathscr{S}\mathscr{R}$ an arbitrary intermixture with subintermixture $\mathscr{S}_1\mathscr{R}_1$ with $\mathbf{w} \in \mathscr{S}\mathscr{R}\backslash\mathscr{S}_1\mathscr{R}_1$ and $\mathbf{v} \in \mathscr{S}_1\mathscr{R}_1$, we have

$$p_{\mathscr{S}\mathscr{R}}(\mathbf{v}, \mathbf{w}) = p_{\mathscr{S}\mathscr{R}\backslash\mathscr{S}_1\mathscr{R}_1|\mathscr{S}_1\mathscr{R}_1}(\mathbf{w}|\mathbf{v}) \cdot p_{\mathscr{S}_1\mathscr{R}_1}(\mathbf{v}). \quad [10]$$

Example 1. Let $\{X_1, X_2, X_3\}$ be a multi-agent system, and let $\mathscr{S} = S_1 S_2$ and $\mathscr{R} = R_3$. Then $\mathscr{S}\mathscr{R} = S_1 S_2 R_3$ and $\mathbf{SR}\backslash\mathscr{S}\mathscr{R} = S_3 R_1 R_2$. The imf is

$$p_{S_1S_2S_3R_1R_2R_3}(v_1, v_2, v_3, w_1, w_2, w_3) = p_{S_3R_1R_2|S_1S_2R_3}(v_3, w_1, w_2|v_1, v_2, w_3) \cdot p_{S_1|S_2R_3}(v_1|v_2, w_3) \cdot p_{S_2R_3}(v_2, w_3). \quad [11]$$

$p_{S_3R_1R_2|S_1S_2R_3}(v_3, w_1, w_2|v_1, v_2, w_3)$ is the conditional selectability and rejectability associated with X_3 selecting action v_3, X_1 rejecting action w_1, and X_2 rejecting action w_2, conditioned on X_1 ascribing its entire unit of selectability mass to v_1, X_2 ascribing its entire unit of selectability mass to v_2, and X_3 ascribing its entire unit of rejectability mass to w_3. The other two conditional interdependence mass functions are interpreted similarly.

An advantage of factoring the imf into conditional components is that preferences conditioned on various situations makes it possible to characterize global preferences in terms of conditional local preferences. As noted by Pearl [18], it is often easier to specify conditional local characteristics rather than unconditional joint global characteristics. Conditional probabilities permit local, or specific responses to be characterized; they possess modularity features similar to logical production rules.

Factorizations make it possible to simplify the structure of the interdependence function when agents or agent groups are indifferent to each other, in which case the interdependence function will factor into products of marginal interdependence functions. When agents are conditionally independent of each other, various Markovian structures may be present, which also results in a significant simplification of the interdependence function. The interdependence function thus provides a parsimonious way of accounting for interagent relationships. This is important, because as the dimensionality of the agent system becomes large, the number of possible inter-agent relationships grows combinatorially. Thus, it is necessary to exploit any simplifying structure that may exist, while at the same time not ignoring dependencies that are critical. Multivariate probability theory is ideally suited to this task. While it can be complex, it is not more complex than it needs to be. It fulfills the dictum offered by Palmer:

> ... complexity is no argument against a theoretical approach if the complexity arises not out of the theory itself but out of the material which any theory ought to handle [16, p. 176].

5. Satisficing games

5.1. Multilateral and unilateral decisions

The role of the joint imf is to characterize all of the logically possible interagent preference relationships that exist between the members of the multi-agent system. Once this function is defined, dichotomies may be extracted from it by computing appropriate marginals. We examine two perspectives. The first is appropriate for *multilateral,* or group decision making, and the second is appropriate for *unilateral,* or individual decision making. We then expose some relationships between group decision making and individual decision making.

To form a multilateral decision, we compute the joint selectability and rejectability marginals as

$$p_{\mathbf{S}}(\mathbf{u}) = \sum_{\mathbf{v}\in\mathbf{U}} p_{\mathbf{SR}}(\mathbf{u}, \mathbf{v}) \quad [12]$$

$$p_{\mathbf{R}}(\mathbf{v}) = \sum_{\mathbf{u}\in\mathbf{U}} p_{\mathbf{SR}}(\mathbf{u}, \mathbf{v}) \quad [13]$$

for all $(\mathbf{u}, \mathbf{v}) \in \mathbf{U} \times \mathbf{U}$. These marginals may then be used to generate the *multilateral satisficing set* according to a straightforward extension of the PLRT:

$$\mathbf{\Sigma}_b = \{\mathbf{u} \in \mathbf{U}\colon p_{\mathbf{S}}(\mathbf{u}) \geq b p_{\mathbf{R}}(\mathbf{u})\}. \quad [14]$$

The *multilateral satisficing equilibrium set* is

$$\mathfrak{S}_b = \mathscr{E} \cap \mathbf{\Sigma}_b, \quad [15]$$

where $\mathscr{E}$ is the set of multilateral satisficing equilibrium options defined in accordance with Eq. [8].

The multilateral satisficing set given by Eq. [14] provides the set of jointly satisficing options as viewed from the group perspective. To obtain the individual perspective, we must compute the univariate satisficing sets for each agent. We first computing the univariate selectability and rejectability marginals, given, for agent X_i, by

$$p_{S_i}(u_i) = \sum_{j\neq i} p_{S_1\cdots S_N}(u_1, \ldots, u_{i-1}, u_i, u_{i+1} \ldots, u_N) \quad [16]$$

$$p_{R_i}(u_i) = \sum_{j\neq i} p_{R_1\cdots R_N}(u_1, \ldots, u_{i-1}, u_i, u_{i+1} \ldots, u_N). \quad [17]$$

The *univariate satisficing set* for X_i is

$$\Sigma_b^i = \{u \in U_i\colon p_{S_i}(u) \geq b p_{R_i}(u)\}, \quad [18]$$

and the *unilateral satisficing equilibrium set* is

$$\mathfrak{S}_b^i = \mathscr{E}_i \cap \Sigma_b^i. \quad [19]$$

We may examine the relationships between group-level satisficing and individual-level satisficing, that is, between $\mathbf{\Sigma}_b$ and the family of individually satisficing sets $\{\Sigma_b^i, i = 1, \ldots, N\}$. Under the special circumstance of all agents acting completely independently of each other, the joint distributions will factor into the product of the marginals, yielding

$$p_{\mathbf{S}}(\mathbf{u}) = p_{S_1}(u_1) \cdots p_{S_N}(u_N) \tag{20}$$

$$p_{\mathbf{R}}(\mathbf{u}) = p_{R_1}(u_1) \cdots p_{R_N}(u_N), \tag{21}$$

in which case

$$\mathbf{\Sigma}_b = \Sigma_b^1 \times \cdots \times \Sigma_b^N. \tag{22}$$

In general, however, the interdependencies among agents will render Eq. [22] invalid, and jointly satisficing actions will not necessarily be individually satisficing for all agents. Jointly satisficing and individually satisficing solutions must, however, obey a weak consistency relationship in that, if u_i is individually satisficing for X_i, then it must be an element of some jointly satisficing vector. To establish this claim, we prove the contrapositive: if u_i is not the i-th element of any $\mathbf{u} \in \mathbf{\Sigma}_b$, then $u_i \notin \Sigma_b^i$. Without loss of generality, let $i = 1$. By hypothesis, $p_{\mathbf{S}}(u_1, \mathbf{v}) < b p_{\mathbf{R}}(u_1, \mathbf{v})$ for all $\mathbf{v} \in U_2 \times \cdots \times U_N$, so $p_{S_1}(u_1) = \sum_{\mathbf{v}} p_{\mathbf{S}}(u_1, \mathbf{v}) < b \sum_{\mathbf{v}} p_{\mathbf{R}}(u_1, \mathbf{v}) = b p_{R_1}(u_1)$, hence $u_1 \notin \Sigma_b^1$. This result says that no one is ever completely frozen out of a deal—every individual has a seat at the negotiating table. This is perhaps the weakest condition under which meaningful negotiations are possible. It is not true, however, that there always exists a jointly satisficing option vector such that each component is simultaneously individually satisficing—it is possible that $\mathbf{\Sigma}_b \cap (\Sigma_b^1 \times \cdots \times \Sigma_b^N) = \varnothing$, as evidenced by the following simple example. Let $U_1 = (u, u')$, $U_2 = (v, v')$, and let $p_{S_1 S_2}$ and $p_{R_1 R_2}$ be given in Table 2. From this we easily obtain, for $b = 1$, that $\mathbf{\Sigma}_b = \{(u, v'), (u', v)\}$ and $\Sigma_b^1 \times \Sigma_b^2 = \{(u, v)\}$. As b is decreased, however, these decision sets change. At the value $b = 0.933$, we have $\mathbf{\Sigma}_b = \{(u, v), (u, v'), (u', v)\}$ and $\Sigma_b^1 \times \Sigma_b^2 = \{(u, v)\}$, so $\mathbf{\Sigma}_b \cap (\Sigma_b^1 \times \Sigma_b^2) = \{(u, v)\}$, and there does indeed exist, at this level of boldness, a jointly satisficing option vector whose components are both individually satisficing. Reducing the boldness, b, is a controlled way to relax the standards of intrinsic rationality, which may be necessary in difficult situations if a compromise is to be reached. The amount b must be reduced below unity is a measure of the amount of compromising needed to reach a mutually acceptable solution.

Table 2. Selectability and rejectability functions

$p_{S_1 S_2}$	v	v'
u	0.28	0.40
u'	0.30	0.02

$p_{R_1 R_2}$	v	v'
u	0.30	0.30
u'	0.20	0.20

5.2. The Prisoner's dilemma

Let us now apply satisficing game theory to the Prisoner's Dilemma game. Our task is to define the imf, from which the smf and rmf can be obtained and compared for each joint option. To generate the imf, we must first define operational notions for success and resource consumption.

The fundamental objective of the players, as expressed by their utility functions, is to get out of jail quickly—self-interest is the primary consideration. Laboratory experiments with randomly selected humans, however, result in the cooperative solution being chosen relatively frequently with single play and no communication between the participants [36]. This evidence suggests that for the game to be a model of human behavior, there may exist motives in addition to self interest. An obvious possibility is group interest, which apparently emerges in repeated play as a result of learned cooperation.[8] We view the players of this game as individuals who are concerned primarily with their own task, but at the same time have a degree of consideration for other players' difficulties, and consider it a cost to them if an action they take makes it difficult for others. Such players are *enlightened liberals*, who are intent upon pursuing their own self interest but give some deference to the interests of the group in general. Accordingly, we proceed by associating success with reducing individual jail time and associating resource consumption with increasing group jail time. Thus, short individual sentences will have high selectability, and long group sentences will have high rejectability.

The notion of group interest can have significance only if the players each acknowledge some form of dependence on the other, however weak it may be. To the extent that these dependencies reinforce each other, the players implicitly forge a joint opinion regarding the relative merits of cooperation and conflict. The success of Rapoport's approach suggests that there may be (at least) two attitudes in the minds of the players that may affect their decisions: (a) a propensity for *dissociation*, that is, for the agents to go their separate ways without regard for coordination, and (b) a propensity for *vulnerability*, that is, for the agents to expose themselves to individual risk in the hope of improving the joint outcome.

Let $\alpha \in [0, 1]$ be a measure of the joint value the players place on rejecting the joint option (S, S). We may identify α as a *dissociation index:* if $\alpha \approx 1$, the agents are completely dissociated and coordination is unlikely. Also, let $\beta \in [0, 1]$ be a measure of the joint value placed on rejecting the joint option (C, C). β may be viewed as a *vulnerability index:* $\beta \approx 1$ means the agents are each willing to risk a long jail sentence in the hope of both obtaining a shorter one. A condition of high dissociation *and* high vulnerability would indicate a contradictory unconcern for possible cooperation while implying a hope for cooperation. We may prohibit this situation by imposing the constraint that $\alpha + \beta \leq 1$. If, for example, $\alpha = 1$ and $\beta = 0$, then self-interest is the only consideration. If, however, $\alpha = 0$ and $\beta = 1$, then the players are willing to assume high risk to achieve cooperation.

If the sentence lengths are independent of the identity of the agents, then the joint rejectability of (C, S) should be equal to the joint rejectability of (S, C). With this assumption and the constraints on α and β, we define the joint rmf:

$$p_{R_1R_2}(S, S) = \alpha, \qquad p_{R_1R_2}(S, C) = \frac{1-\alpha-\beta}{2}$$
$$p_{R_1R_2}(C, S) = \frac{1-\alpha-\beta}{2} \qquad p_{R_1R_2}(C, C) = \beta. \qquad [23]$$

Although selectability deals with individual objectives, it is a joint consideration, since the consequences of the agents' decisions are not independent, and thus preferences cannot be independent. A convenient way to express this dependency is to exploit the probabilistic structure of the imf and to compute joint selectability conditioned on joint rejectability, $p_{S_1S_2|R_1R_2}(v_1, v_2|w_1, w_2)$ for all (v_1, v_2) and (w_1, w_2) in $U \times U = \mathbf{U} = \{(S, S), (S, C), (C, S), (C, C)\}$, from which the imf may be obtained by the product rule:

$$p_{S_1S_2R_1R_2}(v_1, v_2, w_1, w_2) = p_{S_1S_2|R_1R_2}(v_1, v_2|w_1, w_2) \cdot p_{R_1R_2}(w_1, w_2). \qquad [24]$$

The conditional selectability mass function, $p_{S_1S_2|R_1R_2}(v_1, v_2|w_1, w_2)$, characterizes the selectability of the joint option (v_1, v_2) given that the agents jointly place all of their rejectability mass on (w_1, w_2). We may compute the conditional selectability by invoking straightforward and intuitive rules of the form: "If X_1 and X_2 jointly reject (w_1, w_2), then they should jointly select (v_1, v_2)." Let, say, $(w_1, w_2) = (S, S)$, that is, the agents jointly reject stonewalling. Given this situation, it is trivially obvious that the preferred joint option is to both confess.[9] We may encode this rule into the conditional selectability mass function by placing all of the mass of on the joint option (C, C), that is, $p_{S_1S_2|R_1R_2}(C, C|S, S) = 1$. If exactly one agent rejects stonewalling, then it is obvious that, in this case as well, the preferred joint option is to both confess, consequently, $p_{S_1S_2|R_1R_2}(C, C|S, C) = 1$ and $p_{S_1S_2|R_1R_2}(C, C|C, S) = 1$. Finally, if both agents reject confessing, then it is clear that $p_{S_1S_2|R_1R_2}(S, S|C, C) = 1$. Table 3 summarizes the structure of this conditional credibility function.

Substituting the conditional selectability interdependence function given by Table 3 and the joint rejectability given by (23) into (24) and applying Eq. [12]

Table 3. Conditional credibility for the satisficing Prisoner's Dilemma

$p_{S_1S_2\|R_1R_2}(v_1, v_2\|w_1, w_2)$				
	(w_1, w_2)			
(v_1, v_2)	(S, S)	(S, C)	(C, S)	(C, C)
(S, S)	0	0	0	1
(S, C)	0	0	0	0
(C, S)	0	0	0	0
(C, C)	1	1	1	0

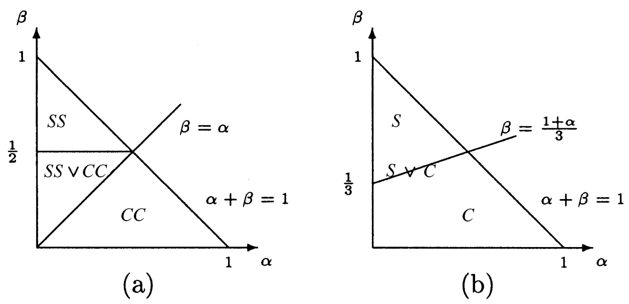

Figure 1. Decision regions for (a) bilateral decisions and (b) unilateral decisions.

yields the joint selectability function:

$$p_{S_1S_2}(S, S) = \beta \qquad p_{S_1S_2}(S, C) = 0 \tag{25a}$$

$$p_{S_1S_2}(C, S) = 0 \qquad p_{S_1S_2}(C, C) = 1 - \beta. \tag{25b}$$

Comparing the bilateral selectability, given by Eq. [25a], with the bilateral rejectability, given by Eq. [23], we obtain the bilateral satisficing set, which consists of all decision pairs, $(w_1, w_2) \in \mathbf{U}$, such that $p_{S_1S_2}(w_1, w_2) \geq bp_{R_1R_2}(w_1, w_2)$. Thus, parameterized by α and β, this set is, for the special case $b = 1$,

$$\mathfrak{S}_b = \begin{cases} \{(S, S)\} & \text{for } \beta \geq \frac{1}{2} \\ \{(C, C)\} & \text{for } \beta \leq \alpha \\ \{(S, S), (C, C)\} & \text{for } \alpha < \beta < \frac{1}{2} \end{cases} . \tag{26}$$

These regions are depicted in Figure 1(a). The bilateral satisficing set coincides with the Pareto-optimal solution, (S, S), when the vulnerability index is at least as large as $\frac{1}{2}$. It coincides with the Nash solution when the vulnerability index is less than the dissociation index. If the vulnerability index is greater than dissociation index but less than $\frac{1}{2}$, then the bilateral satisficing set contains both (S, S) and (C, C). To take action in this situation requires the invocation of a tie-breaker. For example, (C, C) is the satisficing option placing higher emphasis on individual interest (higher selectability), and (S, S) is the satisficing option placing higher emphasis on group interest (lower rejectability).

We may compute the unilateral satisficing equilibrium set by computing the univariate selectability and rejectability marginals in accordance with Eqs [16] and [17], from which the univariate satisficing set for either agent is

$$\mathfrak{S}_b = \begin{cases} \{S\} & \text{for } \beta > \frac{1+\alpha}{3} \\ \{C\} & \text{for } \beta < \frac{1+\alpha}{3} \\ \{S, C\} & \text{for } \beta = \frac{1+\alpha}{3} \end{cases} . \tag{27}$$

The unilateral decision regions are illustrated in Figure 1(b). Note that the set is a singleton except in the special situation of $\beta = \frac{1+\alpha}{3}$.

In contrast to the VNM solution to this game, the satisficing solution accounts for the inclinations of the players. The VNM solution emerges as a special case (e.g., $\alpha = 1$ and $\beta = 0$), but the satisficing solution gives the solution for all admissible (α, β) pairs. Uncertainty regarding these parameters may be handled in several ways. An agent is free to regard them as (a) random variables with known or interval-valued distributions and compute expectations, (b) deterministic interval-valued parameters and perform worst-case analysis, or (c) unknown parameters for which only an ordinal relationship is assumed. Thus, under fairly general circumstances, a decision can be rendered even though there may be considerable uncertainty regarding the values of the dispositional indices.

6. Summary and discussion

6.1. Summary

Current decision-making methodologies have two extreme categories: (1) superlative approaches, such as VNM game theory, that rely on expected utility maximization; and (2) positive, or heuristic approaches, such as searching, planning, and scheduling procedures, that rely on the knowledge and experience of recognized experts. Much current work is devoted to finding hybrid approaches, such as bounded rationality, which blend optimality and heuristics by modifying the performance or stopping criteria. In this paper we do not rely upon optimality, or heuristics, or even a hybrid. Instead, we invoke a comparative paradigm and establish a concept for satisficing based on the logical middle ground between the superlative and positive paradigms. Using this concept, we define satisficing equilibria and establish it as a basis for making decisions that can be systematically justified as "good enough." Table 4 provides a summary of the two extreme paradigms and ours, and compares them in terms of the solution concept, and knowledge requirements.

Substantive rationality—seeking the best and only the best—is perhaps the strongest possible notion of rationality, but it is more important to be good enough

Table 4. Frameworks for decision making

Rationality concept	Decision paradigm	Solution concept	Knowledge requirements
Substantive (extrinsic) rationality	Superlative degree (optimal)	Maximal expectations	Global mathematical models
Intrinsic rationality	Comparative degree (dichotomous)	Acceptable tradeoffs	Local mathematical models
Procedural rationality	Positive degree (heuristic)	Authoritative procedures	Local behavioral rules

than to be best. Intrinsic rationality—getting what you pay for—is a weaker notion, but it is also more fundamental. Stated another way, evaluating dichotomies is a more elementary activity than searching for maxima. It is local, rather than global; it is an internal, rather than external, evaluation.

The primary mechanism for identifying intrinsically rational options is epistemic utility theory. This theory is a cognitive decision-making procedure designed to acquire information and avoid error. Using the same mathematical structure, we provide new semantics for the design of a practical decision-making procedure designed to conserve resources and avoid failure. To distinguish this practical context from the cognitive context, we term this approach praxeic utility theory. This approach applies to both univariate (single-agent) and multivariate (multi-agent) decision making. It provides a compatible interface with the notion of satisficing we use, since it is designed to accommodate dichotomies. Furthermore, the mathematical structure of the interdependence function (as a mass function) permits the characterization of global interdependence in terms of conditional local interdependence, and provides a viable means of constructing the interdependence function to account for whatever dependencies exist among the agents.

In complex and uncertain environments, it is essential that decision-making procedures be parsimonious and robust. The middle ground offered by the comparative paradigm accommodates both of these desiderata in a natural way. By softening the demands of strict optimality, the door is opened for a flexible and economical way of characterizing multi-agent behavior while not abandoning demands for acceptable performance. The comparative paradigm uses whatever modeling information and performance principles are available to formulate the interdependence function and apply the praxeic likelihood ratio test. The resulting satisficing solutions are justified as being good enough in the sense that they result in tradeoffs that favor benefits over costs. This line of research is still in its early stages of development, but substantive results are beginning to emerge. In the single-agent context, [5, 6, 29, 32] demonstrate its applicability to the control of nonlinear dynamic systems that have proven to be difficult to solve using conventional means. In (Stirling and Goodrich, 1999a; Stirling et al., 1996b; Stirling and Goodrich, 1999b) the approach is applied to the multi-agent case, resulting in the formalization of the notion of satisficing games.

6.2. Some questions and reflections

The approach developed in this paper provides a contrast to conventional game theory, which tells us about outcomes we can expect from substantively rational agents. Game theory has been used extensively as a means of modeling human behavior, but there is considerable evidence that people often do not behave in ways consistent with substantive rationality; that is, they are not optimizers, or even constrained optimizers [2, 3, 4, 21, 28]. An important open question is whether or not, and under what conditions, our notion of satisficing based on intrinsic rationality provides a valid model for human behavior; that is, are people satisficers as we have defined the term? An answer to this question may be provided by appropriately designed psychological testing and evaluation.

Another characteristic of conventional game theory is that it employs rationality postulates that are imposed at the individual level, rather than on the group. Game theory does not easily accommodate group interests, since the preferences of each agent are expressed as functions only of the choices of other agents, and are not conditioned on their preferences. With conventional game theory, attitudes such as cooperation and conflict are not built into the utility functions, but become evident only when the utilities are juxtaposed—the linkages are external. A characteristic of our approach, however, is that preference relationships between agents can be expressed via the interdependence function—the linkages are internal. This feature invites further investigation into its significance. Does the explicit linkage of inter-agent preferences provide a basis upon which to construct an artificial society that captures important aspects of human behavior?

Satisficing decision theory may provide a convenient framework for negotiation. One of the problems with VNM game theory as a framework for negotiation is that it is not constructive—it may identify a best compromise, but does not provide a procedure for reaching it. The main trouble is dealing with the dynamic nature of coalition formation. Consequently, the strategic form is used extensively, and much of classical game-theoretic attention has been focused on situations where the process of negotiation can be presumed irrelevant to actual play. In other words, all of the deals, promises, threats, etc., are presumed to take place before the first move is actually played [24]. Heuristic approaches to negotiation are naturally amenable to the development of processes, but they lack the capacity for self-policing, and quality cannot be assured. The principal run-time "decision-making" activity under substantive rationality is searching for an option that possesses the externally defined attribute of "optimality," and under procedural rationality it is rule-following. But under intrinsic rationality, the principal run-time activity is evaluating dichotomies and actually making choices dynamically and interactively according to internal assessments of both group and individual preferences. This feature is potentially a great advantage when designing negotiatory processes. When negotiating, is seeking a good enough compromise a more robust and flexible posture than directly seeking a best compromise?

No realistic decision problem can account for all logically possible options. All decision problems are framed against a background of knowledge and assumptions that result in a subset of options that are deemed relevant by the decision maker. This set may or may not be adequate for the task at hand, and one of the most difficult of all decision-theoretic issues is to decide whether or not this set of options should be enlarged, and if so, how to go about expanding it. Rank-ordering-based techniques, by their very nature, provide only relative information and cannot be used to address this concern. Dichotomy-based techniques, such as praxeic utility theory, may stand a better chance of addressing this issue, since they are grounded in the fundamental properties of the options and permit self policing.

For example, being unable to find a good enough option in a situation may lead an agent to reconsider what it is willing or unwilling to do. If, for another example, there are no options for which the benefit-to-cost ratio provides a clear choice, this is evidence that the decision problem is a tense one for the agent. This is not to say that the agent cannot make a good decision. Rather, it is merely evidence that it may not be well-suited, at present, for the task. This realization may trigger the expansion of the set of options. In practical situations this may require activating additional sensors, applying more computational power, interrogating information sources, or other means of acquiring additional information, perhaps at a cost, in an attempt to better equip the agent to deal with its environment. Once again the praxeic has influenced the epistemic. Our need to act has led our need to know.

Notes

1. It is interesting to note that the story that gives the Prisoner's Dilemma its name was invented post factum to conform, for pedagogical illustration, to the payoff matrix with the given structure, rather than the other way around [35]. This illustrates the mind-set of many game theorists: the actual "game" is the payoff matrix, not the story.
2. Traditional game theorists are quick to point out that Rapoport's solution is not optimal. In fact, however, Axelrod proves [1] that, for repeated-play games where future payoffs are important, there does not exist an optimal strategy that is independent of the strategies used by other players.
3. Praxeology: the branch of knowledge that deals with practical activity and human conduct; the science of efficient action [11].
4. Informational value, as used in this context, is distinct from the notion of "value of information" of conventional decision theory [20], which deals with the change in expected utility that obtains if uncertainty is reduced or eliminated from a decision problem.
5. Such tradeoffs are also made under the superlative paradigm, but in a different way. As every control engineer well knows, specifying the performance criterion and solving for the optimum is only the first step in a subjective exercise of controller design. Consider, for example, the elementary control problem of designing an optimal linear quadratic regulator, where the goal is to maintain the state of a linear dynamical system within acceptable deviations from a desired set point while keeping the cost of control at an acceptable level. The tradeoff between performance and cost is accomplished by adjusting the relative weights of these two desiderata. Each weighting ratio leads to a different optimal control policy, and the task for the designer is to tune the weights iteratively to achieve an acceptable balance. In reality, there is no universally optimal solution—the "optimal" solution technique is nothing more than a convenient and systematic design procedure to achieve a solution that can be defended, at the end of the day, as being "good enough.".
6. It is in this sense only that the superlative paradigm is involved in our usage of satisficing; the members of this set do not necessarily inherit any superlative attributes.
7. Not to be confused with a mixture of distributions, which is a convex combination of probability distributions.
8. Although other psychological factors, such as expected behavior, may contribute to the results of repeated play games, for simplicity we restrict attention to the notion of group interest.
9. This apparent triviality is a consequence of each agent having only two options, since rejecting one implies selecting the other, but the situation is not so trivial when there are more than two options.

References

1. R. Axelrod, *The Evolution of Cooperation*, Basic Books: New York, 1984.
2. M. Bazerman, "A critical look at the rationality of negotiator judgement," *Behavorial Science*, vol. 27, pp. 211–228, 1983.
3. M. H. Bazerman and M. A. Neale, "Negotiator rationality and negotiator cognition: The interactive roles of prescriptive and descriptive research," in P. H. Young (ed.), *Negotation Analysis*, Univ. of Michigan Press: Ann Arbor, MI, 1992, pp. 109–129.
4. G. Gigerenzer and P. M. Todd, *Simple Heuristics That Make Us Smart*, Oxford Univ Press: New York, 1999.

5. M. A. Goodrich, W. C. Stirling, and R. L. Frost, "A Theory of satisficing decision and control," *IEEE Trans. Systems, Man, Cybernet*, vol. 28, no. 6, pp. 763–779, 1998.
6. M. A. Goodrich, W. C. Stirling, and R. L. Frost, "Model predictive satisficing fuzzy logic control," *IEEE Trans. on Fuzzy Systems*, vol. 7, no. 3, pp. 319–332, 1999.
7. J. Harsanyi, *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*, Cambridge Univ. Press: Cambridge, 1977.
8. R. M. Hogarth and M. W. Reder (eds.), *Rational Choice*, Univ. Chicago Press: Chicago, 1986.
9. L. Hogg and N. R. Jennings, "Variable sociability in agent-based decision making," in S. Parsons and M. J. Wooldridge (eds.), *Workshop on Decision Theoretic and Game Theoretic Agents*, University College: London, United Kingdom, 5 July, 1999, pp. 29–42.
10. B. E. Kaufman, "A new theory of satisficing," *The Journal of Behavorial Economics*, vol. 19, no. 1, pp. 35–51, 1990.
11. T. Kotarbiński, *Praxiology: An Introduction to the Sciences of Efficient Action*, Pergamon Press: Oxford, 1965. Translated by Olgierd Wojtasiewicz.
12. D. M. Kreps, *Game Theory and Economic Modelling*, Clarendon Press: Oxford, 1990.
13. I. Levi, *The Enterprise of Knowledge*, MIT Press: Cambridge, MA, 1980.
14. I. Levi, *The Covenant of Reason*, Cambridge Univ. Press: Cambridge, 1997.
15. R. D. Luce and H. Raiffa, *Games and Decisions*, John Wiley: New York, 1957.
16. F. R. Palmer, *Grammar*, Harmondsworth, Penguin: Harmondsworth, Middlesex, 1971.
17. D. W. Pearce, *Cost-Benefit Analysis*, MacMillan: London, second edition, 1983.
18. J. Pearl, *Probabilistic Reasoning in Intelligent Systems*, Morgan Kaufmann: San Mateo, CA, 1988.
19. K. R. Popper, *Conjectures and Refutations: The Growth of Scientific Knowledge*, Harper & Row: New York, 1963.
20. H. Raiffa, *Decision Analysis*, Addison-Wesley: Reading, MA, 1968.
21. A. Rapoport and C. Orwant, "Experimental games: A review," *Behavorial Science*, vol. 7, pp. 1–36, 1962.
22. E. Rasmusen, *Games and Information*, Basil Blackwell: Oxford, 1989.
23. T. Sandholm and V. Lesser, "Coalitions among computationally bounded agents," *Artificial Intelligence*, vol. 94, no. 1, pp. 99–137, 1997.
24. M. Shubik, *Game Theory in the Social Sciences*, MIT Press: Cambridge, MA, 1982.
25. H. A. Simon, "A behavioral model of rational choice," *Quart. J. Econ.*, vol. 59, pp. 99–118, 1955.
26. H. A. Simon, "Rational choice and the structure of the environment," *Psychological Review*, vol. 63, no. 2, pp. 129–138, 1956.
27. H. A. Simon, "Invariants of human behavior," *Annu. Rev. Psychol.*, vol. 41, pp. 1–19, 1990.
28. M. Slote, *Beyond Optimizing*, Harvard Univ. Press: Cambridge, MA, 1989.
29. W. C. Stirling, "Coordinated intelligent control via epistemic utility theory," *IEEE Control Systems Magazine*, vol. 13, no. 5, pp. 21–29, 1993.
30. W. C. Stirling and M. A. Goodrich, "Satisficing equilibria: A non-classical approach to games and decisions," in S. Parsons and M. J. Wooldridge (eds.), *Workshop on Decision Theoretic and Game Theoretic Agents*, University College, London, United Kingdom, 5 July, 1999a, pp. 56–70.
31. W. C. Stirling and M. A. Goodrich, "Satisficing games," *Information Sciences*, vol. 114, pp. 255–280, 1999b.
32. W. C. Stirling, M. A. Goodrich, and R. L. Frost, "Procedurally rational decision-making and control," *IEEE Control Systems Magazine*, vol. 16, no. 5, pp. 66–75, 1996a.
33. W. C. Stirling, M. A. Goodrich, and R. L. Frost, "Toward a theoretical foundation for multi-agent coordinated decisons," in *Proc. Second Int. Conf. Multi-Agent Systems*, Kyoto, Japan, 1996b, pp. 345–352.
34. W. C. Stirling and D. R. Morrell, "Convex bayes decision theory," *IEEE Trans. Systems, Man, Cybernet*, vol. 21, no. 1, pp. 173–183, 1991.
35. P. Straffin, "The Prisoner's Dilemma," *UMAP Journal*, vol. 1, pp. 101–103, 1980.
36. G. Wolf and M. Shubik, "Concepts, theories and techniques: Solution concepts and psychological motivations in Prisoner's Dilemma games," *Decision Sciences*, vol. 5, pp. 153–163, 1974.
37. S. Zilberstein, "Satisficing and bounded optimality," in *Proc. 1998 AAAI Symposium*, 1998, pp. 91–94. March 23–25, Stanford California. Technical Report SS-98-05.