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

Wynn Stirling[1] and Michael Goodrich[2]

[1] Electrical and Computer Engineering Department
Brigham Young University, Provo, UT, 84602, USA
wynn@ee.byu.edu

[2] Computer Science Department
Brigham Young University, Provo, UT, 84602, USA
mike@cs.byu.edu

Abstract. Satisficing, or being "good enough," is an alternative decision-making philosophy when optimal decisions are infeasible due to excessive computational difficulty or modeling complexity. In a departure from the traditional notion of satisficing as an externally specified species of bounded rationality, satisficing is redefined in terms of an agent-centered notion of ***comparative 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 praxeological interagent relationships are modeled by an interdependence function, from which multipartite selectability and rejectability mass functions may be derived and compared via the ***praxeic likelihood ratio test***.

1 Introduction

Multi-agent decision problems are often formulated as mathematical games. One of the more serviceable concepts within this framework is the *principle of substantive rationality:* a decision maker should do the best thing possible, given the available knowledge. This principle is usually instantiated by maximizing expected utility [4]. Equilibrium obtains when each player maximizes expected utility conditioned on the actions of other players, that is, when all agents have selected actions such that no agent can improve performance (for itself, in the case of Nash equilibria, or for the group, in the case of Pareto equilibria) by unilaterally changing. Such solutions are deemed *optimal.*

For complex problems, however, there often exist information, memory, or computing limitations such that obtaining a strictly optimal solution is not feasible. Under these circumstances, a principle of *bounded rationality* is often recommended. Theories of bounded rationality can be generated by relaxing one or more assumptions of the standard optimization approach. One such theory is Simon's well-known *satisficing* idea of searching for an optimal solution until the cost (in terms of computational resources, time requirements, cognitive capacity, etc.) of continuing to search exceeds the expected benefit of doing so [11, 12]. The best solution available when searching stops is implemented, although it is not necessarily globally best. Other notions of bounded rationality, such as augmenting the utility function with computational costs, are closely related to constrained optimization (see, for example, [10]), and yield optimal solutions according to a modified criterion. Regardless of the details of how a boundedly rational decision is obtained, however, the ultimate rationale for adopting it is that it represents an externally imposed 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 is certainly good enough, the issue becomes problematic under circumstances of bounded rationality. A characteristic of boundedly rational approaches is that the modification to strict optimality is *externally* and perhaps arbitrarily imposed by the agent designer according to criteria (such as computational costs) that are independent of actual performance, rather than *internally* determined by the agent itself in terms only of performance. It is possible, even for single-agent decision problems, that modifying one or more assumptions of an optimality-based solution methodology may lead to solutions that are <u>not</u> good enough from an agent-centered perspective. This issue becomes even more delicate for multi-agent systems. Unless a well-defined notion of agent-centered "good enough" equilibria is established, it will be difficult to establish confidence in other-than-optimal multi-agent decisions.

2 Motivation for a New Paradigm

A solution is best if it is superior to all other options according to some performance criterion, but a solution is good enough if the consequences it produces satisfy minimum performance requirements. When clearly defined and commonly accepted criteria are available, being "best" is a statement of objective factuality, while being "good enough" may be a subjective value judgment. Under bounded rationality, these two notions are often conflated, in the sense that the search for good enough solutions is conducted by modifying the way the search for a best solution is conducted. Furthermore, maximizing performance requires a global (perhaps exhaustive) search, but a local, non-exhaustive search may suffice to achieve minimum requirements. (It is not even necessary to consider maximum performance to define minimum requirements.) Since the notions of optimality and adequacy are logically independent, there is no reason to require that these two decision principles be implemented by the same mechanism. Instead, what is required is an independent and defensible agent-centered notion of what it means to meet minimum requirements.

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 on the idea of globally rank-ordering options according to performance criteria. On the other hand, we must not retreat into *ad hocism*, that is, decision making based on vague notions of desirability or convenience without any definitive measures of quality. *Ad hoc* procedures, while perhaps producing good decisions, will not produce decisions that can be systematically established as being adequate in terms of performance.

In matters of grammar, there is a logical gap between the *superlative* paradigm of being “best” and the *positive* paradigm of being “good.” This gap is filled by the *comparative* paradigm of being “better.” The optimality-based literature, particularly that of statistical decision theory, optimal control theory, and game theory, is overwhelmingly vast, reflecting many decades of serious research and development of ideas based on the superlative paradigm. The positive paradigm, manifest in the form of heuristics, procedurally rational decision making methods, and multitudinous *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.

There are at least two motives for considering a comparative paradigm. First, solution techniques based on the superlative paradigm concentrate on searching for the best, and only the best solution. Such techniques typically are “hill-climbing” procedures that iteratively follow the gradient of the performance measure until it approaches zero. Until the gradient actually approaches zero, however, no decision can be rendered, and there is no way to determine how close any intermediate results are to being optimal. Ho invites the metaphor of “skiing downhill in a fog,” to illustrate this situation [2]. The solution technique either yields the best solution or it yields no solution; no alternatives are available for consideration under the superlative paradigm. If, however, a solution technique based on a comparative paradigm could be designed that were capable of identifying justifiably good enough solutions with a reasonable amount of effort, it would be a potentially attractive alternative if the effort required to achieve an optimal solution were exorbitant.

A second reason for considering a comparative paradigm has to deal with the way solutions are justified. When optimizing, the critical task is to decide on the performance criterion. Once this is done, all that remains is to implement an appropriate search procedure to find a solution. Although the details of how the search is conducted may be a significant computational issue, they should not affect the final result. This is a characteristic of the superlative paradigm; the end justifies the means. Decisions made under the positive paradigm, however, are at the other extreme—attention is centered on validating the procedures that are to be followed, with the solution emerging automatically as the consequence; the means justify the end. Particularly with problems for which the models are complex or uncertain but performance issues are still important, it would be desirable to have a methodology that captures some essential features of these two extreme views of decision making. It would use the model, with all of its imperfections, in conjunction with the principles that characterize performance to guide toward, rather than dictate, a solution that can be justified in terms of both its procedural integrity and its qualitative performance.

3 Satisficing Equilibria

As a first step toward the development of a comparative paradigm, we formalize a notion of comparative rationality. The basic idea 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. So doing provides a convenient way of evaluating both the adequacy of and the penalties associated with options under consideration, and conforms with common praxeological[1] behavior. We will say that a decision is *comparatively 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. As a formalized means of decision making, this idea 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 the formal practice of 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."[7]. Cost-benefit analysis has proven to be 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 [6] 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 [1, 13, 14]), 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 added to 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, its informational value[2], that is, the degree to which the proposition supplies the kind of information that is demanded by the question. Thus, for a proposition to be admissible, it must be both true and informative—all implausible or uninformative propositions 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, by Levi) times another probability function, termed the informational-value-of-rejection probability. The set of propositions that maxi-

[1] Praxeology: the branch of knowledge that deals with practical activity and human conduct; the science of efficient action[5].

[2] Informational value, as used in this context, is distinct from the notion of "value of information" of conventional decision theory, which deals with the change in expected utility that obtains if uncertainty is reduced or eliminated from a decision problem.

mizes 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.

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 propositions; 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).

We define a *satisficing option as any option that is comparatively rational.* We retain the "satisficing" terminology under this new paradigm because we are motivated by exactly the same issue—to find options that are good enough. We differ from conventional usage only in the way in which good enough options are defined. Under the conventional notion, satisficing applies an externally supplied form of cost-benefit analysis directly to the stopping criterion, while, under the comparative paradigm, satisficing applies an agent-centered form of cost-benefit analysis directly to the options.

All rational decision making is subject to value judgments, and utility theory is a convenient and proven way to express preferences based on such judgments. 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 proposition. Thus, instead of making one global decision with respect to the entire collection of propositions, the comparative paradigm requires a separate local decision to be made for each proposition. Each proposition is 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, without the need to invoke any notions of optimality. We will say that a proposition is in a state of *satisficing equilibrium* if

S-1 The benefit derived from adopting it at least compensates for the cost incurred.

S-2 No other proposition 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, particularly under circumstances of uncertainty and when functionality is a more pressing issue than strict notions of optimality. Condition S-2 applies the domination principle to the cost-benefit framework and thereby eliminates propositions that needlessly either sacrifice performance or incur expense. These two conditions provide an agent-centered operational definition of being "good enough," or meeting minimum requirements.

Under conditions of time criticality or computational limitations, it may not be possible to establish condition S-2, since that may require testing of all options. condition S-1, however, does not require a global search. Thus, if the decision maker were forced to relax condition S-2 due to time or computational limitations, then it could simply implement the first satisficing solution found via a local search with full confidence that it will qualify, at least weakly, as good enough. Alternatively, in the spirit of the ordinal optimization ideas proposed by Ho [3], the decision maker may stop searching when sufficiently many satisficing options have been found to ensure that, with high probability, a satisficing equilibrium solution has been identified.

4 Praxeic Utility

Our notion of satisficing requires us to obtain decision theoretic 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 epistemic utility theory, originally developed for cognitive 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 3.

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* ... " [9, Page 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.) 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 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; that is, if $u \in U$ is rejected, then the agent conserves $p_R(u)$ worth of its unit of resources. 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 in a decision problem if successful options are not rejected. Therefore, we wish to consider sets of options that are likely to be successful, rather than to concentrate exclusively on finding a singleton best option. We will assume that the agent has a unit of success utility. Suppose that implementing $u \in U$ would lead to success. For any subset $A \subset U$, the utility of not rejecting A in the interest of avoiding failure is the indicator function, $I_A(u) = \begin{cases} 1 & \text{if } u \in A \\ 0 & \text{otherwise} \end{cases}$. This utility function is very even-handed; it is concerned only with the value of not rejecting A if u is successful. It does not count one option to be more desirable (if successful) 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:

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

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.

Generally, the decision maker will not know precisely which u will lead to success, and thus cannot maximize praxeic utility. In the tradition of von Neumann-Morgenstern theory, however, the decision maker may maximize 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 $u \in U$. Let $p_S \colon U \to [0, 1]$ denote the normalized degree of support that the elements of U lend to achieving success. The expected praxeic utility is then

$$\begin{aligned} \bar{\phi}(A) &= \sum_{u \in U} \left[\alpha I_A(u) + (1 - \alpha)\left(1 - \sum_{v \in A} p_R(v)\right)\right] p_S(u) \\ &= \alpha \sum_{v \in A} p_S(v) - (1 - \alpha) \sum_{v \in A} p_R(v) + (1 - \alpha). \end{aligned} \tag{2}$$

Since von Neumann-Morgenstern 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 the *expected praxeic utility:*

$$\bar{\varphi}(A) = \sum_{u \in A} \left[p_S(u) - b p_R(u) \right], \tag{3}$$

where $b = \frac{1-\alpha}{\alpha}$. p_S satisfies the definition of a probability mass function, and evaluates each option with respect to the degree to which it accomplishes the objective, 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.*

We may obtain the largest set of options for which the selectability is greater than or equal to b times the rejectability by maximizing expected praxeic utility[3], resulting in the *satisficing set*

$$\Sigma_b = \arg\max_{A \subset U} \bar{\varphi}(A) = \{u \in U \colon p_S(u) \geq b p_R(u)\} = \left\{ u \in U \colon \frac{p_S(u)}{p_R(u)} \geq b \right\}. \tag{4}$$

Thus, 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 $\Sigma_b \neq \emptyset$, since otherwise $p_S(u) < b p_R(u)$ 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 (4) 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

$$\begin{aligned} B_S(u) &= \{v \in U \colon p_R(v) < p_R(u) \text{ and } p_S(v) \geq p_S(u)\} \\ B_R(u) &= \{v \in U \colon p_R(v) \leq p_R(u) \text{ and } p_S(v) > p_S(u)\}, \end{aligned} \tag{5}$$

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

$$B(u) = B_S(u) \cup B_R(u); \tag{6}$$

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 $B(u) = \emptyset$, 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

$$\mathcal{E} = \{u \in U \colon B(u) = \emptyset\}. \tag{7}$$

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

$$\mathfrak{S}_b = \mathcal{E} \cap \Sigma_b. \tag{8}$$

[3] 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.

5 Multi-Agent Systems

An act by any single member X_i of a multi-agent system, $\{X_1, \ldots, X_N\}$, has possible ramifications throughout the entire community. Some agents may be benefited, some may be damaged, and some may be indifferent. Furthermore, although the single agent may perform the act in its own interest, 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.

Definition 1. *A* mixture[4] *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 $\mathcal{S} = S_{i_1} \ldots S_{i_k}$, is a mixture consisting of agents $X_{i_1}, \ldots X_{i_k}$ 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 $\mathbf{S} = S_1 \ldots S_N$. 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 $\mathcal{R} = R_{j_1} \ldots R_{j_\ell}$, is a mixture consisting of agents $X_{j_1}, \ldots X_{j_\ell}$ 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 $\mathbf{R} = R_1 \ldots R_N$. 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 $\mathcal{SR} = S_{i_1} \ldots S_{i_k} R_{j_1} \ldots R_{j_\ell}$. The* joint intermixture *is the concatenation of the joint selectability and joint rejectability mixtures, and is denoted $\mathbf{SR} = S_1 \ldots S_N R_1 \ldots R_N$. A* myopic *intermixture is a mixture of the form $\mathcal{SR} = S_i R_i$.* $\square$

Definition 2. *Given arbitrary action spaces U_i, $i = 1, \ldots, N$, the* product action space*, denoted $\mathbf{U} = U_1 \times \cdots \times U_N$ is the set of all N-tuples $\mathbf{u} = (u_1, \ldots, u_N)$ where $u_i \in U_i$. The* selectability action space *associated with a selectability mixture $\mathcal{S} = S_{i_1} \ldots S_{i_k}$ is the product space $\mathbf{U}_{\mathcal{S}} = U_{i_1} \times \cdots \times U_{i_k}$. The* rejectability action space *associated with a rejectability mixture $\mathcal{R} = R_{j_1} \ldots R_{j_\ell}$ is the product space $\mathbf{U}_{\mathcal{R}} = U_{j_1} \times \cdots \times U_{j_\ell}$. The* interaction space *associated*

[4] Not to be confused with a mixture of distributions, which is a convex combination of probability distributions.

with an intermixture $\mathcal{SR} = S_{i_1} \cdots S_{i_k} R_{j_1} \cdots R_{j_\ell}$ *is the product space* $\mathbf{U}_{\mathcal{SR}} = U_{i_1} \times \cdots \times U_{i_k} \times U_{j_1} \times \cdots \times U_{j_\ell}$*. The* joint interaction space *is* $\mathbf{U}_{\mathbf{SR}} = \mathbf{U} \times \mathbf{U}$. $\square$

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

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

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

Let $\mathbf{v}$ and $\mathbf{w}$ be two joint options in $\mathbf{U}_{\mathcal{S}} \times \mathbf{U}_{\mathcal{R}}$. Then $p_{\mathcal{SR}}(\mathbf{v}, \mathbf{w})$ 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, $p_{\mathcal{SR}}(\mathbf{v}, \mathbf{w})$ characterizes the disposition of the agent system with respect to 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* $\mathcal{SR} = S_{i_1} \ldots S_{i_k} R_{j_1} \ldots R_{j_\ell}$*, a* subintermixture *of* $\mathcal{SR}$ *is an intermixture formed by concatenating subsets of* $\mathcal{S}$ *and* $\mathcal{R}$*:* $\mathcal{S}_1\mathcal{R}_1 = S_{i_{p_1}} \ldots S_{i_{p_q}} R_{j_{r_1}} \ldots R_{j_{r_s}}$*, where* $\{i_{p_1}, \ldots, i_{p_q}\} \subset \{i_1, \ldots, i_k\}$ *and* $\{j_{r_1}, \ldots, j_{r_s}\} \subset \{j_1, \ldots, j_\ell\}$*. We shall use the notation* $\mathcal{S}_1\mathcal{R}_1 \subset \mathcal{SR}$ *to indicate that* $\mathcal{S}_1\mathcal{R}_1$ *is a subintermixture of* $\mathcal{SR}$*.*

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

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

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

$$p_{\mathcal{SR}}(\mathbf{v}, \mathbf{w}) = p_{\mathcal{SR}\backslash\mathcal{S}_1\mathcal{R}_1|\mathcal{S}_1\mathcal{R}_1}(\mathbf{w}|\mathbf{v}) \cdot p_{\mathcal{S}_1\mathcal{R}_1}(\mathbf{v}). \tag{9}$$

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

$$p_{S_1 S_2 S_3 R_1 R_2 R_3}(v_1, v_2, v_3, w_1, w_2, w_3) = \\ p_{S_3 R_1 R_2 | S_1 S_2 R_3}(v_3, w_1, w_2 | v_1, v_2, w_3) \cdot p_{S_1 | S_2 R_3}(v_1 | v_2, w_3) \cdot p_{S_2 R_3}(v_2, w_3). \quad (10)$$

$p_{S_3 R_1 R_2 | S_1 S_2 R_3}(v_3, w_1, w_2 | v_1, v_2, w_3)$ is the conditional selectability/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 behavior conditioned on various situations makes it possible to characterize global behavior or preferences in terms of conditional local behavior or preferences. As noted by Pearl [8], it is often easier to specify conditional local behavior rather than unconditional joint global behavior. Conditional probabilities permit local, or specific responses to be characterized; they possess modularity features similar to logical production rules.

The role of the joint imf is to characterize all of the logically possible relationships that exist between the members of the multi-agent system. Once this function is defined, the joint selectability and rejectability marginals are obtained as

$$p_{\mathbf{S}}(\mathbf{u}) = \sum_{\mathbf{v} \in \mathbf{U}} p_{\mathbf{SR}}(\mathbf{u}, \mathbf{v}) \quad (11)$$

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

for all $(\mathbf{u}, \mathbf{v}) \in \mathbf{U} \times \mathbf{U}$. The joint selectability and rejectability mass functions may then be used to generate a multipartite satisficing set according to a straightforward extension of the PLRT:

$$\boldsymbol{\Sigma}_b = \{\mathbf{u} \in \mathbf{U}\colon p_{\mathbf{S}}(\mathbf{u}) \geq b p_{\mathbf{R}}(\mathbf{u})\}. \quad (13)$$

The *multipartite satisficing equilibrium set* is

$$\mathfrak{S}_b = \boldsymbol{\mathcal{E}} \cap \boldsymbol{\Sigma}_b, \quad (14)$$

where $\boldsymbol{\mathcal{E}}$ is the set of multipartite satisficing equilibrium options defined in accordance with (7). This satisficing equilibrium property is distinct from both Nash equilibria and Pareto optimality. A joint action is a Nash equilibrium if, were any single agent to change its decision, it would reduce its level of satisfaction. A joint action is Pareto optimal if no single agent can increase its level of satisfaction without lowering the satisfaction level of at least one other agent. Multipartite satisficing equilibrium options, however, possess the property that, if any agents change their decisions, then either the joint selectability decreases or the joint rejectability increases (or both).

Nash equilibrium solutions are justified via individual optimality, and Pareto optimal solutions are justified via group optimality. Since the satisficing concept does not depend upon a superlative paradigm, it is not tied to any particular notion of optimality. Rather, satisficing equilibrium solutions are justified via comparative rationality.

6 Illustrative Example

The game of Prisoner's Dilemma 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 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 in return for dropping the charges. Thus, each prisoner has two options: either to confess (C), or to remain silent (S). If neither prisoner confesses, they both will receive short jail sentences. If only one confesses, the one who remains silent will receive the maximum sentence for a serious crime, with the one who confesses being released from jail. In the event that both prisoners confess, each will receive a prison term of intermediate length. The conventional solution to this game is easy to obtain, but is terribly unsatisfactory. It is well-known that playing C for either player is the dominant strategy. Furthermore, (C, C) is also the unique Nash equilibrium pair. Unfortunately, this solution is inferior to playing the "cooperative," Pareto-optimal solution, (S, S).

Under the superlative paradigm, the mechanism for ordering the options is Nash equilibria—it dictates the joint solution that is in the best interest of each player. Under the comparative paradigm, however, options are related through 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 [15]. 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[5]. Satisficing game theory provides a mechanism to implement multiple attribute praxeological considerations, in contrast to the single-attribute approach of conventional utility theory. Although there is not a unique way to frame this problem, we proceed by

[5] 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.

associating success with reducing individual jail time and 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 acknowledged dependencies reinforce each other, the players implicitly forge a joint opinion regarding the relative merits of cooperation and conflict. There appear to be two aspects of any such relationship: (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 <u>and</u> high vulnerability would indicate an 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 payoffs are symmetric (that is, 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 may define the joint rmf:

$$\begin{aligned} p_{R_1R_2}(S, S) &= \alpha, & p_{R_1R_2}(S, C) &= \frac{1-\alpha-\beta}{2} \\ p_{R_1R_2}(C, S) &= \frac{1-\alpha-\beta}{2} & p_{R_1R_2}(C, C) &= \beta. \end{aligned} \tag{15}$$

Although selectability deals with individual objectives, it is a joint consideration, since the consequences of the agents' decisions are not 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:

$$p_{S_1S_2R_1R_2}(v_1, v_2, w_1, w_2) = p_{S_1S_2|R_1R_2}(v_1, v_2|w_1, w_2) \cdot p_{R_1R_2}(w_1, w_2). \tag{16}$$

Let, say, $(w_1, w_2) = (S, S)$. If X_2 were then to ascribe any conditional selectability mass to S, it would suffer the maximum penalty if X_1 were to choose C; consequently, it would be in X_2's interest to ascribe its entire unit of conditional selectability to C. Assuming that X_2 follows this course, X_1 must do likewise, otherwise it suffers the maximum penalty. Thus, $p_{S_1S_2|R_1R_2}(C,C|S,S) = 1$. By a similar argument for $(w_1,w_2) = (S,C)$ and for $(w_1,w_2) = (C,S)$, we may also conclude that $p_{S_1S_2|R_1R_2}(C,C|S,C) = 1$ and that $p_{S_1S_2|R_1R_2}(C,C|C,S) = 1$.

Now suppose the agents were to place all of their joint rejectability mass on (C,C); that is, $\beta = 1$, so both agents are committed to the notion that what is jointly desirable is also individually desirable. This condition is equivalent to the probabilistic notion of total dependence, meaning that the joint conditional smf is degenerate and the two-dimensional distribution collapses into a one-dimensional distribution:

$$p_{S_1S_2|R_1R_2}(v_1,v_2|C,C) = \begin{cases} 0 & v_1 \neq v_2 \\ p_{S_1|R_1R_2}(v_1|C,C) & v_1 = v_2 \end{cases}, \tag{17}$$

so whatever X_1 chooses, X_2 makes the same choice. Clearly, the option that minimizes joint time in jail is for X_1 (and, hence, X_2) to choose S. This fact can be expressed by setting the conditional selectability function, evaluated at S, to unity: $p_{S_1|R_1R_2}(S|C,C) = 1$. Hence, by (17), $p_{S_1S_2|R_1R_2}(S,S|C,C) = 1$.

Substituting the conditional selectability and the joint rejectability given by into (16) and applying (11) yields the joint selectability function:

$$\begin{array}{ll} p_{S_1S_2}(S,S) = \beta & p_{S_1S_2}(S,C) = 0 \\ p_{S_1S_2}(C,S) = 0 & p_{S_1S_2}(C,C) = 1 - \beta. \end{array} \tag{18}$$

Comparing the joint selectability, (18), with the joint rejectability, (15), we obtain the jointly satisficing set, which consists of all decision pairs, $(w_1,w_2) \in \mathbf{U}$, such that $p_{S_1S_2}(w_1,w_2) \geq b p_{R_1R_2}(w_1,w_2)$. Thus, parameterized by α and β, this set is

$$\mathfrak{S}_b = \begin{cases} \{(S,S)\} & \text{for } \beta \geq \frac{1}{1+b}, \beta > b\alpha \\ \{(C,C)\} & \text{for } \beta < \frac{1}{1+b}, \beta \leq b\alpha \\ \{(S,S),(C,C)\} & \text{for } \beta < \frac{1}{1+b}, \beta > b\alpha \\ \emptyset & \text{for } \beta \geq \frac{1}{1+b}, \beta \leq b\alpha \end{cases}. \tag{19}$$

The jointly satisficing set coincides with the Pareto-optimal solution, (S,S), when the vulnerability index is at least as large as $\frac{1}{1+b}$ and the vulnerability index exceeds the product of rejectivity and dissociation. The jointly satisficing set coincides with the Nash solution when the vulnerability index is not only less than $\frac{1}{1+b}$, but is also less than the product of rejectivity and the dissociation index. If the vulnerability index is greater than rejectivity times the dissociation index and the vulnerability index is less than $\frac{1}{1+b}$, then the jointly 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).

7 Discussion

According to conventional usage, satisficing is an externally specified heuristic means of finding a "good enough" solution when it is not feasible to obtain an optimal one. Essentially, it is a blending of the superlative paradigm (searching for the optimal solution) and the positive paradigm (defining an *ad hoc* rule for stopping the search or modifying the performance criteria according to constraints that are not directly performance related). In this paper we invoke a comparative paradigm and establish an agent-centered 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 justified as being systematically "good enough."

Key properties of this resulting satisficing decision-making approach include: (a) the formalization and generalization of the simple notion of cost-benefit analysis; (b) a concern for both qualitative performance and procedural integrity in making other-than-optimal decisions and an agent-centered formalism to include both features; and (c) the extension of praxeic utility theory to multi-agent decision making.

References

1. M. A. Goodrich, W. C. Stirling, and R. L. Frost. A Theory of Satisficing Decision and Control. *IEEE Transactions on Systems, Man, and Cybernetics*, 28(6):763–779, November 1998.
2. Y. C. Ho. Heuristics, Rules of Thumb, and the 80/20 Proposition. IEEE *Transactions on Automatic Control*, 39(5):1025–1027, 1994.
3. Y. C. Ho. On the Numerical Solutions of Stochastic Optimization Problem. IEEE *Transactions on Automatic Control*, 42(5):727–729, 1997.
4. R. M. Hogarth and M. W. Reder, editors. *Rational Choice*. University of Chicago Press, 1986.
5. T. Kotarbiński. *Praxiology: An Introduction to the Sciences of Efficient Action*. Pergamon Press, 1965. Translated by Olgierd Wojtasiewicz.
6. I. Levi. *The Enterprise of Knowledge*. MIT Press, Cambridge, Massachusetts, 1980.
7. D. W. Pearce. *Cost-Benefit Analysis*. Macmillan, second edition, 1983.
8. J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Mateo, Ca, 1988.
9. K. R. Popper. *Conjectures and Refutations: The Growth of Scientific Knowledge*. Harper & Row, 1963.
10. T. Sandholm and V. Lesser. Coalitions among Computationally Bounded Agents. *Artificial Intelligence*, 94(1):99–137, 1997.
11. H. A. Simon. A Behavioral Model of Rational Choice. *Quart. J. Economics*, 59:99–118, 1955.
12. H. A. Simon. Rational Choice and the Structure of the Environment. *Psychological Review*, 63(2):129–138, 1956.
13. W. C. Stirling and M. A. Goodrich. Satisficing Games. *Information Sciences*, 114:255–280, March 1999.
14. W. C. Stirling and D. R. Morrell. Convex Bayes Decision Theory. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(1):173–183, January/February 1991.
15. G. Wolf and M. Shubik. Concepts, Theories and Techniques: Solution Concepts and Psychological Motivations in Prisoner's Dilemma Games. *Decision Sciences*, 5:153–163, 1974.