Satisficing Negotiation for Resource Allocation with Disputed Resources

Todd K. Moon*
Electrical and Computer Engineering Dept.
Utah State University, Logan, UT 84322-4120
voice: 435-797-2970 fax: 435-797-3054
email: Todd.Moon@ece.usu.edu

Wynn C. Stirling
Electrical and Computer Engineering Dept.
Brigham Young University, Provo, UT 84602
voice: 801-378-7669 fax:810-378-6586
email: wynn@ee.byu.edu

Abstract

Satisficing decision making provides a more malleable framework for negotiation than conventional techniques based on optimization of a utility function. In this paper we summarize satisficing decision theory, which provides a mechanism for determining decision options which are "good enough" as as tradeoff between a selectability function and a rejectability function, with an index of caution as a decision control parameter. Single agent satisficing is extended to multi-agent satisficing, by which group rationality can be represented; option vectors for the entire group are obtained as a result of this decision process. Multi-agent satisficing provides the stage upon which negotiation takes place. Negotiation, in this context, is the process by which all agents determine a set of options which are both individually satisficing and jointly satisficing; through the course of negotiation agents can accommodate the needs of the group — compromise — through lowering their index of caution. An example is presented of resource allocation with disputed resources.

Introduction

Decision making agents acting together should be influenced not only by their own aspirations and budgets but by these aspects of other agents in the system. To represent this interaction among agents, a notion of *group* rationality must be embodied in the decision systems of interacting agents. Group rationality is not necessarily a logical consequence of rationality based on individual self interest. Under a model of rationality in which maximization of utility is the operative notion, group behavior obtained by amalgamation of the individual behaviors is not usually optimized by optimizing each individual behavior, as is typically done in a game-theoretic setting. Those who put their final confidence in the limited perspective of exclusive self-interest may ultimately function disjunctively, and perhaps illogically, when participating in collective activities. Rather than reorient game theory to accommodate situations where coordination is a more natural operational descriptor of the game than is self-interested conflict, we propose to describe notions that are neutral with respect to questions of conflict and coordination.

*Partial support provided by DARPA under F30602-00-2-0532

Beyond simply taking into account the presence of other acting agents in a system, there is frequently some form of sociality that is conducive to at least a weak form of congruity or mutual agreement. In cooperative scenarios, agents agree to work together; in competitive scenarios, agents tacitly agree to oppose each other. The procedures used to arrive at these agreements are not determined simply as a function of the preference structure of the decision makers, whether posed in a framework of self interest or community interest. Agreement among agents is typically obtained via a process of negotiation, in which multiple agents evaluate and share information when they have incentive to strike a mutually acceptable compromise. In the negotiating process, it is not sufficient for a decision maker merely to identify an acceptable joint solution (for the community) according to its own lights. The entire community should "buy into" a joint solution that is mutually acceptable.

In negotiation it is rare that all parties involved will tip their hands to reveal all of the factors influencing their decisions. In a competitive setting, a policy of secrecy may keep competitors from exploiting a weakness, or it may be used to persuade competitors to a more advantageous position. In a cooperative setting, complete disclosure of information might be precluded due to restrictions in communication bandwidth and/or time. Because of a lack of disclosure, negotiation may invoke principles of inference, wherein agents attempt to estimate positions or attitudes of other agents based on the options they bring to the bargaining table.

In light of these observations, some principles of negotiation are suggested:

N-1 Negotiators must typically be concerned with meeting minimum requirements more than achieving maximum performance.

N-2 Negotiations should lead to decisions that are both good enough for the group as a whole (as established by a group rationality) as well as good enough for each individual (as established by local preferences).

N-3 Negotiation is typically an iterative process. Starting from a set of initial joint options, it is natural to iterate toward solutions which are individually acceptable, rather than attempting to move directly to joint options which are a best compromise.

N-4 Negotiation may frequently incorporate elements of inference.

A rich model for negotiation should be able to capture other aspects of the negotiation process, such as recalcitrance (resistance to accede to group preferences), accommodation (openness to group preferences), or annoyance over extended or unchanging negotiation positions.

In this paper, we will briefly review the concept of praxeic utility decision theory as a means of implementing satisficing control, then extend this to multiple agent decision making to model group rationality. Concepts of negotiation consistent with the principles outlined above are established using this multi-agent satisficing framework. As case study of a problem for which a negotiated solution is reasonable, a problem of resource allocation with disputed resources is modeled.

Satisficing decision making: single and multiple agents

Single agent satisficing

Satisficing, a term coined by Simon (Simon 1955), refers to a decision making strategy in which options are selected which are "good enough," differing thereby from conventional approaches which seek only the best. From the satisficing viewpoint, being "good enough" is sufficient; insisting on the best and only the best via an optimizing algorithm may be an overly restrictive luxury. From an operational point of view, however, while establishing that an option is (at least locally) optimal is at least expressible as a optimization problem, establishing what is "good enough" appears to be more elusive. The question of establishing good enough choices is addressed from a philosophical point of view with regard to truth systems by Levi (Levi 1984; 1980; 1967). In this epistemological framework, known as epistemic utility, options are sought for which the amount of information associated with them exceeds the potential for error. All options are deemed acceptable — good enough — that pass a likelihood ratio test comparing a truth valuation (a probability) and an informational value of rejection (also constructed as a probability). Application of epistemic utility to control problems yields *praxeic utility theory* (Stirling, Goodrich, & Packard 2001; Stirling, Goodrich, & Frost 1996c; 1996b; Stirling 1994; Goodrich, Stirling, & Boer 2000; Goodrich, Stirling, & Frost 1998; Stirling, Goodrich, & Frost 1996a; Stirling & Goodrich 1999). (For a discussion of praxeic utility theory in the context of negotiation, and for a more complete development of these concepts, see (Stirling & Moon 2001).) In this theory, a selectability function $p_S(u)$ is formed which, for each option u available from a universe of options U available to a decision making agent X, measures the degree to which u works toward success in achieving the agent's goals. Also, a rejectability function $p_R(u)$ is established which measures costs associated with each option. This pair of measures, called collectively the *satisfiability functions*, are endowed with the mathematical structure of probabilities (e.g., they are nonnegative and sum to 1 on the U).

Definition 1 The *satisficing set* Σ_q is the set of options defined by

$$\Sigma_q = \{u \in U : p_S(u) \geq q p_R(u)\}. \tag{1}$$

□

The satisficing set consists of those options for which the benefit exceeds the cost: the set of alternative which are arguably "good enough." There may be more than one option in Σ_q. Moving away from strict adherence to optimality increases the flexibility, while by not retaining only the best. Ultimate selection of a single option for action is accomplished by means of a tie breaking rule, such as most selectable, least rejectable, or maximally discriminating.

The parameter q in (1) is the *index of caution*. As q is increased, fewer options are accepted into the the satisficing set. As such, the agent exhibits greater caution, accepting only options of higher merit in comparison to their cost. We say that Σ_q is the satisficing set at level q. Because of its similarity to likelihood ratio tests in conventional decision theory, the test in (1) is referred to as the *praxeic likelihood ratio test* (PLRT).

Multiple agent satisficing

Satisficing decision theory extends very naturally to multiple agent systems. Satisficing admits degrees of fulfillment, whereas optimization is an absolute concept. While the statement "What is best for me and what is best for you is also jointly best for us together" may be nonsense, the statement "What is good enough for me and what is good enough for you is also jointly good enough for us together" may be perfectly sensible, especially when we do not have inflexible notions of what it means to be "good enough." Satisficing grants room for compromise, leaving open the opportunity for one or more agents involved to relax standards of individual performance in the interest of the good of the community. A theory of multi-agent satisficing thus provides the stage on which the act of negotiation can reasonably be presented.

Since they possess the mathematical structure of probabilities, selectability and rejectability can be naturally extended to the multivariate case by defining *joint* selectability and rejectability measures, which may be used to determine a jointly satisficing set. In addition, individual decision makers may establish individual notions of satisficing by computing marginal selectability and rejectability functions from the joint expressions.

Let $X_1, \ldots, X_N$ be N interacting agents, where each agent has its own decision space U_i. The joint action space is the space $\mathbf{U} = U_1 \times U_2 \times \cdots \times U_N$. A joint decision is an element $\mathbf{u} = (u_1, u_2, \ldots, u_N) \in \mathbf{U}$. We denote the ith element of $\mathbf{u}$ as $\mathbf{u}(i)$.

An act by any single member of a multi-agent system has potential ramifications throughout the entire community. And, although a participant may perform an act either in its own interest or for the benefit of others, the act is usually not implemented free of cost: resources are expended or risk is taken, perhaps by the single agent, but also perhaps by the entire community. Although these consequences

may be defined independently from the benefits, the measures associated with benefits and cost cannot necessarily be specified independently of each other. In light of this, the object representing the relationships between agents in their systems regarding their individual and joint selectability and rejectability is an *interdependence measure* which combines both rejectability and selectability for all agents as a multivariate probability function of the form

$$p_{S_1,S_2,\ldots,S_N,R_1,R_2,\ldots,R_N}(u_1,u_2,\ldots,u_N,v_1,v_2,\ldots,v_N).$$

This is expressed more briefly as $p_{\mathbf{S},\mathbf{R}}(\mathbf{u},\mathbf{v})$. Values for the interdependence measure are typically obtained by means of factorization into constituent conditional and marginal probabilities. In these factorizations, agents may represent how their selectabilities or rejectabilities are affected by the selectability and rejectability of other agents. From the general interdependence function, the joint selectability function is obtained by a marginalization

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

and similarly

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

Definition 2 The *multipartite satisficing decision rule* defines the set the *multipartite satisficing set* by

$$\mathbf{\Sigma}_q = \{\mathbf{u}\in\mathbf{U}\colon p_{\mathbf{S}}(\mathbf{u}) \geq q p_{\mathbf{R}}(\mathbf{u})\}. \tag{2}$$

□

Joint options in $\mathbf{\Sigma}_q$ are those for which the benefits exceed the costs, as viewed from the perspective of the group and as represented by the joint selectability and joint rejectability. The test in (2) is referred to as the *joint praxeic likelihood ratio test* (JPLRT).

Given joint selectability and joint rejectability, an individual agent can compute marginals by

$$p_{S_i}(u_i) = \sum_{\{\mathbf{u}\in U\colon \mathbf{u}(i)=u_i\}} p_{\mathbf{S}}(\mathbf{u})$$

$$p_{R_i}(v_i) = \sum_{\{\mathbf{v}\in U\colon \mathbf{v}(i)=v_i\}} p_{\mathbf{S}}(\mathbf{v}).$$

The resulting individually satisficing set for X_i is then

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

Alternatively, an agent could employ a p_{S_i} and p_{R_i} not obtained as a marginal, presenting a different face to the public than what it holds for itself, and use these function to compute Σ_q^i.

We will use the notation $u = \mathbf{u}(i)$ to indicate that the option $u \in U_i$ is the ith element of a joint option vector $\mathbf{u}$.

An option u that is jointly satisficing for X_i, is not necessarily individually satisficing for X_i. That is, a joint option $\mathbf{u}\in\mathbf{\Sigma}_q$ does not necessarily have $\mathbf{u}(i)\in\Sigma_q$. The converse, however, is true: if $u\in\Sigma_q^i$, then $u=\mathbf{u}(i)$ for some $\mathbf{u}\in\mathbf{\Sigma}_q$. This is established by the following.

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

Proof Without loss of generality, assume $i=1$. Let $u\in\Sigma_q^1$ (i.e., $p_{S_1}(u)\geq q p_{R_1}(u)$). To establish proof by contradiction, assume that $u\neq\mathbf{u}(i)$ for all $\mathbf{u}\in\mathbf{\Sigma}_q$. It follows that for all $\mathbf{v}\in U_2\times\cdots\times U_N$, $p_{\mathbf{S}}(u,\mathbf{v}) < q p_{\mathbf{R}}(u,\mathbf{v})$. Then

$$p_{S_1}(u) = \sum_{\mathbf{v}} p_{\mathbf{S}}(u,\mathbf{v}) < q\sum_{\mathbf{v}} p_{\mathbf{R}}(u,\mathbf{v}) = q p_{R_1}(u),$$

which contradicts $u\in\Sigma_q$. □

On the basis of the negotiation theorem, it may be argued that each agent has a seat at the negotiation table. No one is necessarily frozen out of a deal.

It is important to emphasize what the negotiation theorem does not provide. If u_1 is individually satisficing for X_1, and u_2 is individually satisficing for X_2, then by the theorem $u_1=\mathbf{u}(1)$ and $u_2=\tilde{\mathbf{u}}(2)$ for some $\mathbf{u},\tilde{\mathbf{u}}\in\mathbf{\Sigma}_q$, However, it is not necessarily the case that $\mathbf{u}=\tilde{\mathbf{u}}$: the options that are both individually and jointly satisficing may be different for different agents. Thus, the negotiation theorem does not establish a "solution" to the problem. However, it provides the basis upon which a solution may be sought through an iterative negotiation scheme. To obtain buy-in from all agents, the options that are individually satisficing for each agent must be elements of the same jointly satisficing options. Such options are the result of negotiation.

The proof of the negotiation theorem makes use of the fact that the same index of caution is used to compute Σ_q^i as is used for $\mathbf{\Sigma}_q$, but an agent could use a higher q to determine $\mathbf{\Sigma}_q$ than it does for its own satisficing set. The negotiation theorem is not necessarily true in this case. Or it may happen that each agent has its own perception of the interdependence function. In this case, we denote X_i's interdependence function as $p^i_{\mathbf{S},\mathbf{R}}$, and the corresponding joint selectability and rejectability as $p^i_{\mathbf{S}}$ and $p^i_{\mathbf{R}}$. Again, the negotiation theorem applies to each agent separately, but not collectively. Another possibility is that an agent may use joint functions $p^i_{\mathbf{S}}$ and $p^i_{\mathbf{R}}$ for establishing $\mathbf{\Sigma}_q^i$, but use individual p_{S_1} and p_{R_1} not computed as marginals of $p_{\mathbf{S}}$ and $p_{\mathbf{R}}$, thereby presenting a different "public face" and "private face." Again, the negotiation theorem does not apply. (This raises the the question for future investigation: to what degree can these private satisfiability functions differ from marginal satisfiability functions and still have reasonable negotiations.)

The negotiation theorem, and these observations about its provisional application, motivate the development of algorithms for negotiation based on satisficing.

Satisficing Negotiation

Multi-agent satisficing is suited to the principles of negotiation outlined in the introduction. By admitting degrees of fulfillment, satisficing agents can explore options which are both mutually and individually good enough. In a negotiation process, however, it is not sufficient for a decision

maker to identify a solution, even one which it views as being jointly acceptable. As mentioned above, other agents in the system may have their own models of the joint interdependence function, and their own individual satisficing functions, or may determine individually and jointly satisficing options not coincident with those of other agents. A negotiated solution should ideally be one in which all members of the community can individually concur. The multipartite satisficing set $\mathbf{\Sigma}_q$ and the individually satisficing set Σ_q^i provide each X_i with a basis for negotiation: an assessment from each agent's point of view of all options that are good enough for the group, and of the assessment of all individual options that are good enough for itself.

Compromise among a group of agents involves a lowering of standards by admitting possible actions that an agent, acting only unilaterally, would not necessarily prefer, but which it is willing to admit in the interest of acting as part of a group. The lowering of standards motivates the use of a satisficing outlook in the decision theory. An approach based on optimization, particularly one formulated on the basis of exclusive self-interest, does not admit grades or degrees. A choice is either optimal, or it is not. Compromise does not necessarily entail complete abolition of any agent's standards. An agent feeling that too much compromise is imposed it may walk away from the negotiating table.

In the formalism of multi-agent satisficing, an agent's index of caution q_i acts as a parameter representing the degree of compromise an agent is willing to adopt. By lowering the degree of caution, an agent is willing to consider placing more options in its satisficing set. As $q_i \to 0$, every option available to X_i is satisficing for X_i. If all agents are willing to sufficiently reduce their standards, a jointly acceptable solution can be obtained.

We will let $\mathbf{q} = (q_1, \ldots, q_N)$ denote the caution vector of the players. The least cautious index is $q_L = \min\{q_1, \ldots, q_N\}$. From an individual perspective, the negotiation theorem applies if an agent uses its own index of caution, q_i to determine the individually satisficing set, but uses q_L to determine its jointly satisficing set. It is assumed hereafter that each agent uses $q = q_L$ to determine $\mathbf{\Sigma}_q$. This reflects the conservative observation that the standards of a group can be no higher than the standards of any member of the group.

The relationship between individually and jointly satisficing sets for an agent is formalized by the following:

Definition 3 The set of all jointly satisficing vectors in $\mathbf{\Sigma}_q^i$ that are also individually satisficing for X_i is the *compromise set* $\mathbf{C}_i$, defined by

$$\mathbf{C}_i = \{\mathbf{u} = \{u_1, u_2, \ldots, u_N\} \in \mathbf{\Sigma}_{q_L}^i : u_i \in \Sigma_{q_i}\}.$$

□

Since $q_L \leq q_i$ the negotiation theorem indicates that $\mathbf{C}_i$ is not empty. By the negotiation theorem, if $u \in \Sigma_{q_i}$, then there is some $\mathbf{u} \in \mathbf{C}_i$ such that $u = \mathbf{u}(i)$.

We define

$$\mathbf{C}_i(j) = \{u_j : u_j \in \mathbf{u} \text{ for some } \mathbf{u} \in \mathbf{C}_i\}$$

as the set of all options for X_j in $\mathbf{C}_i$.

Definition 4 The *joint accord set* $\mathbf{N}_{q_L}$ is the set of all vectors that are jointly (at caution level q_L) and individually (at caution level q_i) satisficing for all agents. That is,

$$\mathbf{N}_{q_L} = \bigcap_{i=1}^{N} \mathbf{C}_i.$$

Any joint option $\mathbf{u} \in \mathbf{N}_{q_L}$ is a *joint accord option*. □

The desired outcome of a negotiation algorithm is a joint accord set. From $\mathbf{N}_{q_L}$, a single element joint accord option may be selected according to some tie breaking rule, such as the rule maximizing joint benefit to cost ratio,

$$\mathbf{u}^* = \arg \max_{\mathbf{u} \in \mathbf{N}_{q_L}} \frac{p_{\mathbf{S}}(\mathbf{u})}{p_{\mathbf{R}}(\mathbf{u})}. \tag{3}$$

This option is called the *rational compromise*.

If $\mathbf{N}_{q_L} = \emptyset$, then there are no decisions which are jointly and individually acceptable to all agents in the system.

Definition 5 In the context of multi-agent satisficing theory, *negotiation* is the process of working toward a solution which is individually and jointly acceptable to each agent in the system. □

Working toward a negotiated solution requires at least one of the agents to lower its standards, then recompute their compromise sets. If no agent is willing to compromise further (by lowering its own standards), then an impasse is reached in the negotiation process. Until that point is reached, however, negotiations may proceed in good faith. The lowering of standards is represented in this context by a lowering of an agents index of caution. Algorithm 1 outlines a negotiation algorithm based on this observation, termed the Enlightened Liberals algorithm ("liberal" in the sense of being tolerant of views other than one's own; "enlightened" in the sharing of information).

Algorithm 1 The Enlightened Liberals Negotiation Algorithm

Step 1: Initialize: $q_i = {}_0q_i$ for $i = 1, \ldots, N$. $q_L = \min({}_0q_i)$.

Step 2: X_i forms $\mathbf{\Sigma}_{q_L}^i$ and $\Sigma_{q_i}^i$.

Step 3: X_i forms $\mathbf{C}_i = \{\mathbf{u} \in \mathbf{\Sigma}_{q_L}^i : u_i \in \Sigma_{b_i}^i\}$.

Step 4: Communicate $\mathbf{C}_i$ and q_i to other agents.

Step 5: Each agent forms $\mathbf{N}_{q_L} = \cap_{i=1}^N \mathbf{C}_i$.

Step 6: If $\mathbf{N}_{q_L} = \emptyset$, each agent determines how much to lower q_i, then communicates q_i with the other agents. Then repeat from step 2.

Step 7: If $\mathbf{N}_{q_L} \neq \emptyset$, form the rational compromise $\mathbf{u} = (u_1, u_2, \ldots, u_N)$ according to (3).

In this algorithm, all agents communicate their choices in the same step. There is no way to use the partial information provided by another agent's compromise set to modify an agent's decisions.

Inference in negotiation

It may be noted that Enlightened Liberals is in accord with the first three principles of negotiation outlined in the introduction. However, no inference is employed in the algorithm as stated, since all agents essentially pass information simultaneously. The inference problem faced by agent X_i is to estimate $p^j_{\mathbf{S}_j}$, $p^j_{\mathbf{R}_j}$, p_{S_j}, and p_{R_j} — the praxeic system employed by X_j — given the offered solutions brought to the negotiating table in the form of $\mathbf{C}_j$. As an estimation problem all the tools of statistical estimation theory can be brought to bear, such as Bayesian estimates, maximum likelihood, minimum variance, maximum entropy, etc. (The method selected is problem dependent.) In the example presented below, a heuristic is illustrated which is similar to maximum likelihood.

Incorporation of the inference aspect of the negotiation is outlined in algorithm 2, which differs from Enlightened Liberals mostly in the sequence nature of the exchange of information.

Algorithm 2 The Inferring Liberals Negotiation Algorithm

Step 1: Initialize: $q_i = {}_0q_i$ for $i = 1, \ldots, N$. $q_L = \min({}_0q_i)$.

Step 2: X_i infers updates for $p^i_{S_j}, p^i_{R_j}, p^i_{\mathbf{S}}$ and $p^i_{\mathbf{R}}$, based on the compromise sets for $\{\mathbf{C}_j, j = 1, 2, \ldots, N, j \neq i\}$.

Step 3: X_i forms $\mathbf{\Sigma}^i_{q_L}$ and $\Sigma^i_{q_i}$

Step 4: X_i forms $\mathbf{C}_i$.

Step 5: X_i communicates $\mathbf{C}_i$ and q_i to all other agents.

Step 6: After all agents have transmitted their information, each agent forms $\mathbf{N}_{q_L} = \cap_{i=1}^N \mathbf{C}_i$.

Step 7: If $\mathbf{N}_{q_L} = \emptyset$, each agent determines how much to lower q_i, then communicates q_i with the other agents. Then repeat from step 2.

Step 8: If $\mathbf{N}_{q_L} \neq \emptyset$, form the rational compromise $\mathbf{u} = (u_1, u_2, \ldots, u_N)$ according to (3).

Example: Disputed Resource Allocation

> Complexity is no argument against a theoretical approach if the complexity arises not out of the theory but out of the material which any theory ought to handle.
>
> — Frank Palmer
> *Grammar* (1971)

We illustrate the negotiation framework outlined above — including interdependence factorization, establishing satisfiability functions and inference — by means of a resource allocation problem. Consider a situation in which N agents are to allocate M resources among themselves in such a way that all resources are allocated to at least one agent, but more than one agent may claim a resource. A resource with more than one claimant is a *disputed* resource. Disputed resources are of lower value than undisputed resources, both because utilization of a resource is attenuated by virtue of sharing, and because of an intrinsic societal valuation that would avoid dispute. The goal of each agent is to obtain as much of the resources as possible (or the resource allocation with the maximum valuation), while working toward having the fewest disputed resources as possible. Starting from some initial allocation of resource, each agent must also sustain a cost of acquisition for each additional resource that is required. While expressed as an abstract "resource allocation" problem, it may be helpful to envision the geographical division of a country among non-aligned factions. The apportionment of the land of Israel among Israeli and Palestinian claimants is a recent motivating example. Interestingly, data that might be adapted for a problem on a larger scale for Europe in the mid-twentieth century have been the subject of research in studies in cooperation and complexity (see, e.g., (Axelrod 1996)).

We will consider specifically only the two agent case; extensions to more agents is straightforward in principle. We will denote the allocation decision vector of agent X_i by a vector $u^i \in \{0,1\}^M$ where $u^i_j = 1$ if resource j is selected by agent i. (A final option denoted as $u^i = \emptyset$ might also be used to indicate that X_i is terminating the negotiation process and walking away from the negotiating table.) The decision vector $\overline{u}^i$ is used to indicate the Boolean complement of the decision vector u^i. For decision vectors u^1 and u^2, we will denote by $d = u^1 \cap u^2$ the disputed resources claimed by both agents[1] We denote by $u^i \backslash d$ the resources claimed exclusively by X_i. As agents begin the process, they have some initial allocation ${}_0u^1$ and ${}_0u^2$, with disputed resources

$$ {}_0d = {}_0u^1 \cap {}_0u^2. $$

Formulation of selectability and rejectability

The joint interdependence function is $p_{\mathbf{S},\mathbf{R}}(\mathbf{u},\mathbf{v}) = p_{S_1,S_2,R_1,R_2}(u_1,u_2,v_1,v_2)$. When we want to represent explicitly that this is the interdependence function as perceived by agent X_i, this will be denoted as $p^i_{S_1,S_2,R_1,R_2}(u_1,u_2,v_1,v_2)$. If this is changing with negotiation iteration number η, we will indicate this with $p^i_{S_1,S_2,R_1,R_2}(u_1,u_2,v_1,v_2;\eta)$.

To formulate specific results, it is expedient to factor the joint interdependence function into conditional and marginal probability measures. Conditional probabilities, as observed by Pearl (Pearl 1988) permit local or specific responses to be characterized. Conditional behavior is behavior at the local level, with all dependencies specified. Such factorizations permit characterization of global behavior in terms of local relationships, which are frequently easier to specify. A variety of factorizations are possible, even for the simple case of two agents, and it is not necessary for all agents to invoke the same factorizations. However, in this example, both agents will factor the joint interdependence function the same way.

We will express the factorization from the point of view of X_1, then express the inference process from the point of view of X_2 (using information from X_1). As a shorthand we will represent the factorization of the probability functions in terms only of their variables, using, for example, $S_2|R_2$ as a representation for $p_{S_2|R_2}(u_2|v_2)$. A reasonable (but not unique) factorization, expressed from the point of view of

[1] The notation $u_1 \cap u_2$ might be more exactly represented as $u_1 \wedge u_2$, where a "bitwise" AND operation is implied in each element. However, the $\cap$ notation seems to be more suggestive.

X_1 is

$$(S_1, S_2, R_1, R_2) = (S_1|S_2, R_1, R_2)(R_1|S_2, R_2)(S_2|R_2)(R_2).$$

Under the assumption that rejectability and selectability are independent for a given agent, we obtain

$$(S_1, S_2, R_1, R_2) = (S_1|S_2, R_2)(R_1|S_2, R_2)(S_2)(R_2) \quad (4)$$

There is an attractive symmetry in the first two factors, being the selectability and rejectability (respectively), conditioned on those quantities for the other agent. The first two terms of this factorization represent X_1's selectability and rejectability, respectively, when X_2 places all of its selectability mass and rejectability mass as the conditioning arguments. The factorization in (4) is expressed more explicitly as

$$p^1_{S_1,S_2,R_1,R_2}(u_1, u_2, v_1, v_2) = p^1_{S_1|S_2,R_2}(u_1|u_2, v_2)p^1_{R_1|S_2,R_2}(v_1|u_2, v_2)p^1_{S_2}(u_2)p^1_{R_2}(v_2). \quad (5)$$

In the sections below, we describe the parameters that affect each of the terms in this factorization.

Goals

Each agent wants to maximize the value of the resources it claims. There is a functional $g^i(u)$ adopted by X_i evaluating its intended allocation. This might be quite simple, as in[2]

$$g^i(u) = \sum_{j \in u} \varepsilon^i(j),$$

where $\varepsilon^i(j)$ measures the intrinsic value of resource j. This may include the size of the resource as well as other attributes. (In the case of land as a resource, it might measure attributes such as a harbor or an airport, mineral or agricultural assets, or historical or religious value.) In the case that the resources are distributed in space, the value may be determined using less localized measures. For example, there may be more value in having the resources as near to each other as possible, or in a contiguous block. Or there may be less value to a resource which is surrounded by resources claimed by the other agent. (In the case of land apportionment, an agent might prefer large pieces contiguously joined, with no islands of other agents' land in the middle.) All of these variations can be incorporated into $g^i(u)$.

Given that the agents' goals are prescribed by the desire to obtain more resources, we simply normalize the allocation value to form a probability mass function:

$$p^1_{S_1|S_2,R_2}(u_1|u_2, v_2) = p^1_{S_1}(u_1) \propto g^1(u_1).$$

That is, the the selectability is conditionally independent of X_2's options. In the joint selectability each agent thus acts independently:

$$p_{S_1,S_2}(u_1, u_2) = p_{S_1}(u_1)p_{S_2}(u_2).$$

In (5), the term $p^1_{S_2}(u_2)$ is X_1's model (or perception) of X_2's selectability. This is determined simply X_1's estimate of $g^2(u)$ (estimated according to X_1's knowledge of X_2).

[2]The set $\{j \in u\}$ appearing in this summation is a shorthand for $\{j\colon u_j = 1\}$

Costs

Several elements of the problem contribute to an agent's perception of the cost of the choice.

Reduce disputed resource Each agent evidences the difficulty of sharing the resource by seeking to eliminate the disputed resources. This not only serves his purposes — since disputed resources may not be enjoyed at full value — but also makes a concession to the society of the agents, which would prefer undisputed allocations.

In general, there is a cost function associated with disputed resources, which for agent X_i is denoted as $\delta^i(u^1 \cap u^2)$. This could depend on a variety of societal or historical factors. In some disputed resources, there may be no cost associated with more than one claim on the resources, whereas for others there is considerable cost.

A simple model for the cost is simply to make the disputation cost function proportional to the value of the disputed resource to each agent,

$$\delta^i(u^1 \cap u^2) = \delta^i(d) \propto \sum_{i \in d} \varepsilon^1_i + \varepsilon^2_i.$$

This cost can be placed in the context of the conditional probability $p^1_{R_1|S_2,R_2}(u_1|u_2, v_2)$ as follows.

- When $u_2 = \overline{v}_2$, then X_2 places all of its selectability and none of its rejectability on the vector u_2, so it is fully committed to the option u_2. Then it may be presumed that there will be a disputation of $d = v_1 \cap u_2$, and the cost becomes $\delta^1(v_2 \cap u_2)$.
- When $u_2 = v_2$, then X_2 places all of its selectability as well as all of its rejectability on u_2, and hence is conflicted. In this conflicted state, X_1 assumes half the cost of disputed territory, $\frac{1}{2}\delta^1(v_2 \cap u_2)$. (Other options are, of course, possible to deal with this conflicted state.)
- For those territories that X_2 has indicated that it doesn't want (no selectability and high rejectability), there is no cost for a disputed territory.

An overall rejectability function based on disputation can be formulated by normalization. We will call this rejectability function $p^1_{R_1|S_2,R_2;\delta}(u_1|u_2, v_2)$.

Cost of Acquisition There is a cost associated with acquiring the resources beyond the initial allocation. (For example, in the case of land resources, simply making the decision to acquire the land does not make it so. It may be necessary to deploy troops to enforce the decision, or to move in colonists, etc.) The cost of acquisition will also depend on the interest that the competing agent has in the new territory. In the context of the conditional probability $p^1_{R_1|S_2,R_2}(u_1|u_2, v_2)$, the following observations can be made.

- In the case that $u_2 = \overline{v}_2$ (that is, X_2 puts all of its selectability and none of its rejectability on the vector u_2), then it may be presumed that there will be a dispute over $d = v_1 \cap u_2$. The cost of the disputed acquisition is denoted by

$$\check{\chi}^1(d; {}_0u^1, {}_0u^2),$$

while the cost of the undisputed acquisition is

$$\hat{\chi}^1(v^1 \backslash d; {}_0u^1, {}_0u^2)$$

The total cost is then the sum of these:

$$\chi^1(v^1; {}_0u^1, {}_0u^2) = \hat{\chi}^1(v^1 \backslash d; {}_0u^1, {}_0u^2) + \check{\chi}^1(d; {}_0u^1, {}_0u^2)$$

- When $u_2 = v_2$; that is, X_2 is conflicted, placing all of its selectability as well as all its rejectability on u_2, then X_1 might assume that X_2 will not be in disputation; then

$$\chi^1(v^1; {}_0u^1, {}_0u^2) = \hat{\chi}^1(v^1; {}_0u^1, {}_0u^2)$$

- For those options on which X_2 places none of its selectability and all of its rejectability on, we will assume that same result as when X_2 is conflicted. (More generally, we could have a reduced cost for acquisition.)
- In the more general case, X_2 may be conflicted in some areas but not in others. In this case, X_1 only counts as disputed those territories which intersect with its interests and for which X_2 is unconflicted.

Combining these costs together and suitably normalizing, the rejectability function $p^1_{R_1|S_2,R_2;\chi}(u_1|u_2, v_2)$ is obtained.

Cost of negotiation An agent may attribute cost to the process of negotiation. If the negotiation must proceed through several iterations, an agent may become sufficiently annoyed at the process that its response is to walk away from the negotiating table. Several factors may be incorporated into the cost of negotiation, including the number of iterations (which we denote by η), or the apparent lack of progress (if the compromise sets coming from other agents appears to be unchanging). A cost based on the number of iterations can also represent determination of an agent to with respect to certain options: while the overall boldness is decreasing, the rejectability of some options can be correspondingly increased to partially offset the reduction. The cost of negotiation is represented by $\alpha^i(u_1, u_2, v_2; \eta)$; suitably normalized it becomes the rejectability function $p^1_{R_1|S_2,R_2;\alpha}(u_1|u_2, v_2; \eta)$

Overall conditional rejectability The conditional rejectability function in (5) is expressed as a convex sum of the rejectability functions described above:

$$p^1_{R_1|S_2,R_2}(u_1|u_2, v_2) = \beta_1 p^1_{R_1|S_2,R_2;\delta}(u_1|u_2, v_2) + \beta_2 p^1_{R_1|S_2,R_2;\chi}(u_1|u_2, v_2) + \beta_3 p^1_{R_1|S_2,R_2;\alpha}(u_1|u_2, v_2; \eta) \quad (6)$$

where $\sum_i \beta_i = 1$.

The marginal $p^1_{R_2}(v_2)$ and joint rejectability The quantity $p^1_{R_2}(v_2)$ in (5) is X_1's model of X_2's marginal (unconditional) rejectability. This is viewed (in this formulation) as separate parameter, not a derived quantity. A variety of factors influence the joint rejectability. Even if the factors could be computed exactly, the weighting factors in the combination may be unknown. The difficulty of estimating this reliably suggests the need to estimate this quantity, if possible, during the negotiating process. Inference of $p^i_{R_j}(v)$ is discussed below.

As the negotiating process begins, some initial condition is needed. One initial condition reflecting this uncertainty is to assume that $p^1_{R_2}(v_2)$ apportions equal rejectability to all options. Another approach is to allow $p_{R_2}(v_2)$ — as an unconditional measure — to reflect those aspects of the problem that are most independent of actions or goals of other agents. In this light, allowing $p^1_{R_2}(v_2)$ to be proportional to the cost of acquisition is reasonable,

$$p^1_{R_2}(v_2) \propto \chi^2(v_2; {}_0u^2) + \delta^2(v_2, {}_0u^1 \cap {}_0u^2).$$

It is straightforward to verify that the joint rejectability can be computed as

$$p_{R_1,R_2}(v_1, v_2) = \sum_{u_2 \in U_2} p^1_{R_1|S_2,R_2}(v_1|u_2, v_2) p^1_{S_2}(u_2) p^1_{R_2}(v_2).$$

Inference

We consider now the question of inference of the parameters of other agents during the course of negotiation, presenting a method which is reasonable in the context of the present problem. After its initial decision-making step, X_1 presents $\mathbf{C}_1$ and q_1 to X_2. Based on this compromise set and caution index, what can be inferred about X_1's satisfiability functions? Because X_2 will be doing its computations based on the factorization (5), it may be assumed that that X_2 has a model of $p_{S_1}(u)$, since this is based primarily on economic questions which are observable by all agents. As mentioned above, however, $p^j_{R_i}(u)$ is difficult to obtain without further information. This parameter influences the joint rejectability $p^2_{R_1,R_2}$, and hence the marginal $p^2_{R_2}$, so its estimation has an extended influence in the decision making process. (This section, for the sake of definiteness, is presented as if X_2 were making inference based on information from X_1.)

Given $p^2_{S_1}(u)$ and $p^2_{R_1}(u)$, consider the joint options in $\mathbf{C}_1$. If $p^2_{S_1}(u) \geq q_1 p^2_{R_1}(u)$ and $u = \mathbf{u}(i)$ for some $\mathbf{u} \in \mathbf{C}_1$, then the compromise set provides no information: it reflects decisions that would be made by X^2 using its estimates. Also, if $p^2_{S_1}(u) < q_1 p^2_{R_1}(u)$ and $u \notin \mathbf{C}_i(i)$ then no additional information is provided: X_2 did not expect the choice, and X_1 did not select it.

However, if $p^2_{S_1}(u) \geq q_1 p^2_{R_1}(u)$ and $u \notin \mathbf{C}_i(i)$ then X_1 has rejected option, both individually and jointly, which according to X_2's model it should have accepted. Furthermore, if $p^2_{S_1}(u) < q_1 p^2_{R_1}(u)$ but $u \in C_i(i)$, then X_1 has accepted options both individually and jointly which, according to X_2's model it should have rejected. Both of these circumstances evince that X_2's model $p^2_{R_1}(u)$ is inaccurate at u and stands updating. Our inference rule is to change the rejectability $p^2_{R_1}(u)$ in such a way that these inconsistencies are resolved, and in such a way that the change at each point is minimized while ensuring that the probability constraint is satisfied.

Define the sets

$$\overline{U} = \left\{ u \in U_1 \colon \begin{array}{c} p^2_{S_1}(u) < q_1 p^2_{R_1}(u) \text{ and } u \notin \mathbf{C}_i(i) \\ \text{or} \\ p^2_{S_1}(u) \geq q_1 p^2_{R_1}(u) \text{ and } u \in \mathbf{C}_i(i) \end{array} \right\}$$

$$\hat{U} = \{u \in U_1 \colon p^2_{S_1}(u) < q_1 p_{R_1}(u) \text{ and } u \in \mathbf{C}_i(i)\}$$

$$\check{U} = \{u \in U_1 \colon p^2_{S_1}(u) \geq q_1 p_{R_1}(u) \text{ and } u \notin \mathbf{C}_i(i)\}$$

Elements in $\overline{U}$ have rejectability consistent with $\mathbf{C}_1$. Elements in $\hat{U}$ have rejectability too high to be consistent with $\mathbf{C}_1$, while elements in $\check{U}$ have rejectability too low. For $u \in \hat{U}$ or $u \in \check{U}$, form updated rejectability functions by

$$p^2_{R_1,\text{new}}(u) = \alpha(u) p^2_{R_1}(u)$$

where

$$\alpha(u) = \begin{cases} \frac{p^2_{S_1}(u) - \epsilon}{q_1 p^2_{R_1}(u)} & u \in \hat{U} \\ \frac{p^2_{S_1}(u) + \epsilon}{q_1 p^2_{R_1}(u)} & u \in \check{U} \end{cases}$$

for some small positive ϵ. This introduces a net change in rejectability

$$\Delta = \sum_{u \in \hat{U} \cup \check{U}} p_{R_1}(u)(1 - \alpha(u))$$

which is to be distributed among the rejectabilities of the elements in U_1 with the smallest change possible without affecting the decision boundaries. Let

$$\overline{U}_{\geq} = \{u \in \overline{U} \colon p^2_{S_1}(u) \geq q_1 p^2_{R_1}(u)\}$$

and let

$$\overline{U}_{<} = \{u \in \overline{U} \colon p^2_{S_1}(u) < q_1 p^2_{R_1}(u)\}.$$

The notation $|U|$ denotes the number of elements in the set U. Then the redistribution is as follows:
If $\Delta > 0$ (i.e., rejectability is added):

- Distribute Δ among $\overline{U} \cup \check{U}$ as equally as possible, such that in $\overline{U}_{<} \cup \check{U}$, $p^2_{R_1,\text{new}}(u) < 1$ and in $\overline{U}_{\geq}$, $p^2_{S_1}(u) \geq q_1 p^2_{R_1,\text{new}}(u)$.

If $\Delta < 0$ (i.e., rejectability is removed):

- Distribute Δ among $\overline{U} \cup \hat{U}$ as equally as possible, such that in $\overline{U}_{\geq} \cup \hat{U}$ $p^2_{R_1,\text{new}}(u) > 0$ and in $\overline{U}_{<}$, $p^2_{S_1}(u) < q_1 p^2_{R_1,\text{new}}(u)$.

This simple-minded inference does not fully exploit the information available. For example, if by $\mathbf{C}_1$ X_1 appears uninterested in some resource, X_2 could parameterize an increase its interest in that area either by lowering a rejectability with respect to its acquisition, or by increasing its selectability. However, the inference described above suffices to demonstrate the concept.

Numerical demonstration

Consider a country with four regions as shown in figure 1. Values are apportioned in such a way that adjacent regions have a value greater than the sum of the constituent areas, and that value increases with more regions, as shown in table 1. The cost of disputed regions is also shown in table 1 (unnormalized). Cost of acquisition is on a region-by-region basis, as shown in table 2. The initial allocation is ${}_0u^1 = [1110]$ (countries 2, 3, and 4) and ${}_0u^2 = [1101]$. Figure 2 illustrates the sequence of estimated probabilities $p^1_{R_2}(u)$ and $p^2_{R_1}(u)$ for six iterations. (The abscissa represents the choice u as a decimal representation of the binary decision vector. The lines spaced within an integer $[u, u+1)$ represent the probability estimates for different iterations of the negotiation algorithm.) After several iterations of negotiation, the compromise sets shown in table 3 are obtained, and the joint accord set $\mathbf{N}$ join in table 4 is obtained, where the rational compromise is indicated with *. (The integers represent the decimal form of the corresponding binary option vectors). The individually satisficing sets are $\Sigma_{q_1} = \{14\}$ and $\Sigma_{q_2} = \{5, 9, 12, 13\}$. The final boldness reached is $\mathbf{q} = (1.3, 1.3)$, after starting at $\mathbf{q}_0 = (1.9, 1.9)$ and decrementing each time by $\Delta q = 0.05$. It is interesting to note that even after negotiation, for the given value/cost data, the two agents end up with disputed regions, and that the initial conditions still remain in the joint accord set. However, having gone through the negotiating process, while regions must be shared, the agents may feel that they have "bought in" to this circumstance, since these options are individually satisficing.

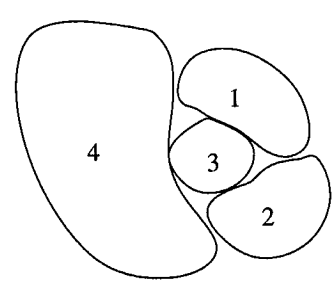

Figure 1: Four resources to be distributed

Discussion

Within this framework for negotiation there are several observations that may be made. In some human negotiations, parties often repeat a position repetitively, without an apparent change of state, until at some point there is an abrupt change in feasible options. The procedure represented here provides a model for such behavior: even when from one iteration to the next there might be no change in compromise sets, each agent is modifying its models of the other agent, lowering its caution, and potentially changing its rejectability as a function of the number of iterations.

u	$\mathbf{u}$	$p^1_{S_1}(u)$	$p^2_{S_2}(u)$	$\delta^1(u)$	$\delta^2(u)$
0	0000	0	0	0	0
1	0001	0.0210	0.0227	3	3
2	0010	0.0252	0.0182	3	3
3	0011	0.0546	0.0500	7	7
4	0100	0.0420	0.0318	6	5
5	0101	0.0714	0.0636	8	9
6	0110	0.0756	0.0500	8	9
7	0111	0.1092	0.0818	15	12
8	1000	0.0252	0.0455	2	6
9	1001	0.0504	0.0727	4	9
10	1010	0.0546	0.0636	4	9
11	1011	0.0588	0.0727	9	5
12	1100	0.0714	0.0818	9	12
13	1101	0.1008	0.1045	13	14
14	1110	0.1050	0.0955	13	14
15	1111	0.1345	0.1455	20	20

Table 1: Valuations for resource allocation

u	$\check{\chi}^1$ (disp.)	$\hat{\chi}^1$ (undisp.)	$\check{\chi}^1$ (disp.)	$\hat{\chi}^1$ (undisp.)
0001	6	3	6	3
0010	4	2	4	2
0100	4	2	6	3
1000	8	4	10	5

Table 2: Cost of acquisition per country

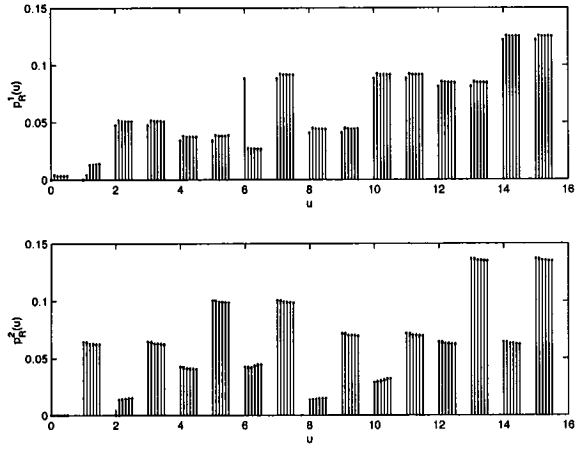

Figure 2: Sequence of estimated probabilities

$\mathbf{C_1}$	$\mathbf{C_2}$			
(14,1)	(2,5)	(4,9)	(15,9)	(2,13)
(14,3)	(4,5)	(6,9)	(2,12)	(4,13)
(14,4)	(5,5)	(7,9)	(4,12)	(6,13)
(14,5)	(6,5)	(8,9)	(6,12)	(7,13)
(14,6)	(7,5)	(9,9)	(7,12)	(8,13)
(14,8)	(8,5)	(10,9)	(8,12)	(10,13)
(14,9)	(12,5)	(11,9)	(10,12)	(12,13)
(14,12)	(14,5)	(12,9)	(12,12)	(13,13)
(14,13)	(2,9)	(13,9)	(14,12)	(14,14)
(14,15)	(3,9)	(14,9)	(15,12)	(15,13)

Table 3: Compromise sets after negotiation

$\mathbf{N}$
(14, 5)*
(14,9)
(14,12)
(14,13)

Table 4: Joint accord set

Other behaviors such as recalcitrance or openness can be modeled depending on how the boldness is changed.

A concern that may be raised regarding this procedure is its computational complexity. A large measure of the complexity arises due to the computation of marginals in the formulation of p_{R_1,R_2}. The complexity can be mitigated somewhat by efficient organization of the computations, using, for example, the factor graph approach described in (Kschischang, Frey, & Loeliger 2001). Another approach is to absorb the normalization used determining in p_{S_1,S_2} and p_{R_1,R_2} into the index of caution q. Once an initial q can be determined which provides for meaningful individual and joint solutions, the index of caution is adjusted until a group accord is established.

In conclusion, the multi-agent satisficing theory provides a means of describing solutions which are individually and jointly satisficing from the perspective of an individual agent in the community of agents. We have provided a definition of negotiation, which is the process of working to achieve accord among the different agents with regard to the solutions they they find acceptable, and provided some algorithms to implement that process. To demonstrate how the theory may be applied to a multi-agent problem, a resource allocation problem was presented in which agents vie for disputed resources.

References

Axelrod, R. 1996. http://www.pscs.umich.edu/Software/-CC/CC4. Univ. of Michigan Center for the Study of Complex Systems.

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

Goodrich, M. A.; Stirling, W. C.; and Frost, R. L. 1998. A theory of satisficing decision and control. *IEEE Trans. Systems, Man, Cybernet.* 28(6):763–779.

Kschischang, F.; Frey, B.; and Loeliger, H. 2001. Factor graphs and the sum-product algorithm. *IEEE Trans. Info. Theory* 47(2):498–519.

Levi, I. 1967. *Gambling with Truth*. Cambridge, MA: MIT Press.

Levi, I. 1980. *The Enterprise of Knowledge*. Cambridge, MA: MIT Press.

Levi, I. 1984. *Decisions and Revisions*. Cambridge: Cambridge Univ. Press.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. San Mateo, CA: Morgan Kaufmann.

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

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

Stirling, W. C., and Moon, T. K. 2001. A praxeology for rational negotiation. In *AAAI Symposiusm on Negotiation Methods for Autonomous Cooperative Systems*. (submitted).

Stirling, W. C.; Goodrich, M. A.; and Frost, R. L. 1996a. Procedurally rational decision-making and control. *IEEE Control Systems Magazine* 16(5):66–75.

Stirling, W. C.; Goodrich, M. A.; and Frost, R. L. 1996b. Satisficing intelligent decisions using epistemic utility theory. In *Intelligent Systems: A Semiotic Perspective. Proceedings of The 1996 International Multidisciplinary Conference Conference on Multi-Agent Systems, Volume II: Applied Semiotics*, 268–273.

Stirling, W. C.; Goodrich, M. A.; and Frost, R. L. 1996c. Toward a theoretical foundation for multi-agent coordinated decisons. In *Proceedings of The Second International Conference on Multi-Agent Systems*, 345–352.

Stirling, W. C.; Goodrich, M. A.; and Packard, D. J. 2001. Satisficing equilibria: A non-classical approach to games and decisions. *Autonomous Agents and Multi-Agent Systems Journal*. Forthcoming.

Stirling, W. 1994. Multi-agent coordinated decision-making using epistemic utility theory. In Castelfranchi, E., and Werner, E., eds., *Artificial Social Systems*. Berlin: Springer-Verlag. 164–183.