

Procedurally Rational Decision-Making and Control

Wynn C. Stirling, Michael A. Goodrich, and Richard L. Frost

Substantive rationality requires a decision-maker to be a utility maximizer; under this paradigm, the decision is paramount, and not dependent on the computational process used to obtain it. Procedural rationality is dependent on the method used to make the decision; reasonableness of the procedure is paramount. Well-formed problems are amenable to substantive rationality; ill-formed problems are not, but are amenable to procedural rationality. To qualify as being procedurally rational, a methodology must possess a sound epistemological basis, it must be amenable to a formal design synthesis procedure, and it must be consistent with substantive rationality. Epistemic utility theory forms the basis of a decision rule that is procedurally rational. This theory adapts to decision-making in the context of control theory, and leads to a specific design procedure that may be applied to single- and multiple-agent ill-formed control problems.

Rational Choice

The dominant paradigm for decision-making in engineering, the physical sciences, statistics, and mathematical economics is self-interested *rational choice*, whereby the decision-maker, or *agent*, seeks to implement a policy that achieves its own greatest good or greatest preference. Simon terms this definition of rational choice *substantive rationality* [1, pp. 26-27]; it is usually implemented by maximizing an appropriate utility function. Two consequences of this paradigm are potentially objectionable. First, the agent is concerned solely with the *outcome* of the decision-making mechanism; it is not concerned with the *process* that leads to the outcome. In other words, though the paradigm establishes the existence of a best solution, it is not constructive, since it is not concerned with the agent's modes of calculation. Second, it never makes sense deliberately to make a choice that is merely good or sufficient for one's purposes rather than to choose what is best for oneself.

In the social sciences, however, the paradigm of substantive rationality is not always appropriate, either because there is not a clearly defined and unique utility, or because the agent is not a utility maximizer. In such cases, attention is focused on the *process* by which choices are made [2], with secondary attention regarding the outcome. Rationality, in this context, refers to the reasonableness of the procedures used to arrive at a decision. In Simon's terminology, this concept of rational choice is called *procedural rationality*. He argues that procedures cannot be ignored: "broadly stated, the task is to replace the global ration-

ality ... with a kind of rational behavior that is compatible with the access to information and the computational capabilities that are actually possessed ..." [3]. Since the resulting outcome may not be one that maximizes utility, it may not be substantively rational. If it has any claim to rationality at all, the claim must be justified through the procedures used to generate it.

Decision-making in the engineering discipline is heavily biased toward substantive rationality. For example, optimal control, optimal estimation, and optimal detection all stem from the common heritage of maximizing utility. Two basic, but often unarticulated, assumptions must be satisfied, however, if an optimization paradigm is to be successful. First, a *best* solution must exist. Second, there must be sufficient information, time, and resources to find it. Meystel terms problems that meet these criteria "well-formed problems" [4]. An ill-formed problem either does not have a unique or well-defined notion of best, or insufficient information, time, or resources exist to find a best solution. For cases where an optimal solution is either impossible, intractable, or impractical, the only recourse is to implement a non-optimal solution.

But the moment the notion of non-optimality is given play, the paradigm switches from substantive rationality to procedural rationality. The problem focus shifts from simply *implementing* the optimal solution that exists independently of procedures to obtain it, to *finding* a solution to substitute for the inaccessible or unimplementable optimal one. A "suboptimal" solution presupposes the existence of an optimal solution and is usually formed as a compromising modification to it (for example, using steady-state, rather than time-varying, gains). Consequently, suboptimal control does not represent a paradigm shift from substantive rationality. If an optimal solution cannot be found or implemented, what is evidently needed is a notion of rational control that does not draw its viability from the concept of optimality.

In the context of an optimality paradigm, the acceptance of non-optimality is pragmatic; no formal justification to legitimize any such form of decision-making as rational can be obtained within the strict purview of substantive rationality. Unfortunately, neither the optimization principle nor the actual optimal solution provides a systematic design synthesis procedure for calculating any of the possibly very good, but not strictly optimal, solutions that may exist. It is often left to the designer's own ingenuity, independently of optimization, to devise a solution.

In this essay we explore the thesis that there are rational alternatives to the extreme notions of optimal control on the one hand, and *ad hoc* control based on imprecise or vague notions of desirability on the other. We first establish criteria for a control synthesis procedure to be rational. We then present a principle-

The authors are with the Electrical and Computer Engineering Department, 459CB, Brigham Young University, Provo, Utah, 84602. (Email: wynn@ee.byu.edu).

based synthesis procedure applicable to ill-formed control problems.

Criteria for Procedural Rationality

Virtually every real-world problem deals with issues that cannot be completely idealized, and solution procedures, as well as the end results, are critical to successful design. Furthermore, even when an optimal solution is found, it is susceptible to fine-tuning of the parameters of the performance index—a procedure which can significantly modify results. Several respected procedural-based techniques are available for nonlinear control problems; well-known examples include feedback linearization, adaptive control, and gain scheduling. These techniques, though highly procedure based, are usually considered part of the conventional control methodology, since the procedure usually contains a step that involves utility maximization.

Increasingly, research in control theory is focusing on purely procedural methodologies as alternatives to conventional control. Fuzzy logic, expert systems, genetic algorithms, and neural networks are examples of procedure-based control concepts. Controllers based on these concepts are sometimes termed “intelligent,” since they are motivated by the types of decision-making and problem representations that are found in human, animal, or biological systems. These alternative approaches sometimes meet resistance from the conventional control community, in part because they are new and different, but in part also because they do not possess a compelling substantively rational component (that is, they do not involve utility maximization).

The fact that a solution arises from a procedure, either from conventional or alternative roots, does not automatically endow it with rationality. Though a conventional procedure may possess an optimal component, that is no guarantee that the entire procedure is reasonable. Furthermore, though an alternative procedure may be based on biological, animal, or human system models, heuristic arguments are not adequate to establish rationality, nor are appeals to biological or anthropomorphological metaphors.

To make a judgment that an action is reasonable requires an understanding of the context in which the agent operates. A procedure is rational if it is appropriate for the situation, is likely to lead to the attainment of the goals, and conforms to the available computational capabilities. Unfortunately, these characteristics are in tension; for example, there may not be sufficient computational capability to assure that the goal can be met. Consequently, procedural rationality cannot provide the same kind of solution existence guarantee that substantive rationality exhibits. What can be established, however, are criteria for evaluating a procedure in terms of its logic, its practicality, and its consistency.

Logic: A rational procedure must be based upon epistemological principles. An epistemological basis will establish the precise aim of the procedure and will provide means for interpreting the results. Without a firm philosophical basis, a procedure, no matter how novel it may be, is *ad hoc*.

Practicality: A rational procedure must be formalized, in that it provides a mechanism to follow that is independent of the problem context. A formalism provides a systematic synthesis recipe. Without a formalism, the procedure cannot be applied to new contexts, and it cannot easily be taught or learned.

Consistency: When applied to problems for which a substantively rational (that is, optimal) solution exists, a rational procedure must be consistent with the substantively rational perspective in that if both solutions use exactly the same information, the optimal solution must be attainable, at least in principle, via the procedure. Consistency provides an important link between procedural and substantive rationality. Without consistency it would be difficult to develop confidence in a procedure-based solution.

Epistemology

Epistemology is the study of classifying propositions on the basis of knowledge and belief. It seeks to define knowledge, to ascertain what it means to know something, and what conditions need to be satisfied before a proposition becomes knowledge. Alston [5] distinguishes between two schools of epistemology: *substantive epistemology* and *meta-epistemology*: “Meta-epistemology is concerned with the basic concepts we employ in epistemology, concepts of *knowledge, truth, belief, justification, rationality*, and so on, and with the methods, procedures, and criteria to be employed in determining how to apply those concepts. Substantive epistemology, on the other hand, consists in our endeavors to use these concepts to arrive at results on such matters as the conditions under which we have knowledge or justified belief of one kind or another, and on what knowledge or justified belief we have.” [5, p. 2].

Substantive epistemology is concerned with results. To justify a proposition is to establish epistemic principles: “Sometimes justification is alleged to require tracing ... causes of belief to legitimating sources ... beliefs become knowledge only if they can be derived from impeccable first premises according to equally noble first principles” [6, p. 1]. The paradigm of substantive rationality is founded upon this epistemology. Acquiring justification, in the epistemological case, and obtaining the optimal solution, in the control case, are of primary importance under this paradigm. Both are based on principles. Appealing to legitimating sources or noble first principles is analogous to claiming that the best solution exists independently of the method used to obtain it. In this sense, neither substantive epistemology nor substantive rationality are constructive approaches to decision-making (that is, they do not provide a specific computational procedure).

Meta-epistemology is deeply concerned with procedures; it is a constructive epistemology. Justification involves establishing the reasonableness of the procedures in light of the information available, not appealing to principles. An essential feature of a meta-epistemology is that it must provide a formal *decision rule* expressed in mathematical language. This decision rule embodies not only the philosophical underpinnings of the epistemology, it also provides the dynamic impetus for the justification procedure. The paradigm of procedural rationality is born of this epistemology. One of the goals of this essay is to apply meta-epistemological principles to the domain of control engineering.

Two main aspects of virtually any control problem are (a) the specification of the ultimate goal of the controller and (b) the design criteria, or cost of control. If cost is disregarded, many control functions exist that will achieve the fundamental goal, and if the goal is disregarded, many low-cost control functions exist. Those controls that perform well at achieving the goal

(regardless of their cost) will be said to possess a great degree of accuracy, and those controls that conform well to the design criteria (regardless of how well they achieve the goal) will be said to possess a low degree of rejectability. Any control function that has a high degree of accuracy and a low degree of rejectability is clearly a good, if not optimal, control.

Both accuracy and rejectability notions are present in modern optimal control formulations, but their presence is often implicit in the structure of the performance index. For example, consider the quadratic regulator problem

$$\mathbf{x}(k+1) = \mathbf{f}[\mathbf{x}(k), \mathbf{u}(k), k], \quad \mathbf{x}(0) = \mathbf{x}_0, \quad k = 0, 1, \dots \quad (1)$$

where \mathbf{f} is a n -dimensional vector function which possess the necessary conditions for the existence of a unique solution, $\mathbf{x}(\cdot)$ is the state vector, and $\mathbf{u}(\cdot)$ is the m -dimensional control. The problem is to choose a control sequence $\{\mathbf{u}(0), \mathbf{u}(1), \dots, \mathbf{u}(k_f-1)\}$, with terminal time, k_f , unspecified, so as to minimize the performance index

$$J = \mathbf{x}^T(k_f)P\mathbf{x}(k_f) + \sum_{k=0}^{k_f-1} [\mathbf{x}^T(k+1)Q(k+1)\mathbf{x}(k+1) + \mathbf{u}^T(k)R(k)\mathbf{u}(k)], \quad (2)$$

where $P \geq 0$, $Q(k) \geq 0$, and $R(k) \geq 0$. Accuracy is represented by the terminal portion of the this index,

$$J_A = \mathbf{x}^T(k_f)P\mathbf{x}(k_f). \quad (3)$$

Forcing J_A to be small coincides with the fundamental goal of the controller, which is to drive the terminal system state to the origin. Rejectability is represented by the transitional portion of the index,

$$J_R = \sum_{k=0}^{k_f-1} [\mathbf{x}^T(k+1)Q(k+1)\mathbf{x}(k+1) + \mathbf{u}^T(k)R(k)\mathbf{u}(k)]. \quad (4)$$

Keeping J_R small represents the desires of the agent to keep the cost of control small.

Whereas the notions of accuracy and rejectability are implicit in the traditional optimization approach of minimizing $J = J_A + J_R$, there are advantages to having them appear explicitly. Accuracy and rejectability are independent notions; knowing the accuracy of a control provides no information regarding its rejectability, and vice versa. Since a highly accurate control may not have low rejectability, however, the two notions may be in tension.

Satisficing: A Foundation for a Formal Procedure

Is there a way to formally legitimize the search for good, but not necessarily best, solutions to a problem? The issue is to find a procedure to fill the void left when the maximization decision engine is removed from consideration. To address this issue, we must undergo a complete rethinking of the basic philosophy underlying the goals of our inquiry. James offers some insight:

There are two ways of looking at our duty in the matter of opinion—ways entirely different, and yet always about whose

difference the theory of knowledge seems hitherto to have shown very little concern. We must know the truth, and we must avoid error—these are our first and great commandments as would-be knowers; but they are not two ways of stating an identical commandment, they are two separable laws ...

Believe truth! Shun error!—these, we see, are two materially different laws; and by choosing between them we may end by coloring differently our whole intellectual life. We may regard the chase for truth as paramount, and the avoidance of error as secondary; or we may, on the other hand, treat the avoidance of error as more imperative, and let truth take its chance. [7, pp. 17-18].

The notions of truth and error are appropriate for epistemology, but analogous notions must be developed if the context changes. In the control context, the imperative to “obtain the best and only the best” control solution is analogous to the epistemological stance to “seek the truth and nothing but the truth.” If the epistemological stance is then changed to “avoid error,” the analogous control imperative is modified to become “eliminate bad solutions.”

Substantive rationality is couched in superlative terms: maximize a utility function. It is more natural, however, to couch procedural rationality in comparative terms: reject all and only those propositions that do not achieve at least some minimum standard of performance. If a minimum standard, or threshold, can be established, then only those controls that do not meet or exceed this threshold will be rejected from consideration. All other controls will be deemed “good enough,” and any of them may be implemented. Economists have coined the term *satisficing* to describe choices that, though perhaps not the best possible, at least meet a minimum aspiration level [3, 8]. We will designate controls that meet a minimum standard threshold as *satisficing controls*. Satisficing is a more general concept than optimizing. An optimal control is clearly satisficing, but the notion of a minimum standard persists even if a best solution either does not exist or is not attainable. If the minimum standard admits an epistemological justification, the resulting solution cannot be dismissed as *ad hoc*. A key result of this essay, therefore, is a definition of the minimum performance standard.

Globality, Locality, and Consistency

If an agent possesses all information relevant to the problem over its full extent, then it has sufficient information, in principle, to obtain a solution. Such a scenario admits a global solution. If an agent either does not possess all relevant information over the full extent of the problem or cannot actually use all of the available information, then the agent must resort to a local, or limited, solution. Locality may refer to temporal locality, spatial locality, or to other aspects, such as the status of some but not all agents in a multiple-agent environment. Well-formed problems are global; ill-formed problems are limited.

We may, in principle, solve a global problem from either the substantive perspective or the procedural perspective. It would be disquieting, however, if the optimal solution were not an element of the set of procedurally rational solutions when both solutions use exactly the same information. There is no reasonable excuse for a procedure designed to find globally good solutions to fail to include the best one. We must therefore establish conditions, when applying a procedure-based approach to a problem that admits an optimal solution and when the same




Fig. 1. Inverted pendulum on a cart.

information is used, for the optimal solution to be included in the set of procedurally rational solutions. When this situation obtains, the approaches will be said to be *consistent*. Consistency applies only in the context of globality. For problems addressed from a local perspective, we cannot guarantee that the globally optimal solution will be contained in the set of limited procedurally rational solutions. In this situation we can only resort to the reasonableness of the procedure to evaluate the solution.

A Procedure for Ill-Formed Control Problems

We present a formal procedure for controller design that is suitable for addressing ill-formed control problems. This procedure establishes an explicit minimum standard required for satisficing control. For a procedure to be useful, it must possess at least the following properties.

1. It must be systematic, in that it permits the development of a body of knowledge or expertise that is transferable between contexts.
2. It must incorporate design principles as characterized by a utility function, performance index, or other appropriate means.
3. It should not require restrictive modeling assumptions solely for the purpose of expediting a solution.
4. It must conform to the amount of information available; problems with only limited, or local, information must be addressable from that perspective.

The philosophical motivation for this development is derived from the work of Levi [6, 9-11]. Although the focus of Levi's work is limited to cognitive decision-making, the basic structure is easily adapted to practical contexts, since it is based in probability theory. We first define the state of nature and the agent's control space, we then present epistemic utility theory as a decision engine, and we next introduce the concept of strongly satisficing control.

States of Nature

Let U denote the *decision space*, or the collection of possible control values that are relevant to the problem. In addition to the control value, many other variables are relevant to a control problem. For example, the state, x , of the system given by (1) is part of the state of nature. We must also include in the state of nature the values of other possible entities that have bearing on the decision, such as the states of other agents. Let Θ denote the set of possible states of nature; Θ will often be a subset of \mathfrak{R}^n

(n -dimensional Euclidean space). In our development it will be important to distinguish between the decision space and the state of nature; this will lead to an extension, due to Kenney [12, 13], of the original epistemic utility theory developed in [6] and first applied in the engineering context in [14-19].

We will regard the state of nature as a random variable (or vector), θ , governed by the probability $P_C: \mathcal{B} \rightarrow [0, 1]$ where \mathcal{B} is the Borel field in Θ . Let F_C denote the cumulative distribution function of θ , and assume that the probability density function $f_C(\vartheta) = \frac{d}{d\vartheta} F_C(\vartheta)$ exists for all $\vartheta \in \Theta$. For a deterministic control problem, the state of the system is observable from the measurements taken of the system. In this situation, the true state of nature, ϑ_0 , is exactly known, so all of the probability mass is assigned to the true state of nature, and $f_C(\vartheta) = \delta(\vartheta - \vartheta_0)$ (the Dirac delta function).

Epistemic Utility

As a prelude to a formal development, we provide a heuristic prescription for the minimum standard we wish to impose on controls to render them satisficing: *eliminate all controls for which rejectability exceeds accuracy*. The only surviving possibilities will then be those controls that both achieve the ultimate goal and are not too expensive. To formalize this procedure, we must establish the basic structure of the accuracy and rejectability utility functions.

Utility theory as it is generally applied in decision theory is used to provide a notion of preference in terms of order relationships. If an agent possesses an accuracy utility, it may order its choices accordingly; it may also order its choices according to its rejectability utility. Let f_A and f_R denote accuracy and rejectability utilities, respectively. Since utility assessments will in general be dependent on the state of nature, both f_A and f_R must depend not only on the control, u , but also on ϑ . We indicate this dependence by the notation $f_A(u; \vartheta)$ and $f_R(u; \vartheta)$. If $u_1 \in U$ and $u_2 \in U$, the condition $f_A(u_1; \vartheta) < f_A(u_2; \vartheta)$ means that u_2 is more accurate than u_1 , given the state of nature, ϑ , and the condition $f_R(u_1; \vartheta) < f_R(u_2; \vartheta)$ means that u_1 is less rejectable (that is, more acceptable) than u_2 , given ϑ .

Difficulties may arise, however, if we attempt to interpret an accuracy preference in terms of a rejectability preference, but that is precisely what our (as yet heuristic) satisficing decision rule requires. To make any such interpretations meaningful, we must impose additional structure on the utility functions. If a utility is required only for ordering relative to a single preference, its scale and bias are immaterial; the utilities $f(u)$ and $g(u) = af(u) + b$, where $a > 0$, are equivalent, so order is preserved. Consequently we may assume, without loss of generality, that the utilities f_A and f_R are normalized such that, for each ϑ , $f_A(u; \vartheta) \geq 0 \forall u \in U$, $f_R(u; \vartheta) \geq 0 \forall u \in U$, and $\int_U f_A(u; \vartheta) du = \int_U f_R(u; \vartheta) du = 1$. In other words, we may assume that f_A and f_R are probability density functions (or probability mass functions if U is discrete) parameterized by the state of nature, ϑ . With this convention, we may compute the accuracy and rejectability of any Borel subset $V \subset U$ as

$$P_A(V; \vartheta) = \int_V f_A(u; \vartheta) du \quad (5)$$

$$P_R(V; \vartheta) = \int_V f_R(u; \vartheta) du. \quad (6)$$

Thus, P_A and P_R are transition probabilities. Since they are normalized utilities, it is possible to relate the amount of accuracy, $P_A(V; \vartheta)$ to the amount rejectability, $P_R(V; \vartheta)$ associated with a given set $V \subset U$ for a given state of nature. A natural way to combine these utilities is to form a single utility that is a function of both desiderata. Levi [6, 9, 10] proposes that such a utility can be formed as a function of the two individual utilities: "... it will often be the case that an individual's system of preferences may be thought of as represented by a utility function which is itself a function of other utility functions each of which represents a desideratum of interest he is committed to taking into account in his deliberations" [10, p. 97]. Levi advocates that a convex combination of the two individual normalized utilities comprise an *epistemic utility* function:

$$\begin{aligned} P_e(V; \vartheta) &= \beta P_A(V; \vartheta) + (1-\beta)(1 - P_R(V; \vartheta)) \\ &= \beta \int_V f_A(u; \vartheta) du + (1-\beta) \int_{U-V} f_R(u; \vartheta) du \end{aligned} \quad (7)$$

for all Borel sets $V \subset U$, where for $0 \leq \beta \leq 1$. This single utility actually employs the converse to rejectability, namely $1 - P_R(V; \vartheta)$. The resulting structure rewards those controls that have high accuracy and low rejectability, with the parameter β acting as a weighting parameter between these two desiderata.

The name, epistemic utility, derives from the original usage by Levi in the context of evaluating propositions on the basis of their truth value and their informational value. Although we have changed the context, we retain the name, since the fundamental purpose of the utility has not changed: it provides a mechanism for combining both desiderata into one utility. In the control context, the desiderata are that controls meet the fundamental goal of the controller (accuracy), but do so in a way that keeps cost low (rejectability).

One key difference between epistemic utility and other utilities is that this utility orders preferences over sets of controls, rather than over the individual controls. This is an important distinction that may be exploited for profit. Before doing so, however, it is convenient to perform a positive linear transformation, resulting in the equivalent epistemic utility

$$\begin{aligned} P_e(V; \vartheta) &= P_A(V; \vartheta) - bP_R(V; \vartheta) \\ &= \int_V [f_A(u; \vartheta) - bf_R(u; \vartheta)] du, \end{aligned} \quad (8)$$

where $b = \frac{(1-\beta)}{\beta}$.

The epistemic utility function, P_e , is parameterized by the state of nature, ϑ . Since we view the state of nature as a random variable, θ , we may take expectations over Θ to obtain the expected epistemic utility:

$$\begin{aligned} \bar{P}_e(V) &= \int_{\Theta} \int_V [f_A(u; \vartheta) - bf_R(u; \vartheta)] f_C(\vartheta) d\vartheta du, \\ &= \int_V \bar{f}_A(u) - b\bar{f}_R(u) du, \end{aligned} \quad (9)$$

where

$$\bar{f}_A(u) = \int_{\Theta} f_A(u; \vartheta) f_C(\vartheta) d\vartheta \quad (10)$$

$$\bar{f}_R(u) = \int_{\Theta} f_R(u; \vartheta) f_C(\vartheta) d\vartheta. \quad (11)$$

Let \mathcal{F} denote the Borel field in U , and define the set

$$S = \arg \max_{V \in \mathcal{F}} \bar{P}_e(V). \quad (12)$$

If $\bar{P}_e(S) > 0$, then

$$\bar{f}_A(u) \geq b\bar{f}_R(u) \quad \forall u \in S; \quad (13)$$

that is, the accuracy of the control u is at least as great as its rejectability (after scaling by b). This likelihood ratio constitutes a reasonable minimum threshold for characterizing u as being satisfying. Thus, S constitutes the set of satisfying controls; we term S the *maximal satisfying set*. This decision rule is termed *Levi's rule of epistemic utility*, suitably adapted to a control context: none of the elements of the set S should be rejected [6, p. 53].

It is easy to see that for $b \leq 1$ each element $u \in S$ has the property that (13) holds. For $b = 1$, our heuristic prescription to eliminate all controls for which rejectability exceeds accuracy is literally true. The parameter, $b \in [0, \infty)$, is the *index of rejectivity*. The smaller b becomes, the lower the satisfying threshold becomes. If $b \leq 1$ ($\beta \in [\frac{1}{2}, 1]$), then we ascribe higher epistemic utility to accuracy than to rejectability. Setting b small (close to zero) will reject very few controls. Setting $b > 1$ may result in $\bar{P}_e(S) \leq 0$, in which case S contains only those controls for which expected accuracy is least dominated by expected rejectability, that is, the controls that maximize $\bar{P}_e(S)$. The condition $b = 1$ ($\beta = \frac{1}{2}$) corresponds to equally weighting accuracy and rejectability. Experience has shown that, nominally, b should be in the neighborhood of unity; slightly increasing or decreasing its value simply alters the performance in a continuous way. The index of rejectivity thus may be viewed as a tuning parameter, similar in function to the specification of the parameters of a traditional performance index.

Strongly Satisficing Control Decisions

Once the set S is obtained from (12), any of the elements of this set are procedurally rational choices to be implemented, since they all perform above a minimum-standard threshold. If S is a singleton set, then it may be implemented without further ado. Generally, however, S will consist of more than one element, in which case, the decision rule is not complete, and a tie-breaking mechanism must be emplaced.

One possible tie-breaker is to choose randomly between the elements of S . Although arguments can be advanced that such a procedure is rational, it may also be reasoned that there is no need

to default to such an arbitrary scheme, and that additional considerations should be explored. For example, two satisficing controls may have nearly the same accuracy but quite different rejectabilities, in which case it would seem prudent to implement the one with the lower rejectability since it can be done with a minimal accuracy reduction and considerable cost savings. This thought leads us to consider only those controls that enjoy a certain equilibrium property. If a control u is perturbed slightly, say by δu , then the accuracy and rejectability values for the new control, $u + \delta u$, will also change. The control will be in a state of *epistemic equilibrium* if, as a result of the small perturbation, accuracy cannot be increased without also increasing rejectability, and rejectability cannot be decreased without also decreasing accuracy. The set of satisficing controls that are in a state of epistemic equilibrium is termed the *strongly satisficing set* [20].

As examples of strongly satisficing controls, we form an analogy to Neyman-Pearson theory, and consider specifying a rejectability level, $\rho \in (0, 1)$ (analogous to size in the Neyman-Pearson context), and implementing the satisficing control that maximizes the accuracy (analogous to power) subject to the constraint that the rejectability does not exceed ρ . Let

$$S_\rho = \{u \in S : \bar{f}_R(u) \leq \rho\}.$$

The *most accurate control of rejectability* ρ is then

$$u_\rho = \arg \max_{u \in S_\rho} \bar{f}_A(u). \quad (14)$$

If $S_\rho = \emptyset$, then a most accurate control of rejectability ρ does not exist. Alternatively, one could implement the satisficing control that minimizes rejectability subject to an accuracy constraint to obtain a *least rejectable control of accuracy* α :

$$u_\alpha = \arg \min_{u \in S_\alpha} \bar{f}_R(u), \quad (15)$$

provided $S_\alpha \neq \emptyset$, where

$$S_\alpha = \{u \in S : \bar{f}_A(u) \geq \alpha\}$$

for $\alpha \in (0, 1)$. A third example of a strongly satisficing control is to maximize the difference between accuracy and rejectability, resulting in a *most discriminating control*

$$u_D = \arg \max_{u \in S} \{\bar{f}_A(u) - b\bar{f}_R(u)\}. \quad (16)$$

Since u_ρ , u_α , and u_D may not be unique, the agent may be required to select arbitrarily from the set of most accurate, least rejectable, and most discriminating controls, respectively. The above criteria represent only three of the many possible tie-breaking mechanisms that are possible with satisficing control.

A Procedurally Rational Solution for the Nonlinear Quadratic Regulator

In this section, we present a satisficing procedure for solving the general discrete-time nonlinear quadratic regulator problem.

Our approach is based on Levi's meta-epistemology, it possesses the properties of a formal procedure, and there exists a consistency theorem.

Control Problems in a Limited Context

We consider the system defined by (1), with performance index given by (2). We will assume that the control input is scalar-valued, and at each time k there is a different decision space, denoted $U_k = \{u_{0k}, u_{1k}, \dots, u_{N_k k}\}$, $k = 0, 1, \dots$. In principle, this problem is solvable via the principle of optimality, and a globally best solution exists. As a practical matter, however, analytical results are available only for special cases (for example, linear systems), and no systematic substantively rational design synthesis procedure is available for this general class of systems. Furthermore, even if a global solution were found for a given special case, the form of the solution would require prior knowledge of U_k , $k = 0, 1, \dots$. We therefore impose an additional restriction on the problem; namely, that U_k is unknown prior to time k . Thus, the problem is ill-formed.

One common method for designing nonlinear controllers is via spatial localization. Usually, a spatially local controller is designed to operate in the neighborhood of an equilibrium point. The system is then linearized about that point and linear control techniques are applied. Many problems are not amenable to spatial localization, however, and for these problems alternatives must be sought.

An alternative way to deal with nonlinear problems is via temporal localization, which may be done by invoking a receding, or rolling, horizon control strategy [21, 22]. This approach consists of implementing a feedback controller through a series of repeated open-loop calculations based on the instantaneous state. For a discrete-time receding horizon of length d , the next d values, $\{u(k), \dots, u(k+d-1)\}$, are computed as functions of the current state, $\mathbf{x}(k)$. The control $u(k)$ is implemented, producing a state $\mathbf{x}(k+1)$, the horizon is shifted forward one time unit, and the process is repeated. Temporal locality is very relevant to time-varying or stochastic problems, especially when knowledge concerning the future structure of the system equations (1) is uncertain. As with spatial localization, the justification of temporal locality is situation-dependent; in this discussion we assume that the use of a receding horizon is justified.

Accuracy and rejectability utility functions are mechanisms to implement the goals and design ideals of the problem. If information is available over the full extent of the problem, then the utilities may be designed from a global perspective. When considering ill-formed problems, however, we assume that a global solution is not available or is not implementable. For such problems we may still incorporate whatever local or global information is available to design the accuracy and rejectability utilities. By permitting the designer to tailor the structure of the utilities according to what is actually known or defensibly assumed, the problem may be cast in its natural setting. This capability frees the designer from the need to make arbitrary and possibly objectionable assumptions simply to invoke a global solution technique.

We first observe that the state of nature is the state of the system at time k , that is, $\theta = \mathbf{x}(k)$. To simplify the current development, we will assume that the state of nature is known with probability one, so $f_C(\mathbf{x}) = \delta[\mathbf{x} - \mathbf{x}(k)]$ and therefore

$\bar{f}_A(u) = f_A[u; \mathbf{x}(k)]$ and $\bar{f}_R(u) = f_R[u; \mathbf{x}(k)]$. For the nondeterministic case, we would obtain $f_C(\mathbf{x})$ as the probability density function of an estimate of $\mathbf{x}(k)$.

Once a receding horizon control strategy is adopted, our next step is to define the notions of accuracy and rejectability in this context. We may assess accuracy by constructing a utility function that takes large values for controls that achieve the fundamental goal of the controller. For the nonlinear regulator the fundamental goal of the controller is to force the terminal condition of the performance index, (3), to zero. From a local perspective, however, we may evaluate accuracy in terms of how well the global goal is met when a control is applied only over the d -length horizon. We therefore associate $k + d$ with k_f , and assess accuracy on the basis of how well the local control sequence $\{u(k) \dots u(k + d - 1)\}$ would achieve the fundamental goal of the controller if $k + d$ were the terminal time. For the quadratic regulator, this assessment is made in terms of values that

$$J_A[u(k) \dots u(k + d - 1)] = \mathbf{x}^T(k + d) P \mathbf{x}(k + d) \quad (17)$$

achieves as a function of the local control sequence. As J_A decreases, the accuracy of the control sequence $\{u(k) \dots u(k + d - 1)\}$ increases.

Rejectability has nothing directly to do with the ultimate goal of the controller. Rather, it has to do with keeping costs down. Considering the quadratic regulator and viewing the costs from a local perspective, the rejectability of a local control sequence $\{u(k) \dots u(k + d - 1)\}$ must be computed from

$$J_R[u(k) \dots u(k + d - 1)] = \sum_{j=k}^{k+d-1} [\mathbf{x}^T(j+1) Q(j+1) \mathbf{x}(j+1) + R(j) u^2(j)]. \quad (18)$$

As J_R decreases, the rejectability of $\{u(k) \dots u(k + d - 1)\} = \mathbf{x}^T(k + d) P \mathbf{x}(k + d)$ also decreases.

Calculating Accuracy and Rejectability Utilities

We present only the solution for $d = 1$ (the case for $d > 1$ may be found in [13]; we therefore associate k_f with $k + 1$. Under this scenario, (17) and (18) become

$$J_A(u) = \mathbf{f}^T[\mathbf{x}(k), u, k] P \mathbf{f}[\mathbf{x}(k), u, k] \quad (19)$$

$$J_R(u) = \mathbf{f}^T[\mathbf{x}(k), u, k] Q(k + 1) \mathbf{f}[\mathbf{x}(k), u, k] + R(k) u^2. \quad (20)$$

We first consider the accuracy utility. To obtain maximum dynamic range in this utility, we bias J_A to assign zero utility to the controls that have the largest values of J_A , obtaining

$$J'_A(u) = \max_{z \in U_k} J_A(z) - J_A(u). \quad (21)$$

The accuracy utility density, \bar{f}_A , may then be calculated by normalizing this quantity:

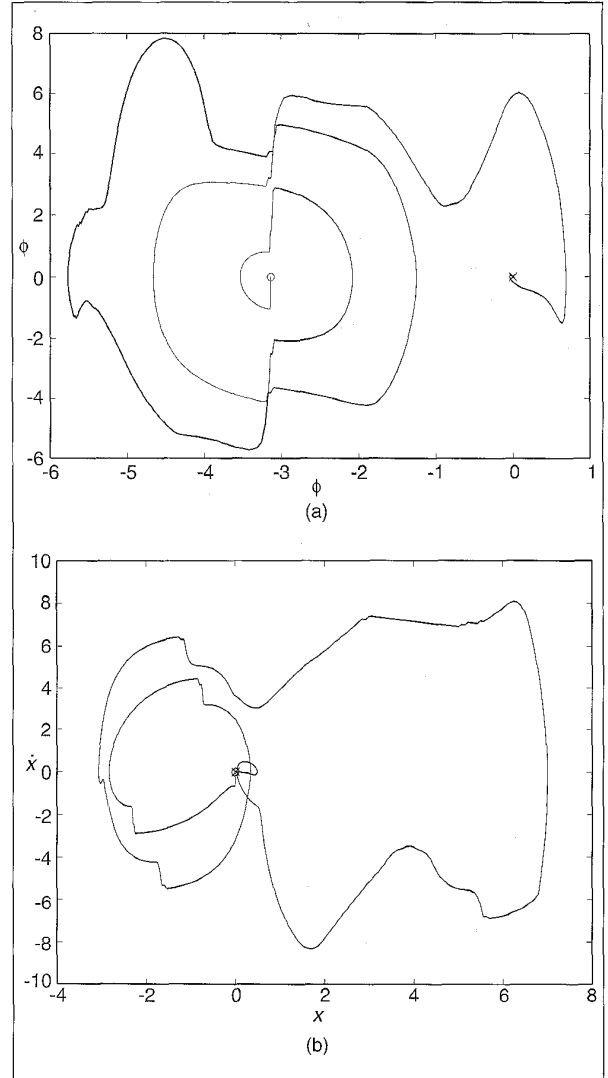


Fig. 2. Phase planes for the inverted pendulum on a cart: (a) rotational phase plane (in radians and radians per second), (b) translational phase plane (in meters and meters per second).

$$\bar{f}_A(u) = \frac{J'_A(u)}{\sum_{u \in U_k} J'_A(u)}. \quad (22)$$

Next, consider the rejectability utility. To obtain maximum dynamic range, we must bias J_R to assign zero utility to controls that have the smallest values of J_R , yielding

$$J'_R(u) = J_R(u) - \min_{z \in U_k} J_R(z). \quad (23)$$

The rejectability utility density, \bar{f}_R , may then be calculated as

$$\bar{f}_R(u) = \frac{J'_R(u)}{\sum_{u \in U_k} J'_R(u)}. \quad (24)$$

Once these functions are obtained, we simply apply (13) to obtain the maximal satisficing set, S_k , of controls that may be applied at time k :

$$S_k = \{u \in U_k : \bar{f}_A(u) - b\bar{f}_R(u) \geq 0\}, \quad (25)$$

where b has been chosen such that $S_k \neq \emptyset$. Each of these controls meets or exceeds the threshold of minimum standard, so all are good, or satisficing, candidates. As a tie-breaker, we invoke the most discriminating satisficing control, u_D . Applying this procedure at each time-increment, k , we obtain a nonlinear feedback control policy. The following theorem establishes consistency [20].

Theorem 1. *For the quadratic regulator problem (1) with quadratic cost (2) and $U = (-\infty, \infty)$, if an optimal control exists and the control horizon spans the full extent of the problem, that is, $d = k_f$, then there exists a rejectivity $b \in (0, \infty)$ such that the most discriminating control sequence,*

$$u_D = \arg \max_{z \in U^d} \{\bar{f}_A(z) - b\bar{f}_R(z)\}, \quad (26)$$

is identical to the optimal control.

Proof: We will prove this result for the scalar control case only. From (22) and (24)

$$\arg \max_{z \in U^d} \{f_A(z) - bf_R(z)\} = \arg \max_{z \in U^d} \{J'_A(z) - b'J'_R(z)\}, \quad (27)$$

where $b' = b \frac{\sum_{\mathbf{u} \in U^d} J'_A(\mathbf{u})}{\sum_{\mathbf{u} \in U^d} J'_R(\mathbf{u})}$. From (21) and (23),

$$\begin{aligned} \arg \max_{z \in U^d} \{J'_A(z) - b'J'_R(z)\} &= \arg \max_{z \in U^d} \{-J_A(z) - b'J_R(z)\} \\ &= \arg \min_{z \in U^d} \{J_A(z) + b'J_R(z)\}. \end{aligned} \quad (28)$$

Using (17) and (18), it is now immediate that, with $k = 0$ and

$d = k_f$,

$u_D =$

$$\arg \min_{z \in U^{k_f}} \left\{ \mathbf{x}^T(k_f) P \mathbf{x}(k_f) + b' \sum_{j=0}^{k_f-1} \mathbf{x}^T(j+1) Q(j+1) \mathbf{x}(j+1) + R(j) z^2(j) \right\}. \quad (29)$$

For $b = \frac{\sum_{\mathbf{u} \in U^d} J'_R(\mathbf{u})}{\sum_{\mathbf{u} \in U^d} J'_A(\mathbf{u})}$, we choose $b' = 1$, and (29) is exactly the

optimal quadratic regulator solution.

Nonlinear Regulator

To demonstrate the operation of our approach, we stabilize a pendulum on a cart while *simultaneously* regulating the cart's position as well. Consider a pendulum in a vertical plane with full circular freedom of motion, as illustrated in Fig. 1. The problem is to bring the pendulum from an arbitrary initial condition to a vertical-up orientation while positioning the cart at a desired point. The control consists of a lateral force applied directly to the cart [23, p. 194]. This nonlinear control problem is difficult since there are two degrees of freedom but only one

input. The control limit at each time k will be computed as a random walk initialized at the endpoints of nominal interval $(-u_m, u_m)$, with a reflecting boundary at $\frac{u_m}{2}$.

The details, including the dynamical equations and the development of the accuracy and reliability utilities for this problem, are in [13, 20]. Figs. 2(a) and 2(b) illustrate the rotational (pendulum) and translational (cart) phase planes obtained by the application of the controller. The "o" symbol represents the initial conditions (the cart at the origin with the pendulum in the vertical-down position) and the "x" symbol represents the terminal conditions (the cart at the origin with the pendulum balanced in the vertical-up position). The system achieves its desired objective of balancing the pendulum at the origin by swinging the pendulum back and forth while the cart oscillates around the origin. As the cart oscillates, the pendulum gathers momentum. In the translational and rotational phase planes, this motion is manifest as growing spirals. When the amplitude increases sufficiently, the oscillation ceases and the pendulum then converges to the vertical-up position. Finally, the cart returns slowly to the origin.

Extensions to Multiple Agents

In multiple-agent control scenarios, the interplay between distinct decision-making agents can be complex. The traditional approach to this problem is game theory, with the goal being to determine the Nash equilibrium, thus solving the game and specifying the optimal decisions for all agents, or players. Traditional game theory, however, is predicated on two assumptions. The first assumption is that all of the players know all logical consequences of their assumed knowledge—the principle of *hyperrationality*. The second assumption is that all players know *only* the logical consequences of their assumed knowledge—the principle of *indeterminacy* [24]. If both of these principles apply, then a condition of logical closure obtains, and a well-formed joint control problem can be formulated and a substantively rational control decision is possible. If either of these assumptions fails, however, the problem is ill-formed.

The satisficing theory we have thus far developed is for the single-agent case, where only one decision-maker is present. We may extend this theory to the multiple-agent case by constructing a *coordination* function, which is a utility characterizing the joint accuracy and rejectability of all agents. Let X_1 and X_2 be agents with decision spaces U and V , respectively, and define the coordination function $Coord: U \times U \times V \times V \rightarrow [0, 1]$ as a joint probability density function

$$Coord(u_a, u_r, v_a, v_r) = \bar{f}_{A_1 R_1 A_2 R_2}(u_a, u_r, v_a, v_r). \quad (30)$$

Once the coordination function has been defined, a joint satisficing decision rule may be implemented by generalizing the single-agent rule (13). To express this rule, we first must isolate the joint accuracy and rejectability functions:

$$\begin{aligned} \bar{f}_{A_1 A_2}(u_a, v_a) &= \int_U \int_V Coord(u_a, u_r, v_a, v_r) du_r dv_r \\ \bar{f}_{R_1 R_2}(u_r, v_r) &= \int_U \int_V Coord(u_a, u_r, v_a, v_r) du_a dv_a. \end{aligned}$$

The joint satisficing set is then given by

$$S = \{(u, v) : \bar{f}_{A_1 A_2}(u, v) \geq b \bar{f}_{R_1 R_2}(u, v)\}, \quad (31)$$

yielding a set of satisficing control decisions for the agents. If the agents possess identical models, they will each arrive at the same satisficing set. The notion of strongly satisficing also generalizes to the multi-agent case (see [20]).

As an example, consider the well-known Prisoner's Dilemma game, in which two players suspected of a crime are interrogated separately. If both remain silent, they will both be fined and released. If only one confesses, the one who does so goes free, with the other receiving the maximum jail sentence. If both confess, each receives a mild jail sentence. The Nash equilibrium solution, in which both players confess, is inferior to the cooperative solution, in which both remain silent, thus the dilemma.

One interpretation of this dilemma is to assert that it exists because agents are disposed to consider various metatheoretical principles, such as dispositions toward benevolence, morality, masochism, and other such parameters [25, 26]. Thus, the players may be influenced by information other than the logical consequences of the knowledge provided by the penalty structure—a condition of indeterminacy exists. Since the game is not logically closed, the efficacy of the substantively rational solution (the Nash equilibrium solution) is in dispute.

In [18], this game is recast in an epistemic utility setting. Accuracy is interpreted in terms of jail-time; the shorter the time served, the higher the accuracy. Rejectability is used to express the metatheoretical aspects of the game, specifically, dispositions toward self-defense, egotism, masochism, and collusion. These parameters are injected into the game via the coordination function.

Since *Coord* is a probability density function, we may invoke the notions of independence and conditioning to synthesize a model of the system under study. For example, we may express the coordination function as the product of conditional density functions (conditional utilities of X_1 conditioned on X_2 and the densities associated with X_2):

$$\begin{aligned} \text{Coord}(u_a, u_r, v_a, v_r) &= f_{A_1 R_1 | A_2 R_2}(u_a, u_r | v_a, v_r) f_{A_2}(v_a) f_{R_2}(v_r) \\ &= f_{A_1 | R_1 A_2 R_2}(u_a | u_r, v_a, v_r) f_{R_1 | A_2 R_2}(u_r | v_a, v_r) f_{A_2}(v_a) f_{R_2}(v_r). \end{aligned}$$

The conditional density function $f_{A_1 | R_1 A_2 R_2}(u_a | u_r, v_a, v_r)$ represents X_1 's accuracy ascribed to u_a given that X_2 places all of its accuracy mass on v_a and all of its rejectability mass on v_r , and X_1 places all of its rejectability on u_r . The conditional density function $f_{R_1 | A_2 R_2}(u_r | u_a, v_a, v_r)$ represents X_1 's rejectability given that X_2 places its accuracy mass on v_a and its rejectability mass on v_r . For example, consider the condition that X_2 places all of its accuracy on $v_a = \textit{confession}$ and all of its rejectability on $v_r = \textit{silence}$. This situation would mean that X_2 is in a state of tension: it wants to minimize jail-time, but it refuses to remain silent, a necessary condition for minimum jail-time. The conditional utility value $f_{R_1 | A_2 R_2}(\textit{silence} | \textit{confession}, \textit{silence})$ represents X_1 's propensity for rejecting *silence* that X_2 adopts the given stance. Using this epistemic utility-based model, we may systematically implement the agents' predispositions and reasoning processes to construct the coordination function. The decision rule given by (31) illuminates both the logical and the metatheoretical

aspects that influence the decision. For a complete discussion of this approach, see [18].

Discussion

One of the great benefits of conventional control theory is the design synthesis methodology that has arisen from solutions to optimal control problems. A systematic synthesis procedure is at least as important as is an idealized notion that the very best solution is being obtained. For an alternative control methodology to receive acceptance, it must provide a systematic design synthesis procedure that is perceived as being reliable and reasonable. If the procedure is not based on utility maximization, then it must at least conform to logical and reasonable criteria. We have presented specific criteria for procedural rationality: the synthesis procedure (i) must be epistemologically defensible, (ii) must be amenable to systematic design synthesis, and (iii) must be consistent.

To satisfy the first of these criteria, we have adopted Levi's meta-epistemology and have drawn a fundamental analogy between the cognitive context and the practical context by likening truth-seeking to optimality and error-avoidance to satisficing. The satisficing constraints are met through the minimum standard made explicit with Levi's rule of epistemic utility.

To address the second criterion we have introduced the accuracy and rejectability utilities and have shown how to obtain them systematically from the performance index. Using these utilities, Levi's rule provides a natural decision engine for obtaining a set of satisficing controls. To refine the set, the notion of strong satisficing is introduced and specific decision rules are presented.

A satisfactory treatment of the third criterion requires the development of a body of consistency theory for each type of problem being treated. In this essay we have provided a consistency theorem for the general nonlinear time-varying quadratic regulator. It is conjectured, though not yet established, that this result extends to a larger class of performance indices, and to a larger class of control problems, such as the general nonlinear tracking problem.

[1] H.A. Simon, "Rationality in Psychology and Economics," in R.M. Hogarth and M.W. Reder, *Rational Choice*, University of Chicago Press, 1986.

[2] R.M. Hogarth and M.W. Reder, *Rational Choice*, University of Chicago Press, 1986.

[3] H.A. Simon, "A Behavioral Model of Rational Choice," *Quart. J. Economics*, 59, 1955, 99-118.

[4] A. Meystel, *Autonomous Mobile Robots*, World Scientific, Singapore, 1991.

[5] *Epistemic Justification*, Cornell University Press, Ithaca, New York, 1989.

[6] I. Levi, *The Enterprise of Knowledge*, MIT Press, Cambridge, Massachusetts, 1980.

[7] W. James, *The Will to Believe and Other Essays*, Dover, 1956.

[8] H.A. Simon, "Theories of Decision-Making in Economics and Behavioral Science," *The American Economic Review*, XLIX, 1959, 253-283.

[9] I. Levi, *Gambling with Truth*, MIT Press, Cambridge, Massachusetts, 1967.

[10] I. Levi, *Decisions and Revisions*, Cambridge University Press, London, 1984.

[11] I. Levi, *Hard Choices*, Cambridge University Press, Cambridge, 1986.

[12] J.D. Kenney, *Stochastic Satisficing Control*, Ph.D. dissertation, Brigham Young University, to appear, 1996.

[13] M.A. Goodrich, W.C. Stirling, R.L. Frost, and J.D. Kenney, A Theory of Satisficing Decision and Control, submitted for publication.

[14] W.C. Stirling and D.R. Morrell, "An Epistemology of Bayes Estimation," *Conference Record of Twenty-Third Asilomar Conference on Signals, Systems, and Computers*, October 1989, 699-703.

[15] W.C. Stirling and D.R. Morrell, "Convex Bayes Decision Theory," *IEEE Transactions on Systems, Man, and Cybernetics*, January/February, vol. 21, no. 1, 1991, 173-183.

[16] W.C. Stirling, "Coordinated Intelligent Control Via Epistemic Utility Theory," *IEEE Control Systems Magazine*, vol. 13, no. 5, October 1993, 21-29.

[17] T.K. Moon, S.E. Budge, W.C. Stirling, and J.B. Thompson, "Epistemic Decision Theory Applied to Multiple-Target Tracking," *IEEE Trans on Systems, Man, and Cybernetics*, vol. SMC-24, no. 2, 1994, pp. 234-245.

[18] T.K. Moon, R.L. Frost, and W.C. Stirling, "An Epistemic Utility Approach to Coordination in the Prisoner's Dilemma," *BioSystems*, vol. 37, nos. 1-2, 1996, pp. 167-176.

[19] R.L. Frost and W.C. Stirling, "An Architecture for Intelligent Control Based on Epistemic Utility Theory," *Proceedings of the 1994 American Control Conference*, June-July 1994, pp. 2580-2584.

[20] M.A. Goodrich, *On a Theory of Satisficing Control*, Ph.D. dissertation, Brigham Young University, to appear, 1996.

[21] D.Q. Mayne and H. Michalska, "Receding Horizon Control of Nonlinear Systems," *IEEE Transactions on Automatic Control*, vol. 35, 1990, pp. 814-824.

[22] O. Hernandez-Lerma and J.B. Lasserre, "Error Bounds for Rolling Horizon Policies in Discrete-Time Markov Control Processes," *IEEE Transactions on Automatic Control*, vol. 35, 1990, pp. 1118-1124.

[23] J.J.E. Slotine and W. Li, *Applied Nonlinear Control*, Prentice Hall, Englewood Cliffs, NJ, 1991.

[24] M. Bacharach, "The Epistemic Structure of a Theory of a Game," *Theory and Decision*, vol. 37, 1994, pp. 7-48.

[25] G. Wolf and M. Shubik, "Concepts, Theories and Techniques: Solution Concepts and Psychological Motivations in Prisoner's Dilemma Games," *Decision Sciences*, 1974, vol. 5, pp. 153-163.

[26] F.M. Frohock, *Rational Association*, Syracuse University Press, Syracuse, NY, 1987.



Wynn C. Stirling received a Ph.D. in electrical engineering from Stanford University, Stanford, CA, in 1983. He received an M.S. degree in electrical engineering and an Honors M.A. degree (magna cum laude) in mathematics from the University of Utah, Salt Lake City, UT, in 1971 and 1969, respectively. From 1972 to 1975 he was with Rockwell International Corp., Anaheim, CA, and from 1975 to 1984 he was employed by ELS Inc., Sunnyvale, CA. Since 1984 he has been with Brigham Young University, where he is a professor in the Department of Electrical and Computer Engineering. His research interests include decision theory, control theory, estimation theory, and stochastic processes. Dr. Stirling is a member of Phi Beta Kappa.



Michael A. Goodrich received an M.S. degree and a B.S. degree (cum laude) in electrical engineering from Brigham Young University, Provo, UT, in 1995 and 1992, respectively. He is currently completing his Ph.D. degree in electrical engineering at Brigham Young University. His research interests include intelligent control, decision theory, control theory, multiple-agent coordination, estimation theory, and fuzzy logic.



Richard L. Frost received a B.S. in physics (1975, magna cum laude), an MSEE (1977), and a Ph.D. (1979) in electrical engineering, all from the University of Utah. He was at the MIT Lincoln Laboratory from 1979-1981, on the faculty of the University of Utah from 1981-1984, and with the Communication Systems Division of Sperry (now Loral) from 1984-1987, when he joined the faculty at Brigham Young University. He is presently an associate professor in the Department of Electrical and Computer Engineering, with principal research interests in quantization and source coding and in intelligent control.