Satisficing Revisited

Michael A. Goodrich,* Wynn C. Stirling,† and Erwin R. Boer‡

June 11, 1999

Abstract

In the debate between simple inference heuristics and complex decision mechanisms, we take a position squarely in the middle. A decision making process that extends to both naturalistic and novel settings should extend beyond the confines of this debate; both simple heuristics and complex mechanisms are cognitive skills adapted to and appropriate for some circumstances but not for others. Rather than ask "Which skill is better?" it is often more important to ask "When is a skill justified?" The selection and application of an appropriate cognitive skill for a particular problem has both costs and benefits, and therefore requires the resolution of a tradeoff. In revisiting satisficing, we observe that the essence of satisficing is tradeoff. Unlike heuristics, which derive their justification from empirical phenomena, and optimal solutions, which derive their justification by an evaluation of alternatives, satisficing decision-making derives its justification by an evaluation of consequences. We formulate and present a satisficing decision paradigm that has its motivation in Herbert Simon's work on bounded rationality. We characterize satisficing using a cost-benefit tradeoff, and generate a decision rule applicable to both designing intelligent machines as well as describing human behavior.

*Computer Science Department, Brigham Young University, Provo, UT, and Nissan Cambridge Basic Research, Cambridge, MA.
†Electrical and Computer Engineering Department, Brigham Young University, Provo, UT.
‡Nissan Cambridge Basic Research, Nissan Research and Development, Inc., Cambridge, MA.

1 Introduction

While driving an automobile, many of us have experienced something similar to the following. We have been following a vehicle for an extended period of time even though there is very little traffic on the road. Suddenly, we realize that not only can we easily pass but also that we want to pass because the lead vehicle is going slower than our desired speed. We decide to pass, and act accordingly.

What are the factors that dictate our behavior in this situation? Can a characterization of the corresponding behavior-generation process be used to design better machines? This paper is written from two perspectives: first, from the perspective of a designer charged with the task of creating a machine capable of some degree of goal-directed autonomy and agency, and second from a perspective of describing goal-directed human behavior generation. Inherent in the resolution of these problems is the need to resolve tradeoffs. For machine intelligence, who resolves the tradeoffs, the designer or the machine? For humans, who resolves the tradeoffs, the human or societal/evolutionary forces?

There appear to be two disparate approaches to solving these problems. The first approach, a “top-down” approach, contends that intelligence is tantamount to normative rationality and optimality. Representatives from this approach cite success in the philosophical foundations of cognitive science (a la psychology) and success in optimal decision theory (a la design) as evidence for a top-down description of intelligence. The second approach, a “bottom-up” approach, contends that intelligence emerges from ecologically adapted behavioral and cognitive skills. Representatives from this approach cite evidence from the usefulness of cognitive heuristics (Gigerenzer and Goldstein, 1996) (a la psychology) and the success of, for example, ecological robotics (Brooks, 1986) (a la design). In both design and description, the top-down approach tends to rely on complex decision mechanisms whereas the bottom-up approach tends to rely on simple inference heuristics. A marriage between these extremes is necessary to ensure behavior that achieves a goal subject to environmental constraints.

1.1 Problem Statement

From a machine intelligence perspective, goal-directed decision-makers capable of situated and continued existence must have the ability to self-police their behaviors. Included in self policing are the abilities to evaluate and anticipate performance internally, and the ability to resolve decision tradeoffs internally. These abilities can be accomplished to a limited extent by allowing a designer to specify complex decision mechanisms capable of handling all but the most subtle (and possibly most treacherous) situations. Alternatively, the designer can specify or identify simple inference heuristics and allow the machine to select among these heuristics as afforded by the environment. From a designer's perspective, the latter approach has the ability to scale to larger domains and is thus a useful approach but, unfortunately, this approach begs the question of how these heuristics are systematically created and managed.

Switching attention to descriptions of human decision making, the distinction between simple inference heuristics and complex decision mechanisms appears to be somewhat artificial. Rational people can use either simple heuristics or complex mechanisms depending on which is more appropriate for the circumstances, and it is an open question as to how people select and obtain these skills. Both simple inference heuristics and complex decision mechanisms are *cognitive skills*, and a rational person naturally (either through instinctual responses, responses learned through external feedback, or responses learned through goal-directed internal feedback) employs tradeoffs and expected performance associated with each skill and chooses appropriately. The question of which skill is <u>more</u> correct is misguided because, from an agent's point of view, the environmentally afforded "means" to reaching the decision are subjected to the goal-contextualized "ends" produced by the decision, and any approach that efficiently uses means to generate productive ends is justifiable. Extending this thought, prescriptive approaches to decision-making should permit either simple heuristics or complex mechanisms provided that the expected result is good enough.

1.2 Solution Approach

Cognitive skills can be treated as agents and organized into a society of Minskian agents (Minsky, 1986). Management of these skill-based agents is tantamount to a meta decision problem that requires an appropriate notion of rationality. By framing the problem as one of skill management, the decision maker formulates a control problem wherein, given certain goals and a certain context, the decision maker controls which cognitive skill agent operates. This control problem is addressed by a meta agent, whence the problem becomes one of coordinating agents in a multi-agent society. Multi-agent societies used to generate rational decisions that use cognitive skills require meta-choices which serve to resolve tradeoffs and assure rational agency.

An appeal to meta-rationality to settle a question of rationality is always risky. Too often, such appeals result in an endless chain of "how do I know that I know that I know ..." Fortunately, if meta-choices are justifiable (from a prescriptive perspective) and produce a useful decision rule (from a descriptive perspective) then such an infinite regression can be avoided. Although in a prescriptive/design sense it may be desirable, such meta-rationality need not be explicitly possessed by the agent, but can instead be (and often is) imposed externally by a designer or through evolutionary forces. Our objective is to identify a decision rule that, from a descriptive perspective, is a useful heuristic in the spirit of Simon's notion of satisficing and that, from a prescriptive perspective, can be justified by an appeal to meta-rationality. In the end, we present a mathematical characterization of satisficing, discuss how Simon's original notion is compatible with this characterization, and describe how this characterization is manifest in observations of human decision making (including the car-following example).

1.3 Outline

This paper is organized as follows. In Section 2, we describe the elements of a decision problem and discuss the limits of both optimality and heuristics with an emphasis on justifiability and practicability. In Section 3, we discuss a decision mechanism for resolving tradeoffs and present a decision theoretic characterization of satisficing. In Section 4, we extend this characterization of satisficing decision making to include the interaction between two independent decision forces and the resulting coordination of Minskian agents. Then, in Section 5, we discuss the implications of this satisficing decision paradigm in the context of the debate between simple inference heuristics and complex decision mechanisms.

2 Elements of Decision Making

***** INSERT FIGURE 1 ABOUT HERE ******

The elements of a decision problem are diagrammed in Figure 1. Given an observation $x \in X$ that is a function of the state of nature $\theta \in \Theta$, the decision task is to select an option $u \in U$ that produces acceptable (according to values and preferences) consequences. In decision making, there are two conventional approaches: complex decision mechanisms based on seeking *superlative* decisions using normative rationality, and simple inference heuristics based on seeking *positive* decisions using empirically derived procedures. Superlative approaches seek to identify options U, estimate states Θ from sensory observations X, determine consequences using some causal model, and then extremize some performance metric that imposes a preference pattern on these consequences. By contrast, positive approaches short circuit some of these stages resulting in, for example, rules of the form "if x then u." The optimality-based literature, particularly that of optimal control theory and game theory, is overwhelmingly vast, reflecting many decades of serious research and development of ideas based on the superlative paradigm. The positive paradigm, manifest in the form of heuristics, procedurally rational decision making methods, and multitudinous *ad hoc* techniques, has also been well-represented in the computer science, social science, and engineering literatures.

There are alternatives to the superlative and positive paradigms. The most well known example of this *comparative* paradigm is Simon's notion of satisficing (Simon, 1996; Simon, 1955). A formally stated comparative paradigm, however, has not been well represented in the literature as a basis for a viable decision-making concept for general application. In this section, we first review the superlative and positive paradigms, and then discuss Simon's notion of satisficing to establish a foundation for our subsequent revisitation of satisficing. In the following subsections, we refer to the utility of accepting a decision and a utility of rejecting a decision. This discussion includes probabilistic inference as a special case where the utility of accepting a decision is unity if the decision is correct, and zero otherwise. Additionally, this allows us to treat optimality as the typical problem in normative rationality without loss of generality.

2.1 Superlative Rationality: Optimal Decisions

When estimates of x and/or θ are distributed according to a known probability distribution, then a decision problem is said to be one of decision under *risk* (Luce and Raiffa, 1957). The conventional approach to decisions under risk is to define a utility function for each of the consequences and then select an option that produces the maximum expected utility (where the expectation is taken with respect to the distribution of states of nature). The option that maximizes expected utility is the optimal option u^* defined as

$$u^* = \arg\max_{u \in U} \sum_{\theta} v(u, \theta) p(\theta | x) \tag{1}$$

where $v(u, \theta)$ is the utility of selecting option u given state θ, and $p(\theta|x)$ is the probability density function for θ given observation x. By contrast to decisions under risk, when the probabilities of x and θ are completely unknown then the decision is said to be one of decision under *uncertainty* (Luce and Raiffa, 1957). The conventional approach to decisions under uncertainty is to use a maximin approach yielding

$$u^* = \arg\max_{u \in U} \min_{\theta \in \Theta'} v(u, \theta) \tag{2}$$

where Θ' is the set of feasible states given observation x. The function $\min_{\theta \in \Theta'} v(u, \theta)$ is called the security level for u and can be interpreted as an expectation with respect to a least favorable distribution of θ given x. Therefore,

u^* is interpreted as the option that maximizes security.

These methods have been tremendously successful for certain applications. However, not all decision problems are optimization problems, nor should they be. Recalling a proverb from control theory, performing a study of nonlinear control problems is analogous to performing a study of "non-elephant animals"; there are simply many more nonlinear control problems than linear control problems. Similarly, there are many problems addressable by "non-optimal" approaches that are not amenable to optimal approaches. Although some may categorize the practice of "non-optimal" choice as a species of irrationality, the quest for the successful development of intelligent machines rests, to some degree, upon the assumption that intelligence extends beyond naive optimization (Slote, 1989). Rationality is not tantamount to optimality.

2.2 Bounded Rationality: The Presence of Tradeoffs in a Superlative World

Many cognitive scientists recognize that insistence on optimality is a misplaced requirement in situations of limited resources and information, and that optimality inadequately describes observed behavior in naturalistic settings (Gigerenzer and Goldstein, 1996; Zsambok and Klein, 1997). For complex problems, there often exist information, memory, or computing limitations such that finding a strictly optimal solution is not feasible because (1) and (2) must be formulated and solved. Under these circumstances, a principle of *bounded rationality* is often recommended. Many such theories are based on Simon's well-known satisficing idea wherein a decision-maker uses "experience to construct an expectation of how good a solution we might reasonably achieve, and halting search as soon as a solution is reached that meets the expectation" (Simon, 1990, Page 9). Satisficing thus becomes a means of addressing when an option is "good enough" in the sense that its utility exceeds an aspiration level. Determination of an aspiration level is based on experience-derived expectations of possible consequences, and a search algorithm is proposed that is compatible with limited computational resources and that terminates when an option is identified that exceeds the aspiration level.

Dissatisfied with this under-specified algorithm, some researchers have proposed other satisficing-like notions of bounded rationality such as augmenting the utility function with computational costs. Such methods are closely related to constrained optimization (see, for example, (Sandholm and Lesser, 1997; Zilberstein, 1996; Kaufman, 1990)), and yield optimal solutions according to a modified criterion. These algorithms appear to abandon Simon's original intention of comparing predicted consequences with expected potential consequences to justify good enough decisions. Instead, these procedures derive their justification by an appeal to optimality with respect to a modified performance criterion. Since no mention is made of how a situated decision maker might choose such a criterion, it appears that advocates of such approaches transform satisficing from a consequence-based justification to a procedure-based justification and thereby make a "virtue out of a necessity" (Levi, 1997, Page vii). However, insofar as justification can be derived through the process of optimization, these approaches are compatible with the goals of the proponents.

Regardless of the details of how a boundedly rational decision is obtained, it is clear that the ultimate rationale for adopting a decision obtained in such a way is that it is the resolution to a tradeoff between the goal-directed capabilities of a decision-maker and the environmental affordances relevant to that goal.

2.3 Positive Rationality: Heuristics

Once defined, approaches based on both optimality and non-Simon-like bounded rationality find the best possible solution (according to an implicit performance metric) given context-dependent constraints and imprecise information about the true state of nature. For real environments, a decision maker must also be able to determine not only the set of possible options U (the search space), but also something about the utility $v(u, \theta)$ of taking action[1] as well as the set of relevant states Θ. This can lead to intractable complexity, especially for the designers of machines. For example, control engineers sometimes use an explicit model to predict the consequences of a sequence of actions using a method termed "model predictive control" (Michalska and Mayne, 1995; Sistu and Bequette, 1996; Richalet, 1993; Mayne and Michalska, 1990; Scokaert et al., 1997)). The extent of the action sequence can be adjusted according to a receding planning horizon, and must often be very limited because of the combinatorial complexity of enumerating multiple action sequences. Often, when faced with such increasing complexity, the designer must resort to heuristics (consider the success of heuristic search techniques).

In effect, *heuristics are empirically derived cognitive shortcuts in a decision problem.* For example, under particular sensory influences x a decision maker might use the rule *if x then u*. A criticism of the use of heuristics is that they are unjustifiable and lead to capricious results because they are essentially *ad hoc* in nature (Kahneman and Tversky, 1996). *Ad hoc* procedures while producing good (maybe even very good) decisions, will not produce decisions that can be reliably established as being adequate in terms of performance, but are instead based on vague notions of desirability or convenience without any definitive measures of quality. Fortunately, some heuristics appear to be ecologically adapted to certain niches, and work is proceeding on identifying these niches and comparing the behaviors produced by these heuristics to more conventional approaches (see, for example, (Chase et al., 1998)).

The appropriate use of heuristics in machines and humans can increase capacity and can help generate solutions to non-optimal decision problems (Ho, 1999).

2.4 Interlude

The use of non-optimal decision mechanisms need not result in *ad hocism*. For example, Lotfi Zadeh, the father of fuzzy logic, can undoubtedly be included as someone who is interested in exploring non-optimal but justifiable choice. Near the beginning of his career he wrote an essay entitled "What is optimal?" (Zadeh, 1958) and four decades later revisited the theme in his paper "Maximizing Sets and Fuzzy Markoff Algorithms" (Zadeh, 1998). In these papers, Zadeh questions the feasibility (and wisdom) of seeking for optimality given limited resources. However, in resisting naive optimizing Zadeh does not abandon the quest for justifiability, but instead resorts to modifications of conventional logic that are compatible with linguistic and fuzzy understanding of nature and consequences. Other researchers, including many who have contributed to the area of optimal decision and control, have explored non-optimal but justifiable solution methodologies as exemplified in work in suboptimal decision making, ordinal optimization (Ho, 1994; Ho and Larson, 1995), probably approximately correct algorithms (Greiner and Orponen, 1996), multi-resolutional intelligence (Albus, 1991; Meystel, 1996), heuristic search, behavior-based/ecological robotics (Brooks, 1986; Brooks, 1991; Duchon et al., 1998), anytime algorithms (Zilberstein, 1996), and satisficing decision-making (Simon, 1996; Sen, 1998). It is interesting that each of these approaches seeks to resolve a tradeoff between the ultimate behavior of the agent or system and the practicable methods for generating this behavior.

2.5 Comparative Rationality: Being "Good Enough"

The notion of being "good enough" is an underlying issue in all decision problems and is an inseparable companion to the notion of a tradeoff. For example, under Simon's satisficing, rejecting an option that does not meet or exceed the aspiration level derives its justification from the observation that the option is rejected in favor of an unknown alternative that produces better consequences; we trade the would-be consequences of the rejected option for the expected consequences of an unidentified option. In machine intelligence, ensuring good enough performance has conventionally been the responsibility of the designer. By contrast, in human intelligence ensuring good enough performance is either the responsibility of the human or, in a much broader sense, the responsibility of the species subject to evolutionary forces. For an individual human, evaluating success in goal-directed behavior requires rational self policing.

Restricting attention to goal-directed behavior, self-policing becomes very important. Self policing must include the ability to determine if a behavior produces good enough consequence and, if not, change or adapt behaviors. As part of this evaluative phase, a decision-maker may need to identify feasible alternatives, coherent beliefs, and consistent values. Another aspect of self-policing is the ability to, in the spirit of Simon's expectation-based aspiration level, anticipate the efficaciousness of an option. Regardless of whether heuristic or optimal, self-policing is essential for robust goal-directed behavior generation. Self-policing allows a decision-maker to evaluate and adapt (possibly context-dependent) "means" subject to (possibly task-specific) "ends" in an effort to produce good enough performance.

Recall the example of following a vehicle for an extended period of time even though passing it is a superior alternative. Unless a driver is a voracious optimizer capable of limitless attention, few would say that the behavior is irrational (although we reflect on the situation with mild amusement). The point is that being good enough is required, and being optimal is optional.

3 Satisficing and Tradeoff

Too often, in a quest to impart intelligence to a machine we resort to one of two extremes. We either require the designer to have sufficient expertise to identify and encode a simple and effectual task-specific algorithm, or to determine and encode a complex context-free algorithm responsible for solving any and all task-specific problems. Similarly, in an effort to describe and prescribe human behavior we often resort to one of these extremes. Thus, we are forced into an artificial and unhealthy separation of task-specific/context-dependent (i.e., simple inference heuristics) and general-purpose/context-independent (i.e., complex decision mechanisms) methods. Both extremes tend to ignore the interdependence of "means" and "ends" (Connolly, 1999) as well as the requirement of simultaneously efficient and robust behavior.

In (Simon, 1990, Page 7), Simon identifies the two factors that determine effectual behavior, "Human rational behavior . . . is shaped by a scissors whose two blades are the structure of task environments and the computational capabilities of the actor." Simon backs up this statement, albeit implicitly, in his development of satisficing. The computational capacities of the decision-maker are means, and the consequences produced by these means, evaluated in the context of overall goal-directed behavior, are evaluated against the standard for good enough ends. We wish to characterize the essence of satisficing as a cost-benefit tradeoff using a justifiable decision theoretic standard for performing rational self policing.

3.1 Some Related Characterizations of Satisficing

Satisficing facilitates the development of a decision theoretic paradigm that differs from the de facto paradigm of optimality. One application of the concept of satisficing is in multi-attribute decision-making. "Aspiration levels provide a computational mechanism for satisficing. An alternative satisfices if it meets aspirations along all dimensions." (Simon, 1996, Page 30). Exploiting a parallelism between multiple attributes and multiple relevant states, this notion of satisficing has been mathematically formalized in (Mesarovic, 1970; Mesarovic and Takahara, 1972; Matsuda and Takatsu, 1979b; Matsuda and Takatsu, 1979a; Takatsu, 1980; Takatsu, 1981). These developments compare a utility, $v(u, \theta)$, defined over the consequences of an option u given state θ, to a decision threshold (or aspiration level), $\rho(\theta)$. Note that this decision threshold depends only on observations and not on decision consequences. An option u is satisficing if and only if $v(u, \theta) \geq \rho(\theta)$ for all feasible θ. Our approach is similar to these other developments in that it is applicable to multiple states or attributes but, by contrast, compares two utilities defined over the consequences of a decision whence our approach mathematically generalizes these decision rules (i.e., the decision threshold $\rho(u, \theta)$ depends upon both control actions and the state of nature).

In this section, we characterize tradeoffs using two utility functions: one to represent the payoff for accepting an option and another for rejecting the same option. In our development and examples, we demonstrate why this generalization to an option-dependent threshold is useful. We then discuss two methods for combining these two utility functions to resolve tradeoffs. In Section 4, we discuss the applicability of each of these methods.

3.2 Epistemic Utility Theory: A Related Characterization

The philosopher Karl Popper made the following insightful comment regarding the goals of scientific inquiry, "*. . . truth is not the only aim of science.* We want more than mere truth: what we look for is *interesting truth*."(Popper, 1965, Page 229). Although this statement is implicitly accepted by philosophers and scientists, most formal descriptions of scientific inquiry only implicitly accommodate this observation. By contrast, the epistemologist Isaac Levi made explicit this observation in his characterization of rational decision-making (Levi, 1980). A decision maker seeking to increase its knowledge is not only trying to learn truth but also trying to gain new and useful information. Such a decision maker is simultaneously playing two games: a game to obtain useful information and a game to preserve truth. Given a set of propositions, U, closed under negation (that is, if u is in the set than so is the negation, $\bar{u}$), information is gained whenever irrelevant or useless propositions $u \in U$ are rejected. This translates into a utility for rejecting proposition u or, equivalently, retaining proposition $\bar{u}$. On the other hand, identifying true propositions is one of the goals of epistemology which translates into a truth-based utility of accepting proposition u. Given these sometimes competing cognitive goals, the decision-maker engages in inquiry to identify true but informative propositions; stated simply, Levi asserts that error should be avoided (that is, truth should not be compromised) in the interest of adopting informative propositions. The lesson we learn from Levi is that truth and information are both essential elements of decision making and can be made explicit in the construction of a tradeoff-centered comparative rationality.

3.3 Comparing Values: Satisficing

Turning attention from the narrow world of epistemology to the broad world of practical decision making, we observe that truth is the epistemological manifestation of the practical decision-maker's goal of achieving success, and that information is the epistemological manifestation of the practical decision-maker's goal of efficiently using resources. In the practical decision-making arena, Popper's injunction can be rephrased to become *we want more than success — what we look for is efficient success*.

Building on Levi's work, tradeoffs can be thought of as a game between competing values. For most decision problems, there are not only reasons for accepting an option, but also reasons for rejecting an option. We need to translate these "pros" and "cons" into a decision rule that resolves these tradeoffs. Thus, we have two independent value functions: a payoff for selecting option u given θ, $J_1(u, \theta)$ similar to Levi's truth support utility, and a payoff for rejecting option u given θ, $J_2(\bar{u}, \theta)$ (similar to Levi's informational value of rejection).

Returning again to Simon's notion of satisficing, we can think of an aspiration level as the utility of rejecting an option. In Simon's formulation, the aspiration level is derived from an expectation of possible consequences. By rejecting option u, the decision maker expects a payoff at least as great as the aspiration level. Thus, $J_2(\bar{u}, \theta)$ (which equals $\rho(\theta)$) encodes the aspiration level when the aspiration level is independent of the option. According to Simon, a decision is good enough only if $J_1(u, \theta) \geq J_2(\bar{u}, \theta) = \rho(\theta)$.

In the more general case when the payoff for rejecting an option depends on the option, we can think of the relationship of J_1 and J_2 as a tradeoff. The conventional approach to resolving tradeoffs is to combine the two utilities into a single utility and then to maximize the resulting hybrid utility; we discuss some aspects of this approach in the next subsection, but in this section we discuss an alternative formulation. Recall that a decision is optimal if and only if, when compared to all other options, no other option is superior whence optimality is determined by comparing options against each other. By contrast, tradeoffs are resolved not by comparing options but rather by comparing values (defined over consequences) against each other, whence comparative rationality requires the evaluation of consequences. This is in the spirit of Simon's satisficing wherein the consequences of an option, encoded in the option's utility, are compared to the consequences of rejecting the option, encoded in the expected utility of an unidentified option.

An alternative to the optimization of an aggregated utility is to treat the resolution of a tradeoff as a meta-decision problem. Note that, in general, heuristics are used as fast and frugal ways to produce decisions and are therefore not decisions themselves but rather decision rules. Such decision rules are produced through a meta-decision process, sometimes the result of evolutionary forces, sometimes the result of external feedback, and sometimes the result of self-directed internal feedback. We construct a satisficing decision rule as the resolution of the tradeoff between J_1 and J_2 in Appendix A. Switching from the awkward notion of utility of rejecting u encoded in $J_2(\bar{u}, \theta)$, we instead choose to think of the cost of choosing u and the benefit of choosing u encoded in, respectively, $\mu_L(u;\theta) = J_2(\bar{u}, \theta)$ and $\mu_A(u;\theta) = J_1(u, \theta)$. The satisficing decision rule, derived in Appendix A and presented as Equation (11) is repeated here for convenience

$$S_b = \{(u, \theta) : \mu_A(u;\theta) \geq b\mu_L(u;\theta)\}. \tag{3}$$

Under this rule, the consequences of decision u given observation θ are evaluated without reference to other decisions; an option is good enough if the consequences it produces are satisficing, and this characterization can be determined without reference to other options.

From (3) we see that the essence of satisficing, as determined from a tradeoff-centered resolution of indeterminate values, is a comparison. Intuitively speaking, this notion of satisficing requires that the payoff of selecting an option outweigh the payoff of rejecting that option. The definition of "good enough" is based on comparing an option's benefit against the option's cost (and noting that the payoff for rejecting an option is equivalent to a cost for accepting the option). This permits an agent-centered characterization of good-enough. An option is "good enough" if benefit (as encoded in μ_A) outweighs cost (as encoded in μ_L). Satisficing therefore becomes a two-attribute decision problem with a benefit attribute (operationally termed *Accuracy*, meaning conformity to a given standard) and a cost attribute (operationally termed *Liability*, meaning susceptibility or exposure to to something undesirable). Simon likened situated rationality to scissors with one blade the structure of the task environments and the other the computational capabilities of the actor. When these scissors operate, they produce two independent evaluations of consequences (the set of consequences are cut in two):a success-based evaluation called accuracy, and an efficiency-based evaluation called liability.

Given the satisficing decision rule, we can characterize the set of all states which are satisficing for a given u, and those skills which are satisficing given the state of nature, respectively defined as

$$\begin{aligned} S_b(u) &= \{\theta : \mu_A(u,\theta) \geq b\mu_L(u,\theta)\} && (4) \\ S_b(\theta) &= \{u : \mu_A(u,\theta) \geq b\mu_L(u,\theta)\}. && (5) \end{aligned}$$

In practice, a decision-maker will not identify all elements of these sets, but will instead rely on the boundaries of these sets to detect when a behavior modification is mandatory. Suppose a cognitive skill $u \in U$ is being used to solve a decision problem. When $\theta \in S_b(u)$ then there is no need to resort to another approach. However, when $\theta \notin S_b(u)$, the current skill is inadequate and must be switched to a different skill. Given the need to switch, any skill $u' \in S_b(\theta)$ can be employed. An evaluative algorithm can be outlined for tradeoff-based skill management as follows: `If` $\theta \in S_b(u)$ `then` $u' = u$`; Else` $u' \in S_b(\theta)$. This algorithm can be used to determine when a switch is mandatory. In other words, when θ is such that u is not satisficing then a new skill $u' \neq u$ must be selected.

3.4 Comparing Alternatives: Domination

Satisficing, as we have defined it, is a notion of rationality determined by comparing two aspects of the consequences of making a decision. Under this rationality, a decision can be admitted or rejected without reference to other decisions. However, learning, memory, and the ability to model the world sometimes permits an agent to compare the consequences of one decision against another. This allows a decision maker to compare the consequences of alternative decisions in an effort to improve performance. For every $u \in U$ let

$$\begin{aligned} B_A(u;\theta) &= \{v \in U : \mu_L(v;\theta) < \mu_L(u;\theta) \text{ and } \mu_A(v;\theta) \geq \mu_A(u;\theta)\} \\ B_L(u;\theta) &= \{v \in U : \mu_L(v;\theta) \leq \mu_L(u;\theta) \text{ and } \mu_A(v;\theta) > \mu_A(u;\theta)\}, \end{aligned} \quad (6)$$

and define the set of actions that are *strictly better* than u (i.e., set of actions that dominate u)

$$B(u;\theta) = B_A(u;\theta) \cup B_L(u;\theta); \quad (7)$$

that is, $B(u;\theta)$ consists of all possible actions that have lower liability but not lower accuracy than u, or have higher accuracy but not higher liability than u. If $B(u;\theta) = \emptyset$, then no actions can be preferred to u in both accuracy and liability, and u is a (weakly) non-dominated action with respect to θ. The *non-dominated* set

$$\mathcal{E}(\theta) = \{u \in U : B(u;\theta) = \emptyset\} \tag{8}$$

contains all non-dominated actions. It is interesting to note (see (Goodrich et al., 1998b)) that the set $\mathcal{E}(\theta)$ is equivalent to the set of those options which maximize the aggregated utility $\alpha\mu_A(u;\theta) - (1-\alpha)\mu_L(u;\theta)$ for some $\alpha \in [0,1]$. In other words, $\mathcal{E}(\theta) = \{u : \exists\alpha \in [0,1]$ for which $u = \arg\max_{v \in U} \alpha\mu_A(v;\theta) - (1-\alpha)\mu_L(v;\theta)\}$. This means that the set of non-dominated options is equivalent to the set of maximizing options when the tradeoff parameter α is completely indeterminate.

It is important to note that the interpretation of $\mathcal{E}(\theta)$ as the set of optimal multi-attribute decisions is inadequate to justify selection of an option. Observe from (8) that $\mathcal{E}(\theta)$ is not a function of the consequences of making a decision, but rather a function of the state of nature. This distinction is important because decisions should be justified on the basis of their consequences and not simply because they are superior to some other decisions according to an arbitrary criterion. An element of $\mathcal{E}(\theta)$ might be optimal with respect to some criterion, but it may also produce unacceptable consequences. Thus, domination should act as a secondary criterion for determining the usefulness of an option and not as the primary criterion whence *domination is discretionary* (which is a companion to the notion that optimality is optional); it is a fact of life that sometimes the best option available to us is still unacceptable.

3.5 Postlude

To summarize the discussion of the preceding sections, two thoughts have emerged. First, decisions in the satisficing set are justified by the consequences, and decisions in the non-dominated set are justified by the alternatives. Second, satisficing is mandatory and domination is discretionary. One more point deserves mention before we end this section. One advantage of the aspiration-based satisficing approach is that multiple attributes (or, analogously, multiple states) decreases the size of the set of options that are satisficing. This implies that searching for a solution that is satisficing may take longer. However, once a satisficing option is identified it is likely to be robustly applicable under many circumstances.

Returning to our automobile driving example, following the vehicle is satisficing because the benefits of following, relative to our goal of reaching our destination, outweigh the costs of following, relative to time loss or risk. The passing behavior dominates the following behavior, but passing is optional so we feel no mandated need to pass. When cognitive resources permit, we may observe that passing dominates, but we need not pass since the current skill produces good enough consequences. Additionally, because car following is a skill that is satisficing under both heavy and light traffic densities, the skill is robust in that it affords safe but productive driving in many driving environments.

4 Intelligence Through a Multiple Agent Society

A large step toward resolving the debate between simple heuristics versus complex decision mechanisms is made by realizing that for goal-directed choice there exist meta choice problems. For example, applying expected utility theory requires a meta choice to determine the set of feasible options, beliefs, consequences, and utilities. From a machine intelligence perspective, the debate is often a discussion of whether these meta problems should be implicitly included in the choice problem (to produce complex decision mechanisms), or if simple skills and heuristics can be efficiently and explicitly (meta-)managed to produce the same intelligent results. From a human intelligence perspective, the debate is concerned with prescriptive versus descriptive models of rational choice; prescriptive models require the decision-maker to solve the meta-problems internally, and descriptive models suppose that these problems are solved through evolutionary or other externally imposed conditions (although there is nothing unnatural about learning to self-police our behaviors).

To justify the managed-skill hypothesis in describing human behavior or to encode this hypothesis in designing machine intelligence, we must address the theoretical issue of meta choices. Satisficing is a tradeoff-centered decision principle that applies to meta decision problems and therefore decreases the gap between mind and machine, or in the quest to settle the debate between simple and complex decision mechanisms. Given that the essence of satisficing is tradeoff, the important issue is how, when, and by whom should tradeoffs be resolved. These questions are questions in meta-agency, that is, questions in self evaluation and self anticipation..

4.1 Situated Decision Makers

***** INSERT FIGURE 2 ABOUT HERE *****

As we understand the philosopher Charles Peirce, meaning and therefore intelligence can only be present in a semiotic triad consisting of some kind of observation (firstness), some kind of consequent (secondness), and some kind of mapping from observation to consequent (thirdness) that turns firstness into secondness. Goal-directed agents capable of continued existence in real environments should have the capacity to respond to, interpret, and evaluate observations in terms of their capacities and skills. The key to doing this is to allow lessons learned from the past (in the form of values, models of causal behavior, etc.) to turn observations from the present into acceptable future consequences. As diagrammed in Figure 2, the past (thirdness) transforms the present (firstness) into the future (secondness).

The lesson we learn from this triad of situated agency is that much of reasoning is done in terms of either past experiences or expected future experiences. This can be extremely complex unless effective coping strategies are developed and used. A remarkably efficient coping strategy is to organize intelligence into modules appropriate for commonly encountered circumstances. We call these modules cognitive or behavioral skills and note that these skills determine the behavior of a situated decision maker. Such a decision maker can reason about the world in terms of the consequences afforded by these skills. With the emergence of multiple skills including the capacity for general-purpose problem solving, a decision maker can be capable of very sophisticated behaviors.

****** INSERT FIGURE 3 ABOUT HERE *****

In this context, an expert is one who has a skill that will produce satisficing consequences for any state θ in the domain of expertise Θ. This is diagrammed in Figure 3. Each closed curve represents a skill that produces satisficing consequences for the θ that it encloses. Note that multiple skills can be satisficing for a particular θ and that the skill set spans almost the entire domain of expertise Θ. In general, an expert in one domain Θ will not be an expert in all domains.

4.2 Multiple Agent Society

Although many behavioral skills can be organized into a stimulus-response loop, cognitive skills require an appropriate organization. Borrowing on Minsky's *Society of Mind* (Minsky, 1986), we can treat each cognitive skill as an agent and organize these agents into a society. Recognizing that these agents must interact, we can include layers of agents managing agents (i.e., meta agents), and agents managing agents managing agents, et cetera. These layers form a multi-resolutional hierarchical society of cognitive agents. Within this society, multiple forces can influence a decision. These forces include top-down forces from agents responsible for accomplishing certain goals, bottom-up forces from skilled agents responsible for acting in a particular context, and lateral forces from neighboring agents interacting to accomplish tasks that accommodate shared goals or require shared resources. Top-down forces evoke success-based evaluations of skills, and bottom-up forces evoke efficiency-based evaluations of skills,

4.3 Decision Forces

Cognitive skills provide affordances for rational behavior. The term *affordance* is a term introduced by Gibson (Gibson, 1979) and extended by Norman to mean "those fundamental properties that determine just how the thing [skill] could possibly be used" (Norman, 1988, Page 9). Skills whose affordances are compatible with top-down goals induce an attractive potential commensurate with their likely usefulness. In terms of the values involved in a tradeoff, $J_1(u;\theta)$ represents this attractive potential. However, in addition to task specific goals there are also context dependent constraints on the efficiency of these skills, and these constraints induce a repulsive potential commensurate with their likely inefficiency. The function $J_2(\bar{u},\theta)$ represents this repulsive potential.

Two independent descriptions of consequences can be thought of as the interplay between the two potential fields. Given the analogy of J_1 and J_2 as attractive and repulsive potentials, we can use this analogy to interpret the notion of satisficing. An option (skill) is satisficing if and only if the attractive potential is greater than the repulsive potential. Partitioning evaluations of consequences into these attributes recalls the generalized potential field (GPF) approach to path planning and obstacle avoidance (see, for example, (Nam et al., 1996; Guldner and Utkin, 1993)). In the GPF methodology, a goal is represented as an attractive potential, obstacles are represented as repulsive potentials, and the path along the negative gradient of the combined potentials is selected as a collision free path. Although GPF approaches have traditionally been used to plan a feasible path (with a corresponding sequence of actions), the basic idea has been extended to dynamic environments wherein individual actions are identified as a function of current and projected future dynamic states (Nam et al., 1996). Unlike such GPF approaches which produce a unique best path (or unique best option), however, a tradeoff is resolved once a single skill is identified with attractive potential greater than repulsive potential. By contrast, non-dominated options are best in the GPF sense.

In this way, satisficing is a companion to a resolved tradeoff emerging from independent values. The interaction between meta-agents resolving meta problems and choice-agents resolving choice problems involves inherent indeterminacy. Simply put, a meta-agent does not know (nor especially care) what option a choice-agent will select, nor is it appropriate for the meta-agent to speculate about the expected choices of the choice-agent (doing so shifts all responsibility to the meta-agent and relegates the choice-agent to a vacuous role). The meta-agent is responsible for abductively framing the problem, and the choice-agent is responsible for inductively solving the problem. Since the consequences of framing a problem are different from the consequences of solving a problem, there is an indeterminate mapping from the consequences evaluated by the meta-agent and the consequences evaluated by the choice-agent. This indeterminacy requires that the meta-agent deal with sets of options and produces a decision rule used by the choice-agent to identify “good enough” consequences.

4.4 Two Stages of Self-Policing: Evaluation and Anticipation

As diagrammed in Figure 2, a minimum requirement for intelligence is a relationship between past, present, and future. To facilitate this relationship, there must be three phases for any choice problem: anticipating consequences, the “moment of truth” when choice is made, and evaluating consequences. Anticipating future and evaluating past consequences are necessary stages in rational self-policing. By evaluating past consequences a decision-maker is evaluating its past choices, and is thus performing a third person (meta) evaluation of a “past self.” If performance is inadequate or if superior alternatives are manifest then the decision maker should adapt its future behavior. By anticipating future consequences, a decision maker is evaluating its future states, and is thus performing a third person (meta) evaluation of a “future self.” If expected performance is inadequate or if superior alternatives are recognized then the decision maker should act accordingly.

Unless anticipation and evaluation are simply re-enactments of the moment of truth, the decision maker should be seeking to identify feasible options. This is done in two ways: by identifying options that resolve tradeoffs and by identifying options that are non-dominated. In order of increasing complexity and necessity, satisficing-based rationality must be satisfied first (unless, for a particular world, non-domination guarantees satisficing) and then, if resources permit, domination-based rationality can be satisfied.

4.5 Problem Solving

Let us now turn attention to a timeline for making rational decisions. Assume that the decision maker is situated, meaning that the decision maker has a known goal and exists in a particular context. A rational decision-maker should begin by identifying the set of possible states of nature Θ. By identifying relevant states, a process aided by familiarity with the situation or previous exposure to similar situations, the decision-maker is able to identify the goal-driven affordances from the suite of cognitive skills that it has available. Additionally, the decision-maker can recognize contextual factors that restrict the applicability of particular skills.

***** INSERT FIGURE 4 ABOUT HERE *****

Given an observation x and a subsequent understanding of the state of nature θ, the decision maker then enters the three phase process of decision making. This process consists of the two self-policing phases of anticipation and evaluation, and the moment of truth (i.e., choice) phase.

Anticipation

Given the state of nature, the meta-agent can determine the expected payoffs μ_A and μ_L of each skill through a deliberation process, through an experience-based identification process, or through a stimulus-response mechanism. The meta-agent initiates a search for an appropriate skill, perhaps beginning at a default skill and the proceeding to other skills using some practiced procedure. Each possible skill is analyzed to see if its expected consequences resolve the tradeoff between goal-driven payoffs and context-dependent costs. If the decision maker has ample time and memory or a wealth of experience, multiple satisficing options might be identified yielding $G \subset S_b(\theta)$. Furthermore, past experience may help the decision maker identify non-dominated options via lateral decision forces whence the search can be restricted to non-dominated option $\mathcal{E}(\theta)$.

Moment of Truth

At the moment of truth, any $u \in G$ can be applied. Selecting among the alternatives can be done (a) via a constrained optimization policy such as selecting $u^* = \arg\max_{v \in G} \mu_A(v;\theta) - b\mu_L(u;\theta)$, (b) via an exploratory policy wherein an unexplored option $u \in G$ is randomly selected, or (c) through an arbitrary process wherein any $u \in G$ is randomly selected. The policy can be adapted to reflect the nuances present in the moment of truth. Regardless of the policy, the skill is expected to produce satisficing consequences and thus resolve the decision tradeoff because of the anticipation phase. In terms of modeling human behavior, it is important to note that not all behavioral variability is a result of noise and uncertainty. Instead, a portion of this variability results because, for a particular context and a particular task, many behaviors may be satisficing. This suggest that constrained optimization, though possibly appropriate for design, may not be readily applicable to describing the moment of truth in naturalistic settings.

Evaluation

Following the moment of truth, the consequences of the choice are evaluated. Any mismatch between anticipated and observed values can be used to tune these values for future use. Additionally, the meta-agent determines if the consequences are satisficing. If they are satisficing ($u \in S_b(\theta)$) and if cognitive resources are available, then the meta-agent might[2] compare u to other (remembered) options to refine the set of feasible options $\mathcal{E}$, or bias the search mechanism to use u again. If the consequences are not satisficing ($u \notin S_b(\theta)$) then the meta-agent might spawn a new search to find a skill $u' \in S_b(\theta)$ that is expected to produce satisficing consequences. If, after exploring and evaluating all known options, $S_b = \emptyset$ then the tradeoff is not resolvable given current options. The decision maker must then either adjust its expectations (decrease b) or acquire a new skill that will be appropriate for the circumstances.

5 Discussion and Examples: Simple Inference Heuristics Vs. Complex Decision Mechanisms

Satisficing rationality provides a simple but justifiable method for determining when simple inference heuristics and or/complex decision mechanisms are justified. This allows the decision maker to perform cognitive and behavioral tasks with an appropriate mechanism by turning the meta decision problem into one of controlling the selection of an appropriate skill. In this section, we give examples that support the hypothesis that intelligent behavior is organized into cognitive and behavior skills, and discuss how satisficing manifests itself in decision making.

5.1 Heuristics and Biases: the Existence of Cognitive Skills

In studies of human cognitive performance, Daniel Kahneman and Amos Tversky have led the way in identifying several heuristics and biases that systematically differ from standards of normative rationality (Kahneman and Tversky, 1979; Gardner, 1985). Among other observations, two seem most relevant to our discussion. The first observation is that people use and misapply cognitive shortcuts in inappropriate situations. The misapplication of cognitive shortcuts (i.e., heuristics) is evidence that people have and use these heuristics, and the fact that some cognitive biases disappear when the problem is reframed (e.g., the overconfidence bias can be overcome when data are presented as frequencies rather than probabilities (Gigerenzer, 1996; Kahneman and Tversky, 1996)) indicates that cognitive skills are ecologically adapted to certain domains. Additionally, the presence of robust but simple heuristics such as "take the best" demonstrate that these heuristics can be very effective (Gigerenzer and Goldstein, 1996). The second observation is that untrained people are not very good at applying methods of normative rationality. One logical conclusion from the heuristics and biases literature is that if researchers want to fool subjects, they can probably succeed by inducing an incorrect cognitive skill.

5.2 Experts and Naturalistic Decision Making: Non-Optimal Choice

The naturalistic decision making community has emerged in response to dissatisfaction with using normative models of rationality in descriptions of expert behavior (Zsambok and Klein, 1997). A characteristic of naturalistic decision making descriptions of human intelligence is that experts organize intelligence into cognitive skills. The majority of time spent in expert decision making is spent searching and understanding the state of nature. Once the state space is accurately deciphered, an appropriate skill is invoked and the problem is efficiently solved. From (Dreyfus and Dreyfus, 1985), "[Experts] reflect upon the goal or perspective that seems evident to them and upon the action that seems appropriate to achieving that goal." They reason about the world in terms of afforded actions, and select action according to their stated goal. Experts do not reason using context independent and general purpose problem solvers, but rather with cognitive skills spanning the range of relevant states of nature (see Figure 3). Developing expertise is the process of spanning the states of nature and learning how to recognize and use the appropriate skill. Since acquiring a set of skills that span Θ is done during the process of becoming an expert, the majority of expert time is not spent in complex decision making but rather in identifying a skill appropriate for the circumstances.

5.3 Human Interaction with Automation: Detailed Example of Explicit Skill Management

In this subsection, we give a detailed example of skill management in the context of human interaction with automation. Automation is ideal for illustrating skill management because when a human initiates automation they are consciously delegating a skill to the machine, and when they terminate automation they are consciously appropriating a skill from the machine. Both of these skill transitions provide a means to demonstrate how humans can manage skills. This section is largely taken from previously published work or from work currently in review (see (Goodrich et al., 1998a) for a summary of the work under review).

A mental model is an internal representation employed to encode, predict, and evaluate the consequences of perceived and intended changes to the system operator's current state within the dynamic environment. Humans interpret and respond to sensory input according to the context established by a mental model through task-specific filtering of the external world. Skilled action is organized into behavioral quanta that correspond to separate mental models each with their own perceptually delineated operational domain (Goodrich et al., 1998a). Many aspects of cognitive decision-making have been described in terms of mental models (Johnson-Laird, 1988; Minsky, 1986).

Formally, a mental model $\mathcal{M}$ is a triple consisting of the perceived state of the environment Θ, a set of decisions or actions U, and a set of ordered consequences C that result from choosing $u \in U$ when $\theta \in \Theta$ obtains. According to this specification, a mental model not only encodes the relation between the input-action pair[3] (θ, u) and the predicted consequence c, but also induces an evaluation of preferences among consequences (see Figure 5, and compare to related figures in (Meystel, 1996; Sheridan, 1992; Albus, 1991)). In words, the mental model $\mathcal{M}$ provides the context for meaningfully interpreting sensory information and generating purposeful behavior and thus represents and encodes past experience within a decision problem.

***** INSERT FIGURE 5 ABOUT HERE *****

***** INSERT FIGURE 6 ABOUT HERE *****

Human behavior can be organized into a set of skilled activities that are applied when afforded by the environment (Gibson and Crooks, 1938; Norman, 1988). In this context, the term *activity*[4] means the human's actions on the system (e.g., a behavioral activity is pushing the brake pedal or turning the steering wheel, and a cognitive activity is adding two numbers or making a simple deduction). Formally, a *skill* can then be defined as *a learned sequence of human activities*. The human must map environmental cues into selected activities; an efficient way to perform this mapping is to employ a pattern of activities specific for a particular task, and then implement this skill when appropriate. This approach uses a task-specific mental model to determine which skill is appropriate for the circumstances. Switches between skills are mandated when target perceptual states are not achievable by the currently enable skill-based behavior or when enabled skills are not satisficing for the given state.

Human cognition can be described using multiple mental models (treated as agents) which can be organized into a society of interacting agents. This societal structure not only determines which agents contribute to human behavior, but also which agents can employ attentional resources. A three level multi-resolutional society of interacting mental models organized into a hierarchical structure (see Figure 6) can be constructed corresponding to Rasmussen's knowledge-based (KB), rule-based (RB), and skill-based (SB) behaviors[5] (Rasmussen, 1976; Sheridan, 1992). At the KB level of this hierarchy, the agent role is supervisory; at the RB level, the agent role is task management; and at the SB level, the agent role is task execution. Intuitively speaking, the KB, RB, and SB agents think, monitor, and control, respectively. These mental model agents operate within the context of overall complex human behavior. SB agents are akin to cognitive skills, and RB agents are akin to meta agents.

Automobile driving is a mix of cognitive and behavioral skills. When a driver delegates a task to automation, the vehicle assumes responsibility for a behavioral skill. However, the driver retains (meta) responsibility to detecting the limits of the automation and responding appropriately. We conducted an experiment in which human subjects were placed in a driving simulator with a cruise control system engaged. At random intervals, a vehicle cut in front of the subject's vehicle and compelled the subject to determine if the automation can safely perform the skill or if the driver needed to intervene. Empirical estimates of accuracy and liability can be obtained as described below. Figure 7 presents the resulting empirical estimates and the best fit curve to these estimates. Note that for this example, the ecologically valid state variables $\Theta = [T_h, T_c^{-1}, v]$, (time headway, time-to-contact, and velocity, respectively) suffice to describe the domain of expertise Θ (Goodrich and Boer, 1998).

***** INSERT FIGURE 7 ABOUT HERE *****

Empirical Estimates

To identify μ_A and μ_L, our objective is to find substates that trigger active braking. We therefore distinguish between nominal behavior $u \in \{\mathrm{SR}, \mathrm{TR}\}$ and active braking behavior $u = \mathrm{BA}$. Our goal is thus to find when $\theta_{\mathrm{RB}} \notin S_b(u)$ for $u \in \{\mathrm{SR}, \mathrm{TR}\}$. Nominal operating conditions occur when the brake pedal is not pressed. For both nominal and braking conditions, we select representative sample points from each experimental trial and create two sets of $[T_c^{-1}, T_h]^T$ points: one set for nominal conditions, denoted NOM, and one set for braking conditions, denoted BRK. For trials when subjects actively brake, the sub-state(s) $[T_c^{-1}, T_h]^T$ when braking is initiated is included in BRK, and the sub-state(s) $[T_c^{-1}, T_h]^T$ when braking is terminated is included in NOM; for trials when subjects do not brake, the initial sub-state $[T_c^{-1}, T_h]^T$ in the trial is included in NOM; and for trials where subjects only brake (by anticipating the cut-in and then coming to a stop), the initial sub-state $[T_c^{-1}, T_h]^T$ in the trial is included in BRK.

For notational purposes in the subsequent sections, let $N(T = \tau|\mathsf{CONDITION})$ denote the cardinality of the set of points $T = \tau$ given CONDITION. For example, $N(T_c^{-1} = \tau|\mathsf{NOM})$ is the number of points in the set $\{\theta \in \mathsf{NOM} : T_c^{-1} = \tau\}$. Under nominal conditions ($\theta \in \mathsf{NOM}$), relative velocity must be considered acceptable to the driver whence the distribution of T_c^{-1} under nominal conditions is an observable entity that provides information about what is accurate. Clearly, if $T_c^{-1} = \tau_2$ is accurate, then $\tau_1 < \tau_2$ must be at least as accurate. This monotonicity property facilitates the computation of the accuracy function as the cumulative distribution function

$$\begin{aligned}\mu_A(T_c^{-1} = \tau) &= 1 - F_{T_c^{-1}}(\tau|\mathsf{NOM}) \\ &= 1 - \frac{N(T_c^{-1} \leq \tau|\mathsf{NOM})}{N(T_c^{-1} \leq \infty|\mathsf{NOM})}.\end{aligned}$$

For classification purposes, we fit (via least squares) a sigma function of the form $1/e^{(-a\tau+b)}$ to $\mu_A(\cdot)$ yielding the function shown in Figure 7(a).

When braking is initiated ($\theta \in$ BRK), time headway values must be considered unacceptable whence the distribution of time headways when the driver initiates braking is an observable entity that provides information about what is rejectable. Clearly, if $T_h = \tau_2$ is rejectable then $\tau_1 < \tau_2$ must be at least as rejectable. This monotonicity property facilitates the computation of the rejectability function as the cumulative distribution function

$$\begin{aligned}\mu_R(T_h = \tau) &= 1 - F_{T_h}(\tau|\mathsf{BRK}) \\ &= 1 - \frac{N(T_h \le \tau|\mathsf{BRK})}{N(T_h \le \infty|\mathsf{BRK})}.\end{aligned}$$

For classification purposes, we fit (via least squares) a sigma function of the form $1/e^{(-a\tau+b)}$ to $\mu_R(\cdot)$ yielding the function shown in Figure 7(b).

Classification Results

For the driver to switch from one skill to another, it is necessary to identify when $u \notin S_b(\theta)$. Using $\mu_A(T_c^{-1})$ and $\mu_L(T_h)$ from Figure 7, we can construct the set of states $\mathcal{S}_b = \{\theta : \mu_A(T_c^{-1}) \ge b\mu_L(T_h)\}$ that support nominal behavior, and the set of states $\mathcal{S}_b^c = \{\theta : \mu_A(T_c^{-1}) < b\mu_L(T_h)\}$ (superscript c denotes complement) that do not support nominal behavior. If $u \in \{\mathrm{TR}, \mathrm{SR}\}$ and $\theta \in \mathcal{S}_b^c$ then $\theta \notin S_b(u)$. Thus, the line $\mu_A(T_c^{-1}) = b\mu_L(T_h)$ determines when behavior must be switched from nominal to braking. In other words, the line is the boundary of $S_b(\mathrm{SR})$.

***** INSERT FIGURE 8 ABOUT HERE *****

Given the empirically derived functions, we can determine the boundary between nominal and braking behaviors as a function of b by finding the perceptual states θ for which $\mu_A(T_c^{-1}) = b\mu_L(T_h)$. This is illustrated in Figure 8 for the data gathered in the simulator experiment, where $\circ$ indicates $\theta \in$ NOM and $\times$ indicates $\theta \in$ BRK. To the northwest of the line, BA is satisficing but TR and SR are not, and to the southeast of the line TR and SR (and, perhaps, BA) are satisficing. Classification can be performed by finding the value of b that optimally separates braking from nominal behavior. Consider the following three performance indices: $J_1(b)$ is the percentage of trials that are incorrectly classified (i.e., the total number of $\circ$'s above the line plus the total number of $\times$'s below the line) , $J_2(b)$ is the percentage of nominal trials that are incorrectly classified as braking (i.e., number of $\circ$'s above the line), and $J_3(b)$ is the percentage of braking trials that are incorrectly classified as nominal (i.e., number of $\times$'s below the line). The value $b = 0.53$ is the minimax value $b = \arg\min_{b \geq 0} \max\{J_1(b), J_2(b), J_3(b)\}$ which attempts to balance the percentage of misclassifications ($J_1(b)$), false alarms ($J_2(b)$), and missed detections ($J_3(b)$). The value $b = 0.20$ minimizes the number of samples misclassified $b = \arg\min_{b \geq 0} J_1(b)$. The classification results for the different values of b are shown in Table 1 and indicate that, on the average, over 85% of samples are correctly classified.

b	% misclassified	% false braking	% missed braking
0.20	10.04	1.95	8.09
0.53	13.25	8.37	4.88

Table 1: Classification accuracies for different values of b.

These results were validated in a separate experiment using professional drivers in real vehicles responding to cut in events on a closed test track. To perform the classification, $\mu_A(T_c^{-1})$ and $\mu_L(T_h)$ were estimated, and the b that minimizes the misclassification error was determined. The experiments generated one false alarm ($\circ$ above the line) and no missed detections ($\times$ below the line) in fifty trials at varying speeds. The results between the test track experiments and driving simulator experiments are very similar. The test track results produce a slightly smaller value of b ($b = 0.21$ for the test track versus an average value of $b = 0.53$ for the driving simulator) and a slight change in the liability function[6].

5.4 Automobile Driver Behavior and Navigation: Implicit Skill Management

We can now return to the driving example given in the introduction. While driving an automobile, we have been following a vehicle for an extended period of time even though there is very little traffic on the road. Because following the vehicle is satisficing, we do not feel a need to consider changing our behavior but rather rest content with following the vehicle. Suddenly, we realize that we can easily pass and that we want to do so because the lead vehicle is going slower than our desired speed. Once we observe that an alternative behavior is still satisficing but dominates our current behavior, we select this behavior and act accordingly. This is possible because much of the behavior associated with speed management can be described by three simple skills: speed regulation, car following, and active braking. Speed regulation applies in the absence of other traffic, and is a simple perceptual regulation task where we manage the vehicle's speed to produce an optic flow consistent with our calibrated estimate of the vehicle's speed. Car following applies in the presence of other traffic, and is a simple perceptual regulation task where we manage the relative rate of optical expansion to stay within a threshold (Lee, 1976). Active braking is less sophisticated , and is tantamount to detecting an anomalous situation and defaulting to a safe behavior. Coordinating these behavioral skills is an exercise in meta-rationality with multiple satisficing skills overlapping for a given θ.

Navigation problems can be precisely solved by appealing to the topological layout of the city and a sense of direction. However, only a portion of the population actually navigates this way (Aginsky et al., 1997). Many drivers use landmarks extensively; they maintain their course until a landmark triggers a behavioral response. Although the first method is more robust to errors for a well-informed driver, it requires complex representations of city and geography. By contrast, the landmark method is remarkably efficient in producing effective navigation, and requires minimal representation of the city. If the landmark is salient the driver can succeed with minimal knowledge of the street topology (Aginsky et al., 1997). Both skills are justifiable for most of driving, and one is a simple heuristic method while the other is a complex decision method.

Both of these examples illustrate that simple heuristics can be used in dynamic operation of automobiles. Both are remarkably efficient in the required use of cognitive resources, and current evidence suggests that they are effective in producing desired behavior. Moreover, we are exploring how these different skills are managed, and preliminary evidence suggests that satisficing meta-rationality provides a useful mechanism for skill management (Boer and Goodrich, 1998; Boer et al., 1998).

6 Conclusions

The essence of satisficing is tradeoff. Based on this theme, we constructed a decision-theoretic characterization of satisficing as a comparison of two independent evaluations of consequences: the consequence of accepting the option and the consequence of rejecting the option. Simon's original descriptions of satisficing fit nicely within this framework, but many variants of his ideas appear to abandon this comparative rationality in favor of variants of superlative rationality via constrained optimization. Since, as we have discussed, optimality is optional (and its companion domination is discretionary), satisficing, which is a mandatory evaluation of the consequences of a decision, deserves prominent attention in the decision-making community because being the best among the alternatives may not be acceptable, attainable, nor unambiguously definable. Additionally, satisficing provides a mechanism for spanning both simple inference heuristics and complex decision mechanisms; the satisficing decision rule manifests itself as a simple heuristic but has a meta-rationality justification.

We have presented several examples of satisficing and cognitive skill management. The most engaging example referred to automobile driver behavior in interacting with other traffic. The example illustrated why we consider it rational to do non-optimal things. In automobile driving, when the driver has limited attentional resources their superlative rationality is bounded and optimality is precluded — even for voracious optimizers. Nevertheless, the driver's comparative rationality permits the justification of rational driving behaviors. Extending from driving to other domains of goal-directed decision making, we observe that skills (be they heuristic or complex) are justified only if satisficing.

A A Characterization of Satisficing Through Meta-Rationality

Resolving a tradeoff is not a decision itself, but rather a decision about how to decide, that is, a *meta decision* or *decision rule*. In words, we want to identify conditions under which the tradeoff is resolved. We do this by deriving the satisficing decision rule by comparing all possible decision rules that combine a payoff for accepting and a payoff for rejecting an option. The result of this derivation is a decision rule which obtains its justification by an appeal to superlative meta-rationality, but which manifests itself as comparative rationality.

Let $\phi : \Theta \rightarrow \mathcal{B}$, where $\mathcal{B}$ is sigma-algebra associated with U, denote a decision rule that maps the set of states of nature Θ to the subset $G \in \mathcal{B}$ of possibilities. To resolve a tradeoff we must find the optimal decision rule ϕ subject to the constraint that u and $\bar{u}$ cannot simultaneously be accepted. The resulting decision rule represents the resolution of the tradeoffs in values J_1 and J_2. We can identify the utility of the decision rule ϕ as an aggregation of the two payoffs

$$\begin{aligned} H(\phi, x) &= E_{\theta|x}[J(\phi(\theta), \theta)] \\ &= E_{\theta|x}\Big\{ \sum_{u \in U} \Big[\alpha J_1(u,\theta) p(u \in \phi(\theta)) \\ &\quad +(1-\alpha) J_2(\bar{u},\theta) p(\bar{u} \in \phi(\theta)) \Big] \Big\} \end{aligned} \tag{9}$$

where $p(u \in \phi(\theta))$ represents the probability that option $u \in U$ is an acceptable resolution to the tradeoff with $p(\bar{u} \in \phi(\theta))$ defined conversely, where α denotes a tradeoff parameter, and where $E_{\theta|x}(H)$ denotes the expectation of J with respect to the conditional probability $p(\theta|x)$, that is

$$E_{\theta|x}(H(\phi(\theta),\theta) = \sum_{\theta\in\Theta} H(\phi(\theta),\theta)p(\theta|x).$$

By choosing the decision rule ϕ we effectually choose $p(\bar{u} \in \phi(\theta))$ and $p(uxs \in \phi(\theta))$. Our objective is to find a decision rule ϕ that maximizes (9) keeping in mind that the rule ϕ is a resolution between J_1 and J_2. Let

$$\begin{aligned} \mu_A(u;x) &= \sum_{\theta\in\Theta} J_1(u,\theta)p(\theta|x) \\ \mu_L(u;x) &= \sum_{\theta\in\Theta} J_2(\bar{u},\theta)p(\theta|x) \end{aligned}$$

Maximizing (9) over all possible tradeoffs yields the decision rule

$$\phi(x) = \begin{cases} u & \alpha\mu_A(u;x) \geq (1-\alpha)\mu_L(u;x) \\ \bar{u} & \text{otherwise} \end{cases} . \tag{10}$$

Without loss of generality, let observation x uniquely determine state θ whence we drop the dependence on x. The set of consequences that survive the resolution of the tradeoff is called the *satisficing set* and is given by

$$S_\alpha = \{(u,\theta) : \alpha\mu_A(u;\theta) \geq (1-\alpha)\mu_L(u;\theta)\}$$

or, equivalently,

$$S_b = \{(u,\theta) : \mu_A(u;\theta) \geq b\mu_L(u;\theta)\}, \tag{11}$$

where $b = \frac{1-\alpha}{\alpha}$. Thus, we see that resolving a tradeoff by maximizing over possible decision rules produces a weak rationality that eliminates obviously bad choices and admits good enough choices. Intuitively speaking, a tradeoff does not produce unique best decisions but rather a suspension of judgment between decisions that are "good enough." *A tradeoff is resolved if and only if at least one skill u can be identified whose expected benefits outweigh the expected costs, that is if and only if $S_b \neq \emptyset$.*

Notes

[1]This can be determined directly by anticipating consequences via a model and then ranking the consequences, or indirectly by using a value function over state-action pairs (as in Q-learning), or some contribution of both.

[2]Note that this approach of continuing to search if results are satisficing is not observed in studies of the design process (Ball et al., 1996) because, simply put, if the results are good enough then there is very little motive to continue dedicating resources to determining another option.

[3]A decision u is often treated as a mapping from Θ into the set of consequences(Fishburn, 1981).

[4]There are many uses of the term activity in pattern recognition and design literature (see, for example, Bobick97,Norman98). For our purposes, we use activity to mean low-level movements which is yet another use of the term.

[5]These layers correspond not only to Saridis *organization*, *coordination*, and *execution* levels, respectively, for intelligent machine design (Saridis, 1989), but also the strategic, tactical, and operational levels of decision-making (Boer et al., 1998).

[6]Let $T_h^{\max} = \arg_{T_h \geq 0}(\mu_L(T_h) = 1)$. For the test track, $T_h^{\max} > 0$ whereas for the driving simulator $T_h^{\max} < 0$. These differences simply indicate that the costs of error are higher when real vehicles are used; in other words, a real collision on the test track is much more costly than a simulated collision in the driving simulator. Though the subjects were sincere and well-motivated, there is simply no substitute for the fear of death to motivate a driver. Additionally, the functional representation (obtained through least squares) of the empirical liability measure introduced a bias in the simulator study because the time headway space was not uniformly sampled.

References

Aginsky, B., Harris, C., Rensink, R. A., and Beusmans, J. (1997). Two strategies for learning a route in a driving simulator. *Journal of Environmental Psychology*, 17(4):317–331.

Albus, J. S. (1991). Outline for a theory of intelligence. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(3):473–509.

Ball, L. J., Maskill, L., and Ormerod, T. C. (1996). Satisficing in Engineering Design: Causes, Consequences and

Implications for Design Support. In *Proceedings of the First International Symposium on Descriptive Models of Design*, pages 317–332, Istanbul.

Boer, E. R. and Goodrich, M. A. (1998). Mental models in micro worlds: Situated representations for the navigationally challenged. In *Proceedings of the 1998 IEEE International Conference on Systems, Man, and Cybernetics*, San Diego, CA, USA.

Boer, E. R., Hildreth, E. C., and Goodrich, M. A. (1998). The role of mental models in a driver's interaction with traffic. In *Proceedings of the 17th European Annual Conference on Human Decision Making and Control '98*, Valenciennes, France.

Brooks, R. A. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2:14–23.

Brooks, R. A. (1991). Intelligence without reason. Technical Report A.I. Memo 1293, Massachusetts Institute of Technology Artificial Intelligence Laboratory.

Chase, V. M., Hertwig, R., and Gigerenzer, G. (1998). Visions of rationality. *Trends in Cognitive Science*, 2(6):206–214.

Connolly, T. (1999). Action as a fast and frugal heuristic. *Minds and Machines*. This edition.

Dreyfus, H. L. and Dreyfus, S. E. (1985). From socrates to expert systems: The limits and dangers of calculative rationality. In Mitcham, C. and Huning, A., editors, *Philosophy and Technology II: Information Technology and Computers in Theory and Practice*, Boston Studies in the Philosphy of Science. Reidel.

Duchon, A. P., Warren, W. H., and Kaelbling, L. P. (1998). Ecological robotics. *Adaptive Behavior*, 6(3/4):471–505.

Fishburn, P. C. (1981). Subjective expected utility: A review of normative theories. *Theory and Decision*, 13:139–199.

Gardner, H. (1985). *The Mind's New Science: A History of the Cognitive Revolution*, chapter 13. Basic Books, Inc., New York.

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Houghton Mifflin, Boston.

Gibson, J. J. and Crooks, L. E. (1938). A theoretical field-analysis of automobile-driving. *The American Journal of Psychology*, 51(3).

Gigerenzer, G. (1996). On narrow norms and vague heuristics: A reply to Kahneman and Tversky (1996). *Psychological Review*, 103(3):592–596.

Gigerenzer, G. and Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4):650–669.

Goodrich, M. A. and Boer, E. R. (1998). Semiotics and mental models: Modeling automobile driver behavior. In *Joint Conference on Science and Technology of Intelligent Systems ISIC/CIRA/ISAS'98 Proceedings*, Gaithersburg, MD.

Goodrich, M. A., Boer, E. R., and Inoue, H. (1998a). Brake initiation and braking dynamics: A human-centered study of desired ACC characteristics. Technical Report TR-98-5, Cambridge Basic Research, Nissan Research and Development, Inc., Cambridge, MA 02142, USA.

Goodrich, M. A., Stirling, W. C., and Frost, R. L. (1998b). Model predictive satisficing fuzzy logic control. Technical Report CBR TR 98-3, Nissan Cambridge Basic Research, Nissan Research and Development, Inc., Cambridge, MA USA.

Greiner, R. and Orponen, P. (1996). Probably approximately optimal satisficing stragies. *Artificial Intelligence*, 82:21–44.

Guldner, J. and Utkin, V. I. (1993). Sliding Mode Control for an Obstacle Avoidance Strategy Based on an Harmonic

Potential Field. In *Proceedings of the 32nd Conference on Decision and Control*, pages 424–429, San Antonio, Texas.

Ho, Y. C. (1994). Heuristics, rules of thumb, and the 80/20 proposition. *IEEE Transactions on Automatic Control*, 39(5):1025–1027.

Ho, Y.-C. (1999). The no free lunch theorem and the human-machine interface. *IEEE Control Systems*, pages 8–10.

Ho, Y. C. and Larson, M. E. (1995). Ordinal Optimization Approach to Rare Event Probability Problems. *Discrete Event Dynamic Systems: Theory and Applications*, 5:281–301.

Johnson-Laird, P. N. (1988). *The Computer and the Mind: An Introduction to Cognitive Science*. Harvard University Press, Cambridge, Massachusetts.

Kahneman, D. and Tversky, A. (1979). Prospect Theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291.

Kahneman, D. and Tversky, A. (1996). On the reality of cognitive illusions. *Psychological Review*, 193(3):582–591.

Kaufman, B. E. (1990). A new theory of satisficing. *The Journal of Behavioral Economics*, 19(1):35–51.

Lee, D. N. (1976). A theory of visual control of braking based on information about time-to-collision. *Perception*, pages 437–459.

Levi, I. (1980). *The Enterprise of Knowledge*. MIT Press, Cambridge, Massachusetts.

Levi, I. (1997). *The Covenant of Reason: Rationality and the Commitments of Thought*. Cambridge University Press.

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

Matsuda, T. and Takatsu, S. (1979a). Algebraic properties of satisficing decision criterion. *Information Sciences*, 17(3).

Matsuda, T. and Takatsu, S. (1979b). Characterization of satisficing decision criterion. *Information Sciences*, 17(2):131–151.

Mayne, D. Q. and Michalska, H. (1990). Receding horizon control of nonlinear systems. *IEEE Transactions on Automatic Control*, 35:814–824.

Mesarovic, M. D. (1970). Systems theoretic approach to formal theory of problem solving. In Banerji, R. and Mesarovic, M. D., editors, *Theoretical Approaches to Non-Numerical Problem Solving*, pages 161–178. Springer.

Mesarovic, M. D. and Takahara, Y. (1972). On a qualitative theory of satisfactory control. *Information Sciences*, 4(4):291–313.

Meystel, A. (1996). Intelligent systems: A semiotic perspective. In *Proceedings of the IEEE International Symposium on Intelligent Control*, Dearborn, Michigan.

Michalska, H. and Mayne, D. Q. (1995). Moving horizon observers and observer-based control. *IEEE Transactions on Automatic Control*, 40(6):995–1006.

Minsky, M. (1986). *The Society of Mind*. Simon and Schuster.

Nam, Y. S., Lee, B. H., and Ko, N. Y. (1996). A View-Time Based Potential Field Method for Moving Obstacle Avoidance. *SICE*, pages 1463–1468.

Norman, D. A. (1988). *The Design of Everyday Things*. Currency Doubleday. Previously published as *The Psychology of Everyday Things*.

Popper, K. R. (1965). *Conjectures and Refutations: The Growth of Scientific Knowledge*. Harper and Row, New York.

Rasmussen, J. (1976). Outlines of a hybrid model of the process plant operator. In Sheridan, T. B. and Johannsen, G., editors, *Monitoring Behavior and Supervisory Control*, pages 371–383. Plenum.

Richalet, J. (1993). Industrial applications of model based predictive control. *Automatica*, 29:1251–1274.

Sandholm, T. and Lesser, V. (1997). Coalitions among computationally bounded agents. *Artificial Intelligence*, 94(1):99–137.

Saridis, G. N. (1989). Analytic formulation of the principle of increasing precision with decreasing intelligence for intelligent machines. *Automatica*, 25(3):461–467.

Scokaert, P. O. M., Rawlings, J. B., and Meadows, E. S. (1997). Discrete-time stability with perturbations: Application to model predictive control. *Automatica*, 33(3):463–470.

Sen, S., editor (1998). *Satisficing Models*, Stanford, California. AAAI Spring Symposium. Technical Report SS-98-05.

Sheridan, T. B. (1992). *Telerobotics, Automation, and Human Supervisory Control*. MIT Press.

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

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

Simon, H. A. (1996). *The Sciences of the Artificial*. MIT Press, 3rd edition.

Sistu, P. B. and Bequette, B. W. (1996). Nonlinear model-predictive control: Closed-loop stability analysis. *AIChE Journal*, 42(12):3388–3402.

Slote, M. (1989). *Beyond Optimizing*. Harvard University Press.

Takatsu, S. (1980). Decomposition of satisficing decision problems. *Information Sciences*, 22(2):139–148.

Takatsu, S. (1981). Latent satisficing decision criterion. *Information Sciences*, 25(2):145–152.

Zadeh, L. A. (1958). What is optimal? *IRE Transactions on Information Theory*, 4(1):3.

Zadeh, L. A. (1998). Maximizing sets and fuzzy markoff algorithms. *IEEE Transactions on Systems, Man, and Cybernetics—Part C: Applications and Reviews*, 28(1):9–15.

Zilberstein, S. (1996). Using anytime algorithms in intelligent systems. *AI Magazine*, pages 73–83.

Zsambok, C. E. and Klein, G., editors (1997). *Naturalistic Decision Making*, Hillsdale, N.J. Lawrence Erlbaum Associates.

Figure 1: The decision problem. Arrows indicate either perceived (by the decision maker) or actual influence. For example, a consequence is the result of taking action $u \in U$ when $\theta \in \Theta$ obtains. Not all decision methodologies explicitly account for all influences. Typically, complex decision mechanisms seek to make all elements and influences explicit, whereas simple inference heuristics discard some of the influences or omit some of the elements.

Figure 2: Interface between past, present, and future. Past experience, through explicit causal models or through utility elicitation, allows a decision-maker to map present observations and options into desirable future consequences.

Figure 3: For a specific domain of expertise Θ, an expert has skills spanning the space that produce satisficing consequences. For any state, multiple skills can suffice.

Figure 4: Three phases of the situated decision making. Choice is the *moment of truth* for a choice agent, and anticipation and evaluation are elements of (meta) rational self-policing.

Figure 5: Working specification of a mental model. The arrows represent perceived or real influence. Consequences are a function of states and actions; behavior is generated through the operation of a mental model, but the mental model is constrained by the set of behavioral affordances; and sensory observation influence the mental model, but the mental model dictates active sensing of the environment.

Figure 6: Interaction within a society of mental model agents. SP=sensor perception, MM=mental model, and BA=behavior actuation. The horizontal arrows are explained in the caption to Figure 5, and the vertical arrows indicate interaction of low level sensors/high level goals with high level representations/low level actions.

Figure 7: Actual (dashed line) and approximated (solid line) functions as a function of perceptually feasible observations: (a) accuracy as a function of time to collision T_c^{-1} and (b) liability as a function of time headway T_h.

Figure 8: Scatter plot of nominal and braking perceptual states. The line represents the boundary of the nominal skill (states to the northwest of the line are unacceptable). The boundary of the braking skill is not identified in this plot. Compare this figure to Figure 3.

Table 1: Classification accuracies for different values of b.

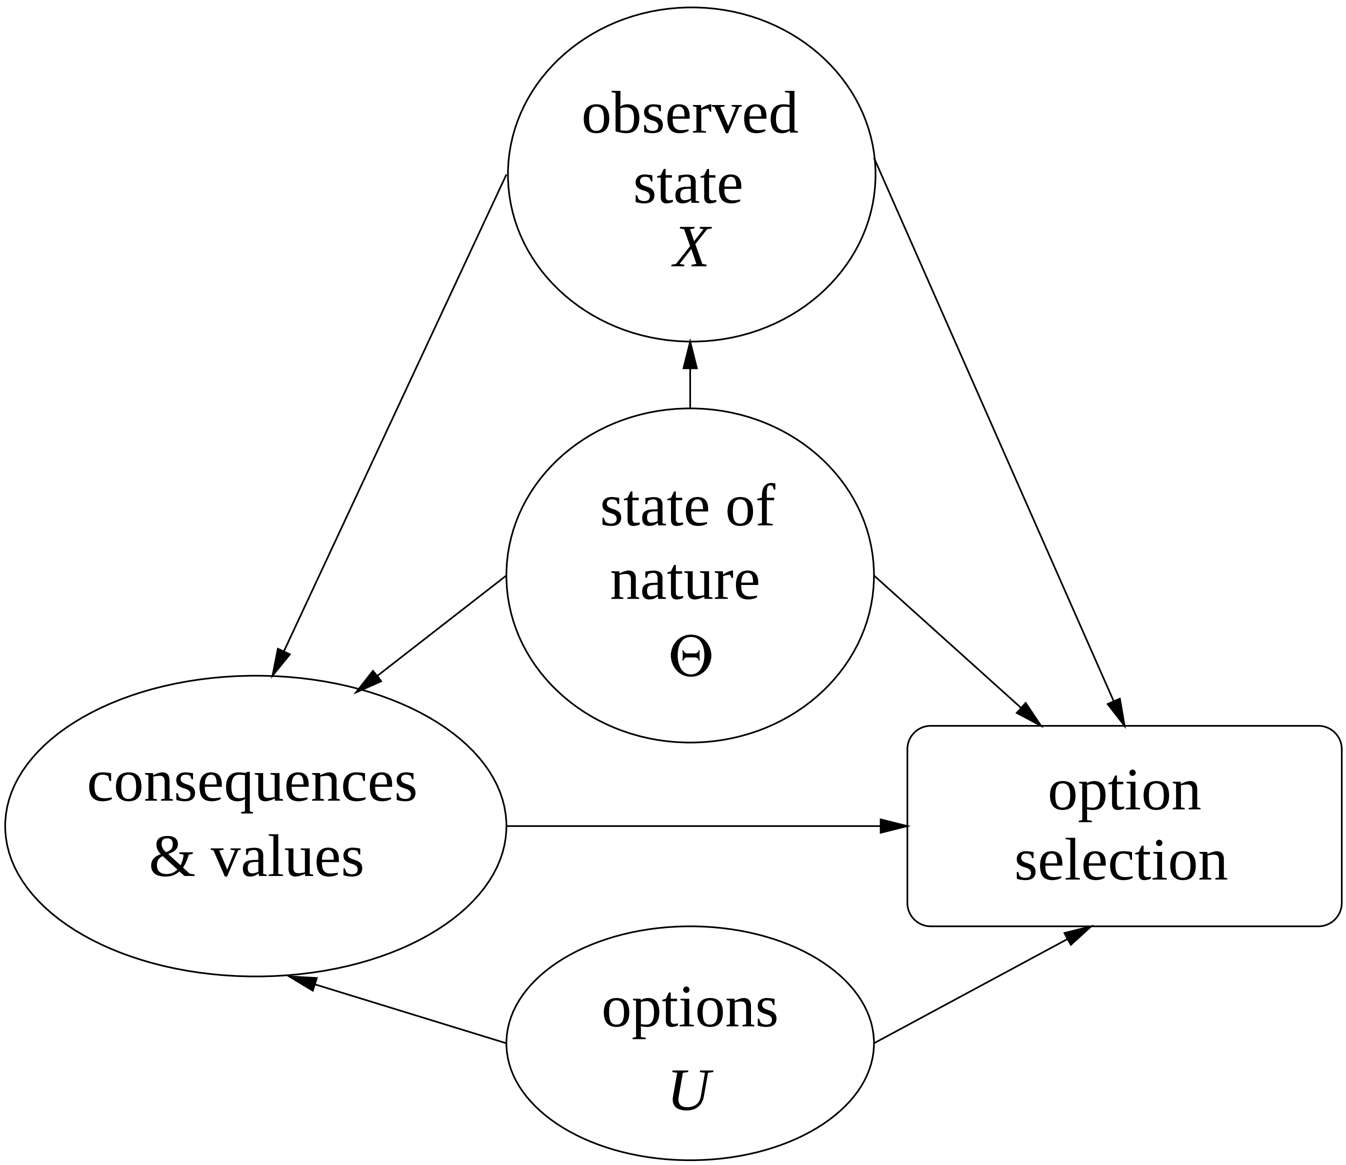
observed state X
state of nature Θ
consequences & values
option selection
options U

Figure 1:

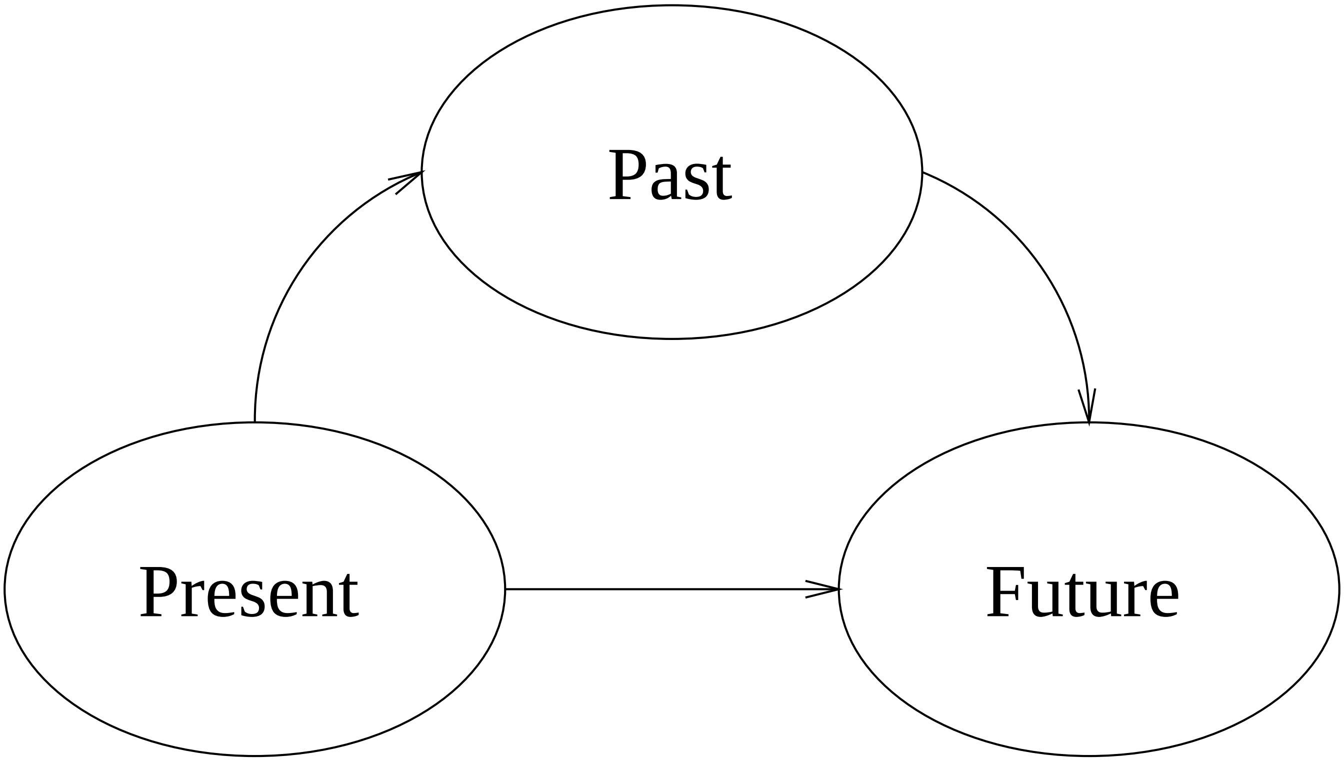

Figure 2:

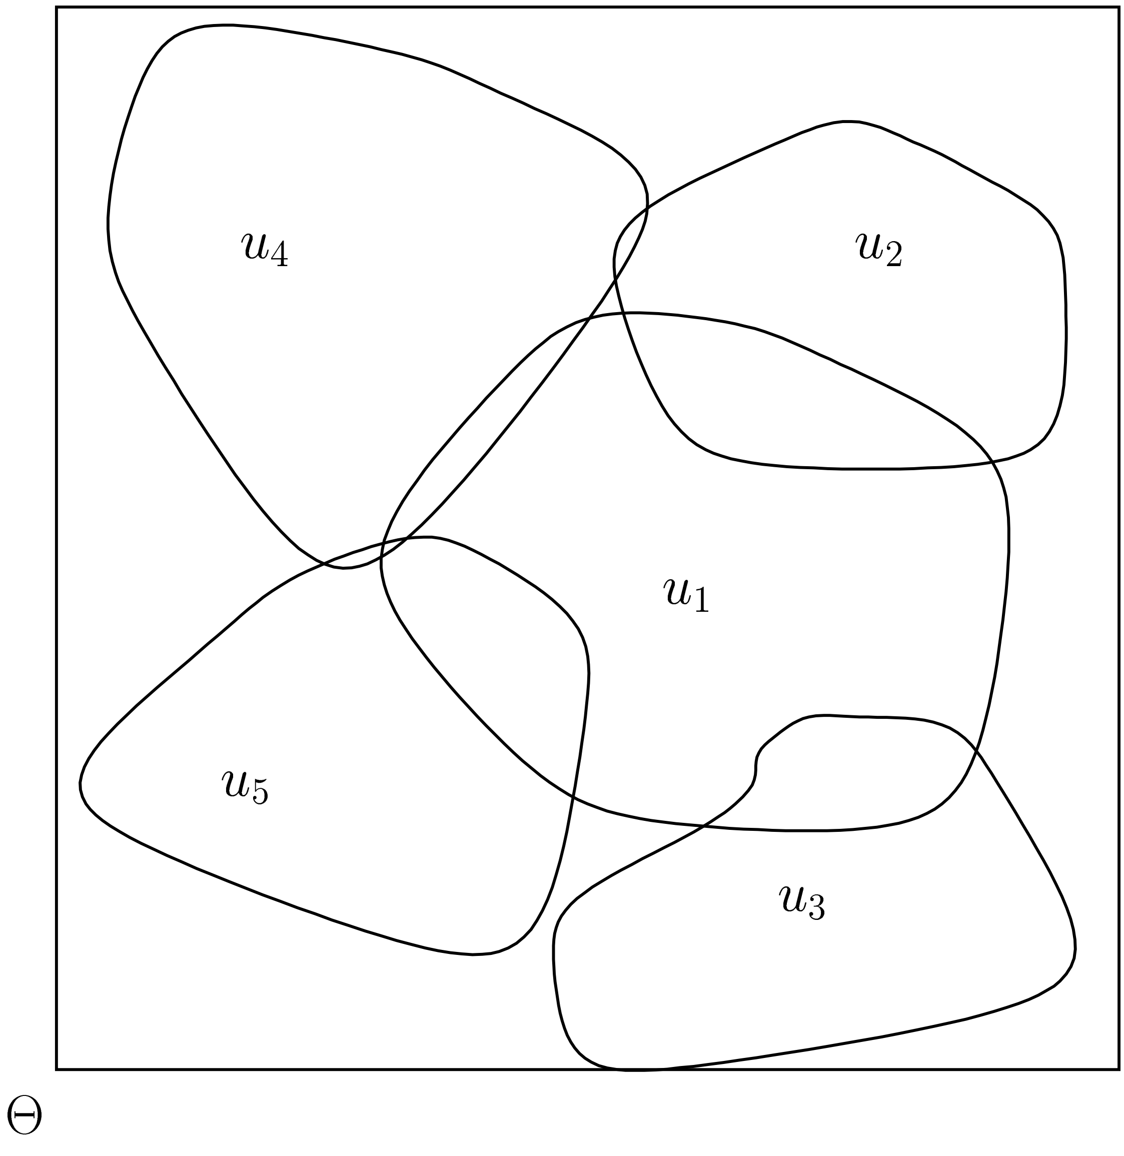

Figure 3:

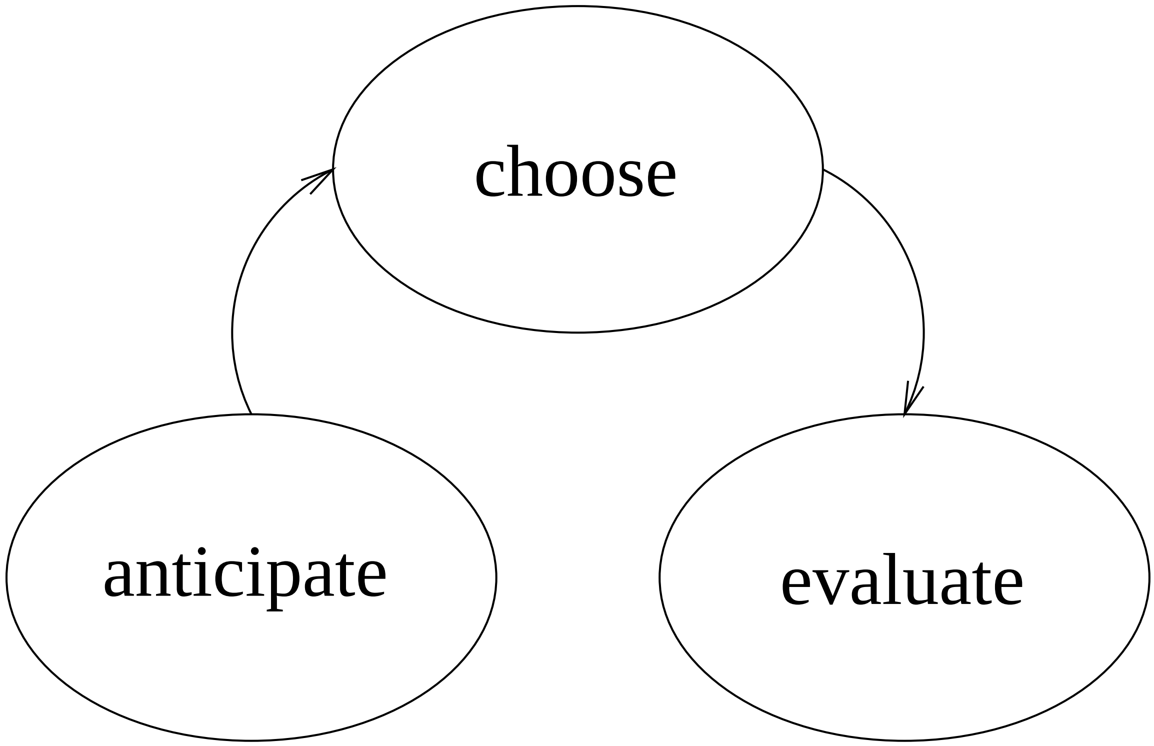

Figure 4:

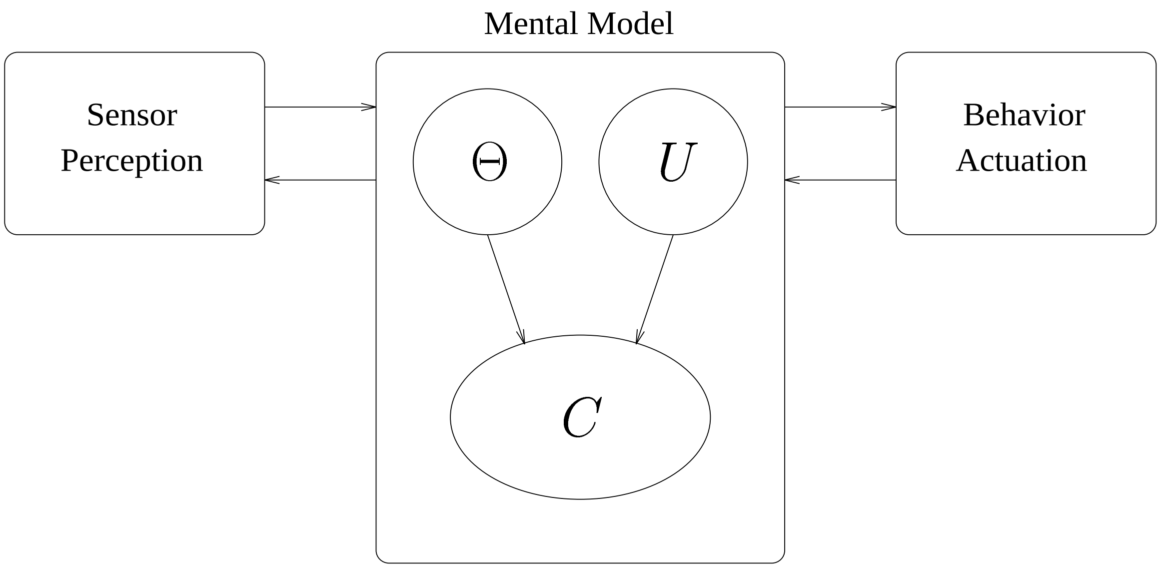

Figure 5:

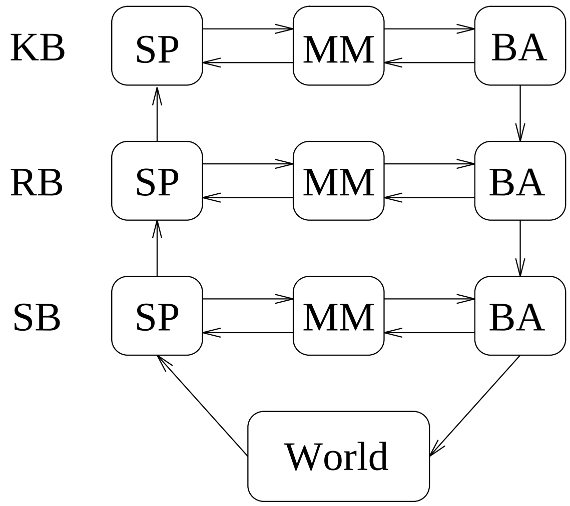

Figure 6:

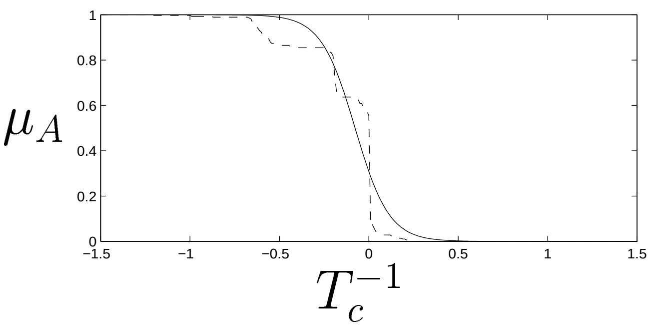

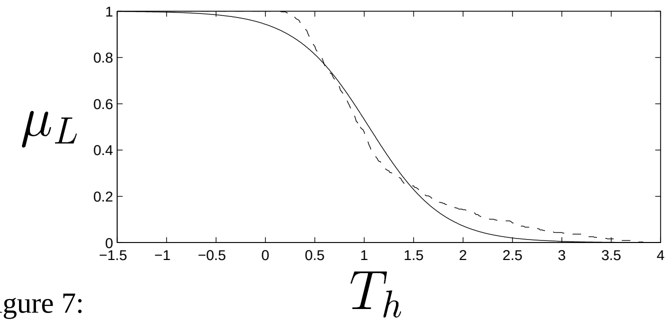

Figure 7:

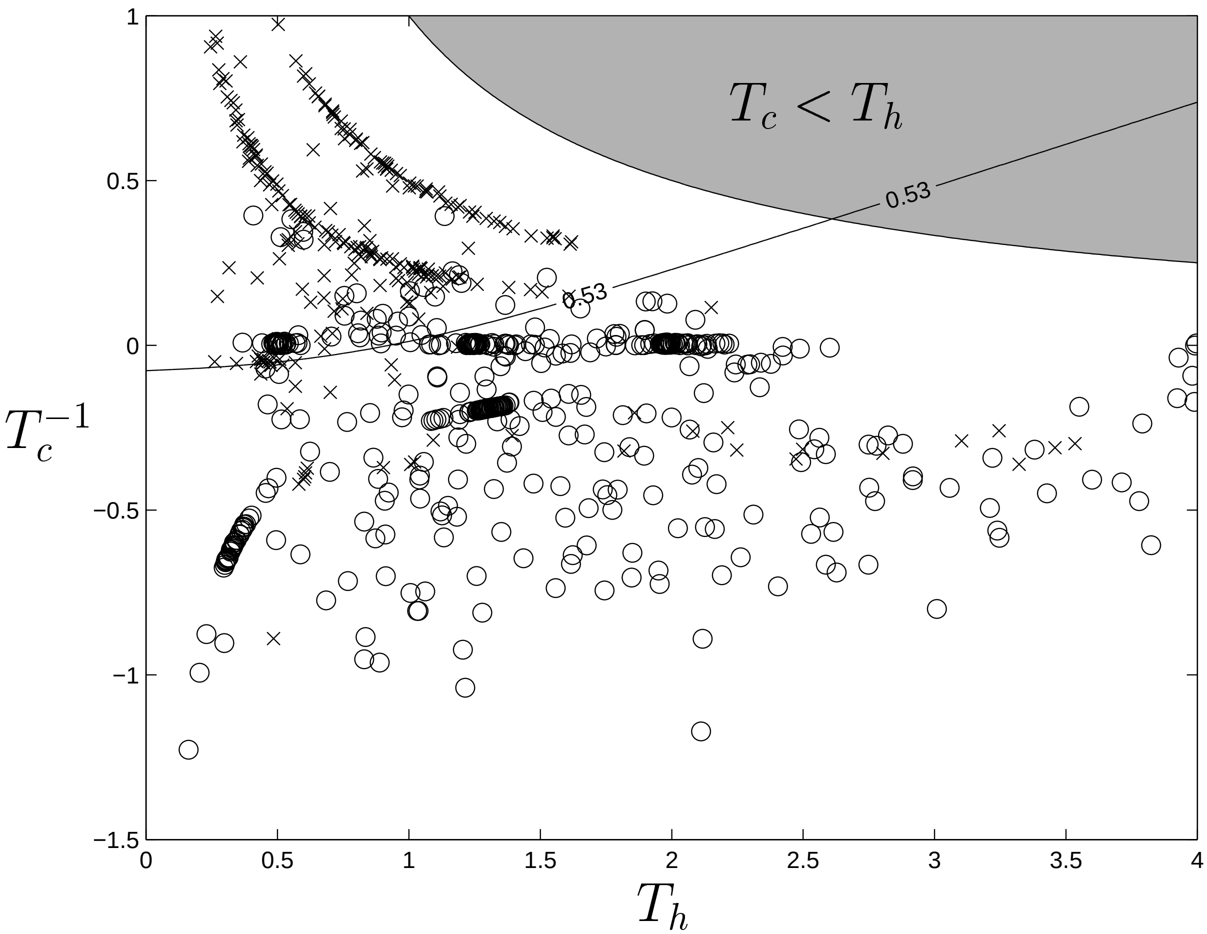

Figure 8: