On the Foundations of Qualitative Decision Theory

Ronen I. Brafman
Computer Science Department
University of British Columbia
Vancouver, B.C., Canada V6T 1Z4
brafman@cs.ubc.ca

Moshe Tennenholtz
Faculty of Industrial Engineering and Mgmt.
Technion – Israel Institute of Technology
Haifa 32000, Israel
moshet@ie.technion.ac.il

Abstract

This paper investigates the foundation of *maximin*, one of the central qualitative decision criteria, using the approach taken by Savage (Savage 1972) to investigate the foundation and rationality of classical decision theory. This approach asks "which behaviors could result from the use of a particular decision procedure?" The answer to this question provides two important insights: (1) under what conditions can we employ a particular agent model, and (2) how rational is a particular decision procedure. Our main result is a constructive representation theorem in the spirit of Savage's result for expected utility maximization, which uses two choice axioms to characterize the *maximin* criterion. These axioms characterize agent behaviors that can be modeled compactly using the *maximin* model, and, with some reservations, indicate that *maximin* is a reasonable decision criterion.

Introduction

Decision theory plays an important role in fields such as statistics, economics, game-theory, and industrial engineering. More recently, the realization that decision making is a central task of artificial agents has led to much interest in this area within the artificial intelligence research community. Some of the more recent work on decision theory in AI concentrates on qualitative decision making tools. For example, Boutilier (Boutilier 1994) and Tan and Pearl (Tan & Pearl 1994) examine semantics and specification tools for qualitative decision makers, while Darwiche and Goldszmidt (Darwiche & Goldszmidt 1994) experiment with qualitative probabilistic reasoning in diagnostics.

There are two major reasons for this interest in qualitative tools. One reason is computational efficiency: one hopes that qualitative tools, because of their simplicity, will lead to faster algorithms. Another reason is a simpler knowledge acquisition process: often, qualitative information is easier to obtain from experts and layman. However, while there is abundant work on the foundations of quantitative approaches to decision making, usually based on the principle of expected utility maximization (e.g.,(Savage 1972; Anscombe & Aumann 1963; Blum, Brandenburger, & Dekel 1991; Kreps 1988; Hart, Modica, & Schmeidler 1994)), we are aware of very little work on the foundations of qualitative methods.[1]

Work on the foundations of decision theory is motivated by two major applications: agent modeling and decision making. Agent modeling is often the main concern of economists and game-theorists; they ask: under what assumptions can we model an agent as if it were using a particular decision procedure? In artificial intelligence, we share this concern in various areas, most notably in multi-agent systems, where agents must represent and reason about other agents. Decision making is often the main concern of statisticians, decision analysts, and engineers. They ask: how should we model our state of information? And how should we choose our actions? The relevance of this question to AI researchers is obvious. The foundational approach helps answer these questions by describing the basic principles that underlie various decision procedures.

One of the most important foundational results in the area of classical decision theory is Savage's theorem (Savage 1972), described by Kreps (Kreps 1988) as the "crowning achievement" of choice theory. Savage provides a number of conditions on an agent's preference among actions. Under these conditions, the agent's choices can be described as stemming from the use of probabilities to describe her state of information, utilities to describe her preferences over action outcomes, and the use of expected utility maximization to choose her actions. For example, one of Savage's postulates, the *sure-thing principle*, roughly states that: if

[1] An interesting related work is the axiomatic approach taken by Dubois and Prade (Dubois & Prade 1995), which proves the existence of a utility function representing a preference ordering among possibility distributions. Many axiom systems that are weaker than Savage's appear in (Fishburn 1988), but we are not aware of any that resemble ours.

an agent prefers action a over b given that the possible worlds are s_1 and s_2, and she prefers a over b when the possible worlds are s_3 and s_4, then she should still prefer a over b when the possible worlds are s_1, s_2, s_3 and s_4. Economists use Savage's results to understand the assumptions under which they can use probabilities and utilities as the basis for agent models; decision theorists rely on the intuitiveness of Savage's postulates to justify the use of the expected utility maximization principle.

Our aim in this paper is to initiate similar work on the foundations of qualitative decision making. Given that we have compelling practical reasons to investigate such tools, we would like to have as sound an understanding of the adequacy of qualitative decision tools as we do of the classical, quantitative tools; both for the purpose of decision making and agent modeling. Our main contribution is a representation theorem for the *maximin* decision criterion.[2] Using a setting similar to that of Savage, we provide two conditions on an agent's choice over actions under which it can be represented as a qualitative decision maker that uses *maximin* to make its choices. One of these conditions is similar to Savage's sure thing principle. It says that if an agent prefers action a over b when the possible worlds are s_1 and s_2 and she prefers a over b when her possible worlds are s_2 and s_3, then she still prefers a over b when the possible worlds are s_1, s_2 and s_3. The other condition is more technical, and we defer its presentation to Section 4.

Beyond qualitative decision theory, the results presented in this paper have another interesting interpretation: There are different ways in which we can encode an agent's behavior (or program). One simple manner is as an explicit mapping from the agent's local state to actions. This is often highly inefficient in terms of space. Alternative, implicit, representations are often used if we desire to cut down on program storage or transmission costs. Probabilities and utilities and their qualitative counterparts can be used to obtain a compact, albeit implicit, representation of programs. Our (constructive) results characterize a class of agent programs that can be represented in $O(n \log n)$ space, where n is the number of states of the world. This is to be contrasted with a possibly exponential explicit representation.

In Section 2 we define a model of a situated agent and two alternative representations for its program or behavior. One is a simple policy that maps an agent's state of information to actions, while the other represents the agent's program (or behavior) implicitly using the *maximin* decision criterion. Our aim is to present conditions under which policies can be represented implicitly using the *maximin* criterion. This will be carried out in two steps: In Section 3 we discuss the case of an agent which has to decide among two actions in various states, while in Section 4 we consider the case where the agent has any finite number of actions to choose from. Proofs of these results, which are constructive, are omitted due to space constraints. Section 5 concludes with a discussion of some issues raised by our results and a short summary.

[2](Hart, Modica, & Schmeidler 1994) presents an axiomatization for maximin in the context of 2-person zero-sum games. However, their axiomatization is probabilistic, and does not fit the framework of qualitative decision theory.

The Basic Model

In this section, we define a fairly standard agent model and the concept of a policy, which describes the agent's behavior. Then, we suggest one manner for implicitly representing (some) policies using the concept of utility and the decision criterion *maximin*.

Definition 1 *States is a finite set of possible states of the world. An* agent *is a pair of sets, (LocalStates, Actions), which are called, respectively, the agent's set of* local states *and* actions.

$PW : LocalStates \rightarrow 2^{States} \setminus \emptyset$ *is the function describing the set of world states consistent with every local state. PW satisfies the following conditions: (1)* $PW(l) = PW(l')$ *iff* $l = l'$, *and (2) For each subset* S *of States, there exists some* $l \in LocalStates$ *such that* $PW(l) = S$.

Each state in the set *States* describes one possible state of the world, or the agent's environment. This description does not tell us about the internal state of the agent (e.g., the content of its registers). These internal states are described by elements of the set of local states. Intuitively, local states correspond to the agent's possible states of information, or its *knowledge* (see (Fagin *et al.* 1995; Rosenschein 1985)). In addition to a set of possible local states, the agent has a set *Actions* of actions. One can view these actions as the basic control signals the agent can send to its actuators.

With every local state $l \in LocalStates$ we associate a subset $PW(l)$ of *States*, understood as the possible states of the world consistent with the agent's information at l. That is, $s \in PW(l)$ iff the agent can be in local state l when the current state of the world is s. In fact, in this paper we identify l with $PW(l)$ and use both interchangeably. Hence, we require that $l = l'$ iff $PW(l) = PW(l')$ and that for every $S \subseteq States$ there exists some $l \in LocalStates$ such that $PW(l) = S$.

Like other popular models of decision making (e.g., (Savage 1972; Anscombe & Aumann 1963)), our model considers one-shot decision making. The agent starts at some initial state of information and chooses one of its possible actions based on its current state of information (i.e., its local state); this function is called the agent's *policy* (see also *protocol* (Fagin *et al.* 1995) and *strategy* (Luce & Raiffa 1957)). This policy maps each state of information of the agent into an action.

Definition 2 *A* policy *for agent* (*LocalStates*, *Actions*) *is a function* $\mathcal{P}$: *LocalStates* $\rightarrow$ *Actions*.

A naive description of the policy as an explicit mapping between local states and actions is exponentially large in the number of possible worlds because $|LocalStates| = 2^{|States|}$. Requiring a designer to supply this mapping explicitly is unrealistic. Hence, a method for implicitly specifying policies is desirable. In particular, we would like a specification method that helps us judge the quality of a policy. Classical decision theory provides one such manner: the policy is implicitly specified using a probability assignment pr over the set *States* and a real valued utility function u over a set O of action outcomes. The action to be performed at local state l is obtained using the principle of expected utility maximization:

$$argmax_{a \in Actions}\{\sum_{s \in PW(l)} pr(s) \cdot u(a(s))\}$$

where $a(s)$ is the outcome of action a when the state of the world is s. We wish to present a different, more qualitative representation. We will not use a probability function, and our utility function $u(\cdot,\cdot)$ takes both the state of the world and the action as its arguments and returns some value in a totally pre-ordered set. (Notice the use of qualitative, rather than quantitative representation of utilities.) For convenience, we will use integers to denote the relative positions of elements within this set. In our representation, the agent's action in a local state l is defined as:

$$argmax_{a \in Actions}\{\min_{s \in PW(l)}(u(a,s))\}.$$

That is, the agent takes the action whose worst-case utility is maximal. *Maximin* is a qualitative decision criterion that seems well-tailored to risk-averse agents.

Definition 3 *A policy* $\mathcal{P}$ *has a* maximin representation *if there exists a utility function on States* $\times$ *Actions such that for every* $l \in$ *LocalStates*

$$\mathcal{P}(l) = argmax_{a \in Actions}\{\min_{s \in PW(l)}(u(a,s))\}.$$

That is, $\mathcal{P}$ has a *maximin* representation if for every local state l, an agent with this utility function that makes her decision by applying *maximin* to the utilities of actions in $PW(l)$, would choose the action $\mathcal{P}(l)$.

Given an arbitrary agent and a policy $\mathcal{P}$ adopted by the agent, it is unclear whether this policy has a *maximin* representation. It is the goal of this paper to characterize the class of policies that have this representation. From this result, we can learn about the conditions under which we can use the *maximin* representation to model agents and understand the rationality of using *maximin* as a decision criterion. Unlike the exponential naive representation of policies, the *maximin* representation requires only $O(logM \cdot |States| \cdot |Actions|)$ space, where $M = \max_{a,a' \in Actions; s,s' \in States} |u(a,s) - u(a',s')|$.

Representing Binary Decisions

This section presents two representation theorems for *maximin* for agents with two possible actions. We start by describing a basic property of *maximin* representable policies.

Definition 4 *We say that a policy* $\mathcal{P}$ *is* closed under union *if* $\mathcal{P}(U) = \mathcal{P}(W)$ *implies* $\mathcal{P}(U \cup W) = \mathcal{P}(U)$, *where* $U, W \subseteq$ *States*.

That is, suppose that the agent would take the same action a when its local state is l or l', and let $\hat{l}$ be the local state in which the agent considers possible all worlds that are possible in l and in l'. That is $PW(\hat{l}) = PW(l) \cup PW(l')$. If the agent's policy is closed under unions, it would choose the action a at $\hat{l}$.

For example, suppose that our agent is instructed to bring coffee when it knows that the weather is cold or warm and when it knows that the weather is warm or hot. If all the agent knows is that the weather is cold, warm, or hot, it should still bring coffee if its policy is closed under unions. This sounds perfectly reasonable. Consider another example: Alex likes Swiss chocolate, but dislikes all other chocolates. He finds an unmarked chocolate bar and must decide whether or not he should eat it. His policy is such that, if he knows that this chocolate is Swiss or American, he will eat it; if he knows that this bar is Swiss or French, he will eat it as well. If Alex's policy is closed under unions, he will eat this bar even if he knows it must be Swiss, French, or American.

Our first representation theorem for *maximin* shows that policies containing two possible actions that are closed under unions are representable using a utility function defined on *Actions* and *States*.

Theorem 1 *Let* $\mathcal{P}$ *be a policy assigning only one of two possible actions at each local state, and assume that* $\mathcal{P}$ *is closed under union. Then,* $\mathcal{P}$ *is maximin representable.*

Notice that this corresponds to a completeness claim, while soundness, which implies that the above conditions hold for *maximin*, is easily verified.

The following example illustrates our result.

Example 1 *Consider the following policy (or precondition for wearing a sweater) in which Y stands for "wear a sweater" and N stands for "do not wear a sweater".*

{cold}	{ok}	{hot}	{c,o}	{c,h}	{o,h}	{c,o,h}
Y	N	N	Y	N	N	N

It is easy to verify that this policy is closed under unions. For example, the sweater is not worn when the weather is ok or when the weather is either hot or cold, hence it is not worn when there is no information at all, i.e., the weather is either cold, ok, or hot.

Using the proof of Theorem 1 we construct the following utility function representing the policy above:

	cold	*ok*	*hot*
Y	*3*	*2*	*0*
N	*1*	*3*	*3*

A slight generalization of this theorem allows for policies in which the agent is indifferent between the two available choices. In the two action case discussed here, we capture such indifference by assigning both actions at a local state, e.g., $\mathcal{P}(l) = \{a, a'\}$. Hence we treat the policy as assigning sets of actions rather than actions. We refer to such policies as set-valued policies, or s-policies. Closure under union is defined in this context as follows:

Definition 5 *An s-policy $\mathcal{P}$ is* closed under unions *if for every pair of local states $U, W \subseteq States$, $\mathcal{P}(U \cup W)$ is either $\mathcal{P}(W), \mathcal{P}(U)$, or $\mathcal{P}(U) \cup \mathcal{P}(W)$.*

We require a number of additional definitions before we can proceed with the representation theorem for s-policies. First, we define two binary relationships on subsets of *States*:

Definition 6 *$U >_{\mathcal{P}} W$, where $U, W \subseteq States$, if $\mathcal{P}(U \cup W) = \mathcal{P}(U)$ and $\mathcal{P}(U) \neq \mathcal{P}(W)$. $U =_{\mathcal{P}} W$, where $U, W \subseteq States$, if $\mathcal{P}(U), \mathcal{P}(W)$ and $\mathcal{P}(U \cup W)$ are all different.*

That is $U >_{\mathcal{P}} W$ tells us that the preferred action in U is preferred in $U \cup W$. $U =_{\mathcal{P}} W$ is basically equivalent to $U \not>_{\mathcal{P}} W$ and $W \not>_{\mathcal{P}} U$. Next, we define a condition on these relations which closely resembles transitivity.

Definition 7 *We say that $>_{\mathcal{P}}$ is* transitive-like *if whenever $U_1 *_1 \cdots *_{k-1} U_k$, where $*_j \in \{>_{\mathcal{P}}, =_{\mathcal{P}}\}$, and $\mathcal{P}(U_1) \neq \mathcal{P}(U_k)$, we have that $U_1 * U_k$. Here, $*$ is $>_{\mathcal{P}}$ if any of the $*_i$ are $>_{\mathcal{P}}$, and it is $=_{\mathcal{P}}$ otherwise.*

Finally, we say that $\mathcal{P}$ respects domination if the action assigned to the union of a number of sets does not depend on those sets that are dominated by other sets w.r.t. $>_{\mathcal{P}}$.

Definition 8 *We say that $\mathcal{P}$* respects domination *if for all $W, U, V \subseteq States$ we have that $W >_{\mathcal{P}} U$ implies that $\mathcal{P}(W \cup U \cup V) = \mathcal{P}(W \cup V)$.*

We have the following representation theorem for s-policies:

Theorem 2 *Let $\mathcal{P}$ be an s-policy for an agent (LocalStates, Actions) such that (1) $|Actions| = \{a, a'\}$, (2) $\mathcal{P}$ is closed under unions, (2) $\mathcal{P}$ respects domination, (4) $>_{\mathcal{P}}$ is transitive-like. Then, $\mathcal{P}$ is maximin representable.*

A General Existence Theorem

In the previous section we provided representation theorems for a class of policies in which the agent chooses between two actions. We would like to generalize these results to represent choice among an arbitrary set of actions. We will assume that, rather than a single most preferred action, the agent has a total order over the set of actions associated with each local state. This total order can be understood as telling us what the agent would do should its first choice became unavailable. The corresponding representation using *maximin* will tell us not only which action is most preferred, but also, which action is preferred to which.

Definition 9 *A* generalized policy *for an agent (LocalStates,Actions) is a function $\mathcal{P}$: LocalStates $\rightarrow TO(Actions)$, where $TO(Actions)$ is the set of total orders on Actions.*

Generalized policy $\mathcal{P}$ is maximin representable *if there exists a utility function $u(\cdot,\cdot)$ on States$\times$Actions such that a is preferred to a' in local state l according to $\mathcal{P}(l)$ iff*

$$\min_{s \in PW(l)}(u(a,s)) > \min_{s \in PW(l)}(u(a',s))$$

for every pair of actions $a, a' \in Actions$ and for every local state $l \in LocalStates$.

The generalization of closure under unions to generalized policies is not a sufficient condition on policies for obtaining a *maximin* representation. The following definition introduces an additional property needed:

Definition 10 *Let $\{\succ_W \mid W \subseteq S\}$, be a set of total orders over Actions. Given $s, s' \in States$ and $a, a' \in Actions$, we write $(s, a) < (s', a')$ if (1) $a' \succ_s a$, $a \succ_{s'} a'$, and $a' \succ_{\{s,s'\}} a$; or (2) $s = s'$ and $a' \succ_s a$. We say that $<$ is* transitive-like *if whenever $(s_1, a_1) < (s_2, a_2) < \cdots < (s_k, a_k)$ and either (1)*

s	s'	$s \cup s'$
a'	a	a'
a	a'	a

	s	s'
a	1	3
a'	3	2

Figure 1: $(s, a) < (s', a')$

$a_k \succ_{s_1} a_1$ *and* $a_1 \succ_{s_k} a_k$ *or (2)* $s_1 = s_k$, *then* $(s_1, a_1) < (s_k, a_k)$.

The left table in Figure 1 helps us clarify this definition. In it, we depict the conditions under which $(s, a) < (s', a')$ holds. There are three columns in this table, each showing the agent's preference relation over actions in different local states. The possible worlds in these local states are s, s', and $\{s, s'\}$. In s the agent prefers a' over a, in s' it prefers a over a', but when all the agent knows is that the world is either in state s or s', it prefers a' over a. Roughly, we can say that $(s, a) < (s', a')$ if the agent dislikes taking action a in state s more than it dislikes taking action a' in state s'.

The following example illustrates the transitivity-like condition.

Example 2 *Suppose that there are three possible states of the world: snowing and cold, raining and cold, neither and warm. I prefer skiing to walking when it is snowing, but prefer walking to skiing when it is raining. However, when I am uncertain about whether it will rain or snow, I'd choose to walk. In this case* $(ski, rain) < (walk, snow)$. *I prefer skiing to jogging when it is warm, and I prefer jogging to skiing when it is raining. However, I really dislike jogging when it is not cold, so I prefer skiing to jogging if I am uncertain whether it is warm or snowing. Hence* $(jog, warm) < (ski, rain)$. *Suppose that, in addition, I prefer walking to jogging when it is warm, and I prefer jogging to walking when it snows. The transitivity-like condition implies that* $(jog, warm) < (walk, snow)$, *and hence I'd prefer walking to jogging if I am uncertain whether it will be warm or it will snow. This seems quite plausible.*

Theorem 3 *Let Actions be an arbitrary set of actions, and let* $\succ_W$, *for every* $W \subseteq S$, *be a total order over Actions such that*

1. *if* $a \succ_W a'$ *and* $a \succ_V a'$ *then* $a \succ_{W \cup V} a'$, *and*
2. $<$ *is transitive-like.*

Then, $\{\succ_W \mid W \subseteq S\}$ *is maximin representable.*

Again, it is easy to see that a preference relation based on *maximin* will have the properties described in this theorem, and this result can be viewed as a sound and complete characterization of the *maximin* criterion for total orders. In addition, this theorem characterizes a class of policies that can be represented using $O(n \cdot log(n))$ space in contrast with the exponentially large naive representation.

Discussion

Decision theory is clearly relevant to AI, and there is little doubt about the need for decision making techniques that are more designer friendly and have nice computational properties. Qualitative decision procedures could offer such an alternative, but the question is: how rational are they? One method of addressing this question is experimentation, as in (Darwiche & Goldszmidt 1994). However, the prominent approach for understanding and justifying the rationality of decision criteria has been the axiomatic approach. This approach characterizes the properties of a decision criterion in a general, domain independent manner. Given a particular domain of application, we can assess the rationality of employing a particular decision criterion using its characteristic properties. Our work provides one of a few results within the axiomatic approach that deals with qualitative decision criteria and helps us understand the inherent properties of *maximin*, assess the rationality of using *maximin*, and understand the conditions under which an arbitrary agent can be modeled *as if* it were a qualitative decision maker.

In classical decision theory, the agent has both a utility function and a probability function. In our representation theorems, the emphasis has been on utilities rather than beliefs. The agent's state of information is modeled by means of the set of worlds consistent with its current local state, $PW(l)$. Most authors (e.g., (Fagin *et al.* 1995)) regard this set as representing the agent's knowledge, rather than belief. However, the concept of belief can be incorporated into this model by imposing additional structure on the set *States* in the form of a ranking function. This model has been suggested by e.g., (Brafman & Tennenholtz 1994; Friedman & Halpern 1994; Lamarre & Shoham 1994). Given a ranking function $r : States \rightarrow N$, we define the agent's beliefs at local state l as:

$$B(l) = \{s \in PW(l) \mid s' \in PW(l) \text{ implies } r(s) \leq r(s')\}.$$

$B(l)$ are often called the agent's *plausible* states at the local state l. We can modify *maximin* by applying it to the plausible states, instead of the possible states (see, e.g.,(Brafman & Tennenholtz 1994)). That is, at state l the agent chooses

$$argmax_{a \in Actions}\{\min_{s \in B(l)} u(a, s)\}.$$

A similar approach is taken in (Boutilier 1994; Tan & Pearl 1994)).

Clearly, any behavior that is *maximin* representable can be represented using the ranked *maximin* representation suggested above. (We would use a ranking function that maps all states to the same integer). We can show that the converse is true as well. That is, if an agent can be represented as using ranked *maximin* it can also be represented as using the standard *maximin* approach discussed in this paper; a formal proof is deferred to the full paper. Therefore, the ranked *maximin* representation is no more expressive than our standard maximin representation, i.e., it can capture the same set of behaviors. Hence, ranked *maximin* is not, *a priori*, a more rational decision criterion.

Our work differs from most other foundational work in decision theory in its definition of the utility function. We define utilities as a function of both the agent's action and the state of the world. Savage, and many others, define the notion of an *outcome*, i.e., a description of the state of the world following the performance of an action. In these works, utilities are a function of outcomes. Savage defines actions as mappings between states and outcomes, and it is possible to obtain the same outcome when two different actions are performed in two different states of the world. Our approach is motivated by the fact that, in practice, an agent chooses an action, not an outcome. That is, the only physically observable aspect of the agent's behavior is its choice of action, e.g., the control signal it sends to its actuators. The outcome of these actions is not directly chosen by the agent. Our representation is identical to the standard representation if it is assumed that the outcomes of different actions on different states are different. Moreover, using utility functions that depend on both the state and the action makes practical sense in our qualitative context: it is reasonable when the manner in which the outcome was received is important, e.g., the cost of an action, and it allows us to use the utility function to encode both the desirability of the action's outcomes and the likelihood of the state in which it is obtained. Nevertheless, obtaining representation theorems for *maximin* in the more standard framework is an interesting challenge.

Acknowledgments: Comments from Joe Halpern, Daniel Lehmann, and the anonymous referees provided much help in improving the content and presentation of this paper. We thank Ehud Kalai, Dov Monderer, and Ariel Rubinstein for useful comments and pointers to related work in other disciplines.

References

Anscombe, F. J., and Aumann, R. J. 1963. A definition of subjective probability. *Annals of Mathematical Statistics* 34:199–205.

Blum, L.; Brandenburger, A.; and Dekel, E. 1991. Lexicographic probabilities and equilibrium refinements. *Econometrica* 59:61–79.

Boutilier, C. 1994. Toward a Logic for Qualitative Decision Theory. In *Proc. of the 4th Int. Conf. on Prin. of Knowledge Rep. and Reas.*, 75–86.

Brafman, R. I., and Tennenholtz, M. 1994. Belief ascription and mental-level modelling. In *Proc. of the 4th Int. Conf.. on Princ. of Knowledge Rep. and Reas.*, 87–98.

Darwiche, A., and Goldszmidt, M. 1994. On the relation between kappa calculus and probabilistic reasoning. In *Proc. 10th Conf. on Uncertainty in AI*, 145–153.

Dubois, D., and Prade, H. 1995. Possibility Theory as a Basis for Qualitative Decision Theory. In *Proc. 14th International Joint Conference on Artificial Intelligence*, 1924–1930.

Fagin, R.; Halpern, J. Y.; Moses, Y.; and Vardi, M. Y. 1995. *Reasoning about Knowledge.* MIT Press.

Fishburn, P. C. 1988. *Nonlinear Preference and Utility Theory.* Johns Hopkins University Press.

Friedman, N., and Halpern, J. Y. 1994. A knowledge-based framework for belief change. Part I: Foundations. In *Proc. of the 5th Conf. on Theoretical Aspects of Reas. About Knowledge.*

Hart, S.; Modica, S.; and Schmeidler, D. 1994. A neo bayesian foundation of the maxmin value for two-parson zero-sum games. *International Journal of Game Theory* 23:347–358.

Kreps, D. M. 1988. *Notes on the Theory of Choice.* Boulder: Westview Press.

Lamarre, P., and Shoham, Y. 1994. Knowledge, certainty, belief and conditionalization. In *Proc. of 4th Intl. Conf. on Principles of Knowledge Rep. and Reas.*

Luce, R. D., and Raiffa, H. 1957. *Games and Decisions.* New York: John Wiley & Sons.

Rosenschein, S. J. 1985. Formal Theories of Knowledge in AI and Robotics. *New Generation Computing* 3(3):345–357.

Savage, L. J. 1972. *The Foundations of Statistics.* New York: Dover Publications.

Tan, S., and Pearl, J. 1994. Specification and Evaluation of Preferences under Uncertainty. In *Proc. of the 4th Int. Conf. on Principles of Knowledge Rep. and Reas.*, 530–539.