On the Axiomatization of Qualitative Decision Criteria

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

Qualitative decision tools have been used in AI and CS in various contexts, but their adequacy is still unclear. To examine this question, our work employs the axiomatic approach to characterize the properties of various decision rules. In the past, we presented a constructive representation theorem for the *maximin* decision criterion, and we characterized conditions under which an agent can be viewed as adopting a qualitative decision-making approach (consisting of beliefs, goals, and a qualitative decision criterion). In this paper we show that the *maximin* representation theorem applies to two additional decision criteria: *minmax regret* and *competitive ratio*, and with slight modifications, to a third one, *maximax*. In addition, we characterize conditions under which an agent with a given qualitative utility function can be ascribed beliefs when we assume it adopts *maximin* as its decision criterion.

Introduction

Decision theory plays a central role in various disciplines, including economics, game theory, operations research, industrial engineering, and statistics. It is widely recognized by now that decision making is crucial to AI as well, since artificial agents are, in fact, automated decision makers (Russel & Norvig 1995). However, many decision making techniques found in the AI literature are quite different from those found in other fields. Work in other disciplines has mostly adopted the view of agents as expected utility maximizers. However, these fields have paid little attention to the automation of mundane decision making with its inherent difficulties: knowledge representation, cost of computation, and knowledge elicitation. AI researchers faced with these difficulties have often resorted to more qualitative decision making techniques because they have felt that such tools could simplify the tasks of knowledge acquisition and may lead to faster algorithms in certain contexts.

The magnitude of the problems we face in automating the process of decision making makes qualitative approaches attractive. Yet, despite their intuitive appeal, little is known about their suitability. In particular, two questions arise: How rational are different qualitative decision criteria? (Or put differently, when should they be employed?) And when can we model an agent as a qualitative decision maker?

Economists, statisticians, and others have made great efforts to address these issues in the context of classical decision theory. In particular, in what can be considered as the most fundamental work in the theory of choice, Savage (1972) shows conditions on the agent's choice among actions under which it can be modeled as an expected utility maximizer. In fact, Savage provides a representation theorem which answers both of the above questions for the case of expected utility maximization. We aspire to provide similar foundations to qualitative decision making.

In previous work (Brafman & Tennenholtz 1996), we provided a representation theorem for the *maximin* decision criterion – a central qualitative decision rule. In addition, in (Brafman & Tennenholtz 1994) we presented a general mental-level model of agents as qualitative decision makers. In the framework of this model, we gave conditions under which one can ascribe qualitative beliefs to an agent. In this paper we extend these studies in two directions:

1. We show that the *maximin*, *minmax regret*, and *competitive ratio* decision criteria have the same expressive power, i.e., they can represent the same sets of preferences. Consequently, we can lift our previous axiomatization of *maximin* to the other two criteria. In addition, we show that a similar set of properties characterizes the *maximax* criterion.

2. We present sound and complete conditions for the ascription of beliefs (captured by an acyclic order on the states of the environment) for *maximin* agents with arbitrary qualitative utilities. Previous work considered only the case of 0/1 utilities.

In the following section, we discuss the decision criteria investigated in this paper. In Section 3, we show that a number of these decision criteria have the same expressive power. Using this equivalence, we lift our axiomatization of *maximin* to the other decision criteria. In Section 4, we follow these results with a discussion of some of the basic properties of these decision criteria in the attempt to understand them better. In Section 5, we discuss the problem of ascribing beliefs to a *maximin* agent based on its policy and goals. Section 6 concludes the paper.

Qualitative Decision Criteria

We start with a general model of decision making with incomplete information.

Definition 1 *An* environment *is associated with a set of (environment)* states S. *An agent is associated with a pair* (L, A), *where L is a set of* local states *and A is a set of* actions *available to the agent. A policy of the agent is a function* $P : L \to TO(A)$, *where* $TO(A)$ *is the set of total orders on actions.*

This model captures a general agent-environment pair. The local state of the agent captures its knowledge state, and its policy captures the action it would select in any given local state; the policy specifies the agent's preferences over actions in each local state. Hence, this generalized notion of a policy describes what the agent would do if its favorite action became unavailable, and so on. Following work in knowledge theory (Halpern & Moses 1990; Rosenschein 1985), we identify each local state l with a subset $PW(l)$ of the set S. $PW(l)$ is the set of *possible worlds* in local state l. For ease of exposition we assume $L = 2^S$. That is, there is a local state corresponding to each subset of environment states. Our study and results can be extended to the case where we replace the total orders on actions by total pre-orders on actions.

A naive representation of the agent's policy might be exponential in the number of elements of S. Moreover, the explicit definition of a policy might not capture the rationale of action selection by the agent. In order to address these problems, one can consider *decision-theoretic* representations of a policy. The classical decision-theoretic representation of a policy is by means of expected utility maximization. According to the expected utility maximization decision rule, the agent has a probability distribution on the set of states and a utility function assigned to the various outcomes of the actions; based on these, it selects an action which maximizes its expected utility. Yet, there are more qualitative decision-making techniques, as well. We now define four central decision criteria which differ from the purely probabilistic and quantitative form of expected utility maximization. Each of these decision criteria takes two parameters: a utility function U defined on $S \times A$, and a local state l; it returns a total order over A describing the agent's preference for actions in l. For convenience, we assume that the utility function maps elements of $S \times A$ into the integers, although we do not always need all of their algebraic structure (e.g., for *maximin* a mapping to any pre-ordered set would do).

Definition 2 *Given a utility function U on* $S \times A$ *and a local state l, the* maximin *decision criterion selects an action*

$$a = \arg\max_{a' \in A}\{\min_{s \in PW(l)} U(s, a')\}.$$

Maximin is a conservative decision criterion. It optimizes the worst-case outcome of the agent's action.

Definition 3 *Given a utility function U on* $S \times A$, *a state* $s \in S$, *and an action* $a \in A$, *define* $R(s,a) = \max_{a' \in A}(U(s,a') - U(s,a))$. *In local state l, the* minmax regret *decision criterion selects an action*

$$a = \arg\min_{a' \in A}\{\max_{s \in PW(l)} R(s, a')\}.$$

$R(s,a)$ measures the agent's "regret" at having chosen a when the actual state of the world is s. That is, it is the difference in utility between the optimal outcome on state s and the outcome resulting from the performance of a in s. *Minmax regret* attempts to minimize the maximal regret value over all possible worlds.

Definition 4 *Given a utility function U on* $S \times A$, $s \in S$, *and* $a \in A$, *define* $R(s,a) = \max_{a' \in A}(\frac{U(s,a')}{U(s,a)})$.[1] *In local state l, the* competitive ratio *decision criterion selects an action*

$$a = \arg\min_{a' \in A}\{\max_{s \in PW(l)} R(s, a')\}.$$

Much like *minmax regret* the *competitive ratio* criterion attempts to optimize behavior relative to the optimal outcome. The only difference is that here we are interested in ratio, rather than difference.

For completeness, we include a treatment of the somewhat less interesting *maximax* criterion:

Definition 5 *Given a utility function U on* $S \times A$ *and a local state l, the* maximax *decision criterion selects an action*

$$a = \arg\max_{a' \in A}\{\max_{s \in PW(l)} U(s, a')\}.$$

To illustrate these rules, consider the following decision matrix, each action of which would be chosen by a different decision criterion:

	s_1	s_2	chosen by:
a_1	60	10	*minmax regret*
a_2	40	20	*competitive ratio*
a_3	30	21	*maximin*
a_4	70	1	*maximax*

Maximin and *minmax regret* are two of the most famous qualitative decision criteria discussed in the decision theory literature (Luce & Raiffa 1957; Milnor 1954). The *competitive ratio* decision rule is extremely popular in the theoretical computer science literature (e.g., (Papadimitriou & Yannakakis 1989)) where it is used as the primary optimization measure for on-line algorithms. As a result, a representation theorem which teaches us about the conditions under which an agent can be viewed as using each of these decision criteria may be a significant step in our understanding of qualitative decision making. Moreover, aside from its direct interest to AI, it may give us better insight as to the validity of current practices in assessing on-line algorithms.

[1] For ease of exposition we assume that utilities are greater than 0; in particular, the division is well-defined.

Notice that some of these decision criteria are more "qualitative" than the others. *Maximin* and *maximax* consider only the order relation between utilities; *minmax regret* and *competitive ratio* are more quantitative, since they care about the actual numbers, their difference, or ratio. Moreover, unlike the expected utility criterion, the four criteria discussed in this paper do not require a quantitative measure of likelihood. In addition, it will be evident from the representation theorems that follow that we can restrict our attention to integer valued utilities when we use these decision criteria. Finally, notice that all four decision criteria use space polynomial in the number of states and actions to represent the agent's preferences.

Axiomatization

Having defined a general agent-environment model and several basic decision-theoretic models, we wish to find conditions under which one can transform a policy into an equivalent decision-theoretic representation.

Definition 6 *A policy P is* maximin representable *if there exists a utility function $u(\cdot,\cdot)$ on $S \times A$ such that a is preferred to a' in local state l iff*

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

for every pair $a, a' \in A$ and for every local state $l \in L$.

The corresponding definition for *maximax* is obtained when we replace the *min* operator with the *max* operator in the definition above.

Definition 7 *A policy P is* minmax regret/competitive ratio representable *if there exists a utility function $u(\cdot,\cdot)$ on $S \times A$ such that a is preferred to a' in local state l iff*

$$\max_{s\in PW(l)} R(s,a) < \max_{s\in PW(l)} R(s,a')$$

for every pair $a, a' \in A$ and for every local state $l \in L$.

Notice that the definitions for the *minmax regret* and the *competitive ratio* representations are similar. The difference stems from the way $R(s,a)$ is defined in these cases. Notice that the utility function assigns natural numbers to the elements of $S \times A$. Given these utilities the agent applies the *min* and *max* operators to select its favorite actions.

The question is under which conditions a policy is *maximin/minmax regret/competitive ratio/maximax* representable. First, we show that the first three representations have the same expressive power. That is, any policy that is representable by one, is representable by the others.

Theorem 1 *A policy is maximin representable iff it is* minmax regret/competitive ratio representable

Proof (partial sketch): In order to prove this theorem it is sufficient to show that if there exists a utility function under which one criterion represents some policy P, we can generate another utility function under which the other criterion will generate P as well.

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

	s	s'
a	1	3
a'	3	2

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

First, notice that *minmax regret* and *competitive ratio* are equivalent. One talks about ratios while the other about differences. Consequently, through the use of logarithms and exponentiation we can transform a utility function for *minmax regret* to an equivalent utility function for *competitive ratio* and vice versa. (Given that S is finite, obtaining an integer valued function is then easy.)

Next, we show that *maximin* and *minmax regret* are equivalent. Notice that *maximin* and a similar (not so rational) criterion *minimax* that attempts to minimize the maximal utility of the action are equivalent. We simply need to multiply utilities by -1. Consequently, it is sufficient to show that *minmax regret* and *minimax* are equivalent. To do this, we have to show that given a regret matrix, we can generate a utility function that has these regret values, and this is straightforward. ∎

Given this result, it follows that the representation theorems for *maximin* from (Brafman & Tennenholtz 1996) hold for the two other criteria, as well. First, we recall the following definition.

Definition 8 *Let $\{\succ_W \mid W \subseteq S\}$, be a set of total orders over A (i.e., a policy). Given $s, s' \in S$ and $a, a' \in A$, 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) $a_k \succ_{s_1} a_1$ and $a_1 \succ_{s_k} a_k$ or
(2) $s_1 = s_k$,
then $(s_1,a_1) < (s_k,a_k)$.

A more intuitive depiction of the $<$ relation appears in Figure 1. The left table describes the agent's preferences given $\{s\}, \{s'\}, \{s,s'\}$ when $(s,a) < (s',a')$. The right table shows one utility function which, under *maximin*, would have led the agent to have such preferences. There, we see that, indeed, $u(s,a) < u(s',a')$.

In the context of *maximax*, a slight change to the definition of the $<$ relation is needed: $(s,a) < (s',a')$ 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 can now prove:

Theorem 2 *Let A be an arbitrary set of actions, and let $\succ_W$, for every $W \subseteq S$, be a total order such that the following holds:*

Closure under unions: *For all $V, W \subseteq S$, if $a \succ_W a'$ and $a \succ_V a'$ then $a \succ_{W\cup V} a'$.*

T: $<$ *is transitive-like.*

Then, the policy described by $\{\succ_W \mid W \subseteq S\}$ *is maximin/minmax regret/competitive ratio/maximax representable.*

Proof: The result for *minmax regret* and *competitive ratio* follows from Theorem 1. For *maximax* an explicit construction is needed based on ideas from (Brafman & Tennenholtz 1996). Details are left to the full paper. ∎

The proof is constructive. That is, given a policy ρ that satisfies the above conditions and a choice of one of these decision criteria, we can construct a utility function which, if adopted by the agent, will lead it to behave as if ρ was its policy.

It is not hard to show that the above conditions hold for all four decision criteria. That is any policy that results from the use of these decision criteria will have these properties. Hence, we have a representation theorem for all four decision criteria. We note that the algorithms used for ascribing utilities to the agent differ depending on the decision criterion chosen. Nevertheless, the similar conditions serve to characterize all the four criteria (although for *maximax*, the definition of $<$ is slightly different). In (Brafman & Tennenholtz 1996) we present a similar representation theorem that covers the case where we allow the agents to express indifference among actions. These results can be extended to cover the three other decision criteria in a similar manner.

Interpreting the Results

What is the significance of these results? First, they imply that from a modeling perspective, all four decision criteria are similar. Any agent whose choice behavior can be modeled using *maximin*, *minmax regret*, or *competitive ratio* can be modeled using any of the other decision criteria. However, these models will differ in the utility function they use. Second, these results expose the fundamental properties of these choice criteria. We see two major characteristic properties: The T property seems very natural to us. It can be viewed as imposing a weak transitivity requirement on the values used to represent utilities. The more central property of the four decision criteria is *closure under unions*: if given a set V of possible worlds the agent prefers action a over a', and given another set W of possible worlds the agent prefers a over a' as well, then it prefers a over a' given $V \cup W$. When V and W are disjoint, we obtain a property analogous to Savage's *sure-thing principle* (Savage 1972). In this restricted form, this property seems essential when we assume that actions are deterministic and all uncertainty about their effects is modeled as uncertainty about the state of the world (as we do here).[2] Closure under disjoint unions is a basic property of another decision criterion, Laplace's *principle of indifference* in which the action maximizing the sum of utilities is preferred.

[2]However, when "actions" represent multi-step conditional plans, during whose execution the agent's state of information can change, this is no longer true.

When the sets V and W are no longer disjoint, closure under unions is a somewhat less natural property of a rational decision maker. To understand this, we define the *column duplication* property (Milnor 1954).

Definition 9 *A decision criterion has the* column duplication *property if whenever it prefers an action a over an action a' given a set V of possible worlds, it prefers a over a' given any set of possible worlds $V \cup \{s\}$, where $s \notin V$ is such that there exists some $s' \in V$ such that $U(\hat{a}, s) = U(\hat{a}, s')$ for all actions $\hat{a} \in A$.*

Intuitively, the column duplication property asserts that the agent's preferences do not change if it considers another state possible which is identical, in terms of its effects, to some existing state. Note that column duplication is just the flip side of an equivalent property allowing for removal of one of two identical columns.

It is easy to see that:

Lemma 1 *Given an infinite set of states, a decision criterion is closed under unions of finite states iff it is closed under disjoint unions of finite states and has the* column duplication *property.*

In the above, we need an infinite number of states in order to allow us to freely duplicate states.

It has been observed that column duplication is a basic property of all these decision criteria (Milnor 1954). For instance, here is the table presented in Section 2 with the column corresponding to s_2 duplicated. As can be seen, the four decision criteria we have discussed choose the same action as before.

	s_1	s_2	s_3	chosen by:
a_1	60	10	10	*minmax regret*
a_2	40	20	20	*competitive ratio*
a_3	30	21	21	*maximin*
a_4	70	1	1	*maximax*

Whether or not column duplication is reasonable depends on the state of information of the agent and the conceptualization of the domain. It has been suggested that this property is characteristic of states of complete ignorance (Luce & Raiffa 1957).

It is interesting to note that another well-known qualitative decision criterion, Hurwicz's criterion, does not satisfy the property of closure under disjoint unions (although it has the column duplication property). Hurwicz's criterion is the following generalization of *maximin* and *maximax*:

Definition 10 *Given a utility function U on $S \times A$ and a local state l, the* Hurwicz *decision criterion selects an action a such that*

$$a = \arg\max_{a' \in A}\{(\alpha \cdot \min_{s \in PW(l)} U(s, a')) + ((1-\alpha) \cdot \max_{s \in PW(l)} U(s, a')\}.$$

When $\alpha = 1$ we obtain the *maximin* criterion, and when $\alpha = 0$ we obtain *maximax*. The following matrix is a counterexample to closure under disjoint unions.

	s_1	s_2	s_3	s_4
a_1	50	50	4	-1
a_2	50	10	1	1

Suppose that $\alpha = 0.5$. Under Hurwicz's criterion, a_1 is preferred over a_2 given either $\{s_1, s_2\}$ or $\{s_3, s_4\}$. However, given $\{s_1, s_2, s_3, s_4\}$, a_2 is preferred over a_1.

In the literature (e.g., (Luce & Raiffa 1957)), one finds various examples of counterintuitive choices made by various qualitative criteria in various settings. For instance, one can argue against *maximin* using the following matrix:

	s_1	s_2	...	s_{99}	s_{100}
a_1	1	1	1	1	1
a_2	1000	1000	1000	1000	0

Under *maximin*, the first action will be preferred, and this seems counterintuitive. While it is not our goal to advocate *maximin* we wish to point out a certain problem with such examples; a problem which lies with the meaning of the numbers used within the decision matrix. If the numbers in the matrix above correspond to dollar amounts, then *maximin* may not make much sense. For each of the qualitative decision criteria, one can construct such counterintuitive matrices. However, in many AI contexts, we are not concerned with monetary payoffs. In that case, one may suppose that the numbers used signify utilities. Yet, the concept of *utility* is meaningless unless it is specified in the context of a decision criterion. For example, the standard notion of utility is *tailored* for expected utility maximizers, and it is somewhat awkward to use it in the context of a *maximin* agent. Of course, once we interpret these values as utilities ascribed to a *maximin* agent, this example is no longer counterintuitive.

Belief Ascription for Maximin Agents

We could improve qualitative decision making by incorporating some notion of likelihood. So far, states could either be possible or impossible. Some authors (e.g., (Boutilier 1994)) have considered qualitative decision making in the context of rankings, which help us distinguish between plausible and implausible possible worlds. Decisions are made by taking into account only the plausible worlds. As we show in the full paper, this approach does not lead to richer choice behaviors. That is, any choice behavior that can be modeled using such rankings and one of the four decision criteria discussed in this paper, can be modeled without using rankings. However, when we consider decision making or modeling given a *fixed* utility function, finer notions of belief can help us make better decisions (or equivalently, model additional behaviors). In this section, we characterize one context in which we can model agents' beliefs using richer belief structures together with the *maximin* decision criterion. Aside from our more theoretical interest in the foundations of qualitative decision theory, it is worth mentioning that belief ascription has various more practical applications, e.g., in predicting agents' future behavior (Brafman & Tennenholtz 1995).

Definition 11 *Let S, L, A, U be defined as in the previous sections; let $R \subseteq S \times S$ be an acyclic binary relation among states. We denote the fact that $R(s, s')$ holds by $s < s'$.*

Given $l \in 2^S$, let $B(l) = \{s \in PW(l) : \nexists s' \in PW(l)$ s.t. $s' < s\}$. A policy P is bel-maximin representable *if we can find U and R such that $a >_l a'$ iff*

$$\min_{s \in B(l)} U(a, s) > \min_{s \in B(l)} U(a', s)$$

for every $a, a' \in A$ and every $l \in L$.

R provides a minimal notion of plausibility on states, where $s < s'$ implies that s is more plausible than s'. $B(l)$ represents the agent's beliefs at l, and it contains the most plausible of its possible worlds, $PW(l)$. The definition of *bel-maximin* representable policies mimics that of *maximin* representable policies. However, only states in $B(l)$ are considered in the minimization process.

This modified agent model raises several basic questions, one of which is the problem of belief ascription: Assuming we are given a policy P and a corresponding utility function U, can we find an acyclic (belief structure) R such that the policy P is bel-maximin representable using U and R? Again, we look for conditions on the agent's policy under which it can be ascribed appropriate beliefs. A representation theorem has been presented in this context only for the case of 0/1 utilities (Brafman & Tennenholtz 1994).

Definition 12 *We write $s <_P s'$ if there exists a, a' such that: (1) $a <_{\{s,s'\}} a'$, (2) $U(a', s') < \min(U(a, s), U(a, s'))$, and (3) $U(a, s) < U(a', s)$.*

Intuitively, $s <_P s'$ is telling us that the policy P treats the state s as more plausible than s' because it is ignoring the effects of the actions on s'. In particular, despite the fact that conditions (2) and (3) would have led a pure *maximin* criterion (that does not ignore any state) to choose action a in $\{s, s'\}$, the policy P prefers a'.

In the sequel we assume that all the elements in the range of U are distinct. Consider the axioms *BEL*:

1. $<_P$ is acyclic.
2. Let $MPW(l)$ be the minimal elements in $PW(l)$ according to $<_P$. Assume $a' <_l a$, and let s be the state in $MPW(l)$ where a' gets the minimal value. Then, for every $t \in MPW(l)$ we have that $a' <_{\{s,t\}} a$.

The first axiom implies that we can talk about minimal elements within sets of possible worlds. The second axiom says that if I prefer a to a' in a particular context l, then I prefer a to a' in any subcontext of l consisting of two possible worlds, one of which is the least desirable world in this context w.r.t. action a.

It is not hard to see that the BEL axioms would hold for an agent using the bel-*maximin* criterion:

Theorem 3 *Given a bel-maximin representable policy P and a corresponding utility function U, then the* BEL *axioms are satisfied.*

That is, if given a utility function U we can find a belief function R such that U and R represent P, it must be the case that P satisfies the *BEL* axioms. The above theorem is a soundness result. It is less obvious that the other direction holds:

Theorem 4 *Let P be a policy and U a utility function such that P and U satisfy the BEL axioms. There exists a relation R on $S \times S$ such that U and R provide a bel-maximin representation of P.*

Thus, we see that the BEL axioms characterize the conditions under which an agent can be ascribed a qualitative model of belief.

Conclusion

The topic of qualitative decision making is receiving growing attention within AI (e.g., see (Boutilier 1994; Tan & Pearl 1994; Darwiche & Goldszmidt 1994; Dubois & Prade 1995)). However, the foundations of qualitative decision making has been pretty much ignored until the recent work of Brafman and Tennenholtz (1994; 1996), Dubois and Prade (1995) (who examine a possibilistic analogue of the von Neumann-Morgenstern theory of utility), and Lehmann (1996) (who provides an axiomatization for generalized qualitative probabilities). These works attempt to understand the foundations of qualitative representations of beliefs and preferences via the axiomatic approach, which has been extensively used in the investigation of expected utility models (e.g., (Fishburn 1988; Savage 1972; Anscombe & Aumann 1963)).

This paper extends previous work on the axiomatization of qualitative decision criteria by showing its applicability to *minmax regret* and *competitive ratio*, two central qualitative decision criteria, as well as to *maximax*. A much earlier effort of Milnor (Milnor 1954; Luce & Raiffa 1957) led to a number of important results on the properties of various qualitative decision criteria. However, Milnor assumes that the decision matrix is given. Our representation theorems do not make this assumption; thus, they are more fundamental. Proofs of the theorems as well as similar theorems for the more general case of total pre-orders are omitted from this abstract, but they are constructive: Given the conditions of Theorem 2, and each particular decision criterion, there exists an efficient algorithm which transforms the agent's policy into a succinct decision theoretic representation. While the algorithms used for the different decision criteria are different, the axiomatizations in all four cases considered in this paper is essentially the same. Although, technically, some of these results were not hard to prove, we did not anticipate them when we set out to understand these criteria. Our study of belief ascription complements previous work by providing sound and complete conditions under which an agent can be ascribed beliefs, given the agent's qualitative utility function and the fact that it uses *maximin* as its decision criterion.

Acknowledgment: We wish to thank the anonymous referees for their useful suggestions. The first author is supported by IRIS project IC-7 and NSERC grants OGPO044121 and A9281. The second author is partially supported by the Gutwirth fellowship and by Technion VPR fund 191-192.

References

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

Boutilier, C. 1994. Toward a Logic for Qualitative Decision Theory. In *Proc. KR&R '94*, 75–86.

Brafman, R. I., and Tennenholtz, M. 1994. Belief ascription and mental-level modelling. In *Proc. KR&R '94*, 87–98.

Brafman, R. I., and Tennenholtz, M. 1995. Towards action prediction using a mental-level model. In *Proc. 14th IJCAI*.

Brafman, R. I., and Tennenholtz, M. 1996. On the Foundations of Qualitative Decision Theory. In *Proc. AAAI '96*.

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

Dubois, D., and Prade, H. 1995. Possibility Theory as a Basis for Qualitative Decision Theory. In *Proc. 14th IJCAI*, 1924–1930.

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

Halpern, J. Y., and Moses, Y. 1990. Knowledge and common knowledge in a distributed environment. *J. ACM* 37(3):549–587.

Lehmann, D. 1996. Generalized qualitative probability: Savage revisited. In *Proc. 12th UAI*, 381–388.

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

Milnor, J. 1954. Games Against Nature. In Thrall, R. M.; Coombs, C.; and Davis, R., eds., *Decision Processes*. John Wiley & Sons.

Papadimitriou, C., and Yannakakis, M. 1989. Shortest Paths Without a Map. In *Automata, Languages and Programming. 16th Int. Col.*, 610–620.

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

Russel, S., and Norvig, P. 1995. *Artificial Intelligence: A Modern Approach*. Prentice Hall.

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. KR&R '94*, 530–539.