Choice under Complete Uncertainty when Outcome Spaces are State Dependent

Clemens Puppe[1] Karl H. Schlag[2]

Revised Version, May 2007
(First Version, September 2006)

[1]Institut für Wirtschaftstheorie und OR, Universität Karlsruhe, puppe@wior.uni-karlsruhe.de

[2]Economics Department, European University Institute, schlag@iue.it

Abstract

One central objection to the maximin payoff criterion is that it focuses on the state that yields the lowest payoffs regardless of how low these are. We allow different states to have different sets of possible outcomes and show that the original axioms of Milnor (1954) continue to characterize the maximin payoff criterion, provided that the sets of payoffs achievable across states overlap. If instead payoffs in some states are always lower than in all others then ignoring the "bad" states is no longer inconsistent with these axioms. Similar dependence on overlap of outcome spaces across states holds for the minimax regret and maximin joy criteria.

1 Introduction

It is difficult to trade-off risks in environments where probabilities are unknown. The maximin payoff criterion, and similarly the minimax regret criterion, have axiomatic foundations that may help us in assessing risk trade-offs in these situations. However, in most of the literature on non-probabilistic decision theory outcome spaces are not explicitly specified. The departure of our paper is to demonstrate that the respective axiomatic foundations implicitly assume identical sets of possible payoffs in all states of nature. For many applications this assumption is not plausible. Outcome spaces differ naturally across states of nature when states influence outcomes (and their associated payoffs or utilities) in a systematic way. Consider, for instance, a firm making decisions when facing either a recession or a boom. In this example it is conceivable that any payoff in the recession lies below any payoff in the state of a boom. Here trading-off risks across states becomes particularly problematic since appropriate points of comparison are missing. The objective of this paper is to explicitly add outcome spaces to the description of a decision problem and to investigate how state dependence of outcome spaces affects the axiomatic foundations of various decision criteria.

To illustrate the role of outcome spaces, consider the following example. Here and later payoffs refer to outcomes evaluated in terms of von Neumann Morgenstern utilities.

	s_1	s_2
A_1	0	20
A_2	1	10

(1)

The maximin payoff criterion generates a preference ordering over acts that does not rely on the decision maker specifying probabilities of the occurrence of the two states s_1 and s_2. Here, we obtain $A_2 \succ A_1$ since the minimal payoff resulting from act A_2 is higher than the minimal payoff resulting from act A_1.

We now consider two possible specifications of the underlying outcome spaces. Assume first that the possible payoffs are only 0 and 1 in state s_1, and only 10, 20 and 25 in state s_2. In this environment, any payoff achievable in state s_2 is strictly larger than any payoff achievable in state s_1 (payoff ranges do not overlap). There are no points of comparison as there is no payoff that is achievable in all states. By consequence, two acts are compared under the maximin payoff criterion solely according to their outcomes in state s_1. Here the use of the maximin payoff criterion seems especially problematic since it neglects available information in a particularly strong way. It implicitly declares some states as *in principle* irrelevant to the decision problem.

Now consider an alternative environment where the set of possible payoffs in both states is equal to $\{0, 1, 10, 20, 25\}$. Clearly, the existence of the additional outcomes in state s_1 does not change the fact that act A_2 is strictly preferred to act A_1 under the maximin payoff criterion. However, it is no longer true for this environment that outcomes in state s_2 are irrelevant. For some pairs of acts the payoffs in both states will matter as we illustrate by adding two more acts to the above example:

	s_1	s_2
A_1	0	20
A_2	1	10
A_3	10	25
A_4	25	10

(2)

Under the maximin payoff criterion we have $A_3 \sim A_4$, a statement that can only be inferred by using information on the outcomes attained in both states. The outcomes in state s_2 are obviously not irrelevant for the maximin payoff criterion when there is overlap in the set of utilities achievable in each state.

When there is no overlap in payoff ranges across different states in decision problem (1), all axioms commonly used to characterize the maximin payoff criterion are consistent with ignoring the outcomes of each act in the "good" state s_2. In particular, this holds for the axioms proposed by Milnor (1954). But not only is ignoring the outcomes in the "good" state consistent with these axioms, ignoring the outcomes in the "bad" state is consistent as well. Specifically, suppose that the states of nature can be partitioned into S^b (the bad states) and S^g (the good states) such that all bad states have identical sets of possible payoffs, all good states have identical sets of possible payoffs, and *all* payoffs in any good state are strictly larger than *all* outcomes in bad states (say the possible payoffs are $\{0, 1\}$ in bad states and $\{10, 20, 25\}$ in good states, as above). In this situation, consider the following decision criterion. If there are only bad states, take the usual maximin payoff criterion. However, if there are good states then rank acts according to their minimal payoff achieved in the *good* states only. Clearly, this criterion differs from the usual maximin payoff criterion, for instance in the decision problem (1) it entails $A_1 \succ A_2$. But it is readily verified that this criterion satisfies all of Milnor's axioms.[1] The crucial observation is that the Symmetry Axiom is the only condition concerned with comparing outcomes between different states, and that it has no bite here when comparing the outcomes in good states with those in bad states. For instance, inferring $A_3 \sim A_4$ in (2) using the Symmetry Axiom presupposes that both payoffs 10 and 25 occur in either state.

In his characterization of the maximin payoff criterion, Milnor (1954) uses the Symmetry Axiom across two arbitrary states given any pair of payoffs and thus implicitly assumes outcome spaces to be identical across states of nature. The above example shows that the axioms no longer uniquely characterize the maximin payoff criterion if the sets of possible payoffs across states do not overlap. On the other hand, our main result below shows that Milnor's (1954) axioms together with Gilboa and Schmeidler's (1989) C-Independence Axiom continue to characterize the maximin payoff criterion provided that there is at least some overlap in the payoff ranges across states. In this case, the C-Independence Axiom allows us to transform all payoffs into the same range in order to apply Milnor's original proof, and in particular to invoke the Symmetry Axiom.[2]

Decision criteria that are based on the comparison of utility *differences* within states, such as the minimax regret criterion (or the maximin joy criterion, see Hayashi, 2006), do not seem to be sensitive to whether or not payoff ranges overlap as only differences matter. For instance, even when there is no overlap in decision problem (1), payoffs in state s_2 will influence preferences.[3] One might thus conjecture that the existence of points of comparison is less important. However, it turns out for the scenario described above, with the separation into good and bad states, that either

[1] These are *Ordering* (i.e. preferences between acts are transitive and complete), *Symmetry* (preferences do not depend on the labeling of acts and states), *Domination* (acts that yield strictly higher payoffs in each state are strictly preferred), *Continuity*, *Row Adjunction* (preferences between existing acts are not affected by adding new acts, commonly referred to as *Independence of Irrelevant Alternatives*), *Column Duplication* (preferences are unchanged if a new state is added provided that payoffs are identical to those in an already existing state) and *Convexity* (randomization between indifferent acts is preferred).

[2] Note also that state dependent outcome spaces limit the set of actions that can be added when invoking the Independence of Irrelevant Alternatives Axiom (Milnor's "Row Adjunction").

[3] We find $A_1 \succ A_2$ with utilities as specified in (1) while $A_1 \prec A_2$ holds if 10 is replaced by 19.5 in the payoff table.

good states or bad states can be ignored. Again the Symmetry Axiom has no bite.

The characterization of Milnor (1954) relies on an axiom called Column Linearity to create an act yielding the same payoff in each state. Such an act does not exist if there is no payoff range overlap. As Column Linearity has little intuitive appeal we follow Chernoff (1954) and replace it by the standard Independence Axiom, which is stronger than C-Independence. Otherwise we follow the framework of Milnor and obtain a characterization of minimax regret with state dependent outcome spaces provided there is sufficient overlap.

We also consider the maximin joy criterion, a dual criterion to minimax regret where payoffs are compared to the minimal payoff achievable in each state.[4] We present a characterization of this criterion that relies on the same type of overlap.

To summarize, we show that maximin payoff, minimax regret and maximin joy are still uniquely characterized when the payoff ranges in each state have a non empty intersection. However, when payoff ranges do not overlap the same axioms no longer have the power to *uniquely* characterize a decision criterion. Thus, we qualify the common belief that the axioms underlying the maximin payoff criterion necessarily lead to focus on a state with lowest payoffs. Moreover, we provide the more subtle insight that the other two criteria based on difference in outcomes within a state (minimax regret and maximin joy) nevertheless require comparability across states.

The paper is organized as follows. In the following Section 2 we lay down our basic notation. In Section 3 we list our axioms. The decision criteria are defined in Section 4, the main characterization result is in Section 5. Section 6 presents some more examples, and Section 7 discusses the relation to the literature.

2 Notation

Consider the following choice setting consisting of states, outcomes for each state and acts. By $S = \{s_1, .., s_m\}$ we denote the finite set of *states*. For each $j \in \{1, ..., m\}$, there is a finite set $\mathcal{X}_j$ of possible *outcomes* in state s_j. We refer to $\mathcal{X}_j$ as the *outcome space of state* s_j. $\mathcal{X} = \cup \mathcal{X}_j$ denotes the set of all possible outcomes. By $L(\mathcal{X}_j)$ we denote the set of all (finite) *lotteries* over outcomes in $\mathcal{X}_j$. A *(pure) act* A_i is a vector $A_i = (A_{i1}, .., A_{im})$ assigning a lottery over outcomes to each state, so $A_{ij} \in L(\mathcal{X}_j)$

[4] In (1) we find $A_1 \sim A_2$ as each act obtains the column minimum in some state.

for all $j \in \{1, ..., m\}$. The state-wise convex combination of two acts A_i and A_k is denoted by $\alpha A_i + (1 - \alpha) A_k$, i.e. $(\alpha A_i + (1 - \alpha) A_k)_j = \alpha A_{ij} + (1 - \alpha) A_{kj}$. A *mixed act* σ is a probability distribution over the set of acts. By σ_j we denote the lottery induced in state s_j by choosing the mixed act σ, thus in particular, $\sigma_j \in L(\mathcal{X}_j)$.

A finite collection of acts is called a *menu* and is denoted by A. A menu A can either be viewed as the vector $A = (A_1, ..., A_n)^T$ of n acts, or alternatively as the matrix $A = (A_{ij})_{i=1,...,n}^{j=1,...,m}$ of (lotteries of) outcomes. By $A^j = (A_{1j}, ..., A_{nj})^T$ we denote the (column) vector of the (lotteries of) outcomes attained by each act in state s_j.

Denote by $\mathbb{L}_{n,m}$ the set of all menus with n acts and m states. The element-wise convex combination of two menus $A, B \in \mathbb{L}_{n,m}$ is denoted by $\alpha A + (1 - \alpha) B$, i.e. $(\alpha A + (1 - \alpha) B)_{ij} = \alpha A_{ij} + (1 - \alpha) B_{ij}$.

By $(A, A^{m+1}) \in \mathbb{L}_{n,m+1}$ we denote the menu in which state s_{m+1} is added and in which the outcome of act i in state s_{m+1} given by A_i^{m+1}. Finally, $A \oplus A_{n+1} \in \mathbb{L}_{n+1,m}$ denotes the menu in which act A_{n+1} is added to menu $A \in \mathbb{L}_{n,m}$. Thus, e.g. $A = (A^1, .., A^m) = A_1 \oplus ... \oplus A_n$.

3 Axioms

Consider a preference relation $\succsim_s$ on $L(\mathcal{X})$. Throughout, we will assume that $\succsim_s$ is complete, transitive and satisfies continuity and (probabilistic) independence. These conditions imply the existence of a von-Neumann-Morgenstern utility function $u : \mathcal{X} \to \mathbb{R}$ such that for all $p, q \in L(\mathcal{X})$,

$$p \succsim_s q \Leftrightarrow Eu(p) \geq Eu(q),$$

where $Eu(p)$ is the expected utility induced by the lottery p.

In the following, we are interested in determining how to compare different acts within a given menu of acts and for given associated outcome spaces. Specifically, given menu A let $\succsim_A$ denote a preference relation over the mixed acts in the menu A. We will consider the following conditions on $\succsim_s$ and the family $\{\succsim_A\}_A$.

3.1 Outcome Spaces

The following condition only pertains to the structure of the outcome spaces.

1. *Payoff Range Overlap*: there exist $x, y \in \mathcal{X}$ such that for every j there exists $x_j, y_j \in \mathcal{X}_j$ such that $x_j \precsim_s x \prec_s y \precsim_s y_j$.

Payoff Range Overlap commonly arises in settings that include state independent outside options or the possibility of insurance across all states. Notice that we do not require that the possible outcomes of each state intersect but that there is intersection when evaluating outcomes in terms of preferences, hence the use of the term "payoff range" instead of "outcome space". An alternative way of stating our condition is that for any two states s_i and s_j, the most preferred outcome in state s_j is strictly preferred to the least preferred outcome in state s_i. In particular, it is important for our results that this preference is strict (see Section 6 below).

3.2 Intra-Menu Axioms

The following conditions refer to preferences among acts A_i *within* a given menu $A \in \mathbb{L}_{n,m}$.

2. *Ordering*: $\succsim_A$ is complete and transitive.

3. *Symmetry*: $\succsim_A$ is invariant to the labeling of acts and states as long as the relabeling of states does not violate the restrictions imposed by the respective outcome spaces. Specifically, label j of state s_j can be replaced with label k if for any i there exists $x_i \in \mathcal{X}_k$ such that $A_{ij} \sim_s x_i$.

4. *Domination*: If $A_{ij} \succ_s A_{kj}$ for all j then $A_i \succ_A A_k$.

5. *Convexity*: $A_i \sim_A A_k$ implies $\frac{1}{2}A_i + \frac{1}{2}A_k \succsim_A A_i$.

We add some comments. Symmetry will play a central role. The Symmetry Axiom of Milnor (1954) postulates that preferences are independent of the labeling of acts and states. Here we adjust the definition to allow for different states to have different outcome spaces.

The Convexity Axiom is sometimes referred to as "ambiguity aversion" in some intuitive sense; however, lacking a formal definition of that concept we avoid this term here.

3.3 Inter-Menu Axioms

The following conditions refer to consistency of preferences across different menus.

6. *Continuity*: If $A^{(t)} \to A$ as $t \to \infty$ and $A_i \succsim_{A^{(t)}} A_k$ for all t then $A_i \succsim_A A_k$.

7. *Column Duplication*: $A_i \succsim_A A_k$ if and only if $A_i \succsim_{(A,A^j)} A_k$ for all $j \leq m$ where $\mathcal{X}_{m+1} = \mathcal{X}_j$.

8. Independence to Menu Enlargement: We present three alternative conditions on a menu $A \in \mathbb{L}_{n,m}$ and an additional act A_{n+1} that can be postulated in order to be able to infer that, for all $i, k \leq n$,

$$A_i \succsim_A A_k \Leftrightarrow A_i \succsim_{A \oplus A_{n+1}} A_k, \tag{3}$$

i.e. to ensure that the ranking between acts is not affected by the availability of additional acts.

   (a) *Independence of Never-Uniquely-Best Alternatives (INUBA)*: (3) holds whenever for each j there is some i such that $A_{n+1,j} \precsim_s A_{i,j}$.

   (b) *Independence of Never-Uniquely-Worst Alternatives (INUWA)*: (3) holds whenever for each j there is some i such that $A_{n+1,j} \succsim_s A_{i,j}$.

   (c) *Independence of Irrelevant Alternatives (IIA)*: (3) always holds.

9. Independence to Inevitable Risk: Two alternative conditions on menus A and B will be considered in order to be able to infer that, for all $\alpha \in (0, 1)$,[5]

$$A_i \succsim_A A_k \Leftrightarrow \alpha A_i + (1-\alpha) B_i \succsim_{\alpha A + (1-\alpha) B} \alpha A_k + (1-\alpha) B_k. \tag{4}$$

   (a) *C-Independence*: (4) holds whenever $B_{ij} \sim_s B_{kl}$ for all i, j, k, l.

   (b) *Independence*: (4) holds whenever B_i is independent of i.

We add some comments.

Column Duplication implies that preferences should not change if states are described in greater detail but without changing any of the outcomes.

[5]It is sufficient to require (4) only for $\alpha = \frac{1}{2}$.

Many different terminologies exist for what we refer to as INUBA, none completely satisfying. As INUBA is consistent with adding an act that is a best alternative in each state, the alternative terminology "Independence of Never-Best Alternatives" would indicate a weaker requirement.

In (4), menu B can be interpreted following Chernoff (1954) as inevitable risk while A captures the part of the risk that can influenced by the choice itself. Under C-Independence ("certainty independence") this inevitable risk is assumed to be state independent while for Independence it measures the inevitable risk within each state. Alternatively one can interpret the independence axioms as time consistency. Let α be the probability of the event that the decision maker is asked (or allowed) to choose an act in menu A. The axiom specifies that preferences do not depend on whether or not choice is before or after this event. B describes what happens when A is not faced. Intuitively, under Independence some state within S occurs, so rows in B are assumed to be identical. Under C-Independence a single state $s_{m+1} \notin S$, so columns in B are assumed to be identical, too.

4 Decision Criteria

Formally, a *decision criterion* is a family of binary relations $\{\succsim_A\}_A$ for all possible menus A. We will consider the following specific decision criteria.

The *maximin payoff criterion* (Wald, 1950) is given by

$$\sigma \succsim_A \sigma' \Leftrightarrow \min_j Eu\left(\sigma_j\right) \geq \min_j Eu\left(\sigma'_j\right),$$

for all menus A and all mixed acts σ, σ'. The *minimax regret criterion* (Savage, 1951) is given by

$$\sigma \succsim_A \sigma' \Leftrightarrow \max_j \left\{\max_l \left\{Eu\left(A_{lj}\right)\right\} - Eu\left(\sigma_j\right)\right\} \leq \max_j \left\{\max_l \left\{Eu\left(A_{lj}\right)\right\} - Eu\left(\sigma'_j\right)\right\},$$

and similarly, the less studied *maximin joy criterion* (Hayashi, 2006) is given by

$$\sigma \succsim_A \sigma' \Leftrightarrow \min_j \left\{Eu\left(\sigma_j\right) - \min_l \left\{Eu\left(A_{lj}\right)\right\}\right\} \geq \min_j \left\{Eu\left(\sigma'_j\right) - \min_l \left\{Eu\left(A_{lj}\right)\right\}\right\},$$

for all menus A and all mixed acts σ, σ'.

5 Characterizations

We show how Milnor's (1954) axiomatizations of maximin payoff and minimax regret can be extended to allow for state dependent outcomes spaces provided there is Payoff Range Overlap. We make the following adjustments. For the maximin payoff criterion we add C-Independence. For the minimax regret criterion we follow Chernoff (1954) and replace the Column Linearity axiom of Milnor (1954) by the Independence Axiom. In addition we provide an axiomatization of the maximin joy criterion (see Hayashi (2006)) for this setting that emerges from the minimax regret criterion by shifting focus in each state from the best to the worst outcome.

Proposition 1 *Assume Payoff Range Overlap, and consider the class of all decision criteria satisfying Ordering, Symmetry, Domination, Convexity, Continuity and Column Duplication.*

1. *The maximin payoff criterion is the only criterion satisfying Independence of Irrelevant Alternatives and C-Independence.*

2. *The minimax regret criterion is the only criterion satisfying Independence of Never-Uniquely-Best Alternatives and Independence.*

3. *The maximin joy criterion is the only criterion satisfying Independence of Never-Uniquely-Worst Alternatives and Independence.*

Proof. Let x and y be outcomes that satisfy the axiom of Payoff Range Overlap. Following continuity there exists $D \in L(\mathcal{X})$ such that $x \prec_s D \prec_s y$ and for each j there exists $D_j \in L(\mathcal{X}_j)$ such that and $D_j \sim_s D$. Consider a menu A.

Part 1. Let B be the menu with $B_{ij} = D_j$ for all j. By C-Independence, we may investigate without loss of generality preferences among acts in $\alpha A + (1 - \alpha) B$ instead of those in A. By continuity, we may consider $\alpha > 0$ small enough so that $x \prec_s \alpha A_{ij} + (1 - \alpha) D_j \prec_s y$ holds for all i, j.

Thus, after compressing lotteries, all elements of the menu under consideration are better than x and worse than y. Now replace outcomes and lotteries over outcomes by utilities, normalizing $u(x) = 0$ and $u(y) = 1$. This is possible since we have assumed an expected utility representation of $\precsim_s$ on $L(\mathcal{X})$. Thus we end up with a menu that only contains elements in $[0, 1]$ and we can act as if there is common outcome space as $[0,1] \subseteq \{u(p) : p \in L(\mathcal{X}_j)\}$. From here on we continue as in the proof of Milnor (1954).

Specifically, the steps in Milnor's proof are: (i) Add rows and use Symmetry and Independence of Irrelevant Alternatives to show that only the set of elements in a row, not their order, matters for preferences. (ii) Then use Domination, Column Duplication and Symmetry to show that only the maximal and minimal elements in a row matter. (iii) Finally, use Convexity to show that it is in fact only the minimal element in a row that matters.

Part 2. As in Part 1 first compress the lotteries and then replace lotteries by utilities so that we only consider menus with elements in $[0,1]$.

The next step is to use the trick of Chernoff (1954) to transform the menu into an isomorphic one that has the same maximal element in each column. Let $z_j = \max_i A_{ij}$. Then there exist $\alpha \in \left[\frac{1}{2}, 1\right]$ and $b_j \in [0,1]$ for each j such that $\min_j z_j + \frac{1-\alpha}{\alpha} = 1$ and $z_j + \frac{1-\alpha}{\alpha} b_j = 1$. Let B be such that $B_{ij} = b_j$ and consider $\alpha A + (1-\alpha) B$. Then $\max_i (\alpha A + (1-\alpha) B)_{ij} = \alpha z_j + (1-\alpha) b_j = \alpha$ for all j. Thus the maximal element of each column of $\alpha A + (1-\alpha) B$ is equal to α. Independence implies that $\succsim_A$ and $\succsim_{\alpha A + (1-\alpha) B}$ are isomorphic.

The remaining proof is identical to that of Milnor (1954) who used an alternative axiom called Column Linearity to perform the above step.

Part 3. The proof is analogous to that of Part 2. Let $w_j = \min_i A_{ij}$ then there exist $\alpha \in \left[\frac{1}{2}, 1\right]$ and $b'_j \in [0,1]$ for each j such that $\min_j w_j + \frac{1-\alpha}{\alpha} = \max_j w_j$ and $w_j + \frac{1-\alpha}{\alpha} b'_j = \max_j w_j$. Given $B'_{ij} = b'_j$ we obtain $\min_i (\alpha A + (1-\alpha) B')_{ij} = \alpha w_j + (1-\alpha) b_j = \alpha \max_j w_j$. Consequently, each column of $(\alpha A + (1-\alpha) B')$ has the same minimal element. The next step is to add a constant row with elements $\alpha \max_j w_j$ and then to continue as the proof of Milnor (1954) for the maximin payoff criterion. This is because the proof of Milnor (1954) for the maximin payoff criterion, when applied to matrices with a constant row consisting of the minimal element of the matrix, only utilizes Independence of Never-Uniquely-Worst Alternatives and does not build on the stronger Independence of Irrelevant Alternatives axiom. ■

6 Further Examples

Payoff Range Overlap requires that intersection of the ranges has a nonempty interior. We illustrate why this is necessary for obtaining our results in the following example:

	s_1	s_2
A_1	0	20
A_2	1	1

Assume that no other payoffs are possible in either state. Then both $A_1 \succ_A A_2$ and $A_2 \succ_A A_1$ is consistent will all of our axioms, the reason being that the Symmetry Axiom is vacuous.

Finally we present an economic application where payoff ranges naturally are different but where there is sufficient overlap. Consider a firm who is competing in prices a la Bertrand with a second firm. Instead of adapting an equilibrium approach we assume that the price charged by the other firm is unknown. In the framework of this paper, each possible price charged by the other firm is associated to a state. To obtain a finite state space, assume that prices are set within a grid $\{0, 1, 2, 3, 4\}$. Assume that the firm is risk neutral, has no fixed cost and marginal cost is equal to 1 and that there is a single consumer with unit demand and a willingness-to-pay equal to 4. Then we obtain the following matrix indicating for each pair of prices the profit of the firm charging the price specified in the row:

	0	1	2	3	4
0	$-1/2$	-1	-1	-1	-1
1	0	0	0	0	0
2	0	0	$1/2$	1	1
3	0	0	0	1	2
4	0	0	0	0	$3/2$

Payoff Range Overlap holds as e.g. -0.1 is in the interior of the utility range of each state. In the following we select among pure acts. Prices in $\{1, 2, 3, 4\}$ are most preferred both under maximin payoff criterion and under the maximin joy criterion while only price 3 is most preferred under the minimax regret criterion.

7 Related Literature and Conclusion

We pay a tribute to Milnor (1954) due to its clear presentation with proofs in which the role of each axiom is clearly highlighted in each of the steps.

In accordance to Milnor (1954) we find it natural to separate uncertainty within and across states belonging to S. Uncertainty within a state is considered objective with common preferences $\succsim_s$ over the set of all possible outcomes across all states admitting an expected utility representation. Thus we focus on how to deal with "complete" uncertainty across states.

All axioms apart from C-Independence and Independence are taken directly from Milnor (1954). As outcomes are the primitive of our model we cannot use Column Linearity of Milnor (1954) and instead follow Chernoff (1954) (see also Luce and Raiffa, 1957, and Stoye, 2007) by requiring Independence. Terminology is changed whenever we strongly believe that original terms are inappropriate. Proofs follow those of Milnor (1954) as close as possible.

Gilboa and Schmeidler (1989) introduced the C-Independence Axiom as a key condition in their characterization of maximin expected utility with respect to some set of priors. The Independence of Irrelevant Alternatives condition is implicit in their framework.

In a recent paper, Stoye (2007) expands on the set of axioms of Gilboa and Schmeidler (1989) to develop a unifying framework for comparing a number of different criteria including Bayes, maximin payoff and minimax regret, providing novel characterizations for the latter two. Compared to Milnor (1954) several axioms are slightly changed in the work of Stoye (2007). The Continuity Axiom is replaced by the stronger version of Mixture Continuity, and Domination is replaced by "Weak Domination". While Column Duplication is missing, some of its features are included in Stoye's symmetry condition which allows to collect states with the same outcomes. C-Independence is used in the characterization of Stoye (2007). Remember that this axiom does not enter the axioms of Milnor (1954), and we only use this axiom to be able to deal with state dependent outcome spaces. Our insights on the connection between Payoff Range Overlap and Symmetry also extend to the framework of Stoye (2007): (i) neither maximin payoff nor minimax regret are uniquely characterized when some outcome space is disjoint from all others, and (ii) the characterizations of maximin payoff and minimax regret extend to state dependent outcome spaces when there is Payoff Range Overlap.

The minimax regret criterion has recently been considered by Hayashi (2006) in a different light. Hayashi (2006) axiomatizes a choice rule that minimizes the expected regret with respect to a not necessarily unique prior. If no prior is excluded, this choice rule reduces to the classical minimax regret criterion; on the other hand, if the prior is unique, the choice rule reduces to the (subjective) expected utility criterion. Since the framework of Hayashi (2006) is quite different from ours, the relevant conditions are not directly comparable. Interestingly however, the Independence of Never-Uniquely-Best Action Axiom (i.e. Milnor's (1954) "Special Row Adjunction") appears also in Hayashi's characterization (as "Irrelevance of Dominated Acts").

The idea to consider state dependent outcome spaces and to investigate the implications seems to be novel.

References

[1] Chernoff, H. (1954), "Rational Selection of Decision Functions," *Econometrica* **22** (4), 422–443.

[2] Gilboa, I. and D. Schmeidler (1989), "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics* **18**, 141–153.

[3] Hayashi, T. (2006), Regret Aversion and Opportunity-Dependence, *Mimeo*, University of Texas, Austin.

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

[5] Milnor, J. (1954), "Games Against Nature," in: *Decision Processes*, ed. R.M. Thrall, C.H. Coombs & R.L. Davis. New York, Wiley.

[6] Savage, L. (1951), "The Theory of Statistical Decision," *Journal of the American Statistical Association* **46**, 55–67.

[7] Stoye, J. (2007), Statistical Decisions under Ambiguity, *Mimeo*, New York University.

[8] Wald, A. (1950), *Statistical Decision Functions*, New York, John Wiley.