A decision rule for imprecise probabilities based on pair-wise comparison of expectation bounds

Sébastien Destercke

Abstract There are many ways to extend the classical expected utility decision rule to the case where uncertainty is described by (convex) probability sets. In this paper, we propose a simple new decision rule based on the pair-wise comparison of lower and upper expected bounds. We compare this rule to other rules proposed in the literature, showing that this new rule is both precise, computationally tractable and can help to boost the computation of other, more computationally demanding rules.

Key words: Maximality, Hurcwitz criterion, E-admissibility, lower previsions, Γ-maximin

1 Introduction

We are concerned here with the problem of making a decision d, which may be taken from a set of N available decisions $\mathscr{D} = \{d_1, \ldots, d_N\}$. Usually, this decision is not chosen arbitrarily, i.e., it should be the best possible in the current situation.

In our case, the benefits that an agent would gain by taking decision d_i depend on a variable X and the knowledge we have about its value. We assume here that the true value of X is uncertain, that it takes its value on a finite domain $\mathscr{X}$ and that the benefit (or gain, reward) of choosing d_i can be modelled by a real-valued and bounded utility function $U_i : \mathscr{X} \to \mathbb{R}$, with $U_i(x)$ the gain of choosing action d_i when x is the value of X. The problem of decision making is then to select, based on this information, the decisions that are optimal, i.e. are likely to gives the best possible gain.

When uncertainty on X is (or can be) modelled by a probability distribution $p : \mathscr{X} \to [0, 1]$, many authors (for example De Finetti [2]) have argued that the

Sébastien Destercke
UMR IATE, Campus Supagro, 2 Place P. Viala, 34060 Montpellier, France e-mail: sebastien.destercke@irsn.fr

optimal decision $\overline{d} \in \mathscr{D}$ should be the one maximising the expected utility, i.e., $\overline{d}_{\mathbb{E}_p} = \arg\max_{d_i \in \mathscr{D}} \mathbb{E}_p(U_i) = \sum_{x \in X} U_i(x)p(x)$. Thus, selecting the optimal decision in the sense of expected utility comes down to considering the complete (pre-)order induced by expected utility, here denoted by $\leq_{\mathbb{E}}$, over decisions in $\mathscr{D}$ ($d_i \leq_{\mathbb{E}} d_j$ if $\mathbb{E}_p(U_i) \leq \mathbb{E}_p(U_j)$), and to choose the decision which is not dominated by others (Given a partial order $\leq$ on $\mathscr{D}$, we say that d dominates d' if $d' \leq d$). In the sequel, we will say that a decision d is optimal w.r.t. an order $\leq$, or a decision rule, if it is non-dominated in the order induced by this decision rule.

However, it may happen that our uncertainty about the value of X cannot be modelled by a single probability, for the reason that not enough information is available to identify the probability $p(x)$ of every element $x \in \mathscr{X}$. In such a case, convex sets of probabilities, here called credal sets [5] (which are formally equivalent to coherent lower previsions [9]), have been proposed as an uncertainty representation allowing us to model information states going from full ignorance to precise probabilities, thus coping with insufficiencies in our information. Formally, they encompass most of the uncertainty representations that integrate the notion of imprecision (e.g., belief functions, possibility distributions, ...).

To select optimal decisions in this context, it is necessary to extend the expected utility criterion, as the expected utility $\mathbb{E}(U)$ is no longer precise and becomes a bounded interval $[\underline{\mathbb{E}}(U), \overline{\mathbb{E}}(U)]$. In the past decades, several such extensions, based on the evaluations of expectation bounds rather than of precise expected values, have been proposed (see Troffaes [6] for a concise and recent review). Roughly speaking, two kinds of generalisations are possible: either using a combination of the lower and upper expectation bounds to induce a complete (pre-)order between decisions, reaching a unique optimal decision, or relaxing the need of a complete order and extending expected utility criterion to obtain a partial (pre-)order between decisions. In this latter cases, there may be several optimal decisions, the inability to select between them reflecting the imprecision in our information.

In this paper, we propose and explore a new decision rule of the latter kind, based on a pair-wise comparison of lower and upper expectation bounds. This rule, which has not been studied before in the framework of imprecise probabilities (to our knowledge), is quite simple and computationally tractable. Section 2 recalls the imprecise probabilistic framework as well as the existing decision rules. We then present in Section 3 the new rule and compare it to existing rules. We will show that this rule is (surprisingly) precise when compared to other rules inducing partial pre-orders between decisions.

2 Imprecise probabilities and decision rules

We consider that our information and uncertainty regarding the value of a variable X is modelled by a credal set $\mathscr{P}$. Given a function $U_i : \mathscr{X} \to \mathbb{R}$ over the space $\mathscr{X}$, the lower and upper expectations $\underline{\mathbb{E}}_{\mathscr{P}}(U_i), \overline{\mathbb{E}}_{\mathscr{P}}(U_i)$ of U_i are such that

$$\underline{\mathbb{E}}_{\mathscr{P}}(U_i) = \inf_{p \in \mathscr{P}} \mathbb{E}_p(U_i) \qquad \overline{\mathbb{E}}_{\mathscr{P}}(U_i) = \sup_{p \in \mathscr{P}} \mathbb{E}_p(U_i)$$

In Walley's [9] behavioural interpretation of imprecise probabilities, $\underline{\mathbb{E}}_{\mathscr{P}}(U_i)$ is interpreted as the maximum buying price an agent would be ready to pay for U_i, associated to decision d_i. Conversely, $\overline{\mathbb{E}}_{\mathscr{P}}(U_i)$ is interpreted as the minimum selling price an agent would be ready to receive for U_i. These two expectation bounds are dual, in the sense that, for any real-valued bounded function f over $\mathscr{X}$, we have $\underline{\mathbb{E}}(f) = -\overline{\mathbb{E}}(-f)$.

When proposing a decision rule based on lower and upper expectations $\underline{\mathbb{E}}, \overline{\mathbb{E}}$, a basic requirement is that this decision rule should reduce to the classical expected utility rule when $\mathscr{P}$ reduces to a single probability distribution. Still, there are many ways to do so, providing $\mathscr{D}$ with a complete or a partial (pre-)order. In the former case, there is a unique optimal non-dominated decision, while in the latter there may be a set of such non-dominated decisions. We will review the most commonly used approaches, dividing them according to the kind of order they induce on $\mathscr{D}$.

Example 1. In order to illustrate our purpose, let us consider the same example as Troffaes [6]. Consider a coin that can either fall on head (H) or tails (T), thus $\mathscr{X} = \{H, T\}$, with our uncertainty given as $p(H) \in [0.28; 0.7]$ and $p(T) \in [0.3; 0.72]$. Different decisions and their pay-off in case of landing on Heads or Tails are summarized in Table 1, together with the lower and upper expectations reached by each decision.

$\mathscr{D}$		U_i	H	T	$\underline{\mathbb{E}}$	$\overline{\mathbb{E}}$
d_1		U_1	4	0	1.12	2.8
d_2		U_2	0	4	1.2	2.88
d_3	$\rightarrow$	U_3	3	2	2.28	2.7
d_4		U_4	$1/2$	3	1.25	2.3
d_5		U_5	$47/20$	$47/20$	2.35	2.35
d_6		U_6	$41/10$	$-3/10$	0.932	2.78

Table 1 Example 1 possible decisions and expectation bounds.

2.1 Rules inducing a complete order

Let us start with the rules pointing to a unique optimal decision.

***Γ*-maximin** The Γ-maximin rule [3], denoted by $\leq_{\underline{\mathbb{E}}}$, consists in replacing the expected value with the lower expectation. The optimal decision under this rule is such that

$$\overline{d}_{\leq_{\underline{\mathbb{E}}}} = \arg \max_{d_i \in \mathscr{D}} \underline{\mathbb{E}}_{\mathscr{P}}(U_i).$$

This rule correspond to a pessimistic attitude, since it consists in maximizing the worst possible expected gain. In example 1, $\overline{d}_{\leq_{\underline{\mathbb{E}}}} = d_5$.

***Γ*-maximax** The optimistic version of the Γ-maximin, denoted by $\leq_{\overline{\mathbb{E}}}$ and consisting in selecting as optimal the decision that maximises the expected outcome is such that

$$\overline{d}_{\leq_{\overline{\mathbb{E}}}} = \arg \max_{d_i \in \mathscr{D}} \overline{\mathbb{E}}_{\mathscr{P}}(U_i).$$

In example 1, $\overline{d}_{\leq_{\overline{\mathbb{E}}}} = d_2$.

Hurcwitz Criterion Hurcwitz criterion in imprecise probabilities [4], denoted here by $\leq_\alpha$, consists in choosing a so-called pessimism index $\alpha \in [0,1]$, and to induce an order where the behaviour of the decision maker range from fully pessimistic ($\alpha = 1$) to fully optimistic ($\alpha = 0$). Once a pessimistic index α has been chosen, Hurwictz rule is such that $d_i \leq_\alpha d_j$ whenever $\alpha\underline{\mathbb{E}}_{\mathscr{P}}(U_i) + (1-\alpha)\overline{\mathbb{E}}_{\mathscr{P}}(U_i) \leq \alpha\underline{\mathbb{E}}_{\mathscr{P}}(U_j) + (1-\alpha)\overline{\mathbb{E}}_{\mathscr{P}}(U_j)$, and the optimal decision $\overline{d}_{\leq_\alpha}$ under this rule is

$$\overline{d}_{\leq_\alpha} = \arg \max_{d_i \in \mathscr{D}} \alpha\underline{\mathbb{E}}_{\mathscr{P}}(U_i) + (1-\alpha)\overline{\mathbb{E}}_{\mathscr{P}}(U_i).$$

Γ-maximin and -maximax are respectively retrieved when $\alpha = 1$ and $\alpha = 0$. In Example 1, the set of optimal decisions $\overline{d}_{\leq_\alpha}$ that can be reached by different values of α is $\{d_2, d_3, d_5\}$

Note that determining optimal decisions for these three criteria requires N comparisons and at most $2N$ computations of expectation bounds.

2.2 Rules inducing a partial order

The other alternative when extending expected utility criterion is to let drop off the assumption that the order on the decisions has to be complete. That is, to allow the order to be partial and to possibly induce a set of optimal decisions rather than a single one. Three rules following this way have been proposed up to now.

Interval dominance A first natural extension to the comparison of precise expectations to the case of interval-valued expectations is the interval dominance order $\leq_{\mathscr{I}}$ such that $d_i \leq_{\mathscr{I}} d_j$ if and only if $\overline{\mathbb{E}}_{\mathscr{P}}(U_i) \leq \underline{\mathbb{E}}_{\mathscr{P}}(U_j)$. That is, d_j dominates d_i if the expected gain of d_j is at least as great as the one of d_i. The resulting set of non-dominated (or optimal) decisions is denoted by $\overline{\mathscr{D}}_{\mathscr{I}}$ and is such that

$$\overline{\mathscr{D}}_{\mathscr{I}} = \{d \in \mathscr{D} |\ \nexists d' \in \mathscr{D}, d \leq_{\mathscr{I}} d'\}.$$

Computing $\overline{\mathscr{D}}_{\mathscr{I}}$ requires the computation of $2N$ expectations and $2N$ comparisons. For Example 1, we have $\overline{\mathscr{D}}_{\mathscr{I}} = \{d_1, d_2, d_3, d_5, d_6\}$. As we can see, this rule has the advantage to be computationally efficient, but is also very imprecise.

Maximality When expectations are precise, we have $d_i \geq_{\mathbb{E}} d_j$ whenever $\mathbb{E}_p(U_i) \geq \mathbb{E}_p(U_j)$ or, equivalently, whenever $\mathbb{E}_p(U_i - U_j) \geq 0$. The notion of maximality consists in extending this notion by inducing a pre-order $\geq_{\mathscr{M}}$ such that $d_i >_{\mathscr{M}} d_j$ whenever $\underline{\mathbb{E}}_{\mathscr{P}}(U_i - U_j) > 0$. In Walley's interpretation, $\underline{\mathbb{E}}_{\mathscr{P}}(U_i - U_j) > 0$ means that we are ready to pay a positive price to exchange U_i for U_j, hence that decision d_i is preferred to decision d_j. The resulting set of optimal decisions $\overline{\mathscr{D}}_{\mathscr{M}}$ is such that

$$\overline{\mathscr{D}}_{\mathscr{M}} = \{d \in \mathscr{D} | \nexists d' \in \mathscr{D}, d \leq_{\mathscr{M}} d'\}$$

Computing $\overline{\mathscr{D}}_{\mathscr{M}}$ requires the computation of $N^2 - N$ lower expectations and $N^2 - N$ comparisons. For Example 1, we have $\overline{\mathscr{D}}_{\mathscr{M}} = \{d_1, d_2, d_3, d_5\}$.

E-admissibility Robustifying the expected utility criterion when uncertainty is modelled by sets of probabilities can simply be done by selecting as optimal those decisions that are optimal w.r.t. classical expected utility for at least one probability measure of $\mathscr{P}$. In this case, the set of optimal decision $\overline{\mathscr{D}}_{\mathscr{E}}$ is such that

$$\overline{\mathscr{D}}_{\mathscr{E}} = \{d \in \mathscr{D} | \exists p \in \mathscr{P} \text{ s.t. } \overline{d}_{\mathbb{E}_p} = d\}$$

Utkin and Augustin [7] have proposed algorithms that allow computing $\overline{\mathscr{D}}_{\mathscr{E}}$ by solving N linear programs whose complexity is slightly higher than the ones usually associated to the computation of a lower expectation. For Example 1, we have $\overline{\mathscr{D}}_{\mathscr{E}} = \{d_1, d_2, d_3\}$. Both E-admissibility and Maximality give more precise statements than Interval dominance, but their computational burden is also higher (hence, they are more difficult to use in complex problems).

3 The new decision rule

The rules presented in the previous section consist, for most of them, in comparing numeric values (expectation bounds) to determine which decisions are dominated by others and are therefore non-optimal. Other ways to order interval-valued numbers can therefore be considered and studied as potential decision rules. One such ordering that has not be studied in imprecise probability theory (as far as we know) is the one where an interval $[a,b]$ is considered as lower than $[c,d]$ if $a \leq c$ and $b \leq d$. This comes down to a pair-wise comparison of the interval bounds.

Using this ordering, we therefore propose a new decision rule, that we call *Interval bound dominance* ($\mathscr{I}\mathscr{B}$-dominance for short), denoted by $\leq_{\mathscr{I}\mathscr{B}}$, and defined as follows

Definition 1 (Interval bound dominance). Given a credal set $\mathscr{P}$ and two decisions $d_i, d_j \in \mathscr{D}$, $d_i \leq_{\mathscr{I}\mathscr{B}} d_j$ whenever $\underline{\mathbb{E}}_{\mathscr{P}}(U_i) \leq \underline{\mathbb{E}}_{\mathscr{P}}(U_i)$ and $\overline{\mathbb{E}}_{\mathscr{P}}(U_i) \leq \overline{\mathbb{E}}_{\mathscr{P}}(U_i)$ ($d_i <_{\mathscr{I}\mathscr{B}} d_j$ when at least one of the two inequalities is strict).

Note that, as for the rules of Section 2.2, the order $\leq_{\mathscr{I}\mathscr{B}}$ is partial and induces a set of optimal decisions. The set of optimal decisions $\mathscr{D}_{\mathscr{I}\mathscr{B}}$ resulting from this decision rule is such that

$$\overline{\mathscr{D}}_{\mathscr{I}\mathscr{B}} = \{d \in \mathscr{D} | \nexists d' \in \mathscr{D}, d \leq_{\mathscr{I}\mathscr{B}} d'\}$$

In Example 1, we have $\mathscr{D}_{\mathscr{I}\mathscr{B}} = \{d_2, d_3, d_5\}$, which is different from any set obtained with other decision rules of Section 2.2.

Computing the set $\overline{\mathscr{D}}_{\mathscr{I}\mathscr{B}}$ requires the computation of $2N$ expectation bounds (the same as for computing $\overline{\mathscr{D}}_{\mathscr{I}}$) and $2N$ comparisons at most. It is therefore as computationally efficient as the interval dominance criterion, and can be more precise (see Example 1). Actually, we will show that it is always at least as precise.

Let us now study the relation of this new decision rule with previous ones. First, we will show that the $\mathscr{I}\mathscr{B}$ decision rule is coherent with the rules inducing a complete order between decisions, before processing to the rules inducing a partial order.

3.1 Relations with complete ordering rules

Let us first start with Γ-maximin and Γ-maximax. As indicates the next proposition, we can easily show that the $\mathscr{I}\mathscr{B}$ decision rule considers as optimal the decisions selected by these two rules.

Proposition 1. *The two optimal decisions $\overline{d}_{\leq_{\underline{\mathbb{E}}}}$ and $\overline{d}_{\leq_{\overline{\mathbb{E}}}}$ in the sense of Γ-maximin and Γ-maximax are also optimal in the sense of $\mathscr{I}\mathscr{B}$ dominance, that is*

$$\{\overline{d}_{\leq_{\underline{\mathbb{E}}}}, \overline{d}_{\leq_{\overline{\mathbb{E}}}}\} \subseteq \overline{\mathscr{D}}_{\mathscr{I}\mathscr{B}}$$

Proof. We will only prove $\overline{d}_{\leq_{\underline{\mathbb{E}}}} \in \overline{\mathscr{D}}_{\mathscr{I}\mathscr{B}}$, proof for $\overline{d}_{\leq_{\overline{\mathbb{E}}}}$ being similar. Let $\overline{d}_{\leq_{\underline{\mathbb{E}}}} = d_i$, as by definition there are no decision $d_j \in \mathscr{D}$ such that $\underline{\mathbb{E}}(U_i) < \underline{\mathbb{E}}(U_j)$, this means that there are no decision that $\mathscr{I}\mathscr{B}$-dominates d_i, hence $\overline{d}_{\leq_{\underline{\mathbb{E}}}} \in \overline{\mathscr{D}}_{\mathscr{I}\mathscr{B}}$.

The next proposition shows that $\mathscr{I}\mathscr{B}$ decision rule can also be seen as a robustification of Hurwictz criterion.

Proposition 2. *Let d_i, d_j be two different decisions. Then, $d_i \leq_{\mathscr{I}\mathscr{B}} d_j$ if and only if $d_i \leq_\alpha d_j$ for every $\alpha \in [0,1]$*

Proof. Let us first prove the "if" part. Since $d_i \leq_\alpha d_j$ for every α, if we consider $\alpha = 1$ and $\alpha = 0$ we respectively have that $\underline{\mathbb{E}}_{\mathscr{P}}(U_i) \leq_1 \underline{\mathbb{E}}_{\mathscr{P}}(U_j)$ and $\overline{\mathbb{E}}_{\mathscr{P}}(U_i) \leq_0 \overline{\mathbb{E}}_{\mathscr{P}}(U_j)$. These two inequalities leading to $d_i \leq_{\mathscr{I}\mathscr{B}} d_j$.

Let us now concentrate on the "only if" part. $d_i \leq_{\mathscr{I}\mathscr{B}} d_j$ means that $\underline{\mathbb{E}}_{\mathscr{P}}(U_i) \leq \underline{\mathbb{E}}_{\mathscr{P}}(U_j)$ and $\overline{\mathbb{E}}_{\mathscr{P}}(U_i) \leq \overline{\mathbb{E}}_{\mathscr{P}}(U_j)$ (these two inequalities covering the case where $\alpha = 0$ and $\alpha = 1$). Hence, for any value $\alpha \in (0,1)$, we also have $\alpha\underline{\mathbb{E}}_{\mathscr{P}}(U_i) \leq \alpha\underline{\mathbb{E}}_{\mathscr{P}}(U_j)$ and $(1-\alpha)\overline{\mathbb{E}}_{\mathscr{P}}(U_i) \leq (1-\alpha)\overline{\mathbb{E}}_{\mathscr{P}}(U_j)$. Summing left and right-hand sides of each equations, we have $\alpha\underline{\mathbb{E}}_{\mathscr{P}}(U_i) + (1-\alpha)\overline{\mathbb{E}}_{\mathscr{P}}(U_i) \leq \alpha\underline{\mathbb{E}}_{\mathscr{P}}(U_j) + (1-\alpha)\overline{\mathbb{E}}_{\mathscr{P}}(U_j)$, hence $d_i \leq_{\mathscr{I}\mathscr{B}} d_j$ implies $d_i \leq_\alpha d_j$ for any α

The $\mathscr{I}\mathscr{B}$ decision rule can thus be seen as a decision rule where a decision dominates another if and only if it dominates it under all different pessimistic/optimistic attitudes, thus safeguarding the decision maker against the need to commit into such an attitude in a first analysis. Actually, it looks possible that $\overline{\mathscr{D}}_{\mathscr{I}\mathscr{B}}$ contains all actions that are optimal in the Hurcwitz sense for some value of α, as is the case in the example. Let us now study the relations with the rules inducing a partial ordering.

3.2 Relations with partial ordering rules

The next proposition indicates that Interval dominance implies $\mathscr{I}\mathscr{B}$ dominance.

Proposition 3. *Given a decision set $\mathscr{D}$ and a credal set $\mathscr{P}$, we have $\overline{\mathscr{D}}_{\mathscr{I}\mathscr{B}} \subseteq \overline{\mathscr{D}}_{\mathscr{B}}$, with the inclusion being usually strict.*

Proof. We need to show that if a decision d_i is not optimal w.r.t. $\leq_{\mathscr{I}}$, then it is also not optimal w.r.t. $\leq_{\mathscr{I}\mathscr{B}}$. If d_i is not optimal w.r.t. $\leq_{\mathscr{I}}$, it means that there is a decision d_j such that $d_i <_{\mathscr{I}} d_j$, hence that $\overline{\mathbb{E}}_{\mathscr{P}}(U_i) < \underline{\mathbb{E}}_{\mathscr{P}}(U_j)$. Since $\underline{\mathbb{E}}_{\mathscr{P}}(U_i) \leq \overline{\mathbb{E}}_{\mathscr{P}}(U_i) < \underline{\mathbb{E}}_{\mathscr{P}}(U_j) \leq \overline{\mathbb{E}}_{\mathscr{P}}(U_j)$, this implies $d_i <_{\mathscr{I}\mathscr{B}} d_j$

The next result concerns the relation of $\mathscr{I}\mathscr{B}$ decision rule with maximality.

Proposition 4. *Given a decision set $\mathscr{D}$ and a credal set $\mathscr{P}$, we have $\overline{\mathscr{D}}_{\mathscr{I}\mathscr{B}} \subseteq \overline{\mathscr{D}}_{\mathscr{M}}$, with the inclusion being usually strict.*

Proof. Let us show that if a decision d_i is not optimal w.r.t. $\leq_{\mathscr{M}}$, then it will also be non-optimal w.r.t. $\leq_{\mathscr{I}\mathscr{B}}$. If $d_i \notin \mathscr{D}_{\mathscr{M}}$, then it means $\exists d_j$ s.t. $\underline{\mathbb{E}}(U_j - U_i) > 0$. Using the properties of lower expectations (see Walley [9, Ch. 2]), we have $\underline{\mathbb{E}}(U_j) + \overline{\mathbb{E}}(-U_i) \geq \underline{\mathbb{E}}(U_j - U_i)$. Using this inequality and the duality between lower and upper expectations, we have $\underline{\mathbb{E}}(U_j) + \overline{\mathbb{E}}(-U_i) = \underline{\mathbb{E}}(U_j) - \underline{\mathbb{E}}(U_i) > 0$, hence $\underline{\mathbb{E}}(U_j) > \underline{\mathbb{E}}(U_i)$. Similarly, we have that $\overline{\mathbb{E}}(U_j) + \underline{\mathbb{E}}(-U_i) \geq \underline{\mathbb{E}}(U_j - U_i)$. using the same reasoning and duality, we have $\overline{\mathbb{E}}(U_j) - \overline{\mathbb{E}}(U_i) > 0$, meaning that $\overline{\mathbb{E}}(U_j) > \overline{\mathbb{E}}(U_i)$. Hence, $d_i <_{\mathscr{M}} d_j$ implies $d_i <_{\mathscr{I}\mathscr{B}} d_j$, and $d_i \notin \mathscr{D}_{\mathscr{I}\mathscr{B}}$

This proposition tells us, among other things, that $\mathscr{I}\mathscr{B}$-dominance can be used as a quick estimate of an inner approximation of the set $\overline{\mathscr{D}}_{\mathscr{M}}$, while interval dominance can be used to estimate an outer approximation of this set. This means that both interval dominance and $\mathscr{I}\mathscr{B}$-dominance, which present a low computational complexity when compared to maximility, can be used to reduce drastically the number of required computations to evaluate $\overline{\mathscr{D}}_{\mathscr{M}}$. In the example, only two decisions that are in $\overline{\mathscr{D}}_{\mathscr{I}}$ but not in $\mathscr{D}_{\mathscr{I}\mathscr{B}}$ would need to be verified: $\{d_1, d_6\}$.

Concerning E-admissibility and $\mathscr{I}\mathscr{B}$-dominance, it is easy to see, from the example, that none imply the other, since the set of optimal actions under these rules only overlap (and their union is the set $\overline{\mathscr{D}}_{\mathscr{M}}$). Figure 1 recalls [6] and summarises the implications relation between the different rules, integrating $\mathscr{I}\mathscr{B}$-dominance into it. Roughly speaking, the figure goes from the most precise decision rules (left) to the most imprecise (right).

4 Conclusion

In this paper, we have proposed a simple new decision rule for imprecise probabilities, based on expectation bound pair-wise comparison, and have studied its relation with other existing decision rules. The interest of this rule is that it remains

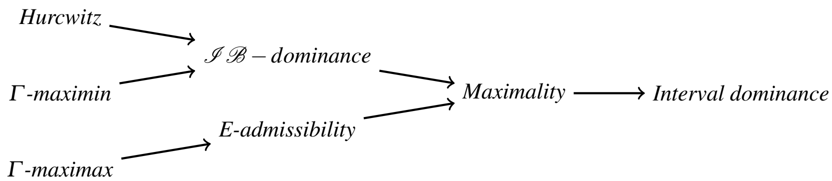

Fig. 1 Relations between decision rules: $A \to B$ means that a decision optimal in the sense of A is also optimal in the sense of B

in the spirit of an imprecise probabilistic approach, since less information will lead to a larger set of optimal decisions, but is both computationally tractable and less conservative than most other rules. Another interesting fact is that this rule implies maximality (i.e. $\mathscr{I}\mathscr{B}$ optimal decisions are also maximal). Therefore, if not used for itself, the $\mathscr{I}\mathscr{B}$ decision rule can boost the computational tractability of $\overline{\mathscr{D}}_{\mathscr{M}}$, using it in conjunction with interval dominance to reduce the number of decision whose optimality under maximality criterion must be checked.

The next step is to evaluate to which extent this decision rule can improve the results of some tasks such as classification [10], and if it is consistent with a dynamic programming approach when dynamics enters the picture [1].

References

1. G. de Cooman and M. C. M. Troffaes. Dynamic programming for deterministic discrete-systems with uncertain gain. *Int. J. of Approximate Reasoning*, 39:257–278, 2004.
2. B. Finetti. *Theory of probability*, volume 1-2. Wiley, NY, 1974. Translation of 1970 book.
3. I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989.
4. J.-Y. Jaffray and M. Jeleva. Information processing under imprecise risk with the Hurwicz criterion. In *Proc. of the fifth Int. Symp. on Imprecise Probabilities and Their Applications*, 2007.
5. I. Levi. *The Enterprise of Knowledge*. MIT Press, London, 1980.
6. M. Troffaes. Decision making under uncertainty using imprecise probabilities. *Int. J. of Approximate Reasoning*, 45:17–29, 2007.
7. L. Utkin and T. Augustin. Powerful algorithms for decision making under partial prior information and general ambiguity attitudes. In *Proc. of the fourth Int. Symp. on Imprecise Probabilities and Their Applications*, 2005.
8. J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.
9. P. Walley. *Statistical reasoning with imprecise Probabilities*. Chapman and Hall, New York, 1991.
10. M. Zaffalon. The naive credal classifier. *J. Probabilistic Planning and Inference*, 105:105–122, 2002.