

Qualitative Decision under Uncertainty: Back to Expected Utility

Hélène Fargier

IRIT, 118 Route de Narbonne
31062 Toulouse Cedex, France
fargier@irit.fr

Régis Sabbadin

INRA-UBIA,
BP 27 - 31326 Castanet-Tolosan Cedex - France
sabbadin@toulouse.inra.fr

Abstract

Different qualitative models have been proposed for decision under uncertainty in Artificial Intelligence, but they generally fail to satisfy the principle of strict Pareto dominance or principle of "efficiency", in contrast to the classical numerical criterion — expected utility. In [Dubois and Prade, 1995] qualitative criteria based on possibility theory have been proposed, that are appealing but inefficient in the above sense. The question is whether it is possible to reconcile possibilistic criteria and efficiency. The present paper shows that the answer is yes, and that it leads to special kinds of expected utilities. It is also shown that although numerical, these expected utilities remain qualitative: they lead to two different decision procedures based on min, max and reverse operators only, generalizing the leximin and leximax orderings of vectors.

1 Introduction and motivation

A decision-making problem under uncertainty is a 4-tuple $(S, X, \mathcal{A}, \succeq)$, where S is a set of states of nature, X a set of consequences, $\mathcal{A} = X^S$ the set of possible acts (in decision under uncertainty, an act is a function $f : S \rightarrow X$) and $\succeq$ is a *preference relation* on $\mathcal{A}$, usually complete and transitive (i.e. is a complete preorder).

A numerical approach is classically advocated (see e.g. [Savage, 1954]) for encoding both the information pertaining to the states of nature and the preferences on X : uncertainty is represented by a probability distribution p and preference is encoded by a utility function $u : X \rightarrow [0, 1]$ ¹. The pair $\langle p, u \rangle$ will be called a *probabilistic utility model*, PU-model for short. Acts are then ranked according to their expected utility $EU_{p,u}$ (written here in the finite setting):

$$f \succeq_{EU,p,u} g \Leftrightarrow EU_{p,u}(f) \geq EU_{p,u}(g)$$

where, $\forall h \in \mathcal{A} \quad EU_{p,u}(h) = \sum_{s \in S} p(s) \cdot u(h(s))$.

Information about preference and uncertainty in decision problems cannot always be quantified in a simple way, but only qualitative evaluations can sometimes be attained. As a

consequence, the topic of qualitative decision theory is a natural one to consider [Pearl, 1993; Dubois and Prade, 1995; Brafman and Tennenholtz, 1997; Dubois *et al.*, 1998b; Doyle and Thomason, 1999; Giang and Shenoy, 2000; Dubois *et al.*, 2000]. Giving up the quantification of utility and uncertainty has led to give up the expected utility (EU) criterion as well: the principle of most theories of qualitative decision making is to model uncertainty by an *ordinal* plausibility relation on events and preference by a *complete pre-ordering* on consequences. In [Dubois and Prade, 1995; Dubois *et al.*, 1998b] two qualitative criteria based on possibility theory, an optimistic and a pessimistic one, are proposed and axiomatized whose definitions only require a finite ordinal scale $L = \{0_L \leq \dots \leq 1_L\}$ for evaluating both utility and plausibility:

- $f \succeq_{OPT,\pi,\mu} g \Leftrightarrow U_{OPT,\pi,\mu}(f) \geq U_{OPT,\pi,\mu}(g)$ where $\forall h, U_{OPT,\pi,\mu}(h) = \max_{s \in S} \min(\pi(s), \mu(h(s)))$
- $f \succeq_{PES,\pi,\mu} g \Leftrightarrow U_{PES,\pi,\mu}(f) \geq U_{PES,\pi,\mu}(g)$ where $\forall h, U_{PES,\pi,\mu}(h) = \min_{s \in S} \max(n(\pi(s)), \mu(h(s)))$,

where $n : L \rightarrow L$ is the order reversing function of L , $\pi : S \rightarrow L$ is a normalized possibility distribution and $\mu : X \rightarrow L$ is a utility function on X . In the following, $\langle S, X, L, \pi, \mu \rangle$ will be called a qualitative possibilistic utility model (QPU-model) and we will assume S, X and L to be finite, as is generally the case in qualitative decision making.

The value $U_{PES,\pi,\mu}(f)$ is high only if f gives good consequences in every "rather plausible" state. This criterion generalizes the Wald criterion, which estimates the utility of an act by that of its worst possible consequence. $U_{PES,\pi,\mu}$ is thus "pessimistic" or "cautious", the pessimism being moderated by taking relative possibilities of states into account. On the other hand, $U_{OPT,\pi,\mu}$ is a mild version of the *maximax* criterion which is "optimistic", or "adventurous".

Although appealing from a qualitative point of view, possibilistic utilities suffer from a lack of decisiveness called the "drowning effect": when two acts give an identical and extreme (either good or bad) consequence in some plausible state, they may be undistinguished by these criteria, although they may give significantly different consequences in the other states. As a consequence the principle of Pareto dominance is not satisfied. That is it may be the case that $\forall s, \mu(f(s)) \geq \mu(g(s))$ and that $\exists s^*, \pi(s^*) > 0$ and $\mu(f(s^*)) > \mu(g(s^*))$ but $g \succeq f$.

¹Since expected utility is not sensitive to linear transformations of u , the choice of $[0, 1]$ as the range for u is made for convenience.

Example 1 Let $S = \{s_1, s_2\}$, $L = \{0, 1, 2, 3, 4, 5\}$. Let f and g be two acts whose utilities in states s_1 and s_2 are listed below, as well as the possibility degrees of the states. One can check that $U_{OPT, \pi, \mu}(f) = U_{OPT, \pi, \mu}(g) = 3$ and $U_{PES, \pi, \mu}(f) = U_{PES, \pi, \mu}(g) = 3$ although f strictly dominates g ($\mu(f(s_1)) = \mu(g(s_1))$ and f has a better consequence in s_2).

	s_1	s_2
f	3	4
g	3	1
π	5	2

Most of the qualitative approaches [Pearl, 1993; Dubois and Prade, 1995; Brafman and Tennenholtz, 1997; Giang and Shenoy, 2000], fail to satisfy Pareto dominance. But this is not the case within expected utility theory, since this model obeys the following *Sure-Thing Principle* (STP) that insures that identical consequences do not influence the relative preference between two acts:

$$STP : \forall f, g, h, h', fAh \succeq gAh \Leftrightarrow fAh' \succeq gAh',$$

where fAh denotes the act identical to f on $A \subseteq S$ and to h on $S \setminus A$. When $\succeq$ is complete and transitive, the principle of Pareto dominance is a direct consequence of the STP. So, is it possible to benefit from the STP in the possibilistic framework in order to satisfy the Pareto principle? Unfortunately, it can be shown that this is generally not possible:

Proposition 1 Let $\langle S, X, L, \pi, \mu \rangle$ be a QPU model.

$\succeq_{OPT, \pi, \mu}$ (or $\succeq_{PES, \pi, \mu}$) satisfies the STP
 $\Leftrightarrow \exists ! s^* : \pi(s^*) = 1_L$ and $\forall s \neq s^*, \pi(s) = 0_L$.

This means that possibilistic decision criteria cannot obey the STP, except in a very particular case: when the actual state of the world is known, i.e. when there is no uncertainty at all! So, we cannot stay in the pure QPU framework and escape the drowning effect altogether. The idea is then to cope with the difficulty by proposing *refinements* of the possibilistic criteria that obey the Sure Thing Principle².

This paper shows (Section 2) that any possibilistic model can be refined by an expected utility. The kind of expected utility that is at work, and the very special probability measure that underlies it, are studied in Section 3 under the light of related work. It is also shown (Section 4) that although numerical, these expected utility criteria remain qualitative, since they lead to a decision procedure based solely on min, max and reverse operators — these new procedures generalize well known leximin and leximax decision procedures.

2 Expected utility refinements of qualitative possibilistic utilities

Recall that a refinement $\succeq'$ of a relation $\succeq$ is a relation perfectly compatible with $\succeq$ (it agrees with $\succeq$ when $\succeq$ provides a strict preference), but can break ties by setting $f \succ' g$ for some f, g that are indifferent w.r.t. $\succeq$. Formally:

²The idea of refining QPU first appeared in [Dubois et al., 2000]: the principle was to break ties through an extra criterion (e.g. refining the pessimistic QPU by the optimistic QPU or by another max-norm aggregation). The use of a max operator kept the approach in an ordinal framework, but forbade the full satisfaction of the STP

Definition 1 (Refinement)

$\succeq'$ refines $\succeq \Leftrightarrow \forall f, g \in \mathcal{A}, f \succ g \Rightarrow f \succ' g$.

Since we are looking for complete and transitive relations it is natural to think of refinements based on expected utility. Savage [1954] has indeed shown that, as soon as a complete preorder is desired that satisfies the STP and some very natural axioms, the EU criterion is almost unavoidable. So, the question is: *are there any expected utility criteria that refine the possibilistic criteria?*

Let $\langle S, X, L, \pi, \mu \rangle$ be a QPU model. When considering the optimistic (resp. pessimistic) criterion, we are looking for a probability distribution p and a utility function u such that $\succeq_{EU, p, u}$ refines $\succeq_{OPT, \pi, \mu}$ (resp. $\succeq_{PES, \pi, \mu}$). The idea is to build the EU criteria by means of a transformation $\chi: L \rightarrow [0, 1]$ that maps π to a probability distribution:

Definition 2 (Probabilistic transformation of a scale)

Let $\langle S, X, L, \pi, \mu \rangle$ be a QPU model. A probabilistic transformation of L w.r.t. π is a mapping $\chi: L \rightarrow [0, 1]$ such that $\chi(0_L) = 0$ and $p = \chi \circ \pi$ is a probability distribution.

Notice the presence of the condition $\chi(0_L) = 0$ that expresses the fact that the impossibility of an event (represented by a degree of 0_L in possibility theory) is expressed by a null probability. But the most plausible events (possibility degrees of 1_L), obviously do not receive a probability degree of 1, since they may be mutually exclusive. Notice also that we are looking for a unique function χ for transforming L — both p and u will be built upon this transformation. This is due to the fact that we assume that preference and uncertainty levels are commensurate and belong to the same scale: it is thus natural to transform the degrees regardless whether they model uncertainty or preference.

Moreover, π and μ originally represent all the information available to the user, both in terms of uncertainty of the actual state of the world and preference over possible consequences. So, no undesirable arbitrary information should be introduced in the refined decision model and p and u must be as close as possible to the original information: we are looking for "unbiased" transformations of L . Formally:

Definition 3 (Unbiased transformation of a scale)

Let $\langle S, X, L, \pi, \mu \rangle$ be a QPU model and χ a probabilistic transformation of L . χ is unbiased iff $\forall \alpha, \alpha' \in L, \alpha \leq \alpha' \Leftrightarrow \chi(\alpha) \leq \chi(\alpha')$

As a consequence, using an unbiased χ ensures that π and $p = \chi \circ \pi$ (resp. μ and $u = \chi \circ \mu$) are ordinaly equivalent:

- $\forall s, s' \in S, p(s) \geq p(s') \Leftrightarrow \pi(s) \geq \pi(s')$,
- $\forall x, y \in X, u(x) \geq u(y) \Leftrightarrow \mu(x) \geq \mu(y)$.

2.1 Expected utility refinements of optimistic QPU

Let us first provide a tractable sufficient condition for a probabilistic transformation to generate an expected utility that refines $\succeq_{OPT, \pi, \mu}$:

Proposition 2 Let $\langle S, X, L, \pi, \mu \rangle$ be a QPU model and χ be a probabilistic transformation of L w.r.t. π . Also let $\mathbf{H}$ be the condition:

$$\mathbf{H}: \forall \alpha, \alpha', \beta \in L \text{ s.t. } \beta \geq \alpha > \alpha', \\ \chi(\alpha) \cdot \chi(\beta) > \chi(\alpha') \cdot \chi(\beta) + (|S| - 1) \cdot \chi(1_L) \cdot \chi(\alpha').$$

Then $\succeq_{EU, \chi \circ \pi, \chi \circ \mu}$ refines $\succeq_{OPT, \pi, \mu}$ whenever χ satisfies $\mathbf{H}$.

H is a sufficient condition to generate an EU-refinement of the optimistic QPU (it is also a necessary one when every degree in L is attained by both π and μ). Importantly, there always exists a probabilistic transformation of L satisfying **H**. Let $n_i = |\{s \in S, \pi(s) = \alpha_i\}|$ and $N = |S| + 1$.

Proposition 3

The function $\chi : L = \{\alpha_0 < \dots < \alpha_k\} \rightarrow [0, 1]$ such that $\chi(\alpha_0) = 0$ and $\chi(\alpha_i) = \frac{v}{N^{2k-i}}, i = 1, k$ satisfies **H** $\forall v > 0$. $\chi \circ \pi$ is a probability distribution iff $v = (\sum_{i=1,k} \frac{n_i}{N^{2k-i}})^{-1}$.

In the sequel, χ^* will denote the function $\chi^*(\alpha_i) = \frac{v}{N^{2k-i}}$ obtained with $v = (\sum_{i=1,k} \frac{n_i}{N^{2k-i}})^{-1}$.

Example 2 Let us take the QPU model of Example 1, where $N = 2$, $L = \{0, 1, 2, 3, 4, 5\}$. $\chi^*(L)$ is the series $(0, \frac{v}{N^{16}}, \frac{v}{N^8}, \frac{v}{N^4}, \frac{v}{N^2}, \frac{v}{N})$ where $v = (\frac{1}{N} + \frac{1}{N^8})^{-1}$. So:

$$EU(f) = \chi^*(5) \cdot \chi^*(3) + \chi^*(4) \cdot \chi^*(2) = \frac{v^2}{N^5} + \frac{v^2}{N^{10}},$$

$$EU(g) = \chi^*(5) \cdot \chi^*(3) + \chi^*(2) \cdot \chi^*(1) = \frac{v^2}{N^5} + \frac{v^2}{N^{24}},$$

f is thus preferred to g .

χ^* is in fact sufficient to generate every unbiased EU-refinement of $\succeq_{OPT, \pi, \mu}$, since all such refinements are equivalent.

Definition 4 Two relations $\succeq$ and $\succeq'$ are said to be equivalent ($\succeq \equiv \succeq'$) iff $\forall f, g \in \mathcal{A}, f \succeq g \Leftrightarrow f \succeq' g$.

Proposition 4 Let $\langle S, X, L, \pi, \mu \rangle$ be a QPU model, χ_1, χ_2 two unbiased probabilistic transformations of L w.r.t. π . $\succeq_{EU, \chi_1 \circ \pi, \chi_1 \circ \mu}$ and $\succeq_{EU, \chi_2 \circ \pi, \chi_2 \circ \mu}$ both refine $\succeq_{OPT, \pi, \mu} \Rightarrow \succeq_{EU, \chi_1 \circ \pi, \chi_1 \circ \mu} \equiv \succeq_{EU, \chi_2 \circ \pi, \chi_2 \circ \mu}$

Notice that Proposition 4 does not mean that the numbers attached to the states by $p_1 = \chi_1 \circ \pi$ and $p_2 = \chi_2 \circ \pi$, nor the ones attached to the consequences by $u_1 = \chi_1 \circ \mu$ and $u_2 = \chi_2 \circ \mu$ are the same - it only means that the two models are ordinally equivalent, that they make the same decisions and order the events and the consequences in the same way. It also implies that the refinements that does not belong to this class (they may exist, e.g. those which introduce a total order in S or in X) cannot be unbiased : they must either introduce a strict preference between equivalent consequences or order equiplausible states.

So, we get the following result for optimistic QPU models:

Theorem 1 For any QPU model $\langle S, X, L, \pi, \mu \rangle$:

- There exists an unbiased probabilistic transformation χ^* of L w.r.t. π such that $\succeq_{EU, \chi^* \circ \pi, \chi^* \circ \mu}$ refines $\succeq_{OPT, \pi, \mu}$.
- If χ and χ' are two unbiased transformations of L s.t. both $\succeq_{EU, \chi \circ \pi, \chi \circ \mu}$ and $\succeq_{EU, \chi' \circ \pi, \chi' \circ \mu}$ refine $\succeq_{OPT, \pi, \mu}$, then $\succeq_{EU, \chi \circ \pi, \chi \circ \mu} \equiv \succeq_{EU, \chi' \circ \pi, \chi' \circ \mu}$.

We have hence obtained what we were looking for: for any QPU model we are able to propose an EU model that refines $\succeq_{OPT, \pi, \mu}$. As a refinement, it is perfectly compatible with but more decisive than the optimistic utility. Moreover, it does not use other information than the original one - it is unbiased. Since based on expected utility, it obviously satisfies the Sure Thing Principle as well as Pareto Dominance.

2.2 EU refinements of pessimistic QPU

When considering the pessimistic qualitative model, the same kind of result can be obtained, noticing that $\succeq_{PES, \pi, \mu}$ and $\succeq_{OPT, \pi, \mu}$ are dual relations:

Proposition 5 Let $\langle S, X, L, \pi, \mu \rangle$ be a QPU model. Then: $\forall f, g \in \mathcal{A}, f \succeq_{PES, \pi, \mu} g \Leftrightarrow g \succeq_{OPT, \pi, \mu} f$.

Proposition 6 Let $\langle S, X, L, \pi, \mu \rangle$ be a QPU model and χ be a probabilistic transformation of L w.r.t. π . Let $p = \chi \circ \pi$, $u = \chi \circ \mu$, $u' = \chi(1_L) - \chi \circ n \circ \mu$. It holds that: $\succeq_{EU, p, u}$ refines $\succeq_{OPT, \pi, \mu}$ iff $\succeq_{EU, p, u'}$ refines $\succeq_{PES, \pi, \mu}$.

Consequently, it is always possible to build a probabilistic transformation χ^* using Theorem 1, a probability $p = \chi^* \circ \pi$ and a utility function $u' = \chi^*(1_L) - \chi^* \circ n \circ \mu$ that define an unbiased EU-refinement of $\succeq_{PES, \pi, \mu}$. This provides the following pessimistic counterpart of Theorem 1:

Theorem 2 For any QPU model $\langle S, X, L, \pi, \mu \rangle$:

- There exist at least one unbiased transformation χ^* of L w.r.t. π , a probability distribution $p = \chi^* \circ \pi$ and a utility function $u' = \chi^*(1_L) - \chi^* \circ n \circ \mu$ such that $\succeq_{EU, p, u'}$ refines $\succeq_{PES, \pi, \mu}$.
- Let χ and χ' be two unbiased probabilistic transformations of L . If both $\succeq_{EU, \chi \circ \pi, \chi(1_L) - \chi \circ n \circ \mu}$ and $\succeq_{EU, \chi' \circ \pi, \chi'(1_L) - \chi' \circ n \circ \mu}$ refine $\succeq_{PES, \pi, \mu}$, then they are equivalent.

At this point in the paper we have proved an important result for bridging qualitative possibilistic decision theory and expected utility theory: we have shown than any optimistic or pessimistic QPU model can be refined by a EU model. Thus, we may conclude that (i) possibilistic decision criteria are compatible with the classical expected utility criterion and (ii) choosing a EU model is advantageous, since it leads to a EU-refinement of the original rule (thus overcomes the lack of decisiveness of the possibilistic criteria), it satisfies the STP and the principle of Pareto. But this does not mean that qualitiveness and ordinality are given up. In Section 4, we will show that, although probabilistic and based on additive manipulations of numbers, these criteria remain ordinal. This is very natural: since we start with an ordinal model and do not accept any bias, we produce another (probabilistic but) ordinal model, in which the numbers only encode orders of magnitude — this is the topic of the next Section.

3 EU refinements and big-stepped probabilities : related work

Both EU refinements of Section 2 are based on the same transformation of the possibility π into a probability distribution $p = \chi \circ \pi$ ³. The corresponding measure P is actually a "big-stepped probability" that is, it satisfies:

³But the optimistic utility $u = \chi \circ \mu$ is not equal to the pessimistic utility $u' = \chi(1_L) - \chi \circ n \circ \mu$: u puts the emphasis on the best consequences, while the pessimistic u' provides a high utility when low consequences are avoided. It is actually hopeless to look for a common refinement of $\succeq_{OPT, \pi, \mu}$ and $\succeq_{PES, \pi, \mu}$ since it may happen that $f \succ_{OPT, \pi, \mu} g$ and $g \succ_{PES, \pi, \mu} f$ altogether.

Definition 5 A probability measure P is said to be big stepped iff: $\forall s \in S, P(\{s\}) > P(\{s' \text{ s.t. } P(\{s'\}) < P(\{s\})\})$

In other terms, for any $s, p(s) > \sum_{s' \text{ s.t. } p(s') < p(s)} p(s')$. Such measures are often encountered in the AI literature. First, they have much in common with [Spohn, 1990]'s κ -functions: these disbelief degrees can indeed be interpreted as the order of magnitude of a ϵ probability [Pearl, 1993; Giang and Shenoy, 1999], which is obviously a big stepped probability. Moreover, big stepped probabilities also form a special class of *lexicographic probabilities* in the sense of [Blume et al., 1991; Lehmann, 1998] — we add the restriction that here all the states within a single cluster are equiprobable. Indeed each cluster corresponds to a class of equipossible states and since we are looking for unbiased transformations, equipossibility leads to equiprobability. Finally, Definition 5 generalizes the notion of big-stepped probability of [Snow, 1999; Benferhat et al., 1999] — which is recovered when each cluster is a singleton. Big stepped probabilities have also been proposed by [Dubois et al., 1998a] as a way to refine any possibility/necessity measure ⁴.

This reasoning on the order of magnitude also applies to utility: in a discrete setting, big-stepped utilities can be defined in the same way:

Definition 6 A utility measure u is said to be big stepped iff: $\forall x, x' \in X, u(x) > u(x') \implies u(x) > (|S| - 1) \cdot u(x')$

The utility functions $\chi^* \circ \mu$ and $\chi^*(1_L) - \chi^* \circ n \circ \mu$ are big stepped utilities. It is also the case of the ϵ utilities that underly κ -utility functions [Pearl, 1993; Wilson, 1995; Bonet and Geffner, 1996; Giang and Shenoy, 2000]. These works have advocated an approach to decision under uncertainty based on κ -functions, but without taking the STP into account (decision is made on the order of magnitude only, with a criterion comparable to optimistic utility). The present work makes a step further: in order to satisfy Pareto optimality, we go back to the underlying ϵ utilities and probabilities, using *double* exponents for epsilons instead of simple ones — we remain “big stepped” on the join scale. The other contribution of our approach is that it can be followed to encode pessimistic utilities as well.

[La Valle and Fishburn, 1992; Hammond, 1998; Lehmann, 1998] have studied decision models of *lexicographic probabilities* or *lexicographic utilities*, but in these models, the lexicographic characteristic is used only on one of the two dimensions (either the likelihood level, or the utility level). We operate on both dimensions simultaneously using a join transformation.

4 QPU refinements are qualitative

Although probabilistic and based on additive manipulations of utilities, our EU criteria remain ordinal, as paradoxical as it may seem at first sight. To establish this claim, this Section relates the previous EU criteria to the ordinal comparison of

⁴The “probabilistic likelihood relation” $p = \chi \circ \pi$ is a refinement of both the necessity (N) and possibility (Π) relations based on π . Reciprocally, it is easy to show that, for any big stepped probability P , there exists a π such that P refines both Π and N .

vectors. When S is finite, the comparison of acts can indeed be seen as a comparison of vectors of pairs of elements of L :

Definition 7 The representative vector of any act $f \in \mathcal{A}$ is the vector $\vec{f} = ((\pi_1, \mu_1), \dots, (\pi_i, \mu_i), \dots, (\pi_N, \mu_N))$ where π_i stands for $\pi(s_i)$ and μ_i for $\mu(f(s_i))$.

Comparing acts thus amounts to comparing elements of $(L^2)^N$. For instance, $\succeq_{OPT, \pi, \mu}$ is a restriction to the case $M = 2$ of the general $\succeq_{Maxmin}$ relation on $(L^M)^N$.

Definition 8 (Maxmin relation) Let $\vec{u}, \vec{v} \in (L^M)^N$. Then $\vec{u} \succeq_{Maxmin} \vec{v} \iff \max_{i=1, N} \min_{j=1, M} u_{i,j} \geq \max_{i=1, N} \min_{j=1, M} v_{i,j}$

where $w_{i,j}$ is the j^{th} element of the i^{th} vector of $\vec{w} \in (L^M)^N$

If $\vec{u}$ and $\vec{v}$ are representative of some acts, $M = 2$ and it is obvious that $f \succeq_{OPT, \pi, \mu} g \iff \vec{f} \succeq_{Maxmin} \vec{g}$. In this Section we will propose a refinement of $\succeq_{Maxmin}$, based only on the ordinal comparison of degrees and we will show the equivalence between this purely syntactical decision rule and the above EU models.

4.1 Case of total ignorance

Let us first consider the degenerate case of total ignorance, where $\forall s \in S, \pi(s) = 1_L$. In this case, the comparison of acts comes down to the comparison of utility degrees: $\vec{f} = ((1_L, \mu_1), \dots, (1_L, \mu_N))$ becomes $\vec{f} = (\mu_1, \dots, \mu_N)$. So, $f \succeq_{OPT, \pi, \mu} g$ iff $\vec{f} \succeq_{Max} \vec{g}$ and $f \succeq_{PES, \pi, \mu} g$ iff $\vec{f} \succeq_{Min} \vec{g}$. In decision making, the comparison of vectors by the max and min operators is well known, as it is known that it suffers from a lack of decisive power. That is why refinements of $\succeq_{min}$ and $\succeq_{max}$ have been proposed [Moulin, 1988]:

Definition 9 (Leximax, Leximin) Let $\vec{u}, \vec{v} \in L^N$. Then

- $\vec{u} \succeq_{lmax} \vec{v} \iff (\forall j, u_{(j)} = v_{(j)} \text{ or } \exists i, \forall j < i, u_{(j)} = v_{(j)} \text{ and } u_{(i)} > v_{(i)})$
- $\vec{u} \succeq_{lmin} \vec{v} \iff (\forall j, u_{(j)} = v_{(j)} \text{ or } \exists i, \forall j > i, u_{(j)} = v_{(j)} \text{ and } u_{(i)} > v_{(i)})$

where, for any $\vec{w} \in L^N$, $w_{(k)}$ is the k -th biggest element of $\vec{w}$ (i.e. $w_{(1)} \geq \dots \geq w_{(N)}$).

In practice, the leximin (resp. leximax) comparison consists in ordering both vectors in increasing (resp. decreasing) order and then lexicographically comparing them.

Example 3 Let $\vec{u} = (3, 2, 4)$ and $\vec{v} = (2, 2, 4)$: $\vec{u} \succ_{lmax} \vec{v}$ since $u_{(1)} = v_{(1)} = 4$ and $u_{(2)} = 3 > v_{(2)} = 2$; $\vec{u} \succ_{lmin} \vec{v}$ since $u_{(3)} = v_{(3)} = 4$ and $u_{(2)} = 3 > v_{(2)} = 2$.

It is obvious that $\succeq_{lmax}$ refines $\succeq_{max}$ and $\succeq_{lmin}$ refines $\succeq_{min}$. Moreover, both relations escape the drowning effect and are very efficient: the only pairs of ties are vectors that are identical up to a permutation of their elements.

4.2 General Case

Since the leximax and leximin comparisons are good candidates in a particular case, we have imagined an extension of these procedures to the case of 2 dimensions ($(L^M)^N$ instead of L^N). The only thing that we need is to use any complete preorder $\succeq$ on vectors of L^M instead of the classical relation $\geq$ on L . It is then possible to order the sub-vectors of any $\vec{w}$ according to $\succeq$ and to apply any of the previous procedures:

Definition 10 (Leximax($\succeq$), Leximin($\succeq$)) Let $\succeq$ be a complete preorder on L^M , $\cong$ the associated equivalence relation and $\triangleright$ is the strict part of $\succeq$. Let $\vec{u}, \vec{v} \in (L^M)^N$. Then

- $\vec{u} \succeq_{\text{leximax}(\succeq)} \vec{v} \Leftrightarrow (\forall j, u_{(\succeq, j)} \cong v_{(\succeq, j)} \text{ or } \exists i \text{ s.t. } \forall j < i, u_{(\succeq, j)} \cong v_{(\succeq, j)} \text{ and } u_{(\succeq, i)} \triangleright v_{(\succeq, i)})$.
- $\vec{u} \succeq_{\text{leximin}(\succeq)} \vec{v} \Leftrightarrow (\forall j, u_{(\succeq, j)} \cong v_{(\succeq, j)} \text{ or } \exists i \text{ s.t. } \forall j > i, u_{(\succeq, j)} \cong v_{(\succeq, j)} \text{ and } u_{(\succeq, i)} \triangleright v_{(\succeq, i)})$.

where, for any $\vec{w} \in (L^M)^N$, $w_{(\succeq, i)}$ is the i^{th} biggest sub-vector of $\vec{w}$ according to $\succeq$.

The leximax procedure can in particular be applied to the preorder $\succeq = \succeq_{\text{leximin}}$. In practice, this comparison consists in first ordering the elements of each sub-vector in increasing order w.r.t $\succeq$, then in ordering the sub-vectors in decreasing order (w.r.t. $\succeq_{\text{leximin}}$). It is then enough to lexicographically compare the two new vectors of vectors.

Proposition 7 ($\succeq_{\text{leximax}(\succeq_{\text{leximin}})$ order)

- i) $\succeq_{\text{leximax}(\succeq_{\text{leximin}})$ is a complete preorder;
- ii) $\succeq_{\text{leximax}(\succeq_{\text{leximin}})$ refines $\succeq_{\text{maxmin}}$;
- iii) If $N = 1$, then $\succeq_{\text{leximax}(\succeq_{\text{leximin}}) \equiv \succeq_{\text{leximin}}$;
- iv) If $M = 1$, then $\succeq_{\text{leximax}(\succeq_{\text{leximin}}) \equiv \succeq_{\text{leximax}}$.

So, $\succeq_{\text{leximax}(\succeq_{\text{leximin}})$ is the refinement of $\succeq_{\text{maxmin}}$ we desire. Let us now compare representative vectors of acts using this relation (letting $M = 2$) – we get a refinement of $\succeq_{\text{OPT}, \pi, \mu}$:

Proposition 8 The relation $\succeq_{\text{leximax}(\succeq_{\text{leximin}})$ defined by $f \succeq_{\text{leximax}(\succeq_{\text{leximin}})} g \Leftrightarrow \vec{f} \succeq_{\text{leximax}(\succeq_{\text{leximin}})} \vec{g}$ refines $\succeq_{\text{OPT}, \pi, \mu}$.

Example 4 The representative vectors of f and g in Example 1 are: $\vec{f} = ((5, 3), (2, 4))$, $\vec{g} = ((5, 3), (2, 1))$. $(5, 3) \cong_{\text{leximin}} (5, 3)$, $(4, 2) \succ_{\text{leximin}} (2, 1)$, so $\vec{f} \succ_{\text{leximax}(\succeq_{\text{leximin}})} \vec{g}$.

The same kind of reasoning can be followed to refine $\succeq_{\text{PES}, \pi, \mu}$. The pessimistic utility of act f is $\min_{s \in S} \max(n(\pi(s)), \mu(f(s)))$. We need to refine a minimax procedure, and this can be done using $\succeq_{\text{leximin}(\succeq_{\text{leximax}})$. Since operator $\max$ does not apply to $(\pi(s), \mu(f(s)))$ but to $(n(\pi(s)), \mu(f(s)))$, we use the π -reverse vectors of acts:

Definition 11 The π -reverse vector of an act $f \in \mathcal{A}$ is $n(f) = ((n(\pi(s_1)), \mu(f(s_1))), \dots, (n(\pi(s_N)), \mu(f(s_N))))$.

Proposition 9 The relation $\succeq_{\text{leximin}(\succeq_{\text{leximax}, n})$ defined by $f \succeq_{\text{leximin}(\succeq_{\text{leximax}, n})} g \Leftrightarrow n(\vec{f}) \succeq_{\text{leximin}(\succeq_{\text{leximax}})} n(\vec{g})$

refines $\succeq_{\text{PES}, \pi, \mu}$.

So, the lexi-refinement of $\succeq_{\text{PES}, \pi, \mu}$ applies the leximin(leximax) comparison to the π -reverse vectors, while the refinement of $\succeq_{\text{OPT}, \pi, \mu}$ applies the leximax(leximin) comparison directly to the representative vectors. Both procedures are purely ordinal: the degrees in L are only compared using min, max and reverse operators — only their relative orders matter. Our final result is that these refinements are equivalent to the EU-refinements identified in Section 2.

Theorem 3 Let $\langle S, X, L, \pi, \mu \rangle$ be a QPU model and χ a probabilistic transformation of L w.r.t. π , $p = \chi \circ \pi$, $u = \chi \circ \mu$, $u' = \chi(1_L) - \chi \circ n \circ \mu$:

- i) $\succeq_{\text{EU}, p, u}$ refines $\succeq_{\text{OPT}, \pi, \mu} \Leftrightarrow \succeq_{\text{EU}, p, u} \equiv \succeq_{\text{leximax}(\succeq_{\text{leximin}})$.
- ii) $\succeq_{\text{EU}, p, u'}$ refines $\succeq_{\text{PES}, \pi, \mu} \Leftrightarrow \succeq_{\text{EU}, p, u'} \equiv \succeq_{\text{leximin}(\succeq_{\text{leximax}, n})$.

So, the probabilistic refinements of possibilistic utilities are equivalent to purely comparative procedures: efficient QPU refinements are probabilistic but remain qualitative. Reciprocally, we can prove that the $\succeq_{\text{leximax}(\succeq_{\text{leximin}})$ and $\succeq_{\text{leximin}(\succeq_{\text{leximax}, n})}$ preference relations over vectors of vectors always admit a representation by a sum (on N) of products (on M), provided that L is discrete – but this is beyond the scope of this paper, that focuses on decision under uncertainty (where $M=2$).

5 Conclusion

The topic of Qualitative Decision Theory has received much attention in the past few years and several approaches, including QPU, have been proposed. This latter model forms a convenient framework for a qualitative expression of problems of decision under uncertainty. However, it suffers from a lack of decisiveness. We have proposed EU-based refinements of QPU which proved to be perfectly compatible with the original qualitative expression of knowledge and preferences: the only difference is that lexicographic (leximax(leximin) or leximin(leximax)) comparisons are used instead of maxmin or minmax. The axiomatization of these decision procedures is out of the scope of this paper. It actually consists in the 5 basic Savagean axioms, together with “mild” versions of the pessimism or optimism axioms of possibilistic utilities (see [Fargier and Sabbadin, 2003]).

κ -utility functions have also received much attention in AI but they too suffer from a lack of decisive power. Indeed, due to the $\min$ operator in $U_\kappa(f) = \min_{s \in S} \{\kappa(s) + \mu(f(s))\}$, they cannot satisfy Savage’s STP. We claim that $\succeq_{\kappa, \mu}$ can be EU-refined, in the same way as we did for QPU, and that the EU-refinement can be expressed syntactically by replacing the $\min$ operator by a leximin one. The proof of this conjecture, which would follow the same line as the one presented here, is left for future research.

The other extension of our work is to refine monotonic utilities [Dubois et al., 1998c] a family of decision criteria that generalizes QPU and are based on the Sugeno integral. Monotonic Utilities admit a maxmin expression $\max_{s \in S} (\min(\sigma(F_s), \mu(f(s))))$, where $\sigma : 2^S \rightarrow L$ is a monotonic measure and $F_s = \{s' \in S \text{ s.t. } \mu(f(s')) \geq \mu(f(s))\}$. Again, we conjecture that replacing $\min$ and $\max$ with their lexicographical versions allows an efficient refinement and lead to Choquet-EU criteria [Gilboa, 1987].

Appendix

Most of the proofs are omitted for sake of brevity and can be found in [Fargier and Sabbadin, 2003] or at the following address : <http://www-bia.inra.fr/T/sabbadin/WEB/FargierSabbadin03Rap.html>.

We provide here the sketches of the most interesting ones.

Proof of Proposition 1 $\Leftarrow$ is trivial to show.

$\Rightarrow$ For any $x, y \in X$, $A \subseteq S$, let us denote xAy the act such that $xAy(s) = x$ if $s \in A$, $xAy(s) = y$ otherwise. Let x_0 and x_1 be two consequences such that $\mu(x_0) = 0_L$ and $\mu(x_1) = 1_L$.

Suppose there exist s_1, s_2 , $s_1 \neq s_2$, $\pi(s_1) \geq \pi(s_2) > 0_L$. Then, we have $x_1\{s_2\}x_0 \succ_{\text{OPT}, \pi, \mu} x_0s_1x_1$ and $x_1\{s_1, s_2\}x_0 \preceq_{\text{OPT}, \pi, \mu} x_1\{s_1\}x_0$, which violates the STP. So, if

$\succeq_{OPT, \pi, \mu}$ satisfies the STP then there exists at most one s^* such that $\pi(s^*) > 0_L$. Since π is normalized, s^* exists and $\pi(s^*) = 1_L$. In the pessimistic case, just notice that $x1 \succ_{PES, \pi, \mu} x0\{s_2\}x1$ whereas $x0\{s_1, s_2\}x1 \succeq_{PES, \pi, \mu} x0\{s_1\}x1$. $\square$

Sketch of the proof of Proposition 4. Let us denote $L_\mu = \{\alpha_i \in L, \exists x \in X, \mu(x) = \alpha_i\}$ — let us rank $L_\mu = (\alpha_0 < \dots < \alpha_i < \dots < \alpha_k)$ and $L_\pi = \{\beta_i \in L, \exists s \in S, \pi(s) = \beta_i\}$ — let us rank $L_\pi = (\beta_0 < \dots < \beta_i < \dots < \beta_{k'})$. It can be assumed (without loss of generality) that $L_\pi \subseteq L_\mu = L$.

Consider any $\alpha_i \in L$ and $\beta_j = \min\{\beta_i \in L_\pi, \beta_i \geq \alpha_i\}$. There exists a pair (s, x) where $\pi(s) = \beta_j, \mu(x) = \alpha_i$. Let us also denote $x0$ (resp. $x1$) a consequence of utility 0_L (resp. 1_L).

The proof of Proposition 4 is based on the observation that act $f : f(s) = x, f(s') = x0$ is always strictly preferred by $\succeq_{OPT, \pi, \mu}$ to act $g : g(s') = x1$ when $\pi(s') < \alpha_i$ and $g(s') = \alpha_{i-1}$ when $\pi(s') \geq \alpha_i$.

From this, it can be shown that (i) the sub products, in the expression of the expected utility, are ranked in the same way by any χ and that (ii) the act with the biggest j -product is surely strictly preferred to the other act (whatever the values of the remaining terms).

Lemma 1 Let χ_1 and χ_2 two unbiased transformation of L refining $\succeq_{OPT, \pi, \mu}$. It holds that: $\forall \alpha, \alpha' \in L_\mu, \forall \beta, \beta' \in L_\pi$: $\chi_1(\alpha) \cdot \chi_1(\beta) > \chi_1(\alpha') \cdot \chi_1(\beta') \iff \chi_2(\alpha) \cdot \chi_2(\beta) > \chi_2(\alpha') \cdot \chi_2(\beta')$

Lemma 2 $f \succ_{EU, \chi_1 \circ \pi, \chi_1 \circ \mu} g \iff \exists j/\forall i < j, \chi(\pi_{\sigma f(i)}) \cdot \chi(\mu_{\sigma f(i)}) = \chi(\pi_{\sigma g(i)}) \cdot \chi(\mu_{\sigma g(i)})$ and $\chi(\pi_{\sigma f(j)}) \cdot \chi(\mu_{\sigma f(j)}) > \chi(\pi_{\sigma g(j)}) \cdot \chi(\mu_{\sigma g(j)})$.

Lemmas 1 and 2 together are enough to prove the equivalence of $\succeq_{EU, \chi_1 \circ \pi, \chi_1 \circ \mu}$ and $\succeq_{EU, \chi_2 \circ \pi, \chi_2 \circ \mu}$, whatever χ_1 and χ_2 : Indeed, $f \succ_{EU, \chi_1 \circ \pi, \chi_1 \circ \mu} g \iff$ (from Lemma 2) $\exists j/\forall i < j, \chi_1(\pi_{\sigma f(i)}) \cdot \chi_1(\mu_{\sigma f(i)}) = \chi_1(\pi_{\sigma g(i)}) \cdot \chi_1(\mu_{\sigma g(i)})$ and $\chi_1(\pi_{\sigma f(j)}) \cdot \chi_1(\mu_{\sigma f(j)}) > \chi_1(\pi_{\sigma g(j)}) \cdot \chi_1(\mu_{\sigma g(j)})$. But, from Lemma 1, the comparison of products are the same if we change χ_1 into χ_2 , so we get: $\exists j/\forall i < j, \chi_2(\pi_{\sigma f(i)}) \cdot \chi_2(\mu_{\sigma f(i)}) = \chi_2(\pi_{\sigma g(i)}) \cdot \chi_2(\mu_{\sigma g(i)})$ and $\chi_2(\pi_{\sigma f(j)}) \cdot \chi_2(\mu_{\sigma f(j)}) > \chi_2(\pi_{\sigma g(j)}) \cdot \chi_2(\mu_{\sigma g(j)})$, which implies (Lemma 2) that $f \succ_{EU, \chi_2 \circ \pi, \chi_2 \circ \mu} g$.

The same reasoning, simply swapping χ_1 and χ_2 implies $f \succ_{EU, \chi_2 \circ \pi, \chi_2 \circ \mu} g \implies f \succ_{EU, \chi_1 \circ \pi, \chi_1 \circ \mu} g$, from which we get the equivalence of $\succeq_{EU, \chi_1 \circ \pi, \chi_1 \circ \mu}$ and $\succeq_{EU, \chi_2 \circ \pi, \chi_2 \circ \mu}$. $\square$

References

- [Benferhat et al., 1999] S. Benferhat, D. Dubois, and H. Prade. Possibilistic and standard probabilistic semantics of conditional knowledge. *Journal of Logic and Computation*, 9:873–895, 1999.
- [Blume et al., 1991] L. Blume, A. Brandenburger, and E. Dekel. Lexicographic probabilities and choice under uncertainty. *Econometrica*, 59(1):61–79, 1991.
- [Bonet and Geffner, 1996] B. Bonet and H. Geffner. Arguing for decisions: A qualitative model of decision making. In *Proc. of UAI'96*, pages 98–105, 1996.
- [Brafman and Tennenholtz, 1997] R. Brafman and M. Tennenholtz. Modeling agents as qualitative decision makers. *Artificial Intelligence*, 94:217–268, 1997.
- [Doyle and Thomason, 1999] J. Doyle and R. H. Thomason. Background to qualitative decision theory. *AI Magazine*, 20(2):55–68, 1999.
- [Dubois and Prade, 1995] D. Dubois and H. Prade. Possibility theory as a basis for qualitative decision theory. In *Proc. of IJCAI'95*, pages 1925–1930, 1995.
- [Dubois et al., 1998a] D. Dubois, H. Fargier, and H. Prade. Possibilistic likelihood relations. In *Proc. of IPMU'98*, pages 1196–1203. EDK, 1998.
- [Dubois et al., 1998b] D. Dubois, L. Godo, H. Prade, and A. Zapico. Making decision in a qualitative setting: From decision under uncertainty to case-based reasoning. In *Proc. of KR'98*, pages 594–605, 1998.
- [Dubois et al., 1998c] D. Dubois, H. Prade, and R. Sabbadin. Qualitative decision theory with sugeno integrals. In *Proc. of UAI'98*, pages 121–128, 1998.
- [Dubois et al., 2000] D. Dubois, L. Godo, H. Prade, and A. Zapico. Advances in qualitative decision theory: Refined rankings. In *Proc. IBERAMIA'00*, pages 427–436, 2000.
- [Fargier and Sabbadin, 2003] H. Fargier and R. Sabbadin. Bridging the gap between qualitative decision theory and expected utility. Technical Report IRIT/03-08-R, IRIT, April 2003.
- [Giang and Shenoy, 1999] P. H. Giang and P. P. Shenoy. On transformations between probability and spohnian disbelief functions. In *Proc. of UAI'99*, pages 236–244, 1999.
- [Giang and Shenoy, 2000] P. H. Giang and P. P. Shenoy. A qualitative utility theory for spohn's theory of epistemic beliefs. In *Proc. of UAI'2000*, pages 220–227, 2000.
- [Gilboa, 1987] I. Gilboa. Expected utility with purely subjective non-additive probabilities. *Journal of Mathematical Economics*, 16:65–88, 1987.
- [Hammond, 1998] P. J. Hammond. *Consequentialism, non archimedean probabilities and lexicographic expected utility*. In: *The logic of strategy*. C. Bicchieri, R. Jeffrey and B. Skyrms, 1998.
- [La Valle and Fishburn, 1992] I. H. La Valle and P. C. Fishburn. State-independent subjective expected lexicographic utility. *Journal of Risk and Uncertainty*, 5:217–240, 1992.
- [Lehmann, 1998] D. Lehmann. Nonstandard numbers for qualitative decision making. In *Proc. of TARK'98*, pages 161–174, 1998.
- [Moulin, 1988] H. Moulin. *Axioms of Cooperative Decision Making*. Wiley, New-York, 1988.
- [Pearl, 1993] J. Pearl. From conditional oughts to qualitative decision theory. In *Proc. of UAI'93*, pages 12–20, 1993.
- [Savage, 1954] L.J. Savage. *The Foundations of Statistics*. Wiley, New York, 1954.
- [Snow, 1999] P. Snow. Diverse confidence levels in a probabilistic semantics for conditional logics. *Artificial Intelligence*, 113(1–2):269–279, 1999.
- [Spohn, 1990] W. Spohn. A general non-probabilistic theory of inductive reasoning. In *Proc. of UAI'90*, pages 315–322. Elsevier Science, 1990.
- [Wilson, 1995] N. Wilson. An order of magnitude calculus. In *Proc. of UAI'95*, pages 548–555, 1995.