Optimizing the Hurwicz criterion in decision trees with imprecise probabilities

Gildas Jeantet, Olivier Spanjaard

Optimizing the Hurwicz criterion in decision trees with imprecise probabilities

Gildas Jeantet and Olivier Spanjaard

LIP6 - UPMC
104 avenue du Président Kennedy 75016 Paris, France
{gildas.jeantet,olivier.spanjaard}@lip6.fr

Abstract. This paper is devoted to sequential decision problems with imprecise probabilities. We study the problem of determining an optimal strategy according to the Hurwicz criterion in decision trees. More precisely, we investigate this problem from the computational viewpoint. When the decision tree is *separable* (to be defined in the paper), we provide an operational approach to compute an optimal strategy, based on a bicriteria dynamic programming procedure. The results of numerical tests are presented. When the decision tree is *non-separable*, we prove the NP-hardness of the problem.

Key words: Sequential decision making; Imprecise probabilities; Hurwicz's criterion; Computational complexity; Exact algorithms

1 Introduction

Decision under uncertainty is one of the main field of research in decision theory, due to its numerous applications (e.g. medical diagnosis, robot control, strategic decision, games...). Decision under uncertainty means that the consequences of a decision depends on uncertain events. In decision under risk, it is customary to assume that a precise probability is known for each event appearing in the decision problem. A decision can thus be characterized by a *lottery* over possible consequences. A popular criterion to compare lotteries (and therefore decisions) is the expected utility (EU) model proposed by von Neumann and Morgenstern [10]. In this model, a *utility function u* (specific to each decision maker) assigns a numerical value to every outcome. The evaluation of a lottery is then performed via the computation of its utility expectation (the greater the better). However, when several experts have divergent viewpoints or when empirical data are missing, it is not obvious to elicit sharp numerical probabilities for each event. A natural way to take into account this difficulty is to use intervals of probabilities rather than scalar probabilities. This is known as decision making under imprecise probabilities.

Comparing decisions amounts then to comparing imprecise lotteries, i.e. lotteries where several possible probability distributions are taken into account. A pessimistic agent will make the decision that maximizes the worst possible expected utility. This is known as the *Γ-maximin* decision criterion. Conversely, an optimistic agent will make the decision that maximizes the best possible expected utility. This is known as the *Γ-maximax* decision criterion. Including these two extremes, Jaffray and Jeleva recently proposed to use the Hurwicz criterion, that enables to model intermediate attitudes by performing a linear combination of both previous criteria [3]. Note that Hurwicz introduced this criterion in the context of decision under complete ignorance (i.e., when absolutely no information is known about the probabilities), but the authors preserved its denomination of "Hurwicz's criterion" since it extends naturally to the case of imprecise probabilities.

To our knowledge, the algorithmic issues related to the use of Hurwicz's criterion in a sequential decision problem with imprecise probabilities have not been studied until now. It is indeed frequent to encounter *sequential decision problems* where one does not make a simple decision but one follows a *strategy* (i.e. a sequence of decisions conditioned by events) resulting in a non deterministic outcome. Several representation formalisms can be used for sequential decision problems, such as decision trees (e.g., [8]), influence diagrams (e.g., [9]) or Markov decision processes (e.g., [7]). A decision tree is an explicit representation of a sequential decision problem, while influence diagrams or Markov decision processes are compact representations and make it possible to deal with decision problems of greater size. It is important to note that, in all these formalisms, the set of potential strategies is combinatorial (i.e., its size increases exponentially with the size of the instance). The computation of an optimal strategy for a given representation and a given decision criterion is then an algorithmic issue in itself. It is well-know that an optimal strategy for EU in a decision tree endowed with scalar probabilities can be determined in linear time by backward induction. This is no more the case when dealing with imprecise probabilities and Hurwicz's criterion. In the particular case of Γ-maximin and Γ-maximax criteria, Kikuti *et al.* [1] have presented algorithms that employ *dynamic feasibility*, that is, one declares infeasible any strategy that includes a suboptimal substrategy (a substrategy is a strategy in a subtree). In the present paper, on the contrary, we consider that all strategies are feasible (i.e., even the ones that include a suboptimal substrategy), and we study the computational complexity of determining an optimal strategy according to Hurwicz's criterion in a decision tree endowed with imprecise probabilities. Furthermore, we propose algorithmic procedures to tackle the problem.

The remainder of the paper is organized as follows. We first give some preliminaries on imprecise probabilities and decision criteria used in such a setting (Section 2). Then, we present the difficulties raised by the use of imprecise probabilities in sequential decision problems, and we distinguish a separable case and a non-separable case (Section 3). The next two sections are devoted to the description of our results in these two cases (Section 4 and 5). Finally, we conclude by giving some avenues for future research (Section 6).

2 Single stage decision making with imprecise probabilities

Several mathematical models of imprecise probabilities have been proposed in the literature [11, 12]. A common point between these models is that they often define a probability interval $[P^-(E), P^+(E)]$ for each event E. Following Jaffray and Jeleva [3], we assume that there exists a *real* probability P_0 such that $P_0(E) \in [P^-(E), P^+(E)]$ for all events E. To compare imprecise lotteries (i.e., lotteries with imprecise probabilities), one must therefore consider a set $\mathcal{P}$ of possible probability distributions. This is close to the approach adopted to compare feasible solutions in discrete optimization with interval data [4], with the difference that the set of possible probability distributions is not the cartesian product of the probability intervals of the events. A probability distribution should indeed satisfy the Kolmogorov axioms ($P(E) \geq 0$, $P(\Omega) = 1$, $P(E_1 \cup E_2 \cup \ldots) = P(E_1) + P(E_2) + \ldots$ for pairwise disjoint events E_i).

Let us present popular decision criteria in such a setting. For instance, consider two lotteries f, g involving three pairwise disjoint events E_1, E_2, E_3. If E_1 (resp. E_2, E_3) occurs, f yields -50 (resp. 0,100). If E_1 (resp. E_2, E_3) occurs, g yields 130 (resp. -30,-50). In the EU model with sharp probabilities, a lottery is evaluated by its expected utility, namely $E(f) = P(E_1)u(-50) + P(E_2)u(0) + P(E_3)u(100)$ for f. Assume now that probabilities are imprecise, e.g. $P_0(E_1) \in [0.2, 0.4]$, $P_0(E_2) \in [0.4, 0.6]$ and $P_0(E_3) \in [0.2, 0.3]$. The set $\mathcal{P}$ of possible probability distributions is therefore defined by $\mathcal{P} = \{P : P(E_i) \in [P^-(E_i), P^+(E_i)] \ \forall i, \text{ and } \sum_i P(E_i) = 1\}$. If the decision maker wants to hedge against the worst possible expected utility, a lottery f is evaluated by $\underline{E}(f) = \min\{E(f, P) : P \in \mathcal{P}\}$ where $E(f, P)$ denotes the expected utility of lottery f according to probability P. This is the so-called *Γ-maximin* decision criterion. The value of the Γ-maximin criterion can be computed by using the following simple result:

Proposition 1 *Consider a lottery f yielding utility u_i if event E_i occurs ($i = 1, \ldots, n$), with $u_1 \leq \ldots \leq u_n$ and $P(E_i) \in [P^-(E_i), P^+(E_i)]$. The probability distribution $\underline{P}_f$ in $\mathcal{P}$ recursively defined by*

$$\begin{cases} \underline{P}_f(E_1) = \min\{1 - \sum_{j=2}^n P^-(E_j), P^+(E_j)\} \\ \underline{P}_f(E_i) = \min\{1 - \sum_{j=1}^{i-1} \underline{P}_f(E_j) - \sum_{j=i+1}^n P^-(E_j), P^+(E_i)\} \quad \forall i \end{cases}$$

yields expected utility $\underline{E}(f)$.

Proof. Consider a probability distribution $P \neq \underline{P}_f$. Let us show that $E(f, \underline{P}_f) \leq E(f, P)$. We denote by i_0 the index such that $P(E_i) = \underline{P}_f(E_i)$ for $i < i_0$ and $P(E_{i_0}) < \underline{P}_f(E_{i_0})$ ($P(E_{i_0}) > \underline{P}_f(E_{i_0})$ is impossible). One should have $\sum_{i=i_0}^n P(E_i) = 1 - \sum_{i=1}^{i_0-1} \underline{P}_f(E_i)$. Consequently, $P(E_{i_0}) < \underline{P}_f(E_{i_0})$ implies that $P(E_i) > P^-(E_i)$ for some $i > i_0$. Let us set $i_1 = \min\{i : i > i_0 \text{ and } P(E_i) > P^-(E_i)\}$ and $\varepsilon = \min\{\underline{P}_f(E_{i_0}) - P(E_{i_0}), P(E_{i_1}) - P^-(E_{i_1})\} > 0$. We denote by P_1 the probability distribution defined by $P_1(E_{i_0}) = P(E_{i_0}) + \varepsilon$, $P_1(E_{i_1}) = P(E_{i_1}) - \varepsilon$ and $P_1(E_i) = P(E_i)$ for $i \neq i_0, i_1$. We have $E(f, P_1) \leq E(f, P)$ since

$E(f,P_1) - E(f,P) = \varepsilon(u_{i_0} - u_{i_1}) \leq 0$. If $P_1 \neq \underline{P}_f$, by the same reasoning one can construct a probability distribution P_2 such that $E(f,P_2) \leq E(f,P_1)$. In this way, one generates a sequence $P_1,\dots,P_k$ of probability distributions such that $E(f,P_{i+1}) \leq E(f,P_i)$ and $P_k = \underline{P}_f$. Therefore $E(f,\underline{P}_f) \leq E(f,P)$. ■

For instance, let us come back to lotteries f, g previously mentioned. We have $\underline{P}_f(E_1) = \min\{1-0.4-0.2, 0.4\} = 0.4$, $\underline{P}_f(E_2) = \min\{1-0.4-0.2, 0.6\} = 0.4$ and $\underline{P}_f(E_3) = \min\{1-0.4-0.4, 0.3\} = 0.2$. Consequently, for $u(x) = x$, we have $\underline{E}(f) = 0.4 \times (-50) + 0.2 \times 100 = 0$. Similarly, one computes $\underline{P}_g(E_1) = 0.2$, $\underline{P}_g(E_2) = 0.5$, $\underline{P}_g(E_3) = 0.3$ and $\underline{E}(g) = -4$. Therefore lottery f is preferred to g for the Γ-maximin criterion.

Conversely, if the decision maker wants to maximize the best possible expected utility, a lottery f is evaluated by $\bar{E}(f) = \max\{E(f,P) : P \in \mathcal{P}\}$. This is the so-called *Γ-maximax* decision criterion. The probability distribution $\bar{P}_f$ yielding $\bar{E}(f)$ is defined by:

$$\begin{cases} \bar{P}_f(E_1) = \max\{1 - \sum_{j=2}^{n} P^+(E_j), P^-(E_j)\} \\ \bar{P}_f(E_i) = \max\{1 - \sum_{j=1}^{i-1} \bar{P}_f(E_j) - \sum_{j=i+1}^{n} P^+(E_j), P^-(E_i)\} \quad \forall i \end{cases}$$

Coming back again to lotteries f, g previously mentioned, we have $\bar{P}_f(E_1) = 0.2$, $\bar{P}_f(E_2) = 0.5$, $\bar{P}_f(E_3) = 0.3$, $\bar{E}(f) = 20$ on the one hand, and $\bar{P}_g(E_1) = 0.4$, $\bar{P}_g(E_2) = 0.4$, $\bar{P}_g(E_3) = 0.2$, $\bar{E}(g) = 30$ on the other hand. Therefore lottery g is preferred to f for the Γ-maximax criterion. This shows that the preferences are of course very dependent on the degree of pessimism of the decision maker.

For this reason, Jaffray and Jeleva [3] propose to extend the Hurwicz criterion for decision under complete ignorance to the case of imprecise probabilities. According to the Hurwicz criterion, a lottery f is evaluated by $\alpha\underline{E}(f) + (1-\alpha)\bar{E}(f)$. In other words, the decision maker will look at the worse and best possible expected utilities and, according to its degree of pessimism, will put more or less weight on the former or the later. It reduces to Γ-maximin for $\alpha = 1$, and to Γ-maximax for $\alpha = 0$. When comparing lotteries f, g previously mentioned according to the Hurwicz criterion, we have f preferred to g for $\alpha > 5/7$, and g preferred to f for $\alpha < 5/7$. Note that the Hurwicz criterion is compatible with dominance, i.e. if a lottery has a greater expected utility than another one for all possible probability distributions, then its evaluation will be better [3]. This property is indeed desirable to guarantee a rational behavior.

3 Multistage decision making with imprecise probabilities

In multistage decision making, one studies problems where one has to take a sequence of decisions conditionally to events. The formalism of decision trees provides a simple and explicit representation of a sequential decision problem under risk. It is a tree with three kinds of nodes: decision nodes (represented by squares), chance nodes (represented by circles) and utility nodes (leaves of the tree). A decision node (resp. chance node) can be seen as a decision variable (resp. random variable), the domain of which corresponds to the labels of the branches starting from that node. When probabilities are imprecise, the sharp probability that a given random variable takes a given value is unknown: one only knows an interval of probabilities in which it is included. The values indicated at the leaves correspond to the *utilities* of the consequences. For the sake of illustration, we now give an example of a well-kown multistage decision problem, and its representation with a decision tree. Note that one omits the orientation of the edges when representing decision trees.

Example 1 (oil wildcatter's problem [8]) *An oil wildcatter has to decide whether to drill or not at a given site. For that purpose, he first has to decide whether to sound or not the geological structure of the site (decision D_1), which costs 10000\$ and gives a better estimation of the quantity of oil to be found. The result of the sounding can be seen as a random variable T that can take three possible values: no if there is no hope of oil, open if some oil is expected, or closed if much oil is expected. Next, he decides whether to drill or not (decision D_2), which costs 70000\$. Finally, if he decides to drill, the result of the drilling can be seen as a random variable S that can take three possible values: the hole is dry (the outcome is 0\$), wet (120000\$) or soaking (270000\$). This problem can be represented by the decision tree on the left side of Figure 1. Note that decision D_2 is duplicated in several nodes (nodes D_2^1, D_2^2, D_2^3 and D_2^4) since it can be taken in several different contexts (a sounding has been performed or not, the result of the sounding is encouraging or not...).*

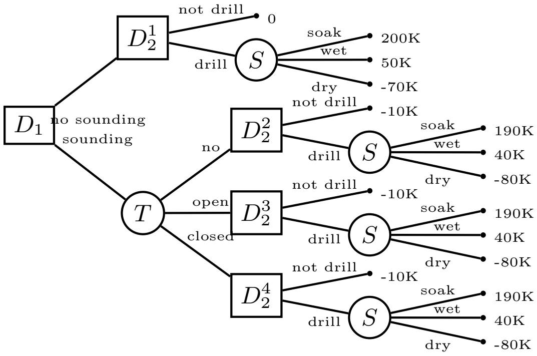

$P(S\|T)$	dry	wet	soak
no	[0.500,0.666]	[0.222,0.272]	[0.125,0.181]
open	[0.222,0.333]	[0.363,0.444]	[0.250,0.363]
closed	[0.111,0.166]	[0.333,0.363]	[0.454,0.625]

T	no	open	closed
$P(T)$	[0.181,0.222]	[0.333,0.363]	[0.444,0.454]

S	dry	wet	soak
$P(S)$	[0.214,0.344]	[0.309,0.386]	[0.307,0.456]

Fig. 1. Decision tree for the oil wildcatter problem.

When sharp probabilities are known, each branch starting from a chance node representing random variable X is endowed with probability $P(X = x|past(X))$, where $past(X)$ denotes all the value assignments to random and decision variables on the path from the root to X. Furthermore, in this paper, we assume that $P(X = x|past(X))$ only depends on the random variables in $past(X)$. For instance, in the decision tree for the oil wildcatter problem, $P(S = soak|D_1 = sounding, T = no) = P(S = soak|T = no)$. When probabilities are imprecise, we assume that a conditional probability table is indicated for each chance node in the decision tree. In each cell of the table, an interval of probabilities is given. For the oil wildcatter problem, the conditional probability tables are presented besides the decision tree in Figure 1. So as to have complete conditional probability tables, we make an assumption of symmetry: the structures of subtrees of a same chance node are identical. Note that this assumption does not imply symmetric decision trees (as those obtained by unfolding an influence diagram [2]). For instance, the decision tree in Figure 1 is not symmetric but the condition holds: the three subtrees of node T have the same structure (the subtrees of nodes S are all leaves).

A *strategy* consists in setting a value to every decision variable conditionally to its past. The decision tree in Figure 1 includes 10 feasible strategies, among which for instance strategy $s = (D_1 = \mathit{sounding}, D_2^2 = \mathit{not\,drill}, D_2^3 = \mathit{drill}, D_2^4 = \mathit{drill})$ (note that node D_2^1 cannot be reached when $D_1 = \mathit{sounding}$). In our setting, a strategy can be associated to a compound lottery over the utilities, where the probabilities of the involved events are imprecise. For instance, strategy s corresponds to the compound lottery yielding $-10K$ if $T = \mathit{no}$, $190K$ (resp. $40K, -80K$) if $T = \mathit{open}$ or $T = \mathit{closed}$ and then $S = \mathit{soak}$ (resp. $\mathit{wet}, \mathit{dry}$). Comparing strategies amounts therefore to compare compound lotteries. Given a decision tree $\mathcal{T}$, the evaluation of a strategy (more precisely, of the corresponding compound lottery) according to the Hurwicz criterion depends on the set $\mathcal{P}_\mathcal{T}$ of possible probability distributions *on decision tree* $\mathcal{T}$ (i.e., the set of assignments of sharp probabilities to the tables coming with $\mathcal{T}$). This evaluation is a combinatorial problem in itself due to the combinatorial nature of $\mathcal{P}_\mathcal{T}$. We distinguish two cases:

Non-separable decision trees. We say that a decision tree $\mathcal{T}$ is *non-separable* when $\mathcal{P}_\mathcal{T}$ is a subset of the cartesian product of possible probability distributions at each chance node. In other words, the fact that the probabilities sum up to 1 at each chance node is not sufficient to ensure the global consistency of the probability distribution on the decision tree. This is the case for the decision tree of Figure 1. Consider for instance the following partial probability distribution on the tree: $P(S = \mathit{dry}|T = \mathit{no}) = 0.55$, $P(S = \mathit{dry}|T = \mathit{open}) = 0.33$, $P(S = \mathit{dry}|T = \mathit{closed}) = 0.12$, $P(T = \mathit{no}) = 0.20$, $P(T = \mathit{open}) = 0.35$, $P(T = \mathit{closed}) = 0.45$, $P(S = \mathit{dry}) = 0.22$. This partial probability distribution can be completed so that the probabilities sum up to 1 at each chance node, but is globally inconsistent since the total probability theorem does not hold: $P(S = \mathit{dry}|T = \mathit{no})P(T = \mathit{no}) + P(S = \mathit{dry}|T = \mathit{open})P(T = \mathit{open}) + P(S = \mathit{dry}|T = \mathit{closed})P(T = \mathit{closed}) = 0.2795 \neq 0.22 = P(S = \mathit{dry})$.

Separable decision trees. We say that a decision tree $\mathcal{T}$ is *separable* when $\mathcal{P}_\mathcal{T}$ is equal to the cartesian product of possible probability distributions at each chance node. In other words, the only requirement to ensure that a probability distribution is globally consistent is that the probabilities sum up to 1 at each chance node. This is for instance the case for the decision tree of Figure 2 as soon as random variables A, B, C, D, E are mutually independent.

Solving a decision tree means finding an optimal strategy according to a given decision criterion (here, Hurwicz and its particular cases). Note that the number of potential strategies grows exponentially with the size of the decision tree, i.e. the number of decision nodes (this number has indeed the same order of magnitude as the number of nodes in $\mathcal{T}$). Indeed, one easily shows that there are $\Theta(2^{\sqrt{n}})$ strategies in a complete binary decision tree $\mathcal{T}$, where n denotes the number of decision nodes. This prohibitive number of potential strategies makes it impossible to resort to an exhaustive enumeration of the strategies when the size of the decision tree increases. For this reason, it is necessary to develop an optimization algorithm to determine the optimal strategy. It is well-known that the rolling back method makes it possible to compute in linear time an optimal strategy w.r.t. EU. Indeed, such a strategy satisfies the optimality principle: any substrategy of an optimal strategy is itself optimal. Starting from the leaves, one computes recursively for each node the expected utility of an optimal substrategy: the optimal expected utility for a chance node equals the expectation of the optimal utilities of its successors; the optimal expected utility for a decision node equals the maximum expected utility of its successors. This is however more difficult to optimize the Hurwicz criterion in decision trees with imprecise probabilities. In Section 5, we will show that this is actually an NP-hard problem in non-separable decision trees. Before that, in the next section, we will study the case of separable decision trees.

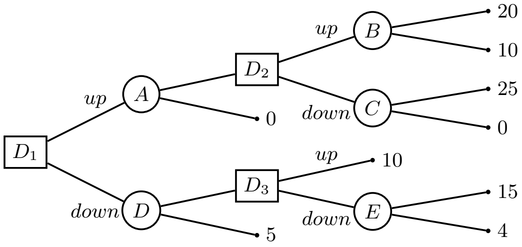

Fig. 2. A separable decision tree.

4 Optimizing the Hurwicz criterion in separable decision trees

When trying to optimize the Hurwicz criterion in a decision tree, it is important to note that the optimality principle does not hold. For instance, consider Figure 2 and assume complete ignorance about probabilities (i.e., all intervals of probabilities are $[0, 1]$). Let us set $\alpha = 0.5$ and perform backward induction on the decision tree with $u(x) = x$. In D_2, the decision maker prefers decision *up* to *down* (the Hurwicz criterion is equal to 15 for $D_2 = up$, compared to 12.5 for $D_2 = down$) and in D_3 he also prefers decision *up* to *down* (a sure utility of 10, compared to 9.5). In D_1, the decision maker has then the choice between a first lottery offering a minimum utility of 0 and a maximum utility of 20 if he decides *up*, and a second lottery offering a minimum of 5 and a maximum of 10 if he decides *down*. The best decision according to the Hurwicz criterion is *up* (10 compared to 7.5). The strategy returned by dynamic programming is therefore $(D_1 = up, D_2 = up)$ with a value of 10. Table 1 indicates the value of every strategy with respect to α. For $\alpha = 0.5$, strategy $(D_1 = up, D_2 = down)$ is optimal with a value of 12.5. In this case, one thus observes that the strategy returned by dynamic programming is suboptimal. For this reason, a decision maker using the Hurwicz criterion should adopt a *resolute choice* behavior [6], i.e. he initially chooses a strategy and never deviates from it later. We focus here on determining an optimal strategy from the root.

D_1	D_2	D_3	$\alpha = 0$	$\alpha = 0.5$	$\alpha = 1$
up	up	–	20	10	0
up	down	–	25	12.5	0
down	–	up	10	7.5	5
down	–	down	15	9.5	4

Table 1. Strategies and their evaluations.

Before showing how to compute an optimal strategy according to the Hurwicz criterion in a separable decision tree, we first show how to compute an optimal strategy according to Γ-maximin and Γ-maximax. It is well-known that the validity of the rolling back method on decision trees relies on the fulfillment of the independence axiom [5]. The *independence* axiom [10] states that the mixture of two lotteries f and g with a third one h should not reverse preferences (induced by the decision criterion used): if f is strictly preferred to g, then $\lambda f + (1-\lambda)h$ (i.e., the compound lottery that yields lottery f (resp. h) with probability λ (resp. $1-\lambda$)) should be strictly preferred to $\lambda g + (1-\lambda)h$. The following result states that the independence axiom holds for Γ-maximin and Γ-maximax *under a separability condition*:

Proposition 2 *Let f, g, h denote lotteries with sets $\mathcal{P}_f, \mathcal{P}_g, \mathcal{P}_h$ of possible probability distributions. If the set $\mathcal{P}_{\lambda f+(1-\lambda)h}$ (resp. $\mathcal{P}_{\lambda g+(1-\lambda)h}$) of possible probability distributions on the compound lottery $\lambda f+(1-\lambda)h$ is the cartesian product of $\mathcal{P}_f$ (resp. $\mathcal{P}_g$) and $\mathcal{P}_h$ (separability condition), then the following properties hold:*

$$\underline{E}(f) \geq \underline{E}(g) \Rightarrow \underline{E}(\lambda f + (1-\lambda)h) \geq \underline{E}(\lambda g + (1-\lambda)h)$$
$$\bar{E}(f) \geq \bar{E}(g) \Rightarrow \bar{E}(\lambda f + (1-\lambda)h) \geq \bar{E}(\lambda g + (1-\lambda)h)$$

Proof. We show that $\underline{E}(\lambda f+(1-\lambda)h) = \lambda\underline{E}(f)+(1-\lambda)\underline{E}(h)$ under the assumptions of the proposition. We have indeed $\underline{E}(\lambda f + (1-\lambda)h) = \min\{E(\lambda f + (1-\lambda)h, P) : P \in \mathcal{P}_{\lambda g+(1-\lambda)h}\}$. By linearity of expectation, it equals $\min\{\lambda E(f,P) + (1-\lambda)E(h,P) : P \in \mathcal{P}_{\lambda g+(1-\lambda)h}\}$. By separability assumption, it equals $\min\{\lambda\, E(f,P_f) + (1-\lambda)E(h,P_h) : P_f \in \mathcal{P}_f,\, P_h \in \mathcal{P}_h\} = \lambda \min\{E(f,P_f) : P_f \in \mathcal{P}_f\} + (1-\lambda)\min\{E(h,P_h) : P_h \in \mathcal{P}_h\}$. By definition of $\underline{E}(\cdot)$, it equals $\lambda\underline{E}(f) + (1-\lambda)\underline{E}(h)$. This implies the validity of the first property. The proof is similar for the second property. ■

In a separable decision tree, the separability condition of Proposition 2 holds at every chance node. For this reason, the rolling back method returns an optimal strategy when used with Γ-maximin or Γ-maximax in a separable decision tree. The computational complexity of this procedure is linear in the number of decision nodes.

Let us now explain our approach for computing an optimal strategy according to the Hurwicz criterion. We recall that the rolling back method does not work when operating directly with the Hurwicz criterion for $\alpha \neq 0, 1$. However, one can use the following simple property: if a substrategy is dominated by another one at the same node for both the Γ-maximin and Γ-maximax criteria (i.e., its value is smaller or equal for both criteria, and strictly smaller for at least one), then it cannot yield an optimal strategy for the Hurwicz criterion. The idea is to compute the set of non-dominated strategies (more precisely, one strategy for each non-dominated vector) by a bicriteria rolling back procedure from the leaves. At the root, one computes then the value of every non-dominated strategy according to the Hurwicz criterion, and one returns the best one. Due to space limitation, we only give here an example to provide an intuitive idea of how the procedure operates.

Example 2 *Let us come back to the decision tree of Figure 2 and assume again complete ignorance about probabilities,* $\alpha = 0.5$ *and* $u(x) = x$ *(for simplicity in the calculation). We describe here, for each node* X *in the subtree rooted at node* A*, how the set* $ND(X)$ *of non-dominated vectors (the first (resp. second) component represents the minimum (resp. maximum) expected utility of a feasible strategy) are inferred from the non-dominated vectors of its successors:*
- *at leaf 20 (resp. 10, etc.)* $ND(20) = \{(20, 20)\}$ *(resp.* $\{(10, 10)\}$*, etc.);*
- $ND(B) = \{(10, 20)\}$ *since combining* $(10, 10)$ *and* $(20, 20)$ *yields* $(10, 20)$*;*
- $ND(C) = \{(0, 25)\}$ *since combining* $(0, 0)$ *and* $(25, 25)$ *yields* $(0, 25)$*;*
- $ND(D_2) = \{(10, 20), (0, 25)\}$ *since both vectors are non-dominated;*
- $ND(A) = \{(0, 25)\}$ *since combining* $(10, 20)$ *(resp.* $(0, 25)$*) and* $(0, 0)$ *yields* $(0, 20)$ *(resp.* $(0, 25)$*), and* $(0, 25)$ *dominates* $(0, 20)$*;*

By proceeding similarly, one obtains $ND(D) = \{(4, 15), (5, 10)\}$*. At the root, one obtains finally* $ND(D_1) = \{(0, 25), (4, 15), (5, 10)\}$*. By evaluating every vector according to the Hurwicz criterion, one finds that* $(0, 25)$ *is an optimal vector (corresponding to optimal strategy* $(D_1 = up, D_2 = down)$*).*

The algorithm has been implemented in C++, and we have carried out numerical tests on a PC with a Pentium IV CPU 2.13Ghz processor and 3.5GB of RAM. Our tests were performed on complete binary decision trees of even depth. The depth of these decision trees varies from 4 to 14 (5 to 5461 decision nodes), with an alternation of decision nodes and chance nodes. Utilities are real numbers randomly drawn within interval $[1, 500]$. The imprecise probabilities were generated by randomly drawning a sharp probability distribution for each chance node, and then randomly generating an interval of probabilities around each probability. The numerical results are summarized in Table 2. Column "Imprecision" (resp. "Ignorance") details results obtained in the case of imprecise probabilities (resp. complete ignorance). Note that some tuning of the bicriteria rolling back method is possible in the case of complete ignorance, that considerably speeds up the procedure. Furthermore, the number of non-dominated vectors at each node is upper bounded by n in this case (where n denotes the number of decision nodes), and therefore the whole procedure performs in $O(n^2)$. For each depth, 500 instances were randomly generated and one indicates the average (Avg) and maximum (Max) computation times (in sec.), as well as the cardinality of the set of non-dominated vectors at the root. Symbol "–" appears when the memory size was not sufficient to execute the algorithm. One can observe that the smaller memory space requirements make it possible to solve larger instances (up to 16 millions of nodes) in the case of complete ignorance.

Algorithms		Imprecision		Ignorance	
Depth (nodes)		Avg	Max	Avg	Max
13 (16, 383)	card.	144	600	24	39
	time	0	0	0	0
15 (65, 535)	card.	940	3, 290	47	68
	time	2.35	44	0.02	1
17 (262, 143)	card.	7, 182	40, 930	90	115
	time	1, 189.97	3, 595	0.14	1
19 (1, 048, 575)	card.	–	–	174	216
	time	–	–	0.58	1
21 (4, 194, 303)	card.	–	–	348	570
	time	–	–	2.09	3
23 (16, 777, 215)	card.	–	–	714	1, 164
	time	–	–	8.31	9
25 (67, 108, 863)	card.	–	–	–	–
	time	–	–	–	–

Table 2. Numerical results.

5 Optimizing the Hurwicz criterion in non-separable decision trees

We now prove that the determination of an optimal strategy according to the Hurwicz criterion in a non-separable decision tree is an NP-hard problem, where the size of an instance is the number of involved decision nodes. Actually, we show a stronger result:

Proposition 3 *The determination of an optimal strategy according to the Γ-maximax criterion in a non-separable decision tree is an NP-hard problem.*

Proof. The proof relies on a polynomial reduction from problem 3-SAT, which can be stated as follows:

INSTANCE: a set X of boolean variables, a collection C of clauses on X such that $|c| = 3$ for every clause $c \in C$.

QUESTION: does there exist an assignment of truth values to the boolean variables of X that satisfies simultaneously all the clauses of C ?

Let $X = \{x_1, \ldots, x_n\}$ and $C = \{c_1, \ldots, c_m\}$. The polynomial generation of a decision tree from an instance of 3-SAT is performed as follows. One defines a decision node for each clause of C. Given c_i a clause in C, the corresponding decision node in the decision tree, also denoted by c_i, has three children (chance nodes), one for each literal in the clause. These chance nodes are denoted by the name of the corresponding literal. Every chance node x_i (resp. $\bar{x}_i$) has two children: a leaf of utility 1 with probability $p_i \in [0,1]$ (resp. $1-p_i$), and a leaf of utility 0 with probability $1-p_i \in [0,1]$ (resp. p_i). Finally, one adds a chance node A as root, predecessor of all decision nodes c_i, with probability $1/m$ on every branch. The obtained decision tree includes m decision nodes, $3m+1$ chance nodes and $6m$ leaves. Furthermore, n probability variables are involved. This guarantees the polynomiality of the reduction. For the sake of illustration, on Figure 3, we represent the decision tree obtained for the following instance of 3-SAT: $(x_1 \vee x_2 \vee x_3) \wedge (\overline{x_1} \vee x_3 \vee x_4) \wedge (\overline{x_2} \vee \overline{x_3} \vee \overline{x_4})$.

Note that, in this kind of decision trees, the Γ-maximax value of any strategy is upper bounded by 1. Furthermore, given an assignment of truth values that makes satisfiable the 3-SAT expression, one can construct a strategy whose Γ-maximax value is 1. There exists indeed at least a literal whose truth value is “true” for every clause c_i. Let us denote by k_i the index of such a literal in c_i. At every node c_i, one makes decision leading to literal whose index is k_i. By setting $p_{k_i} = 1$ (resp. 0) if it is a positive (resp. negative) literal, the expected utility of the corresponding strategy is 1. Conversely, given a strategy whose Γ-maximax value is 1, one can construct an assignment of truth values that makes satisfiable the 3-SAT expression. Indeed, at every decision node c_i the chosen decision necessarily leads to a chance node returning a utility of 1 with a probability set to 1. Let us denote by k_i the index of the chance node chosen at c_i. One obtains a partial assignment by setting x_{k_i} to “true” (resp. “false”) if $p_{k_i} = 1$ (resp. 0). Any completion of this partial assignment makes satisfiable the 3-SAT expression. This concludes the proof. ■

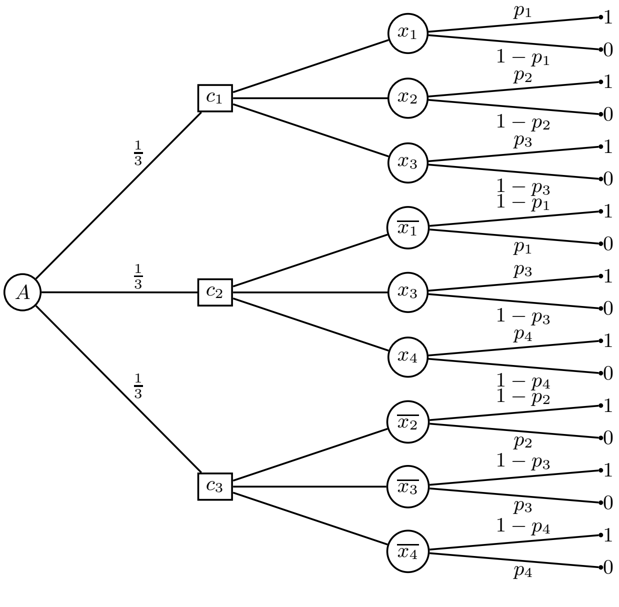

Fig. 3. An example of reduction.

6 Conclusion

In this paper, we have proposed an operational procedure to determine an optimal strategy according to the Hurwicz criterion in a separable decision tree. Furthermore, we have proved that the problem becomes NP-hard in non-separable decision trees. For future research, it would be interesting to propose an algorithm for optimizing the Hurwicz criterion in a non-separable decision tree. In this purpose, a branch and bound is worth investigating. An upper bound easily computable would consist for instance in computing an upper bound of the value of a $\varGamma$-maximin strategy by determining the maximum expected utility for a feasible sharp probability distribution (i.e., consistent with the intervals of probabilities), and an upper bound of a $\varGamma$-maximax strategy by relaxing the non-separability constraints (and therefore using the procedure for $\varGamma$-maximax detailled in Section 4). Combining both upper bounds with α and $1-\alpha$ would provide an upper bound for the Hurwicz criterion.

References

1. C. P. de Campos D. Kikuti, F. G. Cozman. Partially ordered preferences in decision trees: computing strategies with imprecision in probabilities. In *IJCAI Workshop on Advances in Preference Handling*, 2005.
2. R. Howard and J. Matheson. *Influence Diagrams*. Menlo Park CA: Strategic Decisions Group, 1984.
3. J.-Y. Jaffray and M. Jeleva. Information processing under imprecise risk with the hurwicz criterion. In *5th International Symposium on Imprecise Probability: Theories and Applications*, pages 233–242, 2007.
4. A. Kasperski. *Discrete Optimization with Interval Data: Minmax Regret and Fuzzy Approach*. Studies in Fuzziness and Soft Computing. Springer, 2008.
5. I.H. LaValle and K.R. Wapman. Rolling back decision trees requires the independence axiom. *Management Science*, 32(3):382–385, 1986.
6. E.F. McClennen. *Rationality and Dynamic choice: Foundational Explorations*. Cambridge University Press, 1990.
7. M.L. Puterman. *Markov Decision Processes - Discrete Stochastic Dynamic Programming*. Wiley & Sons, 1994.
8. H. Raiffa. *Decision Analysis: Introductory Lectures on Choices under Uncertainty*. Addison-Wesley, 1968.
9. R. Shachter. Evaluating influence diagrams. *Operations Research*, 34:871–882, 1986.
10. J. von Neuman and O. Morgenstern. *Theory of games and economic behaviour*. Princeton University Press, 1947.
11. P. Walley. *Statistical reasoning with imprecise probabilities*, volume 91 of *Monographs on statistics and applied probability*. Chapman and Hall, 1991.
12. K. Weichselberger. The theory of interval-probability as a unifying concept for uncertainty. In *1st International Symposium on Imprecise Probability: Theories and Applications*, pages 387–396, 1999.