

Cost-effectiveness analysis with sequential decisions

M. Arias and F. J. Díez
Dept. Inteligencia Artificial, UNED.
Juan del Rosal, 16
28040 Madrid, Spain
{marias, fjdiez}@dia.uned.es

June 9, 2011

Abstract

In this paper we present a new method for performing cost-effectiveness analyses of problems that involve multiple decisions and probabilistic outcomes. This issue has been ignored by most of the literature on medical decision making, and the few proposed solutions are either wrong or unfeasible, except for very small problems. The method proposed in this paper consists of building a decision tree with several decision nodes and evaluating the tree with a modified roll back algorithm that operates with partitions of intervals.

1 Introduction

In medicine, it is often necessary to determine whether the benefit of an intervention outweighs its cost. By “intervention” we mean either a single action, such as applying a therapy, or a whole strategy, such as “Do the test T ; if it is positive, apply therapy drug D every 8 hours; if it is negative, repeat the test after 24 hours and then...”.

Cost-effectiveness analysis (CEA) is one method of assessing whether the health benefits of an intervention outweigh the economic cost [2, 4]. Cost-utility analysis is a particular form of CEA in which the effectiveness is measured in quality-adjusted life years (QALYs) [19, 20]. In this context, the *net monetary benefit* [?] of an intervention I_i is

$$NMB_{I_i}(\lambda) = \lambda \cdot e_i - c_i, \quad (1)$$

where e_i is the cost and c_i is the effectiveness. The parameter λ is used to convert the effectiveness into a monetary scale. It takes values on the set of positive real numbers, i.e., on the interval $(0, \infty)$. This parameter is measured in effectiveness units divided by

cost units, e.g., in dollars per death avoided or euros per QALY. It is sometimes called *willingness to pay*, *cost-effectiveness threshold* or *ceiling ratio*, because it indicates how much money a decision maker is willing to pay to obtain a certain “amount” of health benefit.

CEA assumes that the cost and effectiveness of each intervention, which are objective magnitudes, are known. Therefore, if we know the value of λ , then we are able to compute the NMB of each intervention, a unicriterion decision problem results, and we can choose the set of interventions that maximize the NMB. The problem is that λ depends on each decision maker; therefore the selection of interventions must be expressed as a function of λ .

When the consequences of an intervention are not deterministic, it is necessary to apply a model that considers the probability of each outcome. The most usual tool for modeling decision problems with uncertainty is a decision tree.

The standard way of evaluating a unicriterion decision tree is the roll back algorithm, which evaluates the tree from its leaves to the root node. The utility associated to a chance node is the average of the utilities of its branches, and the utility of a decision node is the maximum of their utilities. This roll-back algorithm can be adapted to perform CEA on a decision tree: each leaf node represents two separate utilities, the cost and the effectiveness; at each chance, the cost and effectiveness are computed separately by averaging the values of its branches, and the propagation continues backwards until we reach a decision node, in which a standard CEA is performed. However, the result of this analysis cannot be propagated backwards as a pair of numbers, because in general the cost and the effectiveness depend on the value of λ , which is assumed to be unknown. Given that the modified roll-back algorithm stops at a decision node, the tree can contain only one decision node, and that must be its root. This is a strong limitation, because many realistic situations involve more than one decisions, as we can see in the following problem, which we will use as a running example throughout the entire paper.

Example 1 *For a disease whose prevalence is 0.14, there are two possible therapies. The effectiveness of each therapy depends on whether the disease is present, as shown in Table 1.*

There is a test with a sensitivity of 90% and a specificity of 93%, and a cost of 150 €. Is the test cost-effective?

In this hypothetical example, the second therapy is more effective for those who suffer from the disease (effectiveness = 6.5 vs. 4.0 for the first therapy). Therefore, when we are sure that a patient is suffering from that disease, we should apply the second therapy, but only if we can afford the extra cost, 50,000 €. On the contrary, if the patient is not suffering from the disease, we should not apply any therapy, not only because they have a cost, but also because both have secondary effects, especially the second one, which leads to an effectiveness of 9.3, while the effectiveness of the first one is 9.9.

Therefore, it may be desirable to know with certainty whether the disease is present or not. The test may offer useful information, but it has a cost and the value of the information it provides is limited. First, because the test is not completely reliable, and

Therapy	Cost	Effectiveness	Effectiveness
		+disease	−disease
No therapy	0€	1.2	10,0
Therapy 1	20,000€	4.0	9.9
Therapy 2	70,000€	6.5	9.3

Table 1: Cost and effectiveness of each intervention.

therefore there are still some uncertainty even after applying it. And second, because if λ is very small, we can not afford to apply any therapy, not even the first one, let alone the second one. In this case, the test would be useless, even if it were very cheap. Therefore, the net benefit of the test depends on what therapy will be applied later, which in turn depends on λ .

As mentioned above, if we knew the value of λ , we might transform the above example into a unicriterion problem and solve it using standard techniques. For example, the decision tree in Figure 1 shows that the test is cost-effective when $\lambda = 30,000$ €/QALY. A similar analysis would arrive at the opposite result for $\lambda = 10,000$ €/QALY.

If we wish to find the intervals in which the test is cost-effective and those for which it is not, we should repeat this analysis for each single value of λ , however this is clearly unfeasible, because λ can assume infinite values. We might also try to apply standard CEA techniques, but the problem is they can only be applied to decision trees having only one decision node, that must be its root. Another method is to build a decision tree such that each branch of the root node represents an entire strategy, i.e., an intervention that includes both decisions. This method is suitable for solving the above example, but is not appropriate for larger problems, as we will show in Section 5. We will also show that the most commonly used commercial program for decision trees often provides a wrong result for cost-effectiveness decision trees with several decisions. For this reason, the main purpose of the current paper is to propose a scalable CEA method for problems that involve several decisions and probabilistic outcomes, that gives the right answer in all cases but, at the same time, scales up reasonably well. The basic idea of our algorithm is to group the λ 's into intervals having the same cost and effectiveness.

The rest of this paper is structured as follows: in Section 2 we review the standard methods for unicriterion decision analysis and for CEA. In Section 3 we present our method for CEA with multiple decisions; Section 4 explains in detail how to solve the problem stated in the Example 1. In Section 5 we compare our method with those proposed previously and we conclude in Section 6.

2 Overview of standard methods

In this section we review the standard techniques for decision analysis. First, we present unicriterion decision trees, which can be used to find the most cost-effective interventions when λ is known. Second, we review the fundamentals of CEA for the case in which the cost and effectiveness of each intervention are known explicitly. Then, we combine both methods to show how to perform CEA in decision trees without embedded decision nodes, i.e., trees whose only decision node is at the root.

2.1 Unicriterion decision trees

A decision tree is a model for decision analysis of problems involving probabilistic outcomes. It has three types of nodes: chance, decision, and utility. Utility nodes represent the decision maker's preferences. All the leaves must be utility nodes and, conversely, all utility nodes must be leaves. A path from the root to a utility node is called a *scenario*; the utility node represents the reward obtained by the decision maker in that scenario. A decision node represents a choice that the decisor can choose; each branch of a decision variable represents one of the values that the variables can assume and has an associated probability. Chance nodes represent events that are not under the direct control of the decision maker; each branch of a chance variable represents one of the values that the variable can take on, and has an associated probability. For each chance node, the sum of the probabilities of its branches must be one.

Decision trees can be evaluated by applying Algorithm 1, which operates recursively from the leaves to the root.

Algorithm 1: Roll-back algorithm for unicriterion decision trees

Input: A decision tree

Result: The expected utility and
a policy for each decision node

```

1 foreach node  $n$  do
2   if  $n$  is a chance node then
3      $u_n = \sum_i p_i \cdot u_i$ , where  $p_i$  is the probability of the  $i$ -th branch
       and  $u_i$  is its utility
4   if  $n$  is a decision node then
5      $u_n = \max u_i$ 

```

2.2 Fundamentals of CEA

Cost-effectiveness analysis (CEA) consists of finding an intervention that maximizes the net benefit for each value of λ ; in practice, it consists in finding the intervals for which an

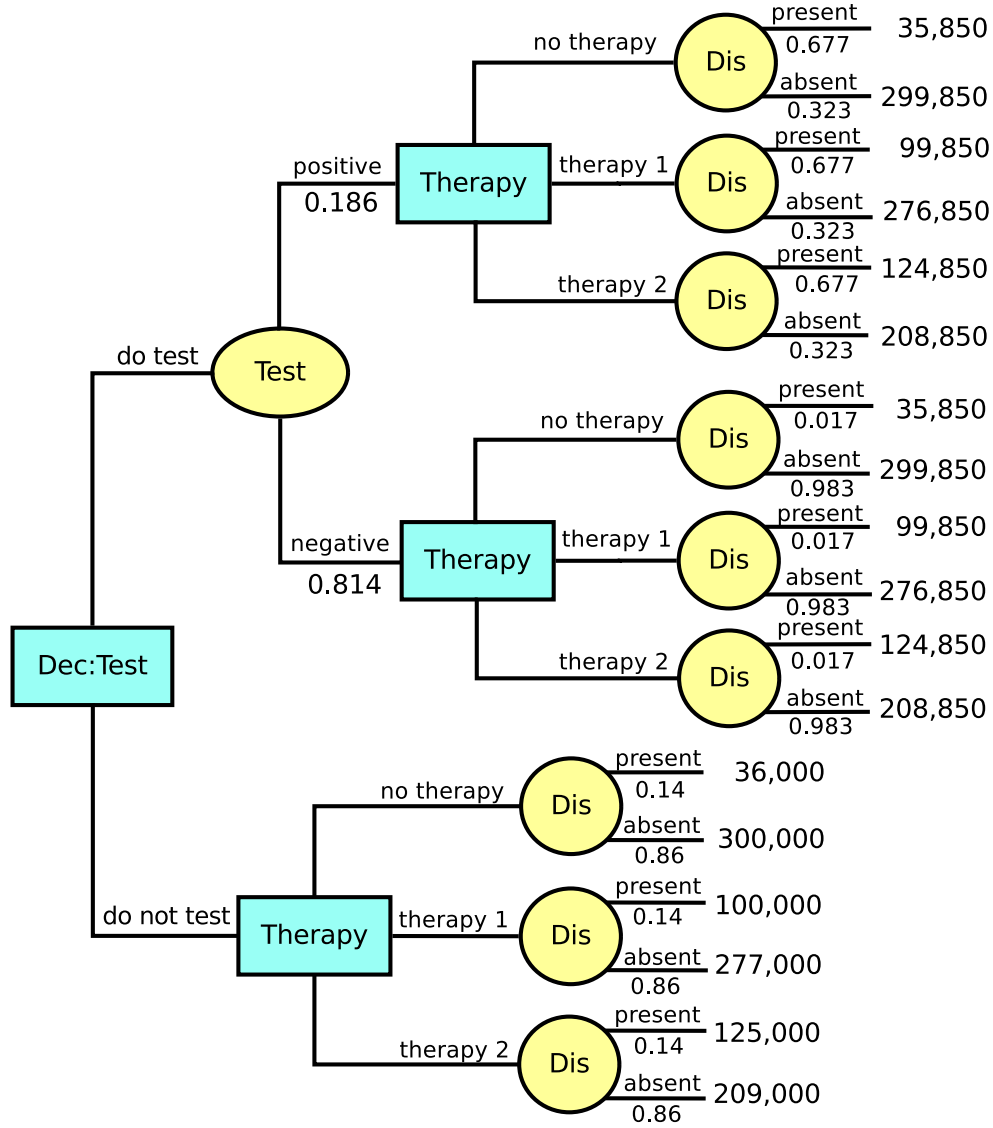


Figure 1: Unicriterion decision tree for the Example 1. We have chosen $\lambda=30,000\text{€}/\text{QALY}$ to turn it into a unicriterion problem.

intervention is more beneficial than the others. We formalize this idea by introducing the concept of *cost-effectiveness partition*, CEP. Then we present two algorithms for finding a CEP when the cost and effectiveness of each intervention are known explicitly. The key idea of those algorithms is to eliminate the interventions dominated by others, i.e., the interventions such that, for every value of λ , there is a more beneficial alternative.

2.2.1 Net benefit, dominance, and ICERs

Each intervention I_i has a cost, c_i , and an effectiveness, e_i . We assume that there are not two interventions having the same cost and the same effectiveness. If that unlikely case occurred, then both interventions would be indistinguishable from the point of view of cost-effectiveness, and it would be indifferent to choose any of them.

Definition 2 (Simple dominance) *An intervention I_i dominates another intervention I_j if $c_i \leq c_j$ and $e_i \geq e_j$.¹*

Proposition 3 *If I_i dominates I_j , then $NMB_{I_i}(\lambda) > NMB_{I_j}(\lambda)$ for all λ .*

The proof of all the propositions are in the appendix.

This proposition implies that when I_i dominates I_j then I_j is never the optimal intervention, whatever the value of λ . This proposition implies that I_k is never the optimal intervention: depending on the value of λ , either I_i or I_j is more beneficial than I_k , or both. This is the reason for discarding some interventions in Algorithm 2, line 1.

Definition 4 (ICER) *Given two interventions such that $e_i < e_j$ and $c_i \leq c_j$, the incremental cost-effectiveness ratio is*

$$ICER(I_i, I_j) = \frac{c_j - c_i}{e_j - e_i}. \quad (2)$$

Proposition 5 *Given two interventions I_i and I_j ,*

$$\lambda < ICER(I_i, I_j) \implies NMB_{I_i}(\lambda) > NMB_{I_j}(\lambda) \quad (3)$$

$$\lambda = ICER(I_i, I_j) \implies NMB_{I_i}(\lambda) = NMB_{I_j}(\lambda) \quad (4)$$

$$\lambda > ICER(I_i, I_j) \implies NMB_{I_i}(\lambda) < NMB_{I_j}(\lambda). \quad (5)$$

The importance of this proposition is that it allows us to determine, for each value of λ , which is the most beneficial intervention by comparing the ICERs. Put another way, the ICERs partition the interval $(0, +\infty)$ into a set of subintervals, each associated with an intervention I_i , such that when λ falls in that subinterval the optimal intervention is I_i . This property is the basis of the definition of cost-effectiveness partition (CEP), that we will present in the next section. But first we introduce the concept of extended dominance.

¹It is usual to define simple dominance by the condition “ $c_i < c_j$ and $e_i > e_j$ ”, which is more restrictive than ours, as it excludes the cases “ $c_i < c_j$ and $e_i = e_j$ ” and “ $c_i = c_j$ and $e_i > e_j$ ”. (We have assumed that the case “ $c_i = c_j$ and $e_i = e_j$ ” never occurs.) We will see in Section 4, when analyzing the Example 1, that these two cases can occur in practice, and the usual definition, more restrictive, would fail to detect in some cases that one intervention clearly dominates another one—see Figure 11(a).

Definition 6 (Extended dominance) A pair of interventions, I_i and I_j , dominate another intervention I_k if the following conditions hold:

- $c_i < c_k < c_j$,
- $e_i < e_k < e_j$, and
- $ICER(I_i, I_k) > ICER(I_i, I_j)$.

Proposition 7 If I_i and I_j dominate I_k , then $\max(NMB_{I_i}(\lambda), NMB_{I_j}(\lambda)) > NMB_{I_k}(\lambda)$ for all λ .

This proposition implies that I_k is never the optimal intervention: depending on the value of λ , either I_i or I_j is more beneficial than I_k , or both. This is the reason for discarding some interventions in Algorithm 2, line 2.

2.2.2 Cost-effectiveness partition (CEP)

Definition 8 A cost-effectiveness partition (CEP) of n intervals is a tuple $Q = (\Theta_Q, C_Q, E_Q, I_Q)$, where:

- $\Theta_Q = \{\theta_1, \dots, \theta_{n-1}\}$ is a set of $n - 1$ values (thresholds), such that $0 < \theta_1 < \dots < \theta_{n-1}$,
- $C_Q = \{c_0, \dots, c_{n-1}\}$ is a set of n values (costs),
- $E_Q = \{e_0, \dots, e_{n-1}\}$ is a set of n effectiveness values, and
- $I_Q = \{I_0, \dots, I_{n-1}\}$ is a set of n interventions..

For the sake of simplifying the exposition, we define $\theta_0 = 0$ and $\theta_n = +\infty$ for every CEP.

Alternatively, a CEP can be denoted by a set of n 5-tuples of the form (interval, cost, effectiveness, intervention),

$$Q = \{((0, \theta_1), c_0, e_0, I_0), \\ ((\theta_1, \theta_2), c_1, e_1, I_1), \\ \dots, \\ ((\theta_{n-1}, +\infty), c_{n-1}, e_{n-1}, I_{n-1})\},$$

meaning that when λ is in the interval (θ_i, θ_{i+1}) the most beneficial intervention is I_i , which has a cost c_i and an effectiveness c_i . When $\lambda = \theta_{i+1}$, there is a tie between I_i and I_{i+1} .

Put formally, we may define the function *index*:

$$index_Q(\lambda) = \min\{i | \lambda < \theta_{i+1}\} \quad (6)$$




Figure 2: Cost-effectiveness partition (CEP).

and the following three functions:

$$cost_Q(\lambda) = c_i \quad (7)$$

$$eff_Q(\lambda) = e_i \quad (8)$$

$$interv_Q(\lambda) = I_i, \quad (9)$$

where $i = index_Q(\lambda)$. The net monetary benefit for a particular value of λ is

$$NMB_Q(\lambda) = \lambda \cdot eff_Q(\lambda) - cost_Q(\lambda) \quad (10)$$

It is easy to see that if two values of λ lie in the same subinterval of a partition, then they have the same cost and the same effectiveness.

2.2.3 Deterministic CEA

When we have a set of interventions such that the cost and effectiveness of each one are known with certainty, we can perform a deterministic cost-effectiveness analysis, which returns a CEP that indicates the optimal intervention for each interval of the possible values of λ . There are at least two algorithms for this operation: Algorithm 2 is an adaptation of the standard method described in the literature ([18]); Algorithm 3 is introduced as a contribution of this paper because it is more efficient than the former and because it does not need to care about the distinction between simple dominance and extended dominance.

Both algorithms take as an input a set of m interventions $\{I_1, \dots, I_m\}$, such that the cost of I_i is c_i and its effectiveness is e_i , and return the same CEP. We can better understand the difference between them by tracing their performance on the set of interventions displayed in Figure 3. Algorithm 2 would first check for simple dominance by analyzing each pair of interventions; it would rule out I_1 and I_2 because they are dominated by I_3 , and then I_4 because it is dominated by I_5 (and by I_7). Then it would check for extended dominance by analyzing each tern of interventions; it would rule out I_6 because it is dominated by I_3 and I_7 (and by I_5 and I_7).

In contrast, Algorithm 3 would proceed as follows. First, it selects the intervention having the lowest cost; as there is a tie between I_1 and I_3 , it selects I_3 because it has a higher effectiveness, which implies that it dominates I_1 . Therefore, $\sigma(0) = 3$, which means that the first intervention in the output set I_Q is $I_{\sigma(0)} = I_3$. The set R_0 , which contains the indices of the interventions that remain after selecting the first one, is $\{4, 5, 6, 7, 8\}$. We have discarded the interventions dominated by $I_{\sigma(0)}$, i.e., I_3 .

Algorithm 2: Deterministic CEA (standard)

Input: a set $I = \{I_1, \dots, I_m\}$ of interventions.

Result: A CEP $Q = (\Theta_Q, C_Q, E_Q, I_Q)$, with $I_Q = \{I_{\sigma(0)}, \dots, I_{\sigma(n-1)}\}$.

- 1 Eliminate the interventions dominated by another interventions
(simple dominance)
 - 2 Eliminate interventions dominated by a pair of other interventions
(extended dominance)
 - 3 $I_Q = \{\text{remaining interventions in increasing-cost order}\}$
 - 4 $\Theta_i = ICER(I_{i-1}, I_i)$
 - 5 $c_i = cost(I_i)$
 - 6 $e_i = eff(I_i)$
-

Algorithm 3: Deterministic CEA (new)

Input: a set $I = \{I_1, \dots, I_m\}$ of interventions.

Result: A CEP $Q = (\Theta_Q, C_Q, E_Q, I_Q)$, with $I_Q = \{I_{\sigma(0)}, \dots, I_{\sigma(n-1)}\}$.

- 1 $\sigma(0) = \arg \min_i cost(I_i)$
 - 2 $R_0 = \{i \mid I_i \in I \wedge eff(I_i) > eff(I_{\sigma(0)})\}$
 - 3 $i := 1;$
 - 4 **while** $R_{i-1} \neq \emptyset$ **do**
 - 5 $\sigma(i) := \arg \min_{j \in R_i} ICER(I_{\sigma(i-1)}, I_j)$
 - 6 $\theta_i := \min_{j \in R_i} ICER(I_{\sigma(i-1)}, I_j)$
 - 7 $c_i := cost(I_{\sigma(i)})$
 - 8 $e_i := eff(I_{\sigma(i)})$
 - 9 $R_i := \{j \mid j \in R_{i-1} \wedge eff(I_j) > eff(I_{\sigma(i)})\}$
 - 10 $i := i + 1$
-




Figure 3: Comparison between interventions. Dominated interventions (I_1, I_2, I_4 and I_6) are shown in red and not dominated (I_3, I_5, I_7 and I_8) in blue.

Then the algorithm enters the while loop with $i = 1$. First it computes the ICER of each remaining intervention with respect to $I_{\sigma(0)}$, i.e., $ICER(I_3, I_j)$. The value of $ICER(I_3, I_j)$ is the slope of the line that connects I_3 and I_j in Figure 3. When looking for the minimum, there is a tie between $ICER(I_3, I_5)$ and $ICER(I_3, I_7)$; therefore, it selects I_7 , because it has a higher effectiveness: $\sigma(1) = 7$. The first threshold is $\theta_1 = ICER(I_3, I_7)$. We have discarded I_4 and I_6 .²

Then the algorithm performs another iteration of the while loop, now with $i = 2$ and $R_1 = \{8\}$. Therefore, $\sigma(2) = 8$ and $\theta_2 = ICER(I_7, I_8)$. In the third iteration, $i = 3$ and $R_2 = \emptyset$, which makes the algorithm terminate, returning a CEP of three intervals: $I = \{I_{\sigma(0)}, I_{\sigma(1)}, I_{\sigma(2)}\} = \{I_3, I_7, I_8\}$ and $\Theta = \{ICER(I_3, I_7), ICER(I_7, I_8)\}$.³

It is easy to prove that this algorithm always terminates, because $R_i \subsetneq R_{i-1}$ and R_0 is finite.

2.3 CEA with uni-decision trees

In some cases, we do not know the cost and effectiveness of each intervention, but we do know that each one may lead to different outcomes with different probabilities and this

²Please note that I_4 is dominated by I_7 , while I_6 is jointly dominated by I_3 and I_7 . However, this algorithm does not make any distinction between both types of dominance.

³In this example, $ICER(I_3, I_5) = ICER(I_5, I_7) = ICER(I_3, I_7)$, which implies that I_5 is never more beneficial than I_3 or I_7 ; but when $\lambda = ICER(I_3, I_5) = ICER(I_3, I_7)$, the three interventions have the same NMB. For this reason, it might be not to discard I_5 , which would have even ethical implications [1]; in this case, we should modify the algorithm so that when it finds a tie in the ICERs, it selects first the intervention having the lower cost and lower effectiveness. In this example, the modified algorithm would return a CEP of four interventions, $I = \{I_3, I_5, I_7, I_8\}$ and three thresholds, $\Theta = \{ICER(I_3, I_5), ICER(I_5, I_7), ICER(I_7, I_8)\}$, but the first and the second thresholds would be the same. In any case, this issue is irrelevant in practice, as the probability of having an absolute tie is virtually null.

may result in other outcomes, each having a known cost and effectiveness. In this case, the standard analysis method consists of building a decision tree such that each node, instead of representing a single utility, represents the cost and effectiveness of the corresponding scenario. The evaluation of this kind of tree is very similar to the unicriterion case: we proceed from the leaves to root, averaging at each chance node, with the only difference that the cost and effectiveness are computed separately throughout the evaluation process; when evaluating the root node, we have a cost and effectiveness for each of its branches, and we can then perform a deterministic CEA, as explained in Section 2.2.3, which returns a CEP.

A limitation of this method is that the tree can contain at most one decision node, that must be the root, because if the tree contained an embedded decision node, its evaluation would not return a cost-effectiveness pair, but a CEP, which can not be propagated backwards. For this reason, we present in the next section a method for combining CEPs, which will permit to evaluate decision trees with embedded nodes.

3 New algorithm: multi-decision CEA

3.1 Combination of cost-effectiveness partitions

When evaluating a unicriterion decision tree, chance nodes are processed by taking the average of the utilities of its branches, and decision nodes by taking the maximum. In this section we generalize the average and maximization operations from single utilities (unicriterion analysis) to cost-effectiveness partitions.

3.1.1 Weighted average of partitions

Definition 9 (Weighted average) *Given a set of m CEPs $\{Q_1, \dots, Q_m\}$, a chance variable X whose domain is $\{x_1, \dots, x_m\}$, and a probability distribution for X , $P(x_j)$, we say that a CEP Q is a weighted average of the CEPs if*

$$\forall \lambda, \quad \text{cost}_Q(\lambda) = \sum_{j=1}^m P(x_j) \cdot \text{cost}_{Q_j}(\lambda) \quad (11)$$

and

$$\forall \lambda, \quad \text{eff}_Q(\lambda) = \sum_{j=1}^m P(x_j) \cdot \text{eff}_{Q_j}(\lambda) . \quad (12)$$

Due to Equation 10, a straightforward consequence of this definition is the following:

$$\forall \lambda, \quad \text{NMB}_Q(\lambda) = \sum_{j=1}^m P(x_j) \cdot \text{NMB}_{Q_j}(\lambda) . \quad (13)$$

These three equalities mean that for a every value of λ , the cost, effectiveness, and NMB of the weighted average partition Q are the same as if we had performed a weighted average

for the values of cost and effectiveness of the Q_j 's. Clearly, Q cannot be computed by performing a weighted average for each single value of λ , because this parameter can take on infinite values, but we can compute Q efficiently by applying the Algorithm 4.

The intervention composed at the fifth line of the algorithm means: “if the chance variable X takes on the value x_j , then follow the policy indicated by the corresponding branch of the tree (for the values of λ corresponding to the interval (θ_i, θ_{i+1}))”.

Algorithm 4: Weighted average of CEPs

Input: A set of m CEPs $\{Q_1, \dots, Q_m\}$, with $Q_j = (\Theta_j, C_j, E_j, I_j)$
a chance variable X whose domain is $\{x_1, \dots, x_m\}$, and
a probability distribution for X , $P(x_j)$.

Result: A new CEP $Q = (\Theta, C, E, I)$.

```

1  $\Theta = \bigcup_{j=1}^n \Theta_j, (n = \text{card}(\Theta))$ 
2 for  $i \leftarrow 1$  to  $n$  do
3    $c_i = \sum_{j=1}^m P(x_j) \cdot \text{cost}_{Q_j}(\theta_i)$ 
4    $e_i = \sum_{j=1}^m P(x_j) \cdot \text{eff}_{Q_j}(\theta_i)$ 
5    $I_i = \text{“If } X = x_1, \text{ then } \text{interv}_{Q_1}(\theta_i); \text{ if } X = x_2, \text{ then } \text{interv}_{Q_2}(\theta_i); \text{ etc.”}$ 

```

3.1.2 Optimal partition

Definition 10 (Optimal partition) *Given a set of m CEPs $\{Q_1, \dots, Q_m\}$ and a decision D whose domain is $\{d_1, \dots, d_m\}$, a CEP Q is optimal if*

$$\forall \lambda, \exists j, \text{NMB}_{\text{interv}_{Q_j}(\lambda)}(\lambda) = \max_{j'} \text{NMB}_{\text{interv}_{Q_{j'}}(\lambda)}(\lambda), \quad (14)$$

$$\text{interv}_Q(\lambda) = \text{“choose option } d_j; \text{ then apply } \text{interv}_{Q_j}(\lambda)\text{”}, \quad (15)$$

$$\text{cost}_Q(\lambda) = \text{cost}_{Q_j}(\lambda), \quad (16)$$

$$\text{eff}_Q(\lambda) = \text{eff}_{Q_j}(\lambda). \quad (17)$$

The interpretation of this definition is as follows: for each value d_j (a possible choice) of the decision D there is a CEP Q_j and for each value of λ there is an intervention $\text{interv}_{Q_j}(\lambda)$ in Q_j which is optimal for d_i . Equation 14 means that we select j such that $\text{interv}_{Q_j}(\lambda)$ is the intervention that attains the highest NMB for that particular value of λ . In fact, there is only one possible choice for j , except in the case of a tie between several interventions. The optimal intervention for decision D , $\text{interv}_Q(\lambda)$, is to choose first the option d_j and then apply the intervention $\text{interv}_{Q_j}(\lambda)$, which maximizes the expected NMB for that value of λ (cf. Eq. 15). The cost and effectiveness associated with intervention $\text{interv}_Q(\lambda)$ are the same as those in the CEP chosen, Q_j . In practice, Q can not be computed by performing a maximization for each single value of λ , but we can compute Q efficiently by applying the Algorithm 5.

The key property of this definition is Equation 14, which states that for every λ the NMB of the optimal partition, Q , is the same as if we had performed a unicriterion maximization of the NMB for each single value of λ .

The optimal CEP for a set of Q_j 's and a decision D can be obtained by applying Algorithm 5, which collects all the thresholds of the Q_j 's and performs a deterministic CEA (cf. Sec. 2.2.3) on each interval. Finally, it fuses some intervals by eliminating the unnecessary thresholds. In Section 4 we show with an example how this algorithm operates and why it is sometimes necessary to fuse intervals.

Algorithm 5: Optimal CEP.

Input: A set of m CEPs $\{Q_1, \dots, Q_m\}$, with $Q_j = (\Theta_j, C_j, E_j, I_j)$ and a decision node

Result: A new CEP $Q = (\Theta, C, E, I)$.

- 1 $\Theta = \cup_i \theta_i$
 - 2 **for** $i \leftarrow 1$ **to** n **do**
 - 3 └ perform a deterministic CEA analysis of interval i with Algorithm 2 or 3
 - 4 Fuse contiguous intervals having the same intervention, the same cost, and the same effectiveness
-

3.2 CEA with multi-decision trees

Finally, we can present the algorithm for performing cost-effectiveness analyses on decision trees that contain several decisions, Algorithm 6, which is based on the above methods for combining CEPs. It is very similar to the method for evaluating unicriterion decision trees (Algorithm 1), with the only difference that instead of operating with single utility values, it operates with CEPs: for each chance node, it computes the weighted average of the partitions (Algorithm 4) and for each decision node, it computes the optimal partition (Algorithm 5).

This algorithm computes a CEP for every node in the tree, but we are interested only in the nodes for decision nodes, because they determine the optimal strategy for each value of λ .

4 Example: CEA of a test

To better explain our algorithm, we apply it to the Example 1, presented in the introduction. First, we build the decision tree shown in Figure 4. It is very similar to that in Figure 1: the structure and the probabilities for the branches of the chance nodes are exactly the same. The first difference is that instead of assigning a single utility value to each leaf node, we assign a CEP with only one interval, $(0, +\infty)$, the same for all branches. The effectiveness of each partition is taken from Table 1. The costs are also taken from Table 1, but in all the leaves of the branch “do test” we have added the cost of the test.

Algorithm 6: Roll-back algorithm for CEA in decision trees with several decisions

Input: A decision tree having a cost and an effectiveness value at each leave.

Result: A CEP $Q = (\Theta_Q, C_Q, E_Q, I_Q)$ for each decision node.

```
1 foreach node  $n$  do
2   if  $n$  is a chance node then
3     Evaluate it using Algorithm 4
4   if  $n$  is a decision node then
5     Evaluate it using Algorithm 5
```

The interventions are null for these nodes, because no decision node has been evaluated yet.

The evaluation begins by performing a weighted average of partitions (Algorithm 4) for each of the chance nodes, numbered from 6 to 14. As the leaf nodes do not contain thresholds—each input partition consisted of only one interval—and the weighted average operation does not introduce new thresholds, the CEPs obtained for those nodes do not contain any thresholds either—see Figure 5.

Now we evaluate the decision nodes 3 to 5 using Algorithm 5 (optimal CEP). For node 3, the input CEPs are those shown in Figure 5. As they do not contain any threshold, there is only one interval on which to perform a deterministic CEA: $(0, +\infty)$. The output CEP contains two thresholds, corresponding to the slopes of the two lines in Figure 6, and three intervals, as shown in Figure 9.

The analysis for node 4 is similar, but in this case it does not produce new thresholds, because *no therapy* dominates both *therapy 1* and *therapy 2*, as we can observe in Figure 7. The output CEP contains only one interval—see again Figure 9.

The analysis for node 5 produces only one threshold, because *therapy 2* is dominated by *therapy 1*, as shown in Figure 8. The output CEP contains two intervals—see Figure 9.

The chance node *Test* (number 2) is evaluated as a weighted average of partitions (Algorithm 4). The CEP for the *positive* branch contains two thresholds, 10,730 and 33,384 €/QALY, as displayed in Figure 9; the CEP for the *negative* branch contains none. Therefore, there is a total of two thresholds, i.e., three intervals. For each one of them, the algorithm computes the weighted average of the costs and that of the effectiveness values. The output CEP is shown in Figure 10.

Finally, the root node, *Dec:Test*, is evaluated with Algorithm 5 (optimal CEP). The branch *do test* contributes two thresholds, and the branch *do not test* contributes one, which makes a total of three: $\{10,739, 33,384, 65,359\}$. Thus, there are four intervals on which to perform a deterministic CEA, as shown in Figure 11.

In the first interval, $(0, 10,739)$, the option *do test* has the same effectiveness as *do not test*—see Figure 11(a)—because when $\lambda < 10,739$ no therapy will be applied, whatever the test result. Given that the option *do test* has a cost and does not increase the effectiveness, it is dominated by the option *do not test*.

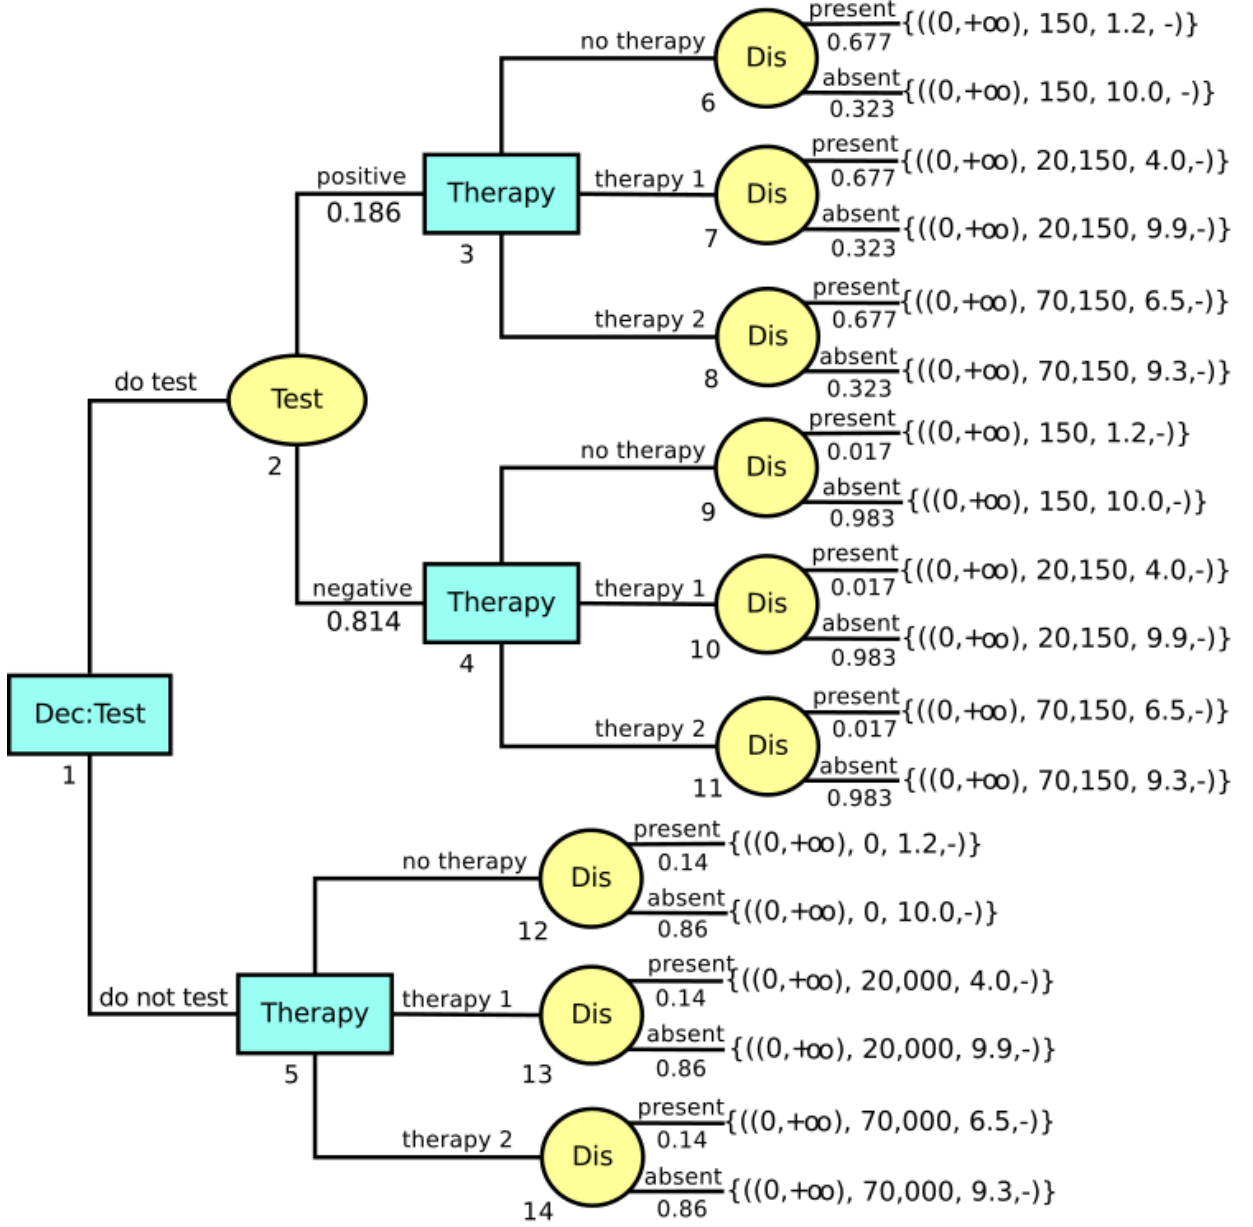


Figure 4: Initialization of the decision tree for the Example 1. Each leaf node has a CEP consisting of only one interval, $(0, +\infty)$. All the interventions for the leaf nodes are null, because no decision node has been evaluated yet.

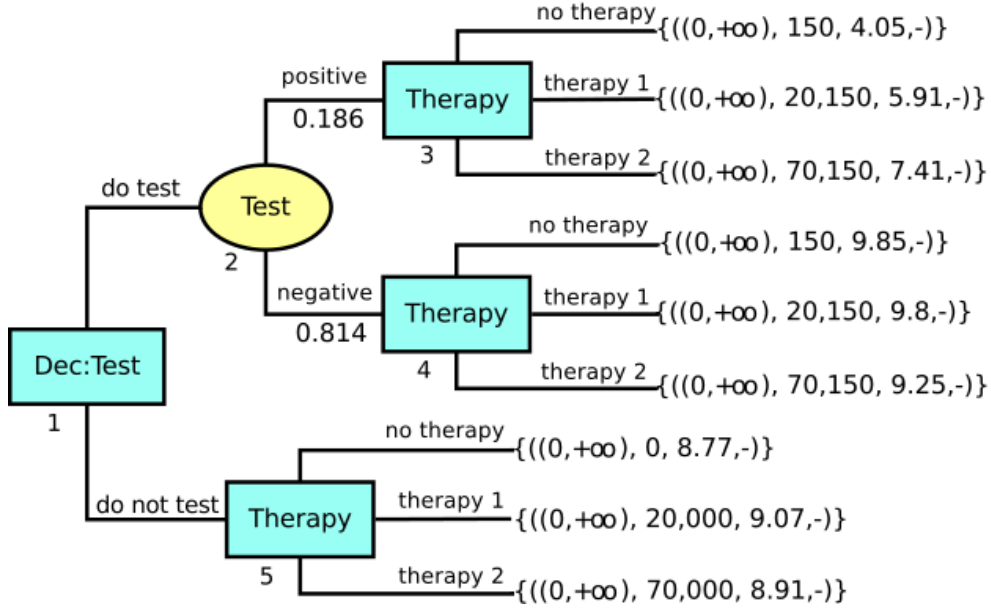


Figure 5: Tree obtained after evaluating the nine *Disease* chance nodes (6 to 14).

In the second interval (10,739, 33,384), the cost and effectiveness of each option are those shown in Figure 11(b). As no option dominates the other, we have to calculate the ICER, which is 11,171 €/QALY. This threshold splits the interval in two: (10,739, 11,171) and (11,171, 33,384). In the first one, the best option is *do not test* and in the second *do test*.

In the third interval (33,384, 65,539), no option dominates the other, as shown in Figure 11(c). The ICER is 21,072 €/QALY, that lies outside the interval being analyzed. Therefore, the interval is not divided. The best option for this interval, in which $\lambda > 21,072$ €/QALY, is *do test*.




Figure 6: Evaluation of decision node 3. The values of cost and effectiveness are taken from the three input CEPs, shown in Figure 5. There is only one interval on which to perform a deterministic CEA: $(0, +\infty)$. The output CEP is shown in Figure 9.




Figure 7: Evaluation of decision node 4, in the interval $(0, +\infty)$. The input CEPs are shown in Figure 5; the output, in Figure 9.




Figure 8: Evaluation of decision node 5, in the interval $(0, +\infty)$. The input CEPs are shown in Figure 5; the output, in Figure 9.

In the fourth interval $(65,539, +\infty)$, shown in Figure 11(d), the ICER is 21,300 €/QALY, which, as in the previous case, lies outside the interval. The best option for this interval is *do test*.

The resulting CEP, which is shown in Table 2, has four thresholds: three proceeding from the input CEPs, plus one arisen when analyzing the second input interval.

However, we can see in that table that the first and second output intervals have the same cost, effectiveness, and optimal intervention. Therefore, the first threshold, 10,730 €/QALY is unnecessary and should be removed. For the same reason, the fourth threshold, 65,539 €/QALY, which separates the fourth and fifth intervals, must be removed. The result is the CEP shown in Table 3.

The result of evaluating the node *Dec:Test* is shown in Figure 10. We can observe that the two first intervals have the same cost, effectiveness and recommended therapy, so we can fuse them in only one. Same happens with the last two intervals.

We may wonder why some of the thresholds that arose in the analysis are later removed from the final CEP. If we look at the first threshold removed, 10,730 €/QALY, we observe that it appeared when evaluating node 3 (branch *do test, positive* result): it determined

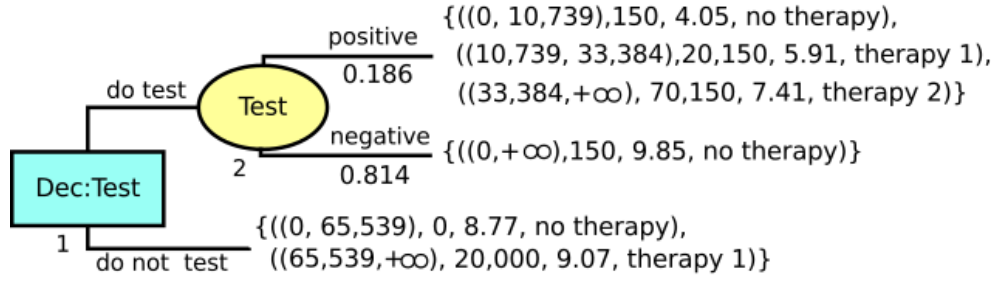


Figure 9: Decision tree obtained after evaluating the three *Therapy* nodes (3 to 5).




Figure 10: Decision tree after evaluating the node 2, *Test*.

Interval	Cost	Effectiveness	Dec:Test	Therapy
(0, 10,739)	0	8.77	Do not test	No therapy
(10,739, 11,171)	0	8.77	Do not test	No therapy
(11,171, 33,384)	3,874	9.11	Do test	{ test:positive→Therapy 1 test:negative→No therapy
(33,384, 65,539)	13,184	9.39	Do test	{ test:positive→Therapy 2 test:negative→No therapy
(65,539, + ∞)	13,184	9.39	Do test	{ test:positive→Therapy 2 test:negative→No therapy

Table 2: CEP obtained after evaluating the four input intervals of the root node, *Dec:Test*.

Interval	Cost	Effectiveness	Dec:Test	Therapy
(0, 11,171)	0	8.77	Do not test	No therapy
(11,171, 33,384)	3,874	9.11	Do test	{ test:positive→Therapy 1 test:negative→No therapy
(33,384, + ∞)	13,184	9.39	Do test	{ test:positive→Therapy 2 test:negative→No therapy

Table 3: Final CEP. It is obtained from the CEP in Table 2 after removing the unnecessary thresholds.

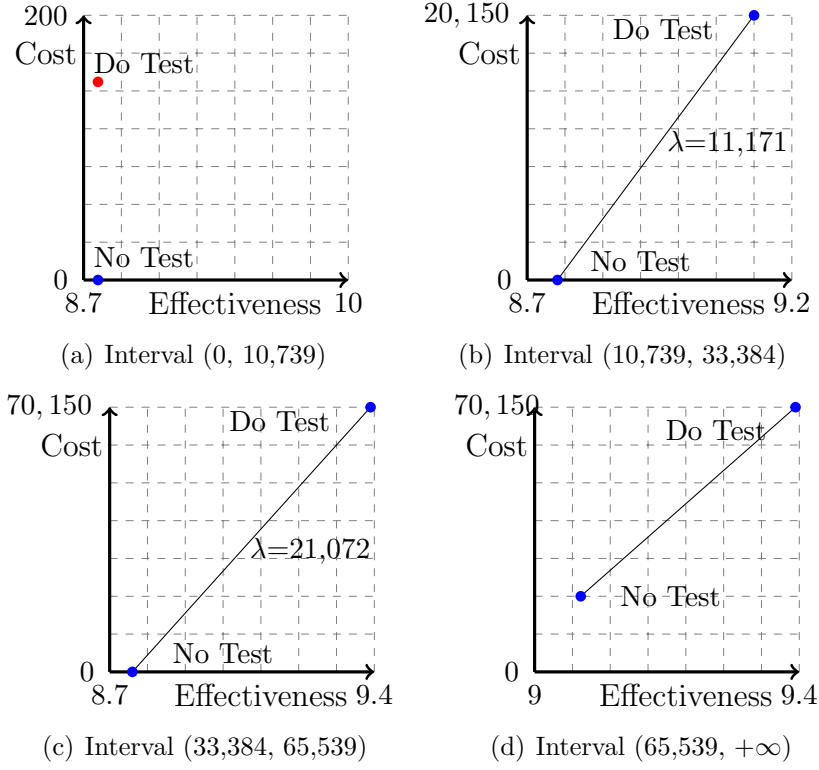


Figure 11: The evaluation of the root node, *Dec:Test*, consists of four deterministic CEAs, one for each interval.

that, in that scenario, when $\lambda > 10,730$ €/QALY *therapy 1* was more beneficial than *no therapy*, and when $\lambda < 10,730$ €/QALY it was the opposite. However, the subsequent analysis of node 1 showed that when $\lambda < 11,171$ €/QALY, the best option is *do not test*; i.e., the optimal path goes through the lower branch of *Dec:Test*, and for this reason the threshold that appeared in the upper branch (*do test*) becomes irrelevant.

The other threshold removed, 65,359 €/QALY, arose when evaluating node 4 (branch *do not test*), as the ICER between *no therapy* and *therapy 1*. However, the subsequent analysis of node 1 showed that when $\lambda > 33,384$ €/QALY, the best option is *do test*; i.e., the optimal path goes through the upper branch of *Dec:Test*, and therefore the threshold that appeared in the lower branch (*do not test*) is irrelevant.

5 Discussion

In Section 3 we have presented an algorithm for performing cost-effectiveness analysis (CEA) in decision trees with several decisions. The result of the evaluation is a cost-effectiveness partition (CEP) for each node. For each value of λ , the CEP of the root node gives, the expected cost, the expected effectiveness, and the optimal intervention. The net

monetary benefit (NMB)—that can be obtained from the cost, the effectiveness, and λ with Equation 10—and the optimal intervention for each value of λ are the same as those that we would have obtained if we had made the tree unicriterion by computing the NMB of each leaf and evaluated it with the standard (unicriterion) method.

We have based our analysis on the NMB. We would have obtained the same results if we had used the *net health benefit* (NHB) [?], defined as

$$NHB_{I_i}(\lambda) = e_{I_i} - c_{I_i}/\lambda . \quad (18)$$

The reason is that $NHB_{I_i}(\lambda) = \lambda \cdot NMB_{I_i}(\lambda)$ and therefore, when comparing two interventions, the one having the higher NMB also has the higher NHB:

$$\forall \lambda, \quad NMB_{I_i}(\lambda) > NMB_{I_j}(\lambda) \iff NHB_{I_i}(\lambda) > NHB_{I_j}(\lambda) . \quad (19)$$

5.1 Related work

To our knowledge, the problem of performing cost-effectiveness with multiple decisions and uncertain outcomes had not been solved satisfactorily in the literature. In general, textbooks on medical decision analysis discuss both CEA *and* decision trees, but not CEA *with* decision trees [4, 8, 10, 11, 13, 14, 15]. There are other books that explain how to do CEA with decision trees, but do not warn that embedded decision nodes are problematic [2, 3, 5, 16]. The only reference that we have found that discusses this issue is a paper by Kuntz and Weinstein [9, Sec. 7.2.1]. This reference states the following:

Cost-effectiveness analyses can be performed with a decision tree that has one decision node at the root. The branches of the initial decision node represent all of the strategies that are to be compared. Embedded, or downstream, decision nodes are not useful in cost-effectiveness analysis because the optimal branch cannot be determined when folding back the tree without an explicit decision rule for comparing costs and consequences.

The difficulty in dealing with embedded decision nodes explains why it is hard to find a case of CEA performed on a decision tree having more than one decisions in the literature. One exception is the example in [2, Sec. 5.2], based on a previous study by Hull et al. [7]. The root node of that tree represents the five options decision about prophylaxis. There are ten embedded decision nodes that represent the decision about the diagnostic technique used to confirm the presence of venous thromboembolism. However, each of these nodes has only one outgoing branch; hence, the evaluation is trivial, because the cost and effectiveness are the same of the only child of that node.

Another exception is the example in [16, Chapter 11, Sec. B.1]. The root node represents the main decision. There are three options: operating all patients, buying an ultrasound machine, or referring patients to a neighbor hospital. The tree is split into three figures (11-1, 11-2, and 11-3), one for each branch of the root not, that is not displayed graphically. There are two embedded decision nodes (see Fig. 11-2), that represent the choice between

surgery and ultrasound therapy. Chance nodes are evaluated by averaging the cost and effectiveness (life expectancy) of their branches, with the algorithm that we have explained in Section 2.3. The parameters of this hypothetical example were chosen so that when evaluating the two embedded decision nodes, ultrasound dominates surgery. (If we applied our method to that tree, the CEA of the embedded decision nodes would not introduce any threshold, and consequently each CEP would have only one interval, i.e., the cost and effectiveness would be the same for all the values of λ .) Unfortunately, that book do not explain how to evaluate embedded decision nodes when there is not an option that dominates all its alternatives.

In contrast, the commercial program TreeAge evaluates embedded decision nodes by applying an algorithm that, in essence, proceeds as follows: first, it eliminates the dominated interventions; second, it discards the interventions whose ICER is higher than the λ value (called WTP in TreeAge, for *willingness to pay*) chosen by the user, and finally it selects the most effective remaining alternative—see the sections “CE roll back optimal path parameters” and “The CE optimal path algorithm” in [17]. Clearly, the aim of this algorithm is to select the intervention with the highest NMB for that λ value, even though it is not explained this way in the user’s manual. The problem with this method is that in many cases the result of the CEA depends on the λ value chosen by the user, and consequently, some of the choices of λ may lead to wrong results. For example, if we build in TreeAge the decision tree shown in Figure 4 and evaluate it with $\lambda = 10,000$ €/QALY, its algorithm will choose *no therapy* for each of the embedded nodes, and as the therapy will not depend on the test, it is counterproductive to do it. This makes the option *do not test* dominate *do test*. On the contrary, if the user sets $\lambda = 50,000$ €/QALY, TreeAge gives an ICER of $\lambda = 21,072$ €/QALY, which is also wrong. If the user sets λ between 10,739 and 33,384, TreeAge returns the correct ICER, 11,171 €/QALY, but the user cannot know whether this value is correct or not. And even when the ICER is correct, the cost and effectiveness are not, because they depend on which therapy will be applied later (see again Table 3).

5.2 CEA with uni-decision trees

There is another method, alternative to the one proposed in this paper, capable of performing cost-effectiveness analyses correctly in problems that involve several decisions, but at the price of having to build much more complex decision trees. It is based on the advice of Kuntz and Weinstein [9], cited above, that the branches of the root node represent all the strategies; we can accomplish this by making each intervention include all the decisions. For instance, in the Example 1, this method would give rise to 12 interventions:

- I_{nn} : do not test; no therapy;
- I_{n1} : do not test; therapy 1;
- I_{n2} : do not test; therapy 2;

- I_{tnn} : do test; no therapy;
- I_{tn1} : do test; if positive, no therapy; if negative, therapy 1;
- I_{tn2} : do test; if positive, no therapy; if negative, therapy 2;
- I_{t1n} : do test; if positive, therapy 1; if negative, no therapy;
- I_{t11} : do test; always therapy 1;
- I_{t12} : do test; if positive, therapy 1; if negative, therapy 2;
- I_{t2n} : do test; if positive, therapy 2; if negative, no therapy;
- I_{t21} : do test; if positive, therapy 2; if negative, therapy 1,
- I_{t22} : do test; always therapy 2.

The corresponding tree would have 12 branches for the root node; each *do test* branch has two leaves and each *do not test* branch has four, which makes a total of 36 leaves. Using some common sense reasoning, we may discard the interventions I_{tnn} , I_{t11} , and I_{t22} because it is counterproductive to do the test when its result will not affect the decision about therapy. In the same way, the interventions I_{tn1} and I_{tn2} are clearly suboptimal because it does not make sense to apply a therapy when the test is negative and no therapy when it is positive. Therefore, the tree can be pruned down to 22 nodes. We have built that tree with TreeAge—see Figure 12—and it returns the same costs, effectiveness values, and thresholds as our method (Table 3), which is not surprising, because TreeAge always gives the right answer when every embedded decision node has only one outgoing branch.

The drawback of this method is that the size of the trees grows very fast. Thus, if we have m therapies and a test with n possible outcomes, there are $m + 1$ *do not test* interventions, one for each therapy option; each of these branches has two leaves, one for the presence of the disease and one for its absence. There are also $(m + 1)^n$ *do test* interventions; each one is a combination of n choices of therapy—one choice for each outcome of the test. We may remove the $m + 1$ interventions that make the same choice for every test outcome, as we did in the example above. Therefore, the tree would contain $[(m + 1)^n - (m + 1)]$ *do test* branches, each having $2n$ leaves, as we can see in Figure 12. This leads to a tree of $2[n(m + 1)^n - (m + 1)(n - 1)]$ leaves.⁴

In contrast, the tree used by our method is much smaller. It contains $n + 1$ *Therapy* nodes: one for each test outcome plus one for the *do not test* branch, as shown in Figure 4; each of them has $m + 1$ *Disease* nodes, which in turn have two branches. This makes a total of $2(n + 1) \cdot (m + 1)$ leaves in the tree.

⁴In some cases, this one-decision tree might be pruned further, in the same way as we eliminated the interventions I_{tn1} and I_{tn2} when building the tree in Figure 12. However, there is no algorithm for deciding which branches can be pruned without risking to miss the optimal intervention, and in any case the proportion of branches that may be pruned is small.

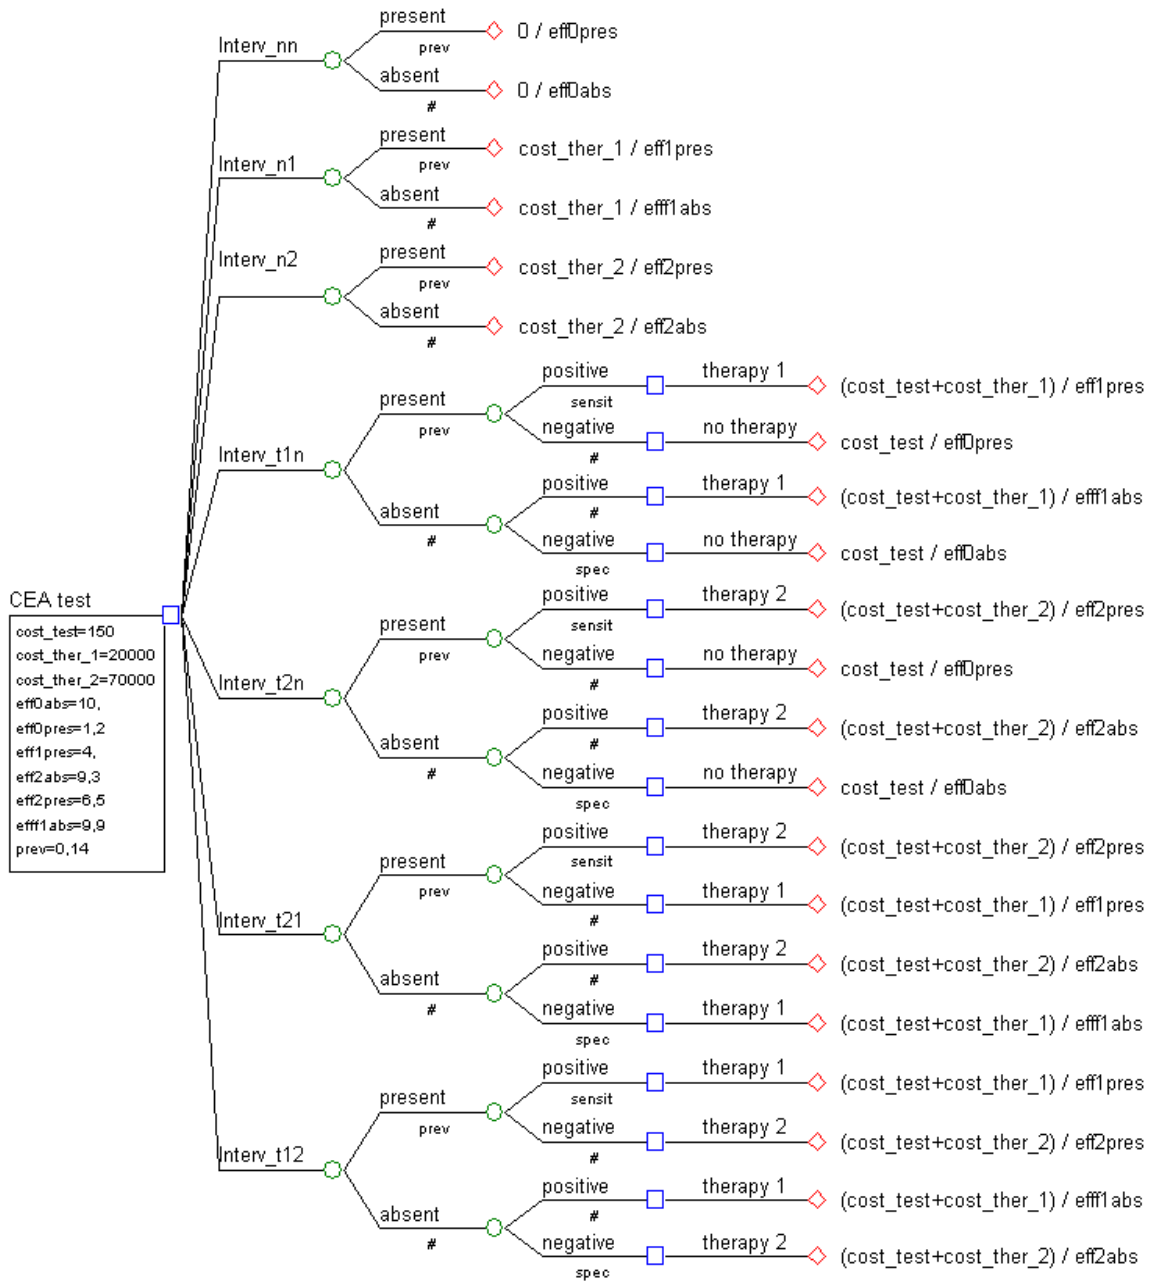


Figure 12: A decision tree for the Example 1, built with TreeAge. Each branch of the root node represents a complex intervention that includes all the decisions.

In a problem involving 3 test outcomes (for example, *low*, *medium*, and *high*) and 4 therapies (which may be, for instance, different dosage patterns for a drug), our method would analyze a tree of 20 leaves, while the one-decision tree of complex interventions would have 730 leaves. With 4 test outcomes and 5 therapies, our method would analyze a tree of 60 leaves, while the one-decision tree would have 10,332 leaves. Building a decision tree with dozens of branches is hard, but feasible, while building a tree with hundreds or thousands of nodes is impracticable. Consequently, our method permits to perform cost-effectiveness analysis for many problems that are out of the reach of traditional methods.

6 Conclusion

In this paper we have proposed a method for performing cost-effectiveness analysis in problems that involve several decisions and probabilistic outcomes. The method consists in building a tree with the same structure as if it were a unicriterion decision problem; this tree can be evaluated with a modified roll back algorithm that, instead of operating on single utilities, operates on cost-effectiveness partitions (CEPs). The CEP obtained for the root node gives, for each value of λ , the expected cost, the expected effectiveness, and the optimal intervention. The net monetary benefit (NMB) and the optimal intervention for each value of λ are the same as those that we would have obtained by computing the NMB of each leaf and evaluating the tree with the standard roll back algorithm.

In the discussion, we have shown that this problem had not been addressed properly in the past: most of the literature on medical decision analysis has ignored it, the algorithm used by TreeAge gives wrong results in many cases, and the solution proposed by Kuntz and Weinstein [9] is unfeasible, except for very small problems.

In a future paper we will show how the combination of CEPs, which is the basis of our method, can be adapted to influence diagrams [6, 12]. Using a software package developed by our group, we have been able to perform CEA on two influence diagrams for medical problems: the first one contains 5 decisions and 8 chance variables; the second, 4 decisions and 11 chance variables in its current version. The equivalent decision trees, which can be obtained automatically from the influence diagrams, have thousands of leaves. Clearly, those problems exceed by far the capabilities of standard CEA methods.

Acknowledgments

This work has been supported by grants TIN2006-11152 and TIN2009-09158, of the Spanish Ministry of Science and Technology, and by FONCICYT grant 85195.

A Appendix: Proofs

Proof of Proposition 3. It stems directly from Equation 1. (We have assumed that either $c_i \neq c_j$, or $e_i \neq e_j$, or both.)

Proof of Proposition 5. From the definition of ICER we have that

$$\begin{aligned}\lambda < ICER(I_i, I_j) &\implies \lambda \cdot (e_j - e_i) < c_j - c_i \\ &\implies \lambda \cdot e_j - c_j < \lambda \cdot e_i - c_i \implies NMB_{I_j}(\lambda) < NMB_{I_i}(\lambda) .\end{aligned}$$

The proof of the other two implications is almost identical.

Before proving Proposition 7, we introduce the following lemma.

Lemma 11 *If I_i and I_j dominate I_k , then $ICER(I_i, I_j) > ICER(I_k, I_j)$.*

Proof. From the definition of ICER, we have that

$$\begin{aligned}ICER(I_j, I_k) &= \frac{e_j - e_k}{c_j - c_k} = \frac{(e_j - e_i) - (e_k - e_i)}{c_j - c_k} \\ &= \frac{ICER(I_i, I_j) \cdot (c_j - c_i) - ICER(I_k, I_i) \cdot (c_k - c_i)}{c_j - c_k} .\end{aligned}$$

The fact that I_i and I_j dominate I_k implies that $ICER(I_i, I_k) > ICER(I_i, I_j)$, and consequently,

$$\begin{aligned}ICER(I_j, I_k) &< \frac{ICER(I_i, I_j) \cdot (c_j - c_i) - ICER(I_i, I_j) \cdot (c_k - c_i)}{c_j - c_k} \\ &= ICER(I_i, I_j) \frac{c_j - c_k}{c_j - c_k} = ICER(I_i, I_j) .\end{aligned}$$

Proof of Proposition 7. We analyze two cases. When $\lambda \leq ICER(I_i, I_j)$, we have, because of Proposition 5, that $NMB_{I_i}(\lambda) \geq NMB_{I_j}(\lambda)$. We also have $\lambda < ICER(I_i, I_k)$ and

$$NMB_{I_k}(\lambda) \leq NMB_{I_i}(\lambda) = \max(NMB_{I_i}(\lambda), NMB_{I_j}(\lambda)) .$$

The proof for the case $\lambda > ICER(I_i, I_j)$ is very similar.

References

- [1] S. B. Cantor. Cost-effectiveness analysis, extended dominance, and ethics: A quantitative assessment. *Medical Decision Making*, 14:259–265, 1994.
- [2] M. F. Drummond, M. J. Sculpher, G. W. Torrance, B. J. O’Brien, and G. L. Stoddart. *Methods for the Economic Evaluation of Health Care Programmes*. Oxford University Press, third edition, 2005.
- [3] T. L. Gift, A. C. Haddix, and P. S. Corso. Cost-effectiveness analysis. In A. C. Haddix, S. M. Teutsch, and P. S. Corso, editors, *Prevention Effectiveness. A Guide to Decision Analysis and Economic Evaluation*, pages 156–177. Oxford University Press, New York, second edition, 2003.

- [4] M. R. Gold, J. E. Siegel, L. B. Russell, and M. C. Weinstein. *Cost-Effectiveness in Health and Medicine*. Oxford University Press, New York, 1996.
- [5] A. M. Gray, P. M. Clarke, J. Wolstenholme, and S. Wordsworth. *Applied Methods of Cost-effectiveness Analysis in Healthcare*. Oxford University Press, New York, 2011.
- [6] R. A. Howard and J. E. Matheson. Influence diagrams. In R. A. Howard and J. E. Matheson, editors, *Readings on the Principles and Applications of Decision Analysis*, pages 719–762. Strategic Decisions Group, Menlo Park, CA, 1984.
- [7] R. D. Hull, J. Hirsh, D. L. Sackett, and G. L. Stoddart. Cost-effectiveness of primary and secondary prevention of fatal pulmonary embolism in high-risk surgical patients. *Canadian Medical Association Journal*, 127:990–995, 1982.
- [8] M. Hunink, P. Glasziou, J. Siegel, J. Weeks, J. Pliskin, A. Elstein, and M. Weinstein. *Decision Making in Health and Medicine*. Cambridge University Press, Cambridge, UK, 2001.
- [9] K. M. Kuntz and M. C. Weinstein. Modelling in economic evaluation. In M. F. Drummond and A. McGuire, editors, *Economic Evaluation in Health Care*, chapter 7, pages 141–171. Oxford University Press, New York, 2001.
- [10] H. M. Levin and P. J. McEwan. *Cost-Effectiveness Analysis*. Sage Publications, Thousand Oaks, CA, second edition, 2001.
- [11] P. Muenning. *Cost-Effectiveness Analysis in Health*. Jossey-Bass, San Francisco, CA, second edition, 2008.
- [12] R. F. Nease and K. D. K. Owens. Use of influence diagrams to structure medical decisions. *Medical Decision Making*, 17(3):263–275, 1997.
- [13] G. Parmigiani. *Modeling in Medical Decision Making*. John Wiley and Sons, Chichester, UK, second edition, 2002.
- [14] D. B. Petitti. *Meta-Analysis, Decision Analysis, and Cost-Effectiveness Analysis*. Oxford University Press, New York, second edition, 2000.
- [15] F. A. Sloan, editor. *Valuing Health Care*. Cambridge University Press, Cambridge, MA, 1995.
- [16] H. C. Sox, M. A. Blatt, M. C. Higgins, and K. I. Marton. *Medical Decision Making*. Butterworth-Heinemann, Woburn, MA, 1988.
- [17] TreeAge Software, Inc. *TreAge Pro User’s Manual*. TreeAge, Williamstown, MA, 2009.
- [18] M. C. Weinstein, H. V. Fineberg, A. S. Elstein, et al. *Clinical Decision Analysis*. Saunders, Philadelphia, PA, 1980.

- [19] M. C. Weinstein and W. B. Stason. Foundations of cost-effectiveness analysis for health and medical practices. *New England Journal of Medicine*, 296:716–721, 1977.
- [20] M. C. Weinstein, G. Torrance, and A. McGuire. QALYs: The basics. *Value in Health*, 12(Supplement I):S5–S9, 2009.