# IDENTIFYING SUBTREE PERFECTNESS IN DECISION TREES 

NATHAN HUNTLEY AND MATTHIAS C. M. TROFFAES


#### Abstract

In decision problems, often, utilities and probabilities are hard to determine. In such cases, one can resort to so-called choice functions. They provide a means to determine which options in a particular set are optimal, and allow incomparability among any number of options. Applying choice functions in sequential decision problems can be highly non-trivial, as the usual properties of maximizing expected utility may no longer be satisfied. In this paper, we study one of these properties: we revisit and reinterpret Selten's concept of subgame perfectness in the context of decision trees, leading us to the concept of subtree perfectness, which basically says that the optimal solution of a decision tree should not depend on any larger tree it may be embedded in. In other words, subtree perfectness excludes counterfactual reasoning, and therefore may be desirable from some philosophical points of view. Subtree perfectness is also desirable from a practical point of view, because it admits efficient algorithms for solving decision trees, such as backward induction. The main contribution of this paper is a very simple non-technical criterion for determining whether any given choice function will satisfy subtree perfectness or not. We demonstrate the theorem and illustrate subtree perfectness, or the lack thereof, through numerous examples, for a wide variety of choice functions, where incomparability among strategies can be caused by imprecision in either probabilities or utilities. We find that almost no choice function, except for maximizing expected utility, satisfies it in general. We also find that choice functions other than maximizing expected utility can satisfy it, provided that we restrict either the structure of the tree, or the structure of the choice function.


## 1. Introduction

In statistical decision problems, one is often faced with too little data for too many parameters, and elicitation of all probabilities and utilities can be a prohibitively expensive process, rendering full modeling infeasible. In such cases, it may still be possible to bound the set of reasonable models, i.e. to bound probability distributions and utility functions [3, 20, 27, 15, 14, 21, The way in which such models can be used for decision making is well studied [26, 11, 1, 2, 22, 12, 19, 23, 25, 13.

In this paper, we consider a very large class of decision models, namely any that can be represented by conditional choice functions on gambles: we merely assume that for any finite set of gambles (functions from the possibility space $\Omega$ to a set of rewards $\mathcal{R}$; these generalize random variables, or horse lotteries), and any conditioning event, the subject can give a non-empty subset of gambles that he considers optimal if the conditioning event were to occur. Gambles that are non-optimal would never be selected, and the subject is unable to express further preference between the optimal ones. Maximizing conditional expected utility is a simple example of such a choice function. General choice functions, however, need neither probability nor utility-not even a total preorder, and in fact, not even a partial one, although obviously these obtain as special cases.

Then, in a single agent sequential decision problem, at any stage, one has two ways of looking at its solution: the problem can be considered either in its simplest form-discarding any past

[^0]

Figure 1. Two-stage problem.


Figure 2. Second stage.
stages, or as part of a much larger problem - possibly considering choices and events that did not actually obtain. A reasonable requirement is that, at any stage, the solution is independent of the larger problem it is embedded in. In this paper, we call this requirement subtree perfectness.

Selten [24] introduced a similar idea for multi-agent extensive form games, called subgame perfectness. Solutions of extensive form games take the form of equilibrium points. If, for every subgame (a part of a game that is again a game), the restriction of the equilibrium point of the full game to that subgame also yields an equilibrium point of that subgame, then the equilibrium point of the full game is called subgame perfect. Selten showed that for games with perfect recall, (perfect ${ }^{11}$ ) equilibrium points are subgame perfect [24, p. 39, Thm. 2], that is, they are independent of any larger game in which they could be embedded.

In this paper we investigate single agent sequential decision making modeled by decision trees. Although such problems differ in many ways from extensive form games, subtree perfectness is clearly analogous to subgame perfectness, as the following example shows.

Consider the decision problem in Fig. 11 (replicated from [9). In the first stage, the subject chooses between taking scones, or proceeding to the second stage. In the second stage, the subject chooses between cake, or ice cream. Suppose the subject prefers to reject scones and to choose ice cream at the second stage. This strategy induces a substrategy in the subtree for the second stage: choose ice cream over cake.

But, as with multi-agent games, we can instead consider the subtree for the second stage separately, as in Fig. 2, If, in this smaller tree, the subject prefers ice cream, then his solution is subtree perfect: his solution for the full tree induces a strategy in the subtree, and this strategy coincides with his solution for the subtree. If the subject states a different preference (either no preference, or clear preference for cake), then his solution lacks subtree perfectness. So, subtree perfectness essentially means that the optimal induced strategies in a subtree do not depend on the full tree in which the subtree is embedded.

Our main result, given in Theorem 2 further in the paper, is an extremely straightforward necessary and sufficient condition on choice functions for subtree perfectness to hold in any decision tree. A similar result was given rigorous treatment in [9], however that treatment was rather technical and opaque, and not as obvious to apply to practical problems. The main contribution of this paper is to provide an overview of the result that is accessible and easy to apply, but still general enough to apply to any choice function. Therefore, there is far more emphasis on rigorous examples, diagrams, and procedures, and far less on detailed definitions and lengthy proofs. In so doing, we establish a useful philosophical and practical link between the work that has been done on the behavior of specific choice functions, such as [11, 2, 22, the more general and theoretical works such as [5, 17, 18, 9, and proposed algorithms for solving decision trees [6, 12, 4, 8, 13, 10]. Readers seeking formal proofs of the results, or those who are interested in proving similar results about sequential decision problems, are directed to 9 .

[^1]In this paper we consider only the traditional normal form method of listing all strategies, finding their corresponding gambles, applying a choice function, and listing all strategies that induce optimal gambles. In doing so, we assume act-state independence - that is, choice functions do not depend on the decision. This type of solution is also called resolute 18, 13. Several related works have considered the alternative extensive form solutions, such as in Hammond [5], Seidenfeld [21, and Machina [17. We do not consider extensive form solutions in this paper, although the results are closely related. Without going into details, subtree perfectness is similar to Hammond's consistency and consequentialism, Machina's separability over mutually exclusive events, and McClennen's separability and dynamic consistency [18]. Further details about the links with extensive form subtree perfectness can be found in 9 . Subtree perfectness also is related to the weaker concept of backward induction. More details on normal form backward induction can be found in (9, 10.

It should also be noted that the algorithm proposed by Kikuti et al. 12, 13 follows exactly our concept of normal form backward induction, although they then use it to construct an extensive form (what they call consequentialist) solution. They note that, for all the choice functions they consider, their solution is "inconsistent" with our normal form solution, which essentially means that subtree perfectness is violated. For two these choice functions, maximality and E-admissibility, it is observed that this inconsistency is not so severe: although arcs that are non-optimal in the full tree may appear in local trees, no arc that is non-optimal in a local tree may reappear in the full tree. In Fig. 1 , for example, it would not be possible under these criteria for the subject to say they preferred cake to ice cream at $N_{2}$ and then at $N_{1}$ include ice cream in the optimal set. This weaker form of subtree perfectness occurs because the two criteria satisfy the conditions of our backward induction theorem [10]. What this means in practice is that for these criteria the work of Kikuti et al. can be used to efficiently construct the consequentialist solution that they are interested in, and also the normal form solution that we are interested in, and although subtree perfectness does not hold, both the solutions are at least somewhat consistent and sensible.

The paper is structured as follows. Section 2.1 provides an introductory example using expected utility to outline what we mean by a solution of a tree, and what we mean by subtree perfectness. Sections 2.2 and 2.3 give rigorous, yet non-technical, definitions of the basic concepts required to formulate our main result. That result, the subtree perfectness theorem, is provided in Section 2.4. We then demonstrate the theorem on numerous choice functions and decision trees: Section 3.1 gives further examples of satisfaction of subtree perfectness for certain choice functions and trees, whilst Section 3.2 gives further examples of failures of subtree perfectness. Section 4 concludes the paper.

## 2. Subtree Perfectness

2.1. Examples: Expected Utility. We begin by considering some examples of solving small decision trees using expected utility. Of course, these solutions are almost trivial to find, but we spend some time on them to outline our ideas in a familiar environment. Consider the decision tree in Fig. 3, and suppose that $P\left(E_{1}\right)=0.6, P\left(E_{2}\right)=0.4$. Square decision nodes are points where the subject chooses the arc to follow. Round chance nodes are points where the arc followed depends on the (initially unknown) state of nature, in this case $E_{1}$ or $E_{2}$. Typically one would approach the problem by backward induction [16], but we consider an alternative method. This is because for more general choice functions, backward induction is more complicated and may not give the desired answer.

Consider a path that the subject may take through the tree. At $N$, she has a choice of arcs and suppose she follows the arc to $N_{1}$. Here again she has a choice; suppose she takes the arc to $N_{11}$. Here, she has no choice over the arc followed: nature decides. Thus, for this combination


Figure 3. Decision tree for first expected utility example.


Figure 4. Normal form solution for first expected utility example.
of her choices, her final reward is 2 if $E_{1}$ occurs and 0 otherwise. We can calculate the expected utility of this strategy to be 1.2.

There are three other strategies in the tree: the paths $N \rightarrow N_{1} \rightarrow N_{12}, N \rightarrow N_{2} \rightarrow N_{21}$, and $N \rightarrow N_{2} \rightarrow N_{22}$. These have expected utilities $1,1.8$, and 2 respectively. We therefore conclude that the solution of the tree is to take the path from $N$ to $N_{2}$ and then to $N_{22}$. This is illustrated in the decision tree in Fig. (4).

Now, subtree perfectness is a property relating solutions of large trees to solutions of subtrees. In this particular example, there are two non-trivial subtrees of the full tree, shown in Fig. 5. We can solve these trees in the same manner; indeed we have already calculated the required expected utilities. The optimal paths are $N_{1} \rightarrow N_{11}$ in the left-hand tree, and $N_{2} \rightarrow N_{22}$ in the right-hand tree. These solutions are displayed in Fig 6 .

Are these local solutions consistent with the solution in Fig. 4]. The first thing to note is that $N_{1}$ does not appear at all in Fig. (4) We can therefore ignore it: behavior in a subtree that will never actually be reached is irrelevant to subtree perfectness. At $N_{2}$, we see that both the solutions (Fig. [4] and Fig. 6) contain only the path to $N_{22}$, so there is no inconsistency here either. We say that subtree perfectness holds for this problem.


Figure 5. Subtrees for first expected utility example.


Figure 6. Solutions of subtrees for first expected utility example.


Figure 7. Decision tree for second and third expected utility examples.

Next we move onto a slightly more involved example, in which the initial node is a chance node, as shown in Fig. 77 We need a few more probabilities for this example: let $P\left(E_{1} \mid A_{1}\right)=0.6$, $P\left(E_{2} \mid A_{1}\right)=0.4, P\left(E_{1} \mid A_{2}\right)=0.4, P\left(E_{2} \mid A_{2}\right)=0.6, P\left(A_{1}\right)=0.5$, and $P\left(A_{2}\right)=0.5$.

In this example, the two available strategies are "if $A_{1}$, take the path $N_{1} \rightarrow N_{11}$ ", and "if $A_{1}$, take the path $N_{1} \rightarrow N_{12}$ ". Note that there is no decision node following $A_{2}$, so the subject does not need to specify anything in that case. These strategies have slightly more complicated gambles than before. Consider the first: if $A_{1}$ and $E_{1}$ both occur, the outcome is 2 ; if $A_{1}$ and $E_{2}$ both occur, the outcome is 0 ; if $A_{2}$ and $E_{1}$ both occur, the outcome is 3 ; and if $A_{2}$ and $E_{2}$ both occur, the outcome is 0 . This gives expected utility $0.5(0.6 \cdot 2+0.4 \cdot 0)+0.5(0.4 \cdot 3+0.6 \cdot 0)=1.2$. The second strategy gives expected utility $0.5 \cdot 1+0.5(0.4 \cdot 3+0.6 \cdot 0)=1.1$. Hence the solution is the first strategy, shown in Fig 8 .


Figure 8. Normal form solution for second expected utility example.


Figure 9. Subtree for second expected utility example.

Again, we look for non-trivial subtrees. In this example, there is only one, at $N_{1}$. This is shown in Fig 9 - note that this is exactly the left-hand tree of Fig. 5. We can solve this as usual, although here we have to remember that the event $A_{1}$ has occurred, and so we must use the suitable conditional probabilities. The expected utilities of the two strategies in this case are $0.6 \cdot 2+0.4 \cdot 0=1.2$ and 1 respectively, so the normal form solution of this subtree is to take the path $N_{1} \rightarrow N_{11}$ (the same as the left-hand tree in Fig. (6).

Once again we check for subtree perfectness. In both the global and the local solutions, the only path from $N_{1}$ is the one to $N_{11}$, hence subtree perfectness is satisfied in this example.

Finally, we provide an example of expected utility failing subtree perfectness. As we show later, this requires a conditioning event with probability zero. Changing the probabilities in the previous example to $P\left(A_{1}\right)=0, P\left(A_{2}\right)=1$ will be sufficient to fail subtree perfectness. Now the expected utilities of the two strategies in the tree of Fig. 7 are $0(0.6 \cdot 2+0.4 \cdot 0)+1(0.4 \cdot 3+0.6 \cdot 0)=$ 1.2 and $0 \cdot 1+1(0.4 \cdot 3+0.6 \cdot 0)=1.2$ respectively. Hence, we cannot choose between the two strategies using expected utility, and the normal form solution must include both. This is illustrated in Fig 10

Because we did not change any of the conditional probabilities, the normal form solution of the subtree at $N_{1}$ will be the same as in the previous example: the left-hand tree of Fig. 6. Now we see inconsistency between the two solutions: one of the strategies in the global solution contains the path $N_{1} \rightarrow N_{12}$, but this path does not appear in the local solution. We say that subtree perfectness does not hold at $N_{1}$.


Figure 10. Normal form solution for third expected utility example.

Having seen how subtree perfectness works for a standard choice function and simple trees, we now explain how to carry out the steps above for any tree and any choice function.
2.2. Definition of Subtree Perfectness. To state our main result, we need to define subtree perfectness for arbitrarily large (but finite) decision trees. In this paper, we will describe and define the concept of subtree perfectness directly on decision trees, without further formal mathematical notation. The tree in Fig. 11 will serve as a leading example. Of course, subtree perfectness can also be defined more formally, using mathematical notation that is very useful for proofs, as we did in 9, however such approach is rather technical, and unnecessary for illustrating the main ideas. Any reader who is interested in rigorous mathematical formulations and proofs of properties relating to choice functions and decision trees is referred to [9]. It should also be noted that, although throughout this paper almost all our examples feature real-valued rewards, the language and results we introduce here will work for any form of rewards.

The solutions we consider are sets of strategies. A strategy corresponds to the subject initially specifying all her actions in all eventualities, that is, an action for each decision node she may possibly reach. Upon arriving at a decision node, she then follows the mandated action. Thus, after the strategy is chosen, the subject no longer has any control. In other words, she is a resolute decision maker [18].

A strategy is easy to represent as a subtree of the initial tree. Consider the decision tree in Fig. 11 and the strategy "at $N_{1}$ choose $d_{S}$, then if $S_{1}$ occurs choose $d_{1}$ and if $S_{2}$ occurs choose $d_{2}$ ". This can be represented by the decision tree in Fig. 12. In this figure, an arc with a double


Figure 11. A more complex example of a decision tree.


Figure 12. One of the strategies of the tree in Fig. 11.
line through means the arc has been deleted. So the decision tree representation of a strategy is the original tree minus all decision arcs that are not in the strategy. Note that the labels of the nodes were retained in this process: this will be an important point later on. Also note that this is now a trivial decision tree, since there is only one option available at each decision node.

In general, we can construct a strategy for a given decision tree by the following method:
(1) For each decision node, pick exactly one decision arc and delete all others.
(2) Delete everything not connected to the original root node of the tree.

By applying this method for all possible combinations of choices of decision arcs, we arrive at the set of all available strategies for a given decision tree $2^{2}$

[^2]

Figure 13. Possible normal form solution of Fig. 1

As in the examples earlier, we consider solutions of decision trees of the following type: first we find all available strategies, then we apply some criterion to them that returns a non-empty subset of strategies. Ideally, we might want our criterion to return a single strategy. However, as argued in the introduction, in many situations the subject may not have enough information to make such a judgment, for instance because their probabilities and utilities are not fully determined. A function that is applied to a set of options to return a non-empty subset is called a choice function. The set of strategies that a choice function returns is called a normal form solution of the decision tree. So in the case of Fig. 11, we would first find the six strategies and then apply our choice function to these. Let us suppose that a particular choice function gives the normal form solution in Fig. 13. Note that here, we apply the choice function conditional on the certain event. Conditioning on other events is only relevant for subtrees, discussed next.

To examine subtree perfectness, we need to know how to restrict trees to smaller subtrees and solve these subtrees, and also how to restrict the solutions of larger trees to smaller subtrees. The first task is very straightforward: for each node on a tree, the subtree at that node is the node itself and all of its descendants. For example, the subtree at $N_{111}$ in Fig. 11 is shown in Fig. 14. Observe that the labels of the nodes are as in the original tree. Then the final step of the procedure is simply to solve this subtree, conditional on the logical conjunction of all events


Figure 14. The subtree of Fig. 11 at $N_{111}$.


Figure 15. Possible normal form solution of Fig. 14 ,
leading up to $N_{111}$ (or, conditional on the certain event, if no chance nodes are visited) $3^{3}$ In our example, we find the two normal form solutions in Fig. 14 and then apply our choice function, conditional on $S_{1}$ (this is the only event leading up to $N_{111}$ ), to this set. Let us suppose this gives the singleton strategy in Fig. 15 ,

This process of restriction and then solution can be summarized as follows.
(1) Choose a node $N$ to which to restrict the tree $T$.
(2) Find the conjunction of events which lead up to $N$ from the root of $T$. Denote this conjunction by $E$. If there are no events leading up to $N$, then take $E$ to be the certain event.
(3) Remove from $T$ all non-descendants of $N$ other than $N$ itself. This is the subtree at $N$.
(4) Find all strategies for this subtree.
(5) Apply the choice function to this set of strategies, conditional on $E$. This is the normal form solution of the restricted tree.
The opposite order, solving the large tree and then restricting the solution, is a little more involved. We saw the first step in Fig. 13 . How can this solution be restricted to, say, $N_{111}$ ? This is done element-by-element. We take the first strategy in the solution, and restrict it to $N_{111}$ exactly as above. This gives the same tree as in Fig. 15, We then attempt to restrict the second strategy to $N_{111}$, but we notice that this node does not appear in the strategy. In such a case, we remove the strategy entirely. Next we move to the third strategy, but again $N_{111}$ is not present, so remove this strategy too. This is the final strategy in the solution, so we find that the restriction of the normal form solution to $N_{111}$ is the singleton in Fig. 15. In this example, we see that restricting the tree to $N_{111}$ and then solving is the same as solving the full tree and restricting the solution to $N_{111}$. We say that subtree perfectness is satisfied at $N_{111}$.

In general, the process for solving a tree and then restricting the solution to a subtree goes as follows.
(1) Find all strategies for the tree.
(2) Apply the choice function to this set of strategies. This is the normal form solution of the tree.

[^3](3) Choose a node $N$ to restrict the solution to.
(4) For each strategy $T^{\prime}$ in the solution:

- If $N$ is not a node of $T^{\prime}$, remove $T^{\prime}$.
- Otherwise, remove from $T^{\prime}$ all non-descendants of $N$ other than $N$ itself.
(5) This gives the restriction of the normal form solution.

Now we have all the tools to define subtree perfectness.
Definition 1 (Subtree Perfectness). Let a decision tree $T$, a node $N$ in $T$, and a choice function be given. Carry out the operations of restricting the solution of $T$ to $N$, and solving the restriction of $T$ to $N$. The choice function is called subtree perfect for $T$ at $N$ if either

- the restriction of the solution is empty (that is, $N$ is not part of any strategy in the solution of $T$ ); or
- the restriction of the solution equals the solution of the restriction.

The choice function is called subtree perfect for $T$ if it is subtree perfect at all nodes $N$ in $T$. The choice function is called subtree perfect if it is subtree perfect for every $T$.

Subtree perfect solutions are useful for several reasons. When solving a decision problem, one can usually see it as being embedded in some larger decision problem that started in the past. With a subtree perfect solution, the rest of this larger problem can be ignored. A choice function that fails subtree perfectness may require consideration of events that did not happen or options that were refused in the past. This is practically and perhaps intuitively unappealing. Similarly, if the decision tree changes midway through the problem (for instance, new options that were not modeled become available), lack of subtree perfectness may require considering the whole problem again, whereas with subtree perfectness only the local problem needs to be solved.
2.3. Choice Functions on Gambles. The procedure introduced for finding a normal form solution involved applying a choice function to the set of available strategies. In practice, it is rare to apply choice functions to sets of strategies, but rather to sets of gambles. A gamble is a map from the set of possible outcomes to a set of rewards - to avoid technicalities, we assume the set of possible outcomes to be finite.

For example, consider the three strategies in Fig. [13. The first of these will give 9 if $S_{1} \cap E_{1}$ occurs, 14 if $S_{1} \cap E_{2}$ occurs, 4 if $S_{2} \cap E_{1}$ occurs, and 19 if $S_{2} \cap E_{2}$ occurs. This covers all possible outcomes, so for any outcome we know what reward the subject will receive. This is exactly what we call a gamble.

We write such a gamble $9 S_{1} E_{1}+14 S_{1} E_{2}+4 S_{2} E_{1}+19 S_{2} E_{2}$, or even as $S_{1}\left(9 E_{1}+14 E_{2}\right)+$ $S_{2}\left(4 E_{1}+19 E_{2}\right)$. Similarly, the other two strategies in Fig. 13 correspond to the gambles $S_{1}\left(10 E_{1}+\right.$ $\left.15 E_{2}\right)+S_{2}\left(10 E_{1}+15 E_{2}\right)$ and $S_{1}\left(5 E_{1}+20 E_{2}\right)+S_{2}\left(5 E_{1}+20 E_{2}\right)$ respectively. Note that which one of $S_{1}$ and $S_{2}$ occurs is irrelevant to both of these gambles, so they could be written as $10 E_{1}+15 E_{2}$ and $5 E_{1}+20 E_{2}$ respectively. These objects are much more natural to work with; for instance, should we know the probabilities of the four events, then we could calculate the expected utility of the three gambles. To work with sequential decision making, we actually need conditional choice functions: the set returned by the choice function depends on what event is being conditioned upon.

Many, if not all, common choice functions used in decision theory are choice functions on gambles, hence restricting attention to them makes sense. This adds two extra steps to the procedure for finding a normal form solution:
(1) Find all available strategies.
(2) For each strategy, find its corresponding gamble.
(3) Apply the choice function to this set of gambles.


Figure 16. Decision trees for which subtree perfectness implies subtree perfectness for any other tree.
(4) For each gamble chosen by the choice function, find all strategies in the original set that induced this gamble. This is the normal form solution. In other words, a strategy is in the normal form solution if and only if its corresponding gamble is chosen by the choice function.

Note that we lose some generality by moving to choice functions on gambles, because there may be several strategies with the same corresponding gamble. A choice function on gambles would be unable to distinguish between these, but a choice function on strategies would. This loss is not so important, given that it would be rare to want to distinguish between strategies that give the same reward in every eventuality.
2.4. Subtree Perfectness Theorem. Now we can present our main result. We are interested in when a particular choice function on gambles is subtree perfect. It turns out that all that is required for this is subtree perfectness for two simple classes of decision trees, shown in Fig. 16. In that figure, the $X_{i}, Y_{i}$ and $Z$ are arbitrary gambles. As part of a tree, they represent a terminal chance node. For instance, the gamble $X=5 E_{1}+3 E_{2}$ would represent a chance node with two arcs, with event $E_{1}$ leading to reward 5 , and event $E_{2}$ leading to reward 3.

Theorem 2 (Subtree Perfectness Theorem). A choice function on gambles is subtree perfect if and only if it is subtree perfect for all decision trees of the two types in Fig. 16. for every conditioning event $B$ at $N_{1}$ such that $A \cap B \neq \emptyset$ and $A^{c} \cap B \neq \emptyset$, and for every non-empty conditioning event $B$ at $N_{2}$.

So, in Fig. 16, we consider solving the full trees conditional not just on the certain event, but conditional on every possible non-empty event that does not cause the decision tree to have a path corresponding to $\emptyset$.

This theorem is equivalent to the subtree perfectness theorem presented and proved in [9, Theorem 22], but is perhaps more straightforward to understand. All we need to check for subtree perfectness is two simple types of decision tree, and this proves very useful when seeking both proofs and counterexamples for a particular choice function. Note that it is not immediately obvious that this theorem is equivalent to [9, Theorem 22], but in fact it appears in the proof [9, A.9].

## 3. Examples

3.1. Examples of Subtree Perfectness. In this section we give some examples of choice functions and decision trees where subtree perfectness holds. We concentrate on two simple trees, shown in Fig. 16. This is not just for simplicity; as we saw in Section 2, these two types of trees are key to subtree perfectness.

### 3.1.1. Expected Utility.

Example 3. First, we consider the simplest case, and investigate the tree in Fig. 16)(b), with conditioning event $B$. Maximizing conditional expectation for this tree is always subtree perfect, simply because

$$
\begin{align*}
\max \left\{E\left(X_{1} \mid B\right), \ldots\right. & \left., E\left(X_{n} \mid B\right), E\left(Y_{1} \mid B\right), \ldots, E\left(Y_{m} \mid B\right)\right\}  \tag{1}\\
& =\max \left\{\max \left\{E\left(X_{1} \mid B\right), \ldots, E\left(X_{n} \mid B\right)\right\}, \max \left\{E\left(Y_{1} \mid B\right), \ldots, E\left(Y_{m} \mid B\right)\right\}\right\}
\end{align*}
$$

Example 4. Consider the tree in Fig. 1ø(a), with conditioning event $B$ such that $A \cap B \neq \emptyset$ and $A^{c} \cap B \neq \emptyset$. On the one hand, the normal form solution of the full tree is obtained by those gambles $A X_{i}+A^{c} Z$ which maximize their expectation conditional on $B$ :

$$
\begin{equation*}
\arg \max _{i=1}^{n} E\left(A X_{i}+A^{c} Z \mid B\right) \tag{2}
\end{equation*}
$$

On the other hand, the normal form solution of the subtree at $N_{11}$ is obtained by those gambles $X_{i}$ which maximize their expectation conditional on $A \cap B$ :

$$
\begin{equation*}
\arg \max _{i=1}^{n} E\left(X_{i} \mid A \cap B\right) \tag{3}
\end{equation*}
$$

Now, note that, by the partition theorem,

$$
\begin{equation*}
E\left(A X_{i}+A^{c} Z \mid B\right)=P(A \mid B) E\left(X_{i} \mid A \cap B\right)+P\left(A^{c} \mid B\right) E\left(Z \mid A^{c} \cap B\right) \tag{4}
\end{equation*}
$$

Because the term $P\left(A^{c} \mid B\right) E\left(Z \mid A^{c} \cap B\right)$ is independent of the index $i$, the solution of both maximization problems is identical:

$$
\begin{equation*}
\arg \max _{i=1}^{n} E\left(A X_{i}+A^{c} Z \mid B\right)=\arg \max _{i=1}^{n} E\left(X_{i} \mid A \cap B\right) \tag{5}
\end{equation*}
$$

provided that $P(A \mid B)>0$. In other words, maximizing expected utility satisfies subtree perfectness on this type of tree whenever the probability of every non-empty event is strictly positive.

From the above analysis, and Theorem 2, we recover a well-known result: expected utility satisfies subtree perfectness for arbitrary decision trees, provided that all chance arcs have strictly positive probability. Interestingly, if some chance arcs have probability zero, subtree perfectness may fail; we already saw an instance of such failure in the last example of Section 2.1.
3.1.2. Choice Functions Induced By A Total Preorder. In Example 3, where we established subtree perfectness of maximizing expectation for the tree in Fig. 1G(b), all we really relied on is the fact that expectation induces a total preorder over gambles.

Let us be a bit more specific. Consider a relation $\succeq$ over gambles, which satisfies transitivity, reflexivity, and completeness; that is, for all gambles $X, Y$, and $Z$, we require that:
(1) $X \succeq Y$ and $Y \succeq Z$ implies $X \succeq Z$;
(2) $X \succeq X$; and
(3) $X \succeq Y$ or $Y \succeq X$.

Any relation which satisfies these properties is called a total preorder. Now, an obvious way to associate a choice function with a total preorder is to select those gambles which dominate all other gambles:

$$
\begin{equation*}
\max _{\succeq} \mathcal{X}:=\{X \in \mathcal{X}: \forall Y \in \mathcal{X}, X \succeq Y\} \tag{6}
\end{equation*}
$$

for any finite set of gambles $\mathcal{X}$. Now, it is easily established that, for any total preorder $\succeq$ over gambles,

$$
\begin{equation*}
\max _{\succeq}\left\{X_{1}, \ldots, X_{n}, Y_{1}, \ldots, Y_{m}\right\}=\max _{\succeq}\left\{\max _{\succeq}\left\{X_{1}, \ldots, X_{n}\right\}, \max _{\succeq}\left\{Y_{1}, \ldots, Y_{m}\right\}\right\} \tag{7}
\end{equation*}
$$

In other words, any choice function induced by a total preorder is subtree perfect for trees of the type of the tree in Fig. 16(b)

In fact, a much stronger result holds: a choice function is subtree perfect only if it is induced by a total preorder [9, Theorem $22 \&$ Lemma 15]. Consequently, many of the choice functions used for imprecise probability and imprecise utility fail subtree perfectness, as they are typically not induced by total preorders; specifically, they typically drop completeness. Interestingly, completeness is not required for some schemes of backward induction to work [12, 8, 13, 10].
3.1.3. Maximin. The maximin criterion values a gamble at its smallest possible reward, and then chooses gambles that give the maximum value. Because it is a total preorder, subtree perfectness will hold for the tree in Fig. 16)(b), as seen in Section 3.1.2,

Maximin does not in general hold for the tree in Fig. 18(a), we shall see this later in Example 9 .
3.1.4. $\Gamma$-Maximin. Suppose that too little information is available to identify a unique probability distribution, but we can specify a set of plausible probability distributions (see for instance [3, 27, 15, 28, 26]). In particular, suppose that

- rewards are expressed in utiles, so $\mathcal{R}=\mathbb{R}$,
- the subject can express her beliefs by means of a set $\mathcal{M}$ of probability distributions $P$ ( $\mathcal{M}$ is called the credal set, and is typically assumed to be closed and convex, however we need no such assumption here), and
- each probability distribution $P \in \mathcal{M}$ satisfies $P(A)>0$ for all events $A \neq \emptyset$.
- Note that the subject does not express beliefs about the relative plausibility of the distributions, so there is no second-order probability over the set. Also note that the positive probabilities are not necessary to define imprecise probability, but this assumption lets us avoid some technical details.
Under the above assumptions, each $P$ in $\mathcal{M}$ determines a conditional expectation $\sqrt{4}$

$$
E_{P}(X \mid A)=\frac{\sum_{x \in \mathbb{R}} x P(X=x)}{P(A)}
$$

and the whole set $\mathcal{M}$ determines a conditional lower and upper expectation

$$
\begin{equation*}
\underline{E}(X \mid A)=\min _{P \in \mathcal{M}} E_{P}(X \mid A) \quad \bar{E}(X \mid A)=\max _{P \in \mathcal{M}} E_{P}(X \mid A) \tag{8}
\end{equation*}
$$

and this for every gamble $X$ and every non-empty event $A$. We consider four choice functions for this uncertainty model. Each is introduced only briefly; see [25] for a more in-depth discussion and comparison.

The first, $\Gamma$-maximin, selects any gamble that maximizes lower expectation $\underline{E}$. Failure of subtree perfectness for $\Gamma$-maximin has been well-documented [21, 22], but for lower expectations with a particular structure, subtree perfectness can hold.

Example 5. Consider the tree in Fig. 10); for simplicity of notation, we assume here that the conditioning event $B$ is the certain event, however the analysis extends trivially to any conditioning event.

Suppose also that $\underline{E}$ satisfies marginal extension [26, §6.7.2], that is,

$$
\begin{equation*}
\underline{E}\left(A X+A^{c} Y\right)=\underline{E}\left(A \underline{E}(X \mid A)+A^{c} \underline{E}\left(Y \mid A^{c}\right)\right) . \tag{9}
\end{equation*}
$$

for any gambles $X$ and $Y$. Note that this is naturally satisfied if beliefs are expressed as an unconditional lower expectation over gambles of the form $A \alpha+A^{c} \beta$ (for arbitrary $\alpha, \beta \in \mathbb{R}$ ), and conditional lower expectations $\underline{E}(\cdot \mid A)$ and $\underline{E}\left(\cdot \mid A^{c}\right)$.

[^4]On the one hand, the normal form solution of the full tree is obtained by those gambles $A X_{i}+$ $A^{c} Z$ which maximize their lower expectation:

$$
\begin{equation*}
\arg \max _{i=1}^{n} \underline{E}\left(A X_{i}+A^{c} Z\right) \tag{10}
\end{equation*}
$$

On the other hand, the normal form solution of the subtree at $N_{11}$ is obtained by those gambles $X_{i}$ which maximize their lower expectation conditional on $A$ :

$$
\begin{equation*}
\arg \max _{i=1}^{n} \underline{E}\left(X_{i} \mid A\right) \tag{11}
\end{equation*}
$$

Now, note that, by marginal extension,

$$
\begin{equation*}
\underline{E}\left(A X_{i}+A^{c} Z\right)=\underline{E}\left(A \underline{E}\left(X_{i} \mid A\right)+A^{c} \underline{E}\left(Z \mid A^{c}\right)\right) . \tag{12}
\end{equation*}
$$

Using the above identity, one can easily show that the solution of both maximization problems is identical:

$$
\begin{equation*}
\arg \max _{i=1}^{n} \underline{E}\left(A X_{i}+A^{c} Z\right)=\arg \max _{i=1}^{n} \underline{E}\left(X_{i} \mid A\right) \tag{13}
\end{equation*}
$$

provided that $\underline{E}(A)>0$. In other words, $\Gamma$-maximin satisfies subtree perfectness on this type of tree whenever the lower probability of every event is strictly positive.

Clearly, $\Gamma$-maximin corresponds to a total preorder, and therefore, it will also satisfy subtree perfectness for the tree in Fig. 1][(b)] Concluding, as long as the lower probability of every event is strictly positive, and marginal extension is satisfied with respect to the way the chance nodes partition the possibility space, $\Gamma$-maximin satisfies subtree perfectness.
3.1.5. $\Gamma$-Maximax. The next criterion, $\Gamma$-maximax, selects any gamble that maximizes upper expectation $\bar{E}$. The treatment of this criterion is very similar to that of $\Gamma$-maximax. Here, we simply mention that $\Gamma$-maximax satisfies subtree perfectness under exactly the same conditions as $\Gamma$-maximin.
3.1.6. Maximality. Maximality starts from a strict partial order between gambles, conditional on an event $A$ :

$$
\begin{equation*}
X \succ_{A} Y \text { whenever } \underline{E}(X-Y \mid A)>0 \tag{14}
\end{equation*}
$$

and then selects those gambles which are undominated with respect to this strict partial order:

$$
\begin{equation*}
\max _{\succ_{A}} \mathcal{X}:=\left\{X \in \mathcal{X}: \forall Y \in \mathcal{X}, Y \nsucc_{A} X\right\} \tag{15}
\end{equation*}
$$

for any finite set of gambles $\mathcal{X}$.
Here, one can show that, under strictly positive lower probability of conditioning events, maximality will satisfy subtree perfectness for trees of the type of Fig. 16(a), essentially because

$$
\begin{equation*}
\underline{E}\left(\left(A X_{i}+A^{c} Z\right)-\left(A X_{j}+A^{c} Z\right) \mid B\right)=\underline{E}\left(A\left(X_{i}-X_{j}\right) \mid B\right) . \tag{16}
\end{equation*}
$$

has the same sign as $\underline{E}\left(X_{i}-X_{j} \mid A \cap B\right)$, because of the generalized Bayes rule 26], provided that $\underline{E}(A \mid B)>0$. Note that we do not require marginal extension.

We will see later that it does not satisfy subtree perfectness for trees of the type of Fig. 16(b)
In other words, maximality satisfies subtree perfectness only in some special cases. Interestingly, there are backward induction schemes which do work for maximality (under strictly positive lower probability) [12, 8, [13, 10].


Figure 17. A tree to demonstrate that subtree perfectness is not as strong as it may first appear.


Figure 18. Solution of the tree of Fig. 17
3.1.7. E-admissibility. E-admissibility, selects any gamble that maximizes expectation under at least one probability distribution in $\mathcal{M}$.

Under strictly positive lower probability, as with maximality, it satisfies subtree perfectness for trees of the type of Fig. 16(a). It does not satisfy subtree perfectness for trees of the type of Fig. 16(b)

The next example uses E-admissibility to demonstrate that subtree perfectness is not as strong a condition as it may first appear.

Example 6. Consider the tree in Fig. 17. Suppose that $\mathcal{M}=\left\{P_{1}, P_{2}\right\}$, with $P_{1}\left(A_{1}\right)=P_{1}\left(A_{2}\right)=$ $P_{2}\left(A_{1}\right)=P_{2}\left(A_{2}\right)=0.5, P_{1}\left(B_{1} \mid A_{1}\right)=P_{1}\left(B_{1} \mid A_{2}\right)=0.6$, and $P_{2}\left(B_{1} \mid A_{1}\right)=P_{2}\left(B_{1} \mid A_{2}\right)=0.4$.

The normal form solution of the full tree is depicted in Fig. 18. The left-hand solution maximizes expectation under $P_{1}$, whilst the right-hand solution maximizes expectation under $P_{2}$. Of the four strategies, two are optimal according to E-admissibility.

The normal form solutions of the subtrees at $N_{1}$ and $N_{2}$ are depicted in Fig. 19. In both subtrees, there are two possible strategies, and both are optimal according to E-admissibility.

Subtree perfectness holds here, but observe that only two of the four global strategies are optimal, whilst all of the local strategies are optimal.

(A)

(B)

Figure 19. Solutions of the subtrees at $N_{1}$ and $N_{2}$, for the tree of Fig. 17 ,


Figure 20. An example involving imprecise utility.

This shows that subtree perfectness is weaker than what one might call normal-extensive form equivalence. For further discussion of the relationship between subtree perfectness and normal-extensive form equivalence, we refer to [9].
3.2. Failures of Subtree Perfectness. In Section 2.1. we already visited one well-known instance subtree perfectness failing, namely expected utility, when events of zero probability are involved. In this section, we provide many more examples.

### 3.2.1. Imprecise Utility.

Example 7. Consider again the very first example from the introduction (Fig. 11), redrawn in Fig. 20; note that this tree is an extremely simple instance of the tree in Fig. 1G) which plays a central role in our subtree perfectness theorem.

Suppose we are unsure about our utility between the three available options, and we are happy to accept any option as optimal that maximizes either of the following two utility functions:

$$
U_{1}\left(r_{2}\right)=1
$$

$$
\begin{equation*}
U_{1}\left(r_{3}\right)=2 \tag{17}
\end{equation*}
$$

$$
\begin{align*}
& U_{1}\left(r_{1}\right)=3 \\
& U_{2}\left(r_{1}\right)=-3 \tag{18}
\end{align*}
$$

$$
U_{2}\left(r_{2}\right)=1,
$$

$$
U_{2}\left(r_{3}\right)=2
$$

Clearly, for the full tree, the optimal choice is $\left\{r_{1}, r_{3}\right\}$. However, for the subtree at $N_{2}$, the optimal choice is $\left\{r_{1}, r_{2}\right\}$. Subtree perfectness is violated, because $r_{2}$ is not optimal in the full tree, whereas it is optimal for the subtree at $N_{2}$.


Figure 21. An example involving a partial order.


Figure 22. Failure of subtree perfectness for E-admissibility.


Figure 23. E-admissible normal form solution of the tree of Fig. 22
3.2.2. Choice Functions Not Induced By A Total Preorder. The previous example easily adapts to choice functions induced by a strict partial order. Indeed, assume that we have a strict partial order $\succ$ on gambles, and three gambles, $X_{1}, X_{2}$, and $X_{3}$, such that:

$$
\begin{equation*}
X_{3} \succ X_{2} \tag{19}
\end{equation*}
$$

but where all other pairs are incomparable:

$$
\begin{array}{ll}
X_{1} \nsucc X_{2}, & X_{1} \nsucc X_{3}, \\
X_{2} \nsucc X_{1}, & X_{3} \nsucc X_{1} .
\end{array}
$$

Considering now the situation depicted in Fig. 21, we are in exactly the same situation as before: clearly, for the full tree, the optimal choice is $\left\{X_{1}, X_{3}\right\}$. However, for the subtree at $N_{2}$, the optimal choice is $\left\{X_{1}, X_{2}\right\}$. Subtree perfectness is violated, because $X_{2}$ is dominated, and whence, not optimal, in the full tree, whereas it is optimal for the subtree at $N_{2}$.

Although this is not a rigorous proof, it does demonstrate why any non-trivial partial order might fail subtree perfectness. In other words, maximality (and also point-wise dominance and interval dominance, which we have not discussed here, but which are also induced by partial orders and are commonly used in the literature [25]) will fail subtree perfectness in general.

Although E-admissibility does not correspond to a partial order, unsurprisingly, it fails subtree perfectness in a very similar way.


Figure 24. E-admissible normal form solution of the subtree at $N_{2}$, for the tree of Fig. 22.


Figure 25. Failure of subtree perfectness for maximin.

Example 8. Consider the decision tree in Fig. 22. Suppose that $\mathcal{M}=\left\{P_{1}, P_{2}\right\}$, where $P_{1}\left(A_{1}\right)=$ $P_{2}\left(A_{2}\right)=0.8$. The expectation of $N_{3}$ is 3 under $P_{1}$ and -3 under $P_{2}$. The remaining two strategies have constant values, namely 1, and 2.

In the full tree, clearly, of the three strategies, only two are E-admissible, as depicted in Fig. 23. However, for the subtree at $N_{2}$, both strategies are optimal, as depicted Fig. 24. Subtree perfectness is violated, because the strategy with constant outcome 1 is dominated by the strategy with constant outcome 2, and whence, 1 is not optimal in the full tree, whereas it is optimal for the subtree at $N_{2}$. In fact, we have mimicked the situation of Example 7 .

### 3.2.3. Maximin.

Example 9. Consider the tree in Fig. 25: note that this tree is an extremely simple instance of the tree in Fig. 10) which plays a central role in our subtree perfectness theorem.

Obviously, there are only two strategies. In the full tree, the worst possible outcome for both strategies is 0 , and therefore both strategies are maximin. However, clearly, at $N_{2}$, only the strategy yielding 2 is optimal. Whence, subtree perfectness is violated.
3.2.4. $\Gamma$-maximin. When Eq. (9) does not hold, $\Gamma$-maximin often fails to be subtree perfect, in a similar way that maximin fails it.

Example 10. Consider the tree in Fig. [26] again note that this tree is a simple instance of the tree in Fig. 10) which plays a central role in our subtree perfectness theorem.

Suppose that $\mathcal{M}=\left\{P_{1}, P_{2}\right\}$ with

| $P_{1}$ | $A_{1}$ | $A_{2}$ |
| :--- | :--- | :--- |
| $B_{1}$ | 0.1 | 0.1 |
| $B_{2}$ | 0.4 | 0.4 | and $\quad$| $P_{2}$ | $A_{1}$ | $A_{2}$ |
| :--- | :--- | :--- |
| $B_{1}$ | 0.4 | 0.4 |
| $B_{2}$ | 0.1 | 0.1 |

The two strategies of the full tree are depicted in Fig. 27, and correspond to the following gambles:

|  | $A_{1}$ | $A_{2}$ |
| :---: | :---: | :---: |
| $B_{1}$ | 2 | -1 |
| $B_{2}$ | -1 | 2 | and |  | $A_{1}$ | $A_{2}$ |
| :---: | :---: | :---: |
| $B_{1}$ | 0 | -1 |
| $B_{2}$ | 0 | 2 |



Figure 26. Failure of subtree perfectness for $\Gamma$-maximin.



Figure 27. The two strategies for the tree of Fig. 26. Only the left-hand strategy is $\Gamma$-maximin optimal.



Figure 28. The two strategies for the subtree at $N_{1}$, for the tree of Fig. 26, Only the right-hand strategy is $\Gamma$-maximin optimal.

The lower expectation of the first strategy is 0.5 , whereas the second strategy has lower expectation -0.2. So, in the full tree, only the first strategy is optimal according to $\Gamma$-maximin.

However, in the subtree at $N_{1}$, the conditional probabilities are

|  | $P_{1}\left(\cdot \mid A_{1}\right)$ |
| :---: | :---: |
| $B_{1}$ | 0.2 |
| $B_{2}$ | 0.8 | and | $P_{2}\left(\cdot \mid A_{1}\right)$ |
| :--- |
| $B_{1}$ |

and two strategies, depicted in Fig. 28, of this subtree, correspond to the following gambles:

$$
\left.\begin{array}{l|ccc|c}
B_{1} & 2 \\
B_{2} & -1
\end{array} \quad \text { and } \quad \begin{array}{c}
B_{1} \\
B_{2}
\end{array}\right)
$$

The lower expectation of the first strategy is -0.4 , whereas the second strategy has lower expectation 0 . So, in the subtree, only the second strategy is optimal according to $\Gamma$-maximin.

Concluding, $\Gamma$-maximin does not satisfy subtree perfectness. In fact, the preference between the two strategies reverses depending on whether we consider the full tree, or only the subtree at $N_{1}$.

We refer to Seidenfeld [22] for further beautiful examples.

## 4. Conclusion

We introduced the concept of subtree perfectness: in essence, we extended Selten's idea of subgame perfectness to resolute normal form solutions of decision trees. Subtree perfectness allows us to solve decision trees without worrying about which larger tree they are embedded in. This is a particularly useful property if one desires to avoid counterfactual reasoning for philosophical reasons. It also allows for efficient computational algorithms to solve decision trees via backward induction, and is tremendously helpful when something unexpected happens in the middle of a sequential decision problem, as subtree perfectness allows us to restrict our investigation to a subtree of the larger problem.

We saw that, when using choice functions on gambles, subtree perfectness for all decision trees is equivalent to subtree perfectness for two simple classes of decision trees. Numerous examples for popular choice functions, demonstrated how this result can be leveraged both for proving subtree perfectness and for finding simple counterexamples. The application of the results to counterexamples is particularly useful, since it shows that, when searching for a counterexample, there is no reason to try complicated trees because a small one will always suffice.

For choice functions that are not total preorders (that is, those that attempt to model indeterminacy or indecision), we highlight a few important consequences. First, such choice functions always fail subtree perfectness in general; see the examples in Sections 3.2.1 and 3.2.2 and 9 , Lemmas $15 \& 21]$ for a full proof. We demonstrated that these counterexamples typically appear when the tree contains paths with more than one decision node. However, and this may come as a surprise, such choice functions can satisfy subtree perfectness for decision trees where a chance node is followed by single decision nodes; see the examples in Sections 3.1.4 3.1.7 Most examples of subtree perfectness in the literature have been total preorders failing subtree perfectness in exactly these trees, which may have created the impression that only expected utility can be used in such problems, and it may not be so well-known that alternatives such as partial orders can be useful here too. Problems in which one will observe some data and then make a single decision are not rare. Therefore, it is worth knowing that, in such cases, expected utility does not have a monopoly on subtree perfectness.

In this paper we have only considered a subset of trees of this type: those in Fig. 16.(a) and the one in Fig. 17. We can consider a wider class of trees of this type: those with any size of partition at the first chance node, with each chance arc leading to a decision node involving any number of gambles. We could then try to find conditions for subtree perfectness in all trees of this type, and to find how one has to restrict particular choice functions to attain this behavior (for instance, an extension of Example (5). Some preliminary results can be found in (7].

Subtree perfectness has strong links with several other properties of the solutions of sequential decision problems. The first is the classical method of backward induction: solving decision trees by "rolling back". This algorithm has not often been employed for criteria other than total preorders, but two similar algorithms have been proposed by Huntley and Troffaes [10] and Kikuti et al. [13]. In particular the latter reference contains all the information required to find
the normal form solution for maximality and E-admissibility via backward induction. We can ask when backward induction gives the same solution as the normal form solution we have defined in this paper. As both the above works note, it is possible for backward induction to work but subtree perfectness to fail. Again, this typically never occurs when using a total preorder, so this result may not be well known. It is, however, quite easy to show that subtree perfectness implies backward induction [9, Corollary 27].

Another related question is the equivalence of our normal form solutions with extensive form solutions. In extensive form solutions, the subject makes decisions only upon reaching a decision node, rather than making all decisions ahead of time. Such a solution can be seen as a subtree of the initial tree, with some decision arcs removed. Upon reaching a decision node, the subject then simply picks one of the arcs that remains. Since subtree perfectness refers to consistency between local and global solutions, one might expect it to be sufficient for the equivalence between the two forms, but consider Example 6. Here, subtree perfectness is satisfied, but there is no extensive form solution available that we could really call equivalent. This behavior arises because there are many more possible normal form solutions than there are extensive form solutions. This topic requires further study; introductory material can be found in [9, § 7.2].

## Acknowledgments

The authors are indebted to Teddy Seidenfeld for suggesting the term 'subtree perfectness'. EPSRC supported the first author.

## References

[1] T. Augustin, On decision making under ambiguous prior and sampling information, ISIPTA 01: Proceedings of the Second International Symposium on Imprecise Probabilities (G. de Cooman, T.L. Fine, S. Moral, and T. Seidenfeld, eds.), 2001.
[2] _ Expected utility within a generalized concept of probability - a comprehensive framework for decision making under ambiguity, Statistical Papers 43 (2002), no. 1, 5-22.
[3] George Boole, An investigation of the laws of thought on which are founded the mathematical theories of logic and probabilities, Walton and Maberly, London, 1854.
[4] G. De Cooman and Matthias C. M. Troffaes, Dynamic programming for deterministic discrete-time systems with uncertain gain, International Journal of Approximate Reasoning 39 (2005), no. 2-3, 257-278, doi:10.1016/j.ijar.2004.10.004
[5] P. Hammond, Consequentialist foundations for expected utility, Theory and Decision 25 (1988), no. 1, 25-78.
[6] David Harmanec, Generalizing Markov decision processes to imprecise probabilities, Journal of Statistical Planning and Inference 105 (2002), no. 1, 199-213.
[7] Nathan Huntley, Sequential decision making for choice functions on gambles, Ph.D. thesis, Durham University, 2011.
[8] Nathan Huntley and Matthias C. M. Troffaes, An efficient normal form solution to decision trees with lower previsions, Soft Methods for Handling Variability and Imprecision (Didier Dubois, M. Asunción Lubiano, Henri Prade, María Ángeles Gil, Przemyslaw Grzegorzewski, and Olgierd Hryniewicz, eds.), Advances in Soft Computing, Springer, Sep 2008, pp. 419-426.
[9] Nathan Huntley and Matthias C. M. Troffaes, Subtree perfectness, backward induction, and normalextensive form equivalence for single agent sequential decision making under arbitrary choice functions, arXiv:1109.3607v1 [math.ST], September 2011, arXiv:1109.3607
[10] Nathan Huntley and Matthias C. M. Troffaes, Normal form backward induction for decision trees with coherent lower previsions, Annals of Operations Research 195 (2012), no. 1, 111-134, arXiv:1104.0191, doi:10.1007/s10479-011-0968-2
[11] J. Jaffray, Rational decision making with imprecise probabilities, 1st International Symposium on Imprecise Probabilities and Their Applications, 1999.
[12] D. Kikuti, F. Cozman, and C.P. de Campos, Partially ordered preferences in decision trees: Computing strategies with imprecision in probabilities, IJCAI-05 Multidisciplinary Workshop on Advances in Preference Handling (R. Brafman and U. Junker, eds.), 2005, pp. 118-123.

13] Daniel Kikuti, Fabio Gagliardi Cozman, and Ricardo Shirota Filho, Sequential decision making with partially ordered preferences, Artificial Intelligence 175 (2011), no. 7-8, 1346-1365, doi:10.1016/j.artint.2010.11.017
[14] H. E. Kyburg, Rational belief, Behavioral and Brain Sciences 8 (1983), no. 2.
[15] I. Levi, The enterprise of knowledge, MIT Press, London, 1980.
[16] D. V. Lindley, Making decisions, 2nd ed., Wiley, London, 1985.
[17] M. J. Machina, Dynamic consistency and non-expected utility models of choice under uncertainty, Journal of Economic Literature 27 (1989), no. 1622-1688.
[18] E. F. McClennen, Rationality and dynamic choice: Foundational explorations, Cambridge University Press, 1990.
[19] T. D. Nielsen and J. Jaffray, Dynamic decision making without expected utility: An operational approach, European Journal of Operational Research 169 (2006), 226-246.
[20] Frank P. Ramsey, Truth and probability, Foundations of Mathematics and other Logical Essays (R. B. Braithwaite, ed.), Routledge and Kegan Paul, London, 1931, Published posthumously, pp. 156-198.
[21] T. Seidenfeld, Decision theory without 'independence' or without 'ordering': What is the difference?, Economics and Philosophy 4 (1988), 267-290.
[22] , A contrast between two decision rules for use with (convex) sets of probabilities: $\Gamma$-maximin versus E-admissibility, Synthese 140 (2004), 69-88.
[23] T. Seidenfeld, M. J. Schervish, and J. B Kadane, Coherent choice functions under uncertainty, 5th International Symposium on Imprecise Probability: Theories and Applications, 2007.
[24] R. Selten, Reexamination of the perfectness concept for equilibrium points in extensive games, International Journal of Game Theory 4 (1975), no. 1, 25-55.
[25] Matthias C. M. Troffaes, Decision making under uncertainty using imprecise probabilities, International Journal of Approximate Reasoning 45 (2007), 17-29, doi:10.1016/j.ijar.2006.06.001
[26] P. Walley, Statistical reasoning with imprecise probabilities, Chapman and Hall, London, 1991.
[27] Peter M. Williams, Notes on conditional previsions, Tech. report, School of Math. and Phys. Sci., Univ. of Sussex, 1975.
[28] , Notes on conditional previsions, International Journal of Approximate Reasoning 44 (2007), no. 3, 366-383.

SYSTeMS Research Group, Universiteit Gent, Belgium
E-mail address: nath.huntley@gmail.com
Dept. of Mathematical Sciences, Durham University, UK
E-mail address: matthias.troffaes@durham.ac.uk


[^0]:    Key words and phrases. Choice function, subtree perfectness, backward induction, sequential, decision tree, consequentialist, maximality, maximin, E-admissibility, imprecise probability, imprecise utility.

[^1]:    ${ }^{1}$ Without going into much detail, a perfect equilibrium point is one which is stable under small perturbations [24 p. 38].

[^2]:    ${ }^{2}$ Of course, in practice, there is a more efficient way to find all strategies-for instance, in Fig. 11 there are $2^{4}=16$ combinations of arc choices, but these lead to only 6 distinct strategies: evidently, one does not need to iterate over arcs of decision nodes in subtrees following deleted arcs.

[^3]:    ${ }^{3}$ We exclude decision trees that have $\emptyset$ for the conditioning event of any of their subtrees, as most theories of conditioning do not allow conditioning on the empty set. This goes without loss of generality, because any decision tree can be 'fixed' simply by pruning its $\emptyset$-conditioned subtrees.

[^4]:    ${ }^{4}$ In this paper, gambles, such as $X$, are always assumed to take only a finite number of values.

