

SRI International

DECISION ANALYSIS USING BELIEF FUNCTIONS

Technical Note No. 472

September 27, 1989

By: Thomas M. Strat

Artificial Intelligence Center

Computer and Information Sciences Division

Approved for public release; distribution unlimited

"The views, opinions, and findings contained in this report are those of the author(s) and should not be construed as an official Department of Defense position, policy, or decision, unless so designated by other official documentation."

The research reported herein was partially supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. N00039-88-C-0248 in conjunction with the U.S. Navy Space and Naval Warfare Systems Command.



Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 27 SEP 1989		2. REPORT TYPE		3. DATES COVERED 00-09-1989 to 00-09-1989	
4. TITLE AND SUBTITLE Decision Analysis Using Belief Functions				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) SRI International, 333 Ravenswood Avenue, Menlo Park, CA, 94025				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 33	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Abstract

A primary motivation for reasoning under uncertainty is to derive decisions in the face of inconclusive evidence. Shafer's theory of belief functions, which explicitly represents the unconstrained nature of many reasoning problems, lacks a formal procedure for making decisions. Clearly, when sufficient information is not available, no theory can prescribe actions without making additional assumptions. Faced with this situation, some assumption must be made if a clearly superior choice is to emerge. In this paper we offer a probabilistic interpretation of a simple assumption that disambiguates decision problems represented with belief functions. We prove that it yields expected values identical to those obtained by a probabilistic analysis that makes the same assumption. We maintain a strict separation between evidence that carries information about a situation and assumptions that may be made for disambiguation of choices. In addition, we show how the decision analysis methodology frequently employed in probabilistic reasoning can be extended for use with belief functions. This generalization of decision analysis allows the use of belief functions within the familiar framework of decision trees.

KEYWORDS: belief functions, decision analysis, decision-making, decision tree, Dempster-Shafer theory, evidential reasoning, reasoning under uncertainty

Contents

1	Introduction	1
2	Expected Value	2
2.1	Expected value using probabilities	2
2.2	Expected value intervals	4
2.3	Expected value using belief functions	6
2.4	Discussion	9
3	Decision Analysis	11
3.1	Decision analysis using probabilities	11
3.2	Decision analysis using belief functions	12
3.3	Generalized decision tree examples	16
3.4	Comparing two choices	20
4	Discussion	21
4.1	On making assumptions	21
4.2	On the limitations of the approach	22
4.3	On the automation of decision analysis	23
5	Summary	24
A	Notation	27
B	Decision-Making with Evidential Reasoning	28

1 Introduction

Decision analysis provides a methodological approach for making decisions. Uncertain states of nature are represented by probability distributions, and each possible state is assigned a value or *utility*. The best decision is the one that yields the greatest *expected utility*. By enumerating in a *decision tree* all available choices and assessing the probabilities and utilities of the states of nature that may result, one can mechanically determine the optimal sequence of actions he should take [4, 8, 9, 14].

In practice, these simple requirements are hard to satisfy [3]. Sometimes, reliable estimates of the probabilities involved are hard to come by. For example, few statistics are available for determining the probability of a nuclear reactor core meltdown. Assessing the utility of many-faceted states of nature is equally challenging. How should one give a unique value to the anticipated quality of married life? These limitations have hindered the more widespread application of decision analysis.

Shafer's theory of *belief functions* [12, 15, 16, 18] allows one to express partial beliefs when it is impossible or impractical to assess complete probability distributions confidently. Using belief functions, one can bound the probabilities of events for which the assignment of a precise probability would be misleading. The theory provides a facility to express one's beliefs only to the degree to which there is supporting evidence, thereby resulting in an appropriate description of an uncertain event. For example, there might be reason to assign a probability to a reactor malfunction, without saying what the chance is that it may lead to a core meltdown.

Despite its representational advantages, the theory of belief functions lacks a formal basis upon which decisions can be made in the face of ambiguity [1]. Computing the expected utility of a random event that has been represented with belief functions results in an *expected utility interval* (EUI). To choose between two actions one must compare their respective EUIs. If they don't overlap, the choice is clear. But when the EUIs overlap, the decision-maker is confronted with a dilemma — the available evidence does not support either choice. Ideally, one should collect more information until the intervals no longer overlap and the choice becomes clear. However, sometimes one is forced to choose without benefit of additional information. What should be done?

In this situation there is no recourse except to make an assumption to eliminate the ambiguity. Various authors have expressed preference for different assumptions (such as renormalization, generalized insufficient reason [2, 20], minimax [22] and optimism/pessimism [6]). More elaborate schemes have been suggested, but they also amount to the introduction of unfounded assumptions [11, 13, 23]. Here we advocate the interpolation of a point-valued utility within the EUI. We make no claim that it leads to superior decisions, but do claim that it is no less viable than the alternative assumptions. We show that it gives the same expected utility (and hence leads to the same decisions) as would be obtained by assuming that there is some probability that ambiguity will be resolved in one's favor.

We further show how decision analysis can be generalized to accommodate a belief function representation of uncertainty. This involves two modifications: allowing an interval as the

utility of a state or set of states, and allowing a belief function in place of a probability distribution. The result is a complete decision analysis procedure compatible with either probabilistic or belief function representations of uncertainty.

We should point out that decision theory (and its associated utility theory) is not the only approach for making decisions under uncertainty. For example, Lesh has proposed a model based on an ignorance-preference coefficient that is empirically derived [10]. Shafer has advocated a “constructive” decision theory which seeks support for actions that achieve goals [17]. Loui *et. al.* suggest representing beliefs not by one distribution, but by a sequence of progressively more decisive distributions [11]. In this paper we are concerned with providing for the use of belief functions within the general framework of decision analysis.

It is worth noting that none of the material described in this paper depends on the use of Dempster’s rule, which is commonly used in Shafer’s theory to combine independent bodies of evidence [16]. The computation of expected utility interval, and the procedure for using EUIs in decision analysis, only requires that a belief function representation of the problem be available. Dempster’s rule *could* be used to construct that belief function, but it is not required for decision analysis.

In the sections that follow we develop the theory and illustrate its use with simple examples. In Section 2 we derive the expected utility interval that results from the use of belief functions. We then show how making an assumption about the probability of nature’s cooperation leads to the same expected utility as interpolation within the EUI. In Section 3, this result is used to generalize decision analysis and is illustrated within a decision problem concerning whether or not to drill for oil. We conclude with a discussion of the benefits and limitations of our approach, and compare its use with other approaches to decision-making under uncertainty.

2 Expected Value

Decision analysis provides a methodological approach for making decisions. The crux of the method is that one should choose the action that will maximize the expected utility. In this section we review the computation of expected utility using a probabilistic representation of a simple example and show how a belief function gives rise to a range of expected utilities. We then show how a simple assumption about the inclination of nature leads to a means for choosing a single-point expected utility for belief functions.

2.1 Expected value using probabilities

Example – Carnival Wheel #1 A familiar game of chance is the carnival wheel pictured in Figure 1. This wheel is divided into 10 equal sectors, each of which is labeled with a dollar amount as shown. For a \$6.00 fee, the player gets to spin the wheel and receives the amount shown in the sector that stops at the top. Should we be willing to play?




Figure 1: Carnival Wheel #1

The analysis of this problem lends itself readily to a probabilistic representation. From inspection of the wheel (assuming each sector really is equally likely), we can construct the following probability distribution:

$$\begin{aligned}
 p(\$1) &= 0.4 \\
 p(\$5) &= 0.3 \\
 p(\$10) &= 0.2 \\
 p(\$20) &= 0.1.
 \end{aligned}$$

The expected value $E(x)$ is computed from the formula

$$E(x) = \sum_{x \in \Theta} x \cdot p(x) \tag{1}$$

where Θ is the set of possible outcomes. The expected value of the carnival wheel is \$5.90 as shown here:

x	$p(x)$	$x \cdot p(x)$
1	0.4	0.4
5	0.3	1.5
10	0.2	2.0
20	0.1	2.0
$E(x) =$		5.90




Figure 2: Carnival Wheel #2

Therefore, we should refuse to play, because the expected value of playing the game is less than the \$6.00 cost of playing.¹ Let us now modify the problem slightly in order to motivate a belief function approach to the problem.

2.2 Expected value intervals

Example – Carnival Wheel #2 Another carnival wheel is divided into 10 equal sectors, each having \$1, \$5, \$10, or \$20 printed on it. However, one of the sectors is hidden from view. How much are we willing to pay to play this game?

This problem is ideally suited to an analysis using belief functions. In a belief function representation, a unit of belief is distributed over the space of possible outcomes (commonly called the *frame of discernment*). Unlike a probability distribution, which distributes belief over *elements* of the outcome space, this distribution (called a *mass function*) attributes belief to *subsets* of the outcome space. Belief attributed to a subset signifies that there is reason to believe that the outcome will be among the elements of that subset, without committing to any preference among those elements. Formally, a mass distribution m_Θ is a mapping from subsets of a frame of discernment Θ into the unit interval:

¹We assume that the monetary value is directly proportional to utility because of the small dollar amounts involved. We could instead have chosen to work with utilities to account for nonlinearities in one's preferences for money.

$$m_{\Theta} : 2^{\Theta} \mapsto [0, 1],$$

such that

$$m_{\Theta}(\phi) = 0 \quad \text{and} \quad \sum_{A_i \subseteq \Theta} m_{\Theta}(A_i) = 1 .$$

Any subset to which nonzero mass has been attributed is called a *focal element*. One of the ramifications of this representation is that the belief in a hypothesis A ($A \subseteq \Theta$) is constrained to lie within an interval $[Spt(A), Pls(A)]$, where

$$Spt(A) = \sum_{A_i \subseteq A} m_{\Theta}(A_i) \ ; \ Pls(A) = 1 - Spt(\neg A) . \quad (2)$$

These bounds are commonly referred to as *support* and *plausibility*.

The frame of discernment Θ for Wheel #2 is $\{\$1, \$5, \$10, \$20\}$. The mass function for Wheel #2 is shown below,

$$\begin{aligned} m(\{\$1\}) &= 0.4 \\ m(\{\$5\}) &= 0.2 \\ m(\{\$10\}) &= 0.2 \\ m(\{\$20\}) &= 0.1 \\ m(\{\$1, \$5, \$10, \$20\}) &= 0.1 , \end{aligned}$$

and its associated belief intervals are

$$\begin{aligned} [Spt(\{\$1\}), Pls(\{\$1\})] &= [0.4, 0.5] \\ [Spt(\{\$5\}), Pls(\{\$5\})] &= [0.2, 0.3] \\ [Spt(\{\$10\}), Pls(\{\$10\})] &= [0.2, 0.3] \\ [Spt(\{\$20\}), Pls(\{\$20\})] &= [0.1, 0.2] . \end{aligned}$$

Before we can compute the expected value of the wheel represented by this belief function, we must somehow assess the value of the hidden sector. We know that there is a 0.1 chance that the hidden sector will be selected, but what value should we attribute to that sector? If the carnival hawker were allowed to assign a dollar value to that sector, he would surely have assigned \$1. On the other hand if we (or a cooperative friend) were allowed to do so, it would have been \$20. Any other assignment method would result in a value between \$1 and \$20, inclusive. Therefore, if we truly do not know what assignment method was used, the strongest statement that we can make is that the value of the hidden sector is between \$1 and \$20. Using interval arithmetic we can apply the expected value formula of Equation 1 to obtain an *expected value interval* (EVI):

$$E(x) = [E_*(x), E^*(x)] \quad (3)$$

where²

$$E_*(x) = \sum_{A_i \subseteq \Theta} \inf(A_i) \cdot m_{\Theta}(A_i)$$

$$E^*(x) = \sum_{A_i \subseteq \Theta} \sup(A_i) \cdot m_{\Theta}(A_i) .$$

The expected value interval of Wheel #2 is

$$E(x) = [0.4(1) + 0.2(5) + 0.2(10) + 0.1(20) + 0.1(1), \\ 0.4(1) + 0.2(5) + 0.2(10) + 0.1(20) + 0.1(20)]$$

$$E(x) = [5.50, 7.40] .$$

2.3 Expected value using belief functions

As many researchers have pointed out, an interval of expected values is not very satisfactory when we have to make a decision. Sometimes it provides all the information necessary to make a decision, e.g. if the game costs \$5 to play, then clearly we should be willing to play regardless of who gets to assign a value to the hidden sector. Sometimes we can defer making the decision until we have collected more evidence, e.g. if we could peek at the hidden sector and then decide whether or not to play. But the need to make a decision based on the currently available information is often inescapable, e.g. should we spin Wheel #2 for a \$6 fee? We will present our methodology for decision-making using belief functions after pausing to consider a Bayesian analysis of the same situation.

If we are to use the probabilistic definition of expected value from Equation 1, we are forced to assess probabilities of all possible outcomes. To do this, we must make additional assumptions before proceeding further. One possible assumption is that all four values of the hidden sector (\$1, \$5, \$10, \$20) are equally likely, and we could evenly distribute among those four values the 0.1 chance that the hidden sector is chosen. This is an example of the generalized insufficient reason principle advanced by Dubois and Prade [2] and by Smets [20]. The resulting computation of expected value with this assumption is shown below; the expected value is \$6.30:

x	$p(x)$	$x \cdot p(x)$
1	0.425	0.425
5	0.225	1.125
10	0.225	2.250
20	0.125	2.500
$E(x) =$		6.30

²We use $\inf(A_i)$ or $\sup(A_i)$ to denote the smallest or largest element in the set $A_i \subseteq \Theta$. Θ is assumed to be a set of scalar values [21].

An alternative assumption is that the best estimate of the probability distribution for the value of the hidden sector is the same as the known distribution of the visible sectors. Using this assumption, the result is \$6.00:

x	$p(x)$	$x \cdot p(x)$
1	4/9	4/9
5	2/9	10/9
10	2/9	20/9
20	1/9	20/9
$E(x) =$		6.00

Rather than making one of these assumptions, we may wish to parameterize by an unknown probability ρ our belief that either we get to choose the value of the hidden sector or the carnival hawker does. Let ρ be the probability that the value assigned to the hidden sector is the one that we would have assigned, if given the opportunity, so $(1 - \rho)$ is the probability that the carnival hawker chose the value of the hidden sector. That is,

$$\begin{aligned} p(\text{hidden sector is labeled } \$20) &= \rho \\ p(\text{hidden sector is labeled } \$1) &= 1 - \rho. \end{aligned}$$

The expected value of Wheel #2 can then be recomputed using probabilities and Equation 1 as illustrated here:

x	$p(x)$	$x \cdot p(x)$
1	$0.4 + 0.1(1 - \rho)$	$0.5 - 0.1\rho$
5	0.2	1.0
10	0.2	2.0
20	$0.1 + 0.1\rho$	$2.0 + 2\rho$
$E(x) =$		$5.50 + 1.90\rho$

To decide whether to play the game, we need only assess the probability ρ . For the carnival wheel it would be wise to allow that the hawker has hidden the value from our view; thus we might assume that $\rho = 0$. So $E(x) = 5.50$, and we should not be willing to pay more than \$5.50 to spin the wheel.

Example – Carnival Wheel #3 A third carnival wheel is divided into 10 equal sectors, each having \$1, \$5, \$10, or \$20 printed on it. This wheel has 5 sectors hidden from view. However, we do know that none of these sectors is a \$20, that the first hidden sector is either a \$5 or a \$10, and that the second hidden sector is either a \$1 or a \$10. How much are we willing to pay to spin Wheel #3?

A probabilistic analysis of Wheel #3 requires one to make additional assumptions. Estimating the conditional probability distribution for each hidden sector would provide enough information to compute the expected value of the wheel. Alternatively, estimating just the expected value of each hidden sector would suffice as well. However, doing so can be both tedious and frustrating: tedious because there may be many hidden sectors, and frustrating




Figure 3: Carnival Wheel #3

because we're being asked to provide information that, in actuality, we do not have. (If we knew the conditional probabilities or the expected values, we would have used them in our original analysis.) What is the minimum information necessary to establish a single expected value for Wheel #3?

The probability, ρ , that we used to analyze Wheel #2 can be used here as well.

Definition 1

*Let ρ = the probability that ambiguity will be resolved as favorably as possible;
 $(1 - \rho)$ = the probability that ambiguity will be resolved as unfavorably as possible.*

Estimating ρ is sufficient to restrict the expected value of a belief function to a single point. It is easy to see that the expected value derived from this analysis as ρ varies from 0 to 1 is exactly the value obtained by linear interpolation of the EVI that results from using belief functions. The following derivation shows that this is true in general.

Theorem 1 *Given a mass function m_Θ defined over a scalar frame Θ of utilities, and an estimate of ρ (the probability that all residual ambiguity will turn out favorably), the expected utility given m_Θ is*

$$E(x) = E_*(x) + \rho \cdot (E^*(x) - E_*(x)). \tag{4}$$

Proof:

Consider a mass function m_Θ defined over a frame of discernment Θ . Now consider any focal element $A \subseteq \Theta$, such that $m_\Theta(A) > 0$. Since ρ is the probability that a cooperative agent

will control which $x \in A$ will be selected, and $(1 - \rho)$ is the probability that an adversary will be in control, then the probability that x will be chosen given that focal element A occurs is

$$p_{\Theta}(x|A) = \begin{cases} \rho & \text{if } x = \sup(A) \\ (1 - \rho) & \text{if } x = \inf(A) \\ 0 & \text{otherwise.} \end{cases}$$

Considering all focal elements in m_{Θ} , we can construct a probability distribution $p_{\Theta}(x)$ as follows:

$$\begin{aligned} p_{\Theta}(x) &= \sum_{A_i \subseteq \Theta} p_{\Theta}(x|A_i) \cdot p_{\Theta}(A_i) \\ p_{\Theta}(x) &= \sum_{A: \sup(A)=x} \rho \cdot m_{\Theta}(A) + \sum_{A: \inf(A)=x} (1 - \rho) \cdot m_{\Theta}(A). \end{aligned}$$

Using Equation 1 we have

$$\begin{aligned} E(x) &= \sum_{x \in \Theta} x \cdot p_{\Theta}(x) \\ &= \sum_{x \in \Theta} x \cdot \left(\sum_{A: \sup(A)=x} \rho \cdot m_{\Theta}(A) + \sum_{A: \inf(A)=x} (1 - \rho) \cdot m_{\Theta}(A) \right) \\ &= \sum_{x \in \Theta} \left(\sum_{A: \sup(A)=x} \sup(A) \cdot \rho \cdot m_{\Theta}(A) + \sum_{A: \inf(A)=x} \inf(A) \cdot (1 - \rho) \cdot m_{\Theta}(A) \right). \end{aligned}$$

The double summations can be collapsed to a single summation because every $A \subseteq \Theta$ has a unique $\sup(A) \in \Theta$ and a unique $\inf(A) \in \Theta$.

$$\begin{aligned} E(x) &= \sum_{A \subseteq \Theta} \sup(A) \cdot \rho \cdot m_{\Theta}(A) + \inf(A) \cdot (1 - \rho) \cdot m_{\Theta}(A) \\ &= \sum_{A \subseteq \Theta} \inf(A) \cdot m_{\Theta}(A) + \rho \cdot \sum_{A \subseteq \Theta} [\sup(A) - \inf(A)] \cdot m_{\Theta}(A) \\ &= E_{\star}(x) + \rho \cdot (E^{\star}(x) - E_{\star}(x)). \end{aligned}$$

□

The important point of the proof is that the probabilistic analysis provides a meaningful way to choose a distinguished point within an EVI that results from the use of belief functions. That distinguished point can then be used as the basis for comparison of several choices when their respective EVIs overlap.

2.4 Discussion

Because of its interval representation of belief, Shafer's theory poses difficulties for a decision-maker who uses it. Lesh has proposed a different method for choosing a distinguished point

to use in the ordering of overlapping choices [10]. Lesh makes use of an empirically derived “ignorance preference coefficient,” τ , that is used to compute the distinguished point called “expected evidential belief” (EEB):

$$EEB(A) = \frac{Spt(A) + Pls(A)}{2} + \tau \frac{(Pls(A) - Spt(A))^2}{2}.$$

A choice is made by choosing the action that maximizes the “expected evidential value” (EEV):

$$EEV = \sum_{A_i \subseteq \Theta} A_i \cdot EEB(A_i).$$

There are some important differences between Lesh’s approach and the present approach for evidential decision-making. The ignorance preference parameter τ can be seen as a means for interpolating a distinguished value within a *belief* interval $[Spt(A), Pls(A)]$, while the cooperation probability, ρ , is used to interpolate within an interval of *expected utilities* $[E_*(x), E^*(x)]$. Secondly, Lesh’s parameter τ is empirically derived and has no theoretical underpinning. In contrast, the cooperation parameter ρ has been explained as a probability of a comprehensible event—that the residual ambiguity will be favorably resolved. It leads to a simple procedure involving linear interpolation between bounds of expected utility, and is derived from probability theory.

The use of a single parameter to choose a value between two extremes is similar in spirit to the approach taken by Hurwicz with a probabilistic formulation [6]. Hurwicz suggested that rather than computing the expected utility of a variable for which a probability distribution is known, one could interpolate a decision index between two extremes by estimating a single parameter related to the disposition of nature. When this parameter is zero, one obtains the Wald minimax criterion—the assumption that nature will act as strongly as possible against the decision-maker [22]. In contrast to the Hurwicz approach in which one ignores the probability distribution and computes a decision index on the basis of the parameter only, in our approach the expected utility interval is computed, and interpolation between extremes occurs only within the range of residual ambiguity allowed by the focal elements of a belief function. Thus our approach is identical to the use of expected utilities when a probability distribution is available; it is identical to Hurwicz’s approach when there are known constraints on the distribution; and it combines elements of both when the distribution is a belief function.

There may be circumstances in which a single parameter is insufficient to capture the underlying structure of a decision problem. In these cases it would be more appropriate to use a different probability to represent the attitude of nature for each source of ambiguity. Let ρ_i be the probability that ambiguity within each focal element A_i will be decided favorably, $(\forall A_i) A_i \subseteq \Theta$. Then we obtain

$$E(x) = \sum_{A_i \subseteq \Theta} \inf(A_i) \cdot m_{\Theta}(A_i) + \sum_{A_i \subseteq \Theta} \rho_i [\sup(A_i) - \inf(A_i)] \cdot m_{\Theta}(A_i) \quad (5)$$

in place of Equation 4.

3 Decision Analysis

In the preceding section we have defined the concept of an expected utility interval for belief functions and we have shown that it bounds the expected utility that would be obtained with any probability distribution consistent with that belief function. Furthermore, we have proposed a parameter (the probability that residual ambiguity will be decided in our behalf) that can be used as the basis for computing a unique expected utility when the available evidence warrants only bounds on that expected utility. In this section we will show how the expected utility interval can be used to generalize probabilistic decision analysis.

Decision analysis was first developed as a means by which one could organize and systematize one's thinking when confronted with an important and difficult choice [4, 14]. Its formal basis has made it adaptable as a computational procedure by which computer programs can choose actions when provided with all relevant information. Simply stated, the analysis of a decision problem under uncertainty entails the following steps:

- List the viable options available for gathering information, for experimentation, and for action.
- List the events that may possibly occur.
- Arrange the information you may acquire and the choices you may make in chronological order.
- Decide the value to you of the consequences that result from the various courses of action open to you.
- Judge the chances that any particular uncertain event will occur.

3.1 Decision analysis using probabilities

First we will illustrate the use of decision analysis on a problem that can be represented with probabilities to acquaint the reader with the method and terminology.

Example – Oil Drilling #1 A wildcatter must decide whether or not to drill for oil. He is uncertain whether the hole will be dry, have a trickle of oil, or be a gusher. Drilling a hole costs \$70,000. The payoffs for hitting a gusher, a trickle, or a dry hole are \$270,000, \$120,000, and \$0, respectively. At a cost of \$10,000 the wildcatter could take seismic soundings that would help determine the underlying geologic structure. The soundings will determine whether the terrain has no structure, open structure, or closed structure. The experts have provided us with the joint probabilities shown below. We are to determine the optimal strategy for experimentation and action [8].

State	No struct	Open	Closed	Marginal
Dry	0.30	0.15	0.05	0.50
Trickle	0.09	0.12	0.09	0.30
Gusher	0.02	0.08	0.10	0.20
Marginal	0.41	0.35	0.24	1.00

In decision analysis, a decision tree is constructed that captures the chronological order of actions and events [8, 9]. A square is used to represent a decision to be made, and its branches are labeled with the alternative choices. A circle is used to represent a chance node, and its branches are labeled with the conditional probability of each event, given that the choices and events along the path leading to the node have occurred.

To compute the best strategy, the tree is evaluated from its leaves toward its root.

- The value of a leaf node is the utility of the state of nature it represents.
- The value of a chance node is the expected utility of the probability distribution represented by its branches as computed using Equation 1.
- The value of a choice node is the maximum of the utilities of each of its sons. The best choice for the node is denoted by the branch leading to the son with the greatest utility. Ties are broken arbitrarily.

This procedure is repeated until the root node has been evaluated. The value of the root node is the expected utility of the decision problem; the branches corresponding to the maximal value at each choice node give the best *strategy* to follow (i.e. choices to make in each situation).

The evaluated decision tree for the oil drilling example is portrayed in Figure 4. It can be seen that the expected value is \$22,500 and that the best strategy is to take seismic soundings, to drill for oil if the soundings indicate open or closed structure, and not to drill if the soundings indicate no structure.

3.2 Decision analysis using belief functions

To use the decision procedure just described, it must be possible to assess the probabilities of all uncertain events. That is, the set of branches emanating from each chance node in the decision tree must depict a probability distribution. In many scenarios, however, estimating these probability distributions is difficult or impossible, and the decision-maker is forced to assign probabilities even though he knows they are unreliable. Using belief functions, one need not estimate any probabilities that are not readily available. The representation better reflects the evidence at hand, but the decision analysis procedure cannot be used with the resulting interval representation of belief. In this section we describe a generalization of decision analysis that accommodates belief functions.

Example – Oil Drilling #2 As in the first oil-drilling example, a wildcatter must decide whether or not to drill for oil. His costs and payoffs are the same as before: drilling costs \$70,000, and the payoffs for hitting a gusher, a trickle, or a dry well are \$270,000, \$120,000, and \$0, respectively. However, at this site, no seismic soundings are available. Instead, at a cost of \$10,000, the wildcatter can make an electronic test that is related to the well capacity as shown below. We are to determine the optimal strategy for experimentation and action.

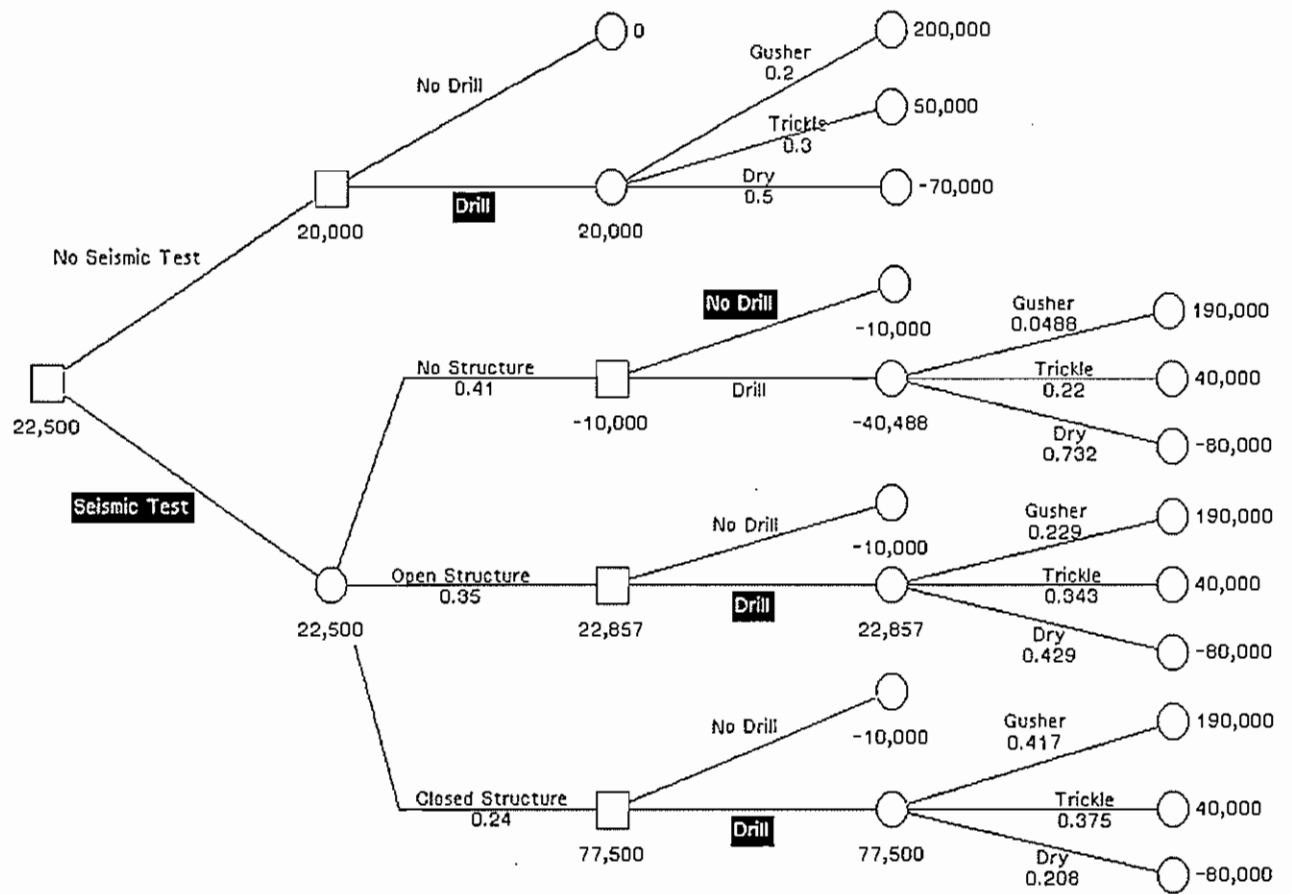


Figure 4: Decision Tree for First Oil-Drilling Example

Prob	Test result	Capacity
0.5	red	dry
0.2	yellow	dry or trickle
0.3	green	trickle or gusher

Several issues arise that prevent one from constructing a well-formed decision tree for this example. First, consider the branch of the tree in which the test is conducted and the result is green (Figure 5). If we drill for oil, then we know we will find either a trickle or a gusher, but we cannot determine the probability of either from the given information. We are tempted to label the branch with the disjunction (Trickle \vee Gusher) with probability 1.0. But what should be the payoff of that branch? All we can say is that the payoff will be either \$40,000 (if a trickle) or \$190,000 (if a gusher). Ordinary decision analysis requires a unique value to be assigned, but we have no basis for computing one. So the first modification we make to the construction of decision trees is to allow disjunctions of events on branches emanating from chance nodes, and to allow intervals as the payoffs for leaf nodes. We will discuss later how to evaluate such a tree.

To see the second issue, consider the branch of the tree in which the test is not conducted. If we drill for oil, there is a chance that we will hit a gusher, a trickle, or a dry well, but what is the probability distribution? We know only that

$$\begin{array}{ll}
 p(\text{Dry} \mid \text{Red}) = 1.0 & p(\text{Red}) = 0.5 \\
 p(\text{Dry} \vee \text{Trickle} \mid \text{Yellow}) = 1.0 & p(\text{Yellow}) = 0.2 \\
 p(\text{Trickle} \vee \text{Gusher} \mid \text{Green}) = 1.0 & p(\text{Green}) = 0.3
 \end{array}$$

There is not enough information to use Bayes' rule to compute the probability distribution for the well capacity. Without adding a new assumption at this point, the strongest statement that can be made is

$$\begin{array}{l}
 0.5 \leq p(\text{Dry}) \leq 0.7 \\
 0.0 \leq p(\text{Trickle}) \leq 0.5 \\
 0.0 \leq p(\text{Gusher}) \leq 0.3
 \end{array}$$

Using belief functions, this can be represented as

$$\begin{array}{l}
 m(\{\text{Dry}\}) = 0.5 \\
 m(\{\text{Dry}, \text{Trickle}\}) = 0.2 \\
 m(\{\text{Trickle}, \text{Gusher}\}) = 0.3,
 \end{array}$$

which yields the required belief intervals

$$\begin{array}{ll}
 [Spt(\{\text{Dry}\}), Pls(\{\text{Dry}\})] & = [0.5, 0.7] \\
 [Spt(\{\text{Trickle}\}), Pls(\{\text{Trickle}\})] & = [0.0, 0.5] \\
 [Spt(\{\text{Gusher}\}), Pls(\{\text{Gusher}\})] & = [0.0, 0.3].
 \end{array}$$

The second modification we make to decision trees is to allow the branches emanating from a chance node to represent a mass function. The masses must still sum to one, but the events

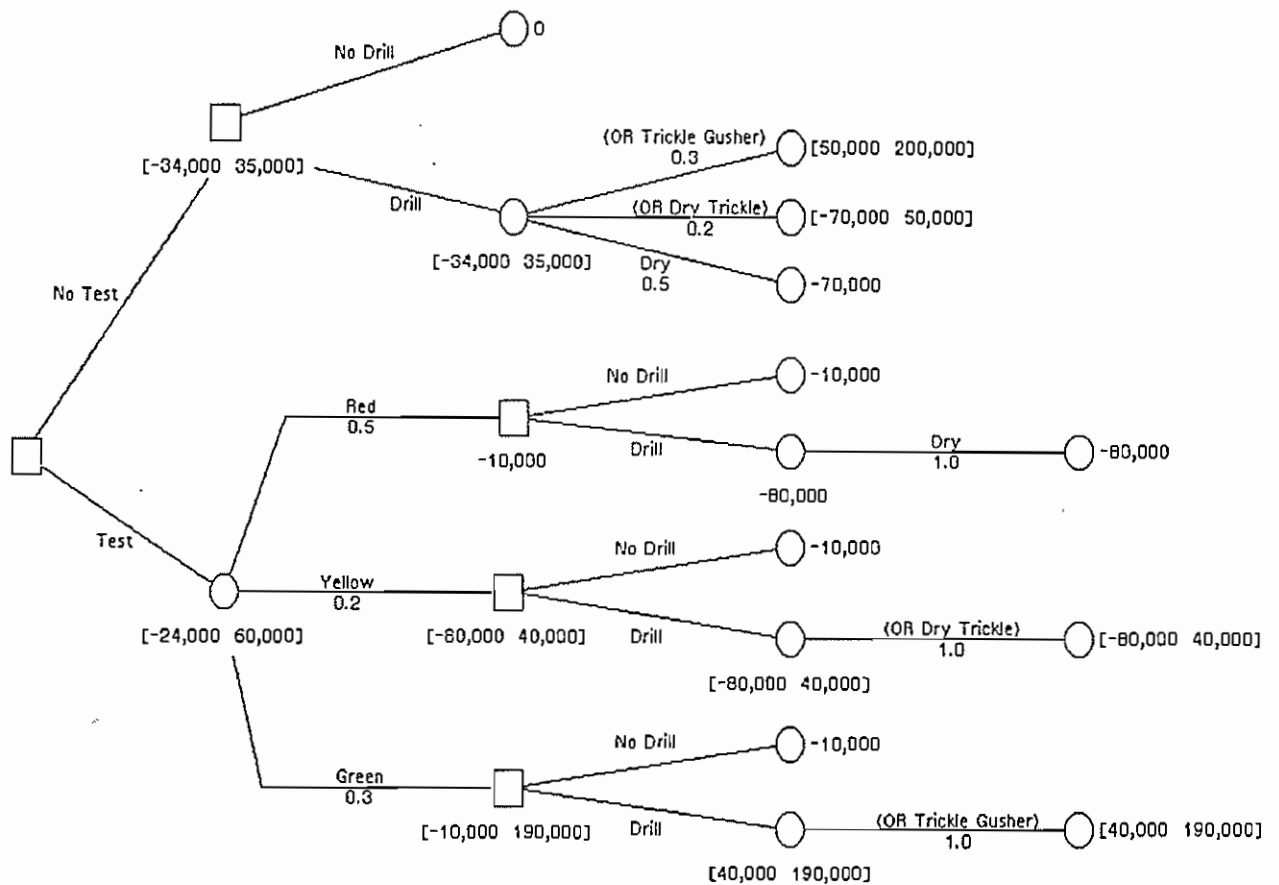


Figure 5: Modified Decision Tree for the Second Oil Drilling Example

need not be disjoint.³ The completed decision tree for Oil Drilling Example #2 is shown in Figure 5.

The tools of Section 2 can be used to evaluate a decision tree modified in this manner.

- The value of a leaf node is the utility of the state of nature it represents. This may be a unique value or, in the case of a disjunction of states, an interval of values.
- A chance node represents a belief function. Its value is the expected utility interval computed with Equation 3 :

$$E(x) = [E_*(x), E^*(x)] .$$

- A decision node represents a choice of the several branches emanating from it. The utility of each branch may be a point value or an interval. The value of a decision node is the expected utility computed using Equation 4 and an estimate of ρ :

$$E(x) = E_*(x) + \rho \cdot (E^*(x) - E_*(x)) .$$

The action on the branch that yields the greatest $E(x)$ is chosen. Ties are broken arbitrarily.

In summary, a decision tree and decision analysis procedure for belief functions have been described. Two modifications were made to adapt ordinary decision trees: intervals are allowed where utilities occur; and belief functions are allowed where probability distributions occur. A unique strategy can be obtained by estimating the probability ρ .⁴

3.3 Generalized decision tree examples

Figures 6, 7, and 8 show the evaluated decision tree for several values of ρ — each node is labeled with its expected value or expected value interval. In the cases where the expected value is an interval, the evidential expected value $E(x)$ is also shown (using the assumed ρ). Preferred decisions are highlighted with a black background.

If we opt not to test, then our choice is either to not drill (expected value 0) or to drill (expected value interval $[-34,000 \ 35,000]$). The better choice depends on what value of ρ is assumed. As can be seen in the figures, if $\rho = 0.0$, then it is better to not drill, but if $\rho = 0.5$ or $\rho = 1.0$, then drilling is the better choice.

If we choose to test and the result is yellow, then our choice is to not drill (expected value -10,000) or to drill (expected value interval $[-80,000 \ 40,000]$). In this case it is better to not drill if either $\rho = 0.0$ or $\rho = 0.5$ and to drill if $\rho = 1.0$.

If the test result is red, then one should not drill regardless of ρ (-10,000 is always better than -80,000). If the test result is green, then one should always drill (-10,000 is never as good as the interval $[40,000 \ 190,000]$).

³Recall that a probability distribution is an assignment of belief over mutually exclusive elements of a set, whereas a mass function is a distribution over possibly overlapping subsets.

⁴When all utilities are point-valued and all belief functions are true probability distributions, no assumption is required and the strategy will be identical to that prescribed by ordinary decision analysis.

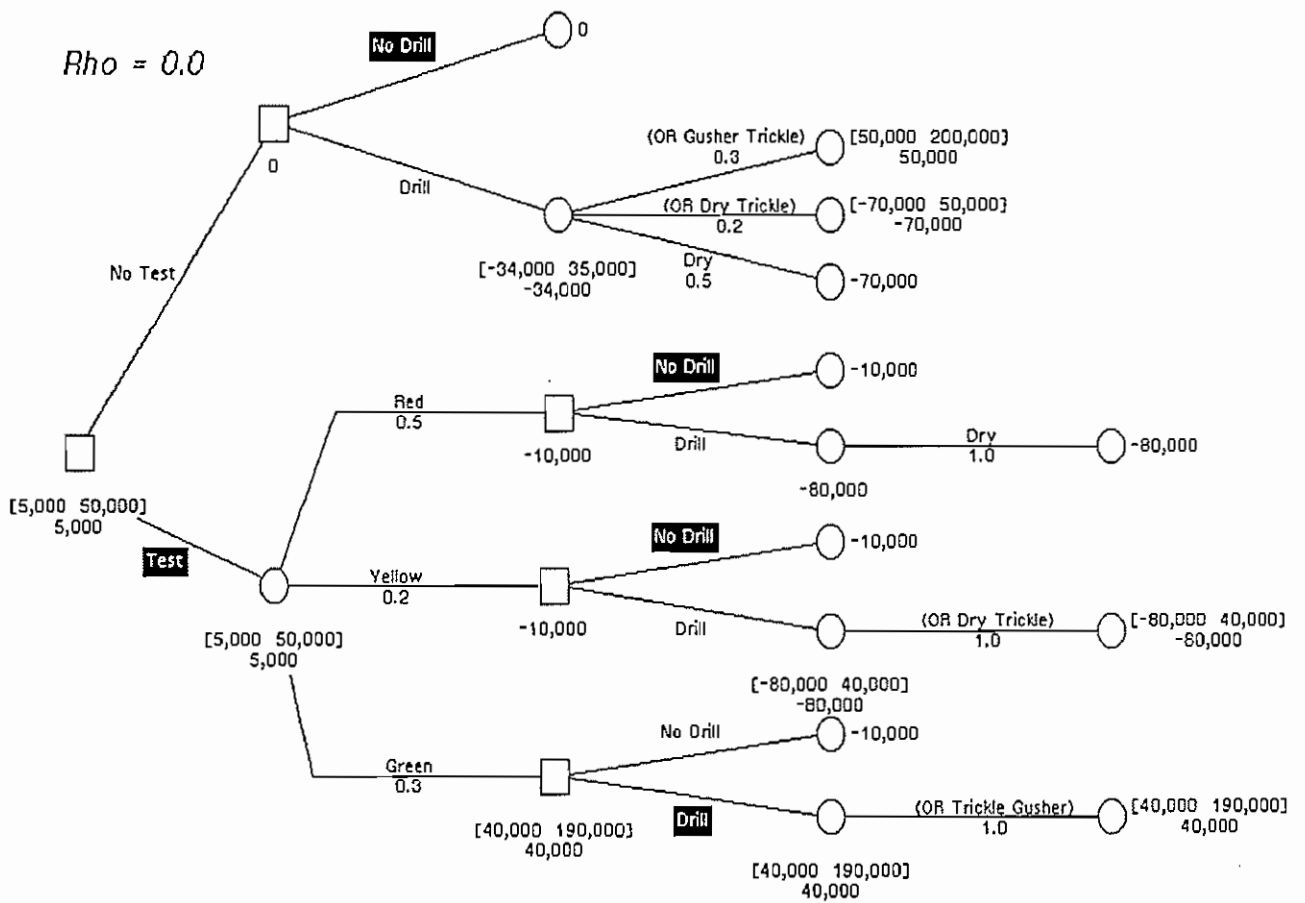


Figure 6: Decision Tree for the Second Oil Drilling Example (assuming $\rho = 0.0$).




Figure 7: Decision Tree for the Second Oil Drilling Example (assuming $\rho = 0.5$).




Figure 8: Decision Tree for the Second Oil Drilling Example (assuming $\rho = 1.0$).

3.4 Comparing two choices

Instead of assuming a value for ρ first, and calculating the choices that result, one may ask the reverse question. At what value of ρ would I change my decision? This can be answered in general by examining a choice between two states having expected utility intervals.

Theorem 2 *Let the expected utility intervals of two choices be as follows:*

$$\begin{aligned} \text{Choice 1: } & [E_{1*}(x), E_1^*(x)] \\ \text{Choice 2: } & [E_{2*}(x), E_2^*(x)]. \end{aligned}$$

Assume without loss of generality that Choice 1 has the smaller interval, i.e. $(E_2^(x) - E_{2*}(x)) > (E_1^*(x) - E_{1*}(x))$. Then Choice 2 is preferred over Choice 1 iff*

$$\rho > \frac{E_{1*}(x) - E_{2*}(x)}{E_2^*(x) - E_1^*(x) + E_{1*}(x) - E_{2*}(x)}. \quad (6)$$

Proof:

Using Theorem 1 and solving for ρ gives the point ρ_c at which one is indifferent between Choice 1 and Choice 2:

$$\begin{aligned} E_1(x) &= E_{1*}(x) + \rho \cdot (E_1^*(x) - E_{1*}(x)) \\ E_2(x) &= E_{2*}(x) + \rho \cdot (E_2^*(x) - E_{2*}(x)) \\ \rho_c &= \frac{E_{1*}(x) - E_{2*}(x)}{(E_2^*(x) - E_1^*(x)) + (E_{1*}(x) - E_{2*}(x))}. \end{aligned} \quad (7)$$

The expected value of both choices at ρ_c is

$$E_c(x) = \frac{E_{1*}(x) \cdot E_2^*(x) - E_1^*(x) \cdot E_{2*}(x)}{E_{1*}(x) - E_{2*}(x) + E_2^*(x) - E_1^*(x)}.$$

Now consider the choice at $\rho = \rho_c + \delta$ where $\delta > 0$:

$$E_1(x) = E_c(x) + \delta \cdot (E_1^*(x) - E_{1*}(x)) \quad (8)$$

$$E_2(x) = E_c(x) + \delta \cdot (E_2^*(x) - E_{2*}(x)). \quad (9)$$

Since $(E_2^*(x) - E_{2*}(x)) > (E_1^*(x) - E_{1*}(x))$ and $\delta > 0$, it must be the case that $E_2(x) > E_1(x)$. Therefore, Choice 2 is preferred. Similar argument shows that Choice 1 is preferred whenever $\rho < \rho_c$.

□

Letting

$$a = E_{1*}(x) - E_{2*}(x) \text{ and } b = E_2^*(x) - E_1^*(x)$$

gives

$$\rho_c = \frac{a}{a + b}. \quad (10)$$

Thus, Choice 1 is preferable if

$$\rho < \frac{a}{a+b},$$

and Choice 2 is preferable if

$$\rho > \frac{a}{a+b}.$$

If $\frac{a}{a+b} > 1.0$ then Choice 1 is always preferred (no assumption of ρ is necessary). If $\frac{a}{a+b} < 0.0$ then Choice 2 is always preferred. It follows that whenever one EUI is slightly “higher” than another, i.e.

$$E_{1*}(x) > E_{2*}(x) \quad \text{and} \quad E_1^*(x) > E_2^*(x),$$

then the action that gives rise to it is always preferred.

Returning to the second oil-drilling example (Figure 5), the decision of whether or not to drill when the test result is yellow involves a choice between

$$\begin{aligned} \text{No Drill: } E(x) &= [-10,000 \quad -10,000] \\ \text{Drill: } E(x) &= [-80,000 \quad 40,000]. \end{aligned}$$

By Theorem 2, $\rho_c = 0.583$, and one should drill only if $\rho > 0.583$.

When $\rho > 0.583$, the decision as to whether or not to conduct the test involves a choice between

$$\begin{aligned} \text{No Test: } E(x) &= [-34,000 \quad 35,000] \\ \text{Test: } E(x) &= [-9,000 \quad 60,000]. \end{aligned}$$

Here, Test is the preferred choice because its EUI is higher.

4 Discussion

The value of the result of an action is frequently measured in money (e.g., in dollars), but people often exhibit preferences that are not consistent with maximization of expected monetary value. The theory of *utility* accounts for this behavior by associating for an individual decision-maker a value (measured in *utils*) with each state s , $u = f(s)$, such that maximization of expected utility yields choices consistent with that individual’s behavior [4]. Utility theory can satisfactorily account for a person’s willingness to expose himself to risk and should be used whenever one’s preferences are not linearly related to value. This attitude toward risk should not be confused with one’s attitude towards ambiguity, which is the quality that is modeled by ρ .

4.1 On making assumptions

It is interesting to compare the types of assumptions made in a probabilistic analysis with the ρ assumption proposed here for belief functions. When using probability, a maximum entropy assumption is often made. Sometimes, this assumption is justified, and it should properly be

considered part of the evidence, not as an assumption. When this is the case, a maximum entropy belief function can be used as well [2]. At other times, the maximum entropy assumption is not justified, but is used simply because *some* assumption must be made, and maximum entropy has some desirable properties [20]. In these cases, the choice of elements in the sample space (the set of possibilities) introduces distortion into the expected value that will result. That is, adding a few more possibilities into the sample space will change the expected value of the maximum entropy distribution over that sample space. For example, if we choose to allow for the possibility of \$2 being among the possibilities for the hidden sector of Carnival Wheel #2, the sample space would be $\{1,2,5,10,20\}$ instead of $\{1,5,10,20\}$, and the expected value of the maximum entropy distribution of that wheel would be \$6.16 instead of \$6.30. On the other hand, for any choice of ρ , the evidential expected value using either of the two preceding sample spaces would be identically $(5.50 + 1.90\rho)$ dollars. Of course, adding possibilities outside the interval $[1, 20]$ would change the evidential expected value. For example, allowing for the possibility of \$50 in the hidden sector would change the maximum entropy expected value to \$7.12 and would change the evidential expected value to $(5.50 + 4.90\rho)$ dollars. The point is that both assumptions introduce bias into the decision criteria. This should not be surprising because *both* are unjustified assumptions. There is no basis on which to prefer one over the other; both assumptions are entirely plausible.

Having made this point, there are some consequently weak arguments for recommending the use of the assumption of the probability of nature's cooperation ρ . Because the EUI spans the range of all expected utilities that could be obtained by adding *any* assumption to a probabilistic analysis, there always exists some value of ρ , $0 \leq \rho \leq 1$ that yields the same expected utility $E(x)$ as a probabilistic analysis. Therefore, the decisions that are prescribed depend only on one's ability to estimate ρ , not on his election to use Equation 3. Furthermore, the use of a single parameter means that the decision-maker is asked to provide only one additional piece of information.

The parameter ρ has been explained as a probability, giving it a formal grounding that earlier decision schemes for belief functions have lacked. Furthermore, we believe that it is the probability of a meaningful event. Selecting $\rho = 0$ is appropriate when an adversary controls the situation (as in game playing, for example) or when a decision-maker wishes only to minimize his expected loss, and is equivalent to the maximin criteria of Wald. An optimistic decision-maker would prefer to choose $\rho = 1$ to maximize his chance of realizing the greatest possible expected payoff without worrying about what losses might be possible. Intermediate values of ρ can be used to compromise between these extremes.

4.2 On the limitations of the approach

Despite the appeal of a computationally efficient decision analysis procedure for belief functions, there remain some issues that are not addressed. As in classical decision analysis, it remains necessary to enumerate the potential states of nature and to assign utilities (actually utility intervals, which should be easier to assign in practice). This task can be overwhelming when complex scenarios are considered. Furthermore, it should not be forgotten that the

assignment of a value to ρ (when it is necessary) remains an assumption unwarranted by the evidence at hand, just as maximum entropy or any other assumption is unwarranted when insufficient information is available.

It is inherent in the methodology described that the determination of what is best or worst is considered *after* the decision-maker's choice is postulated. That is, the reaction of nature is allowed to depend on the decision that is to be taken. This is sometimes reasonable, and sometimes not. For example, conducting a regional test market for a new product may affect the national demand by virtue of publicity or increased competition. As a result, there may be no single underlying probability distribution that can simultaneously give rise to the expected utilities obtained for each choice. This should not be particularly worrisome as long as this consideration suits the problem at hand. If not, the expected utility intervals computed with the method described here may be wider than prescribed by the evidence. In that case, it is necessary to conduct a more complicated case-based analysis that is analogous to the linear-programming problems that arise in game theory. See Jaffray [7] for further discussion of this approach.

4.3 On the automation of decision analysis

A probabilistic analysis of a decision problem (e.g. the second oil-drilling example) follows the paradigm: assess, assume, combine, decide. An *assessment* of a probability distribution is made for each piece of evidence; *assumptions* are made about the distributions of missing pieces of evidence; the assumptions and evidence are *combined* to obtain a distribution of payoffs, and a *decision* is made on the basis of the expected utility of the payoff. In contrast, a belief function analysis follows the paradigm: assess, combine, assume, decide. An *assessment* of a belief function is made for each piece of evidence; these pieces of evidence are then *combined* to obtain a belief function over the possible payoffs; then an *assumption* is made (about the benevolence of nature); and a *decision* is made using that assumption and the expected utility interval of the payoffs.

While the same decisions will be reached whether one makes assumptions first and then combines or combines evidence and then adds those assumptions, the difference in paradigms has important implications for automating the procedure. First, in some decision problems the EU of the top choice will not overlap the EU of any other choice, i.e. the decision follows from what is truly known, and in no way depends upon the accuracy of any assumption that might be made. Using belief functions, the best decision in this case is immediately determinable without additional assumptions. Because Bayes' rule requires a prior distribution, this situation cannot be recognized without a more complex sensitivity analysis when a purely probabilistic representation is used. Second, when an assumption must be made because intervals do overlap, making it as late as possible allows one to maintain the assumption-free intermediate calculations for use in other computations. This is not an issue when the evidence will be used once and discarded, but affords a considerable computational savings when other decisions must be based on the original evidence plus new evidence as it comes along. Third, consider what must be computed if one chooses to use a different

assumption (as needed for sensitivity analysis, for example). In a probabilistic analysis the assumptions and all evidence must be recombined before a decision can be made because the assumptions are needed to combine the evidence. Using belief functions, one need only combine the new assumption with the already combined evidence before selecting the decision. This separation of evidence and assumption is similar in spirit to the distinction between credal and pignistic beliefs described by Smets [19].

5 Summary

We have proposed a decision analysis methodology for Shafer's theory of belief functions. We started by defining the notion of expected utility interval (EUI) and showed it to properly bound the expected utility of any probability distribution that could be obtained by introducing additional assumptions. Because an expected utility interval is often insufficient for decision-making, we recognize that a point-value must be chosen to compare alternative choices. We then showed how a linear interpolation of a distinguished value within the EUI is equivalent to making an assumption of the benevolence or maleficence of nature. Letting ρ be the probability that ambiguity will be resolved favorably, we derived that distinguished point.

We have also shown how the theory can be used to generalize the decision trees used in probabilistic decision analysis. These tools allow a decision-maker to defer unwarranted assumptions until the latest possible moment. In so doing he can sometimes avoid making any assumptions at all. Otherwise, he is forced to provide only enough additional information to allow a clear choice, and has the benefit of all available information to selectively decide where he would like to make that assumption.

We have implemented the techniques and have used that software to generate the decision trees shown in the figures in this paper. In addition a new evidential operator for decision-making has been added to the repertoire of the evidential reasoning technology developed at SRI International (see Appendix B). Decision analysis has been incorporated into Gister,⁵ SRI's evidential reasoning system which uses the Dempster-Shafer theory of belief functions as its underlying representation.

What we have described is by no means a full theory of decision-making for belief functions. Rather, we hope it may provide some insight that will someday lead to a better understanding of decision-making with incomplete information.

Acknowledgment

I thank John Lowrance, Tom Garvey, Leonard Wesley, and Enrique Ruspini for their contributions to the ideas presented in this paper. I am also grateful for the close scrutiny provided by Philippe Smets, who called attention to the difference between this approach and the assumption of a single underlying probability distribution.

⁵Gister is a trademark of SRI International.

References

- [1] Barnett, Jeffrey A., "Computational Methods for a Mathematical Theory of Evidence," *Proceedings VII IJCAI*, Vancouver, Canada, August, 1981, pp. 868-875.
- [2] Dubois, Didier, and Prade, Henri, "On Several Representations of an Uncertain Body of Evidence," in Gupta, M. M., and Sanchez, E., (eds.) *Fuzzy Information and Decision Processes*, North Holland Publishing Company, Amsterdam, 1982, pp. 167-181.
- [3] Horvitz, Eric J., John S. Breese, and Max Henrion, "Decision Theory in Expert Systems and Artificial Intelligence," *International Journal of Approximate Reasoning*, Vol 2, No 3, July 1988, pp. 247-302.
- [4] Howard, R. A., "Decision Analysis: Applied Decision Theory," in *Proceedings of the 4th International Conference on Operational Research*, D. B. Hertz and J. Melese, eds., 1966, pp. 55-77.
- [5] Howard, Ronald A., "The Science of Decision-Making," in R. A. Howard, J.E. Matheson, and K.L. Miller (ed.), *Readings in Decision Analysis*, SRI International, Menlo Park, California, Chapter 6, pp. 147-164.
- [6] Hurwicz, Leonid, "A Criterion for Decision-Making under Uncertainty," Technical Report 355, Cowles Commission, 1952.
- [7] Jaffray, Jean-Yves, "Application of Utility Theory to Belief Functions," in Bouchon, B., Saitta, L., and Yager, R., (eds.) *Uncertainty and Intelligence Systems*, Springer-Verlag, Berlin, 1988, pp. 1-8.
- [8] Lapin, Lawrence L., *Quantitative Methods for Business Decisions*, Harcourt Brace Jovanovich, Inc, New York, 1981.
- [9] LaValle, Irving H., *An Introduction to Probability, Decision, and Inference*, Holt, Rinehart and Winston, Inc., New York, 1970.
- [10] Lesh, Stephen A., "An Evidential Theory Approach to Judgment-Based Decision Making," Ph.D. Thesis, Department of Forestry and Environmental Studies, Duke University, December, 1986.
- [11] Loui, Ronald P., Jerome A. Feldman, and Henry G. E. Kyburg, Jr. "Interval-Based Decisions for Reasoning Systems," *Proceedings of the Workshop on Uncertainty and Probability in Artificial Intelligence*, UCLA, California, August 1985, pp. 193-199.
- [12] Lowrance, John D., Thomas D. Garvey, and Thomas M. Strat, "A Framework for Evidential-Reasoning Systems," *Proceedings AAAI-86*, Philadelphia, Pennsylvania, August 1986.
- [13] Pittarelli, Michael, "Decision Making with Linear Constraints on Probabilities," in *Proceedings of the Fourth Workshop on Uncertainty in Artificial Intelligence*, University of Minnesota, August 1988, pp. 283-290.

- [14] Raiffa, Howard, *Decision Analysis: Introductory Lectures on Choices under Uncertainty*, Addison-Wesley Publishing Company, Reading, Massachusetts, 1970.
- [15] Ruspini, Enrique H., "The Logical Foundations of Evidential Reasoning," SRI International Technical Note 408, December, 1986.
- [16] Shafer, Glenn A., *A Mathematical Theory of Evidence*, Princeton University Press, New Jersey, 1976.
- [17] Shafer, Glenn A., "Constructive Decision Theory," University of Kansas Department of Mathematics Working Paper, December, 1982.
- [18] Smets, Philippe, "Belief Functions," Chapter 9 in *Non-Standard Logics for Automated Reasoning*, P. Smets, E. Mamdani, D. Dubois, and H. Prade, eds., Academic Press, London, 1988.
- [19] Smets, Phillippe, "Decision under Uncertainty," *Proceedings of the Workshop on Uncertainty in Artificial Intelligence*, August 1989.
- [20] Smets, Philippe, and Kennes, Robert, "The Transferable Belief Model: Comparison with Bayesian Models," to be published.
- [21] Strat, Thomas M., "Continuous Belief Functions for Evidential Reasoning," *Proceedings, AAAI-84*, Austin, Texas, August 1984.
- [22] Wald, Abraham, *Statistical Decision Functions*, John Wiley, New York, 1950.
- [23] Yager, Ronald R., "Decision Making under Dempster-Shafer Uncertainties," Iona College Machine Intelligence Institute Tech. Report #MII-915, undated.

A Notation

- $p_{\Theta}(x)$ –Probability distribution over sample space $\Theta, x \in \Theta$
 $m_{\Theta}(A)$ –Mass function defined over frame of discernment $\Theta, A \subseteq \Theta$
 $Spt(A)$ –Support: $Spt(A) = \sum_{A_i \subseteq A} m_{\Theta}(A_i)$
 $Pls(A)$ –Plausibility: $Pls(A) = 1 - Spt(\neg A)$
 $E(x)$ –Expected value of a random variable whose outcome is governed by a probability distribution:

$$E(x) = \sum_{x \in \Theta} x \cdot p_{\Theta}(x)$$

- Evidential expected value – the expected value of a variable governed by a belief function assuming that any residual ambiguity will be decided favorably with probability ρ :

$$E(x) = (1 - \rho) \cdot E_{\star}(x) + \rho \cdot E^{\star}(x)$$

- $E^{\star}(x)$ –Upper bound of expected value:

$$E^{\star}(x) = \sum_{A_i \subseteq \Theta} \sup(A_i) \cdot m_{\Theta}(A_i)$$

- $E_{\star}(x)$ –Lower bound of expected value:

$$E_{\star}(x) = \sum_{A_i \subseteq \Theta} \inf(A_i) \cdot m_{\Theta}(A_i)$$

- EVI –Expected value interval: $[E_{\star}(x), E^{\star}(x)]$
 EUI –Expected utility interval: Same as EVI , when Θ is a frame of utilities
 ρ –The probability that any residual ambiguity will be decided favorably
 $1 - \rho$ –The probability that any residual ambiguity will be decided unfavorably
 ρ_c –The value of ρ at which one would be indifferent between two choices

B Decision-Making with Evidential Reasoning

In this section we reanalyze the second oil-drilling example within the framework of evidential reasoning. First we review some of the tools of evidential reasoning and then introduce a decision operator for belief functions based on the theory described earlier.

In evidential reasoning, domain-specific knowledge is defined in terms of *compatibility relations* that relate one frame of discernment to another. A compatibility relation simply describes which elements from the two frames can simultaneously be true. A compatibility relation $\Theta_{A,B}$ between two frames Θ_A and Θ_B is a set of pairs such that

$$\Theta_{A,B} \subseteq \Theta_A \times \Theta_B ,$$

where every element of Θ_A and every element of Θ_B is included in at least one pair.

Evidential reasoning provides a number of formal operations for assessing evidence, including:

- **Fusion** — to determine a consensus from several bodies of evidence obtained from independent sources. Fusion is accomplished through Dempster's rule of combination:

$$m_{\Theta}^3(A_h) = \frac{1}{1 - k} \sum_{A_i \cap A_j = A_h} m_{\Theta}^1(A_i) m_{\Theta}^2(A_j), \quad (11)$$

$$k = \sum_{A_i \cap A_j = \phi} m_{\Theta}^1(A_i) m_{\Theta}^2(A_j) .$$

Dempster's Rule is both commutative and associative (meaning evidence can be fused in any order) and has the effect of focusing belief on those propositions that are held in common.

- **Translation** — to determine the impact of a body of evidence upon elements of a related frame of discernment. The *translation* of a belief function from frame Θ_A to frame Θ_B using the compatibility relation $\Theta_{A,B}$ is defined by

$$m_{\Theta_B}(B_j) = \sum_{\substack{C_{A \mapsto B}(A_k) = B_j \\ A_k \subseteq \Theta_A, B_j \subseteq \Theta_B}} m_{\Theta_A}(A_k), \quad (12)$$

where $C_{A \mapsto B}(A_k) = \{b_j | (a_i, b_j) \in \Theta_{A,B}, a_i \in A_k\}$.

Several other evidential operations have been defined and are described elsewhere [12].

We now describe a new evidential operation for making decisions. Its operation is analogous to the evaluation of a choice node in probabilistic decision analysis, except that it is defined for belief functions and substitutes the notion of evidential expected utility for the probabilistic expected utility.

- **Decision** — to choose an action based on a body of evidence representing the states of nature believed at the time of the decision and a body of evidence representing the beliefs resulting from any particular decision. Let m_{Θ_A} represent the beliefs in frame Θ_A about the state of nature at the time a decision is to be made. Let frame Θ_D be the possible actions that can be taken. Let $m_{\Theta_U}(U|A, D)$ represent the beliefs over the utility frame Θ_U that result from making decision D when A is true. The mass function representing the best *policy* is

$$m_{\Theta_A \times \Theta_D}(A, D) = \begin{cases} m_{\Theta}(A) & \text{if } E(m_{\Theta_U}(U|A, D)) > E(m_{\Theta_U}(U|A, D_i)), \\ & \forall D_i \in \Theta_D, D_i \neq D \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

Ties are broken arbitrarily. The optimal policy computed by the decision node is given by the focal elements (A, D) in $m_{\Theta_A \times \Theta_D}(A, D_0)$ such that if A is the most precise statement known to be true, then the best decision is D .

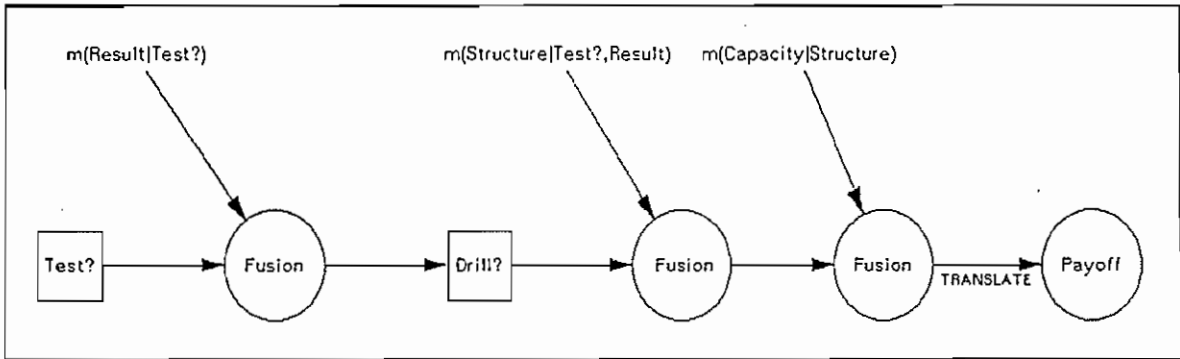


Figure 9: Evidential Decision Analysis of the Second Oil Drilling Example

Within Gister, a decision node is represented by a square. An evidence node leading into the decision node represents what would be known at the time a decision is to be made. The output of the decision node is the optimal policy (as defined above), and is represented as a belief function over the cross-product frame of states of nature and alternative decisions ($\Theta_A \times \Theta_D$). That belief function is then available to other evidential reasoning operations: it may be discounted, translated to a dependent frame, fused with additional evidence that would only be available after the decision is taken, etc. With this definition, a decision node represents a primitive operation that can be included in the data flow represented by an analysis. Figure 9 illustrates the analysis that was constructed for representing the second oil-drilling example within Gister.

The optimal policies computed by the Test? and Drill? nodes are summarized below (assuming that $\rho = 0.5$). The result is identical to the strategy computed using a decision tree as can be verified by comparison with Figure 7.

Decision Test?: Expected value: 27500 Rho=0.5
 1.0 = m((No Test Test)) -- Test
 E(x) = 27500 EVI = [5000 50000] Test
 E(x) = 500 EVI = [-34000 35000] No Test

Decision Drill?: Expected value: 27500 Rho=0.5
 0.50 = m((Test & Red)) -- No Drill
 E(x) = -10000 EVI = [-10000 -10000] No Drill
 E(x) = -80000 EVI = [-80000 -80000] Drill

0.20 = m((Test & Yellow)) -- Drill
 E(x) = -10000 EVI = [-10000 -10000] No Drill
 E(x) = 20000 EVI = [-80000 40000] No Drill

0.30 = m((Test & Green)) -- Drill
 E(x) = 115000 EVI = [40000 190000] Drill
 E(x) = -10000 EVI = [-10000 -10000] No Drill