

AD-A103 879

WISCONSIN UNIV-MADISON MATHEMATICS RESEARCH CENTER

F/G 12/1

THE SUBJECTIVIST VIEW OF DECISION-MAKING, (U)

JUL 81 D V LINDLEY


DAAG29-80-C-0041

UNCLASSIFIED

MRC-TSR-2248

NL

1-1
30-10-81



END

DATE

FILED

081

DTIC

AD A103879

MRC Technical Summary Report #2248

THE SUBJECTIVIST VIEW OF DECISION-MAKING

Dennis V. Lindley

Mathematics Research Center
University of Wisconsin-Madison
610 Walnut Street
Madison, Wisconsin 53706

July 1981

Received May 19, 1981

DTIC
ELECTE

SEP 8 1981

A

Approved for public release
Distribution unlimited

Sponsored by

U. S. Army Research Office
P. O. Box 12211
Research Triangle Park
North Carolina 27709

81 9 08 040

UNIVERSITY OF WISCONSIN - MADISON
MATHEMATICS RESEARCH CENTER

THE SUBJECTIVIST VIEW OF DECISION-MAKING

Dennis V. Lindley

Technical Summary Report #2248
July 1981

ABSTRACT

Accession For	
NTIS	GRAB
DTIC TAB	
Unannounced	
Justification	
By	
Distribution/	
Availability Codes	
Dist	Spec
A	

The subjectivist, Bayesian paradigm for a decision-maker is described. It is shown how the notion of utility, and the principle of maximizing expected utility, both depend on the description of uncertainty through probability. The justification for the necessity of this description due to de Finetti is outlined. The twin, practical problems of the evaluation of the decision-maker's probabilities and utilities are discussed. Probability, as used in the paradigm, is a subjectivist notion which is distinct from the chance, or frequentist, concept and there is discussion of this difference. The calculations for the analysis of a decision tree are described and the notions of the utility of data developed. The statistical analysis of data that flows from the paradigm is described and the basic, likelihood principle derived and discussed. The material is illustrated by a simple example from insurance.

AMS(MOS) Subject Classification: 90D35, 62C05, 90A10, 60A05

Key words: Bayesian statistics, coherence, subjective probability, utility, dynamic programming, insurance premium, likelihood principle, maximization of expected utility, chance, decision-making, risk aversion, scoring rule, utility independence, probability assessment, utility assessment, calibration, reconciliation, indifference curves, rate of substitution, decision trees, prior and posterior probabilities.

Work Unit number 4 - Statistics and Probability

Sponsored by the United States Army under Contract No. DAAG29-80-C-0041.

SIGNIFICANCE AND EXPLANATION

Consider a decision-maker who is required to choose amongst a number of decisions in a situation in which there is some uncertainty. How is he to decide? In the 1920's Ramsey showed that any sensible procedure amounts to describing that uncertainty by a probability distribution, to measuring the quality of possible outcomes by a utility function, and choosing that decision which maximizes expected utility. Any other procedure can be shown to be defective.

The paper discusses this recipe of Ramsey's, outlining a justification due to de Finetti, and then addresses the twin practical problems of assessing the decision-maker's probabilities and utilities. How can we assess the probability of nuclear accidents? How can we evaluate the need for nuclear power stations? The tools we have at the moment are simple but useful, though few attempts have been made to use them.

When uncertainty is present it is natural and sensible to try to reduce it by acquiring more information or more data. This is expensive and loses utility. How can the amount of information be measured; and how can we sensibly handle the data obtained? These statistical questions are answered and the remarkable likelihood principle discussed.

The paper is an invited review for the European Journal of Operational Research. In my view Ramsey's discovery must count amongst the most important advances of this century and a proper appreciation of his argument could greatly improve civilized life because we could make decisions more wisely and also communicate our ideas more easily.

The responsibility for the wording and views expressed in this descriptive summary lies with MRC, and not with the author of this report.

The Subjectivist View of Decision-Making

Dennis V. Lindley

(Invited review paper for the European Journal of Operational Research)

May 1981

THE SUBJECTIVIST VIEW OF DECISION-MAKING

Dennis V. Lindley

1. The Bayesian paradigm.

We begin with a formal statement of the subjectivist, or Bayesian, paradigm for a decision-maker, conveniently called you. A set D of possible decisions d is available and you are required to select one from the set. Complete information is not available and you are uncertain which of a number of possibilities θ in a set H obtains. A pair (d, θ) is called a consequence and effectively describes the outcome for you were you to select d and θ were true. Since all the uncertainty is supposed concentrated in H , a consequence, for given d and θ , is known to you. It is necessary to describe two aspects of the situation: the uncertainty surrounding θ ; and the fact that some consequences are more attractive to you than others. These are expressed numerically as follows. The uncertainty about θ were you to select d is described by a probability distribution P_d over H . The comparison of consequences is effected by a real-valued utility function $u(d, \theta)$. With these two measures available the optimum decision is that d that maximizes the expected (with respect to P_d) utility.

The key ingredient in the paradigm just described is your probabilistic description of the uncertainty surrounding θ . Once that is admitted the utility and expectation results are simply derived as follows. For simplicity in exposition suppose D and H are both finite so that there are a finite number of consequences $c = (d, \theta)$. Select from these the best and worst consequences c_1 and c_0 , and assign them utilities one and zero

respectively. If c is any other consequence, you may consider c as an alternative to a gamble that has probability u of resulting in the best c_1 and probability $1 - u$ of c_0 . For u near one, the gamble is preferred; for u near zero, the sure outcome c is preferred. It is hard to escape the conclusion that there is a unique value of u such that you are indifferent between c and the gamble. This number is called the utility of c , written $u(c)$ or $u(d, \theta)$. To show that only the expected utility is relevant consider any decision d . It will result for you in a consequence (d, θ) with probability P_θ , conveniently described by a density $p(\theta|d)$, the probability of θ were d to be selected (or simply, given d). But (d, θ) is equivalent to a probability $u(d, \theta)$ of the best consequence c_1 (and probability $1 - u(d, \theta)$ of the worst). So, by the rules of probability, the choice of d equivalently results in c_1 with probability $\int_{\theta} u(d, \theta) p(\theta|d)$, and otherwise c_0 . Clearly the best decision is that with the highest probability of the best consequence; but this probability is the expected utility and the MEU principle follows. Notice that utility is defined probabilistically and it is this probabilistic aspect that justifies, by the rules of the probability calculus, the combined probability $\int_{\theta} u(d, \theta) p(\theta|d)$.

2. Example.

An insurance company has to decide the premium d to offer a small airline that wishes to insure two planes, of a new design, for one year. The company's assets are 8 (in suitable units) and the insured value of a plane is 2. (Only total loss is being considered.) The probability of a loss of one aircraft ($\theta=1$) is assessed at 0.0388, and of both ($\theta=2$) at 0.0006, leaving 0.9606 for that of no loss ($\theta=0$). The company's utility for assets x is $1 - e^{-x/5}$. (These values will be discussed further below.)

If the insurance is not undertaken the assets will remain at 8 with utility of $1 - e^{-8/5} = 0.798$. If a premium d is offered and accepted, the assets will be either $8+d$ (if no loss), $6+d$ (one loss) or $4+d$ (two losses) and the expected utility is

$$1 - e^{-(8+d)/5} \times .9606 - e^{-(6+d)/5} \times .0388 - e^{-(4+d)/5} \times .0006.$$

This is equal to the original utility of 0.798 where d is about 0.10.

Hence any premium above this value is sensible. Notice that the expected loss is

$$2 \times .0388 + 4 \times .0006 = 0.08,$$

so that the smallest reasonable premium is 25% above the expected monetary loss. This increase is ascribable to the form of the utility function. In practice, administrative expenses will have to be added to the figure of 0.10 to arrive at a realistic figure.

3. The inevitability of probability.

We return to consider the main feature of the subjectivist paradigm, namely the description of your uncertainty about θ through a probability distribution; your probability for θ . It is important to recognize that it is not assumed that probability is an appropriate description of uncertainty but rather it is proved, starting from other assumptions about uncertainty, that you, to decide sensibly, must have a distribution. The earliest such proof was given by Ramsey. Later proofs were provided by de Finetti and by Savage. We now outline part of de Finetti's demonstration which is both the simplest and the most useful practically.

Suppose you are considering the event A that θ belongs to some set in H : in the example, consider the event that there are no accidents, $\theta=0$. Suppose you agree to describe the uncertainty of A by a real number x , say. Then one possibility is to see how good you are by giving you a penalty score $(x-1)^2$ if A subsequently turns out to be true and x^2 otherwise. The scores are the squares of the discrepancy between x and 1, for a true event, and 0 for a false one. The idea here is simply to provide a measure of your ability in a realizable fashion: to keep a check on your skill as a worker in OR. Then clearly x should lie between 0 and 1 for a value of x in excess of 1 would give larger scores than $x=1$ whatever happened to A : similarly $x=0$ is always better than $x < 0$. Now suppose you assign x to A and y to $\bar{A}$, the negation of A . The total score if A is true is $(x-1)^2 + y^2$, and if A is false ($\bar{A}$ true) $x^2 + (y-1)^2$. These scores are the squares of the distances of (x,y) from $(1,0)$ and $(0,1)$ respectively and can both be reduced by dropping a perpendicular from (x,y) to the line through $(1,0)$ and $(0,1)$ and replacing (x,y) by the coordinates of the foot of the perpendicular. Hence the only reasonable values

of x and y lie on the line which has equation $x+y = 1$. But this is the addition rule of probability that says that the negation of A has 1 minus the probability of A ; $y = 1-x$. The product rule $p(AB) = p(A)p(B|A)$ can be derived by a similar, but more involved, geometric argument. Consequently we have proved that a numeric description of uncertainty, when tested by the quadratic scoring rule, must be a probability. (It turns out that the particular rule used is almost irrelevant.)

In the outline of de Finetti's proof we saw that if values x and y for A and $\bar{A}$ were used that did not add to 1, then the score would always be increased. The result generalizes. If you use a decision procedure which is not equivalent to assignments of probability and utility, followed by MEU, then the decision could always be improved, whatever be θ , by some procedure which did proceed according to the subjectivist paradigm. The EEC has recently enacted weights and measures legislation which uses a t-test: a procedure which does not agree with the paradigm. The community is therefore suffering a sure loss. The subjectivist paradigm is often called coherent because it concerns the way decisions and events fit together, or cohere. The proof above concerned the coherence of x and y for A and for $\bar{A}$. The EEC procedure is incoherent.

A distinction is sometimes made between decision-making when the probabilities are known, and when they are unknown. Such a distinction is void in the subjectivist view because probability is your description of what you know: it always exists for you. There may be a practical problem in your finding it, as will be discussed below, but it is the description of uncertainty. The spurious distinction partly arises through thinking of probability in frequency terms: another point mentioned below.

As a theory of sensible behaviour for a single decision-maker the subjectivist approach is unassailed. No criticism known to me has much substance. The criticism that does warrant serious consideration is that that queries the practicality of the procedure. How, it is argued, can you assess probabilities and utilities: for if you cannot, coherence is not an available option for you. The practical implementation is indeed formidable but in understanding the criticism an analogy may not be out of place. Euclidean geometry was a valid theory for many centuries but was of limited applicability because people had difficulty in measuring angles and distances. It was not until the inventions of triangulation and theodolites that the geometry became fully implementable. As so it is with Bayesian ideas: at the moment we lack good theodolites of uncertainty. Unfortunately OR workers and others, instead of tackling the measurement problems for probability and utility, resort to other, incoherent procedures like minimax. Nevertheless some progress has been made and we now consider practical assessment techniques for utility and probability.

4. Determination of your probabilities:

A simple way to train a person, you, in probability assessment is to use the quadratic scoring rule: to take a series of events, have you assess their probabilities, check on their truth and calculate the total score. It is common to use almanac questions (has Rome a larger population than Paris?) whose answer is unknown to you but can easily be found from an almanac. One context in which this seems to be done is in the training of meteorologists, at least in North America, where the event might be "rain tomorrow". Such exposure can have an effect on your perception of uncertainty. Confident people tend at first to give values near 0 or 1 but learn from the large scores they incur when they are wrong. Overcautious subjects hovering around 0.5 become bolder when they see so many of the events they have given probability 0.6 to become true. Nevertheless the method is not entirely satisfactory if only because the events and the scoring rule are not natural. OR workers are not interested in almanac questions and their rewards are not determined by the scores. The first defect can be alleviated to some extent by replacing the almanac questions by more relevant ones: for example, the casualty assessor in the insurance example above could be scored $(.0388-1)^2$ if one aircraft crashed. But it may be felt that other factors besides the score should be taken into account in evaluating his worth to the insurance company, and if the assessor feels this, he may be led to be motivated to distort his probability evaluations. Thus suppose the penalty scores were greater for true events than false ones by a factor of 2, so that they were $2(x-1)^2$ and x^2 . For an event of probability truly p , the expected score is $2p(x-1)^2 + (1-p)x^2$ which is least at $x = 2p/(1+p) > p$. Hence the assessor will increase his evaluations above their correct values. (Equally $p = x/(2-x)$, so the stated values could be downgraded from x to $x/(2-x)$ to

give probabilities.) Thus we see that it is dangerous to use implicit scoring rules because they may motivate you to give misleading answers. Whilst on the subject of keeping a check on an OR worker or statistician; this is not usually done but it would surely be of interest to do so. For example what about all the hypotheses declared significant at 5% by conventional statistical tests: how many were in fact false?

Another check that is sometimes applied, again with weather forecasters, is to take all the events that were assigned a probability near, say, 0.8; perhaps between 0.75 and 0.85; and see how many were subsequently true. One intuitively feels that 80% should be true; if so, you are said to be well-calibrated. Many people, perhaps most, are not well-calibrated, usually fewer than 80% of the events turn out to be true

There is another method of assessing probabilities that seems promising, though it has been little tried in practice. To appreciate this you need to be clear what is meant by saying that uncertainty is described by probability. It does not just mean that each event is assessed by a number lying between 0 and 1 - which all the methods already mentioned use - but that uncertainties for different events combine according to the rules of the probability calculus. These are the addition and multiplication rules and are the basis of the fundamental idea of coherence between different judgments. The method uses this notion of coherence. For example, if A is the event of interest, you may be asked for $p(A)$ but also for $p(A|B)$ for some appropriate event B, for $p(A|\bar{B})$ and for $p(B)$. These should combine according to the rule

$$p(A) = p(A|B)p(B) + p(A|\bar{B})[1-p(B)].$$

If the stated values do not do this then at least one of them will have to be adjusted. In the aircraft insurance example let A be the event of at least one loss, assessed at 0.0394, or 0.04 to 2D. You may feel that this depends on the usage the aircraft gets in the year and that this in turn depends on B, the event that the airline gets a contract that is on offer. If it does, the value might go up to 0.06, if not it may be as low as 0.03. You assess the probability of getting the contract at 0.6. But the right-hand side of the equation just displayed is

$$0.06 \times 0.6 + 0.03 \times 0.4 = 0.048$$

not 0.04 as originally assessed, but 20% higher. At least one, and usually more, of the four values must be revised. Such a process of revision is called reconciliation: one has to reconcile the different values obtained by looking at different aspects of the problem. The basic idea here is not just to look at the issue of immediate interest - loss of an aircraft - but to look at related matters, like usage, to obtain a coherent picture of the situation. Since MEU is really all about coherence, the incorporation of this idea into probability evaluations seems right in principle. It is not unlike the principle of triangulation already referred to in connection with Euclidean geometry, wherein several measurements are taken and least-squares used to reconcile discrepancies observed. The topic is discussed by Lindley et al. (1979).

There are other issues involved in the assessment of probabilities, including the fact that other factors than the uncertainty may enter into consideration. For example, an event perceived as unpleasant may have its probability underestimated - thus subject's asked to assess their probabilities of death from various causes tend to give values that incoherently add to less than one - or an unfamiliar one exaggerated. A clear statement of

the scoring rule is one possibility, though impractical in considerations of death, but reference to coherence is perhaps a better way. To relate the probability of a nuclear accident to that of an automobile accident is useful: the latter being a familiar risk that we are prepared to tolerate. But issues like this are confused with utility considerations, so we turn to discuss these.

5. Determination of your utilities.

The point was made above that utility is not just a number describing the worth of a consequence but a number measured on a probability scale, and that its rules of combination are essentially probabilistic. Any determination of utility must therefore have a probability ingredient. Consider first the case where the consequences (d, θ) are purely monetary. (This is reasonable in the insurance example.) The required evaluation is of $u(x)$, the utility for monetary assets, x . Notice that it is necessary to speak of assets, not gains or losses, because the paradigm is in terms of consequences or outcomes. A gain of £100 changes a consequence of £1000 into one of £1100. An error is often made of speaking in terms of changes in consequences rather than in terms of the consequences themselves.

There are basically two ways of determining utility; with fixed probability and varying outcomes, or with fixed outcomes and varying probability. Thus you may be asked to consider what sure x is equivalent to equal probabilities of y and z : then $u(x) = \frac{1}{2} \{u(y) + u(z)\}$. Alternatively for $y < x < z$, what probability p makes sure x equivalent to an uncertain situation with probabilities p of z and $1-p$ of y : then $u(x) = pu(z) + (1-p)u(y)$. By asking a series of questions like this the utilities can be determined at a series of values $x, y, z, \dots$ and either a curve faired in or a member of a class of curves fitted by a procedure like least-squares. As with probability it is advisable to ask more questions than are minimally needed to provide a check on coherence.

A phenomenon that often arises in studying monetary utility is that of risk aversion. In the example with equal probabilities in the last paragraph, it often happens that x is less than $\frac{1}{2}(y+z)$, the expected monetary (as distinct from utility) evaluation of the uncertain situation, reflecting a

dislike of the uncertainty. A person giving such an evaluation is said to be risk averse (if $x > \frac{1}{2}(y+z)$, he is risk prone; $x = \frac{1}{2}(y+z)$ is risk neutral). The appropriate measure of risk aversion is $-u''(x)/u'(x)$ where the primes denote differentiation. The function $u(x) = 1 - e^{-ax}$ used in the aircraft insurance example has constant risk aversion of amount a , and we saw how it led to a premium in excess of the monetary amount. Risk aversion that decreases with x is perhaps more reasonable since a risky situation can be more easily tolerated the greater are the assets.

Notice that the discussion just given does not depend on x being money; much of it will be appropriate whenever the consequences are in terms of a single real number. Another example is provided by measures of ability when the decision is whether or not to accept the person for a training programme.

Suppose next that the consequences are described not by one real number, x , but by two, x and y , and, for definiteness that the utility is increasing in both x and y . As before x might be assets, the new y might be inventory. You then have to determine your utility function $u(x,y)$. One possibility is to determine indifference curves in the (x,y) -plane such that the utility is constant on a curve. The problem is then to determine the utility for each curve and the earlier, one-dimensional methods can be used. Another useful device borrowed from economics is the marginal rate of substitution of one quantity for another. If y is decreased by Δ by how much, $\lambda\Delta$, will x have to be increased to keep the utility constant? As $\Delta \rightarrow 0$, λ is the rate, and $-\lambda^{-1}$ the slope, of the indifference curve at (x,y) .

These ideas involve consideration of changes in both x and y . Suppose x is held fixed and variations in y are considered. Then, by supposition, larger y 's are to be preferred. But suppose that more was true; namely that attitudes to any two uncertain situations $p_1(y)$ and $p_2(y)$, defined by probability distributions over y for fixed x , did not depend on x , so that if you preferred one to the other for one x you would prefer it for all x . Then y is said to be utility independent of x . Then it is not difficult to show that $u(x,y) = f(x) + g(x)h(y)$ for suit-able f , g and h . A particularly important case is that of mutual utility independence of x and y , when, in addition the same result holds with x and y interchanged. Then $u(x,y) = F(x)H(y)$ for suitable F and H . Independence notions of this type are increasingly important in higher dimensions. Keeney and Raiffa (1976) provide an excellent account with practical examples.

6. Subjective aspects of probability.

An objection to the subjectivist view of decision-making is its subjectivity. This is correct: it is a view of the world appropriate to a single decision-maker, or subject, that we have called you. You may be an individual, but equally you may be any group that has agreed collectively to act as a decision-maker, whether a company, acting through a board of directors, or a nation, acting through its government. It is not, and does not claim to be, a method of reaching decisions when two or more decision-makers are in conflict. As far as I am aware, there is, outside of the two-person, zero-sum game, no paradigm for conflict decision-making; and that paradigm is deficient in many applications because the zero-sum assumption is inappropriate. Thus in a NATO, Warsaw Pact context the respective utilities are not the same; indeed, that is what the disagreement is about. In default of any sensible theory of conflict decision-making the subjectivist view can make a useful contribution. For example, in a military conflict it is valuable for one side to list the scenarios open to the enemy, the θ 's of the model, and to assess the probabilities of the enemy taking each of them. Certainly this is better than much current military thinking that adopts a minimax strategy, guards against the worst and hence escalates the conflict aspect with a resulting build up of forces that itself threatens peace. (Incidentally, within the context of a single decision-maker, the minimax strategy, not being MEU, is typically incoherent.)

The Bayesian view does not say how the differing opinions are to be reconciled, say within a company. It is however clear that many differences can be ascribed to incoherence on the part of one or more of the directors, and that a sharing of views in the framework of utility and probability can help to resolve many of them. If compromise is finally necessary the theory

does not say how it should be reached. It does say that whatever decision is adopted it should agree with some probability and some utility specifications. An explicit statement of what these are can be enormous help in discussing a position. An insurance company could usefully determine its current utility function for money and instruct its underwriters accordingly.

In the subjective view probability is an expression of your uncertainty concerning the world. Probability, as de Finetti says, does not exist - existing in the sense of being a property of the material world irrespective of you. Rather it is an expression of a relationship between you and that world. An example may clarify this. Consider a conventional die with six faces numbered from 1 to 6. If the die is rolled a large number of times the proportion of times it shows 6 (or any other number) will stabilize around a value ω , say. This value is a property of the die. But before the rolling you may have a probability p that the die, on its first roll, will show 6. There is no suggestion that $p = \omega$. Of course, after all the throws, the new probability (having been revised coherently by the rules of the probability calculus) will equal ω , but this need not be true initially: simply your opinion about the die changes as a result of the rolls. ω is often called a chance: it is a frequency concept, whereas p is not. Probability, as used in this paper has no frequency or repetitive connotation.

The methods described for the practical determination of utility and probability are not entirely satisfactory. Really sound methods will involve training of decision-makers. This training could begin in school. Today our teaching is essentially based on right and wrong; true and false. It ought to be based on a realistic appreciation of the world in which uncertainty is rife. We should be taught to live with uncertainty and to handle it

sensibly: not to say that politician is right, but that he is right with probability 0.7.

Notice that the subjectivist paradigm involves selection amongst a given list D of decisions; and that the uncertainty is amongst the members of $\mathcal{H}$. Essentially the method only compares alternatives. It does not admit the notion of "do something else" (not in D) or the possibility that something else is true (not in $\mathcal{H}$). At first this seems unsatisfactory but reflection suggests that it is reasonable. Things are not good in themselves or unusual in themselves: they are only better or worse than others, more or less common than others. The world is comparative, not absolute. Of course, you should keep your mind open to a possible extension of D or $\mathcal{H}$, for such creativity is tremendously important, and make D and $\mathcal{H}$ as large as possible, but most of the time you merely need to compare the possibilities.

7. Decision Trees and MEU.

We next pass from probability determinations to discuss the implementation of the MEU process. A useful tool is the decision tree (Figure 1) with decision and random nodes. Utilities are associated with the terminal branches and the general procedure is to take expectations at each random node and maximize at each decision node. These two operations are the only ones used in the method. Figure 2 shows a more elaborate tree in which both decisions and events are broken into two groups. The collection and handling of information are important aspects of OR and this tree has first a decision e (for experiment) on what data to collect before a final decision d . It also has two uncertain quantities, θ , relevant to d , and x , the data actually collected (the result of the experiment). Notice how the tree is written in temporal order from the choice of experiment e to the realization of θ , and how the probabilities are always conditional on what has occurred before. The analysis of a tree always proceeds in the order opposite to that of time - begin at the end and go on until you come to the beginning. At the last random node an expectation is evaluated over θ , which is then maximized at the last decision node giving

$$\max_d \sum_{\theta} u(d, \theta; e, x) p(\theta | d, e, x) = U(e, x)$$

say. The reduced tree that remains, with e and x , is of the same structure as that in Figure 1 with d and θ . In particular $U(e, x)$, replacing $u(d, \theta)$, is a utility: namely the utility as perceived by you of the situation where you have performed e with result x . The analysis is completed by taking an expectation over x and then a maximization over e yielding

$$\max_e \sum_x U(e, x) p(x | e).$$

The procedure used in this analysis is often called dynamic programming based on the optimality principle. Both the method and the principle are elementary deductions within the subjectivist paradigm and the failure to recognize this has led to obfuscation and pretentious claims for the principle. The procedure can obviously be generalized to any finite number of stages; the final stage is usually called the horizon. Notice that the method is an algorithm for the evaluation of the optimum decision in that it prescribes a sequence of computations that lead to the answer. Unlike Newtonian mechanics, MEU does not lead to a differential equation for which an algorithmic solution has to be sought. With a long decision tree, it is unfortunately true that the Bayesian calculations become impossibly time-consuming even on the fastest machines and approximations have to be developed.

Notice that the method involves some principles of coherence. One has already been mentioned: the quantity evaluated at the second decision node, $U(e,x)$, is itself a utility and would be appropriate if the horizon were reduced from θ to x . You might find it useful to compare your value of U , calculated from u , with your directly perceived value after x , just as $p(A)$ was compared with its value when B was incorporated. The other coherence concerns the probabilities $p(x|e)$ and $p(\theta|d,e,x)$: the former does not depend on d so that the product is $p(x,\theta|d,e)$, the joint distribution of the two uncertain quantities x and θ . This may also be written

$$p(x,\theta|d,e) = p(x|\theta,d,e)p(\theta|d,e).$$

This alternative presentation of the uncertainty is often convenient because it displays the dependence of x on θ . The purpose of the experiment was to enhance your knowledge of θ so that the result x will depend on θ . Again your perceptions in both approaches may provide a convenient check on coherence. Statisticians favour this last method, calling $p(\theta|d,e)$ the prior (to e) probability of θ , and $p(x|\theta,d,e)$ is the likelihood (of θ)

given x . Then $p(\theta|d,e,x)$ is the posterior (to e) probability of θ . There is no standard nomenclature for $p(x|e)$ which is your perception of what would happen were you to perform e , not knowing θ . We will return to the statistical aspects below, but first we consider the value of an experiment.

The expected utility from performing an experiment e is

$$\sum_x \max_d \sum_{\theta} u(d,\theta,e,x) p(\theta|d,e,x) p(x|e) = U(e)$$

say. One possibility is not to perform an experiment and so collect no data. Consider this as a null experiment e_0 with null data x_0 . Then

$u(d,\theta,e_0,x_0) = u(d,\theta)$ and $p(\theta|d,e_0,x_0) = p(\theta|d)$, evaluations already made, and $p(x_0|e_0) = 1$. Consequently

$$U(e_0) = \max_d \sum_{\theta} u(d,\theta) p(\theta|d)$$

as before. Hence e is only worth performing, or the data worth collecting if $U(e) > U(e_0)$. The difference $U(e) - U(e_0)$ is the expected utility of e . (The expression, expected value of sample information, EVSI, is sometimes used.) In the special case where $u(d,\theta,e,x) = u(d,\theta)$, so that the performance of the experiment never decreases the utility (or is cost-free), and $p(\theta|d,e) = p(\theta|d)$, so that the decision to use e does not change your perception of θ , the expected utility of e is non-negative. Loosely expressed, cost-free data is always expected to be of utility. Notice the use of the word 'expected' in that sentence: it can happen that some data values, x , can reduce utility - but you do not expect that to happen. A conceptually useful experiment is one that tells you the value of θ ; so that $x = \theta$. This is called a perfect experiment, e_1 . Under the two conditions just mentioned, $U(e_1) > U(e) > U(e_0)$ and $U(e_1) - U(e_0)$ is the expected utility of perfect information. It provides an upper bound for the expected utility of all experiments.

An immediate application of these ideas is to sampling inspection. For a batch of uncertain quality θ , the decisions d may be to reject or to accept it. An experiment may be to take a sample of size e (or, in the more usual notation, n) and see how many defectives x it contains. The above analysis deals with both the optimum choice of n (including no sampling, $n = 0$ or $e = e_0$) and the acceptance/rejection problem. A practical difficulty lies in the evaluation of $u(d, \theta)$ when d is acceptance for this involves the disutility of a customer receiving a defective item, which is notoriously difficult to assess. However the principle of coherence can be of help here. Often this disutility will be sensibly constant over a range of products so that if one sampling scheme has been selected for product A , there is an implicit disutility that may be used when discussing product B . The point is that the proper invariant is the disutility, not other quantities like acceptance probabilities. The pioneering work of Hald (1967) seems to have made little impact on sampling inspection.

8. Statistical aspects: the likelihood principle.

We now return to the statistical aspects of experimentation. The basic result is the equality of the alternative expressions for the joint uncertainty of x and θ . Omitting the dependence on d and e (or alternatively thinking of these as fixed) you may write

$$p(x)p(\theta|x) = p(x|\theta)p(\theta),$$

the left-hand side being the original form in the MEU approach. If this is considered as a function of θ for fixed x we may write

$$p(\theta|x) \propto p(x|\theta)p(\theta).$$

In words, the posterior probability of θ given x is proportional to the product of the likelihood of θ for x and the prior probability of θ . This is Bayes theorem, and its ubiquity gives rise to the subject being called Bayesian statistics. Its importance lies in the fact that it tells you how to react sensibly (coherently) to partial information about the quantity of interest, θ , in the form of data, x . In the sampling inspection application, if θ is the fraction defective in a (large) batch, $p(x|\theta) = \theta^x(1-\theta)^{n-x}$, and $p(\theta)$ describes your initial opinions about the batch before sampling. Distinguish between $p(x|\theta)$ as a function of x , for which it is a probability, and $p(x|\theta)$ as a function of θ , called a likelihood, which is not a probability.

We now consider a result which is perhaps the most important discovery made in statistics this century: the likelihood principle. Suppose that you are at the second decision node of Figure 2, having observed x and about to select d . Then the only relevant probability is $p(\theta|d,e,x)$ and by Bayes theorem the only contribution x makes to this probability is through the likelihood $p(x|\theta,d,e)$ (which will not involve d). Hence after observation of x the only relevant feature of x is the likelihood of θ given x :

this is the likelihood principle. For example, in the sampling inspection case with x defects in a sample of n the likelihood is $\theta^x(1-\theta)^{n-x}$ irrespective of whether a single sample of size n was taken, or whether two stages, n_1 and n_2 , with $n_1 + n_2 = n$, were used, or whether the sampling was done completely sequentially and stopped according to some rule that depended only on the sampling experience. It is remarkable that the principle is denied by all of conventional statistics. For example, a significance test of $\theta = \theta_0$ (θ_0 may be the quality limit) is performed by calculating a tail area like $\sum_{y>x} \theta_0^y(1-\theta_0)^{n-y}$ in the single sample case. In general a summation is required over other samples (here (y,n) with $y > x$) besides the one obtained. The other samples will depend on whether the sampling was one-stage, two-stage or sequential and consequently the notion of a significance test violates the likelihood principle, and hence the Bayesian paradigm, and is incoherent. Examples of this misuse abound. Clinical trials, which usually have a strong sequential element, are typically analyzed using significance tests; and the misuse of money in the improper analysis of cancer trials alone must surely be appreciable. Notice that the likelihood principle does not obtain prior to observing x (as indeed, it is then meaningless). This is clearly seen in the analysis at the preceding decisions node where $\sum_x U(e,x)p(x|e)$, a summation over x , as in a significance test, is required. The controversy between sampling-theory and Bayesian statistics really revolves around what happens after the data are to hand when the likelihood principle is the major difference. The argument has been advanced that statistics is concerned with inference, not decision-making, and that the likelihood principle does not there obtain. Ramsey's view seems correct: the purpose of inference is to enable potential decisions to be made. And for any decision problem that involves only θ as the uncertain element, only the

given x matters and that, by Bayes theorem, involves only the likelihood.

We now return to the insurance example and have a more detailed look at the probability calculations involved there. This will illustrate the value of coherence and the relevance of the likelihood principle.

9. Example.

In the earlier example the probabilities of 0, 1 and 2 accidents to the 2 aircraft were assessed at 0.9606, 0.0388 and 0.0006 respectively. You could have assessed these directly (and coherently since they add to one) but it might be advantageous to introduce other probabilities as a further check on coherence. One possibility is to suppose that each aircraft has a chance ω of total loss in a year: ω being an unknown rate for this type of aircraft. (See the earlier discussion of chance for a die.) Then the probability distribution for ω could be assessed. This was done here and a density

$$p(\omega) = 99 \times 98 \cdot \omega(1-\omega)^{97}$$

was selected. This has a mean of $2/100 = 0.02$ or 1 loss per. 50 aircraft per. year. The probability that the rate is below 0.01 is 0.37; 0.02 is 0.59; 0.04 is 0.91 and 0.06 is 0.98. The advantage of thinking in terms of an overall rate ω is that many probabilities can be deduced from it. (Remember, ω is not a probability, but a chance.) Thus, given ω , the probability of no accident is $(1-\omega)$ for one aircraft and, by the independence of chances, for two aircraft is $(1-\omega)^2$. Hence

$$p(\theta=0) = \int_0^1 p(\theta=0|\omega)p(\omega)d\omega = \int_0^1 (1-\omega)^2 p(\omega)d\omega$$

which gives 0.9606 as above. The other values follow similarly.

To further illustrate the value of considering ω suppose that a year later the policy comes up for renewal, no accidents having occurred to the two aircraft. A year later some other data will be available. Suppose that there has during the year been no total loss of any aircraft of the type insured and that the total exposure including the airline's, has been the equivalent of 40 aircraft/years. In the language and notation above, there have been $x=0$ defects in $n=40$ cases. The likelihood for ω is $(1-\omega)^{40}$. Multiplying by

$p(\omega)$ and using Bayes theorem

$$p(\omega|0,40) = (1-\omega)^{40} \omega(1-\omega)^{97}$$

or

$$p(\omega|0,40) = 139 \times 138. \omega(1-\omega)^{137}.$$

Repetition of the calculations for the first year show that your probability of no accidents in the second year is

$$\int_0^1 (1-\omega)^2 p(\omega|0,40) d\omega$$

which is 0.9717, an increase due to the favourable experience with that type of aircraft. The probabilities for one and two accidents are 0.0280 and 0.0003. The new premium is 0.07, a reduction on the original value of 0.10. The revised expected loss is 0.057. Now you have a description which is coherent over both years and is extendable to future experience with the aircraft. The reader might like to do the calculations supposing that an aircraft crashes, providing an additional likelihood of ω , and see how the premium rises in response to the disaster.

Notice that the only aspect of the data used in the analysis was the likelihood $(1-\omega)^{40}$, in accord with the likelihood principle. In particular, it was not necessary to consider whether the exposure of the 40 aircraft/years was fixed or was obtained randomly as would strictly be required by conventional statistical analyses.

10. Further reading.


There are two, elementary treatments of this subject, by Raiffa (1968) and Lindley (1971a). The standard texts on the theory are Raiffa and Schlaifer (1961) and Schlaifer (1969). A good treatment that happily blends theory and practice is Brown et. al. (1974). No one who thinks carefully about probability can afford not to read the brilliant, but difficult, volumes of de Finetti (1974): here is wisdom. The best statistical treatment, again difficult, is Jeffreys (1967). Any one who has read these two books is well-equipped to sort out the wheat from the chaff in the many, other statistical texts. A review is provided by Lindley (1971b).

REFERENCES

- BROWN, R. V.; KAHR, A. S. & PETERSON, C. (1974). Decision analysis for the manager. New York: Holt, Rinehart & Winston.
- DE FINETTI, B. (1974). Theory of Probability. New York: Wiley.
- HALD, A. (1967). Asymptotic properties of Bayesian single sample plans. J. Roy. Statist. Soc. B, 29, 162-173.
- JEFFREYS, H. (1967). Theory of Probability. Oxford: Clarendon Press.
- KEENEY, RALPH L. & RAIFFA, HOWARD (1976). Decisions with multiple objectives: preferences and value tradeoffs. New York: Wiley.
- LINDLEY, D. V. (1971a). Making Decisions. London: Wiley Interscience.
- LINDLEY, D. V. (1971b). Bayesian Statistics: a Review. Philadelphia: SIAM.
- LINDLEY, D. V., TVERSKY, A. & BROWN, R. V. (1979). On the reconciliation of probability assessments. J. Roy. Statist. Soc. A, 142, 140-180.
- RAIFFA, HOWARD (1968). Decision analysis. Reading: Addison-Wesley.
- RAIFFA, HOWARD & SCHLAIFER, ROBERT (1961). Applied Statistical Decision Theory. Boston: Harvard Business School.
- SCHLAIFER, ROBERT (1969). Analysis of decisions under uncertainty. New York: McGraw-Hill

DVL/db


Figure 1



□ : decision node; maximum over d .

○ : random node; expectation over θ .

Figure 2



REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 2248	2. GOVT ACCESSION NO. AD-H103	3. RECIPIENT'S CATALOG NUMBER 87901
4. TITLE (and Subtitle) THE SUBJECTIVIST VIEW OF DECISION-MAKING		5. TYPE OF REPORT & PERIOD COVERED Summary Report, no specific reporting period
7. AUTHOR(s) Dennis V. Lindley		6. PERFORMING ORG. REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS Mathematics Research Center, University of 610 Walnut Street Madison, Wisconsin 53706		8. CONTRACT OR GRANT NUMBER(s) DAAG29-80-C-0041
11. CONTROLLING OFFICE NAME AND ADDRESS U. S. Army Research Office P. O. Box 12211 Research Triangle Park, North Carolina 27709		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 4-Statistics & Probability
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) 134		12. REPORT DATE July 1981
		13. NUMBER OF PAGES 28
		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Bayesian statistics, coherence, subjective probability, utility, dynamic program- ming, insurance premium, likelihood principle, maximization of expected utility, chance, decision-making, risk aversion, scoring rule, utility independence, probability assessment, utility assessment, calibration, reconciliation, indifference curves, rate of substitution, decision trees, prior and posterior probabilities.		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) The subjectivist, Bayesian paradigm for a decision-maker is described. It is shown how the notion of utility, and the principle of maximizing expected utility, both depend on the description of uncertainty through probability. The justification for the necessity of this description due to de Finetti is outlined. The twin, practical problems of the evaluation of the		

ABSTRACT (continued)

decision-maker's probabilities and utilities are discussed. Probability, as used in the paradigm, is a subjectivist notion which is distinct from the chance, or frequentist, concept and there is discussion of this difference. The calculations for the analysis of a decision tree are described and the notions of the utility of data developed. The statistical analysis of data that flows from the paradigm is described and the basic, likelihood principle derived and discussed. The material is illustrated by a simple example from insurance.