Journal of the Operational Research Society (1995) 46, 70-79 0160-5682/95 $9.00

Uncertainty and Imprecision: Modelling and Analysis

SIMON FRENCH
University of Leeds

Any statistical analysis or decision analysis contains numerical inputs of which we are unsure. Some of our uncertainty arises from physical randomness which we can model in various ways, ideally through probability. Some uncertainty relates to judgemental estimates of quantities about which we may be unsure in many different respects. There are other uncertainties involved, however: some relate to ambiguity and imprecision of meaning; some relate to lack of clarity in the objectives which the analysis seeks to meet; some relate to the numerical accuracy of calculations. How should the uncertainty arising from ambiguity be modelled? Other uncertainties can also impact on an analysis. Why is the analysis being conducted? Are the objectives clear?

Key words: Bayesian methods, decision analysis, descriptive and normative models, fuzzy sets and possibility theory, modelling imprecision, prescriptive analysis, probability modelling, requisite modelling, sensitivity analysis

INTRODUCTION

When is it appropriate to model imprecision—which is surely a synonym for uncertainty—with probability? Or, for that matter, with any other mathematical model of uncertainty? To answer this, there is a need to consider the purpose of modelling in an analysis. Why are we seeking to model anything: be it uncertainty, imprecision, the effect of gravity or the spread of AIDS? Only if we keep our attention focused on the objectives of the modelling, are they likely to be achieved. Thus, in the next section, the different objectives of descriptive and normative models and their role in particular analyses, especially prescriptive analyses, will be discussed briefly.

I shall use the term *model* in the sense of one of Savage's small worlds[1]; whereas I shall use the term *analysis* to mean the use we make of models to guide our thinking and actions. Thus, analysis is the mustering and exploration of one or more models to some purpose.

Uncertainty and *imprecision* are portmanteau words used in many different ways according to context. For brevity I shall often use 'uncertainty' to include the various forms of uncertainty that may arise from imprecision or ambiguity or lack of clarity. A number of different causes or categories of uncertainty that may arise and need to be considered in an analysis are suggested and discussed in the third section. I doubt that this categorization is exhaustive, but I hope it is sufficiently representative to inform debate on uncertainty and its modelling. The later sections in the paper consider each type of uncertainty in relation to the descriptive, normative and prescriptive aspects of an analysis.

In any modelling there is a danger of infinite regress. Any model differs in many respects from that which is being modelled—usually reality, but not always. There is always a temptation to introduce more subtleties into a model to reduce these differences. Then, because the more sophisticated model will still differ, to introduce further subtleties and so on. Is it only the finiteness of life that stops us following an infinite regress? Or do we stop when the model is good enough in some sense: to use Phillips' term, *requisite*[2]? This is considered at the end of the paper.

Correspondence: S. French, Division of Operational Research and Information Systems, School of Computer Studies, University of Leeds, Leeds LS2 9JT, UK

To be clear at the outset, I should state that this is not a paper with a basis in empirical fact. It is a paper full of opinion in which I raise issues for discussion. Many authors have suggested techniques for modelling uncertainty and imprecision. I am concerned to ask why we should want to model these; and, in those circumstances that we do wish to do so, why the proposed techniques should be adopted.

DESCRIPTIVE MODELS, NORMATIVE MODELS AND PRESCRIPTIVE ANALYSES

There have been several discussions of the distinction between descriptive and normative models over recent years. In References 3 and 4, I distinguished only between the descriptive and normative, using prescriptive synonymously for normative. Other discussions, such as those in Bell *et al*[5], have suggested that there is further distinction between normative and prescriptive models. I prefer now, however, to suggest that descriptive and normative *models* come together in prescriptive *analyses*, which serve to guide the evolution of our perceptions and our actions.

A *descriptive model* is a conjectured picture of reality, which mirrors possible relations between possible objects or classes of objects in the external world. It may be tentative, a scientific hypothesis, the merest suggestion of how the world might be, and as yet unsupported by comparison with empirical knowledge: or it may be something in which more trust may be placed, i.e. a scientific model that has been tested against data and in some sense has passed. Newton's laws of motion, quantum theory and the rather strange model of our planet adopted by the Flat Earth Society are all descriptive models, some of which we might feel compare better with empirical knowledge than others.

A *normative model* suggests how we might think, choose or act. It seeks to capture a possible set of principles which we might wish our thinking, our judgements, our decision making and our behaviour to obey. Obvious examples in the statistical and decision analytic worlds are subjective probability and utility theories; but others may be found in different logics and in many philosophical discussions.

Descriptive analyses simply explore the implications of descriptive models, although they stop short of comparing these implications with data and do not have the focus of an immediate problem. I have in mind here the sort of exercises many of us enjoyed in our mathematical education. How large an angle of inclination can a rough plane have before a rectangular block begins to slide? What is the probability of drawing a 'five' from a freshly shuffled pack of cards. Similarly, *normative analyses* explore the implications of normative axioms and models. Philosophers conduct these all the time: but so do statisticians, operational research scientists and others. For instance, consider recent discussions of the chance of winning a goat in a certain hypothetical quiz-show[6]. Such hypothetical mind-games help us understand how we should draw inferences.

Most analyses, however, are neither exclusively descriptive nor exclusively normative, but contain aspects of both. As statisticians, we help scientists compare descriptive models with data and we guide that process of comparison by appealing to normative models of inference. As decision analysts, we help decision makers choose strategies by using normative decision theories to guide the comparison of consequences that are predicted by descriptive models of strategies. In recent years it has become common to refer to such analyses as *prescriptive analyses*[5].

The purpose of any analysis is to bring understanding. In descriptive analyses the understanding is of the world about us; in normative analyses it is of norms of behaviour; and in prescriptive analyses it is of our beliefs, perceptions and preferences in relation to the issues before us in a particular inference or decision. Furthermore, in prescriptive analyses the understanding is not of ourselves as we were at the beginning of the analyses, but of ourselves as we are at its end. We learn, our beliefs change and our preferences evolve as a direct result of conducting the analysis.

There are further complexities in prescriptive analyses. They are usually carried out for a client—scientist, decision maker, safety engineer or whoever—by an analyst who possesses the necessary modelling skills. The analyst has to communicate, therefore, with the client, both to extract judgemental input, e.g. belief and preference information, and to convey the results of the analysis to the client. Communication requires that each party has an understanding—a descriptive model—of the other party in order to predict how words and sentences will be understood. The analyst's descriptive psychological models of the client may lead to the adoption of techniques designed to reduce 'biases' in eliciting subjective probabilities, utilities and other judgemental quantities.

The client is seldom a single person. The analyst must deal with a group: and that brings in further problems of communication between the analyst and the group members and between the group members themselves as they discuss the issues. The analyst must deal with group process issues such as 'groupthink', on the one hand, and 'free-riding', on the other. Furthermore, in a decision analysis, questions of equity and fairness may arise bringing in normative models of voting and collective choice[3]. Further problems can arise if expert advice is sought from third parties[7,8]. To complicate matters still further, there are analyses for which there are no clearly identified clients. Much of scientific research falls into this category. Here the objective is to increase the body of knowledge that the scientific community accepts.

The analysis never absolves the user from the responsibility for any inference or decision. In scientific research, for which there may not be clearly identified clients, it is the responsibility of the entire peer group of scientists to examine the analysis and draw from it what understanding they can. If, by and large, they draw the same understanding, that understanding becomes part of the current body of scientific knowledge. Thus, the peer group hold the responsibility for the inference.

I may seem to be getting a long way away from the subject of this paper, but I believe it is important to emphasize the complexity of the contexts in which one might wish to model uncertainty. A modelling approach may be appropriate to one set of circumstances and quite inappropriate to another, which at first sight seems little different. Moreover, I shall argue that there are types of uncertainty which are not usefully modelled in many (any?) contexts.

UNCERTAINTY: AMBIGUITY, IMPRECISION

Most writers on methodology note three major steps in conducting analyses, although they may use different terms and suggest that a single model is built.

(1) Modelling. The construction of a set of models: i.e. a move from the real world to a small universe of small worlds.

(2) Exploration. The exploration of these small worlds.

(3) Interpretation. The interpretation of the conclusions of these explorations into guidance for real-world beliefs, inferences and decisions.

During each of these steps the clients may express uncertainty, be ambiguous or imprecise, or be concerned at a lack of clarity of the issues.

Uncertainties expressed during modelling

- *Uncertainty about what might happen or what can be done*. For example, what might happen to a company's profits and trading position if it launches a new product onto the market? What might go wrong with a nuclear reactor and what set of circumstances might lead to each type of malfunction? How might a university respond if the Government changes the funding structures?
- *Uncertainty about meaning/ambiguity*. For example, a possible outcome of a general election might be a conservative victory: but what is meant by a 'conservative' government? Might a socialist government rule to all intents and purposes like a conservative one? How big a majority does a party need to implement its policies? A hospital may wish to improve its cost-effectiveness: but what is meant by 'cost-effectiveness'?

- *Uncertainty about related decisions*. For example, if I buy this set of furniture and then move house in six months time, how likely is it that the furniture will fit my new house? What will a company's competitors do in response to its launch of a new widget?

Uncertainties expressed during exploration of the models

- *Uncertainty arising from physical randomness or lack of knowledge*. For example, what is the chance of that coin landing heads? How likely is it that both engines on this aircraft fail on the next flight? What is the probability of the Conservative Party winning more than 370 seats in the next election?
- *Uncertainty about the evolution of future beliefs and preferences*. For example, this seems the right course of action now, but how will the clients feel about it in a year's time? Will their beliefs or preferences have changed such that they regret their decision?
- *Uncertainty about judgements, e.g. of belief and preference*. For example, should this subjective probability be set at 0.60 or 0.58? What weight should be placed on a particular attribute: 78% or 85%?
- *Uncertainty about the accuracy of calculations*. For example, given that this regression analysis requires the inversion of a large, sparse, near-singular matrix, how much faith should be placed in the numerical accuracy of the calculations?

Uncertainty expressed during interpretation

- *Uncertainty about the appropriateness of a descriptive model.* For example, is a linear regression model adequate or should a non-linear one be used? Can Newtonian mechanics be used to describe the path of a comet or are relativistic methods necessary?
- *Uncertainty about the appropriateness of a normative model*. For example, should a statistical analysis follow a Bayesian approach with subjective probabilities, etc, or should frequentist hypothesis testing be used? Should a decision be guided by multi-attribute value analysis or an outranking approach?
- *Uncertainty about the depth to which to conduct an analysis*. It is possible to keep refining the models used in an analysis, introducing more and more subtleties. When have sufficient been introduced? When has the analysis reached a point of sufficient sophistication?

I doubt whether this list of categories fully spans the possibilities. Nor is it the only possible taxonomy. Berkeley and Humphreys[9] arrive at seven categories. Moreover, the categories here are not mutually exclusive. But they will, I believe, serve the discussion adequately.

One point that should be noted is that by discussing the modelling of uncertainty within an analysis I am recognizing the context dependence of the process. Imprecision and uncertainty are context dependent[3].

MODELLING AND ANALYSIS OF UNCERTAINTY

How might each of the types of uncertainty above be reflected in the analysis, beginning with those that are expressed about the modelling step?

The modelling step

Uncertainty about what might happen or what can be done. Many of our clients, especially decision makers and safety engineers, spend much of their time worrying about what might happen: not, in the first instance at least, how likely something is to happen; but simply what are the possibilities. Have all eventualities been anticipated and thought through? One of the first questions we are asked, when we try to help structure and analyse their problem, is: how can we be sure that nothing has been overlooked? Are all possible outcomes or states of nature included in the analysis? Is the fault tree complete?

This always seems to me to be one of the questions on which our normative and descriptive models are unable to help. Conceptually, they cannot. Models capture aspects of the clients' (*not* the analyst's) perception. Models cannot capture that which is not perceived. One cannot quantify/model uncertainty about the unimagined. As analysts we can help our clients use various forms of brainstorming and other tools to help widen their perception and tap into their imagination, but at the end of the day we can only model and explore what they perceive. That being said, the very process of modelling can be a powerful catalyst to the imagination, one that brings to the surface possible outcomes which had not been thought of until the analysis began. A good analyst continually asks about the complement of all the events specified up to the current point in the discussions. None the less, it is the clients' imagination, not the model, which identifies further possibilities. Similar comments may be made about client's concerns that they have not thought of all possible strategies.

Thus, I believe that this is one type of uncertainty which cannot be modelled. One simply has to live with it.

Uncertainty about meaning/ambiguity. Zadeh[10] wrote:

> 'More often than not, the classes of objects encountered in the real world do not have precisely defined criteria of membership. For example, the class of animals clearly includes dogs, horses, birds, etc, as its members, and clearly excludes such objects as rocks, fluids, plants, etc. However, such objects as starfish, bacteria, etc, have an ambiguous status with respect to the class of animals.'

So began a vast and growing body of work on fuzzy mathematics, possibility theory and, generally, the modelling of imprecision and ambiguity.

I have written at length elsewhere about my concerns at the assumptions or, rather, lack of explicit assumptions on which many of these theories are based[3,11–13] (see also Walley[14]). I shall not repeat those points here. My concern is, rather, to ask why we might wish to model imprecision and ambiguity.

If the purpose is to build a descriptive model of the ambiguity and imprecision present in a third party's statements, then I have no quarrel with the objectives behind these theories. A 'third party' is someone distinct from the analyst and the clients. Exploring the implications of such models adds to the clients' understanding and perception of the behaviour of others. There are, of course, questions about the appropriateness of any particular model as a description of a particular set of statements, but, in essence, this is no different to the appropriateness of any descriptive model: see later.

However, I do question the need for *normative* models of ambiguity and imprecision. An analyst's purpose in using normative modelling is to help the clients understand and explore their beliefs, perceptions and preferences in relation to the issues before them and to help their judgements evolve. Analysis seeks to bring understanding: in prescriptive analysis, normative modelling seeks to bring the clients understanding of themselves, their judgements and the implications of their judgements. How can any methodology with an emphasis on modelling, rather than resolving, ambiguity and imprecision, serve this aim?

Consider the election example referred to earlier. Suppose that a company director is concerned about the prospects of a conservative victory in a coming election, but is unclear about precisely what she means by 'conservative'. Suppose further that the outcome of the election matters to her in deciding company policy. Instead of modelling her lack of clarity in her use of 'conservative', a decision analyst should explore with her why a 'conservative' victory matters. Rather than deal with nebulous political concepts, the analyst should try to identify clearly defined observable events, such as changes to the law or taxation policies, which a conservative government might introduce and which would affect her company's business.

Similarly, if a client says that he wants to choose the most cost-effective course of action but is not quite sure of what he means by 'cost-effective', will modelling and describing his lack of clarity help his thoughts become clearer? Surely, it is better to provide methods to explore the meaning of concepts such as cost-effectiveness and to help him understand them in the context of his evolving perception of the issues facing him? Approaches such as that described in Keeney's recent book[15] help analysts meet clients' needs. Normative models of imprecision and ambiguity do not.

Uncertainty about related decisions. No decision is taken in isolation. Other decisions can often influence the circumstances in which a decision is taken or the outcomes to which it may lead. Uncertainty about the effects of other decisions can have a number of impacts on the modelling.

Firstly, will the related decisions be taken by those responsible for the current decision or by others? If the analyst is supporting a coherent group of decision makers, uncertainty about related decisions that they themselves may make should be handled by including these decisions in a decision tree, influence diagram or other method of representing interrelated decisions so that interactions can be investigated via dynamic programming. Of course this is an ideal. There may be too many related decisions to introduce explicitly into the analysis. Brown[16] suggests how allowance might be made for future decisions in the modelling of consequences. If the decision will be affected by the decisions of third parties, then uncertainty about their actions is no different to uncertainty about a future event and can be modelled accordingly.

If others are involved, the form of the models may be drawn from the literature of game theory. But there is a need for care. Is the model to be built descriptive or normative? In the descriptive case: are several decision makers involved? If so, one is led into the sort of models developed within economics and game theory[17,18]. Simple maximin approaches may form poor descriptive models, but more complex models may capture facets of actual behaviour[19].

If the model is normative, as perhaps might be used in bargaining or negotiating, the analyst needs to be clear who is the client. When the analyst is employed by one side, a subgroup of those involved, then the form of model may well involve the representation of uncertainty about what the others may do. But in the case that the whole group of decision makers are the clients and the analyst is essentially an arbiter or negotiation facilitator[20,21], it may be unwise to model all the uncertainties. In negotiation, much of the uncertainty arises from lack of trust between the various participants. Modelling that uncertainty will do nothing to reduce that lack of trust. Much as the argument in the previous subsection went, so here it is inappropriate to model this uncertainty. One wishes to reduce it by the creation of contracts to which all parties can agree.

The exploration step

Uncertainty arising from physical randomness or the lack of knowledge. This is the form of uncertainty most—and hotly—debated by statisticians, philosophers and others. Yet it is the one on which there is, in a sense most agreement. There seems to be general agreement that uncertainty arising from physical randomness or lack of knowledge should be modelled using mathematical probability. Our differences concern whether probability has some physical existence as in the propensity theories of Popper[22], whether it is a subjective construct with no objective existence[23,24], whether it is a property of populations and infinite sequences of repeated trials[25] or whatever.

Subjectivists use probability in normative modelling to explore the consistency of beliefs and their evolution in the light of data. Some subjectivists effectively deny the possibility of using probability in descriptive modelling of physical randomness, arguing that 'physical randomness' is a construct arising from exchangeability of certain events in their or their clients' perception. Conversely, objectivists deny the value of probability in normative modelling, arguing that only physical randomness can be modelled. Most of us sit pragmatically somewhere between these two extremes.

One need only consider descriptive models of physical randomness: normative models would seek to prescribe what physical randomness should be—a somewhat grandiose pastime for mere human beings.

Uncertainty about the evolution of future beliefs and preferences. We all expect our beliefs and preferences to change with time. If today (early 1994) I am asked to predict the likelihood of rain on 1 July, 1995 at noon in the centre of Leeds, I might suggest a probability of 0.20. Ask me for the same prediction on 30 June, 1995 at 6.00 pm and I will give you another value based upon my better knowledge of the then recent weather. Ask me at 11.59 am on 1 July, 1995 when I am in the centre of Leeds and I will give you yet another value, much closer to 0.0 or 1.0 depending on how dry or wet I am. As time passes, we acquire new information and our beliefs change.

Similarly, our preferences may change. Assess my utility function for money today and you will discover a certain attitude to risk. In a few years when my fortunes have changed, my risk attitude is also likely to have changed and, in consequence, assessing my utility function then will produce a different result. Working with companies, one expects that they will place different relative importance on short and long term profits depending on their current assets and trading position and on the economic context.

So we expect our beliefs and preferences to change; but how do we model that change and our uncertainty about it? This seems to me to be an area that needs much more research and thought. Apart from a few isolated papers, such as those of Goldstein[26] and Kornbluth[27], there has been relatively little consideration of the problem. Certainly there is a need for normative models to help analyses of strategies which may lead to consequences in the distant future. Consider, for instance, the problem of disposing of nuclear waste. Water seeping into a storage cavern may cause a radiation leak in a century or two: a problem for a future generation with, one presumes, far greater scientific and medical knowledge as well as different values.

Uncertainty about judgements, for example, of belief and preference. The normative models within a prescriptive analysis require judgemental input: subjective probabilites, utilities, weights or whatever. No client is ever happy giving these judgements exactly. There is always an element of uncertainty. Moreover, the uncertainty is one that is not present within the normative model. Savage's theory of subjective expected utility suggests, for instance, that each of us has within us an exact subjective probability for each possible event in the small world under consideration. I certainly do not. Even within the sanitized small world which I use to think about such questions as 'Will it rain in Leeds at noon on 15 December, 1997?' or 'Did it on 15 December, 1897', I cannot give the probability as anything more than about 55%. I would be happy using 53% as the value if you asked me. But I would balk at 78%.

Some normative theories attempt to avoid this problem by introducing axiom systems which effectively lead to interval methods: see Walley[14] for a recent survey. But intervals have hard upper and lower limits and these are as hard to set as the original single number judgements. Even fuzzy sets cannot avoid the problem by the use of set membership functions, because these too must be defined precisely[11].

Moreover, the use of intervals loses the import of Savage's axioms. In an ideal world, I do want to exhibit the coherence implicit in his axioms. Within my finite cognitive powers, however, I am unable to. I have given an example in which a support logic seeking to use interval valued probabilities arrives at an interval for the probability of interest, which prohibits some values that are compatible with precise probability logics and allows some values which are not[28]. Essentially, methods which try to handle intervals instead of precise quantities may lose structural information. For instance, in calculations where they hold each quantity within its interval, they do not ensure that it takes the same value throughout the calculation. Such structural information can be maintained if symbolic values are introduced along with bounds on their values: e.g. if a sensitivity analysis is conducted upon a precise probability model.

Note that sensitivity analysis does not require precise bounds on the ranges of judgemental quantities. The output takes the form: 'The decision (or inference) would be unchanged unless p exceeds 0.63. Does that concern you?'. The clients are not asked to give precise values at which their concern would begin[3,4,29]; see also Lavine[30], Wasserman and Kadane[31] and Rios Insua[32].

Uncertainty about the accuracy of calculations. We often forget the uncertainties introduced by calculation. More and more our analyses rely on the computation of multidimensional integrals or the location of global optima of non-convex functions. That we can even attempt such things is a tribute to the power of modern computers; but we all know we would be fools to take the results as 100% accurate. The algorithms used are iterative and may involve Monte Carlo sampling. Their convergence is not guaranteed. Thus, as analysts we need to consider the uncertainty introduced by the process of calculation. Usually, we can reassure ourselves that it is orders of magnitude less than the other uncertainties in our problems: usually, but not always. O'Hagan[33] suggested some ways of modelling this uncertainty, albeit to help improve the iterative algorithms; Mockus[34] discussed the Bayesian approaches to reflecting uncertainty about global optima.

The interpretation step

Uncertainty about the appropriateness of a descriptive model. Statistical inference is, in large measure, concerned with fitting descriptive models to data, either to summarize the regularity in the data in the form of a physical model or to serve as a basis for prediciting future events. Seldom, if ever, is there a single model which one might consider fitting to the data. A variety of models may be proposed and a choice has to be made between them. Various schools of statistical inference have suggested ways of doing this: hypothesis testing, analysis of variance, mixtures of models, etc. But since no one has produced a definitive answer, the user is always left with a feeling of uncertainty: would it be better or more correct to use an alternative model? Perhaps some of this residual uncertainty can be modelled by introducing a modelling error term into the analysis[35,36]. Although this may have considerable advantages in 'smoothing' by introducing a subtle correlation structure, all it does conceptually is make the descriptive model a little more sophisticated; and there is then further uncertainty about whether this model is appropriate. This further uncertainty may in turn be modelled, and an infinite regress constructed with remarkable ease.

One would be wise to remember the writings of Savage[1] on 'small worlds'. Any model is an abstraction or a limited view of the world. There will always be differences between it and reality. Thus, this form of uncertainty cannot be avoided: it certainly cannot be modelled without recreating it. The analyst and clients simply have to live with this uncertainty.

Uncertainty about the appropriateness of a normative model. In any prescriptive analysis, the client and analyst together have to choose a form of normative model to use in the analysis: e.g. Bayesian or non-Bayesian. At least ideally they do: most analysts—and I number myself among them—are so wedded to a particular school of normative modelling that they seldom admit to the client that there is a choice to make. In which case the client, wittingly or unwittingly, chooses the form of normative model in choosing the analyst. But to return to the ideal: if the clients are aware that there is a choice of normative model to be made, having made it they may have some residual uncertainty. Would it have been better to have adopted a different approach? Should this uncertainty be modelled?

I do not see how it can be.

Normative models are models of how we might think, choose or act. Clients may be unsure about which model to adopt as a guide to their thinking. That lack of conviction is a statement of their indecision. It may arise because each client is undecided; or because each is sure for him or herself, but as a group they differ. In either case modelling that indecision will not help them resolve it: cf. the earlier remarks on ambiguity. In any case, how can one model this indecision? To do so requires a further normative model and that again leads quickly to an infinite regress. One can no more climb out of the small world of a normative model than one can out of that of a descriptive model.

All one can do is to explore the implications of each normative model and see where they differ . . . if, indeed, they do. Those who work regularly with groups of decision makers know that conceptual differences are often entirely diffused on discovering that in a particular context they lead to no difference whatsoever in the suggested course of action. Even if there are differences, exploring the issues from different normative perspectives brings insight and fosters communication, allowing the group to reach a decision.

Uncertainty about the depth to which to conduct a prescriptive analysis. When does one stop an analysis? Phillips[2] has argued that one does so when it is *requisite*: when taking the analysis any further does not promise to bring further insights and when the clients are comfortable with its conclusions.

If one recognizes the interpretation stage of analysis—the point at which one has to move from the small worlds of the models back to the real world of doing things—then it is hard to avoid Phillips' conclusion. Anything else brings up one of those infinite regresses. Thus, Phillips argues that the models should be explored by means of sensitivity analyses, by residual plots, by using the methods proposed by Box[37]. If the exploration causes the clients to doubt the models, to feel that something has been omitted, then the models should be refined: and the cycle of exploration and refinement repeated until no further insights are gained and no new doubts arise. Or rather, until the clients can live with the residual level of doubt.

CONCLUDING REMARKS

My purpose in this paper has not been to say anything new: Christer[38], for instance, has discussed some of these ideas in an OR context. Rather, it has been to emphasize the complexity of the contexts in which one might wish to model ambiguity, imprecision and uncertainty and to suggest that the modelling approach appropriate to one set of circumstances may be inappropriate to another. How one models uncertainty should depend on the reason for doing so. Moreover, in prescriptive analyses there are some forms of uncertainty which should not be modelled. Rather, the analysis should seek to resolve or reduce the uncertainty through its modelling and exploration of other aspects of the problem.

I am aware that I have written the above from a Bayesian viewpoint, but I hope that the issues raised give pause for thought for those from other schools. My prime objective is to stimulate discussion.

Acknowledgements—Many people have helped shape my thoughts— but probably wouldn't wish to admit it having seen the above. I owe an enormous debt of gratitude for discussions to: Roger Cooke, Dennis Lindley, Freerk Lootsma, Larry Phillips, David Rios Insua and my colleagues at Leeds. The British SERC and NATO have provided financial support for related work on sensitivity analysis. An earlier version of this paper was presented at the 4th International Meeting of Statistics in the Basque Country. I am also grateful for the comments from two anonymous referees.

REFERENCES

1. L. J. Savage (1972) *The Foundations of Statistics*, 2nd Edn. Dover, New York.
2. L. D. Phillips (1984) A theory of requisite decision models. *Acta Psychologica* **56**, 29–48.
3. S. French (1986) *Decision Theory: an Introduction to the Mathematics of Rationality*. Ellis Horwood, Chichester.
4. S. French (Ed) (1989) *Readings in Decision Analysis*. Chapman and Hall, London.
5. D. E. Bell, H. Raiffa and A. Tversky (Eds) (1988) *Decision Making: Descriptive, Normative and Prescriptive Interactions*. Cambridge University Press, Cambridge.
6. *OR/MS Today* (1991) June, 9, August, 8–9, October, 8.
7. R. M. Cooke (1991) *Experts in Uncertainty* Oxford University Press.
8. S. French (1985) Group consensus probability distributions: a critical survey. *Bayesian Statistics 2*. In (J. M. Bernardo, M. H. DeGroot, D. V. Lindley and A. F. M. Smith, Eds) pp. 183–202. Oxford University Press, Oxford.
9. D. Berkeley and P. Humphreys (1982) Structuring decision problems and the 'bias heuristic'. *Acta Psychologica* **50**, 201–252.
10. L. Zadeh (1965) Fuzzy sets. *Information and Control* **8**, 338–353.
11. S. French (1984) Fuzzy decision analysis: some criticisms. In *Fuzzy Sets and Decision Analysis*. (H. Zimmermann, L. A. Zadeh and B. R. Gaines, Eds) pp. 29–44. North Holland, Amsterdam.
12. S. French (1989) Fuzzy Sets: the unanswered questions. Research Report 89.6, School of Computer Studies, University of Leeds.
13. S. French (1990) On being Numerate. Inaugral Lecture, *The University of Leeds Rev*. **33**, 119–134.
14. P. Walley (1991) *Reasoning with Imprecise Probabilities*. Chapman and Hall, London.
15. R. L. Keeney, (1992) *Value-Focused Thinking*. Harvard University Press, Cambridge, MA.
16. R. V. Brown (1991) Heresy in decision analysis: modelling subsequent acts without rollback, *Dec. Sci.* **9**, 543–554.

17. M. Bacharach (1976) *Economics and the Theory of Games*. Macmillan, London.
18. L. C. Thomas (1984) *Games: Theory and Applications*. Ellis Horwood, Chichester.
19. A. Colman (1982) *Game Theory and Experimental Games*. Pergamon Press, Oxford.
20. H. Raiffa (1982) *The Science and Art of Negotiation*. Harvard University Press.
21. C. Huxham (1991) Facilitating collaboration: Issues in multi-organizational group decision support in voluntary, informal collaborative settings. *J. Opl Res. Soc.* **42**, 1037–1046.
22. K. Popper (1957) The propensity interpretation of the calculus of probability and the quantum theory. In *Observations and Interpretation*. (S. Körner Ed.) pp. 65–70, Butterworths, London.
23. B. De Finetti (1974) *Theory of Probability*. Vol 1. Wiley, Chichester.
24. B. De Finetti (1975) *Theory of Probability*. Vol 2, Wiley, Chichester.
25. R. Von Mises (1957) *Probability, Statistics and Truth*. 2nd Edn. George, Allen and Unwin, London.
26. M. Goldstein (1985) Temporal coherence. In *Bayesian Statistics 2*. (J. M. Bernardo, M. H. DeGroot, D. V. Lindley and A. F. M. Smith, Eds), pp. 231–248. Oxford University Press, Oxford.
27. J. Kornbluth (1992) Dynamic multi-criteria decison making. *J. Multi-Criteria Decision Analysis* **1**, 81–92.
28. S. French (1991) Recent mathematical developments in decision analysis. *IMA J. Mathematics Applied to Business and Industry* **3**, 1–12.
29. S. French (1992) Mathematical approaches to sensitivity calculations. *J. Opl Res. Soc.* **43**, 813–819.
30. M. Lavine (1991) Sensitivity in Bayesian statistics: the prior and the likelihood. *J. Am. Stat. Assoc.* **86**, 396–399.
31. L. Wasserman and J. B. Kadane (1992) Computing bounds on expectations. *J. Am. Stat. Assoc.* **87**, 516–522.
32. D. Rios Insua (1990) *Sensitivity Analysis in Multi-Objective Decision Making*. Springer Verlag, Berlin.
33. A. O'Hagan (1992) Bayesian numerical analysis. In *Bayesian Statistics 4*. (J. M. Bernardo, J. O. Berger, A. P. Dawid and A. F. M. Smith, Eds) pp. 345–363. Oxford University Press, Oxford.
34. J. Mockus (1989) *Bayesian Approach to Global Optimisation*. Kluwer Academic, Dordrecht.
35. B. J. N. Blight and L. Ott (1975) A Bayesian approach to model inadequacy for polynomial regression *Biometrika* **63**, 79–88.
36. D. Draper (1995) Assessment and propagation of model uncertainty (with discussion). *J. Roy. Stat. Soc.* **B** (in press).
37. G. E. P. Box (1980) Sampling and Bayes inference in scientific modelling and robustness. *J. Roy. Stat. Soc* **A143**, 383–430.
38. A. H. Christer (1975) Innovatory decision making. In *The Role and Effectiveness of Theories of Decision in Practice*. (D. J. White and K. C. Bowen, Eds) pp. 368–377. Hodder and Stoughton, London.

Received June 1993; accepted April 1994 after one revision