



Open

Cynefin: uncertainty, small worlds and scenarios

Simon French*

University of Warwick, Coventry, UK

Uncertainty, its modelling and analysis have been discussed across many literatures including statistics and operational research, knowledge management and philosophy: (i) adherents to Bayesian approaches have usually argued that uncertainty should either be modelled by probabilities or resolved by discussion that clarifies meaning; (ii) some have followed Knight in distinguishing between contexts of risk and of uncertainty: the former admitting modelling and analysis through probability; the latter not; (iii) there are also host of approaches in the literatures stemming from Zadeh's concept of a fuzzy set; (iv) theories of sense-making in the philosophy and management literatures see knowledge and uncertainty as opposite extremes of human understanding and discuss the resolution of uncertainty accordingly. Here I provide a personal perspective, taking a Bayesian stance. However, I adopt a softer position than conventional and recognise the concerns in other approaches. In particular, I use the Cynefin framework of decision contexts to reflect on processes of modelling and analysis in statistical, risk and decision analysis. The approach builds on several recent strands of discussion that argue for a convergence of qualitative scenario planning ideas and more quantitative approaches to analysis. I discuss how these suggestions and discussions relate to some earlier thinking on the methodology of modelling and, in particular, the concept of a 'small world' articulated by Savage.

Journal of the Operational Research Society (2015) 66(10), 1635–1645. doi:10.1057/jors.2015.21

Published online 29 April 2015

Keywords: Cynefin; models; scenarios; small worlds; uncertainty

The online version of this article is available Open Access

Introduction

To this day, I remember the excitement that I felt when I first encountered Bayesian Statistics and Decision Analysis. I found the subjective perspective in which the uncertainty modelled was *my* uncertainty entirely persuasive. The axiomatic bases of probability and utility provided the rigour on which to build quantitative analyses that balanced my uncertainties—or degrees of belief—with my preferences to identify the best inference or course of action. Over the years that view has softened and, influenced by many colleagues, I have come to recognise:

- the variety of forms that uncertainty may take and that not all may or need be modelled by probability, some may need be addressed through sensitivity analysis or resolved through introspection and discussion (French, 1995, 2003);
- the need to balance the harsh clarity of the theory with the limits of human judgement in prescriptive modelling (French and Smith, 1997; French *et al.*, 2009);
- the value of sensitivity analysis in bounding and interpreting the results of an analysis (French, 2003);
- the issues that arise when groups rather than individuals are responsible for inferences and decisions (French *et al.*, 2009; Rios Insua and French, 2010; French, 2011);

- the value of the Cynefin framework in categorising decision contexts and identifying how to address many uncertainties in an analysis (French, 2013).

But I have never really addressed the fundamental question posed by Knight (1921): What do you do in an analysis when an uncertainty is so deep that the range of plausible probabilities that one might use to reflect the views of a group is effective 0–1, meaning that few issues are resolved by an analysis? Knight distinguished circumstances of *Risk*, in which probabilities are known, from those of *Uncertainty*, in which our knowledge of some events or quantities is so meagre that some probabilities are effectively completely unknown. This paper takes a Bayesian perspective to explore analyses in which there are some deep or *Knightian* uncertainties. Sense-making, issue and problem formulation, and the process of modelling will also be major foci. The paper continues the arguments begun in French (2013) (for related discussions, see Cox, 2012; Spiegelhalter and Riesch, 2011).

In the next section I begin the discussion of sense-making, recognising that it takes place as much in our subconscious thoughts and that formalising this process to build models means that we must cross that vague boundary between intuitive thought and formalised analysis. This leads into reflections on the relationship between modelling and analysis, on the one hand, and the real world, whatever that may be, on the other. I then turn to Snowden's Cynefin framework to articulate some further thoughts on the varied contexts of

*Correspondence: Simon French, Department of Statistics, University of Warwick, Gibbet Hill Road, Coventry, Warwickshire CV4 7AL, UK.
E-mail: simon.french@warwick.ac.uk

modelling. Cynefin provides a structure in which to discuss different forms of uncertainty from the deep uncertainty through the growth of knowledge as we learn about the world to stochastic behaviours and randomness. In turn, this will lead us to a discussion of Savage's conception of inference and decision in the face of uncertainty and his introduction of a 'small world' to frame this; and thence to a consideration of whether there is need to consider analyses based on several small worlds rather than just one. We shall discover that while there are ways to develop justifications of the Bayesian models within 'parallel small worlds' and to develop scenario-focused forms of decision analysis, it is not entirely straightforward to do so. Moreover, the modifications required in Savage's approach elucidate the difficulties faced by decision makers in interpreting the output of scenario-focused decision analysis.

Sense-making

Decision making, at least as I understand it, is always a conscious act; unthinking, unconscious choice is not a decision. Hence decisions are invariably preceded by some process of formulation so that the choices are framed sufficiently for the decision makers to be aware of some of the options and able to assess each against their broad values, preferences and uncertainties to be sure that no option stands out as the obvious unconflicted choice (Janis and Mann, 1977). In other words, decision makers need to be aware they face a decision. The cognitive processes by which they frame the choice before them fall into the broader area of sense-making (Weick, 1995; Kurtz and Snowden, 2003).

In many, perhaps most, cases decision making will be intuitive, based on what has become known as *System 1* thinking (Chaiken *et al.*, 1989; Kahneman, 2011). Such forms of thinking tend to be somewhat superficial, using simpler forms of thinking on the fringes or outside of consciousness. System 1 thinking is subject to behavioural biases; indeed, for many years its literature on has been referred under the somewhat pejorative label *heuristics and biases* (Kahneman and Tversky, 1974). In our professional lives, however, we eschew System 1 thinking and adopt more conscious, analytic patterns of thought, known as *System 2* thinking. Oaksford and Chater (2007) make a similar distinction but refer to *Rationality*₁ and *Rationality*₂. In decisions relating to the management of business, industry, communities and society, there is a need for more rational, auditable processes that draw in wider sources of information and evaluate options carefully, attending to details. Thus explicit, analytic System 2 thinking should be the order of the day. It may not be, but it should be. However, whether they use System 1 or System 2 thinking, decision makers must be aware of some options if a *choice* is to be made in any sentient manner.

My concern is to discuss—primarily from a System 2 perspective—the sense-making processes that frame the choice and lead into decision making, particularly how they relate to

the uncertainties, both *aleatory* arising from randomness and *epistemological* arising from lack of knowledge. I shall broaden the discussion to consider statistical inference and risk management processes alongside decision processes, both because of my personal interests and also because I find that the three areas overlap so much that it is difficult to focus on one without reflecting on the others. All require that one develops an understanding of context: what might or might not happen, how much different outcomes matter, what we know and do not know, and so on.

The process of building a picture of the real world through modelling is discussed in several places. There are, for instance, the seminal texts of Ackoff (1962), Churchman (1971), Pidd (1996), Tukey (1977) and White (1975, 1985). The *Journal of the Operational Research Society* has had a tradition of publishing articles on operational research (OR) methodology and philosophy, which include many on the process of problem formulation and modelling. Moreover, the literature on soft systems and soft OR focuses on the sense-making processes (see, eg, Checkland, 2001; Rosenhead and Mingers, 2001; Shaw *et al.*, 2006, 2007). Knowledge management has a long literature on sense-making too (see, eg, Weick, 1995; Kurtz and Snowden, 2003). Notwithstanding these remarks, many discussions of statistical, risk and decision analyses begin with a putative model: maybe quite a generic model, but a model, nonetheless. A collection of well-defined entities, stimuli, relationships and behaviours observed 'out there' in the real world are taken as the starting point. Entities are quickly labelled by variables; stimuli, relationships and behaviours represented by functions. Uncertainties may be recognised and probability models introduced to represent some of these: stochastic behaviour, observational errors, modelling errors and so on.

Note that I am somewhat catholic in what I mean by a 'model'. In most cases, I mean a mathematical relationship; but sometimes the model might be implicit in a computer code, perhaps a simulation of actors and their interactions. Whatever the case, in modelling we focus on a simplified part of reality, which Savage (1972) dubbed a *small world* and which can be represented intuitively by the model. My objective is to discuss processes of *focusing onto* or *constructing* the small world that will form the backdrop for an analysis. I want to ask how small that world can be while still supporting the purpose of the analysis. I also want to reflect on whether we should analyse in the context of one small world or whether several small worlds might better serve our needs.

Modelling and analysis

Discussing the relationship between our understanding of the real world and of modelling and analysis, and of how conducting the latter informs our learning, risk management and decision making inevitably takes us nearer to philosophy than mathematics. Philosophers since the earliest times have debated the so-called mind-body problem, which concerns how our

mental lives, thinking and knowledge relate to the external physical world. Some extreme subjectivists develop their conception of thinking and knowledge without postulating the existence of any real world, arguing that all we can do is seek to represent relationships between our perceptions and stimuli. Although a subjectivist, I am not that extreme and I shall be concerned with our understanding of the external world and how modelling and analysis can guide our actions within it. But I recognise that philosophers have debated the mind-body problem, knowledge and uncertainty for millennia without reaching consensus. Thus much of the following is, at best, a pragmatic view; at worst, personal prejudice.

Figure 1 is typical of many appearing in texts that discuss the relationship between modelling, and analysis and induction. The left-hand side indicates the modelling process in which we first focus on a small part of the real world that we perceive to be of concern, that is, an abstraction from the complexity and detail of the real world that has in its essence all that is relevant to the issues that are being modelled. Of course, in being able to separate out a small world from reality and discuss it, along with behaviours within it, we are effectively forming a model, at least in terms of a broad description. But the models that will concern us are more conceptual and mathematical, and, while mirroring those small world behaviours, are amenable to analysis. These behaviours may be those that we perceive 'out there' in the small world. In such cases we build a purely *descriptive* model. In statistical, risk and decision modelling, however, we sometimes include ourselves in the model and assume that our behaviours are idealised in some sense: that is, we assume that we use System 2 thinking based on conceptions of rational, analytic behaviour, so building a *prescriptive* model to guide our inferences, choices and subsequent behaviours. French *et al* (2009) discuss prescriptive modelling in detail (for related discussions, see Phillips, 1984; Bell *et al*, 1988; Edwards *et al*, 2007). Note also that in more sophisticated studies we seldom use a single model, but a family of models representing different perspectives. Multiple explorations within these models enable us to gain an intuition for how the inputs and outputs are related; and we then broaden this intuition to help us understand the real world—or at least those aspects of the real world that lie within the the small world.

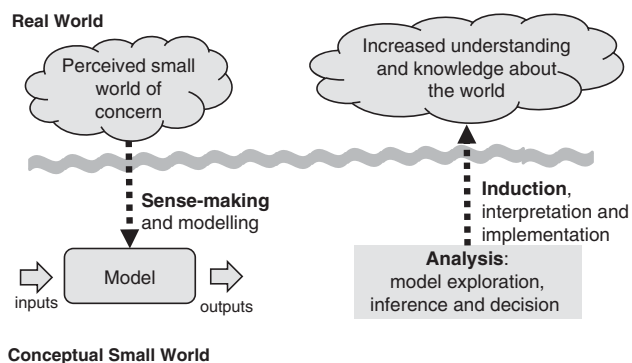


Figure 1 Modelling, analysis and induction.

The right-hand side of Figure 1 represents the step back to the real world on in which understandings of behaviours in the models to *induce* a greater intuitive understanding of the real world. We mean not just that we infer the values of some parameters or derive a hard prescription of what to do, but that we build a wider understanding of the objects and behaviours in the world, how they interact and, in cases where a decision is to be made, we understand better what to do.

This induction step inevitably brings with it uncertainties that arise because the model is not a perfect representation of the small world and that in focusing on the small world some other relevant part of reality may have been ignored. OR, risk and decision studies usually include implementation phases and so face the harsh auditing that the future will bring. Thus, it is usually recognised that actual behaviours may depart from those anticipated in the modelling: that is, this induction step is one that will be accompanied by uncertainty. Professional statisticians too recognise the existence of modelling error, that is, the discrepancies between model and real world behaviours. Too many studies within the applied science and social sciences, however, are published by authors and editors who believe that, say, a 95% confidence interval—even a Bayesian one—relates to a precise 0.95 probability that covers all the potential for error. They do not recognise that the inductive step necessarily introduces further uncertainty. Policy and decision makers, also, can have a tendency to 'believe in the model' too much and be disappointed by what actually happens (see French and Nicolae, 2005 for a discussion of this in the context of crisis management).

For the purposes of our discussion, we will consider three major phases in conducting analyses (cf Holtzman, 1989; French *et al*, 2009).

Sense-making: The process begins with sense-making and modelling in which the context and issues of concern are identified and formulated as models. This phase relates to the dotted downward arrow on the left of Figure 1.

Analysis: In this phase the models are explored and analysed to build an understanding of the behaviours exhibited by the models. This phase relates to the calculations, explorations and studies that take place in the conceptual world at the bottom of Figure 1.

Induction: Through a process of induction the understandings of behaviours within the model are developed into understandings of behaviours in the real world, thus interpreting the results of the analysis and allowing the conclusions to be implemented. This phase relates to the dotted upward arrow in Figure 1.

The overall process is seldom as unidirectional as presented here, but may iterate with the model being elaborated as understanding grows.

Many different types of uncertainty need to be addressed in this process. Table 1 provides a categorisation of these. Note that relating each uncertainty type to the phase of the analytic process encourages an action perspective on how to address and

Table 1 Different forms of uncertainty arising in an analysis

<i>Sense-making</i>	<ul style="list-style-type: none"> ● Uncertainty about meaning/ambiguity ● Uncertainty about what might happen (the science) ● Uncertainty about how much impacts matter (values) ● Uncertainty about related decisions
<i>Analysis</i>	<ul style="list-style-type: none"> ● Uncertainty because of physical randomness ● Uncertainty because of lack of knowledge ● Uncertainty about the evolution of future beliefs and values ● Uncertainty about the accuracy of calculations
<i>Induction</i>	<ul style="list-style-type: none"> ● Uncertainty about the appropriateness of descriptive model (how well we have explained the world) ● Uncertainty about the appropriateness of normative model (principles of modelling beliefs and values) ● Uncertainty about depth to which to conduct an analysis

Source: French (1995)

deal with each category; it is not sufficient just to label them. For a discussion of the majority of these uncertainty types, see French (1995); here we shall discuss the deep uncertainties implicit in some of those in the first and third phases.

Given that statistical, risk and decision analyses are about the development, validation and use of knowledge, there is surprising little cross fertilisation with concepts and perspectives from the literature of knowledge management. Knowledge and uncertainty are polar opposites: the more knowledge we have, the less uncertainty, and vice versa. In French *et al* (2009) and French (2013) we explore some overlaps between these literatures. Snowden's Cynefin framework is particularly informative. He introduced Cynefin to categorise contexts for inference and learning, knowledge management and decision making. Cynefin, while saying little that is new, provides an intuitive backdrop for discussing many analytical processes. I shall use it here to articulate our discussion of small worlds and scenarios. The next section offers a brief introduction to Cynefin and its concepts.

Cynefin: a context for our discussion

Cynefin, see Figure 2, identifies four different, but not entirely distinct contexts for inference and decision. These should not be thought of as providing a hard categorisation; the boundaries are soft and contexts lying near these have characteristics drawn from both sides. But taken with a suitably large 'pinch of salt', Cynefin will serve our discussion well.

The four categories identified by Cynefin are: the Chaotic, Complex, Knowable and Known Spaces. When contexts lie in the *Chaotic Space*, we are unfamiliar with more or less everything. We receive stimuli, but can see no pattern or relationship between them. We cannot yet discern entities, events,




Figure 2 The Cynefin model (Snowden, 2002).

behaviours and so on. So we observe, we act tentatively, 'prodding' where we can to see what happens. Eventually we begin to make sense of things: we see entities and behaviours, we recognise events. As yet we cannot discern any cause and effect relationships. Gradually, though, we do identify putative causes and putative effects. We cannot say that they hold with any certainty, but we recognise potential causes for some effects. Now the context is said to lie in the *Complex Space*, also known as the *Realm of Social Systems*, because typically cause and effect are very difficult to relate with any confidence in such systems. For instance, as I write this, we may be able to identify a number of potential causes that would lead to Greece leaving the Eurozone, but we understand none of them with sufficient certainty to make a confident prediction of whether Greece will be in the Eurozone at the end of 2016. Perhaps a few years later, we will be able to look back and explain what happened and why, but we will need the 20–20 vision of hindsight for that.

Over time, though, as we observe more, for *some* behaviours we see more clearly how the causes and effects are related. We can begin to set up controlled trials to confirm our suspicions. Eventually we are confident in our understanding of cause and effects: we develop scientific laws to encapsulate this understanding. Such behaviours are recategorised as lying in the *Knowable Space*. This space describes contexts in which we have sufficient understanding to build models, though not enough to define all the parameters within those models. For any application of the model we need to collect data and analyse these to estimate the parameters. But again over time, we may gain sufficient experience that we know the parameters well enough for all applications that further data gathering is unnecessary. In this case, the context is recategorised to the *Known Space*, recognising that we fully understand and can predict cause and effect.

In this description of learning, knowledge increases in an orderly, chronological fashion from the *Chaotic Space* through the *Complex* and *Knowable Spaces* to arrive at the *Known Space*. That is, of course, idealised. At any time, as we look at the world some entities and behaviours lie in each space,

recognising that we have learnt nothing about a few, something about some and a lot about others. Moreover, it would be good if progress were always clockwise as shown, but inevitably we get things wrong on occasion and perceive cause and effect where there are none, later learning our mistake and moving back through Cynefin anti-clockwise. In extreme cases, Kuhn (1970) might term such anti-clockwise reversions a paradigm shift.

Almost all the analytic tools used in statistical, operational and risk analysis require that we are working in the Known and Knowable Spaces; this must be the case for they are based on models that assume an understanding of cause and effect. The exceptions to this are techniques such as exploratory data analysis, multivariate statistics, data mining, soft systems and soft OR methods that are designed to catalyse and support processes of sense-making.

There are many caveats that we should make—more than we admit here (see French, 2013 for further discussion).

- Even in the Known Space our uncertainty is not quite zero. We must always admit the possibility that our world may change and our understandings that have served us well in the past no longer apply. Just because the Sun has risen every day in our lives does not mean it will do so tomorrow. Nonetheless, we proceed on the assumption that it will, planning our lives around tomorrow's dawn. Similarly, we accept Newton's Laws of Motion and other well-tried and tested scientific laws without question and ignore the uncertainty that they may cease to hold. Moreover, we accept and live with the uncertainties noted in Table 1.
- We should note that repetition is central to our thinking about the Known and Knowable Spaces. In these cause and effect are understood. In other words, we have experienced the circumstances often enough to understand how different causes or different levels of a cause lead to different effects. We often express this understanding through a scientific law or model, which we validate by repeatedly testing them under controlled circumstances until we are sure that they predict effects from a given set of causes. Repetition is central to the Scientific Method, which expects scientific experiments to be repeatable. This focus on repetition led naturally to the development of the frequency concept of probability and frequentist statistics (French, 2013). It is also worth noting that repetition is also important in thinking about our values. If we have experienced a situation repeatedly, we know what the possible outcomes are and how they impact on us. We do not have to think through and judge how we will feel in completely novel circumstances (French, 2013).
- One should be careful to avoid terminological confusions with complexity science and the Complex Space. Complexity science is concerned with computational issues relating to highly complicated models. Such models and computational issues belong more to Knowable and Known Spaces rather than the Complex.

Our concern in this paper is to discuss how we move from the Complex Space to the Knowable Space and how the uncertainties that we encounter are managed and modelled. Our perceptions of behaviours in the Complex Space recognises entities, events and some putative relationships, but only vaguely, not in sufficient detail to model in anything but a rudimentary manner. We face many uncertainties, some nebulous, too deep to be modelled in a formal sense. As our knowledge and understanding increase, as we approach the boundary between the Complex and the Knowable, we may have a putative model that does capture our broad understanding of cause and effect, but some uncertainties may remain so deep that we cannot usefully encode them as probabilities. Even when conceptually we agree on the structure of probabilities within the model, we may disagree on some of their values, allowing ranges that are effectively 0–1. They remain deep uncertainties. Over time, further observations, experiences and insights bring us much clearer perceptions, ones that we can model in detail and move into the Knowable Space. Uncertainties may indeed will remain, but they can be modelled and analysed in structured, formalised ways, either through probabilities whose values are agreed to lie within a sufficiently small range that they can be analysed through sensitivity and robustness studies.

As I have indicated, once the deep uncertainties have been resolved and we are safely in the Knowable Space, I believe that the Bayesian subjective expected utility model provides the appropriate methodology to articulate, analyse and address uncertainty. The concern of this paper is to discuss in a little more depth how that model might arise as knowledge accumulates sufficiently to move from the Complex to the Knowable Space, and how recent developments in scenario-focused thinking combined with the Bayesian model might provide a methodology to support this process. To do that we need look a little more closely at Savage's thinking on the Bayesian model.

Small worlds and the framing of statistical inference and decision analysis

There are many axiomatic developments of Bayesian subjective expected utility (see French and Rios Insua, 2000 for a survey). We begin by focusing on Savage's development because his approach introduced the concept of a *small world* and, moreover, he discussed in some depth how this abstraction related to reality, and thus how the modelling and analysis could inform inference and decision. Savage's concept of a small world is effectively a single model encoding ideas of cause and effect.

Savage (1972) discussed his concept of a small world in his 1954 monograph. He imagined a decision maker facing a choice that is described by the small world. In a sense, his conception differs from that shown in Figure 1, in that his small world is effectively a mathematical model, whereas in the figure a small world is shown as something more nebulous, a perspective on a part of reality before a model is constructed.

However, the difference is more one of terminology than a real difference of meaning. As Wittgenstein (1921) argued, the use of propositional logic within language acts as model for the part of reality being described or discussed; and the step from propositional logic to a mathematical model is but a small one. Savage's fundamental model relates to a triple $\{\Theta, C, \mathcal{F}\}$:

$$\Theta = \{\theta | \theta \text{ is a state of the world}\}$$

$$C = \{c | c \text{ is a consequence}\}$$

$$\mathcal{F} = \{f: \Theta \rightarrow C | f \text{ is an act which the DM can choose}\}$$

I make no apology for introducing mathematical notation here, though we shall use it little, because its introduction makes quite clear that we are now in the land of mathematical models.

A state of the world is a possible description of the current situation with all uncertainties resolved. Thus Θ is a set of possible descriptions that spans all possibilities. However great our uncertainty, the decision maker is sure that one of the descriptions in Θ is true. The set of consequences C contains all possible outcomes that may arise from the decision-maker's acts and the set \mathcal{F} contains all possible acts, that is, each act relates outcomes to each possible state of the world. Savage modelled acts as functions from Θ to C and he included in \mathcal{F} all conceivable functions. For Savage, the triple $\{\Theta, C, \mathcal{F}\}$ was the small world in which all further analysis was focused. It should be a microcosm in which analysis is possible and relevant to our concerns and understanding of the real world. Note that the small world $\{\Theta, C, \mathcal{F}\}$ encodes the decision-maker's perception of cause and effect. This means that the development of the small world must take place in the context of the Knowable or Known Spaces, almost invariably the former.

Savage further suggested seven postulates, which encode the rationality that the decision maker might demand of her preferences. He showed that these postulates led inexorably to the Bayesian model: the decision maker within the small world should choose as if she had a subjective probability distribution representing her beliefs, a utility function representing her preferences between consequences and then rank the acts according to expected utilities. Since Savage's development, there have been many alternative derivations of the Bayesian model from a set of postulates or axioms, some more constructive, separating the axiomatisation of the decision-maker's beliefs over Θ from the axiomatisation of her preferences over C . Most effectively take $\{\Theta, C, \mathcal{F}\}$ as the small world in which analyses are conducted. Some, however, recognise explicitly that the small world needs a model of the decision maker as well as a model of her external world and include the decision-maker's preference relation, \succ , between acts within the definition taking $\{\Theta, C, \mathcal{F}, \succ\}$ as the small world. I concur with this view, as I take the use of a normative model such as Savage's within a prescriptive analysis as providing a model of how a perfectly

rational decision maker with beliefs and preferences similar to mine would decide in a simplified decision problem, which parallels the one that I face (French, 1986; French *et al.*, 2009). Shafer in his 1986 retrospective on Savage's book takes a similar view describing a prescriptive analysis based on a normative model as providing an 'argument by analogy' (see also Goldstein, 2011).

Essentially, a small world plus the postulates define a model. So we often refer to a small world as model, smearing the distinction implied in Figure 1. Thus we shall write:

$$M = \{\Theta, C, \mathcal{F}, \succ\}.$$

How big should a small world or model be? How much detail should be included? These were questions that Savage worried at but did not resolve. He recognised that if the small world was too small, then any analysis would be too limited to inform the decision maker. But he also recognised that the *grand world*, which included all future conceivable events and possible acts in the decision-maker's future was much too big to analyse, writing:

The point of view under discussion may be symbolised by the proverb 'Look before you leap' and the one to which it is opposed by the proverb 'You can cross that bridge when you come to it.' When two proverbs conflict in this way, it is proverbially true that there is some truth in both of them, but rarely, if ever, can their common truth be captured by a single pat proverb. One must indeed look before he leaps, in so far as the looking is not unreasonably time-consuming and otherwise expensive; but there are innumerable bridges one cannot afford to cross unless he happens to come to them.

Carried to its logical extreme, the 'Look before you leap' principle demands that one envisage every conceivable policy for the government of his whole life (at least from now on) in its most minute details, in the light of the vast number of unknown states of the world, and decide here and now on one policy. This is utterly ridiculous, not—as some might think—because there might latter be cause for regret, if things did not turn out as had been anticipated, but because the task implied in making such a decision is not even remotely resembled by human possibility. It is even utterly beyond our power to plan a picnic or to play a game of chess in accordance with the principle, even when the world of states and the set of available acts to be envisaged are artificially reduced to the narrowest reasonable limits.

Though the 'Look before you leap' principle is preposterous if carried to the extremes, I would none the less argue that is the proper subject of our further discussion, because to cross one's bridges when one comes to them means to attack relatively simple problems of decision by artificially confining attention to so small a world that the 'Look before you leap' principle can be applied there. I am unable to formulate criteria for selecting these small worlds and indeed believe that their selection may be a matter of judgement and experience about which it is impossible to enunciate complete and sharply defined general principles though something more will be said in this connection in §5.5. On the other hand, it is an operation in which we all necessarily have

much experience, and one in which there is in practice considerable agreement. (Savage, 1972, pp 16–17)

Shortly after he says, ‘... I find it difficult to say with any completeness how such isolated situations are actually arrived at and justified’. He then rehearses an argument very similar in flavour to one picked up and extended by Phillips (1984) in developing the theory of requisite decision modelling. Using too small a small world can lead to difficulties in analysis. Bordley and Hazen (1992) show that too small a world can miss correlations and in the presence of dependent multi-attributed preferences lead to apparent ‘irrationalities’. French *et al* (1997) show a similar effect can arise if preferences depend on the resolution of some key event. One can also argue that Allais’ and similar paradoxes arise because the choices are stated too simplistically (French and Xie, 1994).

Savage, unaware naturally of these later writings, approached the issue of how small a small world should be by considering the consistency needed in a sequence of small worlds, each more complex than and containing the previous one. The events in one small world were a set of events in a larger small world containing it; and the largest small world was his grand world. Table 2 gives a simple example with three nested models. The largest model, M_1 , is Savage’s grand world and represents the decision-maker’s best understanding of the part of the Universe on which he is focusing. M_3 is an approximation to this in which the calculations are at least conceptually possible. In the case of highly complex models, it may be possible to evaluate M_3 at given points, but only at great cost and with long calculation times. So M_3 is an emulation of M_2 , which is much more tractable and allows cost-effective evaluation (O’Hagan, 2006; Rougier *et al*, 2009; Goldstein, 2011). While Savage did not interpret his sequence of small worlds in this light, his arguments relating to the consistency needed between the models and the analyses that might be conducted on them provide the justification for current approaches to Bayesian statistics and decision analysis.

Table 2 The use of nested models within the analysis phase

The ‘real world’—whatever that might be, but it is what the decision maker is trying to understand and model

Sense-making: The best current scientific knowledge and understanding of the underlying science together with some broad hypotheses and research questions under investigation (may be entirely qualitative)

Analysis : Nested Models ↓	M_1 , the most complete mathematical model of the system that the scientists can build, perhaps implicit and completely intractable
	M_2 , an approximation to M_1 to make calculations conceptually, if not practically possible
	M_3 , an emulator of M_2 making the calculations yet more tractable

Induction: Interpretation of the results of analysis and understanding the import of the calculations using M_3

While the discussion has focused on Savage’s conception of small worlds, the same thinking applies to other axiomatic approaches to the Bayesian paradigm. All assume that the models are related to reality—the *same* reality—and that that reality provides the data from which we learn through analyses within the models. Moreover, all make a further common assumption: namely that there is a common *reference* or *auxiliary experiment* running through the nested small worlds. The reference experiment in axiomatic terms is simply a sub- σ -field on which the decision maker or scientist perceives an uniform distribution. To give this a practical interpretation, to use Bayesian analysis it is necessary to elicit subjective probabilities and utilities. This is done conceptually by showing the decision maker some randomising device such as a probability wheel. The decision maker is assumed to judge the wheel to be fair and unbiased and thus to judge events of equal size on the wheel to be equally likely. By comparing (i) events on the wheel with events in a small world and (ii) simple gambles constructed on the wheel with possible outcomes in the small world, it is possible to elicit and model the decision-maker’s judgements as probabilities and utilities (French *et al*, 2009). In Savage’s development this is done in his P7 postulate.

In the next section we shall see that an obvious extension of the Bayesian paradigm to fit with recent approaches to scenario-focused thinking means that we must revisit and modify these assumptions.

Scenarios and quantitative risk, and decision analyses

Several authors have begun discussions on how more qualitative forms of analytic discussion may be combined with more quantitative forms and, in particular, the idea of using multiple scenarios to conduct several parallel quantitative analyses. The combination of scenario planning and multi-criteria decision analysis has been a frequent focus (Wright and Goodwin, 1999; Montibeller *et al*, 2006; Ram *et al*, 2011; Schroeder and Lambert, 2011; Stewart *et al*, 2013). Williamson and Goldstein (2012) show how statistical emulation techniques can make the analysis of large complex decision trees tractable and also indicate how their methods can be integrated with scenario planning. Burt (2011) offers a perspective and illustrative case study on the integration of scenario planning and systems modelling. French *et al* (2010) built decision trees in a range of scenarios to explore issues in the sustainability of nuclear power in the United Kingdom.

French (2013) argues that such scenario-focused thinking can be viewed as a stage in moving from the Complex Space to the Knowable Space. The idea is that in making sense of some issues there can be either uncertainties that are so deep or such gross differences in values that a simple Bayesian analysis cannot be used to articulate discussion in any useful way. Experts may disagree on some uncertainties or stakeholders disagree on some societal values so much that sensitivity analysis on any expected utility model will show that some quite disparate alternatives

might all be optimal. The analysis would exhibit the key disagreements, but do little to inform debate and support any move to consensus. Scenario-focused thinking accepts this and begins by focusing on several scenarios. In each, deep uncertainties and key values are fixed to capture an ‘interesting perspective’ on the issues. The remaining uncertainties and values involved are sufficiently understood that informative decision analysis becomes possible *within* each scenario. Participants to the decision will see that, *subject to assuming particular resolution of the deep uncertainties and accepting particular societal values built into a scenario*, there is a reasonable clarity on the way forward. Sometimes, one or more strategies may be dominant in all or most of the expected utility analyses across the scenarios; or there may be a set of strategies that perform poorly in all scenarios. Generally, however, little attempt is made to bring the analyses together *across* scenarios; that is, left to qualitative debate between stakeholders, experts and the ultimate decision makers.

What constitutes an interesting perspective is moot. However, some examples may be given. For instance, in considering the economic viability of an energy portfolio with high levels of nuclear and renewable generation, a deep uncertainty relates to whether some form of energy storage can be developed that allows the slowly variable output of nuclear plants and the vagaries of most renewables to be matched smoothly to relatively fast-changing energy demand. Such storage might come, for instance, from some form of geological heat sink, some novel form of chemical battery capable of taking huge charge or the development of a substantial hydrogen economy. But the development of any of these and the dates by which they might come on stream if developed are deeply uncertain with much disagreement between the relevant experts. One can examine, however, ‘interesting’ scenarios in which each comes to fruition and do so at different dates. Equally the viability of any energy portfolio is also determined by the economic and political climate and such things as whether a low-carbon economy or rapid growth in economic output is pursued by the government. Again interesting scenarios may be established in each of which one of such possibilities is assumed.

There are many parallels between scenarios as they are used in scenario-focused decision analysis and small worlds. Both embody simplified perspectives on possible futures. Both set the bounds of subsequent quantitative analysis delineating what will be modelled and what will be left to intuition and judgement outside the analysis. Reading Savage’s reflections on how a small world may be developed to capture the decision-maker’s understanding of the issues that matter shows many parallels with discussions of the developments of scenarios (Schoemaker, 1993; van der Heijden, 1996; Mahmoud *et al.*, 2009). We have already noted the similarity of some of Savage’s thinking with that of requisite modelling (Phillips, 1984), and scenario need to be developed in a requisite fashion. However, there are differences. As we have seen, Savage developed small worlds as a description of reality. Although there are uncertainties within any of Savage’s small worlds,

there is an assumption that their span contains a perspective on what will ultimately come to pass. There is no such assumption in the development of a set of scenarios: no claim that they span reality in any sense. They are just an interesting set of scenarios, each of which captures some concept of the future that the decision makers wish to discuss. Such a distinction has implications because implicit in Savage’s conception is the idea that as data accumulate, the judgements within prior distributions of belief will be dominated and posterior distributions will become more and more tightly located around the ‘truth’. The Bayesian view of scientific consensus (Box and Taio, 1973; French, 2013) is predicated on the small worlds used in analysis containing reality.

Moreover, there is a significant technical difference. Because Savage essentially considered only one or a nested series of small worlds, his axiomatisation could bury the reference experiment within the axiomatisation of beliefs and preferences within the small world: his P7 implicitly postulates the existence of the reference experiment. Once one begins to consider analyses within non-nested small worlds, that is, scenarios, his approach would lead to several reference experiments, one in each. Moreover, there is nothing in his axiomatisation that would make the quantitative results obtained from analyses *within* each scenario comparable and consistent *across* scenarios. Maybe this is a case of mathematical pedantry; but unless this issue is addressed, many comparisons of the quantitative analyses *across* scenarios would be quantitatively meaningless (Krantz *et al.*, 1971; Roberts, 1979; French, 1986).

Obviously one route out of this conundrum is to create an eighth axiom P8, which makes all the reference experiments essentially the same. A better route is to separate the axiomatisation of the reference experiment from that of beliefs and preferences within each small world (cf French, 1982; Xie and French, 1997; French and Rios Insua, 2000), thus creating a common reference scale against which to elicit the decision-makers’ judgements. This makes the numerical calculations within each scenario comparable across them, without any implication that the scenarios themselves are equally likely or equally important. Indeed, doing so has no implications for any quantitative weighting, equal or unequal of the scenarios.¹ The axiomatic details and further discussion may be found in French (2014).

Separating the axiomatisation of the reference experiment from that of beliefs and preferences in each scenario is particularly useful because it clarifies how the reference experiment forms the basis of elicitation and can help clarify the framing of the judgements that are asked of the decision makers. But doing so makes clear that we may be asking much more difficult judgements from the decision makers than Savage envisaged. We noted that his original approach

¹Note that scenario-focused Bayesian analysis is quite distinct from Bayesian model choice methodologies, which require that we are working in the Knowable Space, identifying a best fit to reality.

assumed a nested sequence of models reaching up to a single reality, one that the decision makers accept. In scenario-focused thinking, we may explore scenarios which all participants believe are effectively impossible, but which are interesting because of the perspective that they offer. For instance, in an environmental debate we might consider an interesting and potentially informative scenario in which all nations agree on a drastic carbon reduction regime and in which all businesses, industries and individuals genuinely seek to achieve this. While this is conceptually possible, I doubt that any party to the debate would consider it to have any chance of becoming reality. Thus in elicitation, we must ask the decision makers to consider the judgements that they would hold in this imaginary world. It may be much harder for the decision makers to make consistent judgements in such an imaginary world, and it will be harder for analysts to constructively challenge these judgements without recourse to reality in testing their consistency. The current literature on scenario-focused thinking does contain suggestions indicating that decision makers find the approach harder and less easy to interpret than the more conventional Bayesian approach: for example, the Italian case in Montibeller *et al* (2006). Moreover, there is little clear agreement yet on how one might display and explore different scenarios with decision makers. That it is hard to deal with deep uncertainties is not surprising, but it should be recognised.

Conclusion

Picking up the various threads of this argument:

- The foundations of Bayesian analysis assume that all aleatory and epistemological uncertainty can be modelled as probabilities.
- In practice, this approach is softened by the use of discussion to resolve ambiguity and sensitivity analysis to address moderate disagreements over the values of particular probabilities and utilities.
- Such approaches have been developed and well-studied for the Known and Knowable Spaces, but do not address the deep uncertainties and deep disagreements that occur in the Complex Space.
- Such deep uncertainties may be explored through the use of a set of scenarios each of which makes assumptions to fix the deep uncertainties at ‘interesting’ values.
- However, the justification of this form of scenario-focused analysis requires that we revisit the axiomatisations of the Bayesian model to allow for several parallel rather than nested small worlds.
- Axiomatising the Bayesian model in parallel small worlds weakens the connection between the model(s) and the real world.

As we noted at the end of the last section, this weakening of the connection between the models and reality means that it may be more difficult cognitively to build understanding and interpret scenario-focused analyses. If we are to use scenario-focused analyses effectively, we need to understand better the justification of the Bayesian model in the context of parallel small worlds and how this may help explore deep uncertainties.

Barankin (1956) wrote ‘... all reality is one grand stochastic process, and any system is a marginal process of this universal process’. In doing so, he caught the mood in mathematical modelling that existed at the time and had influenced Savage in his development Bayesian decision theory. One could conceive of an all-embracing model: a grand world. The recent moves towards scenario-focused thinking may be seen as a step back from that, one that suggests that, in dealing with complex issues, it may be wise to consider several disjoint stochastic processes—several small worlds—each of which captures a different perspective. Fixing deep uncertainties or strong disagreements about societal values in interesting scenarios might help us inform debate and make sense of very complex issues. However, to date developments of scenario-focused analyses have been largely pragmatic. Our discussion has suggested that formal justifications of Bayesian analyses need to be modified to fit with the use of parallel small worlds. Careful study of the required modifications may provide a better understanding of the judgements required from the decision makers, thus elucidating the elicitation process and helping interpret the output of the analyses. That will require much further work.

Acknowledgements—Doug White did much to shape the author’s thinking on decision analysis. In particular, reading and discussing with him his books on *Decision Methodology* and *Operational Research* awoke the author’s interest in the formulation of a ‘mess of incomprehension’ into a model that one can analyse and learn from (White, 1975, 1985). His inspiration and example have remained with the author throughout his career. This paper, inadequate though it be, is dedicated to his memory. Doug was not the only person with whom the author has debated such ideas over the years. The author is grateful to many others and especially to Nikolaos Argyris, Roger Cooke, Roger Hartley, John Maule, Nadia Papamichail, David Rios Insua, Jesus Rios, Jim Smith, David Snowden, Theo Stewart and Lyn Thomas.

References

- Ackoff RL (1962). *Scientific Method: Optimising Applied Research Decisions*. John Wiley and Sons: Chichester.
- Barankin EW (1956). Toward an objectivistic theory of probability. In: Neyman J (ed). *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*. University of California: Berkeley, 5: pp 21–52.
- Bell DE, Raiffa H and Tversky A (1988). *Decision Making*. Cambridge University Press: Cambridge.
- Bordley RF and Hazen GB (1992). Non-linear utility models arising from unmodelled small world intercorrelations. *Management Science* 38(7): 1010–1017.
- Box GEP and Taio GC (1973). *Bayesian Inference in Statistical Analysis*. Addison-Wesley: Reading, MA.
- Burt G (2011). Towards the integration of system modelling with scenario planning to support strategy: The case of the UK energy

- industry. *Journal of the Operational Research Society* **62**(5): 830–839.
- Chaiken S, Liberman A and Eagly AH (1989). Heuristic and systematic information processing within and beyond the persuasion context. In: Uleman JS and Bargh JA (eds). *Unintended Thought*. Guilford: New York, pp 212–252.
- Checkland P (2001). Soft systems methodology. In: Rosenhead J and Mingers J (eds). *Rational Analysis for a Problematic World Revisited*. John Wiley and Sons: Chichester, pp 61–89.
- Churchman CW (1971). *The Design of Inquiring Systems: Basic Concepts of Systems and Organization*. Basic books: New York.
- Cox LA (2012). Confronting deep uncertainties in risk analysis. *Risk Analysis* **32**(10): 1607–1629.
- Edwards W, Miles RF and Von Winterfeldt D (eds) (2007). *Advances in Decision Analysis: From Foundations to Applications*. Cambridge University Press: Cambridge.
- French S (1982). On the axiomatisation of subjective probabilities. *Theory and Decision* **14**(1): 19–33.
- French S (1986). *Decision Theory: An Introduction to the Mathematics of Rationality*. Ellis Horwood: Chichester.
- French S (1995). Uncertainty and imprecision: Modelling and analysis. *Journal of the Operational Research Society* **46**(1): 70–79.
- French S (2003). Modelling, making inferences and making decisions: The roles of sensitivity analysis. *TOP* **11**(2): 229–252.
- French S (2011). Aggregating expert judgement. *Revista de la Real Academia de Ciencias Exactas, Fisicas y Naturales* **105**(1): 181–206.
- French S (2013). Cynefin, statistics and decision analysis. *Journal of the Operational Research Society* **64**(4): 547–561.
- French S (2014). Axiomatising the Bayesian paradigm in parallel small worlds. *Bayesian Analysis*. (in submission).
- French S and Niculae C (2005). Believe in the model: Mishandle the emergency. *Journal of Homeland Security and Emergency Management* **2**(1): 1–16.
- French S and Rios Insua D (2000). *Statistical Decision Theory*. Arnold: London.
- French S and Smith JQ (eds) (1997). *The Practice of Bayesian Analysis*. Arnold: London.
- French S and Xie Z (1994). A perspective on recent developments in utility theory. In: Rios S (ed). *Decision Theory and Decision Analysis: Trends and Challenges*. Kluwer Academic Publishers: Dordrecht: pp 15–31.
- French S, Harrison MT and Ranyard DC (1997). Event conditional attribute modelling in decision making when there is a threat of a nuclear accident. In: French S and Smith JQ (eds). *The Practice of Bayesian Analysis*. Arnold: London.
- French S, Maule AJ and Papamichail KN (2009). *Decision Behaviour, Analysis and Support*. Cambridge University Press: Cambridge.
- French S, Rios J and Stewart TJ (2010). *Decision Analytic Perspectives on Nuclear Sustainability*. Manchester Business School: Manchester.
- Goldstein M (2011). External Bayesian analysis for computer simulators (with discussion). In: Bernardo JM *et al.* (eds). *Bayesian Statistics 9*. Oxford University Press: Oxford.
- Holtzman S (1989). *Intelligent Decision Systems*. Addison-Wesley: Reading, MA.
- Janis IL and Mann L (1977). *Decision Making: A Psychological Analysis of Conflict, Choice and Commitment*. Free Press: New York.
- Kahneman D (2011). *Thinking, Fast and Slow*. Penguin, Allen Lane: London.
- Kahneman D and Tversky A (1974). Judgement under uncertainty: Heuristics and biases. *Science* **185**(4157): 1124–1131.
- Knight FH (1921). *Risk, Uncertainty and Profit*. Hart, Schaffner & Marx; Houghton Mifflin Company: Boston, MA.
- Krantz DH, Luce RD, Suppes P and Tversky A (1971). *Foundations of Measurement Theory. Volume I: Additive and Polynomial Representations*. Academic Press: New York.
- Kuhn TS (1970). *The Structure of Scientific Revolutions*. The University of Chicago Press: Chicago.
- Kurtz CF and Snowden D (2003). The new dynamics of strategy: Sensemaking in a complex and complicated world. *IBM Systems Journal* **43**(3): 462–483.
- Mahmoud M *et al* (2009). A formal framework for scenario development in support of environmental decision-making. *Environmental Modelling & Software* **24**(7): 798–808.
- Montibeller G, Gummer H and Tumidei D (2006). Combining scenario planning and multi-criteria decision analysis in practice. *Journal of Multi-Criteria Decision Analysis* **14**(1–3): 5–20.
- O'Hagan A (2006). Bayesian analysis of computer code outputs: A tutorial. *Reliability Engineering & System Safety* **91**(10): 1290–1300.
- Oaksford M and Chater N (2007). *Bayesian Rationality: The Probabilistic Approach to Human Reasoning*. Oxford University Press: Oxford.
- Phillips LD (1984). A theory of requisite decision models. *Acta Psychologica* **56**(1–3): 29–48.
- Pidd M (1996). *Tools for Thinking: Modelling in Management Science*. John Wiley and Sons: Chichester.
- Ram C, Montibeller G and Morton A (2011). Extending the use of scenario planning and MCDA for the evaluation of strategy. *Journal of the Operational Research Society* **62**(8): 817–829.
- Rios Insua D and French S (eds) (2010). *E Democracy: A Group Decision and Negotiation Perspective*. Group Decision and Negotiation. Springer: Dordrecht.
- Roberts FS (1979). *Measurement Theory*. Academic Press: New York.
- Rosenhead J and Mingers J (eds) (2001). *Rational Analysis for a Problematic World Revisited*. John Wiley and Sons: Chichester.
- Rougier JC, Guillas S, Maute A and Richmond AD (2009). Expert knowledge and multivariate emulation: The thermospher-ionosphere electrodynamic general circulation model (TIE-GCM). *Technometrics* **51**(4): 414–424.
- Savage LJ (1972). *The Foundations of Statistics*. Dover: New York.
- Schoemaker PJ (1993). Multiple scenario development: Its conceptual and behavioral foundation. *Strategic Management Journal* **14**(3): 193–213.
- Schroeder MJ and Lambert JH (2011). Scenario-based multiple criteria analysis for infrastructure policy impacts and planning. *Journal of Risk Research* **14**(2): 191–214.
- Shafer G (1986). Savage revisited. *Statistical Science* **1**(4): 463–485.
- Shaw D, Franco A and Westcombe M (2006). Problem structuring methods I. *Journal of the Operational Research Society* **57**(7): 757–787.
- Shaw D, Franco A and Westcombe M (2007). Problem structuring methods II. *Journal of the Operational Research Society* **58**(5): 545–682.
- Snowden D (2002). Complex acts of knowing—Paradox and descriptive self-awareness. *Journal of Knowledge Management* **6**(2): 100–111.
- Spiegelhalter DJ and Riesch H (2011). Don't know, can't know: Embracing deeper uncertainties when analysing risks. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **369**(1956): 4730–4750.
- Stewart TJ, French S and Rios J (2013). Integration of multicriteria decision analysis and scenario planning. *Omega* **41**(4): 679–688.
- Tukey JW (1977). *Exploratory Data Analysis*. Reading, Mass: Addison-Wesley.
- van der Heijden K (1996). *Scenarios: The Art of Strategic Conversation*. John Wiley and Sons: Chichester.
- Weick KE (1995). *Sensemaking in Organisations*. Sage: Thousand Oaks, CA.
- White DJ (1975). *Decision Methodology*. John Wiley and Sons: Chichester.
- White DJ (1985). *Operational Research*. John Wiley and Sons: Chichester.
- Williamson D and Goldstein M (2012). Bayesian policy support for adaptive strategies using computer models for complex physical systems. *Journal of the Operational Research Society* **63**(8): 1021–1033.

Wittgenstein L (1921). *Tractatus logico-philosophicus*. Routledge & Paul: London.

Wright G and Goodwin P (1999). Future-focused thinking: Combining scenario planning with decision analysis. *Journal of Multi-Criteria Decision Analysis* **8**(6): 311–321.

Xie Z and French S (1997). Towards a constructive approach to act-conditional subjective expected utility models. *TOP* **5**(2): 167–186.

*Received 1 August 2013;
accepted 6 March 2015 after one revision*



This work is licensed under a Creative Commons Attribution 3.0 Unported License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/3.0/>