Journal of Environmental Management (1997) **49**, 53–72

Bayesian Methods for Analysing Climate Change and Water Resource Uncertainties

Benjamin F. Hobbs*

Department of Systems, Control and Industrial Engineering, 610 Crawford Hall, Case Western Reserve University, Cleveland, OH 44106, U.S.A.

The purpose of this paper is to outline the advantages of the Bayesian approach for analysing uncertainties involving climate change, emphasizing the study of the risks such changes pose to water resources systems. Bayesian analysis has the advantage of basing inference and decisions on a coherent and normatively appealing theoretical framework. Furthermore, it can incorporate diverse sources of information, including subjective opinions, historical observations and model outputs. The paper summarizes the basic assumptions and procedures of Bayesian analysis. Summaries of applications to detection of climate change, estimation of climate model parameters, and wetlands management under climatic uncertainty illustrate the potential of the Bayesian methodology. Criticisms of the approach are summarized. It is concluded that in comparison with alternative paradigms for analysing uncertainty, such as fuzzy sets and Dempster–Shafer reasoning, Bayesian analysis is practical, theoretically sound, and relatively easy to understand.

Keywords: climate change, global warming, water resources, Bayesian analysis, decision-making, risk analysis, wetlands, Great Lakes.

1. Introduction

The object of climate change research is to dispel uncertainty concerning the likelihood, magnitude and impacts of climate change. Bayesian analysis is a practical and theoretically appropriate tool for making inferences about climate change and for making decisions based on those inferences. It can help to address questions such as:

- Given present evidence, has the climate already become significantly warmer?
- When might we know if greenhouse gas-forced climate change is really occurring?
- Given the uncertainties inherent in global circulation models (GCMs), what range of mesoscale impacts can be expected in the future?
- Should we commit resources now for preventing or adapting to anticipated climate change, or should we wait to obtain additional information?

* *Current address*: Dept. of Geography and Environmental Engineering, The Johns Hopkins University, 313 Ames Hall, 3400 No. Charles St., Baltimore, MD 21218.

0301–4797/97/010053+20 $25.00/0/ev960116

The purpose of this paper is to summarize the case for Bayesian analysis of climate change uncertainties. Other papers in this special issue present alternative risk analysis methods. Here I address the following questions:

- What can Bayesian analysis do?
- What does Bayesian analysis assume?
- How does Bayesian analysis do what it does?
- How might Bayesian analysis be useful for analysing climate change?
- What criticisms have been made of Bayesian analysis, and what are the responses to those criticisms?

This paper can only touch upon a few aspects of the rich theory that underlies Bayesian analysis. Only a handful of the numerous applications of the approach to climate and weather analysis, hydrology and water resources planning can be mentioned. More thorough introductions to Bayesian analysis and its applications are available elsewhere (e.g. Raiffa, 1968; Berger, 1985; Gatsonis *et al.*, 1992). For an introductory discussion of the use of Bayesian analysis for water resources decision making under climate change, see Fiering and Rogers (1991).

2. What can Bayesian analysis do?

Three features of Bayesian analysis are appealing to practitioners. First are the useful outputs it provides:

- *Inferences.* These include hypothesis tests; expected values for model parameters; credible intervals for those parameters (analogous to, but not exactly the same as, the confidence intervals of classical statistics); and projections and credible intervals for future system behavior. Bayesian analysis can derive these inferences from diverse sources of information, such as historical observations, expert opinions and model simulations. In contrast, classical statistics and information theory can only use historical frequency data and often assume that the system being observed is stationary. Such "frequentist" approaches underlie signal-to-noise ratios and other statistical tests for climate change (Wigley and Barnett, 1990; Schneider, 1994). However, they have little value for projecting the future course of climate change because of their exclusive reliance on historical observations.
- *Optimal strategies.* Given the uncertainties facing us and the information that is available, Bayesian analysis can be used to identify good decisions. Here, a good decision is defined as one that maximizes the expected value of one or more performance indices.
- *Value of information.* Bayesian analysis can be used to trace how additional information would alter the decision we would make and improve the expected performance of those decisions. This enables us to quantify the value of that information in economic terms or whatever performance indices one wishes. The "expected value of perfect information" tells us the most we should be willing to pay for perfect knowledge of the future. Thus, it is a useful measure of the cost of uncertainty. The "expected value of imperfect information" quantifies the value of the imperfect information we might obtain from actual studies. That value can then be compared with the cost of doing such studies. This cost might include the benefits that are foregone if we put off a decision until the information becomes available.

- *Robustness.* Bayesian analysis provides a systematic set of procedures for understanding the sensitivity of our results to the assumptions we make, and the economic consequences of that sensitivity.

These products are of obvious utility to decision-makers and climate-change scientists. A second feature of Bayesian analysis makes these products even more valuable. This feature is the comprehensive framework that Bayesian analysis represents, and the normatively attractive assumptions that underlie that framework. Alternative paradigms, such as fuzzy set theory, lack this advantage. The virtue of comprehensiveness is its coherent encompassing of both inference and decision-making. Since the importance of information cannot be assessed without understanding its impacts on decisions, this is appealing. By "normatively valid assumptions", I mean that the axioms that underpin Bayesian analysis are ones that most rational people would agree ought to be adhered to (even if, in fact, human frailties prevent us from doing so). The comprehensive framework implies that if we believe that we should accept those assumptions, then we should be willing to accept the conclusions derived from Bayesian analysis—or at least be inclined to consider them carefully.

The third and final appealing feature of Bayesian analysis is that its procedures are, in most circumstances, practical and relatively simple compared to other approaches (e.g. Dempster–Shafer reasoning; see Caselton and Luo, 1992), and based on familiar notions of probability. This is an important advantage to the practitioner who must not only obtain results at a reasonable cost, but also be able to explain and justify them to others.

3. What does Bayesian analysis assume?

What assumptions must we accept in order to obtain these benefits of Bayesian analysis? Berger (1985) notes that researchers have developed several dozen alternatives sets of axioms as a basis for Bayesian analysis. The four major assumptions I outline below are loosely based on those of Raiffa (1968), and are similar to other developments. Before presenting the assumptions, some of the notation used must be defined:

x: a performance measure, such as net benefits or environmental impact
$(px_1,(1-p)x_2)$: a lottery with chance p of yielding outcome x_1 and chance $(1-p)$ of x_2.
$\tilde{x}$: a more general lottery, expressed as a probability distribution over x.
$\succ$: "preferred to". For instance "$x_1 \succ x_2$" means "the decision-maker prefers outcome x_1 to outcome x_2".
$\sim$: "indifferent to". As an example, "$x_1 \sim x_2$" means "the decision-maker feels that outcomes x_1 and x_2 are equally attractive (or repulsive!).

The four assumptions are as follows:

1. *Transitivity*: if $x_1 \succ x_2$ and $x_2 \succ x_3$, then $x_1 \succ x_3$.
2. *Certainty equivalence*: if $x_1 \succ x_c \succ x_2$, then the decision-maker must be able to specify some probability p between 0 and 1 such that $x_c \sim (px_1,(1-p)x_2)$.
3. *Substitution*: if $\tilde{x}_1 \sim \tilde{x}_2$ then for any x_3, $(px_3,(1-p)\tilde{x}_1) \sim (px_3,(1-p)\tilde{x}_2)$.
4. *Reduction*: all that matters in making decisions are ultimate outcomes and their probabilities, and not the process involved in getting there. For instance, the following two lotteries should be equally attractive: (i) a 25% chance of winning \$100; (ii) a 50% chance of winning an initial bet. If the initial bet is won, then a second bet is

made with a 50% chance of winning \$100. Both lotteries ultimately involve a 25% probability of gaining \$100, so a decision-maker should not prefer one to the other.

Most reasonable people will agree that these assumptions have normative validity; that is, we ought to try to be consistent with them. Of course, the actual behavior of people often violates those assumptions (von Winterfeldt and Edwards, 1986). But the important point is that these axioms constitute a reasonable definition of "rationality" and that it is reasonable to argue that we ought to strive to be consistent with them. Acceptance of these assumptions has three important and useful implications.

1. "Subjective degree of belief" or "confidence" can be quantified as a subjective probability by the reference lottery method (von Winterfeldt and Edwards, 1986; Morgan and Henrion, 1990). The reference lottery method quantifies a person's subjective probability p of an event A (e.g. "the global mean temperature in the year 2020 will be between 0·5 and 1·5°C higher than in 1990") by finding the p that makes that person indifferent between the following two lotteries: (i) \$X (say, \$100) is won if event A occurs, and \$0 is won if A does not occur (ii) a "reference lottery" in which \$X is won with probability p, and \$0 is won with probability $(1-p)$, with the outcome of the bet occurring at the time that it becomes known whether event A has occurred (in the temperature example, in the year 2020). Spinner dials or other physical devices are often used in actual assessments to visually portray the meaning of different levels of p. When assessing subjective probabilities, the analyst will generally suggest different values of p and ask the decision-maker which lottery is preferred. By trial and error, the p that yields indifference can be determined. There are also many other approaches to eliciting subjective probabilities (e.g. von Winterfeldt and Edwards, 1986; Morgan and Henrion, 1990).
2. Subjective beliefs assessed in this manner have the properties of probabilities:

$$P(A)+P(A^c)=1 \tag{1}$$

$$P(A)+P(B)-P(A\bigcap B)=P(A\bigcup B) \tag{2}$$

$$P(A\mid B)\ P(B)=P(B\mid A)\ P(A)=P(A\bigcap B) \tag{3}$$

where:

A^c:	the complement of A (i.e. "A does not occur"),
$\bigcap$, $\bigcup$:	set intersection and union symbols, respectively, and
$P(A\mid B)$:	the conditional probability of event A, given that event B has occurred.

3. A utility function $U(x)$ exists such that $E(U(\tilde{x}_1))>E(U(\tilde{x}_2))$ if and only if $\tilde{x}_1\succ\tilde{x}_2$, where $E(U(\tilde{x}))$ is the expected value of $U(\tilde{x})$. That is, there exists a function whose expected value can be used to represent the preferences of the decision-maker. This function can be estimated using standard procedures of decision analysis (e.g. Raiffa, 1968).

Of course, people are unsure of their beliefs and preferences and have difficulty expressing them coherently. Consequently, consistency checks generally will show that actual subjective probabilities and utility functions are inconsistent and error-laden. This does not mean that the Bayesian method has no value. Instead, it implies that trustworthy results are only possible if assessments are carefully done, consistency checks are applied, and the users thoughtfully resolve those inconsistencies and have

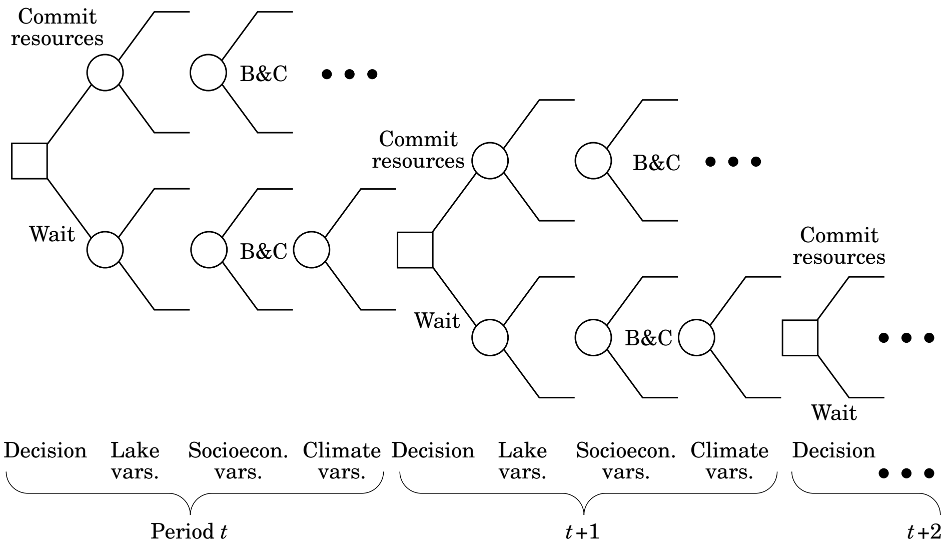

Figure 1. Decision tree for management decisions under climate change.

confidence in the final values. This process can be difficult and time consuming, but users benefit because they have a chance to think through their beliefs and preferences and to make them more coherent.

Equation (3), the third property of subjective probabilities, can be rearranged to yield Bayes' law:

$$P(A \mid B) = P(A)[P(B \mid A)/P(B)] \tag{4}$$

That is, the probability of A given that B has occurred (the "posterior" probability of A), equals the "prior" probability $P(A)$ multiplied by the ratio of (a) the likelihood of observing B given that A will occur ($P(B \mid A)$) to (b) the overall or "marginal" probability of $B(P(B))$. This disarmingly simple equation is the foundation of Bayesian inference.

4. How does Bayesian analysis do what it does?

Bayesian analysis produces the outputs described in Section 2 in four steps:

1. Specify problem structure.
2. Quantify inputs.
3. Perform the analysis.
4. Examine robustness of the results to changes in assumptions.

Each step is discussed below.

4.1. SPECIFY PROBLEM STRUCTURE

This consists of describing the options that are available to the decision-maker (resource commitments, information gathering); when choices from among those options can be made; and when information is obtained that reduces uncertainty. Decision trees and influence diagrams are two useful ways of portraying problem structure (Clemens, 1991); here, trees are used. Figure 1 shows a tree that will be referred to again later in the paper. It represents the general structure of several types of decisions involving management of the Laurentian Great Lakes of North America (Hobbs *et al.*, 1996).

As for any decision tree, the one in Figure 1 has three major elements.

- *Decision nodes*, represented as squares. Here, the options are to make an irreversible commitment of resources, or to put off the decision. Examples of such commitments might be the construction of shore protective works or a control structure at the outlet of Lake Erie, a decision to develop a wetland, or the signing of an agreement to export water out of the basin.
- *Chance nodes*, represented as circles. I show three types of uncertainties. The first includes lake-related variables, such as levels or temperatures, which fluctuate randomly and which also have an uncertain trend due to the possibility of climate change. Socio-economic variables, such as values of shoreline property, make up the second category. The third group includes climate variables, which provide information that can be used to update the decision-makers' beliefs concerning the likelihood and magnitude of climate change. These may overlap with the lake variables. Probabilities, usually containing some subjective element, are associated with each branch on a chance code.
- *outcomes*, which I show as benefits and costs "B&C" that accrue over time. In general, there may be several incommensurable types of benefits and costs, which means that the decision problem is a multicriteria one.

Decision nodes can be omitted in pure inference problems, such as the testing of hypotheses, estimation of model parameters, or making of predictions.

4.2. SPECIFY INPUTS

Once the structure of the problem has been outlined in the form of a tree or influence diagram, then numerical inputs are needed: probabilities and outcomes. First, some more notation:

θ: state of nature, which, at least initially, is unknown. An example would be θ_1 which might represent the state "the global mean temperature will be warm by between 0·25 and 0·5°C per decade over the next century".

z: information, usually in the form of an observation of some random process (e.g. next year's global mean temperature) or a prediction by an expert or model. Obtaining information alters subjective probabilities of θ.

Now the necessary inputs can be summarized as follows:

$P(\theta)$: prior probabilities of the state of nature. These should be assessed based on the best information now available. If information is unavailable or ambiguous, then so-called "non-informative priors", which attempt to avoid favoring some states of nature over others, are appropriate. Priors might also be based on past historical observations, numerical simulations and expert judgment.

$P(z \mid \theta)$: likelihoods, which describe the probability of obtaining evidence z given that the state of nature is actually θ. For instance, if the climate is actually warming by an average of 0·25–0·5°/decade (i.e. the state of nature θ), what is the likelihood that the observed global mean temperature next year will be, say, between X° and Y° (i.e. the observation z)? The values of $P(z \mid \theta)$ are frequently determined by models that, given an assumed set of parameters (the state of nature θ), can generate probability distributions $P(z \mid \theta)$. Expert judgment can also be used to assess likelihoods.

In addition, if we are not merely going to make inferences about the state of nature, but are also going to model the decision-making process, then the following additional information is required:

a: options (resource commitments, information gathering) available at each decision node.

x: measure(s) of performance. In general, x might be a vector, encompassing a range of monetary and non-monetary criteria.

$U(x)$: a utility function translating the measures of performance into a single index of worth. Several $U(x)$'s might be defined, each representing a different interest group's point of view.

4.3. PERFORM THE ANALYSIS

There are two basic types of Bayesian analysis: posterior and preposterior. The purpose of posterior analysis is to update the probabilities of the state of nature after information z is obtained, yielding the posterior probabilities $P(\theta \mid z)$. Bayes' Law (4) is used to accomplish this. If the θ represent different hypotheses, then the posterior probabilities can be used to "test" which hypothesis is most credible. For example, θ_1 might represent the hypothesis that "the climate at location X has not warmed by more than 0·25° in the last 40 years", while θ_2 might instead be "the climate has warmed by more than 0·25°". Their posterior probabilities, given the evidence z (the actual temperature record) and the prior probabilities (say, equal), show which hypothesis is more likely.

If θ is instead a continuous variable, such as the parameter of a flood frequency distribution, then posteriors $P(\theta \mid z)$ can be used to calculate several useful statistics. These include the most likely value of θ (the mode or "maximum likelihood" estimate), its expected value (which might differ greatly from the mode), and credible intervals. A $X\%$ credible interval for θ is defined as a domain D such that the posterior probability of θ falling in that domain is $X\%$:

$$\int_{\theta \in D} P(\theta \mid z)\mathrm{d}\theta = X/100 \tag{5}$$

An example of the calculation of these statistics is given later in this paper.

Posterior analysis focuses on inference. Preposterior analysis, in contrast, emphasizes decision-making. First, considering the distribution of possible outcomes z and θ, preposterior analysis calculates the optimal immediate action a—the so-called "Bayes decision". More generally, the optimal *strategy* is defined, i.e. the option at each decision node that maximizes the expected net benefits from that point on. This is accomplished by the "folding back" method, also called stochastic dynamic programming. Basically, the following calculation is made for each decision node, starting at the last period: choose the action yielding the highest expected benefit. Because this calculation will have already been done for all subsequent decision nodes, the procedure automatically takes into account what decisions might be made later in response to further information.

Preposterior analysis also produces a second important product: the value of further information. The value of *perfect* information, which gauges the cost of uncertainty, is obtained as follows. It equals the difference between the expected net benefits of the Bayes decision, and the expected net benefits if the future could be perfectly forecast—i.e. if there was no uncertainty. More important for decision-making, preposterior analysis can also assess the value of *imperfect* information. Say that an experiment, modeling exercise or other research project would yield information z about the present or future climate. If that information is optimally used to update prior distributions via Bayes' law, then having that information available will usually improve the expected net benefits of the optimal strategy, which in general, will depend on what information is obtained. The expected amount of that improvement is the value of that information. If net benefits are measured using multiple criteria, then the value of information itself will have multiple dimensions (e.g. expected reduction of cost, expected reduction of environmental impact, etc.).

4.4. ROBUSTNESS ANALYSIS

Robustness analysis concerns the question of the dependency of the results on the assumed $P(\theta)$, $P(z \mid \theta)$, and $U(x)$. Such analysis is important because psychologists have shown that assessments of subjective probabilities and utilities can be unreliable and subject to predictable biases (Fischhoff *et al.*, 1979; Kahneman *et al.*, 1982).

Although critics and users of Bayesian analysis both emphasize the importance of prior distributions, likelihoods and utility functions can also significantly alter the conclusions of an analysis. Uncertainties in $P(z \mid \theta)$ are sometimes called modelling and representativeness uncertainties, when the $P(z \mid \theta)$ result from some model (such as a log-normal flood frequency model) that might either be the wrong model or that might be in error because reality is inherently more complex than any model we might realistically formulate (Bernier, 1991).

There are several approaches to robustness analysis (Berger, 1985). The first is traditional sensitivity analysis, in which different assumptions are tested, and variations in the results are noted. In general, if the prior $P(\theta)$ is "sharp" (i.e. high for a few values of θ, and low for the rest), then the conclusions will be relatively insensitive to $P(z \mid \theta)$. That is, evidence will not alter prior beliefs greatly. On the other hand, a sharp $P(z \mid \theta)$ usually implies that the results are relatively insensitive to $P(\theta)$. There evidence matters more than previous beliefs.

In the second approach to robustness analysis, the analyst might attempt to minimize the effect of subjectivity by adopting inherently "robust" prior distributions. These are "flat" or "non-informative" priors, discriminating little or not at all among alternative values of θ. In many cases, use of such priors causes Bayesian analysis to give the same results as traditional statistical methods (maximum likelihood estimators, classic statistical tests) and information theory (Tsao *et al.*, 1993). However, if credible prior information is possessed that does indeed discriminate among the possible θ, then it should be used. Throwing such information away by using a flat prior will, in general, yield inefficient and biased inferences (e.g. parameter estimates) and decisions that fail to maximize expected net benefits (Krzysztofowicz, 1983).

The third approach to robustness analysis calculates the regret that could result from using the "wrong" inputs. Say that there are two alternative sets of assumptions concerning $P(\theta)$, $P(z \mid \theta)$, and $U(x)$: $\{P_1(\theta), P_1(z \mid \theta), U_1(x)\}$ and $\{P_2(\theta), P_2(z \mid \theta), U_2(x)\}$. We start by calculating the optimal action a for each set of inputs: a_1 and a_2. "Regret" is the foregone expected net benefits resulting from assuming that the first set of inputs is correct (and thus implementing its action a_1) when the other set is the right one instead. This results from a failure to choose the true optimal action. If a_1 and a_2 are very similar, or at least yield similar expected net benefits, then this regret will be relatively small. But if the possible regret is large, then alternative assumptions make a practical difference. Additional attention should then be paid to the question of which set of inputs is truly more representative of the decision-makers' knowledge and values.

A fourth approach to robustness analysis is maximum likelihood analysis. It uses observations of z to determine which $P(\theta)$ and $P(z \mid \theta)$ seem most consistent with the evidence. For instance, say that the observed annual floods in the last three years were 12 000, 18 000, and 9000 m^3/sec, respectively. Then, a prior distribution that assigns zero probability to a mean annual flood of over 10 000 m^3/s would seem to be less credible than one that allows such a possibility. A metric of credibility in this case is $P(z) = \Sigma_\theta P(z \mid \theta)P(\theta)$, the overall (marginal) likelihood of the evidence (Berger, 1985). In the above case, this metric is likely to be much smaller for the prior that excludes the possibility of a mean value over 10 000 m^3/sec.

An example of a simple but insightful robustness analysis of the global warming problem is that presented by Lave and Dowlatabadi (1993). They examined three types of uncertainties: the relationship between emissions of greenhouse gases and global climate change; the resulting effects on humans and the environment; and the costs of reducing greenhouse gas emissions. They evaluated three hypothetical greenhouse gas reduction strategies under these uncertainties. The robustness of the results were analysed by considering the sensitivity of the decisions to the inputs, and the regret that results if incorrect inputs are used.

5. How might Bayesian analysis be useful for analysing climate change?

Bayesian analysis can use evidence to make and update climate change predictions and estimates of model parameters. Bayesian methods can also be used to define optimal natural resource management strategies, given the prospect of possible climate change. Three illustrations are summarized here of the use of Bayesian analysis to analyse climate change. The first two focus on inference: what models and parameters are most consistent with evidence? One of the examples examines annual flows in the Senegal River of West Africa, and the other studies sea level rise. The third example concerns optimal management of lacustrine wetlands under climatic uncertainty.

5.1. Flow distributions for the Senegal River (Bernier, 1994a)

There is considerable concern over the possible effect of climate change upon the frequency of floods and droughts (Waggoner, 1990). Duckstein *et al.* (1987) outline how the Bayesian methodology could be used to forecast resulting changes in annual flood frequencies. Bernier (1994b) reviews Bayesian tests for detecting the date of climate change in hydrological time series, and illustrates his methods with applications to the Harricana, St. Lawrence, and Senegal rivers and precipitation in the North American Great Lakes region. In Bernier's (1994a) study of droughts in the Senegal River basin, uncertainties concerning the probability distribution of annual flows Q were analysed with a Bayesian approach. He addressed the following questions:

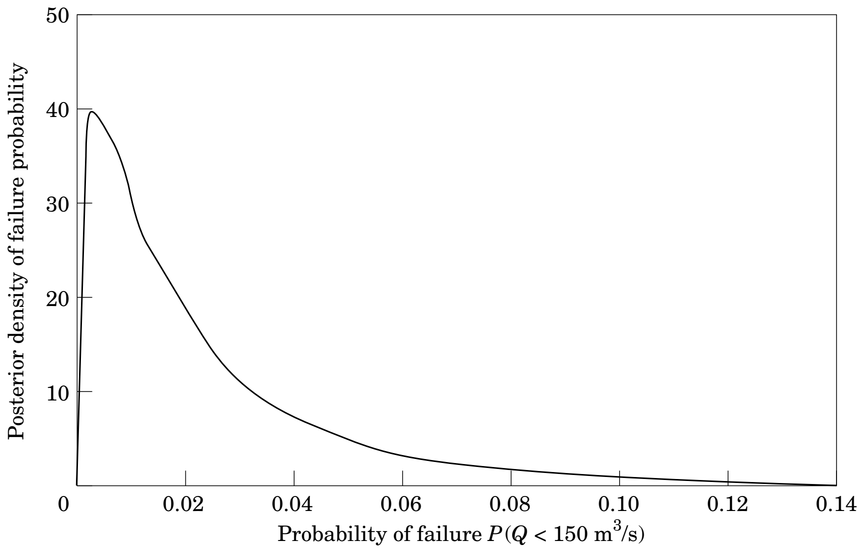

Figure 2. Posterior probability distribution of failure probability, Senegal River [based on data in Berger (1994)].

1. When did the mean annual discharge $E(Q)$ change?
2. What is the future probability of system failure, defined as $P(Q<150\ m^3/s)$, given parameter and model uncertainty? What is its credible interval?
3. What are the parameters of a Markov chain model of the annual flow, and what are their implications for the probability of system failure ($Q<150\ m^3/s$)?

In this analysis, the year in which the change took place is an unknown state of nature, as are the parameters of the Markov chain model. Bernier's results for the first two questions are summarized below.

The Senegal river drains a portion of the Sahel region, and statistical tests showed that the annual flow during the period 1903 to about 1965 was significantly higher on average than flows for subsequent years. A Bayesian analysis to answer the first question was performed as follows. The year in which the mean changed is defined as the state of nature θ, $P(\theta)$ was assumed to be the same for all years, z consisted of the observed flows for 1903–1986, and $P(z \mid \theta)$ was based upon a log-normal distribution whose mean and standard deviation for each year depends on whether the year is before or after the year in which the mean changed. The posterior probabilities of the date of change T were 0·03 for $T<1967$, 0·27 for 1967, 0·07 for 1968, 0·22 for 1969, 0·16 for 1980, 0·17 for 1971 and 0·08 for $T>1971$.

The second question Bernier (1994a) addressed is important because annual flows below 150 m^3/s are assumed to represent a failure of water supply. One answer to this question was derived under the assumption that the mean flow would remain at its estimated post-1967 value. The state of nature θ was defined as the parameters of a log normal distribution from which $P(Q<150\ m^3/s)$ would be calculated. The observations z were the 1967–1986 annual flows, and their likelihood was described by the log-normal distribution. Bayes' Law was then applied, resulting in a posterior distribution of θ, from which a posterior distribution of the estimate of $P(Q<150\ m^3/s)$ could be derived (Fig. 2). The mode of $P(Q<150\ m^3/s)$ is 0·0035, but the expected value is much higher, about 0·025. The short historical record results in a wide credible interval; there is a 90% chance that the true value is between 0·001 and 0·059. Thus, Bayesian analysis has permitted quantification of the uncertainty surrounding the probability of future droughts.

5.2. BAYESIAN MONTE CARLO ANALYSIS FOR UPDATING MODELS OF SEA LEVEL RISE (PATWARDHAN AND SMALL, 1992)

In this study, a formal method was developed for using evidence to update probability distributions for a complex model's parameters, which were treated as the unknown state of nature. The model and the probability distributions of its parameters can then be used to make predictions that reflect the remaining uncertainty. The approach was then applied to the problem of projecting sea level rise. It can also be used to improve models of other impacts of climate change, including hydrologic effects, and to systematically update those models as additional information is obtained. The steps of the method are as follows:

1. Define a prior probability distribution over the parameters θ of the model (in their case, the parameters of a model of sea level rise). A prior distribution could also be provided over alternative models, in which different models are assigned different levels of credibility [in the spirit of Ellis (1988) who considers the implications for SO_2 control of several alternative models to project acid deposition].
2. Sample θ I times by a Monte Carlo procedure. Each realization θ_i, $i=1,\ldots,I$ is then used to obtain a projection $M_i(\theta_i)$ of one or more variables of interest (here, future sea levels). Each projection $M_i(\theta_i)$ is assumed to be equally likely (i.e. having a prior $P(M_i(\theta_i))=1/I$).
3. The variables to be observed, or evidence, are designated as z (in this case, observed changes in sea level at different locations). A likelihood function $P(z \mid M_i(\theta_i))$ is derived, based on likely sources of observational error.
4. z is observed, and then Bayes' law is used to update the posterior distribution of model projections $P(M_i(\theta_i) \mid z)$. These can then be translated back into posterior distributions for the model parameters $P(\theta \mid z)$.
5. Step 2 can then be returned to, this time sampling from $P(\theta \mid z)$ rather than $P(\theta)$. Then, as new evidence is obtained, the distribution can be further updated by repeating Step 4.

To illustrate the application of this procedure, Patwardhan and Small's (1992) application to sea level rise is summarized below.

The model used by them was a simple one-dimensional upwelling-diffusion model of the ocean that simulates the effect of thermal expansion upon sea levels. Within θ are four parameters that were modeled as being uncertain:

- temperature sensitivity (equilibrium surface temperature increase for doubled CO_2, °C)
- vertical diffusion coefficient for heat transfer (cm^2/sec)
- ocean upwelling velocity (m/year)
- ocean mixed layer depth (m)

In addition, the uncertain state of nature θ included growth rates for atmospheric concentrations of four greenhouse gases, which are also treated as being uncertain.

The probability distributions for each of these parameters and variables are assumed to be uniform and independent.

The output $M_i(\theta_i)$ of the model is a projection of mean sea level for a given range of years for an assumed state of nature θ_i. By sampling the parameters from the assumed distribution of θ and then calculating one projection for each sample, a range of predictions $M_i(\theta_i)$ for sea levels can be obtained. For instance, their model projects a mean sea level change of about 0·75 m by the year 2100, with a 90% credible interval of 0·2 m to 1·3 m. As actual sea levels and gas concentrations are observed over time, this evidence can be used to update the parameters of the model, resulting in (presumably, but not necessarily) less uncertainty. However, z consists of local measurements of sea level, which can differ from global means because of factors such as isostatic adjustment and changes in long-term wind patterns. These factors can be considered in the derivation of likelihood functions $P(z \mid M_i(\theta_i))$.

To illustrate how this updating process can work, Patwardhan and Small (1992) first used observed sea levels from 1900 to 1980 as z. As a result of applying their Monte Carlo Bayesian procedure, they refined the distributions of the four model parameters. The largest change was in the temperature sensitivity, whose 90% credible interval shrunk from [0·69,4·3] to [0·96,4·3]. As another illustration of the method, they demonstrated how the method would have been applied sequentially over the period 1900–1980 as additional evidence accumulated. When year 2000 sea levels were forecast using just evidence obtained through the year 1900, the 90% credible interval for the rise relative to 1900 levels was [0·13 m, 0·75 m]. As evidence accumulated through the century, this interval shrunk until by 1980 it was [0·10 m, 0·25 m].

5.3. STOPPING-CONTROL PARADIGM FOR DECISIONS UNDER CLIMATE CHANGE (KRZYSZTOFOWICZ, 1994) APPLIED TO GREAT LAKES WETLAND MANAGEMENT (HOBBS *ET AL.*, 1994)

This example integrates the problems of inference and decision-making in a Bayesian framework. The problem Krzysztofowicz (1994) addresses can be simply stated, as follows. There is uncertainty concerning the future climate; the state of nature θ represents possible future climates. A decision must be made either to commit resources (e.g. invest in a flood control project) or wait for further information z that can be used to revise prior probabilities of the various climates. (In addition, operating decisions that are not irreversible commitments can also be made, but I disregard that complication here.) This commitment of resources is called a "stopping". Waiting can be worthwhile if the net benefits of the resource commitment depend on climate. Waiting can result in, ultimately, a better resource commitment (e.g. design of a hydraulic structure), but there are likely to be opportunity costs in the meanwhile—no benefits are received from a structure that is not built.

Krzysztofowicz (1994) assumes that the evidence Z_t at time t of climate change has a likelihood function of the form $P(z_i \mid \theta, z_{t-1})$. That is, given that the actual climate is θ, z_t is a first order (autoregressive) Markov process. Given this likelihood function, and the prior probabilities of θ, Bayes' law (4) can be applied to the evidence, resulting in an update of the decision-makers' beliefs concerning climate change. Krzysztofowicz (1994) presents two important results for this problem under the assumption that there exists a utility function of the type I discussed in Section 3. The first is that an optimal "stopping-control" strategy can, in theory, be obtained by stochastic dynamic programming. This strategy describes the best decision now (whether or not commit resources, and how to control the system), and the best decisions in the future under all possible realizations of the evidence z_t. The second is that the value of improved forecasts (information) can be calculated.

This general framework is relevant to many problems involving the management of the Great Lakes of North America. Below, I summarize an application to wetlands management as an illustration (Hobbs *et al.*, 1994). Other applications have been made to decisions concerning construction of breakwaters at Presque Isle, PA and regulation of the outflows of Lake Erie (Hobbs *et al.*, 1996).

An increase in greenhouse gas concentrations equivalent to a doubling of CO_2 concentration could drop levels of Lake Erie by 1 to 2 m (Croley, 1990). This decrease would have an enormous effect on coastal wetlands. Many existing wetlands would be deprived of the periodic inundation that they require, while new wetlands would be created from former lake bottom.

To illustrate the application of Figure 1's decision framework, I consider an existing wetland that would disappear if lake levels drop that far. I assume that this wetland is situated so that the land would be valuable for commercial development.

The elements of this problem are:

- a: decisions whether to allow commercial development a_1 or to preserve the wetland ("wait") a_2. I disregard the possibility of maintaining the wetland by diking and pumping by assuming that its expense could not be justified if lake levels drop. This is because diking would cut off access to the lake, thereby destroying the wetland's value for aquatic habitat.
- θ: states of nature indicating whether lake levels are stationary around the historical mean (θ_1) or whether they vary around a mean trend that decreases linearly over 40 years to a level 1·5 m below the historical mean (θ_2). My prior probabilities are $P(\theta_1) = P(\theta_2) = 0{\cdot}5$.
- z_t: evidence obtained in year t concerning whether lake levels are indeed permanently dropping because of climate change. For simplicity, I assume that yearly mean lake levels are the only such evidence, although in reality evidence can be obtained from many sources. Observations of these levels are used in Bayes' law to update the state of nature probabilities. I assume that lake levels are a first-order autoregressive process.
- x: outcomes, including benefits of development, if allowed, and permanent loss of wetlands. I assume, for illustration purposes, that the wetland has a high (but unspecified) ecological and social value if lake levels do not drop permanently (θ_1). However, because the wetland would degrade in quality and other wetlands would be created if levels fall permanently (θ_2), I assume that the wetland would have relatively little ecological value if θ_2 is realized. The particular performance measures I use are:
  - x_1: *worth of commercial development*, equal to the expected present value of the land for commercial purposes at the date of development. An expected value is calculated because it is uncertain when, if ever, commercial development would occur.
  - x_2: *loss of wetlands*, equal to the probability that a valuable wetland is eventually developed (i.e. that both commercial development and θ_2 occur). In a sense, this is the long run ecological cost that could result from making a wrong decision.

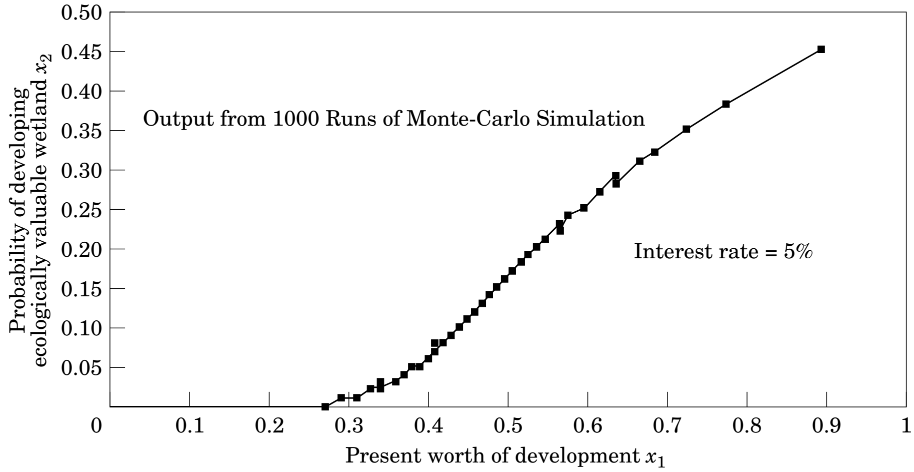

Figure 3. Tradeoff between expected development benefits and probability of mistakenly losing valuable wetlands.

This problem fits into the optimal stopping-control framework because:

1. Commercial development is assumed to be irreversible.
2. Delaying development would allow managers to obtain better information on whether lake levels are going to drop permanently, via Bayes' law.
3. The benefit of not developing the wetland depends on whether lake levels permanently decrease.
4. If development is delayed, economic benefits will be foregone in the meanwhile.

The heart of the problem is thus whether to develop now, and run the risk that a valuable wetland will be destroyed, or to wait, and thereby give up economic benefits in the interim.

One possible approach to solving this problem is to assume a monetary value for the worth of wetlands, and then combine that value with the economic value of commercial development. An alternative approach is multicriteria analysis. The focus of my multicriteria analysis is on generating tradeoffs among competing criteria x_1 (development benefits) and x_2 (risk of ecological loss due to making the wrong decision). Different strategies will result in different levels of performance of each objective, as illustrated in Figure 3. The commercial development axis has been normalized so that "1" represents the maximum possible economic value of development (resulting from development immediately at $t = 0$). The points on the lower left represent conservative strategies that prohibit development unless it is almost certain that lake levels will drop; as a result, there is little chance of developing a valuable wetland but also a relatively low present worth of development. The points on the upper right represent more aggressive strategies that allow development even if climate change is highly uncertain. Consequently, the probability of mistakenly developing a valuable wetland is higher, but so are the development benefits.

The points in Figure 3 are generated as follows. Rather than discretizing probability distributions of lake levels and then solving the resulting (enormous) tree by backwards dynamic programming, I opted to define a range of simple policies and then stimulate their performance using Monte Carlo simulation. The policies are of the form: *If in year t the posterior probability of decline in lake levels $P_t(\theta_2)$ exceeds a threshold T, then develop the wetland. Otherwise, wait.* A low threshold T (e.g. $T = 0{\cdot}7$ or less) yields the upper right points in the figure. In those cases, there are high development benefits x_1 because development is likely to occur relatively soon, but there is also a high probability x_2 of mistakenly developing a valuable wetland. In contrast, a high threshold ($T = 0{\cdot}9$ or more) results in the lower left points. Such a cautious policy lowers the probability of mistaken development, but also diminishes the present worth of development benefits. There may be (indeed, are likely to be) better strategies that lie below and to the right of this curve; they might be discovered either by simulating more complex policies or by backwards dynamic programming.

The exact form of the tradeoff curve depends on the assumed interest rate and the prior probabilities. If the interest rate is zero, then the curve consists of a straight line connecting the points ($x_1 = 0$, $x_2 = 0$) (development never occurs because the threshold is zero) and ($x_1 = 1$, $x_2 = 0{\cdot}5$) (development immediately occurs because of a high threshold, thereby resulting in a high risk of losing a valuable wetland). A change in the prior probability of climate change $P(\theta_2)$ scales the tradeoff curve up or down in rough proportion to that probability.

When managers are presented with a curve such as Figure 3, they can decide what trade-off between development benefits and risk of wetland loss is acceptable. The decision depends on the relative weight given to wetlands and commercial development.

6. Criticisms of Bayesian analysis

The above applications illustrate some potential uses of Bayesian analysis for inference and decision-making. However, critics direct four types of criticisms at Bayesian theory and applications. I discuss each in turn, along with some of the rebuttals offered by advocates of Bayesian analysis.

6.1. SUBJECTIVITY IS UNSCIENTIFIC

The idea that *subjective* judgments of prior probabilities and likelihood functions can—and indeed should—influence the outcome of an inference process is abhorrent to many scientists. The debate between "frequentists" and Bayesians is now an ancient one, dating back more than a century (Poirier, 1988; Rust *et al.*, 1988). Frequentists argue that inference should be based on observed data, and that alone; subjectivity prevents us from viewing that data objectively and can introduce bias.

One reply Bayesians make to this criticism is that subjectivity is inevitable, and is merely swept under the rug by frequentists. For example, many standard statistical tests assume *a priori* that all parameter values are equally likely, even though, for example, physical principles might dictate non-negativity. The Bayesian approach instead makes subjective assumptions explicit, where they can be scrutinized and analysed for robustness.

Subjectivity's inevitability is apparent in the effort to detect whether global warming is occurring. There has been, for example, a statistically significant warming trend during the 20th century. However, whether or not such evidence "confirms" a greenhouse hypothesis depends on the assumptions that are made about:

- other possible hypotheses (e.g. whether the possibility of sun spot-driven climate change is considered) and

- the degree of natural variation, or "noise", that might mask a climate warming "signal" (such as the random shifts in the means of regional temperature and precipitation that occur on a multi-decade time scale; Slivitsky and Mathier, 1994).

These can be viewed as assumptions concerning prior distributions and likelihood functions—assumptions that Bayesians make explicit and can be subject to robustness analysis.

Another reply made by Bayesians is that it is demonstrably suboptimal to ignore prior beliefs when making decisions. That is, decision-makers will, in the long run, realize lower net benefits (utility) if they fail to combine prior beliefs and evidence according to Bayes' law. To just consider past direct observations of a phenomenon when making inferences (a key aspect of the frequentist approach) ignores potentially very valuable information based upon physical reasoning (models) and export opinion. Krzysztofowicz (1983), for instance, shows how an intentional disregarding of such information yields lower expected net benefits for a decision-maker who must use weather forecasts to make decisions about resource allocations (in his case, crop watering).

6.2. PEOPLE DO NOT ACTUALLY THINK LIKE BAYESIANS

There is ample empirical evidence that people act contrary to the assumptions of Bayesian analysis. They fail to update prior beliefs using Bayes' Law, and they have great difficulty specifying utility functions because their preferences are incoherent and inconsistent (Kahneman *et al.*, 1982).

A Bayesians' response to this criticism might be "so what?" The point of Bayesian analysis is to *improve* upon unaided human judgment, not to imitate it. Bayes' Law is not meant to be a psychological theory that can be used to predict behavior; rather it is supposed to be a guide to making more rational, consistent, and defensible decisions. Bayesian analysis is the only integrated approach to inference and decision making that is fully consistent with a set of assumptions that have normative appeal.

6.3. BAYESIAN PROBABILITIES FAIL TO REPRESENT A PERSON'S STATE OF KNOWLEDGE

Critics of Bayesian analysis argue that "uncertainty" encompasses several distinct concepts, and that subjective probabilities fail to do them justice. For instance, linguistic imprecision is one form of uncertainty. What is "warm"? What is a "significant" environmental impact? Another form of uncertainty results from ignorance. As an example, scientists do not know whether higher CO_2 levels will encourage further accumulation of biomass (and thus sequestering of CO_2) in climax forests. Imprecision and ignorance are not the same phenomena as randomness, such as is the situation with random hydrologic variables for which probability distributions have been estimated (e.g. the annual discharge for a river with 100 years of flow data). To boil down these various dimensions of uncertainty into a single number—a subjective probability—strikes many people as distortion and oversimplification (e.g. Tonn and Schaffhauser, 1994).

Fuzzy-set theory has been offered as an alternative in the case where uncertainty is due to the imprecision of language (Zadeh, 1983; Duckstein, 1994a). It represents degrees of, for instance, "warmness" or "significance" by fuzzy set membership functions whose values can range from 0 (definitely *not* warm or *not* significant) to 1 (definitely warm or significant). Although they use the same scale as probabilities, fuzzy set membership functions are subjected to different manipulation procedures. Duckstein (1994b) and Nachtnebel and Duckstein (1994) compare the practical implications of fuzzy set and Bayesian approaches.

When instead ignorance is the problem, Dempster–Shafer (D–S) reasoning has been proposed as a generalization of Bayesian analysis (e.g. Caselton and Luo, 1992). Instead of assessing a precise point probability for each individual event, the D–S method characterizes uncertainty by assigning, in effect, upper and lower bounds upon the probability of an event. These bounds are called "belief functions". It is argued that such bounds are a reasonable characterization of the state of knowledge: if the probability distribution is known, then the bounds are very tight, but if there are few data, then the bounds can be loose. As evidence is obtained, these bounds are updated.

To the criticism that subjective probabilities are an over-simplification, a conservative Bayesian might reply as follows. Under reasonable assumptions (Section 3), it can be proven that subjective probability is all that is needed for rational inference and decision-making. The main requirement is that the decision-maker can respond meaningfully to the reference lottery, which permits reduction of the many dimensions of uncertainty into one. To develop a new, more complex calculus of uncertainty, such as fuzzy sets or D–S reasoning, violates Occam's razor: that a model or theory should be no more complicated than necessary to accomplish its goal (Lindley, 1987). Moreover, if the object is to avoid over-precision resulting from assigning a single probability, it seems counterproductive to replace a single precise number (probability) with two such numbers (upper and lower bounds).

Regarding the issue of linguistic imprecision, some risk analysts argue that imprecision is the result of laziness. They assert that definition of appropriate quantitative scales will, in most cases, get rid of the problem (Morgan and Henrion, 1990).

Another reply to the criticism that Bayesian probabilities do not adequately characterize knowledge is that carefully controlled psychological studies fail to conclude that fuzzy set or D–S representations of uncertainty are any more valid than subjective probability. Indeed, there is some evidence to the contrary (Oden, 1977; Morgan and Henrion, 1990).

A final reply is that non-Bayesian decision procedures are often *ad hoc* in nature. For instance, fuzzy set-based decision procedures aggregate different criteria using formula whose assumptions concerning the structure of people's preferences are unclear and never verified in practice.

Dowlatabadi and Morgan (1993) also criticize the ability of Bayesian probabilities to represent people's knowledge, although for a different reason. They argue that Bayesian value of information studies tend to mischaracterize the value of additional research because analysts do not anticipate the "surprises" that research can reveal. They therefore recommend that stimulation approaches should be used in climate change studies to characterize the range of uncertainty in terms of probability distributions, without engaging in Bayesian preposterior analysis. The retort of a Bayesian advocate might be that the methodology itself is not unsound; rather, people's prior distributions have much too little variance, consistent with the overconfidence bias identified by Kahneman *et al.* (1982). The Bayesian's solution would be to train people to be better calibrated assessors of subjective probabilities, which would in general result in priors that allow for a wider range of possible outcomes.

6.4. BAYESIAN ANALYSIS IS TOO COMPLEX

Bayesian analysis can indeed be complicated. This is especially so if multiple information sources with correlated errors are to be considered; i.e. if evidence z is a vector whose elements have likelihood functions that are not conditionally independent:

$$P(z_1,z_2 \mid \theta) \neq P(z_1 \mid \theta)P(z_2 \mid \theta) \tag{6}$$

It is generally difficult to generate joint likelihood distributions, and the resulting complications can be involved.

Conditional dependence of likelihoods is likely to be a problem when applying Bayesian analysis to climate change. For example, the results of different GCMs, which could be treated as separate z_i, are likely to have correlated errors because the models use more or less similar simplifications of meteorologic and hydrologic process and rely on many of the same data sources. Consequently, the models are likely to have similar biases and cannot be considered to be conditionally independent sources of evidence.

This complication has been one reason why the Bayesian approach has not been widely used in studies attempting to detect climate warming, or to estimate the effect of climate change on hydrologic systems (Duckstein, 1994a). It is also the major reason why most developers of expert systems for medical and other applications do not adopt the full Bayesian philosophy (de Mantaras, 1990). Practical expert systems have used either approximations to Bayesian reasoning that assume, e.g. independence of distinct sources of evidence, or have used non-Bayesian methods, such as D–S reasoning (again, assuming independence).

One possible reply to this criticism is that a correct analysis must sometimes be complicated because the problems themselves are complex. Rival procedures, if they are simpler, take a risk by assuming important aspects of the problem away—meaning that they might yield incorrect inferences or suboptimal decisions. And, in many cases, rival procedures are actually more complex (Lindley, 1987).

Nonetheless, practitioners of the Bayesian approach admit that complexity is a problem, and are working hard at developing practical methods for more difficult problems (e.g. Varis and Kuikka, 1996). Simplified solution procedures based on Bayesian concepts may prove as or more accurate and optimal than non-Bayesian methods, which make their own simplifications while, at the same time, adding unnecessary complications. An example of such a simplified procedure was presented in Section 5.3. There, a simulation method was used instead of backwards dynamic programming to identify strategies for managing the development of wetlands under climate uncertainty.

7. Conclusions

Bayesian analysis is an attractive approach for analysing climate uncertainties for at least three reasons.

- It provides useful outputs: inferences and recommendations for decisions that are consistent with those inferences.
- It is based on a comprehensive and coherent framework rooted in normatively appealing assumptions.

- Compared to alternative paradigms, its concepts are familiar and, as the case studies show, its methods are often practical.

L. Duckstein and R. Krzysztofowicz provided invaluable suggestions. Comments by W. T. Bogart, P. T. Chao, and B. Venkatesh are also gratefully acknowledged, as is computational work by M. Nayal. However, the responsibility for any errors or opinions is solely the author's. This work was sponsored by the National Science Foundation, Grant SES 9223780.

References

Berger, J. O. (1985). *Statistical Decision Theory and Bayesian Analysis.* New York: Springer-Verlag.

Bernier, J. (1991). Bayesian analysis of robustness of models in water and environmental sciences. In *Water Resources Engineering Risk Assessment* (J. Ganoulis, ed.), NATO ASI Series, Vol. G29, pp. 203–229. Berlin: Springer-Verlag.

Bernier, J. (1994a). Quantitative analysis of uncertainties in water resources: aplication for predicting the effects of changes. In *Engineering Risk in Natural Resources Management (With Special References to Hydrosystems under Changes of Physical or Climatic Environment)* (L. Duckstein and E. Parent, eds), NATO ASI Series, Series E: Applied Sciences, 275, pp. 343–358. Amsterdam: Kluwer.

Bernier, J. (1994b). Statistical detection of changes in geophysical series. In *Engineering Risk in Natural Resources Management (With Special References to Hydrosystems under Changes of Physical or Climatic Environment)* (L. Duckstein and E. Parent, eds), NATO ASI Series, Series E: Applied Sciences, 275, pp. 159–176. Amsterdam: Kluwer.

Caselton, W. T. and Luo, W. (1992). Decision making with imprecise probabilities—Dempster–Shafer theory and application. *Water Resources Research* **28**, 3071–3083.

Clemens, R. T. (1991). *Making Hard Decisions: An Introduction to Decision Analysis.* Boston: PWS-Kent.

Croley, T. E. (1990). Laurentian Great Lakes double-CO_2 climate change hydrological impacts. *Climatic Change* **17**, 27–47.

de Mantaras, R. L. (1990). *Approximate Reasoning Models.* Chichester, England: Ellis Horwood.

Dowlatabadi, H. and Morgan, M. G. (1993). A model framework for integrated studies of the climate problem. *Energy Policy* **21**, 209–221.

Duckstein, L. (1994a). Elements of fuzzy set analysis and fuzzy risk. In *Decision Support Systems in Water Resources Management* (H. P. Nachtnebel, ed.), pp. 410–430. Paris: UNESCO Press.

Duckstein, L. (1994b). Engineering risk under non-steady conditions: Bayes and fuzzy logic approaches. In *Proceedings of UNESCO-IHP Conference.* Karlsruhe, Germany, June 28–30.

Duckstein, L., Bobee, B. and Bogardi, I. (1987). Bayesian forecasting of hydrologic variables under changing climatology. In *The Influence of Climate Change and Climatic Variability on the Hydrologic Regime and Water Resources*, Proceedings of the Vancouver Symposium, IAHS Publication 168, Aug. 1987, pp. 301–311.

Ellis, J. H. (1988). Acid rain control strategies. *Environmental Science and Technology* **22**, 1248–1256.

Fiering, M. B. and Rogers, P. (1991). *Climate Change and Water Resources Planning Under Uncertainty*, Draft Reporting, Institute for Water Resources. Ft. Belvoir, VA: US Army Corps of Engineers.

Fischhoff, B., Slovic, P. and Lichtenstein, S. (1979). Knowing what you want: measuring labile values. In *Cognitive Processes in Choice and Decision Behavior* (T. Wallsten, ed.). Hillsdale, NJ: Erlbaum.

Gatsonis, C., Hodges, J., Kass, R. and Singpurwalla, N. (eds) (1992). *Bayesian Statistics and Technology: Case Studies.* New York: Springer-Verlag.

Hobbs, B. F., Chao, P. T., Nayal, M. and Bogart, W. T. (1994). Climate warming: analyzing implications for Great Lakes management. In *Risk Based Decision Making in Water Resources VI* (Y. Y. Haimes, D. Moser and E. Z. Stakhiv, eds), pp. 295–316. New York: American Society of Civil Engineers.

Hobbs, B. F., Chao, P. T. and Venkatesh, B. N. (1996). Using decision analysis to include climate change in water resources decision making. *Climatic Change*, in press.

Kahneman, D., Slovic, P. and Tversky, A. (1982). *Judgment Under Uncertainty: Heuristics and Biases.* New York: Cambridge University Press.

Krzysztofowicz, R. (1983). Why should a forecaster and a decision maker use Bayes theorem. *Water Resources Research* **19**, 327–336.

Krzysztofowicz, R. (1994). Strategic decisions under nonstationary conditions: a stopping-control paradigm. In *Engineering Risk in Natural Resources Management (With Special References to Hydrosystems under Changes of Physical or Climatic Environment)* (L. Duckstein and E. Parent, eds), NATO ASI Series, Series E: Applied Sciences, 275, pp. 359–372. Amsterdam: Kluwer.

Lave, L. B. and Dowlatabadi, H. (1993). Climate change: the effects of personal beliefs and scientific uncertainty. *Environmental Science and Technology* **27**, 1962–1972.

Lindley, D. V. (1987). The probability approach to the treatment of uncertainty in artificial intelligence and expert systems. *Statistical Science* **2**, 3–44.

Morgan, M. G. and Henrion, M. (1990). *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis.* Cambridge: Cambridge University Press.

Nachtnebel, H. P. and Duckstein, L. (1994). Incorporating risk and imprecision into water related decision making: an Austrian case study for instream water requirements. In *Decision Support Systems in Water Resources Management* (H. P. Nachtnebel, ed.), pp. 431–451. Paris: UNESCO Press.

Oden, G. C. (1977). Integration of fuzzy logical information. *Journal of Experimental Psychology: Human Perception and Performance* **4**, 565–575.

Patwardhan, A. and Small, M. J. (1992). Bayesian methods for model uncertainty analysis with application to future sea level rise. *Risk Analysis* **12**, 513–523.

Poirier, D. J. (1988). Frequentist and subjectivist perspectives on the problems of model building in economics. *Journal of Economic Perspectives* **2**, 121–144, 167–170.

Raiffa, H. (1968). *Decision Analysis, Introductory Lectures on Choices under Uncertainty.* Reading, MA: Addison-Wesley.

Rust, J., Pagan, A. and Geweke, J. (1988). Comment on Poirier. *Journal of Economic Perspectives* **2**, 145–166.

Schneider, S. H. (1994). Detecting climatic change signals: Are there any 'fingerprints'? *Science* **263**, 341–347.

Slivitsky, M. and Mathier, L. (1994). Climatic changes during the 20th century on the Laurentian Great Lakes and their impacts on hydrologic regime. In *Engineering Risk in Natural Resources Management (With Special References to Hydrosystems under Changes of Physical or Climatic Environment)* (L. Duckstein and E. Parent, eds), NATO ASI Series, Series E: Applied Sciences, 275, pp. 235–252. Amsterdam: Kluwer.

Tonn, B. E. and Schaffhauser, A. J. (1994). *An Approach to Understanding, Representing, and Managing Uncertainty in Integrated Resource Planning*, ORNL/CON-399. Oak Ridge, TN: Oak Ridge National Laboratory.

Tsao, H.-S. J., Fang, S.-J. and Lee, D. N. (1993). A Bayesian interpretation of the linearly-constrained cross-entropy minimization problem. *Engineering Optimization* **22**, 65–73.

Varis, O. and Kuikka, S. (1996). An influence diagram approach to Baltic Salmon Management. Paper presented at *INFORMS Fall 1996 National Meeting*, Atlanta, November 5.

von Winterfeldt, D. and Edwards, W. (1986). *Decision Analysis and Behavioral Research.* New York: Cambridge University Press.

Waggoner, P. E. (ed) (1990). *Climate Change and U.S. Water Resources.* New York: J. Wiley.

Wigley, T. M. L. and Barnett, T. P. (1990). Detection of the greenhouse effect in the observations. In *Climate Change, The IPCC Scientific Assessment*. Cambridge: Cambridge University Press.

Zadeh, L. A. (1983). The role of fuzzy logic in the management of uncertainty in expert systems. *Fuzzy Sets and Systems* **11**, 199–227.