Coping with uncertainty in ecological advice: lessons from fisheries

John Harwood[1] and Kevin Stokes[2]

[1]Centre for Conservation Science, University of St Andrews, The Observatory, Buchanan Gardens, St Andrews, Fife, UK KY16 9LZ
[2]The New Zealand Seafood Industry Council, Private Bag 24-901, Wellington, New Zealand

All ecologists are familiar with uncertainty, at least at the level of whether they should reject a null hypothesis. Uncertainty is, however, pervasive and its characterization is essential if we are to understand our effects on ecosystems. Traditional fisheries management has a poor track record for confronting uncertainty, but most management authorities are now committed to a precautionary approach. As a result, some of the most interesting methods for taking account of uncertainty in ecological systems have been developed by fisheries scientists. These methods evaluate the relative performance of different management procedures with the use of mathematical and statistical models that synthesize knowledge and speculation about the system of interest. Recent advances in computer-intensive statistics have made it possible to combine this approach with model fitting, so that the uncertainties and risks associated with different outcomes of management can be quantified. We show how this methodology can be applied to a range of ecological problems where the advice that scientists provide to decision makers is likely to be clouded by uncertainty.

Policy makers and environmental managers frequently ask ecologists for advice about the environmental consequences of human activities. However, ecological systems are usually complex, nonlinear and strongly influenced by stochasticity [1]. As a result, it is often impossible to predict their dynamics in any detail. The situation is complicated further because the available information about the way that these systems function is often equivocal.

Ecologists who seek to inform policy makers must distil the results of complex analyses that predict uncertain outcomes into simple and clear advice. They therefore face a dilemma: do they present a simplification of the situation that is persuasive but might pay insufficient attention to the reliability of their conclusions; or do they emphasize the UNCERTAINTIES (see Glossary) inherent in their analysis [2]? The first option is likely to result in the caveats associated with the advice being ignored, the second is likely to result in the advice itself being ignored. Even if the advice is accepted, a high degree of uncertainty about the potential outcomes of management actions provides many opportunities for confrontation among different interest groups, and this can hinder the development of consensus.

Some authors [3,4] have argued that the failure of ecologists to explain adequately the uncertainties associated with their advice has diminished their influence on the decision-making process. Even if the true situation is not as bad as this, ecologists should improve the way that they provide advice about uncertainty, not least because the general adoption of the precautionary principle and precautionary approaches to management (Box 1) requires an assessment of the RISK that serious or irreversible environmental damage will occur as a result of management actions. This risk is an inevitable consequence of the uncertainties that are inherent in our knowledge of ecological systems, and ecologists must develop rigorous methods for evaluating these uncertainties. Biology is not the only, and probably not the most influential, discipline involved in decision making about environmental issues [2], so methods that can take account of the uncertainties associated with advice from all relevant disciplines are needed [5].

The first step in quantifying risk is to identify the sources of uncertainty (Box 2), but understanding the implications of all these sources for particular management actions is a much more challenging problem.

Glossary

Bayesian statistics: considers observations as known quantities that might have been generated by a variety of processes. Distributions are assigned to the parameters of these processes to enable inferences to be drawn about them.
Frequentist statistics: considers observations to be the random results of an underlying process with fixed but unknown parameters. Data are used to estimate the values of these parameters, based on assumptions about the statistical distributions from which the data were drawn.
Likelihood: the probability that a sample has been randomly drawn from a particular probability distribution; it is treated as a function of the parameters of this distribution.
Operating model: a plausible model of an ecological system used to test the robustness of management procedures to uncertain system structures, and to evaluate the tradeoffs between conflicting objectives.
Risk: the probability that a hazardous outcome will occur. It is a consequence of uncertainty: if there is no uncertainty, the concept of risk is irrelevant because the probability of the outcome is 1 or 0.
Uncertainty: incomplete information about a particular subject.
Utility: the value to society associated with a potential outcome of management action; it can be arbitrarily assigned or arrived at by consensus among stakeholders.

Corresponding author: John Harwood (jh17@st-andrews.ac.uk).

http://tree.trends.com 0169-5347/$ - see front matter © 2003 Elsevier Ltd. All rights reserved. doi:10.1016/j.tree.2003.08.001

Box 1. The precautionary principle and precautionary approaches to management

The precautionary principle forms the basis for most European environmental law [33], and is a fundamental component of several international conventions [34]. Yet, it is not defined in the Treaty on European Union (EU), and the most widely used definition [Principle 15 of the Declaration from the 1992 United Nations Conference on Environment and Development (UNCED)] – 'where there are threats of serious or irreversible damage, lack of full scientific certainty shall not be used as a reason for postponing cost-effective measures to prevent environmental degradation' – is vague [33].

A recent review by the Commission of the European Communities (http://europa.eu.int/comm/dgs/health_consumer/library/pub/pub07_en.pdf) has helped to clarify that, at least in the EU, the precautionary principle should be applied within a framework involving the assessment, management and communication of risk. The role of scientific advice in this process is to identify the environmental hazards (see Box Glossary) associated with a particular issue, determine the exposure of sensitive components of the ecosystem to these hazards and establish the potential responses of the system to this exposure [11]. Using this information, the performance of management procedures that reduce exposure, or modify the exposure response, to an acceptable level can be assessed. Establishing what level of risk is acceptable is a political responsibility. The review also recommends that 'when the available data are inadequate or non-conclusive, a prudent and cautious approach…could be to opt for the worst-case hypothesis'. Unfortunately, this can result in the formulation of an escalating series of increasingly pessimistic scenarios, because there is no limit to these [2]. Combining the results of several worst-case scenarios can result in risks being reduced to apparently tiny levels that are hard to justify economically [35]. However, hazards that pose a low risk to an average individual might pose a much higher risk to some segment of the population, so this process might be much less conservative than it appears [5].

We suggest that no scenario should be ignored, provided it can be formulated into an operating model. The importance attached to the results obtained under a particular scenario should be related in some way to its plausibility [2], but it is not clear how plausibility should be assessed. The obvious recourse to expert opinion is likely to depend on the value system of the experts that are chosen [2].

In a parallel development to the refinement by the EU of the precautionary principle, the UN Food and Agriculture Organisation produced a set of guidelines for a precautionary approach to fisheries management [36]. The guidelines recommend that managers:

- Avoid changes that are not potentially reversible;
- Identify undesirable outcomes and measures to avoid or correct these changes;
- Act to conserve productive capacity where there is uncertainty;
- Limit fishing capacity when resource productivity is uncertain.

A crucial distinction between this approach and the precautionary principle, at least in its EU incarnation, is that the undesirable outcomes include economic and social consequences, as well as environmental effects.

Box Glossary

Hazard: a potentially negative effect on the environment or human health as a result of a particular activity.

Box 2. Sources of uncertainty

Although some authors [3,4] follow Smithson [37] in classifying uncertainty as a (small) subset of ignorance, we prefer the standard dictionary definitions of these terms that imply the reverse: that ignorance is an extreme form of uncertainty. A particularly useful distinction [38] is that between epistemic uncertainty (uncertainty in things that can be measured) and linguistic uncertainty (uncertainty in the language used to describe or classify desired states). Here, we focus on the sources of epistemic uncertainty. They are:

(i) Process stochasticity [23], which is a consequence of demographic and environmental stochasticity, and the apparently random behaviour of systems that have chaotic dynamics. It is sometimes referred to as natural variation [38] or natural stochasticity [39].

(ii) Observation error, which is made up of measurement error, a consequence of the way in which observations are taken (e.g. the choice of sampling strategy, or errors in data collection), and estimation (or inference) error, which is the inaccuracy and imprecision introduced by the method of statistical inference used to estimate system parameters from observations.

(iii) Model error. All models are caricatures of reality and thus provide an incomplete, and potentially misleading, representation of system dynamics. Model mis-specification has two major consequences: (a) it can contribute to estimation error through the inferential process; and (b) it will induce further errors if the model is used in forecasting.

(iv) For managed systems, implementation error must be taken into account. In a fisheries context, this might include failure to meet proscribed limits on catches as a result of imperfect policy implementation or changes in market forces that alter the incentives for fishers. The introduction of economic effects [30] adds a human dimension to uncertainty [40]. In a conservation context, sources of implementation error include delays in the establishment of protected areas, or inadequate protection within them. More generally, implementation error attempts to capture the consequences of what John Maynard Keynes called 'the insane and irrational springs of wickedness in most men' [41].

Figure I shows how all these sources of uncertainty are represented in an operating model framework. The text in italics within each box indicates which sources of uncertainty are accounted for by that sub-model.

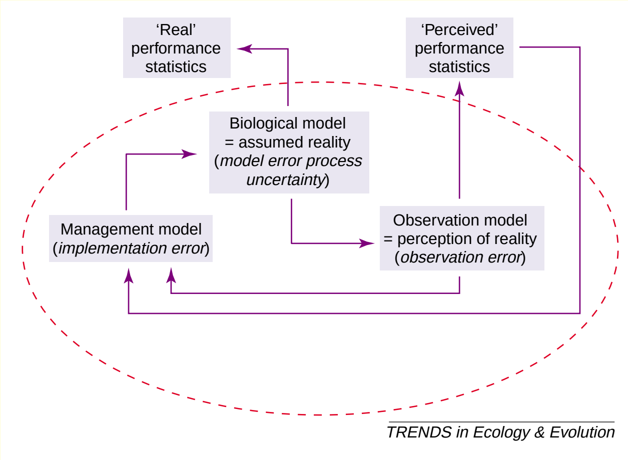

Figure I.

Here, we describe the advances that fisheries scientists have made in dealing with this problem and show how their approach can be generalized into a rigorous statistical framework. We also outline how this framework can be used to provide advice about other ecological problems that takes account of uncertainty in a formal way, yet remains accessible to decision makers.

Lessons learned from fisheries and whaling

The failure of fisheries managers to harvest fish stocks in a sustainable way is often used as a textbook example of what happens when inadequate account is taken of uncertainty in measurement, knowledge and the implementation of regulations [6]. However, fisheries scientists have recently made major advances in addressing these problems by applying a conceptual framework that has its origins in operational research. This has enabled them to address many of the implications of uncertainty in a rigorous and consistent way, and to provide decision makers with information about the risks associated with different management options.

Their approach involves developing computer models of the underlying biological processes (the biological or process model), the way that information is gathered about these processes (the observation model) and the way in which fisheries might affect these processes (the management model). These are combined into an OPERATING MODEL (Box 2). The biological model attempts to describe basic population processes (birth, death, growth and migration) and the way that these are affected by population size and environmental factors. The observation model mimics the activities of a virtual ecologist [7] and reflects the way that data about the system are collected and analyzed. It might be as simple as associating an error structure with estimates of the abundance of the target species but it can, for example, attempt to replicate the way that the age structure of a population is estimated from samples of fish taken by a commercial fishery. In the latter case, it would include sub-models of the nonrandom way that these samples were collected and of the errors involved in determining the age of each individual. A typical management model would include the way in which data that have been processed through the observation model are used to set quotas and the effect of these quotas on the exploited population. These effects would include the impacts of bycatches and the discarding of fish that are caught but not landed.

Often, some of the most important underlying biological processes are either poorly understood or can be described adequately by many rival models. The potential effects of climate change are an obvious example [2]. Rather than trying to identify a single best model, ignorance is accounted for explicitly by incorporating a wide range of biologically plausible complexities into the operating model. This is a form of scenario planning [8]. The robustness of different management strategies to these complexities is then assessed. Those strategies that perform well across all scenarios are preferred over more sensitive ones.

Once the structure of the operating model has been agreed, repeated simulations of the entire model framework are run, using randomized resampling from observed or assumed probability distributions for the model parameters. The results of these simulations are used to generate frequency distributions for the various possible outcomes, including information (known as performance statistics) about the way in which the system performs against the stated objectives of management. Statistics can be extracted directly from the biological model and compared with the equivalent values that are provided by the observation models. This comparison might show, for example, that estimates of population size derived from an observation model based on data collected from the fishing industry are relatively insensitive to changes in actual abundance. In this case, a management approach based solely on observations from the fishery would perform badly if one of the objectives was to ensure that the exploited population was unlikely to fall below a specified reference level.

Management might have many objectives, some of which are contradictory. If UTILITIES can be assigned to the different objectives, decision analysis [9–12] can be used to enable managers to compare the economic and social benefits that are likely to accrue from different management actions.

Decision analysis and the IWC

One of the first attempts to use this approach, which also provides an illustration of what happens when utilities cannot be assigned, was made by the Scientific Committee of the International Whaling Commission (IWC). The Committee was asked to develop a revised management procedure for baleen whales [13] that would achieve three management objectives: (i) stable catches; (ii) no serious increase in the risk of extinction as a result of exploitation; and (iii) highest continuing yields. The IWC did not suggest what utilities should be attached to each of these objectives, probably because member states had very different views on this [1]. Instead, the Committee agreed on a set of performance statistics that were used to measure how well different management approaches might achieve each of these objectives. Several management procedures (a combination of the observation and management models described in Box 2) were then proposed by Committee members and simple operating models were used to filter out those procedures that performed particularly badly. The remaining procedures were then refined and assessed using more complex operating models that incorporated, for example, whale social organization, spatial heterogeneity, the potential effects of climate change and Allee effects. Although little is known about the spatial structure of whale populations or how climate change might affect their dynamics, the Committee was able to formulate a range of plausible models about both aspects of baleen whale population biology. The management procedure that was finally adopted performed well with all of these models.

Decision analysis and fisheries

The same framework is being used increasingly in the southern hemisphere [9]. For example, the performance of a range of management procedures for the eastern stock of gemfish *Rexea solandri* off the coast of south-eastern Australia was evaluated using age-structured operating models [14]. There was considerable uncertainty about the relationship between the recruitment of young fish and abundance for this stock, so a range of forms and error structures for this relationship was incorporated into the operating model. A management procedure based on a simple inference model was found to be preferable to one that used detailed age information from the catch, even though the latter resulted in higher average catches, confirming the results of earlier theoretical analyses [15]. The simple procedure resulted in less variable catches and

Box 3. Bayesian statistical inference

Bayesian statisticians are concerned with the support provided by observation (symbolized by a vector $\mathbf{D}$) for different models (H_i) that have been developed to describe the system of interest, or for the parameters ($\boldsymbol{\theta}$) of these models,. This support is often represented as $P\{H_i|\mathbf{D}\}$ or $P\{\boldsymbol{\theta}|\mathbf{D}\}$. Conventional frequentist statisticians are concerned with $P\{\mathbf{D}|H_i\}$, the probability of obtaining the observed data under H_i. They reject H_i as an explanation of the data if this probability is less than some arbitrarily defined level (usually 5% or 1%). $P\{\boldsymbol{\theta}|\mathbf{D}\}$ (the posterior distribution) can be derived using Bayes theorem from Eqn I:

$$P\{\boldsymbol{\theta}|\mathbf{D}\} = P\{\boldsymbol{\theta}\}P\{\mathbf{D}|\boldsymbol{\theta}\}/k \quad [I]$$

$P\{\mathbf{D}|\boldsymbol{\theta}\}$ is the likelihood for $\boldsymbol{\theta}$ and k is a normalizing constant (Eqn II)

$$k = \int P\{\boldsymbol{\theta}\}P\{\mathbf{D}|\boldsymbol{\theta}\}d\boldsymbol{\theta} \quad [II]$$

$P\{\boldsymbol{\theta}\}$ (the prior distribution) represents our current information about the distribution of $\boldsymbol{\theta}$. It can be based on actual data, or on a subjective assessment of what is likely. In principle, this is an effective way of taking account of ignorance, but it is important to determine the sensitivity of results to the prior distributions that have been used, especially when these are based on opinion rather than data. Unfortunately, this warning is more frequently issued than it is observed. The use of priors that appear to have been chosen arbitrarily is probably one of the major reasons why biologists are suspicious of the Bayesian approach, because it seems to provide something for nothing. When priors are based on opinion rather than data, a variety of different priors must be tried before the results of the analysis are accepted. If those results are sensitive to the choice of prior, and they often are, then every effort should be made to find additional sources of information that can be used to form a more justifiable prior distribution. This can result in more biological information being incorporated into a Bayesian analysis than would ever be possible using a frequentist approach.

To estimate k (Eqn II), the entire area under the likelihood distribution must be evaluated, whereas a frequentist analysis requires only the maximum value of the likelihood. In most cases, the integral cannot be obtained analytically and computer-intensive, numerical techniques must be used (Box 4).

could be implemented with data that were relatively cheap to collect.

Management procedures that have now been implemented for five fish, shellfish and seal stocks along the coast of southern Africa were evaluated using operating models that incorporated environmental change, environmental catastrophes, a range of values for the basic biological parameters and different fishing practice [16]. In four cases, this led to the adoption of simple harvest rules based directly on the data collected from the fisheries rather than on inferred quantities.

Incorporating and evaluating all sources of uncertainty: the state-space framework

Operating models can be viewed as a framework for statistical inference about the system as well as a way of evaluating the performance of management strategies. In this way, the properties of BAYESIAN statistical inference (Box 3) can be used to combine the processes of parameter estimation and risk assessment. Data collected as part of the management procedure are used to refine prior information about the distributions of the parameters of interest to provide a joint probability distribution for all of these parameters. This probability distribution accurately reflects all of the major uncertainties about the system. It can therefore be used to evaluate the uncertainties associated with predictions of the potential outcomes of different management actions. Combining estimation and evaluation avoids many of the problems that are a consequence of the independent estimation of individual parameters. These include the need to make strong assumptions about the underlying probability distributions and covariance structure of the model parameters.

Two recent advances have made this combination possible: the application of a state-space approach to ecological problems [17,18], and the development of computer-intensive statistical methods that can use relatively sparse data to parameterize complex models (Box 4). The basic state-space framework comprises a process model and an observation model, just like the operating models used by fisheries scientists. Model uncertainty is accounted for in the same way, by using a range of process models. In most ecological applications, the process model is an age-structured or stage-structured population model, which is often summarized in matrix notation form [19]. A useful trick is to disaggregate this matrix into a chain of sub-matrices that describe individual processes, such as survival, ageing, reproduction and migration [20,21]. This makes it easier to implement different process models. The final step is to add a series of matrices that represent the management process [18].

The usefulness of Bayesian statistics for addressing problems in fisheries and conservation has been described in detail elsewhere [9,10,22–24]. The ability of Bayesian statistics to take account of model uncertainty and ignorance, and the fact that probabilities can be associated with specific outcomes (e.g. that there is an x% probability that the population will fall below some target level within a particular timeframe) is particularly valuable. In addition, complex process models that contain more parameters than there are data points can be used if appropriate prior distributions can be defined [10]. These prior distributions are most informative if they are based on data rather than on assumptions. They provide a useful way to incorporate valuable biological observations, which would otherwise be unsuitable for model fitting, into the evaluation process. The extent to which the available data support these complex models can be assessed by examining the posterior distributions (Box 3): if there is no support in the data, the posterior distributions will be identical to the priors. This comparison also provides a way to assess the plausibility of different biological models (Box 2). Finally, individual prior distributions can be manipulated to determine the most cost-effective way of reducing the uncertainties associated with predicted outcomes.

There is still considerable debate about the appropriateness of Bayesian statistics in ecology [22,25]. Fortunately, it is possible to implement the same framework using conventional FREQUENTIST or LIKELIHOOD-based approaches [2,26].

Lessons for ecology

The approach used by the IWC in formulating the revised management procedure has also been used to develop a

Box 4. Computer-intensive statistics: a reader's guide

Here, we provide brief descriptions of the techniques that are in common use for evaluating likelihoods and posterior distributions. They are intended to provide readers with some insight into how the techniques work without having to digest the rather daunting specialist literature. More technical detail can be found in [42–45]. Appendix A of [9] provides a particularly helpful guide to the practical implementation of these techniques in fisheries management.

The Kalman Filter

This is an analytical method for evaluating the likelihood if the underlying model is linear and observation errors are normally distributed [46,47].

Grid search

If the likelihood involves less than five or so parameters, it might be possible to carry out a systematic search of the entire parameter space. For more parameters, this is impractical and some efficient way of sampling the parameter space is required.

Markov Chain Monte Carlo (MCMC) methods

These are a class of algorithms that provide correlated samples from a likelihood or posterior distribution [43,44]. Hundreds of thousands of samples might be required for reliable inference. The algorithms move through parameter space using a random or directed walk. An initial vector of parameter values is updated, piece by piece, according to random draws from a pre-specified proposal distribution to create a new sample that forms the starting point for the next set of draws. Providing the appropriate updating strategy is chosen, the distribution of the sample vectors will converge on the required distribution. The choice of algorithm depends on the nature of the proposal distribution that is used.

The Gibbs sampler

If the form of the posterior or likelihood is known, a proposal distribution that is proportional to the likelihood can be chosen. In this case, every new vector chosen by the algorithm is used. The Gibbs sampler forms the basis for the popular WinBUGS (Bayesian analysis Using Gibbs Sampling) software (http://www.mrc-bsu.cam.ac.uk/bugs/winbugs/contents.shtml), although WinBUGS will also implement other algorithms.

The Metropolis-Hastings sampler

If the form of the posterior is not known, a proposal distribution must be chosen. There are many tricks for choosing a proposal distribution so that parameter space is explored at a reasonable rate [43,44]. The probability of accepting a new sample depends on the match between the likelihood and the proposal at both the current and the chosen points in parameter space.

The SIR (sample-importance-resample) algorithm

This is an alternative to MCMC. If a Bayesian analysis is being performed, a large number of candidate vectors are chosen at random from the prior distributions and their importance (i.e. their contribution to the likelihood) is evaluated and stored [9]. Samples of these vectors are then drawn, with replacement and in proportion to their importance, to generate a sample from the posterior distribution.

Sequential importance sampling (SIS)

This is an extension of SIR adapted for time series analysis [45]. Candidate vectors are chosen at random from the prior distributions and their importance is evaluated at each time step. The vectors are resampled at each time step according to their importance. Vectors that make a small contribution to the likelihood tend to be lost, so that eventually only one might be left. This phenomenon is known as particle depletion. There are many tricks to alleviate particle depletion [45], including kernel smoothing [21], in which new samples are created from distributions centred on the samples that survive each time step.

methodology for assessing the conservation status of marine mammals in the USA [27], and to compare the effectiveness of different management strategies for the saiga antelope *Saiga tatarica* [28]. However, the framework that we have described here is applicable to a wide range of ecological problems. For example, it offers an alternative to the methods currently used [29] to take account of uncertainty when assigning species to the categories for conservation status developed by the World Conservation Union (IUCN).

In almost every case where this approach has been used [13–15,28], the most successful management procedures have been those that use observational information directly, rather than using it to estimate the parameters of a detailed biologically model. These findings add weight to recent calls [30] for simple robust conservation strategies that have relatively low data demands. This does not mean, however, that detailed ecological information is unnecessary for effective management. The operating model framework is one of the most efficient ways to assess the robustness of competing management strategies. But, it is only effective if models of the underlying biological processes that take account of the best scientific knowledge, and the uncertainties associated with this knowledge, are available to test the robustness of different management strategies.

What next?

Computer-intensive risk assessment of the kind that we have described here provides a valuable decision support tool that enables the precautionary principle to be implemented in a rigorous and scientific way. Although there is some evidence that decision makers find advice that is couched in terms of probabilities difficult to assimilate [31], this is primarily a problem of presentation. Certainly, advice about the threats posed by variant Creutzfeldt-Jakob disease was readily accepted by decision makers, in spite of the massive uncertainties associated with the original risk assessments [32]. However, there is no doubt that the mathematical and statistical theory associated with the framework that we have described here can be intimidating, and most ecologists would balk at implementing it without expert advice. Computer packages to apply some of the techniques are now available, but even these require considerable experience if they are to be applied effectively (Box 4). We need user-friendly software that will enable ecologists to assemble operating models without having to recast their biological knowledge in matrix form, and that will help them chose appropriate prior distributions and search algorithms for use in their assessments.

Acknowledgements

We thank Mark Bravington, Mark Burgman, Steve Buckland, Mark Mangel, Jason Matthiopoulos, Hugh Possingham and Len Thomas for many useful discussions and comments about the article.

References

1 Mangel, M. *et al.* (1996) Principles for the conservation of wild living resources. *Ecol. Appl.* 6, 338–362
2 Ludwig, D. *et al.* (2001) Ecology, conservation and public policy. *Annu. Rev. Ecol. Syst.* 32, 481–517

3 Dovers, S.R. *et al.* (1996) Uncertainty, ecology, sustainability and policy. *Biodiv. Conserv.* 5, 1143–1167
4 Handmer, J. *et al.* (2001) Managing ecosystems for sustainability: challenges and opportunities. In *Ecology, Uncertainty and Policy: Managing Ecosystems for Sustainability* (Handmer, J. *et al.*, eds), pp. 291–303, Pearson Education
5 Finkel, A. (1997) Disconnect brain and repeat after me: "risk assessment is too conservative". *Ann. N. Y. Acad. Sci.* 837, 397–417
6 Ludwig, D. *et al.* (1993) Uncertainty, resource exploitation, and conservation: lessons from history. *Science* 17, 36
7 Tyre, A. *et al.* (2001) Inferring process from pattern: can territory occupancy provide information about life history parameters? *Ecol. Appl.* 11, 1722–1737
8 Schoemaker, P.J.H. (1995) Scenario planning: a tool for strategic thinking. *Sloan Manage. Rev.*, 25–40
9 Punt, A. and Hilborn, R. (1997) Fisheries stock assessment and decision analysis: the Bayesian approach. *Rev. Fish Biol. Fish.* 7, 35–63
10 Wade, P. (2001) The conservation of exploited species in an uncertain world: novel methods and the failure of traditional techniques. In *Conservation of Exploited Species* (Reynolds, J. *et al.*, eds), pp. 110–143, Cambridge University Press
11 Harwood, J. (2000) Risk assessment and decision analysis in conservation. *Biol. Conserv.* 95, 219–226
12 Maguire, L. (1996) Making the role of values in conservation explicit: values and conservation biology. *Conserv. Biol.* 10, 914–916
13 Cook, J.G. (1995) The International Whaling Commission's Revised Management Procedure as an example of a new approach to fishery management. In *Whales, Seals, Fish and Man* (Blix, A.S. *et al.*, eds), pp. 647–657, Elsevier Science
14 Punt, A. and Smith, A. (1999) Harvest strategy evaluation for the eastern stock of gemfish (*Rexea solandri*). *ICES J. Mar. Sci.* 56, 860–875
15 Ludwig, D. and Walters, C.J. (1985) Are age-structured models appropriate for catch-effort data? *Can. J. Fish. Aquat. Sci.* 42, 1066–1072
16 Geremont, H. *et al.* (1999) Development and application of management procedures for fisheries in southern Africa. *ICES J. Mar. Sci.* 56, 952–966
17 Akoi, M. (1990) *State Space Modelling of Time Series*, Springer-Verlag
18 Trenkel, V. *et al.* (2000) Fitting population dynamics models to count and cull data using sequential importance sampling. *J. Am. Stat. Assoc.* 95, 363–374
19 Caswell, H. (2001) *Matrix Population Models: Construction, Analysis and Interpretation*, Sinauer Associates
20 Thomas, L. *et al.* A unified framework for modelling wildlife population dynamics. *Aust. N.Z. J. Stat.* (in press)
21 Buckland, S.T. *et al.* State space models for the dynamics of wild animal populations. *Ecol. Model.* (in press)
22 Ellison, A.M. (1996) An introduction to Bayesian inference for ecological research and environmental decision-making. *Ecol. Appl.* 6, 1036–1046
23 Hilborn, R. and Mangel, M. (1997) *The Ecological Detective: Confronting Models with Data*, Princeton University Press
24 Wade, P.R. (2000) Bayesian methods in conservation biology. *Conserv. Biol.* 14, 1308–1316
25 Dennis, B. (1996) Discussion: should ecologists become Bayesians? *Ecol. Appl.* 6, 1095–1103
26 Taylor, B.L. *et al.* (1996) A Bayesian approach to classification criteria for spectacled eiders. *Ecol. Appl.* 6, 1077–1089
27 Milner-Gulland, E.J. *et al.* (2001) Competing harvesting strategies in a simulated population under uncertainty. *Anim. Conserv.* 4, 157–167
28 Akçakaya, H.R. *et al.* (2000) Making consistent IUCN classifications under uncertainty. *Conserv. Biol.* 14, 1001–1013
29 Taylor, B.L. *et al.* (2000) Incorporating uncertainty into management models for marine mammals. *Conserv. Biol.* 14, 1243–1252
30 Armsworth, P.R. and Roughgarden, J. (2001) An invitation to ecological economics. *Trends Ecol. Evol.* 16, 229–234
31 Anderson, J.L. (1998) Embracing uncertainty: the interface of Bayesian statistics and cognitive psychology. *Conserv. Ecol.*, 2 http://www.consecol.org/vol2/iss1/art2
32 Ghani, A.C. *et al.* (2003) Factors determining the pattern of the variant Creutzfeldt-Jakob disease (vCJD) epidemic in the UK. *Proc. R. Soc. Lond. Ser. B* 270, 689–698
33 Foster, K.R. *et al.* (2000) Science and the precautionary principle. *Science* 288, 979–981
34 Anon (2000) *Communication from the Commission on the Precautionary Principle*, Commission of the European Communities
35 Burmaster, D. and Harris, R. (1993) The magnitude of compounding conservatisms in superfund risk assessments. *Risk Anal.* 13, 131–134
36 FAO, (1995) *Code of Conduct for Responsible Fisheries*, Food and Agriculture Organisation of the United Nations
37 Smithson, M. (1989) *Ignorance and Uncertainty: Emerging Paradigms*, Springer-Verlag
38 Regan, H. *et al.* (2002) A taxonomy and treatment of uncertainty for ecology and conservation biology. *Ecol. Appl.* 12, 618–628
39 Burgman, M.A. *et al.* (2000) Population viability analysis for conservation: the good, the bad and the undescribed. In *Genetics, Demography and Viability of Fragmented Populations* (Clark, G. and Young, A., eds), pp. 97–112, Cambridge University Press
40 Ralls, K. and Taylor, B.L. (2000) Better policy and management decisions through explicit analysis of uncertainty: new approaches for marine conservation. *Conserv. Biol.* 14, 1240–1242
41 Keynes, J.M. (1971) *Two Memoirs*, Augustus M Kelley
42 Lui, J.S. (2001) *Monte Carlo Strategies in Scientific Computing*, Springer-Verlag
43 Gilks, W. *et al.* (1996) Introducing Markov chain Monte Carlo. In *Markov Chain Monte Carlo in Practice* (Gilks, W. *et al.*, eds), pp. 1–19, Chapman & Hall/CRC
44 Carlin, B.P. and Louise, T.A. (2000) *Bayes and Empirical Bayes Methods for Data Analysis*, Chapman & Hall/CRC
45 Doucet, A., *et al.* eds (2001) *Sequential Monte Carlo Methods in Practice* Springer-Verlag
46 Newman, K.B. (1998) State-space modeling of animal movement and mortality with application to salmon. *Biometrics* 54, 1290–1314
47 Besbeas, P. *et al.* (2003) The efficient integration of abundance and demographic data. *Appl. Stat.* 52, 1–8