

An application of the Bayesian approach to stock assessment model uncertainty

T. R. Hammond and C. M. O'Brien



Hammond, T. R. and O'Brien, C. M. 2001. An application of the Bayesian approach to stock assessment model uncertainty. – ICES Journal of Marine Science, 58: 648–656.

Bayesian methods have a number of advantages that make them especially useful in the provision of fisheries management advice: they permit formal decision analysis, and they facilitate the incorporation of model uncertainty. The latter may be particularly useful in the management of contentious fisheries, where different nations and interest groups may suggest alternative assessment models and management – each likely to imply different findings, even when using the same data. Such situations might be approached in a number of different ways. For example, one might attempt to choose a best model from all those available and to base decisions on it alone. Alternatively, one might make decisions that lead to acceptable outcomes under all envisaged models; or one could reach decisions that are good on average (where average is taken over the set of all competing models and is weighted by a measure of how well each model coheres with available information). This last approach is advocated in this paper, and a Bayesian technique for achieving it is presented and discussed. The main points of the paper are illustrated with a hypothetical application of the technique to the rebuilding of the biomass of haddock by a selective culling of seals.

Key words: decision analysis, Bayesian networks, model uncertainty, ecosystem effects, fisheries management.

Received 20 March 2000; accepted 26 January 2001.

T. R. Hammond and C. M. O'Brien: CEFAS Lowestoft Laboratory, Pakefield Road, Lowestoft, Suffolk NR33 0HT, UK. Correspondence to T. R. Hammond, tel.: +44 (0) 1502 524 213; fax: +44 (0) 1502 524213; e-mail: t.hammond@cefas.co.uk

Introduction

The literature on stock assessment in fisheries has traditionally concentrated on numerical methods and algorithms for estimating parameters of interest to fisheries management using population dynamics models. There are now a wide variety of estimation methods that incorporate uncertainty in both data and model parameters (Deriso *et al.*, 1985; Fournier and Archibald, 1982; Fournier *et al.*, 1998; Gudmundsson, 1994; McAllister and Ianelli, 1997; Richards and Schnute, 1997). Whilst these statistical algorithms allow one to quantify the precision in parameter estimates, and permit formal decision analysis (McAllister and Kirkwood, 1998), results are usually conditional upon a single underlying model. There are published cases in which stock assessments consider structural uncertainty (e.g. Punt and Butterworth, 1995; Patterson, 1999), but this is by no means routine.

In many situations, differences between methods of parameter estimation are trivial when contrasted with

the impact of incorporating different structural models into statistical inference (Buckland *et al.*, 1997). The principal objective of this paper is to show the importance of acknowledging model uncertainty within a hypothetical fishery system.

In contentious multi-national fisheries, consideration of model uncertainty is necessary because different nations and interest groups will tend to propose or support different models. It is only natural that nations, groups and individuals should support the creation and adoption of models that advance their position. It is natural, too, that such groups should be critical of a management system that rejects their advocated models. A well-designed management system should employ techniques that build consensus between groups rather than those that reinforce antagonism. A secondary objective of this paper is to illustrate the use of model averaging as a tool to aid the process of decision making.

This paper advocates the use of Bayesian belief networks as a practical decision making tool. These

networks have already been used in a fisheries context (Kuikka *et al.*, 1999), and they are well established in the statistical literature (Jensen, 1996). In the hypothetical fishery system considered here, the use of Bayesian networks leads to some surprising results and may hold valuable lessons for fisheries modellers and managers.

What are Bayesian networks?

A Bayesian network is always represented visually with a set of vertices (V) and a set of edges. Each vertex represents a specific random variable, and the vertex is labelled with that variable's name. Each variable must have a finite number of mutually exclusive states: examples of such being True or False; Red, Green or Blue; Increasing or Decreasing. Each edge represents a causal relationship between variables, and is depicted with an arrow from cause to effect. One must never be able to return to the same vertex by following causal chains in a Bayesian network. In other words, Bayesian networks are unable to accommodate cycles. It is said that the collection of vertices and edges in a Bayesian network forms a directed acyclic graph (DAG). In this terminology, "directed" results from the orientation of the arrows, "acyclic" from the prohibition on feedback loops, and "graph" from the mathematical term for a collection of vertices and edges.

The wording of family relationships is used in discussing edges, so if there is an edge from variable A to variable B , one describes A as a parent of B or B as a child of A . To each variable B with parents A_1, \dots, A_n there is attached a conditional probability table $P(B|A_1, \dots, A_n)$. Should B have no parents, the table reduces to an unconditional one, denoted simply $P(B)$.

Using only the conditional probabilities specified above, Bayesian networks allow the rapid computation of arbitrary probabilities on the variables in V . Thus, given any set of evidence E on any subset of V , one can rapidly determine $P(B|E)$ for any variable B in V . This calculation can be accomplished using the algorithm of Lauritzen and Spiegelhalter (1998). The networks derive their name in part from the repeated use of Bayes' theorem in this algorithm.

Bayesian model-averaging

Ignoring the uncertainty in model selection can lead to over-confident inferences and decisions (Hodges, 1987). Bayesian model averaging provides a way of accounting for this model uncertainty by averaging over competing models (Madigan and Raftery, 1994).

The conditional probability table $P(B|E)$ introduced earlier was defined with respect to a given Bayesian network (or model). For a finite number of plausible models M_1, \dots, M_k , Bayesian model averaging necessi-

tates the specification of the prior distribution over competing models. When there is little prior information about the relative plausibility of the models, the assumption that all models are equally likely *a priori* can be a reasonable neutral choice (cf. Kass and Wasserman, 1995). Given any set of evidence E , one can determine the conditional probability $\text{Prob}(M_k|E)$ for model M_k . Bayesian model averaging is then an average of the conditional distribution $P(B|E)$ under each M_k weighted by $\text{Prob}(M_k|E)$.

The problem

Consider the following hypothetical situation: a number of haddock fishermen have requested the right to cull a local seal population in order to increase the biomass of haddock available for fishing. The fishermen claim that, since seals compete with them for haddock within the local ecosystem, their catch rates would improve with fewer seals. Despite public opposition, fisheries managers take the proposal seriously. In fact, so effective are the fishermen in mustering political support, that an expert panel is convened with a mandate to consider whether or not culling seals would improve the local haddock fishery. In the event that the cull is deemed advantageous, the fishermen will be allowed to implement it.

Materials and methods

A group of individuals opposed to seal culling (referred to as "Group A" in this paper) has kept abreast of recent seal culling studies and is encouraged by them. They believe that the seals-eat-haddock story oversimplifies the trophic interactions in the local ecosystem (fishery). Furthermore, they suspect, along with Lavigne (1995), that complex ecosystem interactions could well influence the outcome. This certainly was the case in Yodzis (1998), where it was predicted that seal culling would have an adverse affect on the hake biomass in the Benguela ecosystem. Group A agrees to propose a complex ecosystem approach along the lines advocated by Yodzis (1998), in the hopes of obtaining a similar result and thus of averting the need for a seal cull.

For their part, the fishermen are looking for a method that is easier to understand, and they opt to press the case for seal culling using Bayesian networks (Jensen, 1996). The fishermen create a predictive model using Hugin software, based on the algorithm of Lauritzen and Spiegelhalter (1988), which allows them to present their approach in an interactive fashion.

Figure 1 shows a graphical representation of the simple model that the fishermen initially propose. It has four nodes that represent components of their simplified view of the ecosystem. The rightmost is a rectangular




Figure 1. This figure depicts the initial model proposed by the fishermen. The Cull node represents the decision on whether or not to selectively kill seals. This decision affects whether or not the seals are increasing which in turn affects whether the haddock are increasing (since seals eat haddock). Finally, the utility of haddock for the fishermen is improved by increases in the biomass of haddock.

Table 1. A conditional probability table for the Seals node given the Cull decision; i.e. $P(\text{Seals}|\text{Cull})$. The fishermen indicate that if they kill seals, the seal biomass will be decreasing with probability 1.0. On the other hand, if they do not kill the seals, the fishermen indicate ignorance about what will happen to the seals by assigning equal probability to the two states Increasing and Decreasing.

Cull	Cull seals	Don't
Decreasing	1.0	0.5
Increasing	0.0	0.5

decision node labelled Cull that represents the decision between two alternative actions – “cull seals” and “do not cull seals”. An arrow leads from the Cull decision to an elliptical node labelled Seals that represents the state of seal biomass in the local area. It has two states – either Increasing or Decreasing. The arrow from Cull to Seals indicates that the state of the seal stock; i.e. whether their biomass is decreasing or not, is directly affected by the decision about whether or not to cull seals. Another arrow leads from Seals to Haddock and the latter node has the same two states as the former – either Increasing or Decreasing. The arrow implies that the probability that haddock are increasing is affected by the state of Seals; that is, the arrows imply a causal relationship. Finally, the diamond shaped node labelled Utility represents the value fishermen place on changes in the biomass of haddock. The fishermen ascribe a utility level of 100 when the haddock are increasing and a utility level 0 when the haddock are decreasing. Although other participants may have different utility functions (cf. IWC, 1993), this one was consistent with the mandate of the expert panel; namely, to consider only the effects on the haddock fishery.

The effect of the decision to cull on Seals is specified with the conditional probabilities given in Table 1. The fishermen feel secure about defending Table 1, but they are unsure about the importance of the interaction between Seals and Haddock. The impact of a predator on a prey species depends, amongst other things, upon the biomass of the predator, the fraction of its diet that consists of the prey and on the size classes of prey consumed. A detailed trophic analysis that incorporates all these factors (cf. Punt and Butterworth, 1995; Yodzis, 1998) would require substantially more informa-

tion than is available to the fishermen. Thus the fishermen decide to base their probability table for haddock on seal dietary data, reasoning that the importance of the seal–haddock interaction should be reflected in the importance of haddock in the seals’ diet.

Dietary data are generally available for most species (Rice *et al.*, 1991; Anon., 1994), but in Table 2 hypothetical diet data are presented for a selection of species. In reality, such diet information would be based upon examination of predator gut contents and behavioural observations. The information in the Table 2 categorises the importance of each predatory relationship using terms ranging from Very High, through High, Medium, Low and down to Very Low. The interpretation of these categories is represented by relative weights that also accompany Table 2. For example, the importance of the interaction between seals and haddock is assessed as Low with a relative weight of 2.0. The importance of the interaction between seals and whiting is assessed as Very High with a relative weight of 5.0; i.e. the dietary data suggest that whiting are 2.5 times as important to the seals’ diet as haddock are.

In order to model the probability that a particular species is Decreasing given knowledge about what is happening to its predators, the fishermen create a predation pressure model. The primary assumption of the model is that predation pressure is additive; i.e. the model assumes that predators act independently of one another – they neither interfere with one another nor do they cooperate. The model consists of two parts – a rule for determining an index of predation pressure (x) on a species given the state of all its predators; namely, whether they are increasing/decreasing, and an equation for determining the probability that the species is decreasing given the predation pressure (x). The two components are combined to determine the probability of decrease, given the state of the predators. In this way, the model allows the generation of conditional probability tables for use in a Bayesian network.

Suppose that a particular prey species has K predators in a particular network. Let predator k (for $k=1, 2, \dots, K$) have relative importance weight x_k (as defined in Table 2). The state of the predator is either increasing or decreasing, and this is indicated by $s_k=1$ in the former case and $s_k=-1$ in the latter. The total predation pressure (x) on the prey species is determined by summing over the predators:

Table 2. Hypothetical dietary data for a number of the trophic interactions. The interpretation of the importance categories is expressed in terms of relative weight. For example, a predator-prey interaction of High importance is 4.0 times more significant than one of Very Low importance.

Predator	Prey	Importance
Cod	Whiting	Medium
Cod	Sandeels	Very Low
Seals	Haddock	Low
Seals	Cod	Medium
Seals	Whiting	Very High
Seals	Saithe	Very Low
Saithe	Cod	Medium
Saithe	Haddock	High
Saithe	Sandeels	Very High
Whiting	Haddock	Very High
Whiting	Sandeels	Very High
Sea birds	Sandeels	Very High

Importance	Very High	High	Medium	Low	Very Low
Relative weight	5.0	4.0	3.0	2.0	1.0

$$x = \sum_{k=1}^K x_k s_k \quad (1)$$

The equation for determining the probability (p) that the prey are Decreasing is inspired by logistic regression. As in logistic regression, p is determined from the natural logarithm of the odds ratio:

$$\text{logit}(p) = \log_e \{p/(1-p)\} \quad (2)$$

The model is as follows:

$$\text{logit}(p) = \lambda x \quad (3)$$

where $\lambda > 0$ is a parameter that is assumed to hold for all species. For $\lambda > 0$, greater predation pressure (x) implies a greater probability that the prey are decreasing in abundance. There is a biological interpretation for λ , namely $\lambda = \text{logit}(p_0)$, where p_0 is the probability that a prey species is Decreasing, if it has a single predator of Very Low importance and that predator is Increasing. A value of $p_0 = 0.54$ is assumed throughout this paper with the implication that $\lambda = 0.1603$, but the implications of this choice are discussed subsequently.

A number of properties of the predation pressure model are readily apparent. Firstly, given a prey species with two predators of equal importance, the cancellation property implies that when one predator is Increasing and the other is Decreasing, the resulting predation pressure is zero and the probability that the prey are Decreasing is 0.50. Secondly, using the relative importance weights given in Table 2, one can note that the predation pressure from a predator of High importance can be cancelled by two predators of Low importance, or by four of Very Low importance.

Table 3. A conditional probability table for the Haddock node of Figure 1 given the state of the Seals node. The values are based on the predation pressure model and on the fact that Low importance was assigned to the seal predation of haddock in the dietary data of Table 2.

Seals	Decreasing	Increasing
Decreasing	0.421	0.579
Increasing	0.579	0.421

Table 3 shows the conditional probabilities for the Haddock node of Figure 1, obtained using the predation pressure model and the seal dietary data. In the Table 3, a probability of 0.421 is assigned to haddock Decreasing given that seals are Decreasing. In order to compute this value, first note that the importance of seal predation on haddock is considered Low (see Table 2), so $x = (-1)/(2.0)$. Upon multiplying x by $\lambda = 0.1603$, one obtains $\text{logit}(p) = -0.3206$, which in turn implies $p = \exp(-0.3206)/(1 + \exp(-0.3206)) = 0.421$.

The optimal decision is taken to be the one that leads to the highest expected utility. This expectation is computed by multiplying the probability that haddock are increasing (conditional upon any data introduced at network nodes) by the utility derived from such an increase (previously ascribed a level of 100 by the fishermen). The fishermen have ascribed a utility level 0 when the haddock are decreasing so this term is eliminated from the calculation and may be ignored.

Results

The fishermen present their initial model and run using Hugin, giving an expected utility of 57.9 if seals are

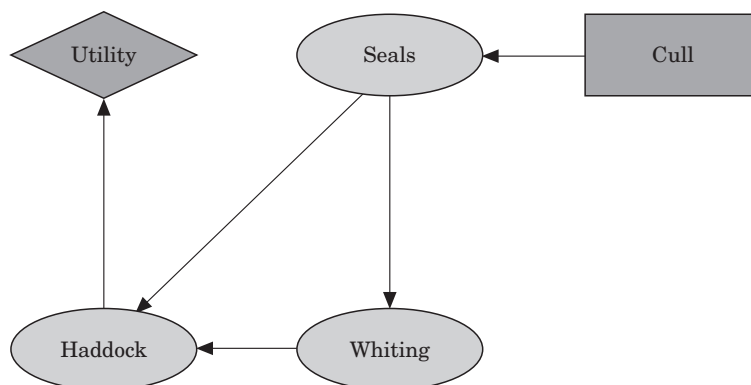


Figure 2. Seals eat whiting and whiting eat haddock.

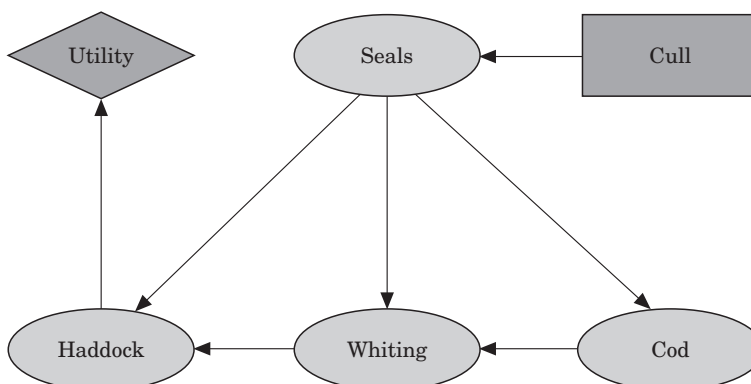


Figure 3. Cod are known to eat whiting and seals eat cod. The utility function remains unchanged and does not include cod because the expert panel's mandate is to consider impacts on haddock alone.

culled, and 50.0 if they are not. They demonstrate that their model is robust against uncertainty in p_0 in the sense that the optimal decision is not affected as long as p_0 is greater than 0.5. Indeed, the higher the value of p_0 in this model, the more attractive seal culling seems. Group A, having neither the time nor the data necessary to conduct a sophisticated analysis (cf. Yodzis, 1998), adopt the simple Bayesian network approach advocated by the fishermen. They raise the point about ecosystem effects, and claim to be able to use the fishermen's own method and model to demonstrate that culling seals would have an adverse effect on the expected utility.

Group A examines the seal dietary data in Table 2 and notes that seals also eat whiting. They construct the network shown in Figure 2 that includes seal predation on whiting, together with whiting predation on haddock. The Whiting node has the usual two states – Decreasing and Increasing. The network is parameterised using the predation pressure model exactly as specified by the fishermen, and the Seals node is entered from Table 1 as before. As anticipated by Group A, the expected utility for culling the seals is 49.732, whilst the expected utility

associated with doing nothing is 50.000. This result suggests culling is counterproductive, so the Group display it to the fishermen. In this model, the optimal decision is not affected by increasing p_0 , but lowering p_0 below 0.53 does lead to the conclusion that culling would increase haddock biomass. Note that, in contrast to the situation in the previous model, increasing p_0 makes the culling option seem less attractive.

The fishermen respond with the network shown in Figure 3, which includes cod as well as whiting. This network is parameterised using the predation pressure model and dietary data, as before. It favours the culling option, with expected utility 50.962, over the alternative of doing nothing, with expected utility of only 50.000. This particular network is sensitive to increases in p_0 , as raising this quantity to 0.55 or higher reverses the optimal decision. On the other hand, decreases in p_0 have no effect on the optimal decision, provided only that p_0 that does not fall below 0.5 (a possibility excluded *a priori*). Thus, in this model, higher values of p_0 are unfavourable to the culling option.

Group A are undeterred and respond with the network shown in Figure 4 in which the expected utility of




Figure 4. Saithe prey upon both cod and haddock; seals eat saithe.

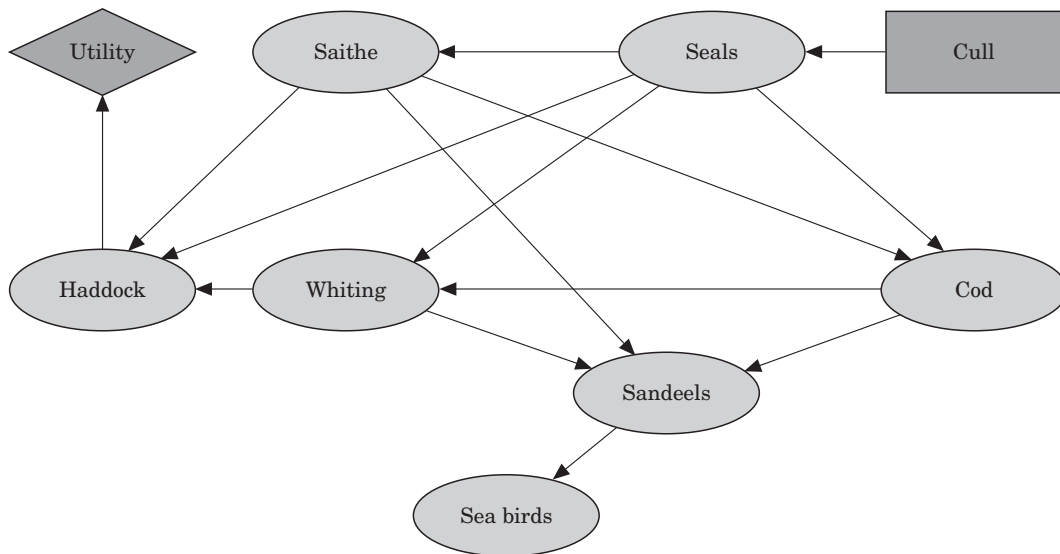


Figure 5. The Trojan sea bird model that includes a Sea birds node with two states – Gathering and Dispersing.

culling seals is 49.800, whilst the expected utility of restraint is 50.000. In this network, the optimal decision is not affected by raising p_0 , but decreasing it below 0.53 does change the best decision. As in the two previous models, lower values of p_0 are good for the fishermen's position, higher values for Group A's.

At this point in the development of ever more complex models, the fishermen have exhausted the choice of species that are eaten by seals. They seem defeated but construct the Trojan sea bird model shown in Figure 5. This model includes a Sandeels node with the usual two states – Decreasing and Increasing, and a node called Sea birds, with two new states – Gathering and Dispersing. Sea birds are known to gather over areas of high sandeel abundance, so the fishermen reason that sea birds will be more likely to gather when sandeels are Increasing. They adapt their predation pressure model to produce a table of probabilities that sea birds are Gathering given the state of the Sandeels. The resulting table is shown as Table 4.

Table 4. A conditional probability table for the Sea birds node of Figure 5 given the state of the Sandeels. The values are based on an adaptation of the predation pressure model and on the fact that Very High importance was assigned to the sea bird predation of sandeels (see Table 2).

Sandeels	Decreasing	Increasing
Gathering	0.31	0.69
Dispersing	0.69	0.31

The fishermen present their new model with every sign of defeat, saying they would agree to be bound by its recommendations if Group A would also. Group A examine the model and note that it says the expected utility of a seal cull remains exactly the same as the situation in their own model depicted by Figure 4. In fact, for all the nodes that are common to both, the models of Figures 4 and 5 lead to identical marginal

Table 5. A summary of the effects of structural changes in the operating model on the decision at hand and an investigation of the effects of modifying parameter p_0 .

Model	Optimal decision at $p_0=0.54$	Effect of increasing p_0	p_0 to change decision
Figure 1	Cull seals	Favours culling	Insensitive
Figure 2	Do not cull seals	Favours restraint	0.53
Figure 3	Cull seals	Favours restraint	0.55
Figure 4	Do not cull seals	Favours restraint	0.53
Trojan seabird model (with seabirds Dispersing)	Cull seals	Favours restraint	0.55




Figure 6. The model averaging network.

probabilities. In other words, the Trojan sea bird model seems to favour Group A's position. Group A agree to the model and congratulate themselves on their victory. They fail to notice, at first, that the fishermen have all moved outside on to the beach and are scanning the sea with binoculars for, on the horizon, a storm is coming.

Mystified, Group A follow the fishermen out on to the beach, and are asked if the sea birds are dispersing. "Of course the sea birds are dispersing", say Group A, "there's a storm coming!" A fisherman enters the evidence that the sea birds are Dispersing into a Hugin representation of the Trojan sea bird model. The expected utilities change: 46.281 for culling the seals and 46.274 for doing nothing. "You can't be serious", says Group A, "the sea birds are dispersing because of the storm. Besides, tomorrow the weather will probably clear up, and the birds will be back." The fishermen enter the evidence that the sea birds are Gathering into the Bayesian network. Once again, the expected utilities change: 53.735 for culling the seals but 53.726 for doing nothing. "You see", say the fishermen, "we should cull the seals either way!"

Though such considerations would probably be overlooked in a storm, we may examine the effect of changes in p_0 on the results from the Trojan sea bird model. Before the introduction of evidence, the model behaves exactly like the model of Figure 4. In other words, increasing values of p_0 favour Group A's position. After the introduction of evidence of any kind on the sea birds, increases in p_0 still make Group A's position more

attractive. Whether the sea birds gather or disperse, increasing p_0 to 0.55 or above will lead to a reversal of the optimal decision, but lowering it only makes culling look more advantageous.

Table 5 summarises the optimal decision in each model and the sensitivity of the models to changes in p_0 . In the Trojan sea bird model, it is the mere presence of evidence about the sea birds that changes the decision, and not the actual nature of that evidence. Jensen (1996) provides further explanation of this property of Bayesian networks.

Group A interprets the events on the beach as clear evidence of the inadequacy of the Trojan sea bird model. They say that (besides its temporal difficulties) it fails to incorporate other causes for sea bird dispersal (like storms), and they withdraw their acceptance of it. The fishermen are outraged, and relations between the two groups deteriorate. Support for models becomes polarised: the fishermen supporting the Trojan sea bird model (with the Sea birds Dispersing), and Group A backing the model in Figure 4.

In an attempt to break the impasse, model averaging is proposed, because it would allow evidence to shed light on which model is most credible. This is agreed to because there seems to be no other way forward. The model-averaging network is shown in Figure 6. It contains a node called Which Model? with two states – Fishermen and Group A. These states represent the fishermen's and Group A's models, respectively, and each state is assigned probability 0.5 to reflect total model uncertainty. Arrows link the Cull decision

Table 6. A conditional probability table for the Haddock node given the state of the Cull decision and given which model holds true. The values are marginal probabilities for Haddock states, computed in Hugin, for each decision and each contending model.

Cull Which model?	Cull seals Group A	Cull seals Fishermen	Don't Group A	Don't Fishermen
Decreasing	0.502	0.537	0.500	0.537
Increasing	0.498	0.463	0.500	0.463

directly to the Haddock since model-specific predictions about the state of the Haddock under each decision are entered on this node. The resulting probability table is shown in Table 6.

The model-averaging network shows an expected utility of 48.041 for culling seals, and 48.137 for restraint, so in the end it is decided not to implement a cull of the seals. Later that year, a stock assessment is conducted on the haddock and they appear to show a marked increase in biomass so the fishermen are appeased. When this information is incorporated into the model-averaging network, the probability in favour of the Group A model is increased from 0.500 to 0.519.

Discussion

Despite the emphasis within fisheries literature on parameter estimation, it is often structural (model) uncertainty that is the true source of controversy, and the choices of parameter estimation method are of secondary importance. This paper presents a case where model structure affects the optimal decision. Though the value of a model parameter (p_0) also plays a role, the decision favoured by increases in p_0 was shown to change with modifications in model structure. In view of such an example, fisheries science and management may wish to place greater emphasis in the future on approaches and techniques for dealing with structural model uncertainty.

It must be admitted that addressing model uncertainty can often seem particularly arduous: it will typically more than double the amount of work required to produce a result. The temptation not to consider it will be strong. Indeed, in this paper we have succumbed to it ourselves, having, for example, omitted consideration of link functions other than the logit in our predation pressure model. Consideration of model uncertainty must remain a matter of scientific discretion; however, certain questions will force such considerations on the scientist. Emotionally or politically charged topics, like culling seals, will demand consideration of structural uncertainty because opposing groups are likely to have created their own models. After all, the power of models

in shaping decisions has not gone unnoticed. Having devoted time to such creation, the groups are not likely to be pleased if their models go unused.

Model averaging should be explored further because it is useful both as a conceptual framework for addressing model uncertainty and as a means for achieving consensus amongst opposing groups. Though the initial probabilities attached to different models in such averaging may depend partly on political considerations, the example shown here illustrates that data can shed light on which model should receive highest weight. If consensus on a single model cannot be reached, model averaging may still be acceptable to all because no group needs to be excluded. After all, the more models have been put forward, the more people's work one must reject in choosing a single model on which to base decisions. Those one angers by rejection may well outnumber the people one pleases by acceptance! The shrewd decision-maker would do well to keep such political considerations in mind. On the other hand, completely caving in to interest group pressure is not advisable either. If one can reach a decision that is acceptable under all envisaged models, such an option is compelling, but experience suggests such options will be the exception rather than the rule. Furthermore, such decision making does not allow the data to decide which model is preferable, and thus provides no mechanism for learning from experience. Simply put, model averaging may provide a compelling compromise between political and scientific considerations.

Acknowledgements

A version of this paper was presented at the 1999 ICES ASC held in Stockholm under the Theme Session on Bayesian Approach to Fisheries Analysis. Thanks are due to the co-convenors of that session – Ray Conser and Maria Azevedo – and to André Punt, John Pope and Kevin Stokes for helpful comments on the manuscript. This paper was prepared with funding support provided by the Ministry of Agriculture, Fisheries and Food (contract MF0316). Hugin software is available from the World Wide Web at <http://www.hugin.dk>.

© 2001 Crown Copyright

References

- Anon. 1994. Report of the multispecies assessment working group. ICES CM 1994/Assess: 9, 177 pp.
- Buckland, S. T., Burnham, K. P., and Augustin, N. H. 1997. Model selection: an integral part of inference. *Biometrics*, 53: 603–618.
- Deriso, R. B., Quinn, T. J., and Neal, P. R. 1985. Catch age analysis with auxiliary information. *Canadian Journal of Fisheries and Aquatic Sciences*, 42: 815–824.

- Fournier, D., and Archibald, C. P. 1982. A general theory for analyzing catch at age date. *Canadian Journal of Fisheries and Aquatic Science*, 39: 1195–1207.
- Fournier, D. A., Hampton, J., and Sibert, J. R. 1998. MULTIFAN-CL: a length-based, age-structured model for fisheries stock assessment, with application to South Pacific albacore, *Thunnus alalunga*. *Canadian Journal of Fisheries and Aquatic Science*, 55: 2105–2116.
- Gudmundsson, G. 1994. Time series analysis of catch-at-age observations. *Applied Statistics*, 43: 117–126.
- Hodges, J. S. 1987. Uncertainty, policy analysis, and statistics. *Statistical Science*, 2: 259–291.
- IWC 1993. Report of the Scientific Committee. Reports of the International Whaling Commission, 43: 57–64.
- Jensen, F. V. 1996. An Introduction to Bayesian Networks. UCL Press, London. 178 pp.
- Kass, R. E., and Wasserman, L. 1995. A reference Bayesian test for nested hypotheses with large samples. *Journal of the American Statistical Association*, 90: 928–934.
- Kuikka, S., Hilden, M., Gislason, H., Hansson, S., Sparholt, H., and Varis, O. 1999. Modelling environmentally driven uncertainties in Baltic cod (*Gadhus morhua*) management by Bayesian influence diagrams. *Canadian Journal of Fisheries and Aquatic Science*, 56: 629–641.
- Lavigne, D. M. 1995. Interactions between marine mammals and their prey: unravelling the tangled web. In *Studies of high-latitude homeotherms in cold ocean systems*. Ed. by M. Montevecchi. Canadian Wildlife Series Occasional Paper (available via <http://www.cws-scf.ec.gc.ca/pub/list/pub/order.html>).
- Lauritzen, S. L., and Spiegelhalter, D. J. 1988. Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society, Series B*, 50: 157–224.
- Madigan, D., and Raftery, A. E. 1994. Model selection and accounting for model uncertainty in graphical models using Occam's window. *Journal of the American Statistical Association*, 89: 1535–1546.
- McAllister, M. K., and Ianelli, J. N. 1997. Bayesian stock assessment using catch-age data and the sampling-importance resampling algorithm. *Canadian Journal of Fisheries and Aquatic Science*, 54: 284–300.
- McAllister, M. K., and Kirkwood, G. P. 1998. Using Bayesian decision analysis to help achieve a precautionary approach for managing developing fisheries. *Canadian Journal of Fisheries and Aquatic Science*, 55: 2642–2661.
- Patterson, K. R. 1999. Evaluating uncertainty in harvest control law catches using Bayesian Markov chain Monte Carlo virtual population analysis with adaptive rejection sampling and including structural uncertainty. *Canadian Journal of Fisheries and Aquatic Science*, 56: 208–221.
- Punt, A. E., and Butterworth, D. S. 1995. The effects of future consumption by the Cape fur seal on catches and catch rates of the Cape hakes. 4. Modelling the biological interaction between Cape fur seals *Arctocephalus pusillus pusillus* and the Cape hakes *Merluccius capensis* and *Merluccius paradoxus*. *South African Journal of Marine Science*, 16: 255–285.
- Rice, J. C., Daan, N., Pope, J. G., and Gislason, H. G. 1991. The stability of estimators of suitabilities in MSVPA over four years of data from predator stomachs. *ICES Marine Science Symposium*, 193: 34–45.
- Richards, L. J., and Schnute, J. T. 1997. Model complexity and catch-age analysis. *Canadian Journal of Fisheries and Aquatic Science*, 55: 949–957.
- Yodzis, P. 1998. Local trophodynamics and the interaction of marine mammals and fisheries in the Benguela ecosystem. *Journal of Animal Ecology*, 67: 635–658.