



Parameterisation and evaluation of a Bayesian network for use in an ecological risk assessment

Carmel A. Pollino^{a,*}, Owen Woodberry^b, Ann Nicholson^b,
Kevin Korb^b, Barry T. Hart^a

^a Water Studies Centre, Monash University, 3800, Australia

^b Faculty of Information Technology, Monash University, 3800, Australia

Received 1 June 2005; received in revised form 7 March 2006; accepted 14 March 2006

Abstract

Catchment managers face considerable challenges in managing ecological assets. This task is made difficult by the variable and complex nature of ecological assets, and the considerable uncertainty involved in quantifying how various threats and hazards impact upon them. Bayesian approaches have the potential to address the modelling needs of environmental management. However, to date many Bayesian networks (Bn) developed for environmental management have been parameterised using knowledge elicitation only. Not only are these models highly qualitative, but the time and effort involved in elicitation of a complex Bn can often be overwhelming. Unfortunately in environmental applications, data alone are often too limited for parameterising a Bn. Consequently, there is growing interest in how to parameterise Bns using both data and elicited information. At present, there is little formal guidance on how to combine what can be learned from the data with what can be elicited. In a previous publication we proposed a detailed methodology for this process, focussing on parameterising and evaluating a Bn. In this paper, we further develop this methodology using a risk assessment case study, with the focus being on native fish communities in the Goulburn Catchment (Victoria, Australia).

© 2006 Elsevier Ltd. All rights reserved.

Keywords: Bayesian network; Ecological risk assessment; Ecology; Fish

1. Introduction

In environmental management, decisions are often based on either expert judgement or on complex quantitative models that consider only a small subset of environmental processes (e.g. sediment transport) within a complex system. These judgements and models typically focus on environmental processes, rather than ecological values, and there is little effort in understanding and quantifying uncertainties associated with complex systems. The risk assessment framework is an iterative process that seeks to address limitations in

environmental management by offering a formal and adaptive approach to decision-making (Hart et al., 2005). The framework aims to improve our understanding of how a system functions, and how decisions to manage a system affect ecological assets. Being an adaptive approach to environmental management, the process acknowledges that often uncertainties in our understanding of a complex system may be typically large at first, but with further data collection and analysis, these uncertainties can be reduced. Unfortunately, at present the modelling tools available for ecological risk assessments are limited.

Traditionally, models used in ecological risk assessments have tended to be restricted to single hazard assessments, with poor quantification of uncertainties, and poor capacity for fitting into an iterative and adaptive management approaches (Pollino and Hart, 2005). One tool that has shown potential in meeting the modelling needs of risk assessments

* Corresponding author. Water Studies Centre, PO Box 23, Monash University, Vic 3800, Australia. Tel.: +61 3 9905 4198; fax: +61 3 9905 4196.

E-mail address: carmel.pollino@sci.monash.edu.au (C.A. Pollino).

is Bayesian networks (Bns). Bns offer a pragmatic and scientifically credible approach to modelling complex ecological systems, where substantial uncertainties exist.

Increasingly, Bns are being used to model diverse problems of high complexity (Laskey and Mahoney, 2000; Korb and Nicholson, 2004), including environmental applications (Borsuk et al., 2004; Bromley et al., 2005; Castelletti and Soncini-Sessa, submitted for publication-b; Dorner et al., in press; Henriksen et al., in this issue; Little et al., 2004; Martin de Santa Olalla et al., in this issue; Ticehurst et al., in this issue; Varis and Fraboulet-Jussila, 2002). In many Bns, variables have been parameterised using either knowledge or data (Borsuk et al., 2004; Bromley et al., 2005; Rieman et al., 2001; Ticehurst et al., 2005), but rarely have both these information sources be combined in order to parameterise one variable. Elicitation of parameters in complex knowledge-based variables can be a difficult and time-consuming task, and often the knowledge of experts is incomplete (Cooke, 1991; Morgan and Henrion, 1990). Conversely, often significant data gaps exist for parameterising variables with data (Dorner et al., in press).

In the computer science and environmental modelling literature, there has been little guidance on how to approach combining knowledge elicitation and data for developing a Bn (Nicholson et al., 2001; Onisko et al., 2000). In computer science, an overarching methodology that can combine diverse information sources has been proposed, which is referred to as Knowledge Engineering of Bayesian Networks or KEBN (Korb and Nicholson, 2004). KEBN is an iterative or spiral approach to model prototype development, based on software development processes (Korb and Nicholson, 2004; Laskey and Mahoney, 2000). To date, the methodologies and associated support tools that accompany KEBN (which include parameterisation tools) have been poorly developed.

In an earlier paper (Woodberry et al., 2004b), we presented methods for the parameterisation and evaluation aspects of KEBN. In this paper, we aim to (a) formalise a process for combining different information sources to parameterise a Bn, (b) define how to focus parameterisation efforts on the more influential or 'sensitive' parts of the Bn, and (c) formalise a process for evaluating a Bn. These aims are achieved by applying KEBN to parameterise and evaluate a Bn designed to assist in the management of native fish in the Goulburn Catchment (Victoria, Australia).

2. Case study: fish communities in the Goulburn Catchment (Victoria, Australia)

The Goulburn River is the largest tributary of the Murray-Darling Basin in the State of Victoria (Australia). The lowland part of the river extends from Eildon to the confluence with the Murray River at Echuca (Fig. 1, Table 1). Many tributaries enter the 436 km lowland stretch of the Goulburn River. A detailed description of the catchment is given in Pollino et al. (2004).

A workshop (workshop 1) conducted in the region identified native fish communities as being at risk. The stakeholders involved in this phase of the project included catchment

managers and scientists from a range of fields. Native fish communities in the Goulburn River have declined over the past 100 years and are highly impacted by irrigation activities in the catchment (Pollino et al., 2004). The aim of the case study is to produce a quantitative model that can support future decisions for the management of the fishery.

3. Knowledge Engineering Bayesian Networks (KEBN)

A Bn is a graphical representation of a joint probability distribution over a set of statistical variables (Castelletti and Soncini-Sessa, submitted for publication-a; Korb and Nicholson, 2004; Pearl, 1988). The structure consists of a directed acyclic graph (DAG), made up of nodes that represent variables. Arrows between variables can represent direct causal dependencies based on process understanding, statistical, or other types of associations. A conditional probability table (CPT) is used to describe the probability of each value of the child node, conditioned on every possible combination of values of its parent nodes. These describe the strength of the causal relationships between variables. If a variable has no parents, it is described by a marginal probability distribution. The posterior probability distribution for a variable is calculated given new observations. Bns exploit the distributional simplifications of a network structure by calculating how probable events are, and how these probabilities change given subsequent observations or external interventions (Korb and Nicholson, 2004).

3.1. A quantitative knowledge engineering methodology

A knowledge engineering based method for parameterising and evaluating a Bn is shown in Fig. 2, and outlined in Woodberry et al. (2004b). Pathways through the cycle are indicated by arrows. These follow through different knowledge engineering processes (rectangular boxes), which can be executed either by humans (represented by the expert (DE) and knowledge engineer (KE), and indicated by clear boxes) or by a computer program (indicated by shaded boxes). Major choice points are indicated by hexagons.

The initial stage in the development spiral is *Structural Development and Evaluation*, which on the first iteration will produce an unparameterised causal network. This phase of model development can be undertaken via a knowledge or data-based approach. Approaches for undertaking this step are not described in this paper, but have been described elsewhere (Nicholson et al., 2001).

In Fig. 2, the next stage is *parameter estimation*, which involves specifying the CPTs for each node. The parameter estimates can be elicited from experts, and/or obtained from expert literature (path 1). Parameter estimates can also be learned (via automated learning) from data (path 2). In this study, we propose a method to combine expert and data sources (shown in path 3). As Bn development is an iterative process, the parameter estimates in early prototypes need not be exact, and uniform distributions (which represent unbiased

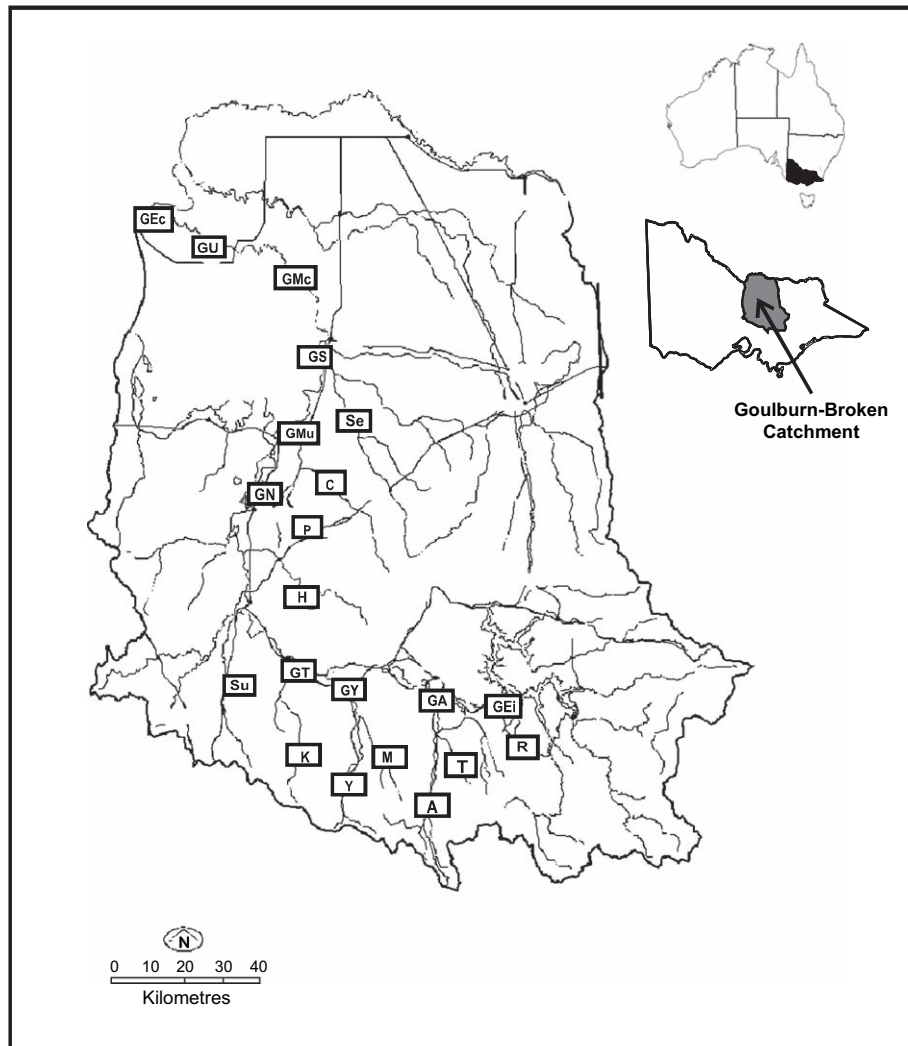


Fig. 1. Map showing location of Goulburn Catchment (see Table 1 for definition of site symbols).

parameter estimates) can be used if knowledge or data are not readily available.

The second major aspect in the KEBN process, as shown in Fig. 2, is *quantitative evaluation*. Evaluative feedback can be generated using either experts or data or both, as we have done here. When data are available, several measures can be used to evaluate the Bn, including predictive accuracy, expected value computations and information reward (Korb and Nicholson, 2004). Expert evaluation techniques include elicitation reviews and model walkthroughs.

Another kind of evaluation is sensitivity analysis. This involves analysing how sensitive the network is by determining how responsive the probabilities of query nodes are to changes in parameters and inputs. Measures for these can be computed using the tools, sensitivity to parameters and sensitivity to findings (Fig. 2). The output of these tools requires evaluation by an expert.

A comprehensive description of this KE process is provided in Woodberry et al. (2004b). In the following sections, we describe the application of the process to the case study.

4. Bn structural development and evaluation

To assess the impacts of human-related activities on native fish communities, it was important to begin constructing the Bn by identifying study endpoints, and establishing linkages between important system variables and endpoints. To do this, experts were consulted in a workshop (workshop 2). Attendees included fish ecologists and local catchment managers.

The assessment endpoint of the case study was the management need to assess the conditions required to establish sustainable native fish communities. The model endpoints (also referred to as query variables) selected for this study were native fish abundance and diversity. Experts also collaborated to develop a conceptual model of the relationships between physical, chemical and biological factors, and the study endpoints. The causal structure had to be flexible, being able to represent multiple spatial scales (21 sites/six regions in the catchment) and two temporal scales (1 year and 5 years).

This conceptual model formed the basis of the Bn structure, shown in Fig. 3. For the purposes of this paper, variable names

Table 1
Site information and location of lowland river sites in the Goulburn Catchment

Region	Stream	Monitored?	Site code
Upper main	Goulburn River	Yes	GEi
	Goulburn River	No	GA
	Goulburn River	No	GY
	Goulburn River	Yes	GT
Upper tributary	Rubicon River	Yes	R
	Taggerty River	No	T
	Acheron River	Yes	A
	Murrindindi River	Yes	M
	Yea River	Yes	Y
	King Parrot Creek	Yes	K
	Sunday Creek	Yes	Su
	Hughes Creek	Yes	H
Mid main	Lake Nagambie	No	GN
Lower main	Goulburn River	Yes	GMu
	Goulburn River	Yes	GS
	Goulburn River	Yes	GMc
	Goulburn River	No	GU
	Goulburn River	No	GEc
Lower tributary	Pranjip Creek	Yes	P
	Castle Creek	Yes	C
	Seven's Creek	Yes	Se

in Fig. 3 are simplified. The Bn consists of five interacting sub-models – water quality, hydraulic habitat, structural habitat, biological potential and species diversity, and two query variables – Future Abundance and Future Diversity. The sub-models are grouped and labelled in Fig. 3. Note that nodes consisted of those parameterised using data only, a combination of data and elicited information, and where no data were available (summary intermediate nodes), elicited information alone.

The development of the model structure is not described in any more detail in this paper, other than to acknowledge that it was an iterative process. Details can be found in Pollino (2004).

5. Parameter estimation

To obtain CPTs for a node, the model needs to be parameterised. It was clear from the outset of this study that the number of parameters in our Bn (Fig. 3) outweighed the data available to us. For this reason, expert elicitation of parameters was to be used in conjunction with data.

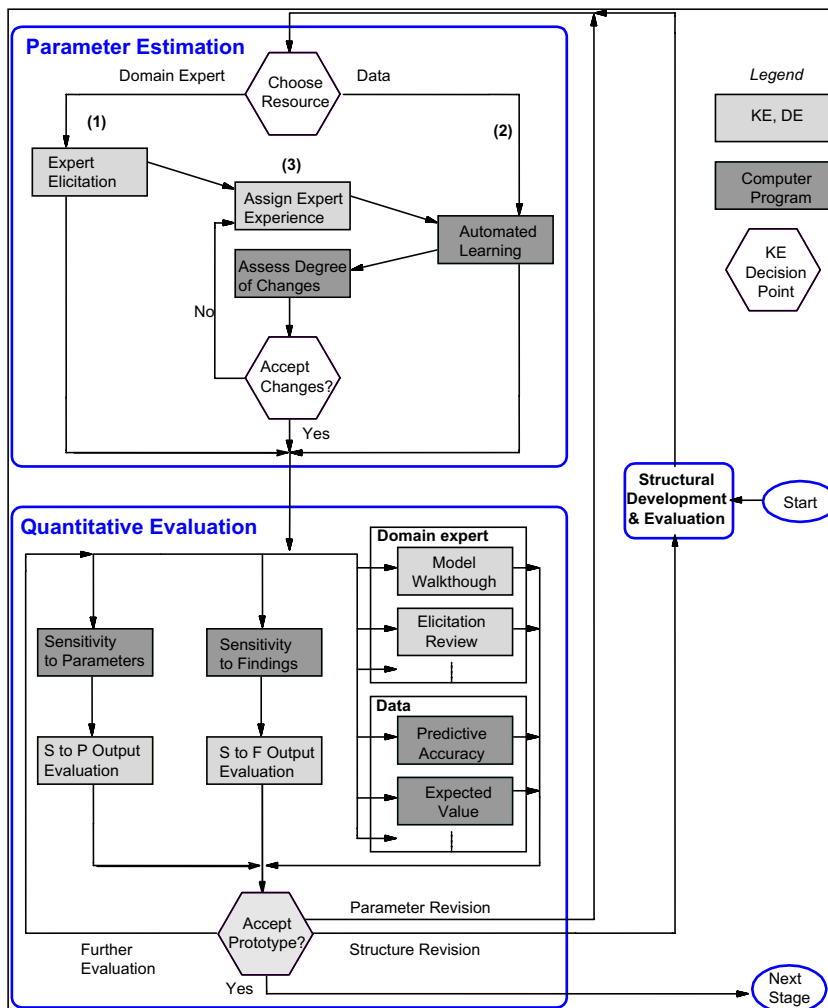


Fig. 2. Proposed methodology for Knowledge Engineering of Bayesian Networks (KEBN).

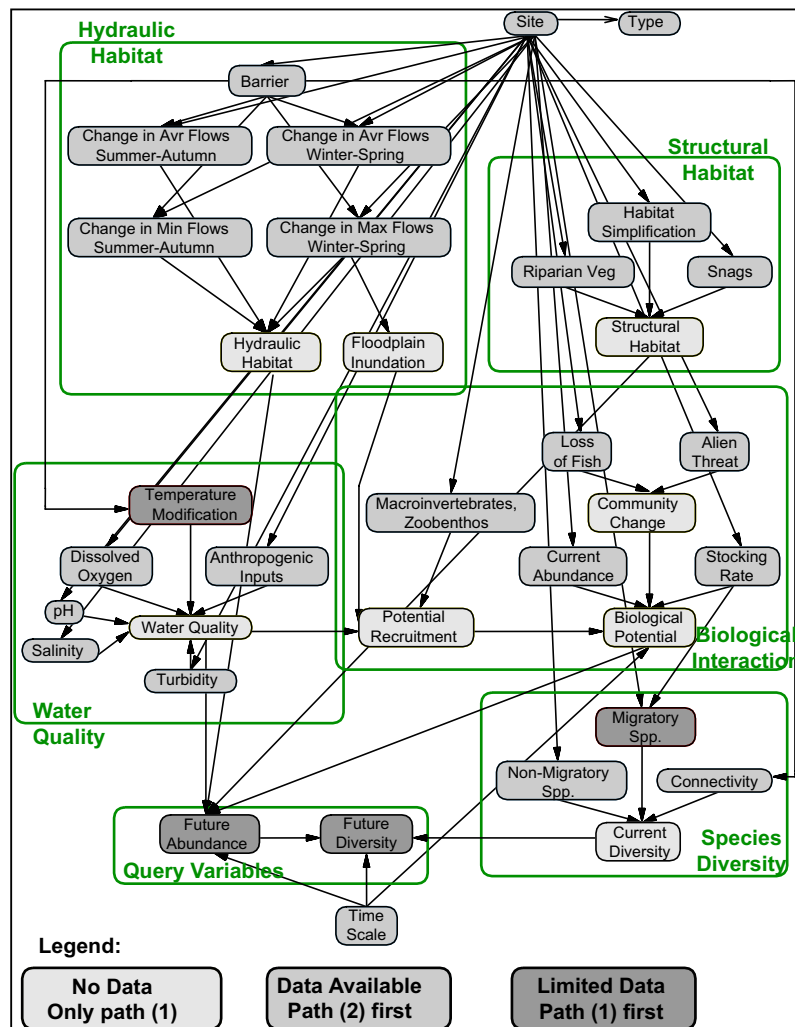


Fig. 3. Native fish Bn structure for the Goulburn Catchment (Victoria, Australia).

Prior to parameterisation, all variables were discretised into states. For continuous variables, states were further discretised into sub-ranges. Where possible, states were established using recognised classifications, management thresholds or guidelines (see Table 2 for details). Where these were not available, sub-ranges were specified with the guidance of experts (workshop 3). The number of ‘states’ or ‘classes’ assigned to each variable were not pre-determined, but evaluated and assigned on an individual basis. While discretisation of continuous variables is not desirable, it facilitates parameterisation process by simplifying expert elicitation, and it acknowledges that our understanding of many parameters, and the data available to support such relationships, is often quite rudimentary.

After discretisation, the variables were split into two groups: those with data, and the remainder that were elicited from experts and expert literature. In Fig. 3, data input variables are specified, although some sites have data missing for select water quality and structural habitat variables (missing data were predominately due to an absence of monitoring data). Variables were not elicited where this occurred. Where no data or limited data were available for remaining variables

(Fig. 3), probability distributions were elicited from experts (workshop 3 and individually).

5.1. Expert elicitation

The parameter elicitation process involved fish ecologists and local catchment managers. Elicitation of distributions employed such questions as “What is the probability that variable A takes state X given information Y ?”. Answers were in the form of probabilities. The experts were also asked to report their confidence (low or high) in these estimates. Confidence applied to the whole CPT, rather than individual parameters (although this does not need to be the case). In this study, guidance for elicitation was sought from Morgan and Henrion (1990). Other elicitation methods, such as the use of reference lotteries, gambling analogies and scoring rules could also be used (Cooke, 1991; Savage, 1971).

When eliciting precise parameters, it is also useful to elicit an acceptable range for the parameter, as these intervals can be used later to identify other parameters needing further attention. To communicate this concept to experts, they were asked

Table 2
Methodology used to discretise data input variables, and the states of these variables

Variable	Discretisation methodology	States
Barrier	Based on classification in ISC	None Complete deep Complete shallow Inundated
Anthropogenic inputs	Expert knowledge	Low Medium High
Dissolved oxygen	Expert knowledge	Extreme low (0–40%) Normal (40–110%) Extreme high (110–200%)
Salinity	Clunie et al. (2002)	Low (0–1000 mg/L) Medium (1000–5600 mg/L) High (5600–10 000 mg/L)
Turbidity	Expert knowledge	Low (0–100 NTU) Medium (100–1000 NTU) High (1000–10 000 NTU)
Temperature modification (°C) from natural	Modelling of pre-regulation temperatures, and relating to temperatures required for spawning (Koehn and O'Connor, 1990)	No change (0–2 °C) Moderate (2–4 °C) Major (4–10 °C)
Seasonal flows – percentage change from natural	Collection of pre-dam and post-dam data, calculating %change in means. Classification in ISC. Data accessed from the Victorian Water Data Warehouse (http://www.vicwaterdata.net/ ; accessed January 2006).	Extreme decrease (0–25%) Decrease (25–75%) No change (75–125%) Increase (125–175%) Extreme increase (175–1000%)
Floodplain inundation	Gippel and Finlayson (1993)	Yes No
Macroinvertebrates	AUSRIVAS score in ISC >0.80: 4 0.79–0.6: 3 0.59–0.4: 2 0.39–0.2: 1 <0.2: 0	Low (0–1) Medium (1–3) High (3–4)
Loss of fish	Expert knowledge	Low High
Stocking rate	DPI	None Low High
Snags Habitat simplification Native riparian vegetation	Classification in ISC	Low/complete/degraded Medium/some/moderate High/none/intact
Current & Future Abundance	Percentiles of populations in the Goulburn Catchment Percentile 60%: native fish strategy – native spp. composed of 60% or more of population	<60%, low (0–54) >60%, high (54–500)
Alien threat	Percentiles of populations in the Goulburn Catchment Percentile 40%: native fish strategy	<40%, low (0–15) >40%, high (15–500)
No. of migratory spp.	Percentiles of populations in the Goulburn Catchment	<50%, low (0–3) >50%, high (3–15)

to consider intervals as 95% credible intervals (often incorrectly referred to as confidence limits). It was observed that experts tended to be more confident in estimating physical and chemical variables, rather than biological relationships. This is likely to be due to the uncertainties associated with our poor knowledge and the inherent variability of biological variables.

5.2. Data-based estimation

Fisheries data were obtained from locally held databases of the Department of Sustainability and Environment and Department of Primary Industries. The Lake Nagambie Angling Club supplied additional data for Lake Nagambie (McGuckin, 2002).

Unfortunately, there are no fishery records for the pre-regulation period, and no data were used for the pre-1970 period, as sampling was sparse and poorly documented. Water quality, flow and stream condition data were accessed from the Victorian Water Data Warehouse (<http://www.vicwaterdata.net/>; accessed January 2006). Further description of the data can be found in Pollino et al. (2004).

Data files were arranged into a series of cases (cases are a set of observations based on the date of a fish survey) for data learning (i.e. calibrating/training the Bn). The number of cases per site ranged between 3 and 272, with a total of 949 cases. The more case data available for a site, the greater the certainty in specifying CPTs. Cases consisted of fisheries data matched to available water quality, flow and structural habitat data and information. At routinely monitored sites (Table 1), time-based monitoring data (pH, salinity, turbidity and dissolved oxygen) were matched to the long-term fisheries monitoring data where possible. Hydraulic parameters were calculated using historical changes in flow. The remaining physical data incorporated into variables were related to the Index for Stream Condition (ISC) survey conducted in 1999 (<http://www.vicwaterdata.net/>; accessed January 2006), and to the authors' knowledge, no other analogous surveys had been conducted at the time of this study.

For parameter estimation using data, the software used was Netica (Norsys, 2005). Netica has three automated algorithms: the Lauritzen Spiegelhalter method (LS) (Lauritzen and Spiegelhalter, 1990); the expectation maximization algorithm (EM) (Dempster et al., 1977); and the gradient descent algorithm (GD) (Norsys, 2005). The simplest method is LS, which uses frequency counts of child states given each possible parent instantiation. Problems arise when the data lack coverage of many states. The LS method was not used, as it is unable to update parameters if there was missing data. EM deals with missing data by finding the parameterisations that yield the greatest likelihoods given the available data. Consequently, the EM algorithm was selected for this study. Preliminary Bn parameterisation trials with the GD method showed it was susceptible to local maxima (Woodberry et al., 2004a). This has also been observed when training neural networks using GD (Gori and Tesi, 1992).

Automated learning trials were then carried out using the EM algorithm to investigate the effects of different weightings of elicited CPTs, and to combine expert and quantitative estimations of CPTs. Where there was no physical and chemical monitoring data for a site, probability distributions were assumed to be uniform, representing an unbiased parameter estimate.

5.3. Combining expert and data-based estimations

For combining expert elicited and data-based parameterisations, the expert elicited information was weighted. As described above, during the expert elicitation process, a weight was assigned to the parameter estimates, based on the expert confidence in the elicited CPT estimate. Subsequently, these weightings (W_i) were treated as being equivalent to the size

of a hypothetical data sample (i.e. an equivalent sample size). This process is shown in Fig. 2 as Assign Expert Experience.

Two sets of trials, each with different initial weighting, were conducted 'blind' with an expert (a fisheries ecologist). Where there was high confidence in an elicited CPT, a higher experience weighting was assigned (e.g. weighting of 10 or 20), and only minimal changes in the CPT updating process would be observed. If the confidence in a parameter was low, a lower experience weighting was assigned (e.g. weighting of 10 or 5). Table 3 shows how the weightings were used in two Bn parameterisation trials, using the EM method.

After incorporating the data into the parameter estimate, the next step was to compare the new CPT to the original (Assess Degree of Changes decision point in Fig. 2). In this study, we used Bhattacharyya distance to assist in the assessment process. This method compares the distance between two probability distributions (Bhattacharyya, 1943). Where the new parameter estimates were outside the acceptable parameter intervals obtained during expert elicitation, these were flagged for attention. The expert then reviewed the model to determine if the model was behaving as expected. Where parameter estimates required further adjustment, a new weighting value (W_{i+1}) was assigned and the process repeated (see Table 3). If the change to the CPT was unacceptable, the expert weighting was increased, and vice versa.

The Adjusted Weightings assessment process was iterated using the algorithm introduced in Woodberry et al. (2004b), and shown below.

ALGORITHM: Adjusting Weightings
Loop until weighting values converge
Parameterize network with current weights

Table 3

To update CPTs of expert elicited nodes with data, two sets of trials were conducted, using the EM algorithm

Node	$H = 10, M = 5$			$H = 20, M = 10$			Combined trial no.	
	Trial no.			Trial no.				
	1	2	3	1	2	3	4	5
Water quality	10	15	18	20	25	25	25	24
Hydraulic habitat	10	15	18	20	25	25	25	24
Structural habitat	10	7	4	20	15	10	1	1
Biological potential	5	2	4	10	5	5	5	5
Temperature modification	10	5	1	20	15	10	1	1
Community change	5	1	1	10	5	1	1	1
Floodplain inundation	10	5	1	20	15	10	1	1
Potential recruitment	5	1	3	10	5	2	3	3
Connectivity	10	10	10	20	17	14	12	12
Migratory spp	5	10	15	10	15	15	15	16
Current diversity	5	1	1	10	7	4	1	2
Future Abundance	5	2	4	10	7	4	5	6
Future Diversity	5	5	5	10	5	5	5	5
Remaining nodes	0	0	0	0	0	0	0	0

Each set had different initial experience weightings, which applied to the weightings on case files being used for model calibration/training. Sets of trials were combined as experience weightings converged. H = high, M = moderate.

Switch

Case changes unrealistic: $W_{i+1} \leftarrow W_i + upLarge$

Case would allow greater changes: $W_{i+1} \leftarrow W_i + upSmall$

Case little OR no change: $W_{i+1} \leftarrow W_i + downLarge$

Case insignificant change: $W_{i+1} \leftarrow W_i - downSmall$

Case changes become unrealistic: $W_{i+1} \leftarrow W_{i-1} + bounceup$

Case changes disappear: $W_{i+1} \leftarrow W_{i-1} - bouncedown$

Case final trials, small adjustments needed: $W_{i+1} \leftarrow W_{i-1} \pm tweak$

Case: changes acceptable: $W_{i+1} \leftarrow W_i$

End Loop

In the native fish cash study, an *upLarge* or *downLarge* change was equivalent to a weighting of five, *upSmall/downSmall* was equivalent to a weighting of three, *bounceup/bouncedown* was equivalent to a weighting of two, and a *tweak* was equivalent to a weighting of one.

The process described seeks to identify inconsistencies between expert elicited and data-derived parameter estimates. If the weightings for the two trials converged (Table 3), we hoped that this process was able to confirm the consistency and reliability of the expert. In this study, convergence occurred over only five iterations. We do not consider this process as being too arduous, but we do recommend that this process be undertaken with multiple experts where possible.

6. Quantitative evaluation

After parameterisation of the Bn, the second major aspect of the knowledge engineering cycle is evaluation (Fig. 2). These results can be used to further guide Bn development.

6.1. Evaluation using data

Where possible, data should be used for evaluation. A common method of evaluation for a Bn is to measure predictive accuracy. This method measures the frequency with which the predicted node state (that with the highest probability) is observed, relative to the actual value. If data are also being used to parameterise CPTs, it is necessary to divide data into a calibration/training set and a test set.

In the case study, available data were split randomly so that 80% of data were used for model training (calibration) and 20% used for testing. Model error rates are generated by withholding data of selected variables, and using the model to predict the outcome of the variable of interest. The error rates of the query nodes, *Future Abundance* and *Future Diversity*, were only 6 and 0%, respectively. The low error rates reflect the lack of variability in the data set.

Another test is to assess whether trends in model prediction are consistent with past field observations. For the fish Bn, trends between predictions and data were maintained (Fig. 6).

6.2. Evaluation using experts

Bn evaluation with experts is also important. This can be done via a structured review of the model. For the native fish Bn, a semi-formal model walkthrough was conducted with experts. Overall, the model received positive feedback. The model was regarded as being a reasonable representation of a complex environment, but the need for routine updating was emphasised. The authors also recognise the need for more data to improve the robustness of model predictions, and to reduce parameter uncertainties. A model deficiency was the inability to differentiate between the responses of different native fish groups, particularly as different groups can respond differently to environmental conditions, such as flow.

6.3. Sensitivity analysis

Sensitivity analysis is used to measure the sensitivity of changes in probabilities of query nodes when parameters and inputs are changed. The query nodes in this study were model endpoints. Two types of sensitivity analyses were used in evaluating the Bn. The first, “sensitivity to findings”, considers how the Bn’s posterior distributions change under different conditions, while the second, “sensitivity to parameters”, considers how the Bn’s posterior distributions change when parameters are altered. To date researchers appear to have employed only one or the other of these methods in any one study (e.g. Coupe and van der Gaag, 2002; Laskey and Mahoney, 2000; Rieman et al., 2001). Both are needed for a careful and thorough investigation of the properties of a network.

In the native fish Bn, sensitivity analysis was used to identify variables that are highly sensitive to change, so that quantification efforts in subsequent model iterations are well focused. Where the accuracy of parameters cannot be improved, they may represent knowledge gaps. Less effort is required in quantifying variables that were identified as contributing little to improving the predictive accuracy of the model.

6.3.1. Sensitivity to findings

Sensitivity to findings can use the properties of d-separation to determine whether evidence about one variable may influence belief in a query variable (Korb and Nicholson, 2004). The d-separation occurs when nodes in a causal graph are conditionally independent, given evidence (for more information see Korb and Nicholson, 2004). Using sensitivity to findings, it is possible to rank evidence nodes. This process allows the expert to identify whether a variable is sensitive or insensitive to other variables in particular contexts, which in turn may help to identify errors in either the network structure or the CPTs. The information can also be used to provide guidance for collecting further data or to direct expert elicitation and evaluation efforts.

Sensitivity to findings can be quantified using two types of measures, entropy and mutual information. Both measures were implemented using algorithms in Woodberry et al. (2004a).

Entropy, H , is commonly used to evaluate the uncertainty or randomness of a variable (X) characterised by a probability distribution, $P(x)$ (Korb and Nicholson, 2004; Pearl, 1998):

$$H(X) = - \sum_{x \in X} P(x) \log P(x) \quad (1)$$

Entropy measures assess the average information required in addition to the current knowledge to specify a particular alternative (Das, 2000).

Mutual information is used to measure the effect of one variable (X) on another (Y) (Korb and Nicholson, 2004):

$$I(X, Y) = H(X) - H(X|Y) \quad (2)$$

where $I(X, Y)$, is the mutual information between variables. This measure reports the expected degree to which the joint probability of X and Y diverges from what it would be if X were independent of Y (Korb and Nicholson, 2004). If $I(X, Y)$ is equal to zero, X and Y are mutually independent (Pearl, 1988).

Table 4
Sensitivity analysis for posterior network (Eildon), showing calculated entropy

Entropy of Future Abundance	0.386506
Future Diversity	0.046391
Water quality habitat descriptor	0.01806
Hydraulic habitat descriptor	0.00818
Natives biological potential descriptor	0.00073
Temperature modification	0.00044
Barrier	0.00043
Change in avr flows summer–autumn	0.000231
Change in min flows summer–autumn	0.000223
Change in max flows winter–spring	0.000202
Change in avr flows winter–spring	0.000198
Diverse structural habitat descriptor	0.000127
Floodplain inundation	0.000025
Connectivity	0.000025
Potential recruitment	0.000017
Community change	0.000009
Current abundance	0.000005
pH	0.000002
Turbidity	0.000001
Native riparian veg.	0.000001
Salinity	0.000001
Snags	0.000001
Anthropogenic inputs	0.000001
Dissolved oxygen	0.000001
Habitat simplification, aquatic veg.	0
Current diversity	0
Site	0
Alien threat	0
Loss of fish	0
Time-scale	0
Migratory spp.	0
Macroinvertebrates, zoobenthos	0
Type	0
Stocking rate	0
Non-migratory spp.	0

The Future Abundance variable in italics is the variable of interest, acting as a reference for the remaining nodes.

A typical output of sensitivity analysis, using entropy measures, is shown in Table 4. In this example, the site of interest is Eildon, the uppermost site within the catchment where native fish communities are highly stressed (Pollino et al., 2004). Results from this site indicate that the *hydraulic habitat* and *water quality habitat* are the variables having the greatest influence on the *Future Abundance* variable. Given that *Future Diversity* is a function of *Future Abundance*, this has the greatest influence on findings of the parent node. Results for mutual information can be found elsewhere (Woodberry et al., 2004a), where again the *hydraulic habitat* and *water quality habitat* variables are having the greatest influence on the findings of the *Future Abundance* variable.

Sensitivity to findings can also be graphically represented. The probability of each parent node can be altered over the probability space, and changes in the endpoint node observed. For example, Fig. 4 shows the model sensitivity for all sites in the catchment to the *Future Abundance* node, over a 1-year time frame. Once again, clearly *Future Diversity* (a child node of *Future Abundance*) is the most influential node, and the variables for *hydraulic habitat* and *water quality habitat* have the greatest influence on the *Future Abundance* variable at the site, Eildon.

Outcomes of sensitivity to findings concur with field observations. At Eildon the seasonal flow regime has been reversed as a result of irrigation demand, and water quality is altered as water temperatures are as much as seven degrees below modelled ‘no dam’ temperatures, due to bottom water releases from Eildon Dam (Pollino et al., 2004).

6.3.2. Sensitivity to parameters

Sensitivity analysis can also be performed using an empirical approach, in which each of the parameters of the query nodes are altered and the related changes in the posterior probabilities of the query node (such as the endpoint) are observed. To examine a complex Bn, this type of analysis can be extremely time consuming. Coupe and van der Gaag (2002) seek to address this limitation by identifying a ‘‘sensitivity set’’ of variables, which are defined as being the most influential in a Bn. This is done by calculating the posterior probability of a node by systematic changing conditional probabilities. It is these parameters that are most influential in calculating posterior probabilities, and it is on these parameters that quantification efforts should be focussed (Coupe et al., 1999). If the plotted sensitivity function does not behave as the expert expects (e.g. its slope, direction or range is unexpected), this may indicate errors in the network structure or CPTs.

A sensitivity set of nodes can be found using an adapted d-separation algorithm (see Eq. (3)). When evidence is entered into a Bn (i.e. a Bn is instantiated) our algorithm identifies the type of function of the parameters by checking whether the query node has any child nodes. Parameter changes are represented as linear if there are no child nodes, or hyperbolic if there are child nodes.

A revised probability distribution of the test node is set by first selecting a new value, P_{new} for the parameter under investigation, P_j . The remaining parameters, P_i , are normalized to retain relative values by the updating function:




Fig. 4. Sensitivity of the mean probability for high future abundance, using a 1-year time scale. Nodes are listed from the least influential (top) to the most influential. Bars represent the range in variation observed 'Future Abundance' node when the values of the nodes on the y-axis are varied from 0 to 1.

$$P_i \leftarrow P_i \frac{1 - P_{\text{new}}}{1 - P_j}, \quad i \neq j \quad (3)$$

before the parameter under study is updated.

$$P_j \leftarrow P_{\text{new}} \quad (4)$$

In the native fish Bn, conditional probabilities were systematically varied over the entire probability space (0–1), and the effects on model endpoints were examined. The sensitivity patterns varied for each site and for each time frame. Generally, these changes were small, causing only minor changes in model predictions. The node state (state with the highest probability) was unchanged for all scenarios tested.

The most sensitive scenario for the model was identified as: *probability that future abundance is low given; water quality was low, structural habitat was low, biological potential was low, and hydraulic habitat was low over a one year time scale.* According to experts, this scenario represents the “worst case scenario” for native fish. The change in the posterior probabilities for Future Abundance for the site Eildon, given alterations of CPTs over 0–1 in the above scenario, is shown in Fig. 5. The alterations in CPTs of the scenario did not change the predicted node state, with the probability of Future

Abundance at Eildon being predicted to remain at low (i.e. $P(\text{Future Abundance}) = \text{Low}$ regardless of change in CPTs for scenario tested above).

6.4. Model predictions

Fisheries data from each site in the catchment were plotted against model predictions (Fig. 6). Relative comparisons show that the trends between the fisheries data set and the predictive outputs are generally consistent.

Using sensitivity analyses (see Table 4 and Fig. 5 as an example for one site only), water quality (temperature) and changed hydrology were identified as the variables primarily influencing fish abundance. At the remaining sites, fish abundance is primarily influenced by the biological potential (potential recruitment and current abundance), water quality (turbidity, dissolved oxygen and pH), and flow/hydraulic habitat.

6.5. Accept model prototype

After evaluation, the Bn can be accepted for the next stage of development. This decision is not intended to be the end of

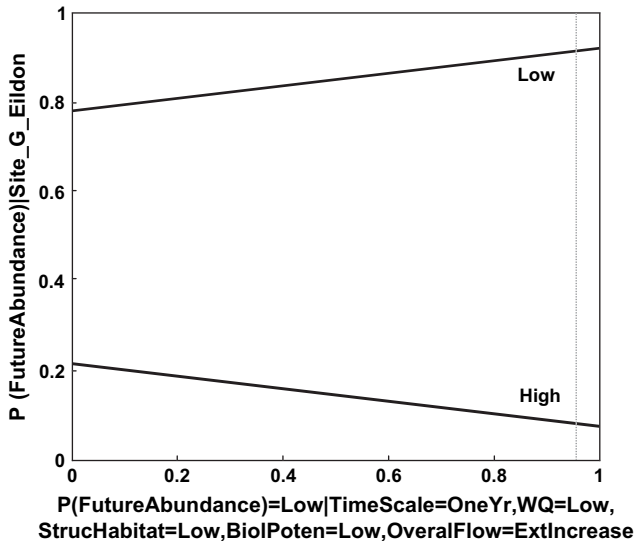


Fig. 5. Sensitivity to parameters output showing slope of change for high and low future native fish abundances at site, Eildon. Observed posterior probabilities for low and high future abundance at Eildon are shown on the y-axis, and the conditional probabilities being altered for the specified scenario are on the x-axis.

the knowledge engineering or prototype development process, rather the KEBN and ERA processes call for continued development of the model.

7. Proposed use of the Bn

Using the Bn developed, risks to native fish communities in the Goulburn Catchment can be prioritised. Important risks to fish can represent key management priorities or important




Fig. 6. Relative abundance data (left axis — bars) versus Bn model predictions (right axis — line) for sites in the Goulburn main channel. See Table 1 for a definition of sites names on the x-axis.

knowledge gaps. The predicted outcomes of alternative management scenarios in the catchment, at specific sites or reaches, can also be tested. An example of how the Bn can be used is described below.

Using Eildon as an example, Table 4 shows that currently the major risks to fish are changes to water quality (due to temperature changes) and seasonal flows/hydraulic habitat. The probability of having a high native fish abundance (defined as ≥ 60 th percentile of entire fish community of that site, see Table 2) is only 0.07. If management actions were targeting at restoring natural water temperatures and flows at that site, predictions for high native fish abundance increase to 0.31. If native fish communities were to become established at the site, improving biological potential, predictions for high native fish abundance increase to 0.45. Taking into consideration these management actions, further sensitivity analyses indicate that future management actions should target rehabilitating the structural habitat in the area and targeting alien (introduced) species.

The next stage of this work is to further explore the use of fish Bn in multi-criteria decision support, particularly in the analysis of tradeoffs in environmental management (Bromley et al., 2005; Dorner et al., in press; Varis, 1997).

8. Conclusions

8.1. Knowledge engineering

In this paper we have described the use of a methodology for combining expert elicitation and data for parameterisation of Bns, an important research topic that has been widely acknowledged in the Bn field but little developed. In many ecological applications, including our case study, information sources are often poorly documented, poorly understood, and generally incomplete. Although other causal network structures (Borsuk et al., 2004; Ticehurst et al., 2005) have been developed using such information sources, unlike this study, parameter estimates for a variable were obtained from only one source (i.e. experts or data). To parameterise the native fish Bn model, we directed our efforts towards combining multiple information sources, each with associated uncertainties, and undertaking an iterative process to derive acceptable parameter estimates.

A suite of evaluative methods was used to investigate the uncertainties and inaccuracies in model structure, relationships and outputs (Coupe and van der Gaag, 2002). This process enabled a more targeted approach to the identification of parameters that needed to be accurately quantified and also to recommendations for targeted monitoring and studies to collect further information and data.

Using a Bn knowledge engineering spiral, we developed a model prototype for use in risk management. Future studies will test and refine this methodology in other domains, and iteratively assess and develop the native fish model.

8.2. Risk assessment

The development of quantitative decision-support systems in catchment management is of high priority as they enable more robust, defensible and tractable decisions. For managing ecological assets, it is preferable that models be integrative, representing the range of hazards that can potentially harm that asset, allowing risks and management actions to be identified and prioritised. However, given that both the understanding of many complex ecological systems is limited, and the existing modelling technologies for describing complex and variable systems are poor, progress has been limited. In this study, a Bn was developed to assess the suitability of using the modelling approach in an ecological risk management case study. The model needed to characterise a complex ecological system and, common to many catchments, there were high uncertainties associated with the lack of data and poor knowledge of the quantitative relationships between variables in the model.

The Bn parameterisation process described in the study enabled expert knowledge and data to be combined using a robust, iterative approach. The current model prototype can be used to inform future management decisions at multiple spatial scales, while taking into account associated uncertainties. To test the robustness of this model further, it is essential that the model be further field-tested to determine its accuracy pre- and post-management interventions or system changes.

In risk assessment, models need to fit into an adaptive management context. Adaptive management involves learning from management actions, and using that learning to improve the next stage of management. The appeal of adaptive management is driven by both our rudimentary knowledge of ecological systems and the fact that these systems are dynamic. Bns can incorporate new information into the model as it becomes available, and allow model parameters to be continually adapted and refined, enabling innovative responses to novel situations, and assisting in the learning process. This process can be conducted using the parameterisation and evaluation process described in this paper.

References

- Bhattacharyya, A., 1943. On a measure of divergence between two statistical populations defined by their probability distributions. *Bull. Calcutta Math. Soc.* 35, 99–110.
- Borsuk, M.E., Stow, C.A., Reckhow, K., 2004. A Bayesian network of eutrophication models for synthesis, prediction and uncertainty analysis. *Ecol. Model.* 173, 219–239.
- Bromley, J., Jackson, N.A., Clymer, O.J., Giacomello, A.M., Jensen, F.V., 2005. The use of Hugin to develop Bayesian networks as an aid to integrated water resource planning. *Environ. Model. Softw.* 20, 231–242.
- Castelletti, A., Soncini-Sessa, R. Bayesian networks in environmental management and planning. *Environ. Model. Softw.*, submitted for publication-a.
- Castelletti, A., Soncini-Sessa, R. Coupling real time control and socio-economic issues in river basin planning. *Environ. Model. Softw.*, submitted for publication-b.
- Clunie, P., Ryan, T., James, K., Cant, B., 2002. Implications for Rivers from Salinity Hazards: Scoping Study. Arthur Rylah Institute, Department of Natural Resources and Environment, Victoria.
- Cooke, R.M., 1991. *Experts in Uncertainty: Opinion and Subjective Probability in Science*. Oxford University Press, New York, 321 pp.
- Coupe, V.M.H., Peek, N., Ottenkamp, J., Habbema, J.D.F., 1999. Using sensitivity analysis for efficient quantification of a belief network. *Artif. Intell. Med.* 17, 223–247.
- Coupe, V.M.H., van der Gaag, L.C., 2002. Properties of sensitivity analysis of Bayesian belief networks. *Ann. Math. Artif. Intell.* 36, 323–356.
- Das, B., 2000. *Representing Uncertainties Using Bayesian Networks*. DSTO Electronics and Surveillance Research Laboratory, Salisbury, South Australia.
- Dempster, A., Laird, N., Rubin, D., 1977. Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. B* 39, 1–38.
- Dorner, S., Shi, J., Swayne, D. Multi-objective modelling and decision support using a Bayesian network approximation to a non-point source pollution model. *Environ. Model. Softw.*, in press. doi:10.1016/j.envsoft.2005.07.020.
- Gippel, C.J., Finlayson B.L., 1993. Downstream environmental impacts of regulation of the Goulburn River, Victoria. *Hydrology and Water Resources Symposium*. 31 June-2 July 1993, Newcastle, Australia.
- Gori, M., Tesi, A., 1992. On the problem of local minima in backpropagation. *IEEE Trans. Pattern Anal. Mach. Intell.* 14, 76–86.
- Hart, B.T., Burgman, M., Grace, M., Pollino, C., Thomas, C., Webb, J.A., Allison, G.A., Chapman, M., Duivenvoorden, L., Feehan, P., Lund, L., Carey, J., McCrear, A., 2005. *Ecological Risk Management Framework for the Irrigation Industry*. Land and Water Australia, Canberra (Technical Report).
- Henriksen, H.J., Rasmussen, P., Brandt, G., von Bulow, D., Jensen, F.V. Public participation and modelling using Bayesian networks. *Environ. Model. Softw.*, in this issue.
- Korb, K.B., Nicholson, A.E., 2004. *Bayesian Artificial Intelligence*. Chapman and Hall/CRC Press, London, 364 pp.
- Koehn, J.D., O'Connor, W.G., 1990. *Biological Information for Management of Native Freshwater Fish in Victoria*. Parks, Flora and Fauna Division, State Government of Victoria, Victoria.
- Laskey, K.B., Mahoney, S.M., 2000. Network engineering for agile belief network models. *IEEE Trans. Know. Data Eng.* 12, 487–498.
- Lauritzen, S.L., Spiegelhalter, D.J., 1990. Local computations with probabilities on graphical structures and their application to expert systems. In: Shafer, G., Pearl, J. (Eds.), *Readings in Uncertain Reasoning*. Morgan Kaufmann, pp. 415–458.
- Little, L.R., Kuikka, S., Punt, A.E., Pantus, F., Davies, C.R., Mapstone, B.D., 2004. Information flow among fishing vessels modelled using a Bayesian network. *Environ. Model. Softw.* 19, 27–34.
- Martin de Santa Olalla, F.J., Dominguez, A., Ortega, J.F., Artigao, A., Fabeiro, C. Bayesian networks in planning a large aquifer. The case study of the Eastern Mancha aquifer. *Environ. Model. Softw.*, in this issue.
- McGuckin, J., 2002. *An Investigative Study of the Fish Fauna of the Nagambie Lakes and Chateau Tahbilk Lagoon*. Streamline Research Pty. Ltd. for the Nagambie Angling Club, Melbourne, Victoria.
- Morgan, M.G., Henrion, M., 1990. *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge University Press, Cambridge, UK, 332 pp.
- Nicholson, A.E., Boneh, T., Wilkin, T., Stacey, K., Sonenberg, L., Steinle, V., 2001. A case study in knowledge discovery and elicitation in an intelligent tutoring application. In: Breese, K. (Ed.), *UAI01*, Seattle, pp. 386–394.
- Norsys, 2005. *Netica*. <www.norsys.com>.
- Onisko, A., Druzdel, M.J., Wasyluk, H., 2000. Learning {B}ayesian network parameters from small data sets. *Workshop on Bayesian and Causal networks*, ECAI-2000.
- Pearl, J., 1988. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Mateo, CA.
- Pearl, J., 1998. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA.
- Pollino, C.A., 2004. *Ecological Risk Associated with Irrigation in the Goulburn-Broken Catchment - Phase 2 - Adverse Changes to Abundance and Diversity of Native Fish*. Water Studies Centre, Monash University.

- Pollino, C.A., Feehan, P., Grace, M.R., Hart, B.T., 2004. Fish communities and habitat changes in the highly modified Goulburn Catchment, Victoria, Australia. *Mar. Freshw. Res.* 55, 769–780.
- Pollino, C.A., Hart, B.T., 2005. Bayesian approaches can help make better sense of ecotoxicological information in risk assessments. *Aust. J. Ecotoxicol* 11, 57–58.
- Rieman, B.E., Peterson, J.T., Clayton, J., Howell, P., Thurow, R., Thompson, W., Lee, D.C., 2001. Evaluation of potential effects of federal land management alternatives on trends of salmonids and their habitats in the interior Columbia River basin. *For. Ecol. Manag.* 153, 43–62.
- Savage, L.J., 1971. Elicitation of personal probabilities and expectations. *J. Am. Stat. Assoc.* 66, 783–801.
- Ticehurst, J.L., Newham, L.H.T., Rissik, D., Letcher, R.A., Jakeman, A.J. Bayesian network approach for assessing the sustainability of Coastal Lakes. *Environ. Model. Softw.*, in this issue.
- Ticehurst, J.L., Rissik, D., Letcher, R.A., Newham, L.H.T., Jakeman, A.J., 2005. Development of decision support tools to assess the sustainability of Coastal Lakes. In: Zenger, A., Argent, R.M. (Eds.), MODSIM 2005 International Congress on Modelling and Simulation. Modelling and Simulation Society of Australia and New Zealand, (December 2005); Melbourne, Australia, pp. 2414–2420.
- Varis, O., 1997. Bayesian decision analysis for environmental and resource management. *Environ. Model. Softw.* 12, 177–185.
- Varis, O., Fraboulet-Jussila, S., 2002. Water resources development in the Lower Senegal River Basin: conflicting interests, environmental concerns and policy options. *Water Res. Dev.* 18, 245–260.
- Woodberry, O., Nicholson, A.E., Korb, K.B., Pollino, C.A., 2004a. A Methodology for Parameterising Bayesian Networks. School of Computer Science and Software Engineering, Monash University (Technical Report).
- Woodberry, O., Nicholson, A.E., Korb, K.B., Pollino, C.A., 2004b. Parameterising Bayesian networks. In: Webb, G.I., Xinghuo, Y. (Eds.), *Lecture Notes in Computer Science. AI 2004: Advances in Artificial Intelligence: 17th Australian Joint Conference on Artificial Intelligence*, Cairns, Australia, pp. 1101–1107.