A DECISION ANALYSIS APPROACH TO TMDL IMPLEMENTATION DECISIONS: MERCURY TMDLS IN THE SAN FRANCISCO BAY AREA

William B. Labiosa[1, 2], James O. Leckie[1], Thomas Mumley[3], James Rytuba[4], and Richard Bernknopf[4]

1. **Department of Civil and Environmental Engineering, MC 4020, Stanford University, Stanford, California USA 94305**
2. **Contact: labiosa@stanford.edu**
3. **California Regional Water Quality Control Board, San Francisco Bay Region**
4. **US Geological Survey, Menlo Park, California, 94025**

ABSTRACT

This paper describes a decision analysis approach to TMDL implementation decisions for mercury using a hypothetical mine-impacted tributary in the San Francisco Bay as an example. Decision analysis is a theoretically sound approach for making significant decisions under uncertainty (see, e.g., Howard, 1968; 1988; Keeney and Raiffa, 1976; Clemen, 1996; Merkhofer, 1999). The Bayesian probabilistic nature of decision analysis makes it ideal for integrating diverse information, including the results from scientific and engineering models, cost and benefit models, empirical data, and expert judgment. One significant advantage of a decision analysis approach is its explicit separation of a decision problem into alternatives, information, and preferences. This, in theory, allows decision makers and stakeholders to separate "what we know" from "what we want". It is hypothesized that a more explicit separation of information and values/preferences will focus the debate. While traditional decision analysis assumes a single rational decision maker (where "single" may also denote a group that agrees on information and preferences), it can be extended to multiple decision maker situations in a variety of ways. Evaluating various extensions of decision analysis in a TMDL implementation stakeholder context is one of the primary goals of this on-going study. It is hypothesized that, in general, decision analysis provides a helpful decision framework for a TMDL implementation planning/stakeholder process in many circumstances.

KEYWORDS

Decision analysis, TMDLs, mercury, water quality predictions, influence diagrams, Bayesian networks, probabilistic networks, information gathering decisions, load allocation decisions, mitigation decisions, sensitivity analysis, value of information

INTRODUCTION

Background

There are a number of examples of the use of decision analysis for environmental decision making in the literature, often in the area of site selection or choosing between remediation, restoration, or technology alternatives (e.g., Keeney, 1980; Merkhofer and

Keeney, 1987; Maguire and Boiney, 1994; Reckhow, 1994a; Merkhofer *et al.*, 1997; Perdek, 1997; Kruber and Schoene, 1998; Freeze and Gorelick, 1999; Merkhofer, 1999; Bonano *et al.*, 2000; Anderson and Hobbs, 2001). Environmental decision situations are often rife with uncertainty and controversy, requiring the integration of diverse kinds of information and compromises between diverse interests. TMDL load allocation decisions are typical in this regard (NRC, 2001; Boese, 2002). Common TMDL decision issues include dealing with appreciable scientific uncertainty and information gaps in understanding the relationships between loadings, mitigation, and effects, determining whether to make allocation decisions based on what is currently known or whether to collect new data and perform new analyses before making those decisions, and prioritizing pre- and post-implementation monitoring activities. Decision analysis provides a normative (as opposed to descriptive) framework for providing decision clarity for these kinds of decision problems. The theory behind decision analysis does not attempt to predict decision strategies that people *will* choose, but rather, it attempts to predict decision strategies people *should* choose, given a set of beliefs, alternatives, and preferences (the decision basis). In a group decision situation, if consensus is achieved on the decision basis, decision analysis can be used to determine optimal decisions. If consensus is not achievable, decision analysis may be used to highlight areas of agreement and disagreement, allowing insights into potential compromises and/or defining positions for negotiation.

Decision analysis makes use of the Bayesian (subjective) definition of probability, which treats uncertainty as a probability and allows the decision maker to combine various kinds of information into a unified probabilistic framework. For decisions that involve perturbations to natural systems, Bayesian (probabilistic) networks that are built up from the best available scientific models, data, and expert judgments can be used to predict the consequences of those decisions (Borsuk et al., 2001, 2002; Stow et al., 2003; Reckhow et al, 1999). In practice, empirical models and expert judgment are the most straightforward means to creating the needed probabilistic relationships. While uncertainty analyses of mechanistic models can be used for this purpose, the computational burden can be excessive (Reckhow, 1999).

Bayesian (probabilistic) networks are by design "as simple as they need to be, but no simpler". This approach allows the modeler to focus on predictive accuracy for the time and spatial scales desired for the variables of interest to the decision makers, removing details that are determined to be extraneous to the decision problem. As pointed out by Reckhow (1999), this approach often leads to superior predictive accuracy compared to deterministic scientific models of water quality impacts. The loss of mechanistic descriptive power is compensated by the ability to perform sensitivity analyses, explore scenarios probabilistically, and estimate credibility of compliance predictions. Recent work has demonstrated that water quality management effects can be effectively modeled using Bayesian (probabilistic) networks (e.g., Reckhow, 1999; Borsuk *et al.*, 2001; Borsuk *et al.*, 2002; Stow *et al.*, 2003). Since compliance is predicted probabilistically, a margin of safety (MOS) can be explicitly considered in terms of credibility of compliance predictions.

From a decision analysis perspective, the Bayesian network model of interest is the influence diagram, which combines decisions (“what you can do”) with a model of key uncertainties (“what you know”), subject to a valuation model (“what you care about”) (Howard and Matheson, 1984; Shachter, 1986; 1988). If consensus is achieved on preferences, influence diagrams allow determinations of optimal decisions, sensitivity of the optimal decision to key uncertainties and assumptions, and value of information on uncertainties, which may be used to plan future information gathering activities. Value of information refers to the fact that improvements in the state of information before a decision is made can lead to a change in the predicted optimal policy. It is the potential for changing the optimal policy that generates economic value (see Howard, 1968; Lawrence, 1999).

Even without consensus on preferences, sensitivity analysis can be performed to explore relationships between key uncertainties and variables of interest (e.g., water quality endpoints), allowing the decision makers to explore “what-if” scenarios of interest. When preferences are ignored (i.e., the value model is removed), the underlying Bayesian network may be referred to as a “belief network”. In the context of water quality management decisions, belief networks can be thought of as modeling the response of the natural system to management strategies. For example, one could use a belief network to probabilistically explore the relationship between mercury load reductions and mercury fish tissue levels under a variety of scenarios, in essence demonstrating how *beliefs* about future mercury fish tissue levels change with load reductions. This research project makes extensive use of influence diagrams and belief networks as tools for performing TMDL decision analysis. However, the emphasis is on the use of these tools for supporting decisions, not as water quality models *per se*.

This paper demonstrates an influence diagram model of mitigation/load allocation decisions for a simple mercury TMDL example. Such a model can be used throughout the TMDL decision process, including initial information gathering decisions, load allocation/mitigation decisions, and post-implementation monitoring decisions. The essential insight is that information gathering/monitoring decisions, whether made before or after allocation decisions, draw their value from making better load allocation/mitigation decisions. For this reason, information gathering decision models *build on* load allocation/mitigation decision models. Our load allocation/mitigation decision model integrates a Bayesian (probabilistic) network model of environmental system response to mitigation decisions with a valuation model, allowing insights into the credibility of compliance with multiple numerical standards, insights into sensitivity of conclusions to small changes in model parameters, and, if a value model can be defined, the determination of optimal strategies.

It is emphasized that decision analysis applied to group decision situations should be thought of as a *process* by which groups may discover useful insights that highlight where consensus may be achieved and where obstacles requiring clarification, negotiation, mediation, or litigation may lay. There are many competing versions of decision analysis with variations on how alternatives are generated, uncertainty is represented, preferences are elicted, etc. In this paper we describe a decision analytic approach that is based on small group elicitation of goals, objectives, and alternatives, a probabilistic model of natural system response, and several potential methods for eliciting and representing preferences. Other related approaches may be just as appropriate, depending on circumstances. One of the focuses of this paper is dealing with the problem of competing preferences between stakeholders, both from the perspectives of making decisions and representing preferences.

At the highest level, decision analysis divides the decision problem into *alternatives*, *information*, and *preferences*. In the context of public environmental decision making, these could be cast as: 1) decision framing/strategy generation; 2) information modeling/synthesis/forecasting; and 3) multiattribute utility analysis, negotiation among interest groups, or other methods of eliciting and representing preferences. Each of these aspects of decision analysis will be described further through examples, with the goal of showing how decision analysis can create clarity in a complex decision problem. But first, we discuss the importance of considering uncertainty in the TMDL decision-making process.

Uncertainty in TMDL Decisions

Models play and will continue to play a central role in the TMDL development and implementation process (Reckhow, 1999; NRC, 2001; Lung, 2001; USEPA, 2002). Whether the models are empirical (statistical) or mechanistic, they represent the best scientific understanding of how contaminant loadings relate to water body impairment of designated beneficial uses (NRC, 2001). Once a waterbody is listed as impaired, predictive models are used to assess the relative contributions of various pollution sources, to predict the total load reduction required to meet ambient water quality standards, and to predict the relationships between specific control measures (e.g., point source load reductions) and water quality targets (e.g., ambient water concentration of a particular pollutant) in the load allocation process.

Decision-making related to TMDL development and implementation requires one to answer questions related to determining the reasons for non-attainment of beneficial use and evaluating strategies for mitigating those determined causes. Neither of these questions can be answered with certainty. Uncertainty, whether the source is incomplete knowledge about the natural system, analytical error, or the stochastic variability inherent in natural systems, is a reality that any water quality management decision framework must recognize, assess, and, when possible, reduce (NRC, 2001). The decision analytic framework proposed in this paper specifically addresses model uncertainty in the context of decision making, using Bayesian network models to integrate predictive uncertainty about the response of the natural system to proposed mitigation strategies with stakeholder valuations of the strategies being considered.

Uncertainty in model predictions can be large and, when explicitly considered, can confound interpretation of results in terms of the decisions that need to be made (Reckhow, 1994b). Uncertainty has been, however, often treated superficially in water quality management decisions, which can be a major source of contention between stakeholders and regulatory agencies (Ortolano, 1997; NRC, 2001). Historically, this occurred because the ability to analyze uncertainty was limited by computing power and, in some cases, by a lack of understanding of how to feasibly model and propagate uncertainty in large mechanistic water quality models. Besides the technical aspects, even when uncertainty analysis is performed well, the political reality is that discussions of the estimated uncertainty often get bogged down with arguments that have more to do with preferences than information. In fact, the use of decision analysis is an attempt to incorporate uncertainty directly into TMDL modeling and decision-making in a manner that separates information and preferences. In effect, this attempts to separate the *estimation* of uncertainty from the *interpretation* of uncertainty. Disagreements about particular beliefs and preferences can be expected to remain, but decision analysis may be able to focus the argument on those sources of disagreement, reducing confusion about the impact of uncertainty on decisions. Downplaying uncertainty to avoid these confrontations may make for an easier stakeholder process in the short term, but that strategy runs of the risk of resulting in poorly informed decisions. The National Research Council (*ibid*.) suggests the use of adaptive management to deal with the significant uncertainty involved in TMDL decisions, an approach that is being employed in many TMDLs. As discussed by Reckhow et al. (2002), an adaptive management approach may be modeled with Bayesian networks, but further discussion of adaptive management is beyond the scope of this paper.

Preferences

If decision outcomes can be valued in terms of a single attribute (e.g., an exchangeable resource like dollars), and consensus can be reached regarding those values and attitudes toward risk, decision analysis can be applied straightforwardly to determine an optimal decision policy, sensitivity analysis can be used to determine the value of information, etc. (e.g., Howard, 1968; 1988; Marshall and Oliver, 1995; Clemen, 1996; Merkhofer, 1999). The optimal decision policy for an uncertain decision situation is the policy that maximizes expected utility, a measure of value. By making maximum expected utility the decision criterion, the utility of a particular outcome is weighted by its probability of occurrence, so that the strategy that yields the highest expected utility can be thought of as promising the “highest probability of achieving the best outcome”.

When a group agrees to cooperate and work towards consensus on information beliefs and preferences, the s*ingle decision maker* decision analysis approach may be used. Single decision maker problems involving utility over uncertain monetary outcomes are solved in terms of expected utilities, incorporating risk attitudes. Non-monetary outcomes can be accommodated in decision analysis using the “preference probability” interpretation of utility, in which the utility of an outcome is interpreted as the probability of obtaining the best outcome instead of the worst outcome. The approach we explore in this paper is the use of multiattribute utility analysis to directly define a mapping from either monetary or non-monetary outcomes to utilities (Howard, 1984b; Marshall and Oliver, 1995; Clemen, 1996; Lawrence, 1999).

However, the assumptions applying in single decision maker situation obviously would not describe many TMDL decision situations, which instead can be expected to have multiple goals with multiple associated attributes with perhaps no obvious consensus on valuing the various possible outcomes. Note that there are two important issues at stake here: 1) TMDL goals have multiple attributes that may not be expressible in terms of a single measure like dollars; and 2) based on experience, we can expect disagreements between work group/stakeholder group members about valuing outcomes even within an agreed-upon multiattribute framework. Each of these issues can be dealt with, if the TMDL decision-making group is willing to cooperate. This does not require that consensus in preferences is achieved, but it does require that group members agree to faithfully participate in the decision analysis process.

Work groups and stakeholder groups may use decision analysis in a number of ways, including a "competing models" approach in which the work group/stakeholder group partitions into sub-groups that agree to act cooperatively in determining mutual preferences and preferred alternatives for the purpose of arriving at negotiating positions for each sub-group (Chechile, 1991). In other words, the sub-groups agree to effectively behave as a "single decision maker" to determine recommended strategies according to the sub-group's viewpoint. While major differences may be found between the various recommended approaches, numerous points of agreement are expected. At this point, decision analysis may be used further with mediated compromises on preferences and information that allows a group "best compromise strategy" to be formulated, but it may be necessary to resort to a purely negotiated or political compromise at this point. The advantage of applying decision analysis in this latter case is that the sources of disagreement can be more easily identified and that potential compromises may become more apparent. However, if a sub-group is non-cooperative and misrepresents their beliefs and preferences in the analysis, decision analysis may not be a useful tool for the TMDL decision-making process. Note that other analytical approaches are similarly hobbled by deliberate attempts to misrepresent positions (e.g., cost-benefit analysis). In such cases, political solutions may be inevitable. In cases in which group members are willing to cooperatively state their beliefs and preferences, decision analysis is a robust process that should be considered.

Decision Analysis Process

Figure 1 shows a flow diagram representing the decision analysis cycle (Howard, 1984a). In a real application of decision analysis, individual steps may be emphasized or de-emphasized, depending on the particular situation. Also, a particular step may be accomplished using very different tools and some tools may be used in more than one step. So, from a "tools perspective", two different decision analysis applications may appear to be very different, so much so that it may be difficult to see the relationships between the two approaches. However, taking a decision analysis cycle perspective, one can see how the seemingly different approaches accomplish the basic steps in decision analysis. For the purposes of this paper, an influence diagram approach to TMDL decisions, we will focus on 1) decision framing/structuring; 2) probabilistic modeling of

Figure 1 - a) Decision Analysis Cycle

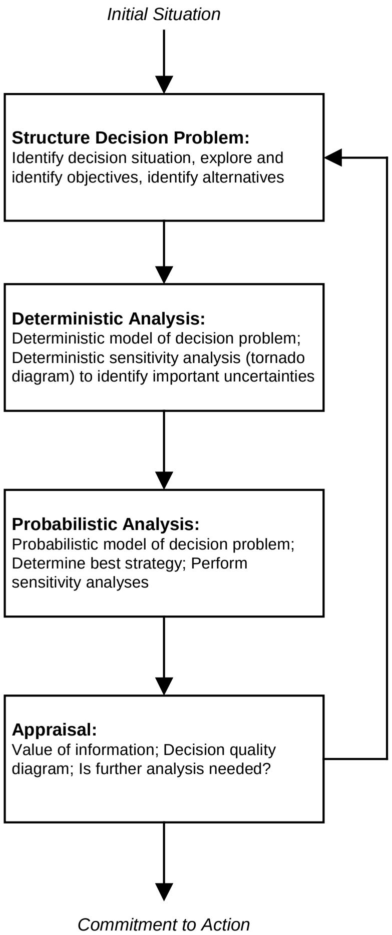

the natural system response; 3) sensitivity analyses; and 4) dealing with preferences and potentially determining optimal strategies and value of information. In particular, we will explore the application of decision analysis to load allocation/mitigation decisions and information gathering decisions.

The initial step in the decision analysis cycle is preliminary framing and structuring of the decision situation in such a way that decision analysis may be used to evaluate the various alternatives. In particular, framing involves identifying and discussing decision performance measures, e.g., decision objectives and their associated measurable/ predictable attributes. Performance measures (attributes) can be thought of as gauging the consequences that the decision maker cares about, so that the range of possible outcomes may be represented in a meaningful way (Keeney, 1992).

The influence diagram (as a Bayesian network) is a powerful tool that allows the decision analyst to perform the deterministic analysis phase, the probabilistic analysis phase, and, if a value model can be determined, to estimate the value of information on key uncertainties and assumptions. In brief, the deterministic analysis phase translates the results of the framing analysis into a mathematical model for the purpose of determining which uncertainties are important enough to warrant probabilistic modeling in the subsequent probabilistic analysis phase. The probabilistic analysis phase assigns probabilities to the identified key uncertain variables. Input variables that have little effect on the value model output are assigned nominal (base) values and, thus, are treated deterministically. The required probability distributions are either modeled empirically from data or assessed from experts and/or decision makers. In some cases, it may be advantageous to probabilistically combine empirical models with expert opinion. In the probabilistic analysis phase, optimal decisions may be determined if a value model can be constructed. The influence diagram may be further manipulated to perform probabilistic sensitivity analysis to determine how sensitive the optimal policy is to current beliefs about key uncertainties. The decision analyst may find that the optimal policy may change given small changes in probability distributions for a key uncertainty, in which case further analysis may be recommended. Performing and reporting the results of sensitivity analysis may be critical in achieving the degree of "decision transparency" that promotes buy-in from stakeholders.

Value of information analysis may also be performed as this stage to determine if additional information may have the potential to change the optimal policy. From a decision analysis perspective, new information only has value when the optimal policy may change in response to the new information (Howard, 1968; Lawrence, 1999). Since value of information analysis requires consensus in preferences, it may not resolve disagreements about information gathering or technical review activities between sub-groups. However, it can provide the basis for positions on information gathering and technical review activities within sub-groups and can shed light on the sources of agreement and disagreement regarding these activities.

Influence Diagrams (Bayesian Networks) for Environmental Decision Analysis: Graphical Tools for Decision Problem Framing

Influence diagrams are often used as framing tools for graphically representing the decision problem in terms of the relationships between decisions, uncertainties, and performance measures (Howard and Matheson, 1984; Shachter, 1988; Howard, 1990; Merkhofer, 1990; Marshall and Oliver, 1995). The influence diagram can be constructed as a group exercise in decision framing, focusing attention on the relationships between the important variables in the decision situation, including decision strategies, uncertain variables describing the state and response of the natural system, and variables related to valuing outcomes. In addition to graphically representing important aspects of the decision problem, the influence diagram can be used to determine information/ forecasting requirements, probability assessment order, and, if decision trees are to be used, decision tree structure. Deterministic sensitivity analysis may later determine that one or more uncertainties can be treated deterministically and hence the influence diagram may evolve during the decision analysis problem. The role of the influence diagram in determining information and modeling/forecasting needs is very important: this approach helps decision makers and technical experts/scientists communicate about what information is important *in terms of the decisions to be made*.

Influence Diagrams as Bayesian Networks for Solving Decision Problems

In addition to decision framing, influence diagrams can also be used directly as Bayesian network models by adding to the graph the requisite probability structures needed for modeling consequences and value. In this use, an influence diagram is a class of Bayesian networks that may include nodes representing uncertain system variables, deterministic system variables, decision variables, and a value variable. Optimal decisions are those that maximize expected utility through relationships between the value variable and the other variables. Thus, influence diagrams can be used in lieu of or in parallel with decision trees to solve for optimal decisions, to evaluate sensitivity of the optimal decision to information and model assumptions, to estimate the value of information and control, and to make inferences from the available data important to the decision situation (Howard and Matheson, 1984; Shachter, 1986; Oliver and Smith, 1990; Pearl et al., 1990). In the approach described in this paper, decision trees (Chechile, 1991; Marshall and Oliver, 1995) are avoided altogether and the Bayesian network is used as the primary analytical tool.

To emphasize the point, the Bayesian network version of the influence diagram may be used to make *predictions* about the response of the natural system to changes in those variables over which the decision maker has some control. Bayesian networks without decision or value nodes ("belief networks") can be also used to model reasoning under uncertainty and may be used as predictive tools in decision situations, e.g., water quality management decision situations (Reckhow, 1999; Borsuk et al., 2001). One of the advantages of using a Bayesian network approach is that the model evolves as new information is collected, yielding an updated model that reflects the current state of knowledge about the system of interest, synthesizing prior information and new evidence using theoretically sound probabilistic calculus (Jensen, 2001; Shachter 1986, 1988; Reckhow, 1999; Pearl *et al*., 1990; Varis, 1995).

Figure 2 shows an example of a Bayesian (belief) network representing causal relationships between precipitation, creek flow, river flow, mine mercury load, creek mercury load, and total mercury in water (Hg_T). The belief network consists of a graph and probabilistic data associated with the nodes in the graph. The graph consists of nodes (ovals) connected by arrows. Ovals represent chance (uncertain) nodes and associated with each chance node is a random variable. The random variables in the Bayesian network represent the attributes of interest to decision makers. Arrows represent potential conditional probabilistic dependence between the various random variables and can be drawn in a causal direction. Graphically, the arrow points from the "parent node" to the "child node", which intuitively indicates that the child node somehow "depends" on the parent node. More precisely, an arrow from a parent node to an uncertain variable (child) means that the probability distribution in the uncertain variable (child) is conditioned on the state of the parent node. The absence of an arrow between two variables indicates that the variables are conditionally independent. If there is a directed path between two variables (i.e., there exists a set of arcs between them which can be traversed in the direction of the arcs) which do not have a direct parent/child relationship, those variables may or may not be relevant to one another, depending on the state of information. For example, Figure 2 asserts that precipitation may be relevant to total mercury in water (Hg_T) if at least one of the values for “Mine Hg Load”, “Creek Hg Load”, and “River Flow” has not been observed. But, it also asserts that, given observations for “Mine Hg Load”, “Creek Hg Load”, and “River Flow”, precipitation and total mercury in water are conditionally independent of each other. These assertions of conditional independence are very important in terms of understanding information needs and performing decision analysis.

The variables included in a network may be included for a variety of reasons, including the decision makers’ direct interest in the state of a variable (e.g., Hg_T) or because the variable helps to interpret or predict those variables of direct interest (e.g., precipitation). It is important to understand that variables needed from a technical perspective for modeling a particular complex system do not need to be shown in the version of the Bayesian network used for decision analysis, communicating with decision makers or stakeholders, etc. Variables needed only for modeling reasons can be probabilistically absorbed into the network, which yields the same results as before the nodes were removed. The local representation is changed, but the global probabilistic relationships are not affected (Shachter, 1988; Pearl et al., 1990). As there may indeed be variables of interest to scientists about the natural system being modeled that are not important to decision makers, this is an important point to understand.

The conditional probabilistic relationships between conditionally dependent variables can be quantified in a modular fashion using an approach suitable to the kind and amount of information available, where this modularity follows from the conditional independence relationships in the model (Reckhow, 1999; 2002). This allows various kinds of statistical and subjective probabilistic information (from data, models, and expert judgment) to be integrated into a single probabilistic network model that can be used for predictions and inferences of use in decision-making situations. Prediction refers to following an arrow in the forward direction, i.e., predicting the probability distribution of

Figure 2 - Bayesian Network Example

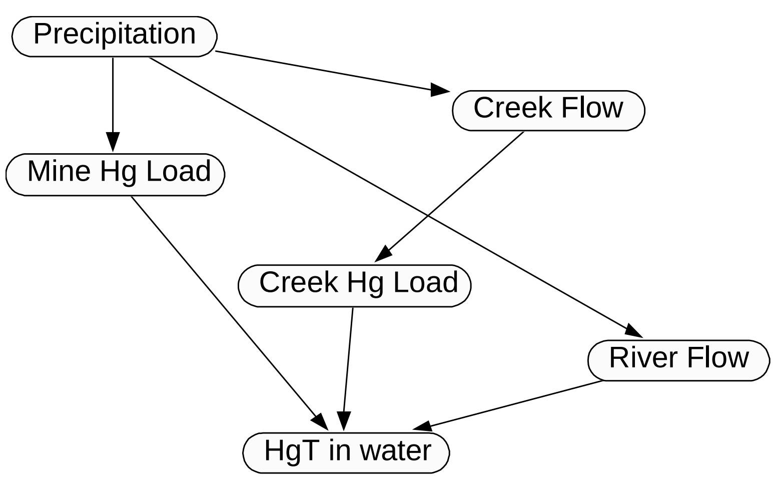

a child node based on the values or distributions of its parent nodes. Inference refers to following an arrow in the reverse direction, i.e., inferring the probability distribution of the parent nodes based on evidence about the value of the child node(s) (Jensen, 2001). The ability of a Bayesian network to make predictions is useful, for example, when trying to model the effects of particular mitigation strategies on water quality attributes of interest. Learning what new evidence about a predicted variable means in terms of hypotheses about cause and effect between parent and child nodes is an example where the ability of Bayesian networks to perform inference may be useful.

A point that needs to be emphasized here is that while such an influence diagram model (or any other model) is an imperfect representation of the real system, it should faithfully represent how the decision maker *believes* the real system will behave, given the available data and current scientific understanding. The decision maker can do no better than this when making a decision. In particular, if the optimal strategy is sensitive to slight changes in the underlying probability distributions, then a value of information analysis may determine that re-framing the load allocation/mitigation decision problem as an information gathering decision problem may be the best course of action. If there are regulatory constraints that prevent this re-framing, using the existing model to suggest optimal strategies is the best course of action, from the decision analytical perspective.

METHODOLOGY AND DISCUSSION

Framing the Decision: Objectives Hierarchy

Decision framing tools are used to represent the decision situation in a way that enables the evaluation and comparison of alternatives according to criteria that are meaningful to the decision maker. The objectives hierarchy is a common framing tool that identifies and organizes decision outcome performance measures (attributes), which are used as evaluation criteria (Keeney, 1992) and can be used as variables in an influence diagram representation of the decision problem. Keeney (*ibid.*) organizes the objectives hierarchy with: 1) an overarching decision goal at the top of the hierarchy; 2) a set of issue-specific objectives consistent with and subject to this top goal; and 3) a set of attributes consistent with and subject to the specific objectives. There may be more than one level of objectives between the top decision goal and the decision attributes, depending on the framing desired by the decision makers. Attributes are ideally the performance measures that the decision makers care about, and they should be chosen to be well-defined, measurable (at least in theory) and predictable. When the objectives hierarchy is complete, the attribute set formed should be comprehensive (capture all of the aspects of value at stake), minimal (as small in number as possible), independent of one another, and operationally feasible (Keeney and Raiffa, 1976; Reckhow, 1994a). These requirements ensure that there are no “holes” or “double-counting” in the analysis and that a suitable value model can be constructed (Merkhofer, 1999).

Figure 3 shows a hypothetical objectives hierarchy for a mercury TMDL for a tributary to the south San Francisco Bay. In practice, the objectives hierarchy would be developed by the TMDL work group and/or stakeholder group, with the help of the decision analyst.

The top goal in this example objectives hierarchy is to manage mercury contamination in the watershed, with objectives pertaining to protecting human health and wildlife (sub-objectives of reducing mercury methylation potential and methylmercury levels), meeting Basin Plan water quality objectives for total mercury, maintaining adequate flood control, meeting the ten and twenty year total mercury load reductions under the San Francisco Bay mercury TMDL, and controlling compliance costs. Each objective is translated into one or more attributes, and this is shown graphically by arrows pointing from a given objective to its attributes (performance measures). Attributes serve multiple roles in decision analysis. They form the basis of the value model, since they are the performance measures that matter to decision makers, and they define information needs for decision modeling. In this latter role, TMDL decision situation attributes help define which natural system variables need to be modeled for relating management strategies to value, as will be demonstrated below.

For example, the "mercury fish tissue levels" attribute might be defined as the average mercury fish tissue burden of a particular fish species (with perhaps specified weight range, sex, etc.) within the watershed over some time scale. An attribute that might be less obvious in the context of mercury mitigation, but that may be very important to some stakeholders is flood capacity. The use of flood capacity as an attribute allows decision makers to keep track of the impact of mitigation strategies on flood capacity, while

Figure 3 - Hypothetical Objectives Hierarchy for Managing Mercury in a Small Mine-Impacted Tributary to the South Bay.

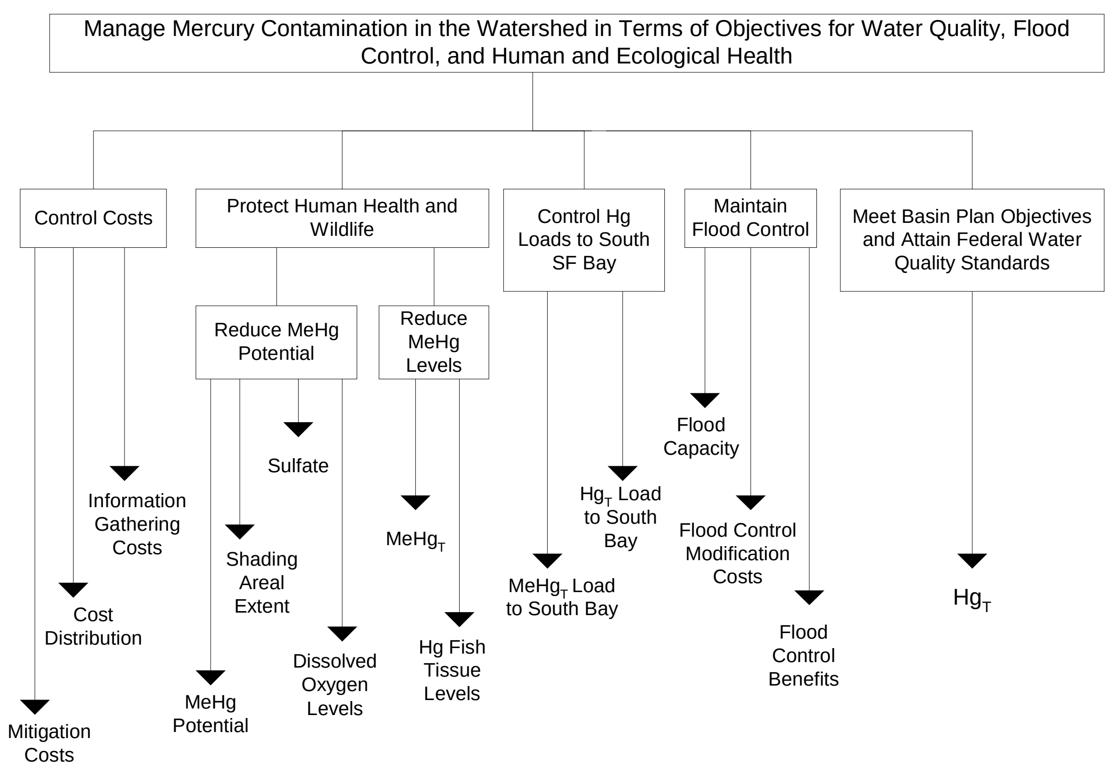

simultaneously evaluating those strategies in terms of other attributes. While establishing TMDL performance measures is an explicit activity in the TMDL process, it is important to identify a list attributes that capture *all* stakeholder values that may be significantly affected. To emphasize the point, it is important to frame the problem not just in terms of "technical TMDL endpoints", but also in terms of attributes that characterize objectives that matter to stakeholders in terms of idiosyncratic preferences. In fact, capturing this latter class of attributes may make the difference between understanding *why* stakeholder values lead to disagreements about acceptable strategies later in the decision process and finding a situation in which there are arguments that are seemingly about "technical information", but that really reflect unstated preferences.

It is common, but often unheeded, advice in the decision analysis literature to appropriately focus attention at this step since careless framing can lead to "solving the wrong problem", leading to inappropriate or incomplete consideration of alternatives, a short-sighted understanding of the decision situation, and a misappropriation of resources (Howard, 1968; 1988; Reckhow, 1994a; Clemen, 1996; Merkhofer, 1999). Nevertheless, decision makers often treat this stage cursorily and plunge quickly into more familiar territory: technical problem framing, information gathering, modeling, and analysis. Decision makers often have a good understanding of many aspects of the decision problem "going in" to a particular decision situation, which can sometimes lead to the misapprehension that detailed decision framing exercises are unneeded. However, extensive decision framing can lead to better planning and resource allocation and to evaluating the "right alternatives" in terms of the "right attributes" for making good TMDL decisions. To a significant degree, TMDL guidance documents already promote this activity from the technical perspective.

Identifying Alternatives and Generating Strategies

Keeney (1992) describes a number of methods for using the attributes and objectives from the objectives hierarchy to explore and generate decision alternatives. Clemen (1996) provides a basic and useful summary of various techniques, including some of the methods discussed in Keeney (*ibid*). The methods build on the identified goals, objectives, and attributes, stressing the importance of flexibility and creativity. One tool in particular may be useful for generating TMDL strategies: the strategy table. Figure 4 shows a simple example of a strategy table with two strategies. Strategy tables are fairly intuitive and the tool can be used in a group setting without much introductory material required. The basic idea is to capture the possibilities, then to select a manageable number of strategies as alternatives for further decision analysis. A strategy consists of a set of single elements from each column in which the combination of those elements makes sense as an approach. There will, of course, be combinations that are incoherent and these combinations would not represent a viable strategy. In a real TMDL allocation decision situation, the strategy table would be expected to have more elements (columns), making such an approach useful for brainstorming and organizing complexity. In Figure 4, two strategies are shown for illustration: 1) a methylmercury potential mitigation strategy that includes "medium reductions" for mine site and creek mercury loads and an aggressive reduction of mercury methylation potential and 2) a mine site mercury load reduction strategy that includes a large reduction requirement for mine site mercury loading and minimal reductions for creek mercury loading and mercury methylation potential.

Once an objectives hierarchy has been created and alternatives have been generated, the decision analyst will work with the group to create an influence diagram from the chosen attributes, alternatives, and variables representing identified important uncertainties. Attributes may become variables in an influence diagram or may become part of the value model, depending the nature of the attribute. It may be worthwhile to revisit the objectives hierarchy and strategy table after the influence diagram has been created to see if revisions are necessary.

Multiattribute Utility Analysis

Multiattribute utility analysis (MUA) is designed to deal with the complexity of eliciting and representing the values at stake in complex decision problems like environmental decision situations (Keeney and Raiffa, 1976; Gregory, 1999; Merkhofer, 1999; Prato, 2003). In particular, multiattribute utility analysis (MUA) or other approaches (e.g, the analytic hierarchy process) may be used to elicit and represent preferences when multiple decision attributes/criteria are important (Chechile, 1991; Marshall and Oliver, 1995; Merkhofer, 1999). In general, MUA is conceptually simple, but may become operationally complex with details that should burden the decision analyst and not the

Figure 4 - Strategy Table Example.

Strategies ↓	**Mine Site Hg Load Allocation**	**Creek Hg Load Allocation**	**MeHg Potential Mitigation**
MeHg Potential Mitigation Focus	Small Reduction	Small Reduction	Small Reduction
Mine Site Load Reduction Focus	***Medium Reduction***	***Medium Reduction***	Medium Reduction
	Large Reduction	Large Reduction	***Large Reduction***

decision makers/work group. One problem is that lay people may see this apparent complexity as being suspect, so great care should be taken to ensure that decision makers/work group members sufficiently understand the concepts being used so that the preference representation approach is trusted (Morgan and Henrion, 1990).

While there may be consensus that one alternative appears superior to the others in terms of one particular attribute, it may appear inferior in terms other attributes. Trade-offs between attributes is thus usually necessary and this idea is at the core of multiattribute decision making (Keeney and Raiffa, 1976; MacCrimmon and Wehrung, 1977; Merkhofer *et al*., 1997). MUA can be used directly to rank alternatives in terms of weighted utilities (Prato, 2003), ignoring probabilities of outcomes in the decision making process. However, it can also be used in a decision analysis framework that includes a probabilistic treatment of uncertainty and that determines best policies based on expected utility.

Once the decision analyst elicits preferences between outcomes among the various decision attributes for each decision maker sub-group, the consensus preferences for each sub-group can be aggregated into a multiple-attribute utility function and maximum expected utility can be determined. Note that MUA does not require monetization of preferences, one of the appeals of the technique. Other approaches like probabilistic cost-benefit analysis, cost-effectiveness analysis, minimization of chance of worst possible outcome, etc. could also be used, depending on the situation (Morgan and Henrion, 1990).

To illustrate the MUA process, Table 1 shows an example multiattribute utility analysis (MUA) for a few outcomes for the two strategies from Figure 4, MeHg Potential mitigation focus (Strategy 1) and mine site Hg load reduction focus (Strategy 2), as evaluated by a hypothetical decision maker sub-group. Other sub-groups could be expected to have different results. This analysis assumes that the decision problem is being modeled with discrete probabilities and that there are a finite number of possible outcomes. MUA can then be used to define a utility function over those outcomes, as suggested by this example. Here the sub-group chose a weighting scheme of 0.3 for cost, 0.6 for credibility of compliance with total mercury load to the south bay target, and 1.0 for credibility of compliance with mercury fish tissue level target. Composite utilities are shown in the rightmost column. Other approaches to defining a multiattribute utility function or analogous scoring functions could be used, depending on the wishes of the subgroup.

Credibility of compliance refers to the conditional probability (not “confidence” in the statistical sense) that a particular attribute has a value that meets a particular target (threshold) value, as computed within the network. The threshold could itself be uncertain, but need not be. The credibility of compliance in essence becomes a node in the Bayesian network conditioned on the attribute (target) of interest. The concept is similar to the “confidence of compliance” described in the literature, but is referred to

Table 1. Multiattribute Utility Analysis for Several Outcomes for Strategies 1 and 2 for a Particular Sub-Group

Possible Outcomes	**Utility on Cost**[1]	**Utility on COC Load**[2]	**Utility on COC fish**[3]	**Composite Utility Using Weighting Scheme**
Strategy 1, Cost = 15, COC Load = 30%, COC fish = %20	10	2	1	5.2
Strategy 1, Cost = 30, COC Load = 35%, COC fish = 50%	6	3	9	12.6
Strategy 2, Cost = 30, COC Load = 40%, COC fish = 30%	6	4	4	8.2
Strategy 2, Cost = 50, COC Load = 60%, COC fish = 45%	2	8	6	11.4
.... other outcomes	...	...	...	...

[1] Mitigation cost

[2] Credibility of compliance with total mercury load to the south bay target

[3] Credibility of compliance with mercury fish tissue level target

here as a "credibility" since it is not statistical confidence to which we are referring. The concept may prove to be useful for evaluating mitigation/allocation strategies since strategies that yield higher probabilities of success would naturally be more appealing.

While the sub-groups may well arrive at different conclusions, their respective positions should be well-defined in terms of beliefs about probabilities of outcomes and their preferences. Consensus building exercises that attempt to arrive at compromises may be performed or negotiation between the various sub-groups may follow. Again, the advantage of using decision analysis is that the positions of each sub-group should be clear and the various sources of differences in positions should be apparent. The price that must be paid to get to this point is that the work group members must agree to accurately state their beliefs and preferences. If trust is lacking, then appropriate measures may be required (e.g., allowing sub-groups to develop their positions privately without sharing analyses) or non-analytical approaches may be required. Exploring the possibilities and determining "what works and why" is an active area of research.

Such a multiattribute utility function could be used in an influence diagram model of TMDL decisions to determine optimal decisions, perform sensitivity analysis for expected utility, and estimate the value of information in terms of utility. If preferences can be expressed in monetary terms, a monetary value of information can be estimated for the uncertainties. One could argue whether or not this appropriate, but the choice reflects the wishes of the sub-group cooperating in the analysis. Whether the scale is dollars or utility, value of information provides a useful signal for prioritizing information gathering activities and technical review needs.

Influence Diagrams for Forecasting Allocation Decision Consequences

Designing and implementing a Bayesian network model occurs in three stages: 1) development of the graphical model linking the identified variables in terms of conditional independence relationships; 2) assessment of the required conditional or marginal probability distributions for each variable; 3) entering evidence/observed data (if applicable) on observable nodes in the compiled model to see how beliefs in unobserved nodes are affected (Jensen, 2001). In the example that follows, the Netica™ Application for Belief Networks and Influence Diagrams (Norwys Software Corp, 1996) was used to implement the model. Other Bayesian network development environments include the MatLab Bayes Net Toolbox (Murphy, 2002), Microsoft® Bayes Networks (MSBN), and Analytica® (Lumina Decision Systems, Inc.). Russell Almond at the University of Washington maintains a website listing and reviewing Bayesian network software: http://www.stat.washington.edu/bayes/almond/belief.html#MSBN. Morgan and Henrion (1990) discuss considerations and issues in choosing a computing environment for probabilistic analysis.

Figure 5 shows an influence diagram that describes a decision situation building on the belief network from Figure 2. This influence diagram includes a decision variable (Load Allocation/Mitigation Decisions), new chance variables for mercury methylation potential, total mercury in water ($MeHg_T$), mercury fish tissue levels, annual total mercury load to the bay, mitigation cost, credibility of compliance (discussed below) for methyl mercury levels in fish, and, credibility of compliance for total mercury load to the bay. It also includes a multiattribute utility node (value node) defined in terms of mitigation cost, credibility of compliance with a mercury fish tissue target, and credibility of compliance with the annual total load to the bay target. By eliciting decision maker (sub-group) preferences over outcomes in terms of these three attributes with multiattribute utility analysis, optimal load allocation/mitigation decisions can be determined for the sub-group using this model.

The influence diagram in Figure 5 states that given total mercury concentration in water and the river flow, the annual total mercury load to the south bay is independent of the mine and creek loads. Precipitation and creek flow are modeled with marginal (unconditional) distributions based on the available historical data. The mine mercury load is modeled as being conditional upon precipitation and the mine site mercury load allocation. The creek mercury load is modeled as being conditional upon creek flow and the creek mercury load allocation. To further illustrate the concept of conditional independence, note that *given observations* for Hg_T in water and river flow (e.g., annual average values over the waterbody), the annual total mercury load to the south bay is conditionally independent of the creek Hg load and mine site Hg load. This does *not* mean that creek Hg load and mine site Hg load do not impact the total annual Hg load to the bay, but rather that the influence is through the Hg_T in water variable. From a causal

Figure 5 - Influence Diagram for Mercury Load Allocation/Mitigation Decisions for a Small Watershed Impacted by a Mercury Mine Site and Mine Wastes.

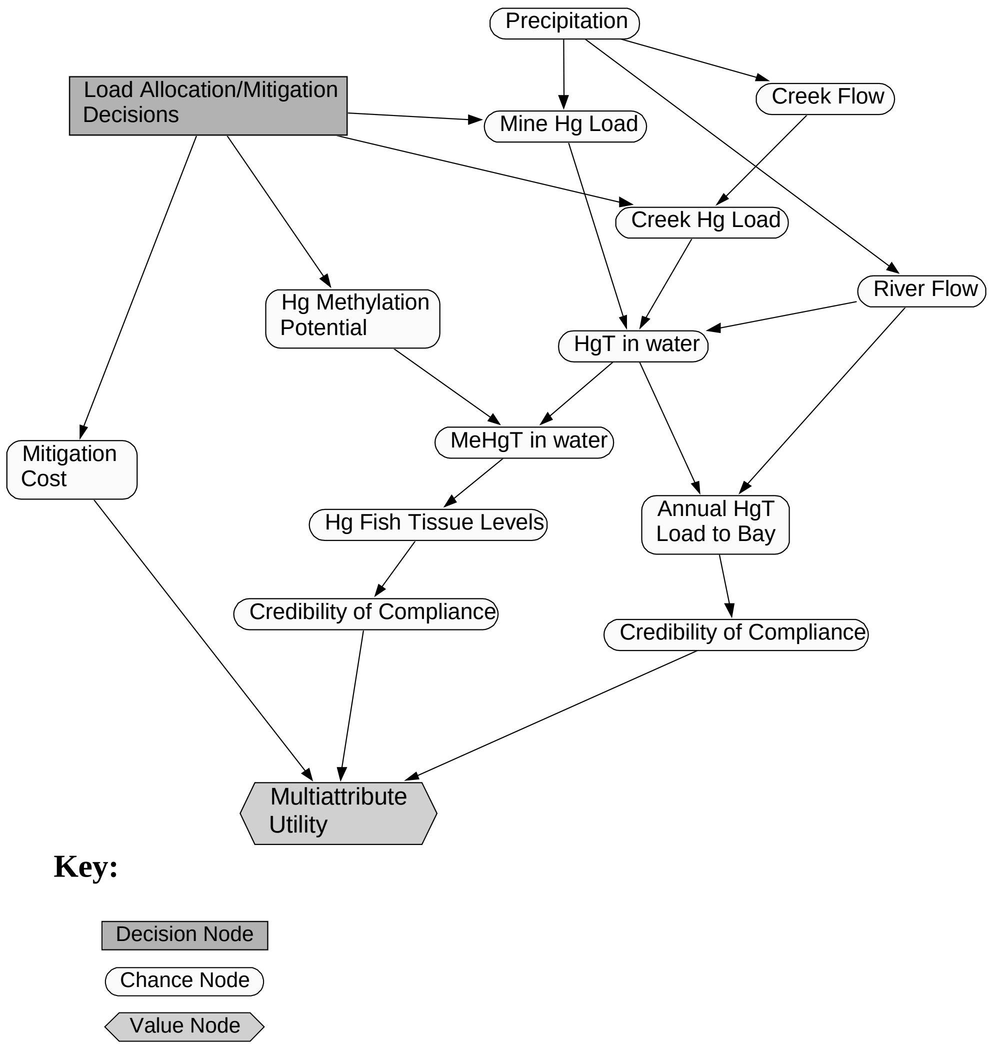

perspective, the observed Hg_T in water value would "reflect" any influence from the mine site and creek Hg loadings, which is why the loadings become irrelevant upon observation of Hg_T. When Hg_T in water is not observed, the mine site Hg load and creek Hg load variables are relevant to the annual total Hg load to the bay through their collective influence on the Hg_T in water variable.

Assumptions about spatial and temporal averaging are built into the Bayesian network, as appropriate to the particular decision problem. For this simple example, precipitation and creek flow probability distributions represent the available data over annual cycles and total mercury in water (Hg_T in water) refers to annual average concentration over the waterbody. These assumptions were made to keep the number of variables manageable for illustration.

The model in Figure 5 represents that fact that decision makers have influence over the natural system, even though the decision outcomes are uncertain. Control is represented in the graph by arrows from the decision node to the Mine Hg Load, Creek Hg Node, and Hg Methylation Potential nodes. In an influence diagram, if the parent node is a decision variable, the probability distribution for the uncertain variable child is conditioned on the decision made. This control can be thought of in causal terms as the ability of the decision maker to require mitigation actions that reduce mercury loadings or alter environmental factors such that mercury methylation potential should be reduced (e.g., river or creek shading, reservoir aeration).

To understand how that influence over loading and methylation potential propagates through the other variables in the network, ultimately influencing value (multiattribute utility), we must understand how information "flows through" the network. For example, the Creek Hg Load node is the parent of Hg_T in water node. When the parent of a chance node is another chance node, the child's probability distribution is conditioned on the state of the parent chance node, which may either have an observed value or may itself be represented by a conditional probability distribution. In this example, the load allocation/mitigation decision alters the Creek Hg Load conditional probability distribution, which in turn alters the Hg_T in water conditional probability distribution. In this manner, the uncertain impacts from the chosen allocation/mitigation strategy propagate through the network, influencing the conditional probability distributions for the attributes of interest to decision makers. A chance node with no parents is described by an unconditional (or marginal) probability distribution, typically created from historical data (e.g., precipitation).

Bayesian networks can accommodate a mixture of continuous and discrete probability distributions for uncertain variables. In special cases decision variables can be continuous (e.g., Gaussian influence diagrams), but in general decision variables are discrete. In the implementation for the influence diagram shown in Figure 5, precipitation, creek flow, mine and creek loads, total mercury concentration in water, annual mercury load to the south bay, and mitigation costs are continuous. Mercury fish tissue levels are modeled as discrete, given the high uncertainty.

Optimal Decisions and Sensitivity Analysis Using Influence Diagrams Without a Value Model

To illustrate how influence diagrams can be used to perform decision analysis *without a value model*, Figure 6 simulates predictions for the hypothetical strategy focusing on reducing mercury methylation potential ("MeHg Potential Mitigation Focus") from the strategy table shown in Figure 4. For this strategy, "medium reductions" are chosen for mine site and creek Hg load reductions and a "high reduction" is chosen for MeHg potential reduction. This simplified hypothetical model predicts mine Hg load, creek Hg load, Hg_T in water, $MeHg_T$ in water, Hg fish tissue levels, the annual total mercury load to the south bay, and credibility of compliance measures for fish tissue levels and the load to the bay. For this strategy, the predicted credibility of compliance with mercury fish

Figure 6 – Example Predictions for "MeHg Potential Mitigation Focus" Strategy.

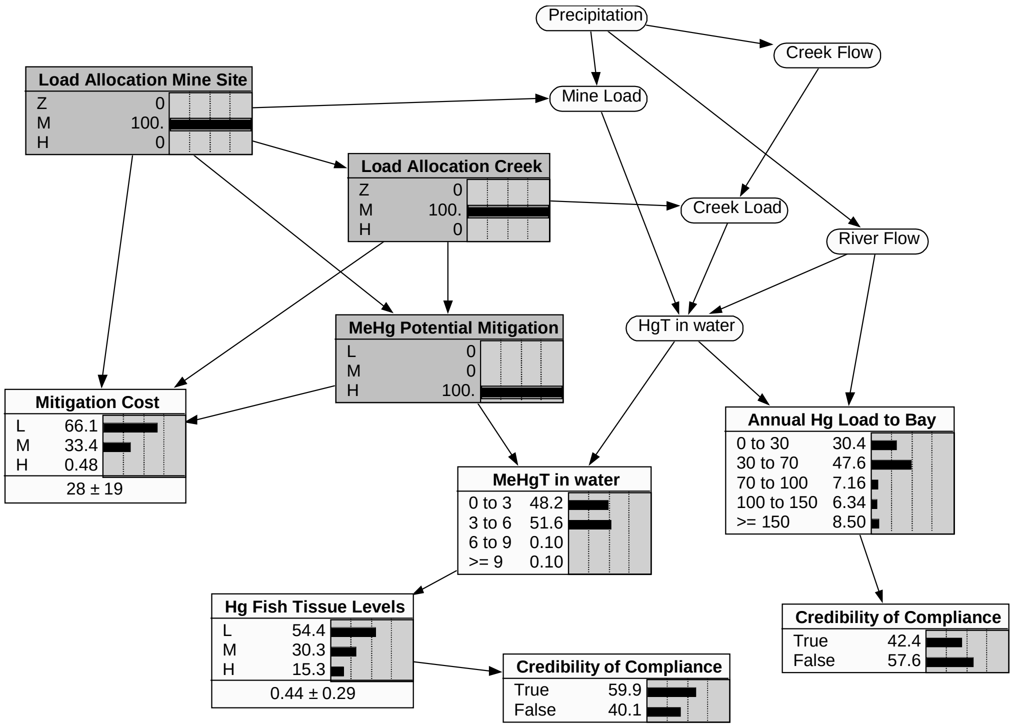

tissue targets is around 60% and the predicted credibility of compliance with the annual Hg load to the south bay is around 42%. The average predicted cost for this strategy is 28 with a standard error of 19, where the units are arbitrary (e.g., $10,000).

For a comparison, Figure 7 simulates predictions for the other hypothetical strategy from the strategy table, "Mine Site Load Reduction Focus", in which a "high reduction" is chosen for the mine site Hg load and "low reductions" are chosen for the creek Hg load and MeHg potential. For this strategy, the predicted credibility of compliance with mercury fish tissue targets is around 50% and the predicted credibility of compliance with the annual Hg load to the south bay is around 40%. The average predicted cost for this strategy is 48 with a standard error of 20 in the same arbitrary units. In this simple example, the "MeHg Potential Mitigation Focus" strategy is clearly superior in terms of predicted credibility of compliance for both endpoints (fish tissue levels and annual load to the bay) and mitigation cost. The first question that arises at this point is how robust is this conclusion? Another question that arises is, what would happen if the results were "mixed", in the sense that one strategy was superior in terms of one attribute and the another was superior in terms of another attribute? In most real world cases, "mixed results" would be anticipated. The first question may be addressed with sensitivity analysis and the second with multiattribute utility analysis, which will be explored next.

Figure 7 – Example Predictions for "Mine Site Load Reduction Focus" Strategy.

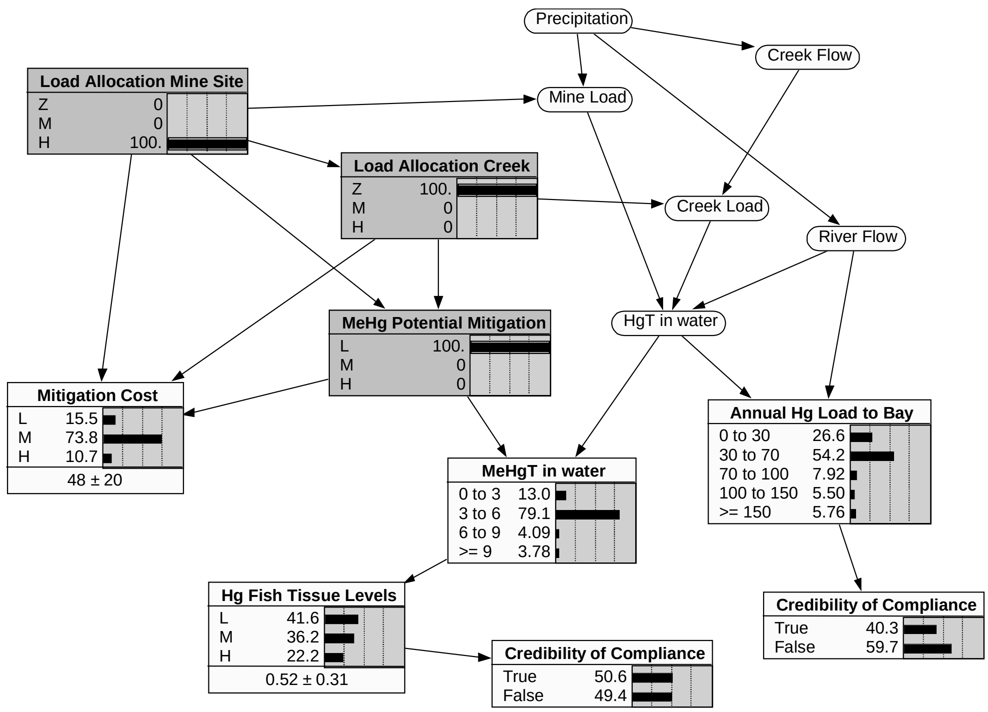

Sensitivity Analysis in Decision Analysis Using Bayesian Networks

Sensitivity analysis within the framework of influence diagrams and decision analysis has several meanings and purposes. In general, the idea is to analyze how sensitive conclusions are to the various pieces that make up the model. In the context of influence diagrams, sensitivity analysis refers to analyzing how sensitive conclusions (probabilities of interest or expected utility) are to small changes in the conditional probabilities that influence those conclusions (Jensen, 2001). Sensitivity analysis may be used, for example, to "tweak" the probability distributions within the network to meet constraints imposed by expert judgment or observations. In this paper, we will focus on some aspects of sensitivity analysis dealing with the robustness of conclusions in the context of influence diagrams describing decision situations. For more details, see Nielsen and Jensen (2003), Laskey (1995), Jensen (2002), and Castillo et al. (1997).

Table 2 illustrates an analysis of the sensitivity of credibility of compliance for the total mercury load to the bay to small changes in the conditional distributions for mine Hg load for the "Mine Site Load Reduction Focus" strategy. Note that while the numbers are based on actual output from the model in Figure 7 implemented in Netica, the underlying distributions are fictitious. The sensitivity analysis output shows that the credibility of compliance for total mercury to the bay ranges from around 9% to 44% for changes to mine Hg load, where the current value is around 40%. "Quadratic scoring" and "Entropy

Table 2. Sensitivity of "Credibility of Compliance for Total Mercury Load to Bay" to Changes to Mine Hg Load for the Mine Site Load Reduction Focus Strategy

Example Output from Netica				
Probability ranges	Min Value	Current Value	Max Value	RMS[1] Change
In Compliance	0.09365	0.4028	0.4394	0.09382
Out of Compliance	0.5606	0.5972	0.9064	0.09382
Quadratic scoring = 0.008803				
Entropy reduction = 0.005048 (0.519 %)				

[1] Root Mean Square

reduction" refer to scoring rules that summarize how sensitive credibility of compliance is to mine site load (Jensen, 2001). These scorings can be used to rank sensitivity of a particular attribute to the variables in the model, allowing the work group to focus attention on those variables that contribute the most uncertainty to conclusions. This information could be used to support, for example, information gathering activities and prioritization of technical review.

The above represents "one-way" sensitivity analysis, in which sensitivity to individual attributes can be explored. "Two-way" analysis can also be performed, in which the sensitivity of an attribute of interest is computed by varying two variables simultaneously. "Three-way", etc., sensitivity analysis can be performed, but the computational burden grows exponentially and quickly becomes burdensome. Other methods of sensitivity analysis can be performed, including the conversion of the influence diagram into a decision tree by discretizing (if necessary) the probability distributions in the network, then performing probabilistic sensitivity analysis using the decision tree. There are many possibilities (see, e.g., Morgan and Henrion, 1990 and Clemen, 1996) and a lot can be learned about the decision problem using relatively simple methods.

Optimal Decisions and Sensitivity Analysis Using Influence Diagrams With a Multiattribute Utility Value Model

Figure 8 shows an influence diagram model with a multiattribute utility function used to determine optimal decisions, using utility values similar to those in Table 1 for some hypothetical sub-group. For this simple example, a strategy with a "high reduction" for the mine site Hg load, a "medium reduction" for creek Hg load, and a "medium reduction" for mercury methylation potential is optimal in terms of maximizing expected utility. All of the other strategies yield lower expected utilities. Sensitivity analysis similar to that presented in the previous section could be used to determine the sensitivity of this conclusion to the various uncertainties (e.g., mine site Hg load and creek Hg load). Value of information could then be determined from this sensitivity analysis, enabling a

Figure 8 – Example Optimal Load Allocation/Mitigation Strategy Using an Influence Diagram with a Multiattribute Utility Model.

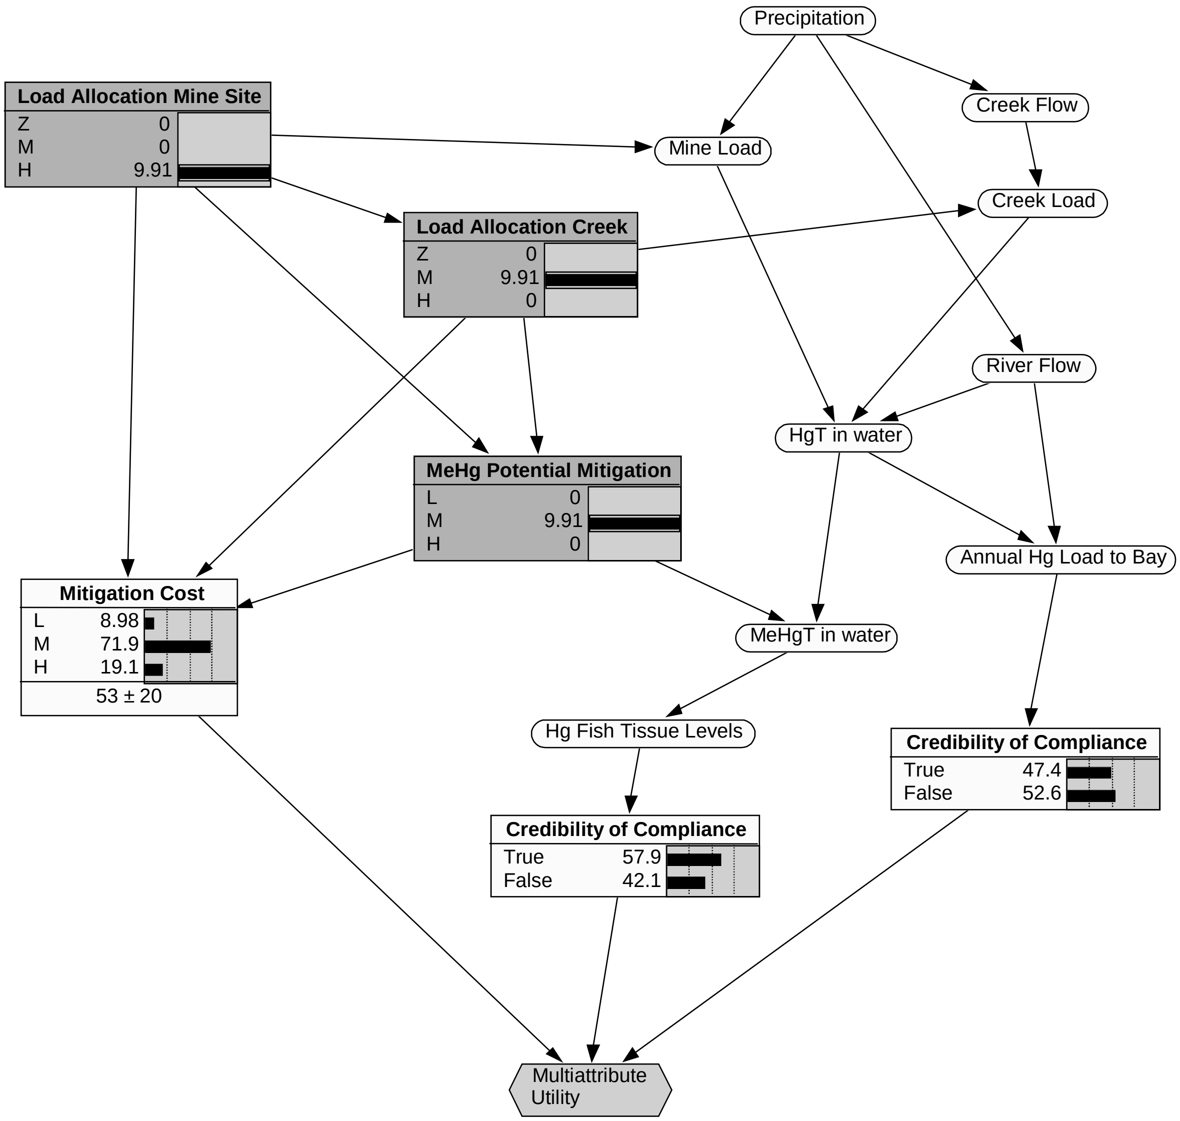

ranking of the uncertainties in terms of importance from the point of view of the preferences of the sub-group.

CONCLUSIONS

This paper describes a decision analysis approach to TMDL load allocation decisions using the Guadalupe River mercury TMDL as an example. Decision analysis is a rigorous and robust common sense approach that, in many circumstances, is an attractive alternative to other decision analytical tools like cost/benefit analysis and what-if analysis. Decision analysis makes use of approaches for eliciting and representing preferences that makes it capable of accounting for non-monetary concerns, which is an appealing characteristic in environmental decision making. While decision analysis does require active involvement of decision makers relative to many other decision making approaches, one could argue that this fact is responsible for much of the power of the decision analysis process. When decision analysis is properly performed, decision makers (or sub-groups) should *believe* the insights, given that the expertise and knowledge represented in the model should reflect trusted information and that the preferences expressed should be their own. While the application of decision analysis in group decision making situations can be problematic, since individual group members may have significantly different beliefs and preferences that cannot be simultaneously modeled, decision analysis can be used to generate sub-group negotiating positions and can shed light on the sources of disagreement (Merkhofer, 1999).

The various decision analysis tools, including objectives hierarchies, strategy tables, influence diagrams, and decision trees, can be very useful aids for communicating, eliciting knowledge and preferences, organizing a complex decision situation, and generating insights that can highlight sources of disagreement and areas of agreement. When properly applied, decision analysis can help decision makers make better decisions in terms of the consideration of uncertainty and value.

The approach highlighted in this paper makes extensive use of Bayesian networks for forecasting the response of the natural system to TMDL load allocations. As shown by Borsuk *et al.* (2001, 2002), Reckhow (1999), and Stow *et al.* (2003), Bayesian network models of water quality and ecological response are competitive with complex mechanistic models in terms of goodness-of-fit statistics and other indications of forecasting ability. They are superior in terms of model updating, since the Bayesian nature of the network allows new monitoring information to be incorporated directly into the existing network, generating an updated model that integrates the old and new information using robust probability calculus. By using a Bayesian network as the basis of the decision analysis (i.e., for more than forecasting water quality and ecological response), the potential for consensus on allocation decisions can be explored, sources of differences can be analyzed for potential compromise, and, at the very least, negotiating positions for sub-groups of stakeholders can be rigorously defined in terms of information and preferences. In addition, sensitivity analysis can be performed using the Bayesian network to inform information gathering priorities and peer review activities.

ACKNOWLEDGEMENTS

This research project is sponsored by the National Science Foundation, Grant Number EIA-0085998. The authors would like to thank Professors Kincho Law and Ross Shachter from Stanford University for their valuable suggestions in this project. The authors would also like to thank Carrie Austin and Richard Looker from the San Francisco Bay Regional Water Quality Control Board for their helpful comments on this paper.

REFERENCES

Abraham, B. and Ledolter, J., 1983. *Statistical Methods for Forecasting*. New York: John Wiley & Sons.

Anderson, R.M. and Hobbs, B.F., 2001. Using Decision Analysis to Choose Phosphorus Targets for Lake Erie. *Environmental Management*, 27(2): 235-252.

Boese, G.W., 2002. *Approaches to the Allocation of Loads in TMDLs*. National TMDL Science and Policy Specialty Conference, November 13 – 16, 2002. Phoenix, Arizona.

Bonano, E.J., Apostolakis, G.E., Salter, P.F., Ghassemi, A., and Jennings, S., 2000. Application of Risk Assessment and Decision Analysis to the Evaluation, Ranking and Selection of Environmental Remediation Alternatives. *Journal of Hazardous Materials*, 71(1-3), SI, 35-57.

Borsuk, M., Clemen, R., Maguire, L., and Reckhow, K., 2001. Stakeholder Values and Scientific Modeling in the Neuse River Watershed. *Group Decision and Negotiation*,10: 355-73.

Borsuk, M., Stow, C.A., and Reckhow, K., 2002. Integrative environmental prediction using Bayesian networks: A synthesis of models describing estuarine eutrophication. *IEMSS Conference proceedings*. Lugano, Switzerland. Available from the World Wide Web: (http://www2.ncsu.edu/ncsu/CIL/WRRI/BayesIEP.pdf)

Burn, D.H. and Lence, B.J., 1992. Comparison of Optimization Formulations for Waste-Load Allocations. *Journal of Environmental Engineering*, 118(4): 597-612.

Castillo, E., Gutierrez, J., and Hadi, A., 1997. *Expert Systems and Probabilistic Network Models*. New York: Springer-Verlag.

Chechile, R.A., 1991. Probability, Utility, and Decision Trees in Environmental Decision Analysis. In *Environmental Decision Making: A Multidisciplinary Perspective,* R.A. Chechile and S. Carlisle, eds. New York: Van Nostrand Reinhold.

Churchman, C.W., 1968. *The Systems Approach*. New York: Dell.

Clemen, R.T., 1996. *Making Hard Decisions: An Introduction to Decision Analysis, 2nd edition*. Pacific Grove, CA: Brooks/Cole Publishing Co.

Domagalski, J.L., Slotton, D.G., Alpers, C.N., Suchanek, T.H., Churchill, R., Bloom, N., Ayers, S.M., and Clinkenbeard, J., 2002. *Summary and Synthesis of Mercury Studies in the Cache Creek Watershed, California, 2000-2001.* Portion of the Final Report to the CALFED Bay-Delta Program. August 2002 draft.

Edwards, W., 1977. Use of Multiattribute Utility Measurement for Social Decision Making. In *Conflicting Objectives in Decisions*, D.E. Bell, R.L. Keeney, and H. Raiffa, eds. New York: John Wiley & Sons.

Freeze, RA and Gorelick, SM, 1999. Convergence of Stochastic Optimization and Decision Analysis in the Engineering Design of Aquifer Remediation. *Ground Water*, 37(6): 934-954.

Gregory, R., 1999. Identifying Environmental Values. In *Tools to Aid Environmental Decision Making*, V. Dale and M. English, eds. pp. 231-281. New York: Springer.

Haith, D.A., 2003. Systems Analysis, TMDLs, and Watershed Approach. *Journal of Water Resources Planning and Management*. July/August 2003: 257-260.

Heckerman, D. and Shachter, R.D., 1995. Decision-Theoretic Foundations for Causal Reasoning. *Journal of Artificial Intelligence Research,* 3: 405-430.

Howard, R.A., 1968. The Foundations of Decision Analysis. *IEEE Transactions on Systems Science and Cybernetics*, SSC-4(3): 581-589.

Howard, R.A., 1984a. The Evolution of Decision Analysis. In *Readings on the Principles and Applications of Decision Analysis*, R.A. Howard and J.E. Matheson, eds. Menlo Park, California: Strategic Decisions Group.

Howard, R.A., 1984b. Risk Preference. In *The Principles and Applications of Decision Analysis*, R.A. Howard and J.E. Matheson, eds. Menlo Park, California: Strategic Decisions Group.

Howard, R.A., 1988. Decision Analysis: Practice and Promise. *Management Science*, 34(6): 679-695.

Howard, R.A., 1990. From Influence to Relevance to Knowledge. In *Influence Diagrams, Belief Nets, and Decision Analysis*, R.M. Oliver and J.Q. Smith, eds. Chichester: Wiley.

Howard, R.A. and Matheson, J.E., 1984. Influence Diagrams. In R.A. Howard and J.E. Matheson, eds., *The Principles and Applications of Decision Analysis*. Menlo Park, CA: Strategic Decisions Group.

Jensen, Finn V., 2001. *Bayesian Networks and Decision Graphs*. New York: Springer-Verlag.

Keeney, R.L., 1980. *Siting Energy Facilities*. New York: Academic Press.

Keeney, R.L., 1992. *Value Focused Thinking*. Cambridge, MA: Harvard University Press.

Keeney, R.L. and Raiffa, H., 1976. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. New York: John Wiley & Sons.

Kruber, S. and Schoene, H., 1998. The Application of Decision Analysis in the Remediation Sector. *Environmental Engineering and Policy*, 1: 25-35.

Laskey, K.B., 1995. Sensitivity Analysis for Probability Assessments in Bayesian Networks. *IEEE Transactions on Systems, Man, and Cybernetics*, 25(6): 901-909.

Lawrence, D.B., 1999. *The Economic Value of Information*. New York: Springer-Verlag, Inc.

Lung, W.-S., 2001. *Water Quality Modeling for Wasteload Allocations and TMDLs.* New York: John Wiley & Sons, Inc.

MacCrimmon, K.R. and Wehrung, D.A., 1977. Trade-off Analysis: The Indifference and Preferred Proportions Approach. In *Conflicting Objectives in Decisions*, D.E. Bell, R.L. Keeney, and H. Raiffa, eds. New York: John Wiley & Sons.

Maguire, L.A., and Boiney, L.G., 1994. Resolving Environmental Disputes - A Framework Incorporating Decision-Analysis and Dispute Resolution Techniques. *Journal of Environmental Management*, 42(1): 31-48.

Marshall, K.T. and Oliver, R.M., 1995. *Decision Making and Forecasting*. New York: McGraw-Hill, Inc.

Merkhofer, M.W., 1990. Using Influence Diagrams in Multiattribute Utility Analysis – Improving Effectiveness Through Improving Communication. In R.M. Oliver and J.Q. Smith, eds., *Influence Diagrams, Belief Nets, and Decision Analysis*. pp. 297-317. Chichester: Wiley.

Merkhofer, M., 1999. Assessment, Refinement, and Narrowing of Options. In V. Dale and M. English, eds., *Tools to Aid Environmental Decision Making*. pp. 231-281. New York: Springer.

Merkhofer, M.W., Conway, R., and Anderson, R.G., 1997. Multi-Attribute Utility Analysis as a Framework for Public Participation in Siting a Hazardous Waste Management Facility. *Environmental Management*, 21(6): 831-839.

Merkhofer, M.W. and Keeney, R.L., 1987. A Multiattribute Utility Analysis of Alternative Sites for the disposal of nuclear waste. *Risk Analysis*, 7(2): 173-194.

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

Murphy, K., 2002. *How to Use the Bayes Net Toolbox*. Available on the World Wide Web: (http://www.ai.mit.edu/~murphyk/Software/BNT/usage.html).

National Research Council (NRC), 2001. *Assessing the TMDL Approach to Water Quality Management.* Washington, D.C.: National Academy Press.

Neilsen, T.D. and Jensen, F.V., 2003. Sensitivity Analysis in Influence Diagrams. *IEEE Transactions on Systems, Man, and Cybernetics—Part : Systems and Humans*, 33(1): 223-234.

Norwys Software Corp., 1997. Netica™: Application for Belief Networks and Influence Diagrams, User's Guide, Version 1.05 for Windows®. Vancouver, Canada: Norwys Software Corp.

Oliver, R.M. and Smith, J.Q.,1990. Preface. In R.M. Oliver and J.Q. Smith, eds., *Influence Diagrams, Belief Nets, and Decision Analysis*. pp. xv-xxiii. Chichester: Wiley.

Ortolano, L., 1997. *Environmental Regulation and Impact Assessment*. New York: John Wiley & Sons, Inc.

Pearl, J., 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Series in Representation and Reasoning. San Francisco, CA: Morgan Kaufmann Publishers.

Pearl, J., Geiger, D., and Verma, T., 1990. The Logic of Influence Diagrams. In R.M. Oliver and J.Q. Smith, eds., *Influence Diagrams, Belief Nets, and Decision Analysis*. pp. 67-87. Chichester: Wiley.

Perdek, J.M., 1997. Decision Analysis of Incineration Costs in Superfund Site Remediation. *Environmental Engineering Science*, 14(1): 15-26.

Prato, T., 2003. Multiple-Attribute Evaluation of Ecosystem Management for the Missouri River System. *Ecological Economics*, 45: 297-309.

Reckhow, K.H., 1994a. A Decision Analytic Framework for Environmental Analysis and Simulation Modeling. *Environmental Toxicology and Chemistry*, 13(12): 1901-1906.

Reckhow, K.H., 1994b. Water Quality Simulation Modeling and Uncertainty Analysis for Risk Assessment and Decision Making. *Ecological Modelling*, 72: 1-20.

Reckhow, K.H., 1999. Water Quality Prediction and Probability Network Models. *Canadian Journal of Fisheries and Aquatic Sciences*, 56:1150-1158.

Reckhow, K. H. 2002. (in press) Bayesian Approaches in Ecological Analysis and Modeling . In *The Role of Models in Ecosystem Science*. Charles D. Canham, Jonathan J. Cole, and William K. Lauenroth, eds. Princeton, NJ: Princeton University Press.

Reckhow, K.H., Borsuk, M.E., and Stow, C.A., 2002. *Adaptive Implementation of TMDLs Using Bayesian Analysis*. National TMDL Science and Policy Specialty Conference, November 13 – 16, 2002. Phoenix, Arizona.

RWQCB (California Regional Water Quality Control Board, San Francisco Bay Region), 2003. Mercury in San Francisco Bay, Total Maximum Daily Load (TMDL) Project Report. June 6, 2003. Available on the World Wide Web: (http://www.swrcb.ca.gov/rwqcb2/TMDL/SFBayMercury/SFBayMercuryTMDLProjectReport.pdf).

Shachter, R.D., 1986. Evaluating Influence Diagrams. *Operations Research*, 34(6): 871-882.

Shachter, R.D., 1988. Probabilistic Inference and Influence Diagrams. *Operations Research*, 36(4): 589-604.

Stow, C., Roessler, C., Borsuk, M., Bowen, J., and Reckhow, K., 2003. Comparison of Estuarine Water Quality Models for Total Maximum Daily Load Development in Neuse River Estuary. *Journal of Water Resources Planning and Management*, 129(4): 307-314.

Thomann, R.V., 1974. *Systems Analysis and Water Quality Management*. New York: McGraw-Hill.

USEPA, 2002. The Twenty Needs Report: How Research Can Improve the TMDL Program. EPA841-B-02-002. Washington, DC: Office of Wetlands, Oceans and Watersheds. Available from the World Wide Web: (http://www.epa.gov/owow/tmdl/20needsreport_8-02.pdf)

Varis, O., 1995. Belief Networks for Modelling and Assessment of Environmental Change. *Environmetrics*, 6: 439-444.