



Research

Many-objective robust decision making for managing an ecosystem with a deeply uncertain threshold response

Riddhi Singh^{1,2}, Patrick M. Reed³ and Klaus Keller^{2,4,5}

ABSTRACT. Managing ecosystems with deeply uncertain threshold responses and multiple decision makers poses nontrivial decision analytical challenges. The problem is imbued with deep uncertainties because decision makers do not know or cannot converge on a single probability density function for each key parameter, a perfect model structure, or a single adequate objective. The existing literature on managing multistate ecosystems has generally followed a normative decision-making approach based on expected utility maximization (MEU). This approach has simple and intuitive axiomatic foundations, but faces at least two limitations. First, a prespecified utility function is often unable to capture the preferences of diverse decision makers. Second, decision makers' preferences depart from MEU in the presence of deep uncertainty. Here, we introduce a framework that allows decision makers to pose multiple objectives, explore the trade-offs between potentially conflicting preferences of diverse decision makers, and to identify strategies that are robust to deep uncertainties. The framework, referred to as many-objective robust decision making (MORDM), employs multiobjective evolutionary search to identify trade-offs between strategies, re-evaluates their performance under deep uncertainty, and uses interactive visual analytics to support the selection of robust management strategies. We demonstrate MORDM on a stylized decision problem posed by the management of a lake in which surpassing a pollution threshold causes eutrophication. Our results illustrate how framing the lake problem in terms of MEU can fail to represent key trade-offs between phosphorus levels in the lake and expected economic benefits. Moreover, the MEU strategy deteriorates severely in performance for all objectives under deep uncertainties. Alternatively, the MORDM framework enables the discovery of strategies that balance multiple preferences and perform well under deep uncertainty. This decision analytic framework allows the decision makers to select strategies with a better understanding of their expected trade-offs (traditional uncertainty) as well as their robustness (deep uncertainty).

Key Words: *a posteriori* decision making; deep uncertainty; lake eutrophication; many objective; robustness analysis; utility

INTRODUCTION

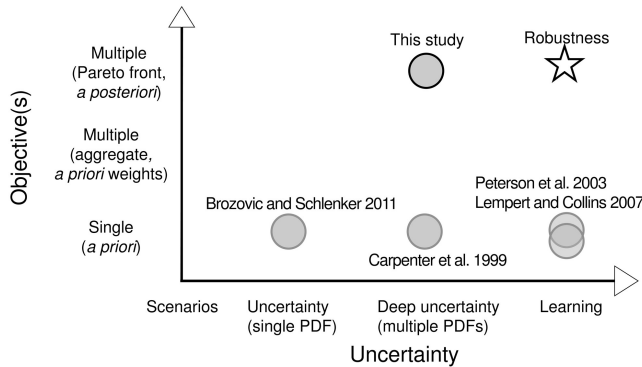
Managing ecosystems poses challenging decision analysis problems because of (1) the presence of several stakeholders with conflicting objectives, (2) the potential for highly nonlinear threshold responses, and (3) the underlying deep uncertainties. Deep uncertainties emerge when planners are unable to agree on or identify the full scope of possible future events, including their associated probability distributions (Knight 1921, Lempert et al. 2006). In cases of conflicting objectives of stakeholders, explicit trade-offs between alternative management actions exist. Mapping these trade-offs poses nontrivial conceptual and computational challenges when defining ecosystem management problems, e.g., choosing objectives, management decisions, planning horizons, etc., as well as predicting the impacts of actions, e.g., imperfect knowledge of system dynamics, external forcings, or environmental thresholds (Clark 2005, Lempert and Collins 2007, Keller and McInerney 2008, McInerney et al. 2012). Representing the potential for highly nonlinear threshold responses, or tipping points, can be crucial because surpassing such thresholds can lead to regime shifts that significantly degrade the objectives of some or all stakeholders (Bennett et al. 2008). Current decision support tools for ecosystem management are limited in their ability to address these interacting challenges (Lempert and Collins 2007). Throughout this study, we use the term “stakeholder” to refer to the diverse set of stakeholders in a decision-making process.

Figure 1 classifies two main challenges that emerge when defining ecosystem management problems: (1) appropriately accounting for uncertainty, and (2) accounting for diverse stakeholder objectives. Approaches for representing uncertainty are arranged on the x-axis with increasing conceptual and computational complexity: beginning from deterministic, i.e., perfect foresight, assumptions, moving toward well-characterized uncertainty captured using a single probability distribution function (PDF), next transitioning to deep uncertainties represented with multiple PDFs, and finally, learning based on new observations to actively reduce uncertainties. The y-axis arranges methods for incorporating stakeholders' objectives in increasing complexity: starting from the evaluation of management actions using a single a priori abstraction of preferences in a utility function (Ramsey 1928, von Neumann and Morgenstern 1944), moving toward weighting schemes aggregating multiple objectives into a single measure as typified by traditional multicriteria decision analysis (Brink 1994, Ralph 2012, Köksalan et al. 2013), and finally, many-objective analysis in which stakeholders' objectives are elicited after an explicit mapping of trade-offs between alternatives.

Prior decision analytical frameworks have been mapped onto the uncertainty and objective axes based on their underlying approaches when addressing the lake problem as illustrated in Figure 1. The lake problem has a long history and represents a hypothetical lake with a nonlinear uncertain threshold response. The lake problem seeks to proxy a broad class of environmental

¹Department of Civil Engineering, Indian Institute of Technology Hyderabad, ²Earth and Environmental Systems Institute, Pennsylvania State University, ³School of Civil and Environmental Engineering, Cornell University, ⁴Department of Geosciences, Pennsylvania State University, ⁵Department of Engineering and Public Policy, Carnegie Mellon University

Fig. 1. The decision analytic approaches used in literature to identify management strategies for a multistate lake with deeply uncertain variables. Studies are classified on the basis of the number of objectives in their problem formulation (y-axis) and characterization of uncertainty (x-axis). The star represents the ideal location of a management strategy. PDF = probability distribution function.



management problems, such as fishery collapse and abrupt climate change (Carpenter et al. 1999, Brozovic and Schlenker 2011). The problem's narrative describes a lake near a town. The citizens have to identify a pollution strategy that meets conflicting objectives, such as maximizing economic gain while maintaining a healthy lake. Here, we build on the lake problem's legacy as a test bed for comparing and contrasting decision-making frameworks for environmental management. Through this example, we first show how standard approaches limit the stakeholders' ability to navigate the objective-uncertainty landscape. We then discuss avenues to overcome these limitations. Finally, we propose a decision-analysis framework that enables the stakeholders to tackle these challenges while contrasting it with the more traditional approaches.

Limitations of expected utility maximization framework in analyzing ecosystem management strategies

The ecosystem management literature in general, and analyses of the lake problem in particular, focus on the economic value represented through the maximization of expected utility (MEU) as the sole objective to evaluate alternative strategies (Carpenter et al. 1999, Maler et al. 2003, Bond 2010, Brozovic and Schlenker 2011, Horan et al. 2011, Johnson et al. 2013). These standard approaches position themselves in the lower left side of Figure 1. Despite its long legacy of applications, MEU faces at least two limitations when applied to ecosystem management problems. First, MEU evaluates alternative actions with an a priori aggregate utility function. This limits the ability to provide insights into the trade-offs between the often-divergent objectives of different stakeholders, e.g., maximizing net economic benefits versus ensuring sustainability, (Common and Perrings 1992, Ludwig 2000, Morse-Jones et al. 2011, Admiraal et al. 2013). Value functions are also insensitive to ecosystem attributes that do not affect them directly, e.g., loss of species, and can suggest strategies resulting in ecosystem collapse for nonlinear, threshold-based systems (Admiraal et al. 2013).

Secondly, the standard implementation of the MEU approach assumes that all uncertainties can be represented with a single joint probability density function, and stakeholders choose alternatives that maximize the expected value of the merit function across this PDF (Carpenter et al. 1999, Peterson et al. 2003, Bond and Loomis 2009). This is a poor assumption in the presence of deep uncertainties, which emerge when unique distributions of system parameters or forcings cannot be specified (Knight 1921, Lempert et al. 2006). Also, it has been shown that stakeholders are ambiguity averse. They prefer a profit of known probability instead of a profit of unknown probability (Ellsberg 1961). Hence, stakeholders do not always choose alternatives that maximize their expected utility. Aversion to ambiguity exists in many application areas, including health, environment, negotiation, and more (Becker and Brownson 1964, Curley and Yates 1985, Hogarth and Kunreuther 1989, Kuhn and Budescu 1996, Budescu et al. 2002).

Potential solutions to limitations of expected utility maximization framework

Moving from the single-objective formulation in the MEU framework to a multiobjective framework can provide insights about tensions and tradeoffs (Brill et al. 1990). If two or more objectives are in conflict, optimization for multiple objectives yields many nondominated solutions that comprise the Pareto front, or a trade-off curve. Nondomination implies that solution performance in any given objective cannot be improved without a performance loss in at least one other objective. This provides stakeholders with several alternatives instead of a single optimal strategy (Pareto 1896, Cohon and Marks 1975). An explicit understanding of the trade-offs across alternatives can significantly change stakeholders' preference of alternatives and broaden the scope of values encompassed in the planning process (Brill et al. 1990, Fleming et al. 2005, Kasprzyk et al. 2009, 2012, Woodruff et al. 2013, Zeff et al. 2014).

Identifying Pareto-optimal strategies poses significant computational challenges, which have severely limited the number of studies that move beyond the MEU approach to explore a richer set of objectives (McInerney et al. 2012, White et al. 2012). White et al. (2012) presented such a trade-off analysis for management of shared marine resources between sectors, such as offshore wind energy, commercial fishing, and whale watching. They showed that including several objectives in the analysis prevents significant losses (> \$1 million), generates extra value (> \$10 billion), provides managers with a means to incorporate sectors that cannot be measured in monetary units (conservation), increases transparency in the decision-making process, and helps avoid unnecessary conflicts caused by perceived but weak trade-offs. McInerney et al. (2012) illustrated an approach to mapping the trade-off curve defined by a transition of preferences between two values that enables stakeholders to choose their desired level of compromise a posteriori. Both studies generate trade-off curves by varying the weights on various objectives to identify one point on the Pareto front with each optimization (Chankong and Haimes 1983). However, this becomes computationally expensive, and arguably computationally infeasible, for three or more objectives because of the factorial growth in the number of subproblems solved when increasing objective counts (Teytaud 2007).

Beyond the challenge of quantifying key trade-offs, it is also critical to evaluate the solution's robustness to deep uncertainties (Lempert et al. 2003, Dixon et al. 2007, Lempert and Collins 2007). When selecting a preferred trade-off solution, stakeholders may prefer a strategy with reduced but consistent multiobjective performance across deeply uncertain states-of-the-world (SOWs) versus a strategy that is optimal in the expected SOW attained by assuming well-characterized uncertainty (Lempert and Collins 2007). Robustness analysis often employs minimization of average loss of utility (regret) under different SOWs representing deep uncertainty (Lempert et al. 2003, Dixon et al. 2007, Lempert and Collins 2007). The above cited robustness studies all adopt a univariate objective function and consistently show that single objective performance degrades in the presence of deep uncertainty. However, these studies are typically silent on the implicit degradation of other potential objectives, e.g., multiobjective regrets. Robustness analysis based on a single objective allows stakeholders to progress along the uncertainty axis on Figure 1 but limits movement along the objective axis. Defining robustness in a multiobjective context allows stakeholders to identify solutions that perform well across multiple objectives as well as under deep uncertainties, thus allowing them to traverse Figure 1 along both axes and reach the ideal location represented by the star.

MANY-OBJECTIVE ROBUST DECISION MAKING (MORDM): A FRAMEWORK FOR CONSTRUCTIVE DECISION AIDING

There is a growing recognition that decision analysis frameworks need to move toward a more flexible process of problem formulation wherein the stakeholders can add/remove objectives as needed, explore the consequences of each formulation, and identify novel alternatives that were not visible before (Brill et al. 1990, Tsoukias 2008, Kasprzyk et al. 2013, Woodruff et al. 2013). We demonstrate many-objective robust decision making (MORDM) as a potential framework to facilitate such a constructive decision to aid in the lake problem. Many-objective robust decision making combines the strengths of many-objective evolutionary search and robust decision making (Lempert and Collins 2007, Kasprzyk et al. 2013, Reed et al. 2013). Multiobjective evolutionary algorithm searches help the stakeholders to search the space of feasible strategies and discover alternatives that compose optimal trade-offs whereas robust decision making helps to test the performance of strategies under changing assumptions of uncertainty. Many-objective robust decision making also exploits recent technological advances in high-dimensional visual analytics (Thomas and Cook 2005, Kollat and Reed 2007, Keim et al. 2010) and multiobjective optimization. The framework employs the BORG multiobjective evolutionary algorithm (MOEA), which can provide high quality approximations of the trade-offs for problems with nonlinear, threshold-based dynamics across a range of objectives and representations of uncertainty (Coello et al. 2005, Fleming et al. 2005, Hadka and Reed 2012, Reed et al. 2013, Woodruff et al. 2013).

If uncertainties are well characterized, a posteriori trade-off analysis can be used to elicit which management action stakeholders prefer based on its expected performance under baseline uncertainty. However, when uncertainties are deep,

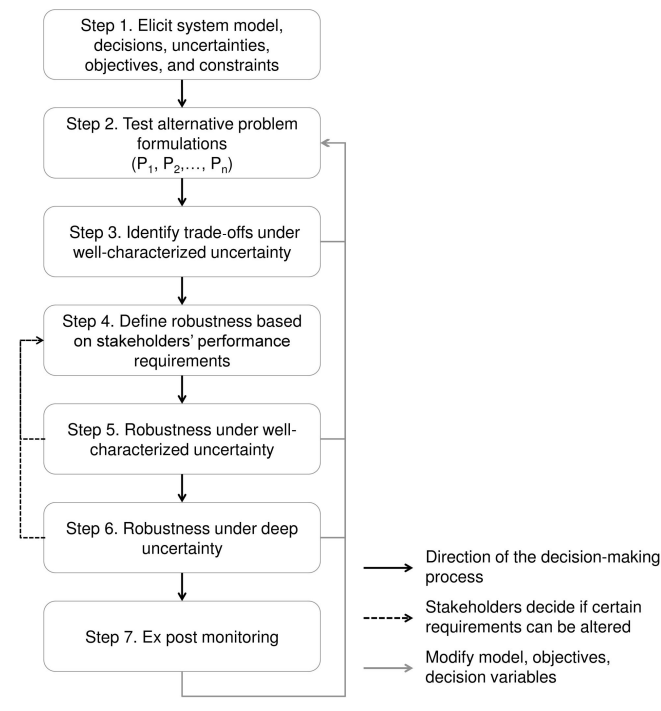
stakeholders might change their preferred management strategies if it is necessary to trade-off high performance under the baseline well-characterized uncertainty with satisficing multiobjective performance across broader ensemble of SOWs representing deep uncertainties (Simon 1959). We refer to Pareto satisficing strategies that perform acceptably across their component objectives as well as under deep uncertainties as robust. There is a third possibility that none of the strategies identified using the particular problem framing perform well under deep uncertainty, which may indicate the need to fundamentally alter the problem formulation. In this way, MORDM facilitates well-informed a posteriori trade-off decisions under well-characterized uncertainties, minimization of multiobjective performance losses under deep uncertainties, and problem falsification when no robust strategies can be identified. Other robustness approaches, such as info gap (Ben-Haim 2001) and decision scaling (Brown et al. 2012) represent local sensitivity analyses around a small set of user specified actions. They lack global multiobjective exploration of alternatives and have limited ability to quantify deeply uncertain trade-offs.

In engineering applications, MORDM has been found to be successful. For example, Kasprzyk et al. (2013) found that including robustness in a multiobjective context as a decision criterion dramatically altered the choice of an urban water supply planning problem's formulation, as well as its trade-off alternatives. Woodruff et al. (2013) discovered two families of aircraft designs using a ten-objective problem formulation that were not evident by using prior highly aggregated goal programming formulations. In a real multistakeholder application, Herman et al. (2014) identified key trade-offs and demand-based failure modes for four cities seeking to coordinate the mitigation of their water supply risks to drought. The recent growth in availability of open source software and comprehensive testing of MORDM in a range of engineering applications sets the stage for its use in a wider array of environmental problems (Haasnoot et al. 2013, 2014, Kwakkel et al. 2014). For example, it could be used for analyzing highly challenging environmental management problems, such as geoengineering, climate change adaptation, etc.

Summary of the framework

Many-objective robust decision making (MORDM) is typically implemented in an iterative approach (Fig. 2). The first step is to elicit stakeholders' objectives, potential management decisions, constraints, uncertainties, and system-planning models through surveys or interviews. The second step is to generate alternative problem formulations by experimenting with different objectives or their combinations, the decisions, and system uncertainty characterizations. Once alternative problem formulations are identified and baseline well-characterized uncertainties are specified, the potential trade-offs between different objectives can be assessed in step 3. At this step, the results are presented to the stakeholders. On viewing the trade-offs between objectives and understanding the compromise required in one objective to maintain a satisfactory performance in another, stakeholders can choose one or more alternatives. If none of the alternatives are acceptable, stakeholders may choose to alter the problem formulations as represented by the gray arrows between step 3 and step 2.

Fig. 2. Flowchart for implementing many-objective robust decision making (MORDM). Each box represents a step in the analysis. Solid arrows show the direction of the decision-making process. Gray arrows connect steps with the potential to alter the problem formulations with step 2. Dashed black arrows connect steps with the potential to alter stakeholders' performance requirements with step 4.



Following this, the robustness component of the framework requires an elicitation step in which the stakeholders identify the minimum levels of system performance that they are willing to accept. These measures can be related to modeled objectives or constraints as well as to other measures that were not included in the optimization process, e.g., new constraints or measures that emerge after viewing the trade-offs. Because performance requirements are elicited after the trade-off analysis, stakeholders have an understanding of the consequences of their preferred management strategies, i.e., a posteriori decision making. Building on this, robustness of the solutions with respect to stakeholders' performance requirements can be assessed in step 5. Carefully elicited measures of robustness provide stakeholders with an understanding of the fragility of management actions under well-characterized uncertainties. At this stage, if very few or none of the solutions are found to be robust, the stakeholders can be asked to either alter their performance thresholds, e.g., reduce their risk aversion and go back to step 4, or the problem formulations by adding/removing objectives/decision variables, i.e., go back to step 2.

Once robust solutions are identified under well-characterized uncertainty, they are tested under deep uncertainty scenarios, and their robustness is reassessed in step 6. At this stage, stakeholders can again go back to step 2 or step 4 if they are not satisfied with

the solution's performance and robustness under deep uncertainty. At the end of this process, the stakeholders should have identified a problem formulation and associated solution strategies that satisfy various objectives and are robust under both well-characterized and deep uncertainty. Finally, the strategy is implemented and the ecosystem is monitored for the impact of the selected strategy. In case of unexpected changes in stakeholders' objectives, the entire process can be repeated, i.e., the constructive decision aiding as discussed by Tsoukias (2008).

It is important to stress that visual analytics are a key component of the MORDM process, without which the goal of enabling stakeholders to explore diverse alternatives and consequences of their different choices would not be feasible (Stump et al. 2003, Kollat and Reed 2007, Castelletti et al. 2010, Reed et al. 2013, Woodruff et al. 2013). In the beginning of the decision-making process, stakeholders may only have a limited understanding of the complex interactions between their objectives. They may also have limited insights in the ranges of performance that are possible across different measures. Selecting candidate strategies after visualizing the trade-off curve avoids the biases associated with a priori weighting methods, which can hide key insights.

APPLICATION OF MANY-OBJECTIVE ROBUST DECISION MAKING (MORDM) TO THE LAKE PROBLEM

In this study, we used MORDM to address four main questions for the lake problem:

1. Does the incorporation of more planning objectives help to avoid eutrophic collapse of the lake?
2. What are the main conflicts in stakeholders' objectives across time and impacts on water quality?
3. How do the strategies attained under well-characterized uncertainty perform when re-evaluated for their robustness?
4. What are the implications of alternative stakeholders' preferences and definitions of robustness for managing the lake?

Our methods and results descriptions follow the MORDM steps illustrated in Figure 2. In the method description, we define the system model, associated uncertainties, etc. (step 1). Following this, three problem formulations, P1, P2, and P3, are defined (step 2). Optimal strategies for these multiple problem formulations are identified using the optimization framework described in the supplementary text "Evolutionary multiobjective optimization" (step 3; Appendix 1). We discuss the methods of estimating robustness in the subsection on assessing robustness (step 4).


In the results, we first present the baseline expected trade-offs in the result subsection "Trade-off analysis under well-characterized uncertainty" (step 3). Next, the subsection "Navigating trade-offs to explore diverse alternatives" illustrates how stakeholders can visually navigate the trade-off space and select solutions that represent competing preferences. We then use the simulated stakeholders' performance requirements to assess the robustness of solutions under well-characterized uncertainty (step 5), and then under deep uncertainty (step 6). The various choices made for each step in the lake problem are summarized in Appendix 2 (Table A2.1).

Step 1. Elicit system model, decisions, uncertainties, objectives, and constraints

The lake model

The lake problem is a hypothetical decision problem faced by a town located near a lake (Fig. 3). The model was developed by Carpenter et al. (1999) by analyzing the behavior of lakes. The inhabitants of the town have to decide on the amount of allowable pollution that can be emitted into the town's lake for a given planning horizon, i.e., they have to formulate a pollution control strategy. At each time step, the town releases a certain amount of anthropogenic pollution into the lake in the form of phosphorus. There is also some amount of uncontrolled natural inflow to the lake. The lake is able to remove a part of this pollution (phosphorus) based on its properties represented by parameters b and q . The phosphorus dynamics in the lake is approximated by:

Fig. 3. Schematic of the lake problem with variables highlighted in bold italic letters, see text for equations and variable descriptions.



$$X_{t+1} = X_t + a_t + \frac{X_t^q}{1 + X_t^q} - bX_t + \ln(\mu, \sigma) \quad (1)$$

where, X_t represents the (dimensionless) phosphorus in the lake at time t (in years); a_t (dimensionless) represents the allowed anthropogenic pollution in the lake at time t ; b and q are parameters of the lake model. Together, they represent the lake's ability to remove phosphorus through sedimentation, outflow, and sequestration in biomass of consumers or benthic plants, and recycle phosphorus, primarily from sediments. The last term represents natural inflows into the lake, which are outside the stakeholder's control. They are represented in this study by lognormal distributions with specified mean and standard deviation. The allowed anthropogenic pollution flow a_t is the only decision variable that can be altered to achieve various pollution strategies.

The lake problem formulation of Carpenter et al. (1999) is simple, yet conceptually rich. It has the ability to represent tipping points, nonlinearity, and deep uncertainties (Carpenter et al. 1999, Lempert and Collins 2007). Carpenter et al. (1999) demonstrated the impact of uncertainty, lags, etc. on the optimal strategy. In a groundbreaking study, Peterson et al. (2003) showed that stakeholders maximizing their expected utility can cause periodic collapses of the lake ecosystem if there is uncertainty about the lake's eutrophication thresholds. Lempert and Collins (2007) analyzed the decision problem to identify robust solutions that compromise optimality for acceptable performance under a

Table 1. Parameters related to the lake model, uncertainty characterization, reliability estimation, and stochastic sampling in BORG.

Category	Name	Parameter	Value	Dimensions
Lake model	Phosphorus removal rate	b	0.42	dimensionless
	Steepness factor	q	2	dimensionless
	Number of years	T	100	years
Utility estimation	Cost multiplier	α	0.4	dimensionless
	Damages multiplier	β	0.08	dimensionless
	Discount factor	δ	0.98	dimensionless
Uncertainty estimation	Number of stochastic samples per distribution	N	10000	dimensionless
Reliability estimation	Critical phosphorous level	X_{crit}	0.5	dimensionless
BORG algorithm	Stochastic optimization sampling frequency	-	100	dimensionless
	Discretization of decision	-	5	years
	Significant precision of decision	-	10^{-2}	dimensionless
				#%160;

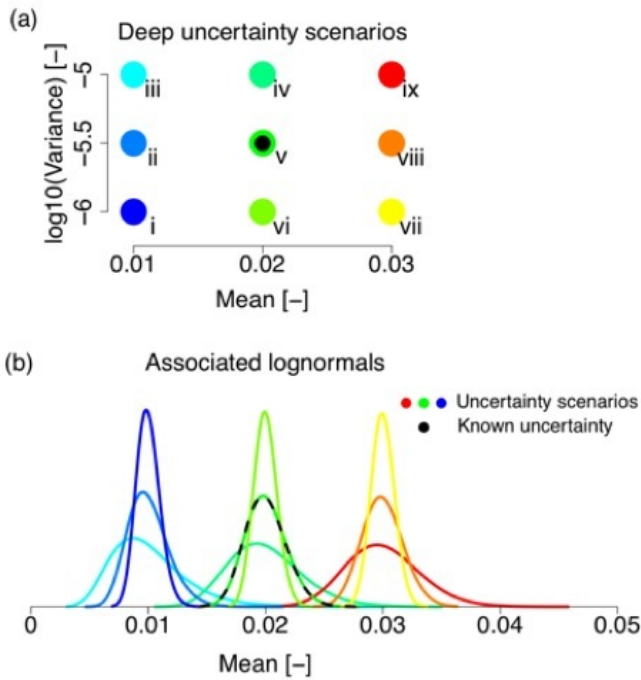
broader envelope of uncertainty assumptions. Refer to Appendix 3 for details on the dynamics of the lake model. Table 1 lists the parameter values related to the lake problem set up.

Representing key uncertainties

Uncertainty may arise in the lake model because of several sources, e.g., parameters governing the lake dynamics and economics, stochastic random inflows, identifying objective functions, etc. For our study, we focused on the uncertainty in the inflow of natural pollution into the lake. We analyzed the impact of two types of uncertainties on the selected strategy, well-characterized uncertainty and deep uncertainty. Uncertainty in random inflows to the lake is modeled stochastically using lognormal distributions (Equation 1). When the parameters of the lognormal distribution are prespecified, it represents the case of well-characterized uncertainty. On the other hand, if the stakeholders are uncertain about the characterization of the uncertainty, in this case the parameters of the lognormal distribution, they face a simple case of deep uncertainty.

Deep uncertainty is incorporated into the analysis to demonstrate the impact of a classic risk-based analysis vs. a multiple scenarios-based robustness analysis. The use of multiple probability distribution functions (PDFs) is similar to the approaches used by previous studies for characterizing deep uncertainty (Epstein and Wang 1994, Krishnamurti et al. 1999, Lempert and Collins 2007). We chose nine distributions for representing deep uncertainty using a standard sampling technique: by moving in increments above and below the baseline uncertainty distribution. The mean of the baseline uncertainty distribution was chosen such that it is half of the mean pollution inflow under the strategy obtained from single objective optimization. The variance of the baseline distribution was chosen such that the variance values that produced estimates of pollution between the next higher/lower

Fig. 4. (a) Parameters of the lognormal distribution used for generating deep-uncertainty scenarios. (b) Associated lognormal distributions. The lognormal curves in (b) are colored based on the associated circle in (a), which locates their mean and variance. The distribution shown by the black circle and the black dashed line in (a) and (b), respectively, is used to represent well-characterized uncertainty. Each lognormal distribution is used to generate time series of natural inflows to the lake. 10,000 such random time series are generated per lognormal distribution resulting in 10,000 states-of-the-world (SOWs) for well-characterized uncertainty, and 90,000 SOWs for deep uncertainty.



mean were selected. The lognormal distributions, which represent well-characterized uncertainty, or the baseline risk scenario, and deep uncertainty are illustrated in Figure 4. Corresponding parameters are listed in Table 2. Each PDF is used to generate 10,000 random SOWs resulting in the 90,000 total SOWs used in our study.

Objectives and constraints

A sole focus on economic valuation when identifying environmental management strategies can lead to a restrictive problem framing, which can limit the types of alternatives that stakeholders' explore (Brill et al. 1990). This concern has been termed cognitive myopia in the literature exploring biases in decision making (Hogarth and Kunreuther 1989, Brill et al. 1990). Recent examples of how cognitive myopia can negatively influence decisions have been published for water management as well as for the design of complex engineered systems (Kasprzyk et al. 2012, Woodruff et al. 2013). Many-objective robust decision making provides a broader trade-off context, as well as the potential to discover a more diverse suite of management policies to overcome cognitive myopia as suggested by Brill et al. (1990).

Table 2. Mean and standard deviation of the lognormal distributions used in the analysis of deep uncertainty.

Color	Mean	Standard deviation
Dark blue	0.01	10^{-6}
Blue	0.01	$10^{-5.5}$
Light blue	0.01	10^{-5}
Dark green	0.02	10^{-6}
Green	0.02	$10^{-5.5}$
Light green	0.02	10^{-5}
Yellow	0.03	10^{-6}
Orange	0.03	$10^{-5.5}$
Red	0.03	10^{-5}

Given its widespread use in the analysis of the lake problem, we selected the expectation of the net present value of utility as the first objective. We identified four additional objectives, which represent different stakeholders in the imaginary town. Some of these objectives are a proxy for ecosystem services, i.e., phosphorus levels in the lake, whereas others serve as proxies for temporal distribution of economic services, i.e., present vs. future utility. This approach can be interpreted as representing five hypothetical stakeholder groups whose preferences are mapped to an objective. Thus, the following objectives were considered in the analysis: discounted net present value of expected utility (maximize), average levels of phosphorus in the lake (minimize), expected utility of the present stakeholder (maximize), expected utility of the future stakeholders (maximize), and reliability of keeping the lake below the eutrophication threshold (maximize). These objectives are labelled O_1 to O_5 , respectively. A constraint was set on the reliability objective because it captures the strong aversion to irreversible losses of key economic and ecosystem services from the lake that will result on eutrophication. Refer to Appendix 3 for a detailed definition of each objective and constraint.

Step 2. Alternative problem formulations

These five objectives are used to generate three different problem formulations that represent three potential frameworks to identify a suitable pollution strategy for the lake.

- Deterministic single objective (P1): this formulation maximizes the net present value of utility (O_1) ignoring all uncertainties and identifies a single optimal management strategy with the assumption of perfect foresight.
- Stochastic single objective (P2): this formulation maximizes the net present value of expected utility (O_1) assuming well-characterized uncertainties in natural inflows and identifies a single optimal management strategy. The formulations P1 and P2 are summarized as follows:

$$\begin{aligned}
 F(x) &= O_1 \Big|_N, \\
 \forall x &\in \Omega \\
 x &= (a_1, a_2, \dots, a_{20}) \\
 N &= \begin{cases} 1 & \text{for P1} \\ 10000 & \text{for P2} \end{cases}
 \end{aligned} \tag{2}$$

In these equations, $F(x)$ is the objective function vector to be optimized, x is the decision vector, and N is the number of SOWs over which O_i is optimized.

- Stochastic five objectives (P3): this formulation identifies a Pareto approximate front that is obtained after considering all the objectives ($O_1, O_2, O_3, O_4,$ and O_5) simultaneously. The Pareto approximate front consists of all strategies that are nondominated, i.e., it is not possible to maximize one objective without degrading the value of another across these strategies. The formulation is as follows:

$$\begin{aligned}
 F(x) &= (O_1, O_2, O_3, O_4, O_5) \Big|_N, \\
 \forall x &\in \Omega \\
 x &= (a_1, a_2, \dots, a_{20}) \\
 N &= 10000 \\
 \text{subject to, } c_{rel} &: O_5 > 0.9
 \end{aligned} \tag{3}$$

Here, c_{rel} is the constraint on the reliability metric that allows only those solutions to be feasible that perform above 0.9 in the reliability objective.

Step 4. Define robustness based on stakeholders' performance requirements

Several robustness metrics have been proposed to evaluate if a management action remains viable across alternative SOWs sampled for deeply uncertain factors (Starr 1963, Schneller and Sphicas 1983, Lempert et al. 2003, Dixon et al. 2007, Lempert and Collins 2007, Kasprzyk et al. 2013, Herman et al. 2015). Lempert et al. (2003:52) define robustness as one that “performs reasonably well compared to the alternatives across a wide range of plausible futures.” They propose metrics such as regret, defined as the difference between performance of a selected strategy in an uncertain scenario and the optimal strategy for that scenario. Strategies that have small average regrets are then categorized as robust. Lempert and Collins (2007:1009) further compared various alternatives to assess robustness: “trading some optimal performance for less sensitivity to assumptions, satisficing over a wide range of plausible futures, and keeping options open.” Herman et al. (2015) provide a comprehensive assessment of robustness measures. Although most studies agree that robustness is related to the performance of a strategy across uncertain states of the world, to our knowledge only a few studies cast robustness as multicriterion regrets whereas others focus on a single objective measure (Kasprzyk et al. 2013, Herman et al. 2014).

In our analysis, the presence of multiple objectives necessitates that robustness be linked with the performance across stakeholder-defined variables as well as multiple SOWs. This forms the basis of our robustness metric, which estimates robustness as the percentage of total SOWs in which a strategy satisfies performance requirements. The estimation of robustness index requires the specification of two criteria:

1. Performance thresholds: stakeholders identify the minimum acceptable performance for variables of interest, a deterioration of performance below this threshold is

considered unacceptable. These variables can be the same as or different from objectives based on stakeholders' requirements.

2. Uncertainty represented through a set of states-of-the-world (SOWs): we consider two types of uncertainty representations: well-characterized and deep uncertainty. Both representations of uncertainty are quantified through a set of scenarios or SOWs sampled from lognormal distributions of fixed (variable) parameters for well-characterized (deep) uncertainty.

A strategy that satisfies performance thresholds for all variables in a given SOW scores 1 for that SOW, otherwise it scores 0. Strategies that achieve acceptable performance across a range of variables instead of being acceptable or high achieving in one variable, while being suboptimal in others, are termed “satisficing” Robustness is then defined as the fraction of the SOWs in which a strategy achieves satisficing performance. Based on this framing, strategies that perform reasonably well across all variables will achieve higher robustness than those that achieve optimal performance in one or more (but not all) variables or those that deteriorate heavily in any variable under changing scenarios. The detailed mathematical description of the robustness index is provided in Appendix 3. Table 3 lists the objective functions and problem formulations used in the study.

Table 3. The five objective functions and the three problem formulations used in the study.

Objectives	Uncertainty	Problem formulations
Expected utility	Deterministic	P1
	Stochastic	P2
Expected utility of present stakeholders	Stochastic	
Expected utility of future stakeholders	Stochastic	
Average levels of phosphorous in the lake	Stochastic	
Reliability	Stochastic	
All of the above	Stochastic	P3

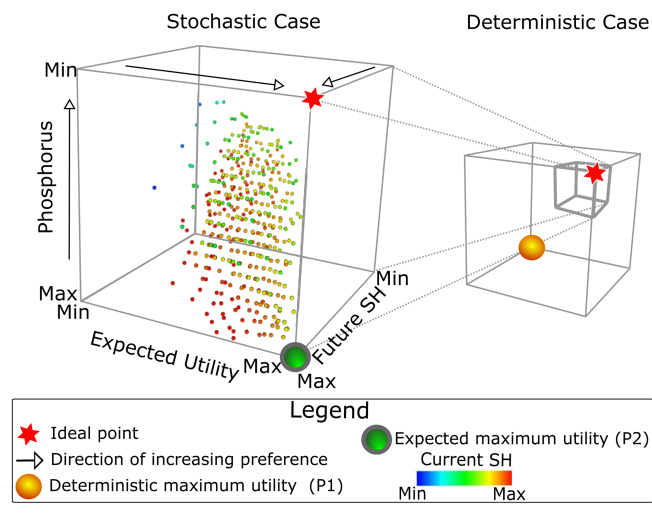
RESULTS

Step 3. Trade-offs under well-characterized uncertainty

The decision myopia inherent in the single objective problem formulations as well as the consequences of not including uncertainty in the decision analysis is illustrated in Figure 5. We obtained single optimal pollution control strategy for P1 and P2, whereas 399 Pareto-approximate strategies were obtained for P3. Thus, we assessed 401 strategies under well-characterized uncertainties and across all objectives in Figure 5. The red star represents the ideal point that shows the best possible values that can be attained across all the objectives simultaneously. Although the ideal solution is not actually feasible for the problem formulations considered because of conflicting objectives, it provides a visual reference to assess potential compromises. The larger box in the right-hand side cube represents the range of objective function values spanning P1, P2, and P3. The smaller

box spans the range of objective function values spanning P2 and P3 only. For clarity, the left-hand side cube shows a zoomed view of the performance trade-offs for P2 and P3. Comparison between the left and right-hand side panels indicates that the entire Pareto-approximate front along with the green solution falls within a much smaller region when viewed with respect to the performance of the deterministic solution.

Fig. 5. Visualizing the alternative solution strategies for the lake problem obtained by using the three problem formulations (P1, P2, and P3). The formulations are: (1) maximization of utility with no uncertainty (P1), (2) maximization of expected utility under stochastic uncertainty (P2), and (3) simultaneous maximization of five objectives under stochastic uncertainty (P3). The deterministic strategy from P1 is plotted in the right panel and the stochastic strategies from P2 and P3 are plotted in the left panel. All solutions have been re-evaluated under 10,000 states-of-the-world (SOWs), which represent well-characterized baseline uncertainty. Each cube represents a three-dimensional space with each dimension mapping performance of an objective function. Color represents the fourth dimension, the performance of the fourth objective function. The five objectives are maximization of the expected utility (x-axis), maximization of expected utility of the future stakeholders (y-axis), minimization of average phosphorus in the lake (z-axis), maximization of expected utility of current stakeholders (SH; color), and maximization of reliability (not shown here since all solutions achieved reliability > 98%). The total number of alternative strategies plotted is 401 (one for P1 and P2 and 399 for P3). The arrows indicate direction of increasing preference for each objective and the red star shows the ideal solution.



The strategy based on P1, i.e., deterministic utility, deteriorates drastically in performance across all objectives when re-evaluated under the broader scope of the uncertainty and additional objectives. This is of potential concern because it is quite common to assume a deterministic problem formulation in environmental planning studies (Maler et al. 2003, Chen et al. 2012). Both P2 and P3 perform relatively well across objectives because they are optimized over the baseline uncertainty. The deterministic P1

formulation yielded an extremely poor performance of 11% for the reliability objective. Alternatively, the MEU solution from P2 is able to maintain high reliability (98%) under all 10,000 SOWs representing well-characterized uncertainty. P3 strategies have the highest performance in the reliability objective, maintaining a reliability of 99%-100% across all solutions.

The trade-offs between different stakeholders' objectives are well evident in the P3 strategies plotted under the stochastic case in Figure 5 (left panel). Recall that each strategy in the Pareto set represents a nondominated solution, i.e., it is not possible to improve one objective without degrading another objective. Hence, they encapsulate many potential objective combinations for managing the lake that are mathematically noninferior, i.e., optimal trade-offs. For example, the trade-off between expected utility and phosphorus in the lake can be visualized by the sloping surface of the Pareto-approximate set. Increasing expected utility simultaneously increases the levels of phosphorus in the lake, causing the Pareto front to move away from the ideal red star. This is anticipated because the utility function requires some pollution emissions to obtain economic gains.

Simultaneous maximization of expected utility and the expected utility of future stakeholders is only feasible in our analysis if phosphorous levels are not actively reduced. This is indicated by the presence of solutions near the maximum of both objectives only on the bottom right of the plot. The intertemporal tension between expected utility of the present stakeholders and expected utility is also highlighted by the change of color from red to yellow/green shades (high to low values of expected utility of current stakeholder) along the increasing direction of the expected utility axis. Thus, the strategy maximizing the expected utility of the present stakeholders aims to pollute the lake at its maximal capacity in the first time period so as to maximize near term economic gains. Consequently, this leads to a reduction in the expected utility, which is estimated across the entire planning horizon. Readers are directed to Appendix 4, which illustrates the trade-offs between each pair of objectives for P3 explicitly.

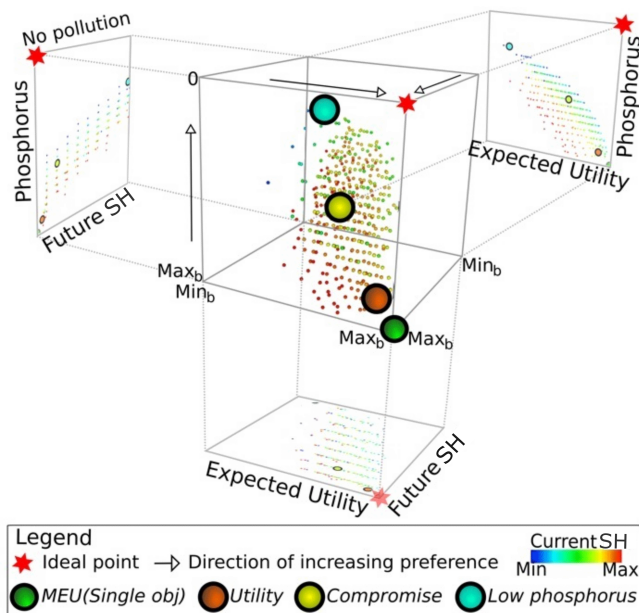
The potential for cognitive myopia to influence stakeholders focused solely on the utility-based formulations is evident in Figure 5. The strategy based on P2 (termed MEU) places the stakeholders in an extreme region of the available trade-off space, which can only be comprehended by contrasting it with the multiobjective formulations. Despite a range of possible compromises between expected utility and phosphorus in the lake, the P2 strategy focuses solely on expected utility maximization in exclusion of understanding environmental or explicit intertemporal trade-offs. The P3 formulation permits exploration of trade-offs and evaluation of the broader consequences of the single objective P1 and P2 formulations. Had we used only P1 or P2 to perform the optimization, their highly restrictive definitions of optimality would severely reduce stakeholders understanding of lake management alternatives.

Navigating trade-offs to explore diverse alternatives

We illustrate a posteriori constructive decision aiding in which stakeholders with different problem conceptions and preferences can falsify formulations while selecting candidate strategies of interest. The identification of compromise solutions will invariably involve negotiations between stakeholders with different objectives and can also involve potentially multiple

iterations on the problem formulations considered. Because it is not possible to mimic the exact process of arriving at potential compromises without a set of real stakeholders, we used a conceptual abstraction of hypothetical stakeholders. The strategy obtained using the deterministic utility maximization (P1) has severe multiobjective regrets even under well-characterized uncertainty (Fig. 5), thus this formulation has been falsified and we do not assess it any further. The MEU strategy obtained using P2 performed relatively well in the utility-based objectives and is retained in further analysis as representative of stakeholders focused solely on this objective. The 5 objective P3 formulation yields a total of 30 subproblems (5 single objective, 10 2 objective, 10 3 objective, and 5 4 objective) whose solutions are plotted in Figure 6. These solutions capture the trade-offs in all the subproblems, and the projection plots show how these multidimensional trade-offs can be projected back to reduced dimensions to assess individual tensions between each pair of objectives. Three solutions are selected from the 399 solutions obtained using P3 to represent diverse stakeholders' preferences.

Fig. 6. Illustrating the process of selecting solution strategies from the set of 401 obtained in Figure 5. The central plot is identical to the left panel plot in Figure 5 with identical axis and color. Highlighted solutions are MEU (green), low phosphorus (light blue), compromise (yellow), and utility (orange). Marginal projections showing pair-wise relationships between objectives are shown on three sides of the cube. The projection plots highlight three pair-wise trade-offs between expected utility, average phosphorus in the lake, and the expected utility of the future stakeholders (SH).



The highlighted light blue lake management strategy, termed low phosphorus, represents stakeholders whose sole emphasis is on maintaining low phosphorus levels in the lake (Fig. 6). The highlighted orange lake management strategy, termed the utility solution, represents stakeholders that seek high values of expected utility, expected utility of the present generation, and expected

utility of future generations. To represent a balance between the opposing preferences captured by the low phosphorous and the utility management strategies, we highlighted a compromise strategy (marked in yellow) that lies at the center of the trade-off region and represents midrange performance in all objectives. This strategy is termed the compromise lake management strategy. The compromise strategy was selected by maximizing the worst performance across all objectives and attempts to ensure that the performance across all objectives is well represented. Although this is a simplified abstraction of the negotiation process, the broader conceptual contribution is that Figure 6 encapsulates a suite of potential objective combinations that can be explored explicitly, i.e., stakeholders can choose what they prefer with the knowledge of what is possible. In real decision contexts, stakeholders are likely to have a far more nuanced evaluation of their objectives while seeking to facilitate compromises (Basdekas 2014). Table 4 lists the expected objective function values obtained for the four selected strategies under well-characterized uncertainties.

Table 4. Performance of four distinct types of solutions from the approximate Pareto set for the five-objective problem (P3) and the single-objective stochastic formulation (P2) under well-characterized uncertainty. The compromise strategy is obtained by using the max-min approach discussed in the text. The utility solution is identified as the one for which performance around 95% can be achieved for the expected utility (phosphorous) objectives. The low phosphorus solution has the lowest average concentration of phosphorus in the lake. MEU = maximization of expected utility.

Objective	Performance (% of range)			
Average levels of phosphorus	11.9	59.0	100.0	0.8
Expected utility	92.6	62.1	5.1	103.7
Expected utility of present stakeholders	86.1	72.1	22.1	52.6
Expected utility of future stakeholders	93.6	78.1	4.5	98.2
Reliability	98.6	100.0	100.0	0.0
Assigned code	Utility	Compromise	Low phosphorus	MEU

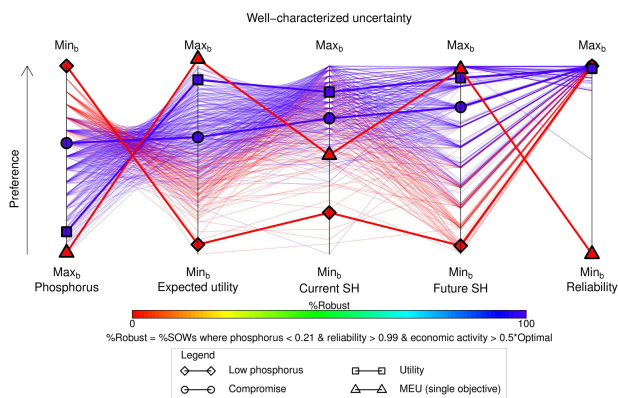
Step 5. Robustness analysis under well-characterized uncertainty

As a consequence of the a posteriori trade-off analysis on P1, P2, and P3, we selected four strategies (MEU, low phosphorus, compromise, and utility) to represent diverse stakeholders in the lake problem. In particular, the compromise strategy was selected as an abstraction of the negotiation process using a simple, but ad hoc best-worst heuristic rule. The difficulty with this rule and other commonly employed simple multicriterion weighting strategies, e.g., minimizing the distance from the ideal solution, is that they assume every stakeholder would be willing to make substantial sacrifices across all performance objectives equally (Keeney and Raiffa 1993, Brink 1994, Köksalan et al. 2013). This is a strongly questionable assumption for many resource-management applications that show high levels of risk aversion and a strong emphasis on the reliability of services (Brown et al. 2012, Admiraal et al. 2013, Kasprzyk et al. 2013, Giuliani et al.

2014, Zeff et al. 2014). In this study, we present a robustness analysis that assesses strategies with respect to diverse performance requirements under well-characterized and deep uncertainty.

We assessed the robustness of strategies under baseline uncertainty using the performance requirements discussed in Appendix 3. Note that in a real decision-support application, this process of setting performance requirements can be repeated several times as new insights are gained through the MORDM analysis. Visualizing the performance of strategies with respect to their robustness revealed that the strategies can be classified into two disparate groups, one with very high and another with very low robustness (Fig. 7). We found that both low phosphorus and MEU have 0% robustness even under baseline uncertainty. The low phosphorus solution is not robust because it fails to satisfy the performance requirements on economic activity whereas the MEU solution fails to satisfy the 99% reliability requirement for avoiding the eutrophic threshold in the lake. Both the utility and compromise solutions have 100% robustness under baseline well-characterized uncertainty. In reference to the broader set of trade-off solutions, two families of solutions are identified, distinguished by the red vs. blue colors. We found that

Fig. 7. Parallel coordinate plots showing the baseline expected performance of solutions from P2 and P3 across five objectives averaged across 10,000 states-of-the-world (SOWs) under well-characterized uncertainty. The horizontal axis represents the five objectives and the vertical axis represents the objective function values (normalized across the range of the objective). Solutions are colored based on their robustness, which is defined using the performance thresholds. The best (worst) strategy would be represented by a horizontal line running through the top (bottom) of the plot. Note that the vertical scale is reversed for the phosphorus objective, as minimization is preferred. Diagonal lines represent explicit trade-offs between objective pairs. The expected utility maximization (MEU; single objective), low phosphorus, compromise, and utility strategies from Figure 6 are highlighted through markers. The ranges of objectives shown as (min_b, max_b) are (0.06, 0.21), (0.03, 0.36), (0.00, 0.04), (0.00, 0.01), and (0.98, 1.00), rounded to two decimal places, for average phosphorus in the lake, expected utility, expected utility of current stakeholders (SH), expected utility of future stakeholders, and reliability, respectively.



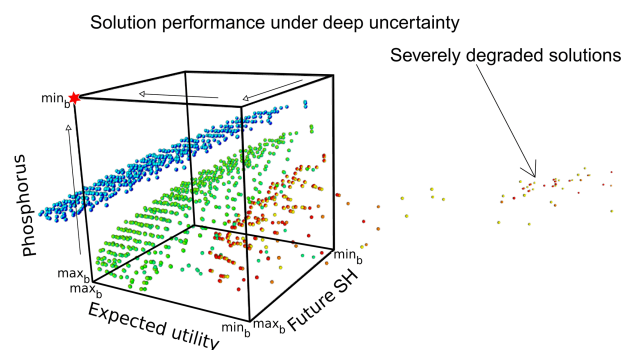
all solutions were able to satisfy the phosphorus and reliability performance requirements, except MEU. It is the requirement on economic activity that mainly distinguishes the two families of solutions.

Step 6. Robustness analysis under deep uncertainty

Once the robustness of solutions is assessed for baseline well-characterized uncertainty, stakeholders can either go back to step 2 or step 4 (of Fig. 2) in case they do not find solutions with satisfactory robustness. However, in this case, we are able to identify a significant number of solutions with high robustness. Following this, we re-evaluated the performance of all solutions across eight varying assumptions of uncertainty that characterize our representation of deep uncertainty (Fig. 4a). We explored the impact of these varying distributions on the Pareto-approximate front in Figure 8. We found that the baseline Pareto-approximate front transitions significantly for alternative distributions representing uncertain pollution inflow changes. The front moves closer to the ideal point as the mean of the uncertain pollution inflow decreases (distributions i, ii, and iii), and the impact of altering the variance is negligible as seen from the overlapping fronts from these distributions. When the mean of the pollution inflow increases (distributions vii, viii, and ix), we observed severe degradation in all objectives, not just the levels of phosphorus in the lake, accompanied by an increased variation across alternative assumptions of variance. This implies that there are multiobjective opportunity costs if the mean of the uncertain phosphorous inflows is over-estimated and multiobjective regrets otherwise. This information can be helpful in providing stakeholders with an understanding of the effects of their assumptions and inform their risk attitudes.

The significant shift in the Pareto front across varying assumptions of uncertainty motivates the need for assessing robustness of solution under deep uncertainty, i.e., across 90,000

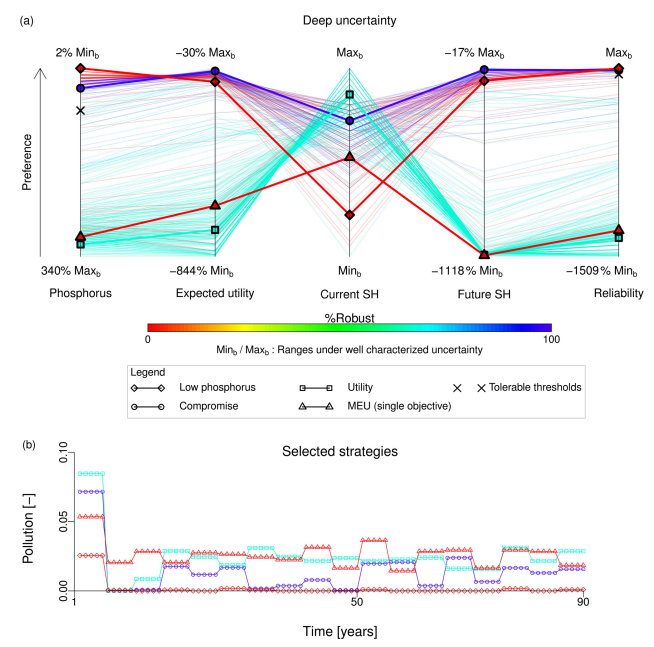
Fig. 8. Re-evaluating 399 solutions from P3 under deep uncertainty. The color of the Pareto-approximate fronts corresponds to the color of the lognormal distributions in Figure 4 that the solutions are re-evaluated against. Expected utility, expected utility of future stakeholders (SH), and average phosphorus in the lake are plotted on x-axis, y-axis, and z-axis, respectively. The red star shows the location of the ideal solution and arrows indicate direction of increasing preference. Solutions with severe performance degradation under changing assumptions of uncertainty are indicated by an arrow.



SOWs. The robustness analysis is carried out using the same performance requirements as before, except for the expected levels of phosphorus in the lake. While assessing robustness under well-characterized uncertainty, we found that all solutions were able to satisfy the phosphorus requirements, indicating that the chosen performance level was too relaxed. It is possible that under deep uncertainty, solutions remain robust despite a significant deterioration in performance of phosphorus when compared to the performance under well-characterized uncertainty. Because stakeholders are likely to prefer solutions that maintain phosphorus levels better than the worse performing solution in the baseline uncertainty, a stricter requirement for phosphorus was fixed at the maximum level of phosphorus obtained across the solutions under the baseline uncertainty. This is one example of how stakeholders can go back to step 4 from step 5 in Figure 2.

By assessing the robustness of strategies under deep uncertainty, we discovered that there were three families of solutions with very low, moderate, and high robustness, respectively (Fig. 9a). For four out of five objectives, there is a significant change in the objective function ranges when compared to the solution

Fig. 9. (a) Parallel coordinate plots showing the performance of solutions from P2 and P3 across five objectives averaged across 90,000 states-of-the-worlds (SOWs) under deep uncertainty. The horizontal axis represents the five objectives and the vertical axis represents the objective function values (normalized across the range of the objective). Solutions are colored based on their robustness, which is defined using the performance thresholds. Expected utility maximization (MEU; single objective), low phosphorus, compromise and utility strategies from Figure 6 are highlighted through markers. The ranges for objectives are shown as deviation from the (min_b, max_b) ranges in Figure 7. (b) The amount of allowed anthropogenic pollution in the lake as a function of time for selected solution strategies.



performance in the baseline case (Fig. 7). There is an overall deterioration in the best and worst objective function values when assessing performance under deep uncertainty. Only a small set of solutions with high robustness (blue color) remain under deep uncertainty, indicating that it is much harder to identify solutions that can preserve performance under changing assumptions of uncertainty. While under baseline uncertainty, only the economic-activity requirement distinguished robust solutions, under deep uncertainty it was hard to satisfy performance requirements for all three variables. Many solutions failed to satisfy the performance requirements for phosphorus and reliability under deep uncertainty, therefore these requirements are now shown as crosses on the vertical axis associated with the respective objective. Solutions that were nearly 100% robust under baseline uncertainty now deteriorated to reduced levels of robustness (approx. 67%). Under deep uncertainty, three out of four selected strategies failed to maintain high levels of robustness, whereas only the compromise strategy remained highly robust (approx. 100%). We also find that the compromise solution is surrounded by medium to low robust solutions. Therefore, it would be incorrect to assume that the solutions near the robust solutions are likely to be robust. The pollution strategies associated with selected solutions are shown in Figure 9b.

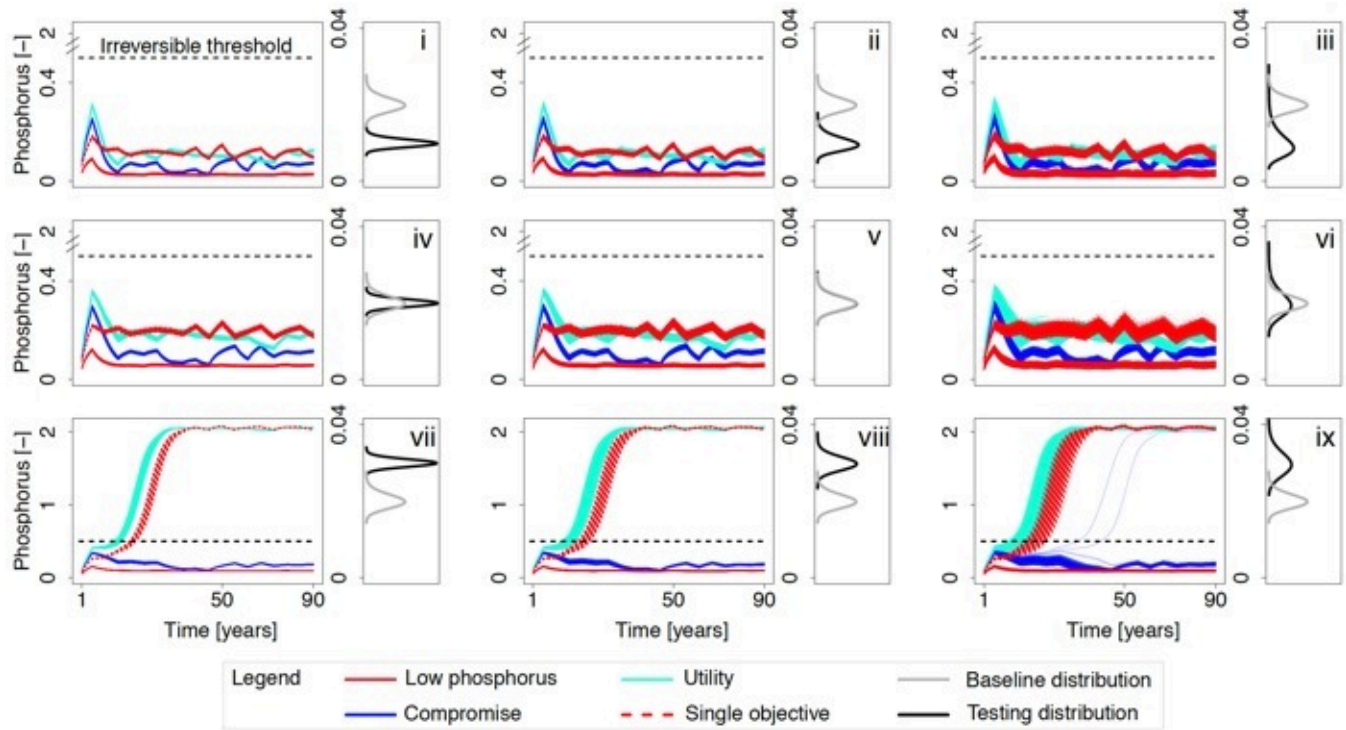
Finally, we show the performance of selected strategies under deep uncertainty in terms of the water quality dynamics for the lake, i.e., the amount of phosphorus in the lake as a function of time (Fig. 10). This shows whether or not selected strategies prevent eutrophication of the lake under various parameterizations of the lognormal distributions representing deep uncertainty. We found (no) considerable impact of altering variance of the lognormal distribution at (low to medium) high mean values of random pollution inflow. The compromise strategy outperforms utility and MEU as the mean and variance of the random pollution inflow increase. Eutrophication occurs for only 2 out of 90,000 SOWs for the compromise strategy. These cases occur when the mean and variance of the lognormal distribution describing random inflows are highest. These two strategies also fail much farther into the planning period, after year 40. On the other hand, the MEU and utility strategies fail under all SOWs very early into the planning horizon, i.e., before year 20, when the mean of the lognormal distribution is the highest (30,000 out of 90,000 SOWs considered), irrespective of the variance.

DISCUSSION

Many-objective robust decision making was successful in identifying the compromise strategy, which had high robustness in preventing eutrophic collapse of the lake in almost all cases (Figs. 9 and 10). The multiobjective (MO) aspect of MORDM provides a richer context for evaluating alternative problem formulations, revealing the myopic location of P1 and P2 when compared to P3 (Fig. 5). The trade-offs evident in Figure 6 clarify key conflicts while providing a diverse suite of solutions to aid stakeholders when selecting a strategy for implementation. The trade-offs can also facilitate negotiated compromise across conflicting objectives.

Then using robust decision making (RDM), we were able to test the strategies further and demonstrate additional weaknesses in the P1 and P2 problem framings. The re-evaluation of strategies under deep uncertainties highlighted the multiobjective regrets

Fig. 10. (i-ix) Phosphorus in the lake for the expected utility maximization (MEU; single objective), low phosphorus, compromise, and utility strategies under different assumptions of uncertainty in natural inflows. Main plots show the levels of phosphorus in the lake as a function of time in years. The strategies are colored based on their robustness in Figure 9. A dashed line is used to represent the single objective (MEU) whereas solid lines represent the three multiobjective solutions. The side plots on the right of each main plot show the lognormal distribution of natural inflows used for estimating the phosphorus in the lake (black) against the baseline well-characterized distribution used for optimization (gray). The plots are numbered according to the distributions in Figure 4. Also note that the y-axis contains an axis break in the main plot for distributions i-vi.



that are likely to occur if utility-based frameworks are used for managing ecosystems. The utility-based strategies (utility and MEU) not only lead to eutrophic collapse of the lake in a significant number of cases (Fig. 10), they also failed to maintain high levels of robustness (Fig. 9). Because economic activity is also included in the performance requirements for robustness, this implies that the utility-based strategies ultimately lead to a loss of both environmental and economic goals of the stakeholders. This is expected because the eutrophic lake will also have low utility.

Thus, MORDM allowed the falsification of the utility-based problem framing in this study by exposing its failure to maintain performance across objectives and under deep uncertainty. Only a small set of initially available strategies were identified as robust across a range of stakeholders' objectives and deep uncertainties. This shows how MORDM provides a means for the stakeholders' to not only falsify problem formulations but also to identify new strategies that will satisfy their diverse objectives and risk attitudes. This has advantages over approaches that rely on single objective functions, or focus on simplified representation of the ecosystem to be analytically tractable (Crépin 2007, Webster et al. 2012). Many-objective robust decision making also enhances the robust decision-making approach by Lempert and Collins (2007) by providing a systematic approach to generate strategy

alternatives based on multiple criteria. To summarize, MORDM can do the following:

- employ a large number of objective functions (up to 10) without the need to aggregate them, thereby presenting stakeholders with the trade-off curve, instead of an optimal strategy. It was not possible to visualize the trade-off surfaces shown in the current study using any other existing methods.
- be independent of the choice of ecosystem models, more complicated models with multiple processes can be easily substituted for simpler ones.
- allow stakeholders to define multiobjective robustness measures according to the level of compromises they are willing to accept. Other approaches either do not have any provision to account for robustness, or employ single-objective robustness measures.

Visual analytics form a crucial aspect of MORDM because it would not have been possible to identify the myopia of the P1 and P2 formulations without the high dimensional visualization in Figure 5. The figure presents a range of Pareto-optimal solutions that comprise solutions maximizing each individual objective along with solutions spanning the entire range of possible trade-

offs between these optimal values. This is in stark contrast with the a priori weighting of different objectives, which results in a single a solution, thus, predetermining the location of the compromise in the multidimensional objective space (Woodruff et al. 2013). A typical weighting-based solution would form just one point among the large number of strategies shown in Figure 5. Similarly, a closer inspection of the Pareto front in Figure 6 allowed the selection of strategies representing a broad range of stakeholder objectives. The transition of the Pareto front under varying assumptions of uncertainty set the stage for a comprehensive robustness analysis (Fig. 8). The identification of solution groups, two for baseline uncertainty and three for deep uncertainty, would not have been feasible without the parallel coordinate plots (Figs. 7 and 9a). Finally, visualizing the lake dynamics for selected strategies across all 90,000 SOWs allowed us to understand the potential for eutrophication under various assumptions of uncertainty.

CAVEATS AND FUTURE RESEARCH NEEDS

Many-objective robust decision making can provide several additional insights compared to the classic expected utility maximization approach. However, the MORDM approach in general, and our approach in this study in particular, still face several limitations. First, the analyzed Pareto set is approximate. Even though this can be categorized as a weakness of this approach, it is well-known that most real-world problems are “wicked” and ill defined (Rittel and Webber 1973). Consequently, small refinements of optimality are likely less important than substantial structural changes in alternative problem framings. Second, the strategy obtained through our analysis is not adaptive, i.e., there does not exist a way to incorporate new knowledge about the system within this framework yet. This forms the basis of ongoing work, which aims to incorporate learning within the MORDM framework. Third, our problem framings still do not explore the full suite of uncertainties relevant for the lake problem and rely heavily on the utility function. For example, we do not consider the uncertainties in the lake model itself (structural as well as parametric), the parameters of the utility function, etc. Although we did not perform an extensive uncertainty analysis to incorporate all possible sources of uncertainties, our proposed framework is well-equipped to tackle a wide array of uncertainties as had been demonstrated elsewhere (Kasprzyk et al. 2013, Woodruff et al. 2013).

Finally, communicating the high dimensional Pareto fronts as well as deep uncertainty are key challenges for the MORDM approach. There have been many real application successes that have used high-dimensional a posteriori trade-off analysis in complex environmental and engineering design contexts, which is encouraging (Fleming et al. 2005, Feringer et al. 2009, Reed et al. 2013, Basdekas 2014). We hypothesize that MORDM can change the decision making process and outcomes when stakeholders have varied objectives, do not know a priori the levels of compromise they can accept, and when there are deep uncertainties about the system’s future. To be clear, we are not aware of any formal, i.e., academic, peer reviewed, or laboratory based, study in the field of judgment and decision making that assesses the effect of confronting stakeholders with the kind of multiobjective trade-off displays shown here and that tests this hypothesis.

CONCLUSIONS

Our results revealed the failure of utility-based frameworks in satisfying multiple objectives of diverse stakeholders. We showed that the strategies that aim only to maximize the utility function will place stakeholders in an extreme region of the trade-off space. We term this failure as myopia and visualize this effect. We also found that the utility strategy degrades severely in robustness under deep uncertainty by causing eutrophication in the lake for 33% of the deep uncertainty SOWs. Our results are also in agreement with the recent findings by Admiraal et al. (2013) that show how traditional problem framings can fail to identify sustainable management policies for environmental systems. Admiraal et al. (2013) attributed this to the inability of the utility function to capture the key features of a natural system that make it sustainable. Our framework allows environmental planners to move beyond utility-based approaches to methods that incorporate (potentially multiple) indicators of sustainability into the problem framing.

We also demonstrated the crucial role visual analytics can play in the decision-making process. Each high dimensional visual allowed us to communicate the insights gained from the a posteriori decision-making process. They allowed the exploration of multiple formulations and the understanding of key trade-offs under well-characterized uncertainties (Fig. 5). This also enabled the communication of the myopia inherent in the MEU formulation. The set of feasible alternatives were examined to identify candidate solutions from different regions of the trade-off space (Fig. 6). Visualizing the robustness of strategies under well-characterized and deep uncertainties provided a broader understanding of strategy performance (Figs. 7 and 9). In addition, a re-evaluation of trade-offs under deep uncertainty motivated the need for exploring the assumptions of well-characterized uncertainty (Fig. 8).

We implement a constructive decision-aiding approach to identify a robust strategy for the management of an irreversible lake. We present a systematic procedure to implement constructive decision aiding through MORDM. Many-objective robust decision making discovers the compromise strategy that balances various environmental, economic, and intertemporal objectives not only under well-characterized uncertainty, but also under deep uncertainty. This was made possible by employing a comprehensive definition of robustness that assesses strategies against minimum acceptable performance requirements under various assumptions of uncertainty. Thus, MORDM allows stakeholders to explore the entire trade-off space and select a posteriori the level of compromises they wish to attain. It was impossible for stakeholders to discover these implicit choices without the facility of a flexible problem framing provided by MORDM. Embedding the requirements of sustainability within the definition of robustness, and then applying MORDM to search for robust (hence, sustainable) strategies, provides a promising avenue for environmental management problems.

Responses to this article can be read online at:
<http://www.ecologyandsociety.org/issues/responses.php/7687>

Acknowledgments:

We are grateful for the constructive comments given by the two anonymous reviewers; their insights have significantly improved the manuscript. We thank David Hadka, the developer of the BORG MOEA for his insights on implementation of the algorithm for the lake problem. This work was partially supported by the National Science Foundation through the Network for Sustainable Climate Risk Management (SCRiM) under NSF cooperative agreement GEO-1240507 and the Penn State Center for Climate Risk Management. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. We thank Stephen R. Carpenter and W. A. Brock for their useful inputs on the initial stages of the analysis. We also thank Robert J. Lempert, Nancy Tuana, Gregory G. Garner, the Reed and Keller research groups, and the SCRiM external advisory team for their inputs via many discussions. Any errors or omissions are the responsibility of the authors alone. The first author performed the analysis and wrote the first draft of the paper while affiliated with the Earth and Environmental Systems Institute, Pennsylvania State University. The entire review process was completed by the first author at the Indian Institute of Technology Hyderabad, India. All coauthors contributed to the interpretation of results, writing, and revision of the manuscript.

LITERATURE CITED

- Admiraal, J. F., A. Wossink, W. T. de Groot, and G. R. de Snoo. 2013. More than total economic value: how to combine economic valuation of biodiversity with ecological resilience. *Ecological Economics* 89:115-122. <http://dx.doi.org/10.1016/j.ecolecon.2013.02.009>
- Basdekas, L. 2014. Is multiobjective optimization ready for water resources practitioners? Utility's drought policy investigation. *Journal of Water Resources Planning and Management* 140:275-276. [http://dx.doi.org/10.1061/\(ASCE\)WR.1943-5452.0000415](http://dx.doi.org/10.1061/(ASCE)WR.1943-5452.0000415)
- Becker, S. W., and F. O. Brownson. 1964. What price ambiguity? or the role of ambiguity in decision-making. *Journal of Political Economy* 72(1):62-73. <http://dx.doi.org/10.1086/258854>
- Ben-Haim, Y. 2001. *Information-gap decision theory: decisions under severe uncertainty*. Academic, Linn, Missouri, USA.
- Bennett, E. M., S. R. Carpenter, and J. A. Cardille. 2008. Estimating the risk of exceeding thresholds in environmental systems. *Water, Air, and Soil Pollution* 191:131-138. <http://dx.doi.org/10.1007/s11270-007-9612-9>
- Bond, C. A. 2010. On the potential use of adaptive control methods for improving adaptive natural resource management. *Optimal Control Applications and Methods* 31:55-66. <http://dx.doi.org/10.1002/oca.910>
- Bond, C. A., and J. B. Loomis. 2009. Using numerical dynamic programming to compare passive and active learning in the adaptive management of nutrients in shallow lakes. *Canadian Journal of Agricultural Economics/Revue canadienne d'agroeconomie* 57:555-573. <http://dx.doi.org/10.1111/j.1744-7976.2009.01170.x>
- Brill, E. D., J. M. Flach, L. D. Hopkins, and S. Ranjithan. 1990. MGA: a decision support system for complex, incompletely defined problems. *IEEE Transactions on Systems, Man, and Cybernetics* 20:745-757. <http://dx.doi.org/10.1109/21.105076>
- Brink, T. L. 1994. Book review: R. L. Keeney, H. Raiffa: decisions with multiple objectives—preferences and value tradeoffs, Cambridge University Press, Cambridge & New York, 1993, 569 pages, ISBN 0-521-44185-4 (hardback), 0-521-43883-7 (paperback). *Behavioral Science* 39:169-170. <http://dx.doi.org/10.1002/bs.3830390206>
- Brown, C., Y. Ghile, M. Laverty, and K. Li. 2012. Decision scaling: linking bottom-up vulnerability analysis with climate projections in the water sector. *Water Resources Research* 48:W09537. <http://dx.doi.org/10.1029/2011wr011212>
- Brozovic, N., and W. Schlenker. 2011. Optimal management of an ecosystem with an unknown threshold. *Ecological Economics* 70:627-640. <http://dx.doi.org/10.1016/j.ecolecon.2010.10.001>
- Budescu, D. V., K. M. Kuhn, K. M. Kramer, and T. R. Johnson. 2002. Modeling certainty equivalents of imprecise gambles. *Organizational Behavior and Human Decision Processes* 88:748-768. [http://dx.doi.org/10.1016/S0749-5978\(02\)00014-6](http://dx.doi.org/10.1016/S0749-5978(02)00014-6)
- Carpenter, S. R., D. Ludwig, and W. A. Brock. 1999. Management of eutrophication for lakes subject to potentially irreversible change. *Ecological Applications* 9:751-771. [http://dx.doi.org/10.1890/1051-0761\(1999\)009\[0751:MOEFLS\]2.0.CO;2](http://dx.doi.org/10.1890/1051-0761(1999)009[0751:MOEFLS]2.0.CO;2)
- Castelletti, A., A. V. Lotov, and R. Soncini-Sessa. 2010. Visualization-based multi-objective improvement of environmental decision-making using linearization of response surfaces. *Environmental Modelling and Software* 25:1552-1564. <http://dx.doi.org/10.1016/j.envsoft.2010.05.011>
- Chankong, V., and Y. Y. Haimes. 1983. *Multiobjective decision making: theory and methodology*. North Holland, Amsterdam, The Netherlands.
- Chen, Y., C. Jayaprakash, and E. Irwin. 2012. Threshold management in a coupled economic-ecological system. *Journal of Environmental Economics and Management* 64:442-455. <http://dx.doi.org/10.1016/j.jeem.2012.04.003>
- Clark, J. S. 2005. Why environmental scientists are becoming Bayesians. *Ecology Letters* 8:2-14. <http://dx.doi.org/10.1111/j.1461-0248.2004.00702.x>
- Coello, C., A. Hernandez Aguirre, E. Zitzler, P. Fleming, R. Purshouse, and R. Lygoe. 2005. Many-objective optimization: an engineering design perspective. Pages 14-32 in C. M. Fonseca, X. Gandibleux, J.-K. Hao, and M. Sevaux, editors. *Evolutionary multi-criterion optimization: 5th international conference, EMO 2009, Nantes, France, April 7-10, 2009, proceedings*. Springer, Berlin, Germany.
- Cohon, J. L., and D. H. Marks. 1975. A review and evaluation of multiobjective programming techniques. *Water Resources Research* 11:208-220. <http://dx.doi.org/10.1029/WR011i002p00208>
- Common, M., and C. Perrings. 1992. Towards an ecological economics of sustainability. *Ecological Economics* 6(1):7-34. [http://dx.doi.org/10.1016/0921-8009\(92\)90036-r](http://dx.doi.org/10.1016/0921-8009(92)90036-r)
- Crépin, A.-S. 2007. Using fast and slow processes to manage resources with thresholds. *Environmental and Resource Economics* 36:191-213. <http://dx.doi.org/10.1007/s10640-006-9029-8>

- Curley, S. P., and J. F. Yates. 1985. The center and range of the probability interval as factors affecting ambiguity preferences. *Organizational Behavior and Human Decision Processes* 36:273-287. [http://dx.doi.org/10.1016/0749-5978\(85\)90016-0](http://dx.doi.org/10.1016/0749-5978(85)90016-0)
- Dixon, L., R. J. Lempert, T. LaTourrette, and R. T. Reville. 2007. *The federal role in terrorism insurance: evaluating alternatives in an uncertain world*. RAND Corporation, Santa Monica, California, USA. [online] URL: <http://www.rand.org/pubs/monographs/MG679.html>
- Ellsberg, D. 1961. Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics* 75:643-669. <http://dx.doi.org/10.2307/1884324>
- Epstein, L. G., and T. Wang. 1994. Intertemporal asset pricing under Knightian uncertainty. *Econometrica* 62:283-322. <http://dx.doi.org/10.2307/2951614>
- Ferringer, M. P., D. B. Spencer, and P. Reed. 2009. Many-objective reconfiguration of operational satellite constellations with the large-cluster epsilon non-dominated sorting genetic algorithm-II. Pages 340-349 in *IEEE congress on evolutionary computation, Trondheim, 18-21 May 2009*. IEEE, New York, New York, USA. <http://dx.doi.org/10.1109/cec.2009.4982967>
- Fleming, P. J., R. C. Purshouse, and R. J. Lygoe. 2005. Many-objective optimization: an engineering design perspective. Pages 14-32 in *Evolutionary multi-criterion optimization: proceedings of third international conference, EMO, Guanajuato, Mexico, March 9-11, 2005*. Springer-Verlag, Berlin, Germany. http://dx.doi.org/10.1007/978-3-540-31880-4_2
- Giuliani, M., J. D. Herman, A. Castelletti, and P. Reed. 2014. Many-objective reservoir policy identification and refinement to reduce policy inertia and myopia in water management. *Water Resources Research* 50:3355-3377. <http://dx.doi.org/10.1002/2013WR014700>
- Haasnoot, M., J. H. Kwakkel, W. E. Walker, and J. ter Maat. 2013. Dynamic adaptive policy pathways: a method for crafting robust decisions for a deeply uncertain world. *Global Environmental Change* 23:485-498. <http://dx.doi.org/10.1016/j.gloenvcha.2012.12.006>
- Haasnoot, M., W. P. A. van Deursen, J. H. A. Guillaume, J. H. Kwakkel, E. van Beek, and H. Middelkoop. 2014. Fit for purpose? Building and evaluating a fast, integrated model for exploring water policy pathways. *Environmental Modelling and Software* 60:99-120. <http://dx.doi.org/10.1016/j.envsoft.2014.05.020>
- Hadka, D., and P. Reed. 2012. Diagnostic assessment of search controls and failure modes in many-objective evolutionary optimization. *Evolutionary Computation* 20:423-452. http://dx.doi.org/10.1162/EVCO_a_00053
- Herman, J., P. Reed, H. Zeff, and G. Characklis. 2015. How should robustness be defined for water systems planning under change? *Journal of Water Resources Planning and Management*. [http://dx.doi.org/10.1061/\(ASCE\)WR.1943-5452.0000509](http://dx.doi.org/10.1061/(ASCE)WR.1943-5452.0000509)
- Herman, J. D., H. B. Zeff, P. M. Reed, and G. W. Characklis. 2014. Beyond optimality: multistakeholder robustness tradeoffs for regional water portfolio planning under deep uncertainty. *Water Resources Research* 50:7692-7713. <http://dx.doi.org/10.1002/2014WR015338>
- Hogarth, R., and H. Kunreuther. 1989. Risk, ambiguity, and insurance. *Journal of Risk and Uncertainty* 2:5-35. <http://dx.doi.org/10.1007/BF00055709>
- Horan, R. D., E. P. Fenichel, K. L. S. Drury, and D. M. Lodge. 2011. Managing ecological thresholds in coupled environmental-human systems. *Proceedings of the National Academy of Sciences* 108:7333-7338. <http://dx.doi.org/10.1073/pnas.1005431108>
- Johnson, F. A., B. K. Williams, and J. D. Nichols. 2013. Resilience thinking and a decision-analytic approach to conservation: strange bedfellows or essential partners? *Ecology and Society* 18 (2): 27. <http://dx.doi.org/10.5751/es-05544-180227>
- Kasprzyk, J. R., S. Nataraj, P. M. Reed, and R. J. Lempert. 2013. Many objective robust decision making for complex environmental systems undergoing change. *Environmental Modelling and Software* 42:55-71. <http://dx.doi.org/10.1016/j.envsoft.2012.12.007>
- Kasprzyk, J. R., P. M. Reed, G. W. Characklis, and B. R. Kirsch. 2012. Many-objective de Novo water supply portfolio planning under deep uncertainty. *Environmental Modelling and Software* 34:87-104. <http://dx.doi.org/10.1016/j.envsoft.2011.04.003>
- Kasprzyk, J. R., P. M. Reed, B. R. Kirsch, and G. W. Characklis. 2009. Managing population and drought risks using many-objective water portfolio planning under uncertainty. *Water Resources Research* 45:W12401. <http://dx.doi.org/10.1029/2009wr008121>
- Keeney, R. L. 2012. Value-focused brainstorming. *Decision Analysis* 9:303-313. <http://dx.doi.org/10.1287/deca.1120.0251>
- Keeney, R. L., and H. Raiffa. 1993. *Decisions with multiple objectives: preferences and value trade-offs*. Cambridge University Press, Cambridge, UK. <http://dx.doi.org/10.1017/cbo9781139174084>
- Keim, D., J. Kohlhammer, G. Ellis, and F. Mansmann. 2010. *Mastering the information age: solving problems with visual analytics*. Eurographics Association, Goslar, Germany.
- Keller, K., and D. McInerney. 2008. The dynamics of learning about a climate threshold. *Climate Dynamics* 30:321-332. <http://dx.doi.org/10.1007/s00382-007-0290-5>
- Knight, F. H. 1921. *Risk, uncertainty and profit*. Houghton Mifflin Company, Boston, Massachusetts, USA.
- Köksalan, M., J. Wallenius, and S. Zionts. 2013. An early history of multiple criteria decision making. *Journal of Multi-Criteria Decision Analysis* 20:87-94. <http://dx.doi.org/10.1002/mcda.1481>
- Kollat, J. B., and P. Reed. 2007. A framework for visually interactive decision-making and design using evolutionary multi-objective optimization (VIDEO). *Environmental Modelling and Software* 22:1691-1704. <http://dx.doi.org/10.1016/j.envsoft.2007.02.001>
- Krishnamurti, T. N., C. M. Kishtawal, T. E. LaRow, D. R. Bachiochi, Z. Zhang, C. E. Williford, S. Gadgil, and S. Surendran. 1999. Improved weather and seasonal climate forecasts from multimodel superensemble. *Science* 285:1548-1550. <http://dx.doi.org/10.1126/science.285.5433.1548>
- Kuhn, K. M., and D. V. Budescu. 1996. The relative importance of probabilities, outcomes, and vagueness in hazard risk decisions. *Organizational Behavior and Human Decision Processes* 68:301-317. <http://dx.doi.org/10.1006/obhd.1996.0107>

- Kwakkel, J. H., M. Haasnoot, and W. E. Walker. 2014. Developing dynamic adaptive policy pathways: a computer-assisted approach for developing adaptive strategies for a deeply uncertain world. *Climatic Change* 2014:1-14. <http://dx.doi.org/10.1007/s10584-014-1210-4>
- Lempert, R. J., and M. T. Collins. 2007. Managing the risk of uncertain threshold responses: comparison of robust, optimum, and precautionary approaches. *Risk Analysis* 27:1009-1026. <http://dx.doi.org/10.1111/j.1539-6924.2007.00940.x>
- Lempert, R. J., D. G. Groves, S. W. Popper, and S. C. Bankes. 2006. A general, analytic method for generating robust strategies and narrative scenarios. *Management Science* 52:514-528. <http://dx.doi.org/10.1287/mnsc.1050.0472>
- Lempert, R. J., S. W. Popper, and S. C. Bankes. 2003. *Shaping the next one hundred years: new methods for quantitative, long-term policy analysis*. RAND Corporation, Santa Monica, California, USA. [online] URL: http://www.rand.org/pubs/monograph_reports/MR1626.html
- Ludwig, D. 2000. Limitations of economic valuation of ecosystems. *Ecosystems* 3:31-35. <http://dx.doi.org/10.1007/s100210000007>
- Mäler, K.-G., A. Xepapadeas, and A. de Zeeuw. 2003. The economics of shallow lakes. *Environmental and Resource Economics* 26:603-624. <http://dx.doi.org/10.1023/B:EARE.0000007351.99227.42>
- McInerney, D., R. Lempert, and K. Keller. 2012. What are robust strategies in the face of uncertain climate threshold responses? *Climatic Change* 112:547-568. <http://dx.doi.org/10.1007/s10584-011-0377-1>
- Morse-Jones, S., T. Luisetti, R. K. Turner, and B. Fisher. 2011. Ecosystem valuation: some principles and a partial application. *Environmetrics* 22:675-685. <http://dx.doi.org/10.1002/env.1073>
- Pareto, V. 1896. *Cours d'économie politique*. Tome I. F. Rouge, Lausanne, France.
- Peterson, G. D., S. R. Carpenter, and W. A. Brock. 2003. Uncertainty and the management of multistate ecosystems: an apparently rational route to collapse. *Ecology* 84:1403-1411. [http://dx.doi.org/10.1890/0012-9658\(2003\)084\[1403:UATMOM\]2.0.CO;2](http://dx.doi.org/10.1890/0012-9658(2003)084[1403:UATMOM]2.0.CO;2)
- Ramsey, F. P. 1928. A mathematical theory of saving. *Economic Journal* 38:543-559. <http://dx.doi.org/10.2307/2224098>
- Reed, P. M., D. Hadka, J. D. Herman, J. R. Kasprzyk, and J. B. Kollat. 2013. Evolutionary multiobjective optimization in water resources: the past, present, and future. *Advances in Water Resources* 51:438-456. <http://dx.doi.org/10.1016/j.advwatres.2012.01.005>
- Rittel, H. W. J., and M. M. Webber. 1973. Dilemmas in a general theory of planning. *Policy Sciences* 4:155-169. <http://dx.doi.org/10.1007/BF01405730>
- Schneller, IV, G. O., and G. P. Sphicas. 1983. Decision making under uncertainty: Starr's domain criterion. *Theory and Decision* 15:321-336. <http://dx.doi.org/10.1007/BF00162111>
- Simon, H. A. 1959. Theories of decision-making in economics and behavioral science. *American Economic Review* 49:253-283.
- Starr, M. K. 1963. *Product design and decision theory*. Prentice-Hall, Upper Saddle River, New Jersey, USA.
- Stump, G. M., M. Yukish, T. W. Simpson, and E. N. Harris. 2003. Design space visualization and its application to a design by shopping paradigm. Pages 795-804 in *ASME 2003 international design engineering technical conference and computers and information in engineering conference*. American Society of Mechanical Engineers (ASME), New York, New York, USA. <http://dx.doi.org/10.1115/detc2003/dac-48785>
- Teytaud, O. 2007. On the hardness of offline multi-objective optimization. *Evolutionary Computing* 15:475-491. <http://dx.doi.org/10.1162/evco.2007.15.4.475>
- Thomas, J. J., and K. A. Cook. 2005. *Illuminating the path: the research and development agenda for visual analytics*. IEEE Computer Society, Los Alamitos, California, USA.
- Tsoukiàs, A. 2008. From decision theory to decision aiding methodology. *European Journal of Operational Research* 87:138-161. <http://dx.doi.org/10.1016/j.ejor.2007.02.039>
- von Neumann, J., and O. Morgenstern. 1944. *Theory of games and economic behavior*. Princeton University Press, Princeton, New Jersey, USA.
- Webster, M., N. Santen, and P. Parpas. 2012. An approximate dynamic programming framework for modeling global climate policy under decision-dependent uncertainty. *Computational Management Science* 9:339-362. <http://dx.doi.org/10.1007/s10287-012-0147-1>
- White, C., B. S. Halpern, and C. V. Kappel. 2012. Ecosystem service tradeoff analysis reveals the value of marine spatial planning for multiple ocean uses. *Proceedings of the National Academy of Sciences* 109:4696-4701. <http://dx.doi.org/10.1073/pnas.1114215109>
- Woodruff, M. J., P. M. Reed, and T. W. Simpson. 2013. Many objective visual analytics: rethinking the design of complex engineered systems. *Structural and Multidisciplinary Optimization* 48:201-219. <http://dx.doi.org/10.1007/s00158-013-0891-z>
- Zeff, H. B., J. R. Kasprzyk, J. D. Herman, P. M. Reed, and G. W. Characklis. 2014. Navigating financial and supply reliability tradeoffs in regional drought management portfolios. *Water Resources Research* 50:4906-4923. <http://dx.doi.org/10.1002/2013WR015126>

Appendix 1

The appendix includes:

- 1.1 A detailed description of the multi-objective optimization procedure.
- 1.2 The impact of random seeds on the results.
- 1.3 Runtime dynamics of the multi-objective evolutionary algorithm.

1.1 Evolutionary Multiobjective Optimization

We use MOEAs to optimize the three problem formulations (P1, P2, and P3). The optimization finds the optimal strategies for single objective formulations (P1 and P2) and the Pareto approximate front that captures the trade-offs between all objectives for P3. Here, we use the recently developed and benchmarked MOEA, the BORG MOEA (Hadka and Reed 2013). While the advantage of using MOEAs lies in their ability to identify Pareto approximate fronts for multi-objective problems, the BORG MOEA can also be used to solve the single objective formulations such as P1 and P2. Solving the single objective formulations (P1 and P2) yields a single approximately optimal pollution strategy per formulation while solving the multi-objective formulation (P3) yields a wide range of strategies that span the trade-off between various objectives. If n_{p3} solutions form the Pareto approximate set for P3, optimizing P1, P2, and P3 together present the stakeholders with a total of $n_{p3}+2$ potential management strategies to choose from. This process of generation of alternative strategies for stakeholders to compare and contrast forms the first part of MORDM.

MOEAs are being used to solve many-objective problems in a rapidly growing body of literature as their population-based search enables the direct approximation of problems' Pareto frontiers in a single optimization run (Purshouse and Fleming 2003, Fleming et al. 2005, Aguirre et al. 2013, Lygoe et al. 2013, Morino and Obayashi 2013, Reed et al. 2013, Woodruff et al. 2013). For example, while solving a five-objective problem, the MOEAs simultaneously solve the five single-objective problems, ten two-objective problems, ten three-objective problems, and five four-objective problems. This has significant computational advantages over weight-based approaches to solve multi-objective problems that assign different weights to each objective and solve a single objective problem by maximizing the aggregated weighted objective. In order to capture the entire trade-off space, such approaches need to perform the optimization many times by varying the weights, which becomes intractable as the number of objectives goes beyond three (Teytaud 2007).

The BORG MOEA is able to solve high dimensional problems for non-linear threshold based models by adaptively using multiple (in this case six) search strategies. Search strategies refer to the operators that the algorithm uses to search the space of feasible solutions. Most multi-objective optimization algorithms employ a single search operator. But the BORG MOEA simultaneously uses six search strategies and assigns higher probability of use to a search strategy based on its ability to identify solutions on or close to the Pareto approximate front. Other features include random re-start of the search when no significant progress is observed (or the algorithm is trapped in a local optima), epsilon-dominance archiving (Laumanns et al. 2002), adaptive population sizing (Kollat and Reed 2007), and a steady state algorithm structure (Deb et al. 2005). Comparative analysis has demonstrated the efficacy of the BORG MOEA on a range

of test problems as well as real world problems many of which require optimization under uncertainty (Hadka and Reed 2012, Reed et al. 2013). The BORG MOEA is also relatively easy-to-use, has an underlying theoretical proof of convergence, and is highly scalable on parallel computing systems, thus increasing its potential usage across a wide variety of disciplines.

The framework used for performing multi-objective optimization of the lake problem is shown in Figure A1.1. In order to optimize under stochastic uncertainty, we randomly sample 100 SOWs out of 10000 SOWs for each evaluation of the lake problem in a function call of the BORG MOEA. In any given evaluation of a lake's pollution strategy (parameter 'a' in Equation 1), the relatively small number of SOWs (100) sampled is likely to yield 'noisy' objective function values. However, the small sample size drastically reduces computational demands. Evolutionary heuristics of the Borg MOEA underlying the Darwinian selection reward those solutions whose performance minimally varies across these small number of samples and have been demonstrated to be capable of maintaining high quality search in uncertain spaces (Miller and Goldberg 1996, Smalley et al. 2000, Reed et al. 2013). In other words, even if the BORG MOEA optimizes the objective function values across only 100 randomly sampled SOWs from the total space of 10000 SOWs at every function call, its underlying structure enables it to identify strategies that perform well across the entire ensemble of uncertain futures.

This framework therefore allows for multi-objective optimization under uncertainty (i.e., identifies a strategy that performs well over many SOWs drawn from the baseline well-characterized uncertainty distribution). The robustness analysis under well-characterized uncertainty tests each solution against all 10000 SOWs, serving as a validation of this computational savings using small number of samples. Moreover, all results presented in the study are objective performance re-evaluated against all 10000 SOWs. We employ a parallelized multi-master version of the BORG MOEA that runs on a cluster with 8 islands (different evolving populations that search the space) across 64 nodes to speed up the optimization procedure. The algorithm's search operator parameters were set at default values based on recommendations discussed in (Hadka and Reed 2013).

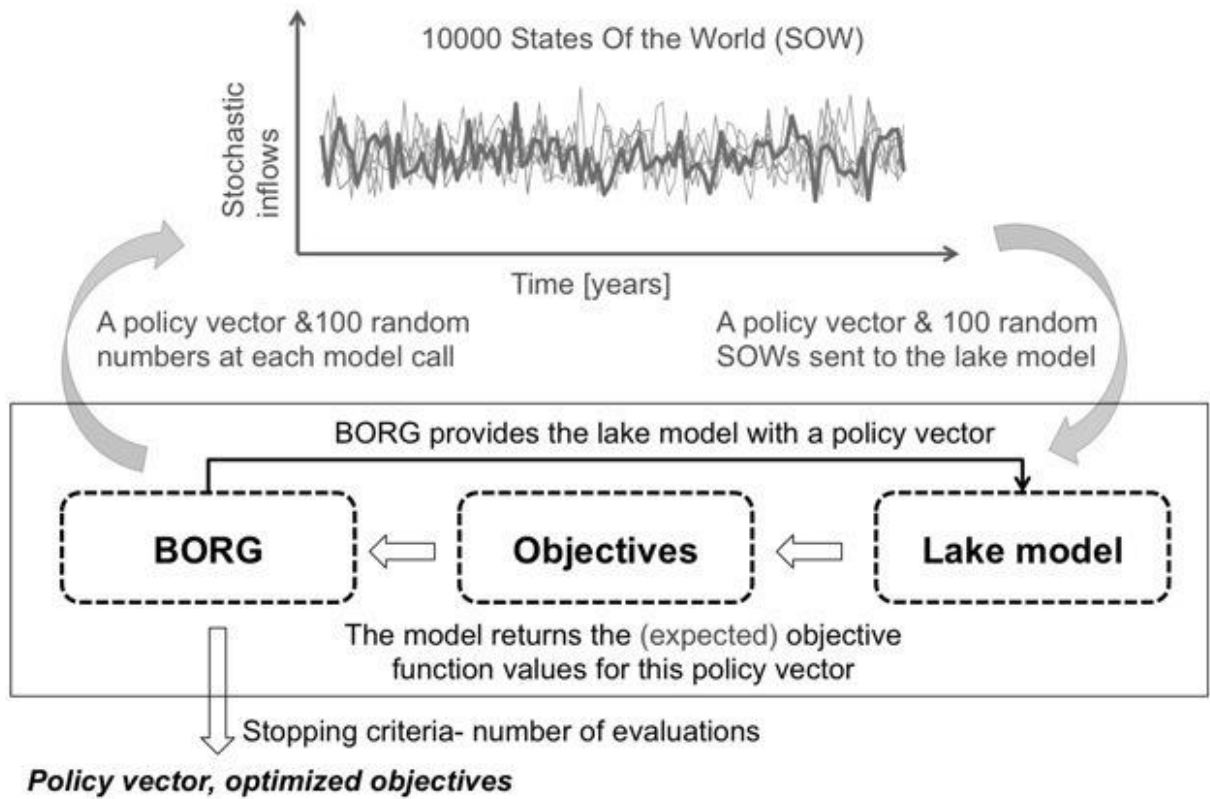


Figure A1.1 The robust optimization framework implemented using the BORG MOEA. This framework is used to optimize the various formulations of the lake model. The procedure within the solid black rectangle is repeated until a stopping criterion is met. Gray components show how the BORG MOEA implements optimization in the presence of stochastic uncertainty.

1.2 Impact of Random Seeds on Resulting Compromise Strategy

This analysis was carried out in order to establish the reliability of the results, i.e., to ensure that the conclusions of the study are independent of a random start of the algorithm.




Figure A1.2 Analyzing the impact of random seeds on the results. Using ten random seeds to start the BORG MOEA, ten different Pareto approximate fronts are obtained for the stochastic multi-objective formulation (P3). Each Pareto approximate front is evaluated to arrive at the *compromise* strategy using the minimum tolerable windows approach. Each strategy is associated with five objectives - expected utility, utility of current generation, utility of future generation, phosphorus in the lake, reliability. The strategy that maximizes the minimum objective among the across all strategies is identified as the *compromise* strategy. Figure shows the identified *compromise* strategies across ten random seeds, and the envelope bounding their lower and upper limits.

1.3 Runtime Dynamics of the BORG MOEA

The appendix presents the runtime dynamics of the algorithm. Runtime dynamics show that the algorithm converged to the Pareto approximate set within one million function evaluations.

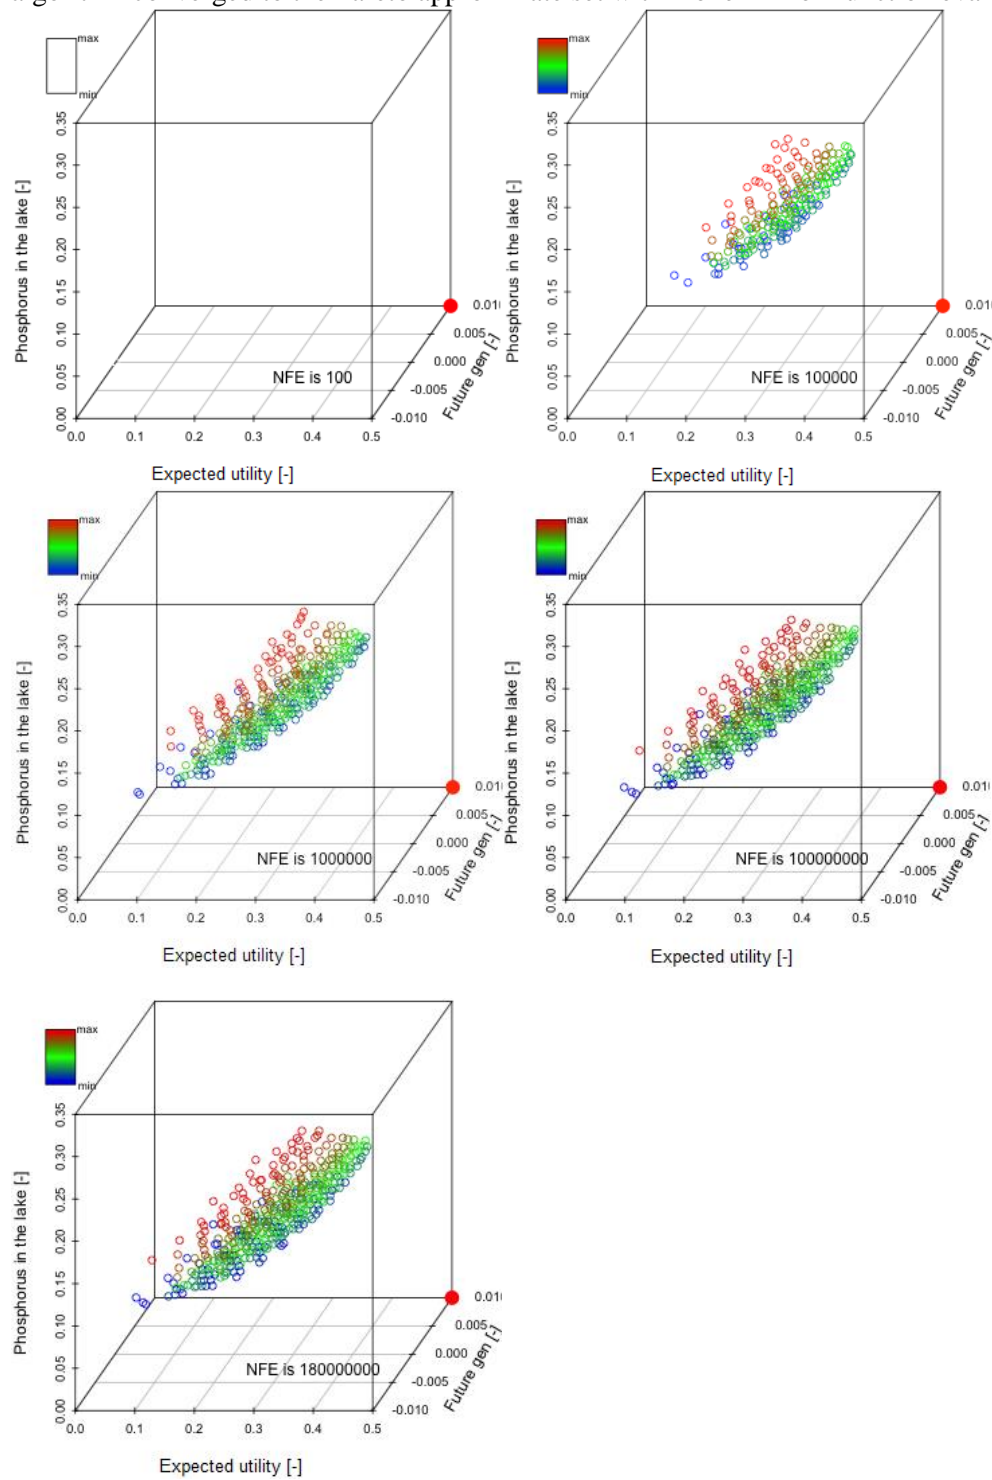


Figure A1.3 Runtime dynamics for the BORG MOEA - The figure shows the evolution of the Pareto approximate front with number of function evaluations (NFE) as the BORG MOEA explores the objective space. Each plot shows expected utility on the x-axis, utility of future generations on the y-axis,

amount of phosphorus in the lake on the z-axis. The color represents the utility of the first generation. Reliability objective is not shown here. The ideal point is shown as the red point on the bottom right corner. Snapshots of the algorithm's search are plotted at the following NFEs – one hundred, hundred thousand, one million, hundred million, and 180 million. The figure shows that the algorithm converges to the Pareto approximate front within 1 million function evaluations. For the results presented in this study, we ran the parallel version of the algorithm on eight nodes, each node running approximately 180 million evaluations.

References

- Aguirre, H., A. Oyana, and K. Tanaka. 2013. Adaptive ϵ -Sampling and ϵ -Hood for Evolutionary Many-Objective Optimization. *in* R. C. Purshouse, P. J. Fleming, C. M. Fonseca, S. Greco, and J. Shaw, editors. *Evolutionary Multi-Criterion Optimization*. Springer Berlin Heidelberg, Sheffield.
- Deb, K., M. Mohan, and S. Mishra. 2005. Evaluating the epsilon-Domination Based Multi-Objective Evolutionary Algorithm for a Quick Computation of Pareto-Optimal Solutions. *Evolutionary Computation* 13:501-525.
- Fleming, P., R. Purshouse, and R. Lygoe. 2005. - Many-Objective Optimization: An Engineering Design Perspective. - 3410.
- Hadka, D., and P. Reed. 2012. Diagnostic assessment of search controls and failure modes in many-objective evolutionary optimization. *Evol. Comput.* 20:423-452.
- Hadka, D., and P. Reed. 2013. Borg: An auto-adaptive many-objective evolutionary computing framework. *Evol. Comput.* 21:231-259.
- Kollat, J. B., and P. Reed. 2007. A framework for Visually Interactive Decision-making and Design using Evolutionary Multi-objective Optimization (VIDEO). *Environmental Modelling & Software* 22:1691-1704.
- Laumanns, M., L. Thiele, K. Deb, and E. Zitzler. 2002. Combining convergence and diversity in evolutionary multiobjective optimization. *Evol. Comput.* 10:263-282.
- Lygoe, R., M. Cary, and P. Fleming. 2013. A Real-World Application of a Many-Objective Optimisation Complexity Reduction Process. 7811.
- Miller, B. L., and D. E. Goldberg. 1996. *Optimal Sampling for Genetic Algorithms*. ASME Press.
- Morino, H., and S. Obayashi. 2013. Knowledge Extraction for Structural Design of Regional Jet Horizontal Tail Using Multi-Objective Design Exploration (MODE). 7811:656-668.
- Purshouse, R. C., and P. J. Fleming. 2003. Evolutionary many-objective optimisation: an exploratory analysis. Pages 2066-2073 Vol.2063 *in* *Evolutionary Computation, 2003. CEC '03. The 2003 Congress on*.
- Reed, P. M., D. Hadka, J. D. Herman, J. R. Kasprzyk, and J. B. Kollat. 2013. Evolutionary multiobjective optimization in water resources: The past, present, and future. *Advances in Water Resources* 51:438-456.
- Smalley, J. B., B. S. Minsker, and D. E. Goldberg. 2000. Risk-based in situ bioremediation design using a noisy genetic algorithm. *Water Resources Research* 36:3043-3052.
- Teytaud, O. 2007. On the hardness of offline multi-objective optimization. *Evol. Comput.* 15:475-491.
- Woodruff, M., P. Reed, and T. Simpson. 2013. Many objective visual analytics: rethinking the design of complex engineered systems. *Structural and Multidisciplinary Optimization* 48:201-219.

Appendix 2

The appendix includes:

1. Table summarizing the different choices made for each step in the lake problem (Table A2.1)
2. Table listing the various objectives analyzed before arriving at the five objective problem formulation (Table A2.2)

Table A2.1 Comparing various approaches towards managing a threshold-based ecosystem

SNo.	Steps in Figure 2	Choice(s) made for the lake problem
1.	Elicit system model, decisions, uncertainties, objectives, and constraints	Model: Lake model by Carpenter et al. (1999) Decision: time series of anthropogenic phosphorus input to the lake Uncertainties: standard (10000 SOWs) and deep (90000 SOWs) uncertainty Objectives: five objectives described in Appendix 3 Constraints: single constraint on reliability
2.	Test alternative problem formulations	Three alternative problem formulations selected for testing
3.	Identify tradeoffs under well characterized uncertainty	Tradeoffs identified by using the BORG MOEA described in Appendix 1
4.	Define robustness based on stakeholders' performance requirements	Definition adapted to satisfy multiple performance requirements under two assumptions definitions of uncertainty (Appendix 3)

Table A2.2 Various objectives analyzed before arriving at the five-objective formulation

SNo.	Objective	Rationale for including/excluding	References/ motivations
1.	Bentham's formulation of utility (expectation based approach)	Used in most analysis of the lake problem in literature	Carpenter et al. (1999), Brozovic and Schlenker (2011)
2.	Rawl's formulation of utility (max-min approach)	An alternative definition of utility used in some studies, removed due to mathematical challenges in optimizing this objective as it tends to solely focus on the worst case causing it to depend upon the chosen uncertainty representation	Rawls (1971), Tol (2000)
3.	Discounted financial benefits	An attempt to break up the utility function into its components, later discarded as: a. discounted losses are heavily correlated with objective (7), b. the standard MEU approach is lost	
4.	Discounted losses		
5.	Undiscounted expected utility of present stakeholders	Represent stakeholders separated in time without any discounting	Brundtland and Development (1987), Holling (1973)
6.	Undiscounted expected utility of future stakeholder		
7.	Average levels of phosphorus in the lake	Represents the preference to solely focus on the ecosystem under analysis	Admiraal et al. (2013)
8.	Reliability	Represents the preference to prevent irreversible changes in multistate ecosystems	Bennett et al. (2008), Carpenter and Lathrop (2008)

References

- Admiraal, J. F., A. Wossink, W. T. de Groot, and G. R. de Snoo. 2013. More than total economic value: How to combine economic valuation of biodiversity with ecological resilience. *Ecological Economics* 89:115-122.
- Bennett, E., S. Carpenter, and J. Cardille. 2008. Estimating the Risk of Exceeding Thresholds in Environmental Systems. *Water, Air, and Soil Pollution* 191:131-138.
- Brozovic, N., and W. Schlenker. 2011. Optimal management of an ecosystem with an unknown threshold. *Ecological Economics* 70:627-640.
- Brundtland, G. H., and W. C. o. E. a. Development. 1987. *Our common future : report of the World Commission on Environment and Development*. Oxford University.
- Carpenter, S., D. Ludwig, and W. A. Brock. 1999. Management of Eutrophication for Lakes Subject to Potentially Irreversible Change. *Ecological Applications* 9:751-771.
- Carpenter, S. R., and R. C. Lathrop. 2008. Probabilistic Estimate of a Threshold for Eutrophication. *Ecosystems* 11:601-613.
- Holling, C. S. 1973. Resilience and Stability of Ecological Systems. *Annual Review of Ecology and Systematics* 4:1-23.
- Rawls, J. 1971. *A Theory of Justice* (original edition). Harvard University Press.
- Tol, R. J. 2000. Equitable Cost-Benefit Analysis of Climate Change. Pages 273-290 *in* C. Carraro, editor. *Efficiency and Equity of Climate Change Policy*. Springer Netherlands.

Appendix 3

The appendix includes:

- 3.1 The lake model dynamics: A detailed description of the lake model dynamics.
- 3.2 Objective functions for the lake problem: A detailed description of the objective functions used to evaluate candidate strategies.
- 3.3 Robustness metric description: A mathematical description of the robustness metric.

3.1 The lake model dynamics

The lake can exist in two states: oligotrophic or eutrophic. In the oligotrophic state, the lake has low concentrations of phosphorus with clear water. In the eutrophic states, the phosphorus concentration is high and algae can bloom. In the eutrophic state, the lake is assumed to be unable to support fisheries, or tourism, and also to be severely degraded in aesthetic value. It is also much harder to revert the lake from the eutrophic to the oligotrophic state in a short time by reducing pollution alone. Therefore, the transition to eutrophic state has multiple disadvantages, besides loss of economic activity.

Depending upon the ease with which a lake in eutrophic state can be brought back to its oligotrophic state, lakes can be classified as reversible, hysteretic, or irreversible. Irreversible lakes are most vulnerable since it is impossible to bring them back to an oligotrophic stage by reducing phosphorus concentrations alone once they exceed a threshold. In this study, the parameters of the lake model are such that the lake is irreversible and therefore represents an ecosystem with two possible states. Once the lake turns eutrophic, it is not possible to return it to an oligotrophic state by reducing phosphorus inputs alone. In reality, these conditions are most likely to occur in shallow lakes, lakes in phosphorus rich regions, or lakes that have received extreme phosphorus inputs for an extended period of time.

In the simple model, the parameter b determines whether the lake is reversible, hysteretic or irreversible for a given value of the recycling parameter q . Higher values of b suggest a lake that has a high capacity to remove pollution and vice-versa. If recycling occurs, the concentration of phosphorus in the lake increases suddenly over a period of time, the rate of this change is governed by the recycling parameter q . Higher values of q correspond to fast transitions and vice-versa. We adopt this formulation from the pioneering study by Carpenter et al. (1999). Carpenter et al. (1999) also provides a very careful and much more detailed description for the lake (model) system. Table 1 lists the parameter values for the lake model used in our study.

3.2 Objective functions for the lake problem

We begin with a widely used objective in the analysis of the lake problem - the expectation of discounted net present value of utility (O_I) given by,

$$O_1 = \frac{1}{N} \sum_{i=1}^N V_i, \text{ where,} \quad (3.1)$$

$$V_i = \sum_{t=0}^T \delta^t u_{t,i}, \text{ and,} \quad (3.2)$$

$$u_{t,i} = \alpha a_{t,i} - \beta X_{t,i}^2. \quad (3.3)$$

In these equations, V_i (to be maximized) is the discounted net present value of utility for the i^{th} SOW, $u_{t,i}$ is the utility, $a_{t,i}$ is the allowed anthropogenic pollution and $X_{t,i}$ is the level of phosphorus in the lake at time step t and i^{th} SOW. The economic parameters α and β capture the willingness to pay for pollution and the compensation lake users are willing to accept to tolerate a given state of the lake respectively. α and β are fixed at 0.4 and 0.08 respectively following the analysis in Carpenter et al. (1999). For simplicity, we neglect the considerable uncertainty about the values for the discount rate and the economic parameters (Chichilnisky 1996, Dasgupta 2008). Note that the case for uncertain discounting was analyzed for the lake model by Ludwig et al. (2005). The discount factor, δ translates future to present utilities. The shortened term ‘expected utility’ is used to refer to this objective in the text and figures.

The time index, t , varies from 1 to T years ($T = 100$ years), and there are N SOWs. The SOWs are sampled from the lognormal distribution in Equation (1) and their total number (N) varies from 0 to 90000 based on the type of uncertainty being considered as described in the section on ‘Uncertainty’. N is 0 for the deterministic case, 10000 for well-characterized uncertainty and 90000 for deep uncertainty. The allowed anthropogenic pollution flow a_t is only decision variable that controls the objective function. The stakeholder can change a_t only every 5 years. As a result, there are 20 planning periods across a planning horizon of 100 years and the optimization framework needs to identify the 20 values of a_t that satisfy selected stakeholders’ objectives.

To contrast the strategy that maximizes the expected utility (O_1), we introduce additional objectives that represent stakeholders that more strongly focus on the long term environmental quality of the lake or are varied in their inter-temporal presence. Stakeholders often assess outcomes using a diverse set of objectives (Kasprzyk et al. 2009, McInerney et al. 2012, White et al. 2012, Herman et al. 2014). Farber et al. (2006) for example, argue that the linking of ecology and economics requires identification of ecosystem services that are likely to be in conflict. Our objective formulation is to a large part motivated by this assessment.

Identifying key objectives that represent diverse stakeholders is challenging and a potentially iterative process. Some of these objectives are a proxy for ecosystem services (recreation, fishery), while others serve as proxies for alternative perspectives with regard to valuing economic services (utility). This approach can be interpreted as representing the perspectives of five hypothetical stakeholder groups in the fictitious town. This resulted in the following objectives considered in our analysis:

1. Minimize the average level of phosphorus in the lake (O_2) – Admiraal et al. (2013) point out that the utility function is strongly biased towards anthropogenic services which is a key limitation in identifying ecosystem management strategies that adequately protect environmental values. Here, we introduce this objective to represent a regulatory

perspective related to an indicator of the health of the lake. This objective can be interpreted as one key concern of individuals that aim to preserve the lake as it is and therefore their sole goal is to reduce the levels of phosphorus in the lake. The objective function is,

$$O_2 = \frac{1}{NxT} \sum_{i=1}^N \sum_{t=1}^T X_{t,i}, \quad (3.4)$$

where, $X_{t,i}$ is the phosphorus in the lake at time step, t and i^{th} SOW. This objective aims to minimize the average levels of phosphorus in the lake.

2. Maximize the expected utility of the present stakeholders (O_3) – This objective represents utility of the current stakeholders. The objective function is

$$O_3 = \frac{1}{N} \sum_{i=1}^N U_{1,i}, \quad (3.5)$$

where, $U_{1,i}$ is the utility of the first year in the 100 year planning horizon for the i^{th} SOW and is to be maximized.

3. Maximize the expected utility of the future stakeholders (O_4) – The objective was motivated by the definition of sustainability adopted by past studies (Holling 1973, United Nations 1987, Cato 2009). These definitions represent the interest of present and future generations quite differently than the discounted expected utility framework. While discounting has been the classic approach to analyze inter-temporal trade-offs, several reports, even governmental decisions have been based on objectives that are not subject to discounting. One simple example is the design of flood defenses in the Netherlands that are subject to an acceptable level of risk (Jonkman 2013). Therefore, we explicitly model inter-temporal stakeholders in separate objective functions. To approximate this perspective, we choose two example stakeholder groups (i) current generation and (ii) generations in the far future. (Far here is represented as the second half of the planning horizon of the problem). This objective represents the utility of future stakeholders who exist in the last 50 years of the 100-year planning horizon. The objective function is

$$O_4 = \frac{1}{N} \sum_{i=1}^N U_{50-100,i}, \quad (3.6)$$

where, $U_{50-100,i}$ is the sum of undiscounted utilities for the generations spanning years 50 to 100 in the 100-year planning horizon for the i^{th} SOW. This objective function is to be maximized.

4. Maximize reliability (O_5) – One of the goals of this study is to capture the behavior of multi-state ecosystems when some states are far less preferable to the stakeholder. The reliability objective seeks to ensure that the lake remains below critical pollution levels to avoid eutrophication. This objective also represents key concerns of stakeholders who either depend directly on the ecosystem services provided by the lake, or those who aim to maintain the ecosystem itself while being able to accept some levels of pollution. In

addition, this formulation approximates a common risk-based engineering metric that has been widely employed across many contexts (Hashimoto et al. 1982). Maximizing the reliability of avoiding a tipping point response captures the strong aversion to irreversible losses of key economic and ecosystem services. The objective is,

$$O_5 = \frac{1}{NxT} \sum_{i=1}^N \sum_{t=1}^T \theta_{t,i} \quad (3.7)$$

$$\text{where, } \theta_{t,i} = \begin{cases} 1 & \text{if } X_{t,i} < X_{crit} \\ 0 & \text{if } X_{t,i} \geq X_{crit} \end{cases}$$

In these equations, $\theta_{t,i}$ is the reliability index which is 1 if the level of phosphorus in the lake ($X_{t,i}$) is below the specified critical threshold (X_{crit}) and 0 otherwise. The critical threshold is set at 0.5 based on the parameters of the lake model. X_{crit} is the minimum steady state pollution value at which the lake transitions from an oligotrophic to eutrophic state. A reliability of 1 represents a pollution strategy that successfully keeps the phosphorus levels in the lake below the specified critical thresholds across the entire planning horizon and across all SOWs. Table 3 lists the objectives used in this study.

3.3 Robustness metric description

We define performance requirements for key variables and a strategy that equal or exceeds these requirements across a range of uncertain scenarios is considered to be robust. An overall measure of robustness is thus defined as –

$$\%Robust = \left(\frac{1}{N} \sum_{i=1}^N r_i \right) \times 100, \quad (3.8)$$

$$\text{where, } r_i = \begin{cases} 1 & \text{when } P_{o,i} \geq \text{requirement}_o \quad \forall o \\ 0 & \text{otherwise} \end{cases}$$

Here, r_i is 1 if the strategy performs above all requirements otherwise 0, $P_{o,i}$ is the value of the o^{th} variable of interest under the i^{th} SOW, requirement_o is the performance requirement for the o^{th} variable, and N is the total number of SOWs (10000 for well-characterized uncertainty and 90000 for deep uncertainty).

The performance requirements represent criteria that actual stakeholders may consider as performance levels that could not be compromised further. For example, stakeholders are likely to have a factor of safety associated with the critical phosphorus levels. Here, we fix that factor of safety at 0.75. Similarly, a high reliability of keeping the lake in the oligotrophic state is binding due to obvious economic and environmental consequences. Thus, the performance requirement on reliability was fixed at 99%. While maintaining the lake in an oligotrophic state is important, a minimum level of economic activity is also required. This level was set at 50% of the value of expected utility obtained in the optimal strategy for expected utility maximization (P2). Our proposed definition of robustness is illustrative and the MORDM framework is highly flexible in accommodating alternative definitions. The primary intent of our example is to emphasize that system performance requirements are themselves likely to be multi-objective,

complex in their effects on filtering solutions, and should be carefully elicited in any real application of MORDM.

Note that our proposed definition of robustness spans multiple objectives and hence, prevents a stakeholder heavily biased towards one objective (say utility) from selecting strategies that favor their preferred objective. For the robustness index of a strategy to be high, all performance criteria need to be simultaneously satisfied across multiple SOWs. So, if a high performance requirement is selected for the utility function, strategies satisfying it may not satisfy the reliability or phosphorus requirements. Thus, it is likely that none of the strategies emerges as robust forcing stakeholders to revise their threshold specifications.

References

- Admiraal, J. F., A. Wossink, W. T. de Groot, and G. R. de Snoo. 2013. More than total economic value: How to combine economic valuation of biodiversity with ecological resilience. *Ecological Economics* 89:115-122.
- Carpenter, S., D. Ludwig, and W. A. Brock. 1999. Management of Eutrophication for Lakes Subject to Potentially Irreversible Change. *Ecological Applications* 9:751-771.
- Cato, M. S. 2009. *Green Economics: An Introduction to Theory, Policy and Practice*. Earthscan.
- Chichilnisky, G. 1996. An axiomatic approach to sustainable development. *Social Choice and Welfare* 13:231-257.
- Dasgupta, P. 2008. Discounting climate change. *Journal of Risk and Uncertainty* 37:141-169.
- Farber, S., R. Costanza, D. L. Childers, J. Erickson, K. Gross, M. Grove, C. S. Hopkinson, J. Kahn, S. Pincetl, A. Troy, P. Warren, and M. Wilson. 2006. Linking Ecology and Economics for Ecosystem Management. *BioScience* 56:121-133.
- Hashimoto, T., J. R. Stedinger, and D. P. Loucks. 1982. Reliability, resiliency, and vulnerability criteria for water resource system performance evaluation. *Water Resources Research* 18:14-20.
- Herman, J. D., H. B. Zeff, P. M. Reed, and G. W. Characklis. 2014. Beyond optimality: Multistakeholder robustness tradeoffs for regional water portfolio planning under deep uncertainty. *Water Resources Research* 50:7692-7713.
- Holling, C. S. 1973. Resilience and Stability of Ecological Systems. *Annual Review of Ecology and Systematics* 4:1-23.
- Jonkman, S. N. 2013. Advanced flood risk analysis required. *Nature Clim. Change* 3:1004-1004.
- Kasprzyk, J. R., P. M. Reed, B. R. Kirsch, and G. W. Characklis. 2009. Managing population and drought risks using many-objective water portfolio planning under uncertainty. *Water Resources Research* 45:W12401.
- Ludwig, D., W. A. Brock, and S. R. Carpenter. 2005. Uncertainty in discount models and environmental accounting. *Ecology and Society* 10(2): 13.
- McInerney, D., R. Lempert, and K. Keller. 2012. What are robust strategies in the face of uncertain climate threshold responses? *Climatic Change* 112:547-568.
- Nations, U. 1987. *Our common future - Brundtland Report*. Page 204. Oxford University Press, Oxford; New York.
- White, C., B. S. Halpern, and C. V. Kappel. 2012. Ecosystem service tradeoff analysis reveals the value of marine spatial planning for multiple ocean uses. *Proceedings of the National Academy of Sciences* 109:4696-4701.

Appendix 4

The appendix illustrates the trade-off between each pair of objectives in the P3 problem formulation.

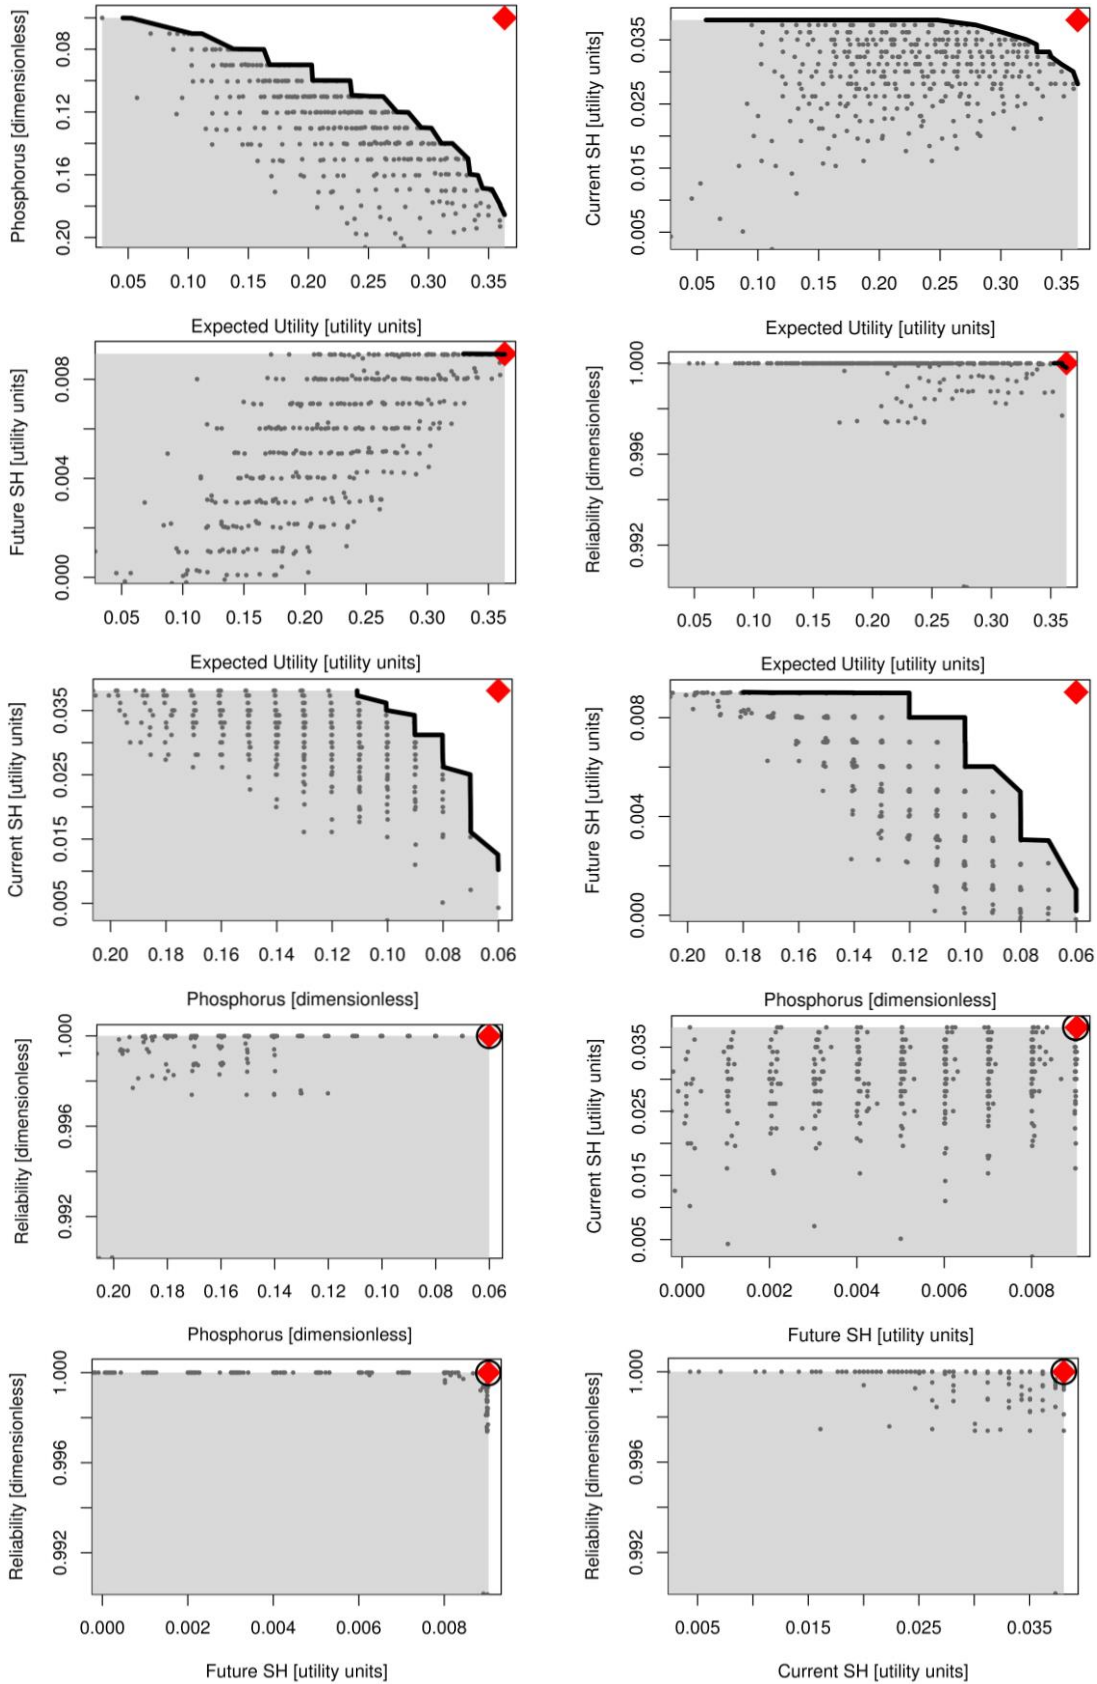


Figure A4.1 Trade-offs between each pair of objectives for the five-objective formulation P3. The objectives are – phosphorus in the lake (minimize), expected utility (maximize), expected utility of the present generation (maximize), expected utility of the future generations (maximize), and reliability (maximize). The 399 solutions from the five objective optimization are plotted as gray points. Nondominated sorting is carried out for each pair of objectives across the 399 points to identify the Pareto approximate front for each pair of objectives. The red diamond represents the ideal point. If there is tension between the two objectives, a front is identified and plotted as the black line. If there is no tension between two objectives, it implies that both can be simultaneously optimized and the ideal point is attainable, shown by the black circle around the red diamond. Gray shading represents the dominated region; solutions in this region are inferior to those at the Pareto approximate front. Note that all 399 solutions are nondominated w.r.t. each other in the five-dimensional objective space, but when the space is collapsed to two dimensions, the nondominated sorting is carried out again to identify the new Pareto front for two objectives.