

Is it Ethical to Use a Single Probability Density Function?

Robert J. Lempert, David G. Groves, Jordan R. Fischbach

WR-992-NSF

May 2013

Prepared for the National Science Foundation

RAND working papers are intended to share researchers' latest findings and to solicit informal peer review. They have been approved for circulation by the RAND Justice, Infrastructure and Environment Division but have not been formally edited or peer reviewed. Unless otherwise indicated, working papers can be quoted and cited without permission of the author, provided the source is clearly referred to as a working paper. RAND's publications do not necessarily reflect the opinions of its research clients and sponsors. RAND® is a registered trademark.



Abstract

Addressing climate change requires both quantitative analysis and ethical reasoning. But the character of many climate-related decisions – in particular deep uncertainty, competing values, and complex relationships among actions and consequences – can make it difficult to use the most common types of quantitative analysis to support appropriate ethical reasoning. This essay will describe Robust Decision Making (RDM), one of a new class of methods for quantitative analysis that can help resolve some of these difficulties, and situate RDM in the framework for ethical reasoning and deliberation laid out in Amartya Sen’s *Idea of Justice*. Two example applications will illustrate these ideas: one examining how greenhouse gas mitigation policies might appropriately address the threat of abrupt climate change and the other addressing management of the Colorado River in the face of climate change and other changing conditions.

Acknowledgements

The authors thank Nancy Tuana, Director of the Rock Ethics Institute at Penn State, and RAND colleague Evan Bloom for his advice and comments in preparing this essay. This paper was prepared for the Princeton Institute for International and Regional Studies workshop on “The Ethics of Risk and Climate Change,” April 12-13, 2013 and was supported in part by the Network for Sustainable Climate Risk Management (SCRiM) under NSF cooperative agreement GEO-1240507.

INTRODUCTION

Addressing climate change requires both quantitative analysis and ethical reasoning. Ethical reasoning is required because climate change will affect many people in current and future generations. Quantitative analysis is required to understand how today's actions might influence those potential consequences. But the character of many climate-related decisions – in particular deep uncertainty, competing values, and complex relationships among actions and consequences – can make it difficult to use the most common types of quantitative analysis to support appropriate ethical reasoning. New approaches to quantitative analysis have, however, become available that resolve some of these difficulties.

In particular, quantitative analyses traditionally summarize scientific understanding with best-estimate predictions of the future. In the more sophisticated analyses such predictions take the form of a single probability density function (pdf) over plausible future states of the world. A pdf describes the relative likelihood of these different states of the world. Such pdfs, combined with a utility function that describe peoples' preferences, can then be used to rank the desirability of alternative decision options. But under conditions of deep uncertainty, such best-estimate predictions can promote overconfidence, so that decision makers fail to give sufficient attention to important, but perceived unlikely, threats and opportunities. Best-estimate predictions can also inhibit deliberations among individuals with differing expectations and values, because by definition best-estimate predictions privilege some expectations and values over others. Nonetheless, such deliberation lies at the core of the ethical reasoning appropriate to address climate change.

This essay will describe Robust Decision Making (RDM) (Lempert *et al.* 2003; Lempert *et al.* 2007), one of a new class of methods for quantitative analysis that support good decisions without predictions. It will situate these methods in the framework laid out in Amartya Sen's *Idea of Justice* (Sen 2009), arguing that for climate change, RDM and related methods can support Sen's proposed framework for ethical reasoning more effectively than can prediction-based quantitative analyses. This essay will then describe two example RDM applications: one a "table top" exercise on how greenhouse gas mitigation policies might appropriately address the threat of abrupt climate change (Hall *et al.* 2012) and the other an analysis used to facilitate discussions among Colorado River users and stakeholders on how the river might be managed in

the future in the face of climate change and other changing conditions as a part of the U.S. Bureau of Reclamation’s Colorado River Basin Study (Reclamation 2012; Groves *et al.* 2013). The essay will close with thoughts on its title, arguing that in most cases ethically informed decision making about climate change is better facilitated by analytic methods that use a set of probability density functions rather than any single pdf.

I. RUNNING THE ANALYSIS BACKWARDS¹

RDM rests on a simple concept. Rather than using computer models and data to describe a best-estimate future, RDM runs models over hundreds to thousands of different sets of assumptions to describe how plans perform in many plausible futures. The approach then uses statistics and visualizations on the resulting large database of model runs to help decisionmakers distinguish those future conditions where their plans will perform well from those in which they will perform poorly. This information can help decisionmakers identify, evaluate, and choose robust strategies—ones that perform well over a wide range of futures and that better manage surprise.

Traditionally, quantitative analysis provides decision makers with information about the future by making predictions. For example, computer simulation models might predict whether a water agency’s investments can ensure its customers reliable, cost-effective supply. Or a model might predict the greenhouse gas mitigation path that best balances the cost and benefits of emission reductions. Such approaches, sometimes referred to in the literature as “science first” or “predict-then-act,” proceed with a particular ordering of steps: Assemble the available evidence into best-estimate predictions of the future and then suggest the best course of action given these predictions. Such methods can work well when predictions are accurate and noncontroversial. Otherwise, the approach can encourage either gridlock or overconfidence: the former when participants to a decision argue over projections rather than solutions, and the latter when they choose solutions that fail when the future turns out differently than expected (Herrick *et al.* 2000; Rayner 2000; Weaver *et al.* 2013).

¹ This section draws, often verbatim, from RAND Research Highlight RB-9791, *Making Good Decisions Without Predictions: Robust Decision Making for Planning Under Deep Uncertainty*.

RDM avoids these problems by running the analysis “backwards.” Rather than start with predictions, the approach begins with one or more plans under consideration (often a current or best-estimate plan) and uses multiple runs of the model to identify the futures most relevant to the plan’s success. The concept of stress-testing policies against scenarios is clearly not new. But new capabilities in information technology help RDM facilitate such reasoning. Just as people routinely use search engines like Google to scan millions of websites to find information that interests them, so RDM uses computer search and visualization on databases of many model runs to help decisionmakers to address such questions as: What are the key characteristics that differentiate those futures in which a plan succeeds from those in which it fails? What steps can be taken so a plan may succeed over a wider range of futures?

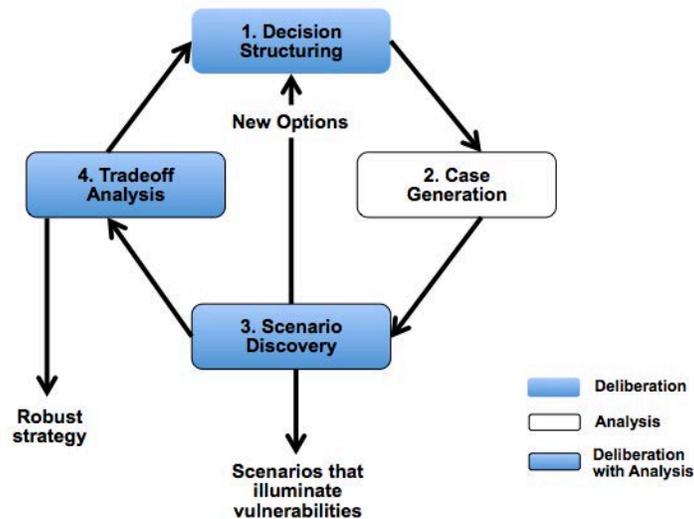
RDM is one of a set of decision support approaches that begin with a specific decision under consideration by a specific community of users and then use questions relevant to these decisions to organize information about future climate and socio-economic conditions . The literature offers several names for such approaches, including “context-first” (Ranger *et al.* 2010), “decision-scaling” (Brown 2010), and “assess risk of policy” (Lempert *et al.* 2004; Carter *et al.* 2007; Dessai *et al.* 2007). All share the central idea of beginning with a proposed policy or policies; identifying future conditions where the policy fails to meet its goals; and then organizing available information about the future to help policy makers to identify potential policy responses to those vulnerabilities and to decide whether and when to adopt these responses.

The decision support literature make clear that analytics –the discovery and communication of meaningful patterns in quantitative information – proves most effective when closely linked to users’ needs and values (NRC 2009). RDM’s analytics are specifically designed to improve conversations among decisionmakers under conditions of deep uncertainty. Deep uncertainty occurs when the parties to a decision do not know—or agree on—the best model for relating actions to consequences nor the likelihood of future events (Lempert *et al.* 2003).

RDM follows a “deliberation with analysis” process of decision support, which can prove particularly appropriate with diverse decision makers whose goals emerge from discussions among them and which may change over time (NRC 2009). The process begins with many participants to a decision working together to define their objectives and other parameters.

Participants next engage with experts to generate and interpret decision-relevant information and then revisit choices and objectives based on this information. RDM adds to this general approach the concepts of running the analysis backwards – that is, beginning with a proposed decision – and testing plans against many different plausible futures.

Figure 1
Iterative, Participatory Steps of an RDM Analysis



RDM’s process, as shown in Figure 1, begins with a decision structuring exercise in which decisionmakers define the goals, values, uncertainties, and choices under consideration (Step 1). Analysts next use computer models to generate a large database of runs (Step 2), where each such case represents the performance of a proposed policy in one plausible future. In a process called “scenario discovery,” (Groves *et al.* 2007; Bryant *et al.* 2010; Lempert 2013) computer visualization and statistics on this database then help decisionmakers identify clusters representing scenarios that illuminate vulnerabilities of the policies (Step 3). These scenarios can then help decisionmakers identify potential new ways to address those vulnerabilities (back to Step 1) and/or evaluate through tradeoff analysis whether these choices are worth adopting (Step 4). In particular, robust strategies are often adaptive, that is, designed to adjust over time to new information (Rosenhead 2001; Walker *et al.* 2010). By providing detailed understanding of the vulnerabilities of proposed strategies, RDM often helps decision makers to design such adaptive strategies (Lempert *et al.* 2003; Lempert *et al.* 2010).

The process shown in Figure 1 continues until decisionmakers settle on a robust strategy. As key deliverables, the analysis provides scenarios illuminating the vulnerabilities of proposed strategies, insights into the design of more robust strategies, and visualizations highlighting the irreducible tradeoffs among alternative robust strategies.

The examples in Section III below illustrate this process.

RDM fosters a new conversation among decisionmakers by combining the best features of two traditional approaches to uncertainty management: scenarios and probabilistic risk analysis. Scenarios describe potential future conditions. By presenting a set of plausible and contrasting futures worthy of consideration, scenarios can help break down the cognitive and organizational barriers that often limit decisionmakers' ability to consider a sufficiently wide range of potential futures and alternative decision options (Schoemaker 1993). But with traditional approaches, the choice of a small number of scenarios to summarize many futures can appear arbitrary or biased (Vantklooster *et al.* 2006; Parson *et al.* 2007; Parson 2008).

In contrast, probabilistic risk analysis, a predict-then-act approach, uses quantitative predictions of risk (often defined as the predicted probability multiplied by the predicted consequence of an event) to systematically inform decisions about the allocation of efforts to reduce risk. While often useful when uncertainties are well understood, the approach faces the perils of prediction when uncertainties are deep.

RDM draws from both scenarios and probabilistic risk analysis by running the latter backwards to ask which policies reduce risk over which range of assumptions, inquiring, for example, "What assumptions would we need to believe were true for us to reject option A and instead choose option B?" Doing so leads to the emergence of high- and low-risk scenarios that, as our example cases below show, can prove useful in contentious policy debates. These scenarios can also inform creative thinking on robust, often adaptive, plans and help decisionmakers assess tradeoffs among these plans, as shown in Figure 1.

RDM uses computer models in a fundamentally different way than predict-then-act analyses. The latter regard models as representations of reality that are sufficiently accurate to recommend the best response to an uncertain future. In contrast, RDM regards models as mapping assumptions to consequences without privileging one set of assumptions over another

(Bankes 1993).² Often RDM can significantly enhance the value of decisionmakers' existing models (designed for predict-then-act analysis) by running them numerous times to identify vulnerabilities and to find plans robust over many combinations of assumptions.

By embracing many plausible futures within a quantitative analysis, RDM can help reduce overconfidence and the deleterious impacts of surprise, can systematically include imprecise information in the analysis, and can help decisionmakers and stakeholders with differing expectations about the future reach a well-grounded consensus on action, even when uncertainties are deep.

II. FACILITATING DELIBERATION

In his book, the *Idea of Justice*, Amartya Sen (2009) lays out principles of ethical reasoning that seem particularly appropriate for the discussion here. Sen argues that ethnically defensible policy judgments must pay close attention to their actual effect on the way people live. That is, while ethnically defensible policy judgments should consider a society's institutional arrangements and the processes by which it makes choices (for instance, the extent to which it respects individual rights), such judgments cannot fail to emphasize the consequence of choices on peoples' opportunities and well-being. Sen measures such consequences by focusing on what he calls 'capabilities,' that is, the ability of people to live meaningful lives. Sen's work on capabilities has led to measures such as the Human Development Index, widely used to measure social well-being. Important schools of utilitarian philosophy reject any ability to make meaningful comparisons of the well-being of different individuals, because one individual can never know for certain what another is thinking and feeling. In response, Sen notes that while capabilities measures are intrinsically uncertain and generally unable to produce a definitive ranking of more or less just outcomes, they can nonetheless distinguish some outcomes as more just than others, and support meaningful debates about why some outcomes are more just than others.

² Bankes distinguishes between what he calls "consolidative" models that make verifiable statements about world and thus are expected to reliably predict real world outcomes and what he calls "exploratory" models that make unverifiable statements about the world (for instance about political, economic, and climate conditions fifty years hence) and thus can only be judged by their consistency with available information and utility for decision making.

The *Idea of Justice* critiques Sen's teacher John Rawls. Rawls' *The Theory of Justice* (1971) describes how a diverse community could pursue justice by first constructing a common vision of a perfectly just social arrangements. While agreeing with Rawls' principle of justice as fairness, Sen argues that establishing such a transcendental vision of justice is generally impossible. The level of agreement and commonality of values it presupposes does not (and should not) exist in a diverse society where people are free to pursue their lives as they see fit and will likely have varying priorities, goals, and values.

Fortunately, Sen also argues that transcendental agreements are neither necessary nor sufficient to the pursuit of justice. Transcendental agreement is not sufficient because even if people could agree on a common vision of justice, human knowledge is too fallible and the uncertainties too deep to ever know what specific set of institutions and rules would steer today's decidedly imperfect world towards that perfectly just society. In our complex, uncertain world knowing one's ultimate goal provides little information about which of the available next steps is most promising. We must therefore make choices between different non-ideal options. Sen claims transcendental agreement is not necessary because people with diverse views and highly imperfect knowledge can nonetheless pursue a more just society by comparing the actual choices before them. This relational approach allows individuals to consider to the best of their abilities the consequences of proposed actions on people's capabilities, debate which set of consequences are distinguishably more just than others, and thus step-by-step proceed from less to more just states of the world.

Public deliberation proves central to this process of social choice. Citizens exchange views and debate their supporting reasons for policy preferences. Debate is important because the reasons and values underlying individuals' preferences are often as important as the preferences themselves in any judgments about just social choices. Such deliberations work best when they focus on the relational, not the transcendental; recognize the inescapable plurality of competing views; facilitate re-examination and iterative assessments; demand clear explication of reasoning and logic; and recognizes an "open impartiality" that accepts the legitimacy and importance of the views of others, both inside and outside the community of interest to the immediate policy discussion.

If applied appropriately, RDM, and the class of “context-first” methods it represents, appears well-tailored to provide analytic support to this approach to ethical reasoning.

The deliberation with analysis process – which aims to provide quantitative information that helps decision makers reflect on their goals, options, and the connections among them – appears compatible with the deliberative process advocated by Sen. RDM’s use of computer simulation models as devices for relating assumptions to consequences, rather than for making predictions based on universally agreed upon certitudes, appears consistent with Sen’s view of the multiplicity of legitimate views and the uncertainties surrounding any discussion linking actions to social realizations, that is to “what emerges in the society, including the kind of lives that people can actually lead, given the institutions and rules, but also other influences, including actual behavior, that would inescapably affect human lives” (Sen 2009, p. 10). RDM’s focus on presenting decision makers with near-term steps that, as part of a long-term adaptive strategy, could lead robustly towards desired outcomes is not only consistent with Sen’s emphasis on a relational rather than transcendental approach to justice, but could provide a method for helping to achieve it.

As a key feature, RDM enables systematic quantitative reasoning about the consequences of and tradeoffs among alternative decision options using multiple values and multiple expectations about the future, whether each of these expectations are embodied in point projections or single pdfs over multiple future states of the world (Kasprzyk *et al.* 2013). Climate change, as with many contentious policy debates, entangles facts and values. Such debates often intermix individuals’ values, expectations, and policy preferences, which can make it difficult to discuss one without also engaging the others (Kahan *et al.* 2006). Predict-then-act analyses require as a starting point consensus on expectations and model structure (that is, how actions relate to consequences) before engaging with the analytics. Beginning with predictions can prove appropriate when the science is inarguable. For many pollutants, for instance, the scientific evidence of their impacts is sufficiently strong to provide a fixed point in any debate over proper policy responses. But many climate-related policy judgments depend on uncertain, incomplete, and fast-changing science – for instance, whether residences ought to relocate from particular coastal communities because of changing storm frequencies, or whether society ought to risk potential economic dislocation from government action that seeks to hold atmospheric

greenhouse gas concentrations below some threshold level. In such situations, demanding a consensus and agreement on model structure prior to deliberation can disenfranchise those whose views differ from the consensus or, perhaps more frequently, disconnect the quantitative analysis from the deliberation.

In contrast, RDM respects the fact of differing preferences and value orientations by embracing the concept of a multiplicity of plausible futures. As in scenario planning, this concept provides people with differing values and expectations multiple entry points into the analytics, and allows them to explore the consequences of potential futures before committing themselves to acting on those consequences. RDM implements the decision criterion of robustness, which allows people to agree on near-term actions without agreeing on values and expectations. The approach thus appears able to provide systematic quantitative reasoning in support of deliberative processes such as those advocated by Sen.

III. EXAMPLES

Two examples help illustrate these ideas.

Responding to the Risks of Abrupt Climate Change

The potential for presumably low probability, high impact threshold changes in the climate system (the so-called “fat tails”) plays an important role in debates over climate change ethics. This first example thus uses RDM to evaluate alternative policies for reduction of greenhouse gas (GHG) emissions in the face of a potential large-scale and economically costly collapse of the North Atlantic Meridional Overturning Circulation (MOC). The analysis, drawn from Hall *et al.* (2012), assumes that the scientific community has generated a best-estimate probability distribution for the most relevant uncertainties and seeks to help policy-makers consider the implications of their choices in light of the possibility that these estimates may be wrong.

The example compares four alternative greenhouse gas emission reduction paths of varying degrees of stringency. One path, Expected Utility Maximization (EUM), derives from an expected utility optimization using the best-estimate probability distribution for an MOC collapse and other important climate and economic parameters. The other paths include the less-stringent Business as Usual (BAU) and two more-precautionary options called Safety First

(SF), and Limited Degree of Confidence (LDC). Each path, except BAU, includes learning. The paths prescribe the emission reduction level through the year 2075, at which point perfect information becomes available about whether or not the MOC is vulnerable to collapse and the emission paths respond accordingly.³

The analysis compares the consequences of these four alternative policies using a version of the Dynamic Integrated model of Climate and Economy (DICE) model (Nordhaus 2008), modified to include the possibility of a MOC collapse triggered if and when atmospheric CO₂ levels exceed an uncertain threshold (Keller *et al.* 2004). The analysis uses present value expected utility to compare the four options and calculates this utility contingent on the values of four DICE simulation model input parameters that capture key climate, technology, and economic uncertainties. These parameters are: 1) the climate sensitivity, describing the equilibrium increase in mean near-surface air temperatures associated with a doubling of atmospheric CO₂ concentrations; 2) the initial growth rate of carbon intensity, which is related to the cost of reducing GHGs; 3) the economic damages associated with MOC collapse, expressed as a proportion of global economic output; and 4) a binary parameter indicating whether or not the MOC will actually collapse if the critical CO₂ concentration threshold, which depends on the climate sensitivity (Stocker *et al.* 1997; Keller *et al.* 2004), is exceeded.

Using this simulation, the analysis calculates the present value expected utility for each emissions path in each of 2,662 plausible future states of the world. Each future is defined by a particular numeric value for each of the four uncertain parameters. The analysis samples a wide range of values for each uncertainty, to capture the full range of plausible values and ensure that the analysis can adequately stress test the proposed emissions paths.⁴

Figure 2 uses this database of model results to begin to compare the mitigation strategies, by showing the distribution of regret for each strategy over each of the 2,662 plausible futures.

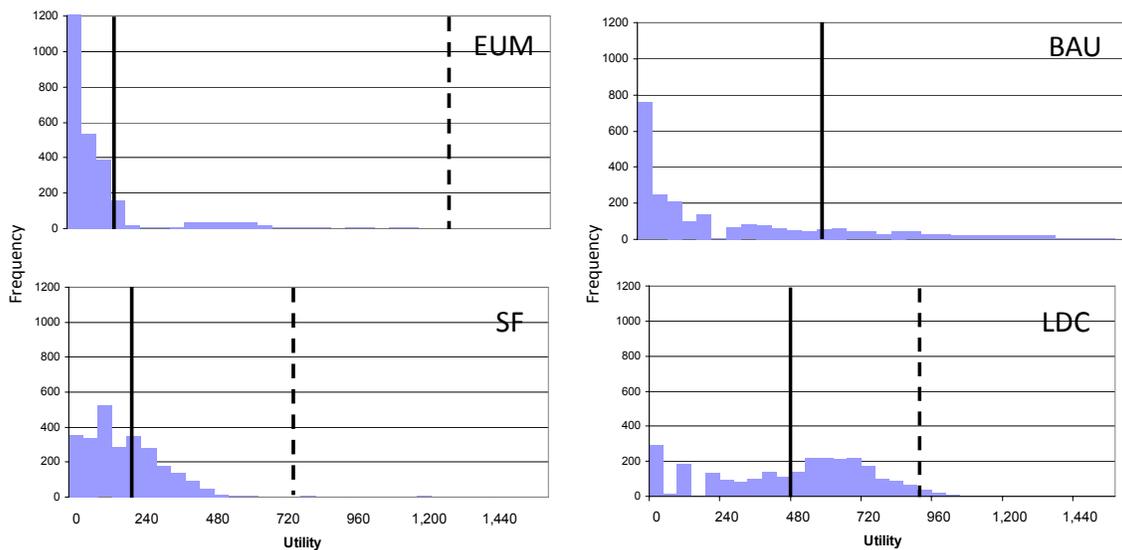
³ These paths derive from an earlier study that compared in the climate change context various “value-at-risk” decision rules from the finance literature (McInerney *et al.* 2012).

⁴ This sample of cases was generated by using the best-estimate probability distributions for the climate sensitivity, carbon intensity growth rate, and economic damage to identify eleven equally likely intervals for each of these uncertain parameters. A full factorial experimental design over these three, in combination with the two possible values for MOC Vulnerable, yields $113 \times 2 = 2662$ cases. The results of the analysis do not depend strongly on the use of the best-estimate distribution to construct the experimental design because the vulnerability analysis shown in Fig 3 does not depend on the probability weightings assigned to the cases and the tradeoff analysis in Fig 4 renormalizes the probabilistic weightings based on the results in Fig 3.

The regret, or distance from optimality, for any strategy in any future state of the world is the difference between that strategy's performance in that future and the performance of the best strategy in that future.⁵ Focusing on regret proves useful because it emphasizes the strategies' comparative, rather than absolute, performance. Some futures are unpleasant and some futures are pleasant no matter what decisions we make today. Regret directs attention to those futures where today's choices make a difference.

Figure 2

Regret of each of four strategies over the 2,662 cases examined with the DICE model. Solid lines show average and dashed lines 5% worst regret using best-estimate probability distributions. Tails of BAU distribution extend further than the range shown.



Notes: EUM = Expected Utility Maximization; BAU = Business as Usual; SF = Safety First; LDC = Limited Degree of Confidence

Not surprisingly, Figure 2 shows many futures where the optimum EUM strategy has low regret and is thus the best choice, or close to it. But in a small number of futures, EUM has large regret, that is, is far from the best choice. In contrast, and as one would expect from a more precautionary strategy, SF has more cases with non-zero regret than EUM, but has fewer

⁵ DICE is a Ramsey-style model of intertemporal choice in which the utility in a give time-period is the product of the population size and the per-capita consumption adjusted by the marginal elasticity of utility. The utilities in each time-period are then weighted by a discount factor that approximates the effects of a social rate of time preference and then summed over the entire considered time-horizon. The regret is the difference between such sums for the given strategy and the optimum strategy calculated in any given future.

cases with very large regret. The solid vertical lines in Fig 2 shows the expected regret for each strategy calculated using the best-estimate probability distribution.⁶ The dashed vertical lines mark for each strategy the five percent of futures with the highest regret.

Figure 2 suggests that SF sacrifices only a small increase in expected regret in order to achieve significant reductions in regret among the worst cases. It thus seems a plausible initial candidate for a robust strategy. The discussion so far, structuring the decision challenge and creating a large database of simulation runs, constitutes the first two steps in the decision support process shown in Figure 1. Step 3 now seeks to stress test SF by describing the future conditions in which the strategy might prove a poor choice.

To do so, the analysis applies statistical “scenario discovery” algorithms (Bryant *et al.* 2010) to the database of cases. These cluster-finding algorithms seek ways to parse the database to provide a concise description of the factors that best distinguish those futures where SF meets decision makers’ objectives from those futures where it doesn’t. These algorithms seek to balance between the competing goals of a providing a description that is both simple and accurate.⁷ The resulting clusters merit the name “scenario” because of the ways in which they can help focus decision makers’ attention on the uncertain future conditions most important to the challenges they face.

This scenario discovery process requires establishing a criterion for whether or not the SF strategy meets decision makers’ objectives. Some applications, such as the Colorado River example below, offer a natural choice for such criteria based, for instance, on legal requirements for sufficient flows from one Basin to another or physical infrastructure considerations like maintaining a minimum volume in a storage reservoir to ensure an uninterrupted water supply. Other examples include a budget constraint, an organization’s requirements for return on investment, or well-defined stakeholder preference. Here, however,

6 Note that the ordering of strategies ranked by expected regret is identical to the ordering of strategies ranked by expected value.

7 To measure accuracy, the algorithms use coverage and density. Coverage, also known as precision or positive predictive value, is the fraction of all cases in the database in which SF fails to meet objectives that are contained within the cluster(s). Density, also known as sensitivity or recall, is the fraction of cases within the cluster(s) in which SF fails to meet objectives. Simplicity, or interpretability, represents the ease with which the information in the clusters can be communicated to and understood by decision makers and is typically measured heuristically as the number of clusters and the number of uncertain parameters used to define the cluster(s). A small number of clusters and defining parameters is generally more interpretable.

we infer a threshold that might approximate decision makers' preferences by defining SF to have unacceptable performance in any future state of the world that yields a regret greater than 272, which characterizes the worst 15% (402 of 2662) SF cases in the database.

Figure 3 shows the results. Two of the four uncertain input parameters – the climate sensitivity and the possibility of MOC collapse—dominate in explaining the high regret cases for the SF path. These two uncertainties combine to form two scenarios, one with very high climate sensitivity (greater than 9°C), and the second with a combination of very low climate sensitivity (less than 1.2°C) and an MOC vulnerable to collapse. Figure 3 helps readers to judge the accuracy of these descriptions by denoting all the 2,662 cases considered with white circles and dark circles to denote those cases in which SF has unacceptably high regret. (The area of the circles is proportionally to the number of cases.) Ideally, all the cases within the two shaded regions, and none of the cases outside of those regions, would be dark circles.⁸

Figure 3
Two scenarios illuminating the vulnerabilities of SF strategy. Dark circles show number of cases where the SF strategy has high regret compared to the number of all cases (white circles).

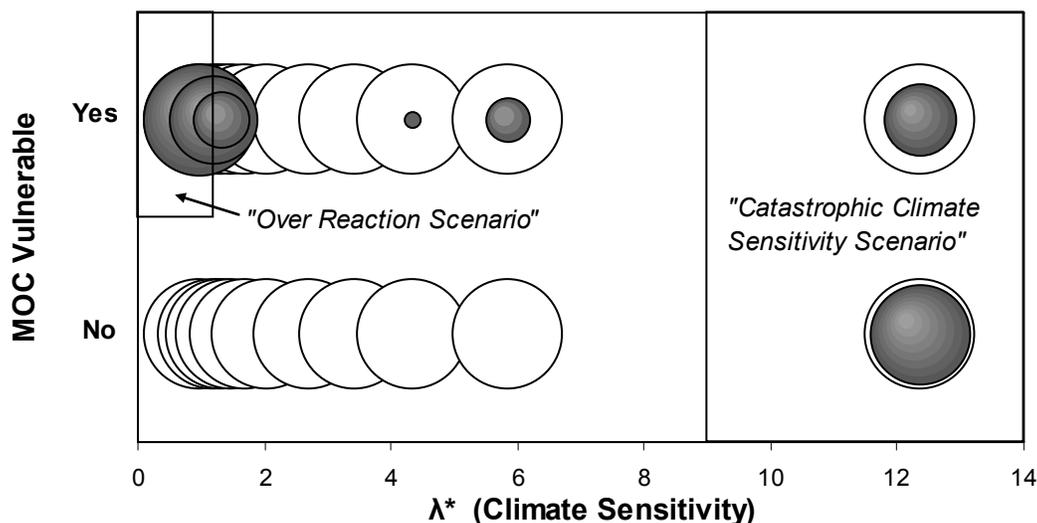


Figure 3, along with the underlying database of simulation runs, helps explain these

⁸ The quantitative measures suggest these scenarios offer a relatively accurate description. The scenario in the upper left has coverage/density = 49% /82%, the one on the right has coverage/density = 37% /62%, and combined the two have coverage/density = 86% /72%.

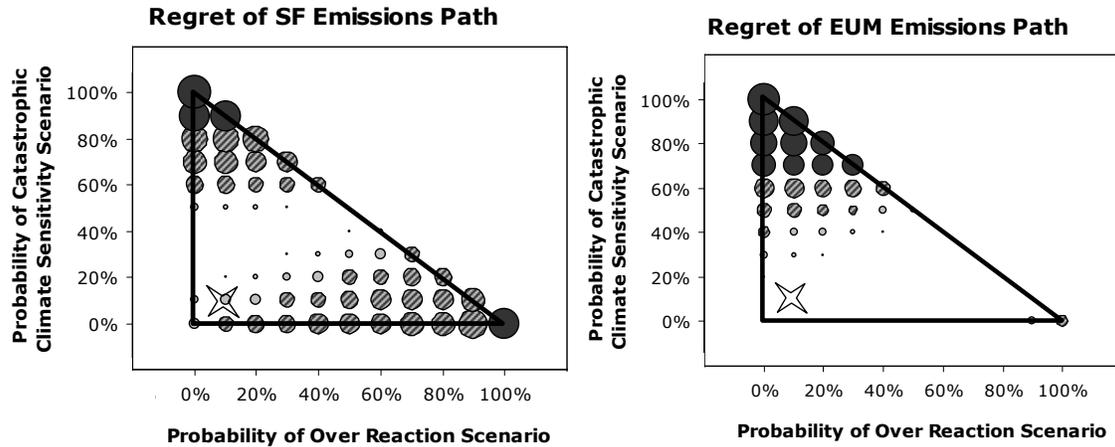
patterns of consequences for the SF strategy. In the first scenario, which we call *Catastrophic Climate Sensitivity*, SF fails because its near-term emissions reductions are too low compared to those of the most-precautionary LDC, which of the four paths considered does the best job of limiting climate impacts in these high sensitivity cases. In the second scenario, SF fails because it has high regret compared to BAU. In these cases, SF fails because the strategy is designed to begin with moderate near-term emissions reductions and then to increase its reduction rate if and when it learns that the MOC is vulnerable. However, in those futures with very low climate sensitivity the MOC is highly unlikely to collapse whether or not it is vulnerable because, as noted above, the critical threshold rises with declining climate sensitivity. In such circumstances SF unnecessarily increases its rate of emissions reductions and thus suffers high regret compared to BAU. We label this cluster of cases the *Over Reaction* scenario.

In brief, SF fails in the *Catastrophic Climate Sensitivity* scenario because it reduces emissions too little in the near-term and in the *Over Reaction* scenario because it reduces emissions too much in the long-term.

Figure 4 now uses this information to describe the tradeoffs involved with choosing the SF strategy. The figure shows the expected regret for the SF and EUM strategies as a function of the probability ascribed to the *Over Reaction* and *Catastrophic Climate Sensitivity* scenarios. The left hand panel shows that SF has high expected regret when the probability *Catastrophic Climate Sensitivity* is greater than about 60% and while the right-hand panel shows EUM has high regret when this probability is greater than about 40%. But SF's expected regret is strongly dependent on the probability of the *Over Reaction* scenario while EUM's is not. Thus relative to EUM, SF increases the range of probability for which it can successfully address the catastrophic scenario, but at the cost of performing more poorly when the probability of the over reaction case is anything but negligible.

Figure 4

Expected regret of SF and EUM strategies as a function of the probability ascribed to the “Over Reaction” and “Catastrophic Climate Sensitivity” scenarios.



Note: The star at (9%, 9%) indicates the best-estimate probability. The size and shading of the dots indicates the level of expected regret.

The expert best-estimate suggest that the probability of these two scenarios is approximately 9%, as shown by the star in Figure 4. The figure thus suggests that if decision makers are confident in the experts’ probability estimate they should choose EUM. If they worry that the probability of catastrophic climate sensitivity is significantly higher than the expert estimate, then SF may be a reasonable choice, but only if decision-makers believe the probability of the *Over Reaction* scenario is sufficiently low.

In addition, the patterns of vulnerability shown in Figure 3 suggest simple modifications that might significantly improve the performance of the SF strategy. If it learned by observing both the climate sensitivity and MOC sensitivity (as opposed to observing only the MOC sensitivity), the SF strategy might eliminate its high regret in the *Over Reaction* scenario without degrading its performance in the *Catastrophic Climate Sensitivity* scenario and thus become a more desirable choice relatively to EUM. This hypothesis would need to be tested, however, by repeating the steps in Figure 1 because changing the learning algorithm could also potentially improve the performance of the EUM and LDC strategies and introduce new vulnerabilities for SF.

Figure 4 thus provides an example of how considering a set of probability distributions,

rather a single distribution, can facilitate deliberations over an appropriate response to a deeply uncertain threat. Nonetheless, this example clearly represents a stylized sketch of an immensely complex decision challenge. The example lacks, for instance, any diversity of values or richness in its consideration of intergenerational equity. Rather than discuss these many limitations, it seems useful here to turn to a contrasting example, a recent effort organized by the U.S. Bureau of Reclamation, in which the same analytic approach was used to support a large study of the future of the Colorado River and facilitate deliberations among key stakeholders.

Colorado River Management

The Colorado River supplies water to forty million Americans in seven states,⁹ helps irrigate five million acres of U.S. agricultural lands, supports ecosystems, provides hydropower and recreational opportunities, and is shared with Mexico. The future reliability of these supplies, and the institutions that currently distribute them, face significant challenges. The 1922 Colorado River Compact initially allocated rights to the river's waters, with successive agreements and court decisions ultimately yielding a complete "Law of the River." Since then the number of users has grown by roughly an order of magnitude and their distribution shifted in important ways. In 1922, only California had any significant population in the regions served by the Colorado River, and for decades the other six states did not use their full annual allocation. Today many millions of people live in the other six states and increasingly seek to use their apportioned share. Recent paleoclimate research has made clear that the river flows used to allocate the river in 1922 were among the highest in 1,000 years, so it is reasonable to expect that in the future the river will rarely have as much flow as has been allocated. In addition, anthropogenic climate change is also expected to reduce future river flows.

The U.S. Department of Reclamation, the federal agency that manages the major infrastructure that stores and distributes Colorado River water to its users, recently conducted a three-year study—*the Colorado River Basin Water Supply and Demand Study*—in collaboration with the seven Basin states, a contractor team, and other stakeholders.¹⁰ This

⁹ The seven states are California, Colorado, Utah, Nevada, Arizona, New Mexico, and Wyoming.

¹⁰ The contractor team included CH2MHill, Black and Veatch, and for the last year of the study, RAND Corporation. Other key stakeholders include the Metropolitan Water District of Southern California, which holds

“Basin Study” sought to: 1) assess water supply and demand imbalances over the next fifty years and 2) develop and evaluate opportunities for resolving imbalances (Reclamation 2012). An RDM analysis, described in greater detail by Groves *et al.* (2013), played a key role in facilitating the deliberations underlying this study, which was released by Secretary of Interior Ken Salazar in December 2012.

As an input to the process, Reclamation and its partners assembled sets of projections to represent various uncertainties, including six alternative projections for potential future demand and nearly two thousand alternative projections of the future hydrology reflecting a wide range of potential climate and streamflow conditions. The project team considered such a large number of projections both to represent large uncertainty and because the spatial and temporal details of the projections could have important implications for the operations of the river. The projections derived from four different sources:¹¹ (1) recent, directly-measured hydrologic conditions, which assume that future climate patterns will be similar to those over the 20th century; (2) the thousand-year paleoclimate record, which documents past conditions significantly more stressing than those in the last century; (3) a blend of historical and paleoclimate reconstructions, which captures the longer-term climate patterns in the paleoclimate record while preserving the extremes of the past century; and (4) downscaled climate projections from climate models, which represent plausible future conditions and are translated into hydrologic conditions using a rainfall/runoff model..

The study team did not attempt to place any single probability distribution over these demand and climate projections. In part, the available scientific information would not support assigning any specific distributions with high confidence, in the sense of term as used by the IPCC uncertainty guidance (Mastrandrea *et al.* 2010). In addition, any such distribution would prove controversial among the participants to the discussion, who would clearly understand how any best-estimate might favor one alternative management strategy over another.

The project team also identified roughly two dozen performance measures for assessing the significance of any supply and demand imbalance, each of interest to one or more parties to

senior rights to Colorado River supplies, the City of Las Vegas, which holds junior rights, and representatives from environmental organizations, Native American tribes, and other organizations.

¹¹ The Basin Study refers to these four sources as four supply scenarios.

the discussions. These measures include those related to water deliveries at various points along the river, electric power generation, recreation, ecosystem health, water quality, and flood control.

Reclamation began soliciting suggestions from members of the project team and other interested parties throughout the region on ways to address the challenge raised by increasing demand and falling supply. The agency collected over a hundred proposals, ranging from conservation efforts, to ocean desalination plants to reduce California's dependence on the Colorado, to weather modification, to towing icebergs from Alaska to Southern California. The study team used Reclamation's simulation model of the Colorado River, called CRSS (for Colorado River Simulation System), to evaluate the performance of the current river management plans, and various proposed improvements, using these performance measures in many alternative futures each defined by a combination of the demand and climate projections.

At this point, Reclamation turned to RDM to help organize this vast body of information and the deliberations the information aimed to inform.¹² The study team generated a database with nearly 24,000 CRSS simulation runs. Working with this data in a series of workshops, Reclamation and its partners identified threshold performance levels that would represent satisfactory performance for each of the two-dozen measures. A scenario discovery analysis on the database of simulation model results helped to concisely define the conditions where the current river management plan would fail to meet one or more of these objectives. The resulting scenarios¹³ for two key measures—associated with Upper Basin and Lower Basin reliability¹⁴—each depended primarily on two factors: the average annual streamflow over the fifty-year time horizon and the lowest streamflow over any 8-year drought period. Irrespective of assumptions about future demand and other uncertainties, if the level for these two factors dropped below specific thresholds, one or more of the Basin's objectives would likely not to be

¹² After successful completion of a pilot study demonstrating the RDM approach, RAND was invited to join main study team in the last year of the Basin study in order to help organize and employ the large amount of information that had been generated in the first two years.

¹³ In the Basin Study these are called “vulnerable conditions”. In this paper we refer to them as “scenarios” as defined earlier.

¹⁴ The main Upper Basin reliability metric is a “Lee Ferry Deficit,” defined as the amount of curtailment of water to Upper Basin users that would be required to ensure that 75 million-acre-feet of supply flows from the Upper Basin to the Lower Basin over any consecutive 10-year period. The main Lower Basin reliability metric is the volume or pool elevation of Lake Mead, which determines how much water can be delivered to the Lower Basin users.

met under the current river management plan. For the Lower Basin reliability metric, these thresholds are close to the historical flow levels. Without the need to agree on any specific probability distribution the project team could acknowledge that this vulnerable scenario is sufficiently likely to warrant consideration of alternative management plans. The scenario illuminating vulnerabilities for the Upper Basin however, lie well below historic flow levels. The Lees Ferry flow would need to decrease significantly to threaten Upper Basin water use.

The analysis next grouped the hundreds of proposed options into four different portfolios, each reflecting an alternative view of how to best manage the river over the coming decades. One portfolio included only those options that were deemed to be highly feasible and reliable. A second included only those options with neutral or positive environmental attributes. A third included only the options common to the first two. A fourth portfolio included all options in the first and second portfolios. Each portfolio was configured as an adaptive decision strategy, that is, one designed to adjust over time in response to new information. Each portfolio ranked included options by cost-effectiveness, as measured by the expected yield divided by net present costs of the option, and specified signposts, such as specific threshold decrease in stream flow that, if observed in the future, would trigger the implementation of the portfolio's options.

The analysis used CRSS to examine the performance of each portfolio in each of the many thousands of future states of the world and then identified the scenarios that illuminate the vulnerabilities of those portfolios. The portfolios all significantly reduced the vulnerabilities of the system. Not surprisingly, the portfolios also show important tradeoffs between cost and the ability to cope with significantly reduced future stream flow.

The analysis then examined how decision makers might prioritize the actions within the portfolios, by tracing the portfolios as they evolved along the many thousands of future paths. The analysis identified some actions, for instance urban and agricultural conservation measures that are adopted in the near-term in almost every future. This information suggests that these options might represent important near-term priorities. The analysis identified other actions, for instance ocean water desalinization, that are only adopted in the long-term and only in some futures, but not in others. This information suggests that final decision on such actions might be deferred until later. The analysis helped the parties to begin discussions today regarding the

specific signposts that, if observed in the future, would indicate the need to take such actions. Finally, the analysis identified some actions, such as municipal wastewater reuse, that are adopted in the near-term in some futures, but not in other futures. The choice depends in part on the probability one ascribes to the scenarios that illuminate the vulnerabilities of the river's current management strategy. These actions may require additional deliberation and analysis.

This RDM analysis appeared to play an important role in helping the parties to this potentially contentious debate reach broad areas of consensus on statements about the future, about high priority near-term steps, and about some of the tradeoffs involved with potential subsequent steps.

In the context of this essay, it is worth noting some practical constraints that arose in implementing this analysis. The CRSS simulation model is relatively computationally intensive. Both the RAND Corporation and the Bureau of Reclamation have some cluster computing capabilities. Nonetheless, the amount of time needed to run tens of thousands of CRSS cases on RAND's and Reclamation's (relatively small) cluster-computing machines was sufficiently long compared to the time between stakeholder workshops that run time became an important consideration in the design of the cases that could be considered.

IV. IS IT ETHICAL TO USE A SINGLE PDF?

Ethically defensible policy judgments require examining arguments from multiple points of view. Such judgments also require close attention to their actual effect on social realizations, that is, on the lives people actually lead. Thus ethical judgment demands a reasonable level of due diligence towards assembling available evidence and using it to judge the consequence of proposed actions.

Uncertainty can complicate this exercise of due diligence by proliferating the potential consequences that flow from an action and making differences in rankings of values or preferred outcomes more salient. Given this complexity, some type of formal process is often required to help people sort through the many combinations of values, attitudes toward risk, and consequences of different policy choices. Traditional probabilistic risk analysis provides a framework that can significantly improve the outcomes that would result if people made their decisions unaided. But such analyses require as a starting point consensus on the likelihood of

alternative futures (a probability density function) and consistency across values and expectations. This demand for a consensus pdf can inhibit deliberation among individuals with differing expectations and values. A single pdf can also promote over-confidence in individual decision makers.

If no other options were available, using a single pdf might be justified. But there now exist means to systematically compare proposed actions in the face of uncertain consequences using a set of plausible, alternative pdfs. This allows the analytics to engage multiple points of view while at the same time enabling systematic due diligence to the extent allowed by the available quantitative information.

These new methods are not costless. At a minimum, they require more computational resources (run time and memory) than would otherwise be the case. More broadly, they generate information that needs to be individually customized for each decision context—reporting, for instance, stream flow thresholds at Lee Ferry where the current Colorado River management plan fails to meet the Upper Basin state’s supply objectives, rather than a single best-estimate probabilistic stream flow forecast intended to be broadly relevant for many decision makers independent of context. These methods can also intermingle the producers and users of uncertainty estimates into a single process rather than retain a clear institutional separation between the two that in some circumstances may prove useful.

When the scientific evidence proves overwhelming, the range of views relatively narrow, or the consequences of choosing unwisely sufficiently small, single best-estimate distributions may prove simpler and appropriate. But in situations where policy choices may have significant consequences for many people, the science is changing and uncertain, and where there exist significant differences in values and expectations, ethical reasoning may be served better by considering a set of multiple pdfs.

References

- Bankes, S. C. (1993). "Exploratory Modeling for Policy Analysis." Operations Research **41**(3): 435-449.
- Brown, C. (2010). "The End of Reliability." Journal of Water Resources Planning and Management: 143-145.
- Bryant, B. P. and R. J. Lempert (2010). "Thinking inside the box: A Participatory, computer-assisted approach to scenario discovery." Technological Forecasting and Social Change **77**: 34-49.
- Bureau of Reclamation (2012). Colorado River Basin Water Supply and Demand Study: Study Report. United States Bureau of Reclamation: 89.
- Carter, T. R., R. N. Jones, S. B. X. Lu, C. Conde, L. O. Mearns, B. C. O'Neill, M. D. A. Rounsevell and M. B. Zurek (2007). New Assessment Methods and the Characterisation of Future Conditions. Climate Change 2007: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. M. L. Parry, O. F. Canziani, J. P. Palutikof, P. J. v. d. Linden and C. E. Hanson. Cambridge, UK,, Cambridge University Press, 1: 33-171.
- Dessai, S. and M. Hulme (2007). "Assessing the Robustness of Adaptation Decisions to Climate Change Uncertainties: A Case Study on Water Resources Management in the East of England." Global Environmental Change **17**(1): 59-72.
- Groves, D. G., J. Fischbach, E. Bloom, D. Knopman and R. Keefe (2013). Adapting to a Changing Colorado River: Making Future Water Deliveries More Reliable Through Robust Management Strategies. Santa Monica, CA, RAND Corporation.
- Groves, D. G. and R. J. Lempert (2007). "A New Analytic Method for Finding Policy-Relevant Scenarios." Global Environmental Change **17**: 73-85.
- Hall, J. M., R. Lempert, K. Keller, A. Hackbarth, C. Mijere and D. McInerney (2012). "Robust Climate Policies under uncertainty: A comparison of Info-Gap and RDM methods." Risk Analysis.
- Herrick, C. and D. Sarewitz (2000). "Ex Post Evaluation: A More Effective Role for Scientific Assessment in Environmental Policy." Sci Technol Hum Values **25**: 309.
- Kahan, D. M. and D. Braman (2006). "Cultural Cognition and Public Policy." Yale Law and Policy Review **24**: 147-170.
- 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 Modeling and Software **42**: 55-71.

- Keller, K., B. M. Bolker and D. F. Bradford (2004). "Uncertain Climate Thresholds and Optimal Economic Growth." Journal of Environmental Economics and Management **48**: 723-741.
- Lempert, R. (2013). "Scenarios that illuminate vulnerabilities and robust responses." Climatic Change **117**: 627-646.
- Lempert, R. and D. G. Groves (2010). "Identifying and Evaluating Robust Adaptive Policy Responses to Climate Change for Water Management Agencies in the American West." Technological Forecasting and Social Change **77**: 960-974.
- Lempert, R., N. Nakicenovic, D. Sarewitz and M. Schlesinger (2004). "Characterizing climate-change uncertainties for decision-makers - An editorial essay." Climatic Change **65**(1-2): 1-9.
- Lempert, R. J. and M. Collins (2007). "Managing the Risk of Uncertain Threshold Responses: Comparison of Robust, Optimum, and Precautionary Approaches." Risk Analysis **27**(4): 1009-1026.
- 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. Santa Monica, CA, RAND Corporation.
- Mastrandrea, M. D., C. B. Field, T. F. Stocker, O. Edenhofer, K. L. Ebi, D. J. Frame, H. Held, E. Kriegler, K. J. Mach, P. R. Matschoss, G.-K. Plattner, G. W. Yohe and F. W. Zwiers (2010). Guidance Note for Lead Authors of the IPCC Fifth Assessment Report on Consistent Treatment of Uncertainties, Intergovernmental Panel on Climate Change.
- McInerney, D., R. Lempert and K. Keller (2012). "What are Robust Strategies in the Face of Uncertain Climate Threshold Responses?" Climate Change.
- National Research Council (2009). Informing Decisions in a Changing Climate. T. N. A. Press. Washington, DC, Panel on Strategies and Methods for Climate-Related Decision Support, Committee on the Human Dimensions of Climate Change, Division of Behavioral and Social Sciences and Education.
- Nordhaus, W. D. (2008). A Question of Balance. New Haven & London, Yale University Press.
- Parson, E. A. (2008). "Useful global-change scenarios: current issues and challenges." Environ. Res. Lett. **3**: 1-5.
- Parson, E. A., V. Burkett, K. Fischer-Vanden, D. Keith, L. O. Mearns, H. Pitcher, C. Rosenweig and M. Webster (2007). Global-Change Scenarios: Their Development and Use, Synthesis and Assessment Product 2.1b, US Climate Change Science Program.
- Ranger, N., A. Millner, S. Dietz, S. Fankhauser, A. Lopez and G. Ruta (2010). Adaptation in the UK: a decision making process. Grantham/CCEP Policy Brief.
- Rawls, J. (1971). A Theory of Justice, Harvard University Press
- Rayner, S. (2000). Prediction and Other Approaches to Climate Change Policy. Prediction: Science, Decision Making, and the Future of Nature. D. Sarewitz, J. Pielke, R and R. Byerly. Washington DC, Island Press: 269-296.

- Rosenhead, J. (2001). *Robustness Analysis: Keeping Your Options Open*. Rational Analysis for a Problematic World Revisited. Johnathan Rosenhead and J. Mingers, John Wiley and Sons.
- Schoemaker, P. J. H. (1993). "Multiple Scenario Development: Its conceptual and Behavioral Foundation." Strategic Management Journal **14**(3): 193-213.
- Sen, A. (2009). The Idea of Justice. Cambridge, Massachusetts, Belknap Press.
- Stocker, T. F. and A. Schmittner (1997). "Influence of CO2 Emission Rates on the Stability of the Thermohaline Circulation." Nature **388**: 862-865.
- Vantklooster, S. and M. Vanasselt (2006). "Practising the scenario-axes technique." Futures **38**(1): 15-30.
- Walker, W., V. Marchau and D. Swanson (2010). "Addressing deep uncertainty using adaptive policies." Technology Forecasting and Social Change **77**: 917-923.
- Weaver, C. P., Robert J. Lempert, Casey Brown, John A. Hall, David Revell and Daniel Sarewitz (2013). "Improving the contribution of climate model information to decision making: the value and demands of robust decision frameworks." WIREs Climate Change **4**(1): 39-60.