Chapter 2
Robust Decision Making (RDM)

R. J. Lempert

Abstract

- The quest for predictions—and a reliance on the analytical methods that require them—can prove counter-productive and sometimes dangerous in a fast-changing world.
- Robust Decision Making (RDM) is a set of concepts, processes, and enabling tools that use computation, not to make better predictions, but to yield better decisions under conditions of deep uncertainty.
- RDM combines Decision Analysis, Assumption-Based Planning, scenarios, and Exploratory Modeling to stress test strategies over myriad plausible paths into the future, and then to identify policy-relevant scenarios and robust adaptive strategies.
- RDM embeds analytic tools in a decision support process called "deliberation with analysis" that promotes learning and consensus-building among stakeholders.
- The chapter demonstrates an RDM approach to identifying a robust mix of policy instruments—carbon taxes and technology subsidies—for reducing greenhouse gas emissions. The example also highlights RDM's approach to adaptive strategies, agent-based modeling, and complex systems.
- Frontiers for RDM development include expanding the capabilities of multi-objective RDM (MORDM), more extensive evaluation of the impact and effectiveness of RDM-based decision support systems, and using RDM's ability to reflect multiple world views and ethical frameworks to help improve the way organizations use and communicate analytics for wicked problems.

R. J. Lempert (✉)
Rand Corporation, Santa Monica, CA, USA
e-mail: lempert@rand.org

2.1   Introduction

Toward the end of the Cold War, in the early 1980s, the RAND Corporation invested much effort toward using computer combat simulation models to inform national security decisions regarding weapons procurement and strategy. Designed to predict the course of future military conflicts, these models were obviously imperfect. Everyone recognized the aphorisms about the military's propensity to plan for the last as opposed to the next war. But growing computational capabilities offered new opportunities to simulate complex military campaigns with unprecedented levels of fidelity and realism. The efforts to harness these capabilities provided a focal point for a growing debate on the value of using newly powerful but clearly imperfect analytic tools to inform real and consequential decisions.

The 1982 Falklands/Malvinas war provided a crystalizing moment. Argentine forces landed on the islands in early April and a British task force set sail a few days later for the two-month voyage to the South Atlantic. At that point, all knew the order of battle for the coming conflict. At RAND, those who distrusted the combat simulations used the opportunity to confront the model advocates, challenging them to predict how the conflict would unfold. Two months would provide ample time to conduct the analysis. But the advocates demurred. The models, they claimed, despite all their fidelity and realism, existed to provide "insight" not predictions. The skeptics wondered about the utility of insights gleaned from models whose predictions could not be trusted.

Decisionmakers often seek predictions about the future to inform policy choices. This seems natural because prediction is the bedrock of science, enabling researchers to test their hypotheses and demonstrate understanding of complicated systems. Decisionmakers also find predictions attractive because, when good ones are available, they unquestionably provide valuable input toward better choices. Society's vast enterprise in the physical, biological, and social sciences continually improves its predictive capabilities in order to test and refine scientific hypotheses. Recent analytic innovations, such as big data and machine learning, enhance decision-relevant predictive capabilities. New processes, such as crowd sourcing, prediction markets, and super-forecasting (Tetlock and Gardner 2016), provide new and more reliable means to aggregate human judgments into probabilistic forecasts.

But, as argued throughout this book, the quest for predictions—and a reliance upon analytic methods that require them—can prove counter-productive and sometimes dangerous in a fast-changing, complex world. Prediction-focused analysis risks over-confidence in organizations' decisionmaking and in their internal and external communications (Sarewitz and Pielke 2000). Prediction-focused policy debates can also fall victim to the strategic uses of uncertainty. Opponents may attack a proposed policy by casting doubt on the predictions used to justify it, rather than engaging with the merits of the policy itself, knowing that the policy may be more sound than the predictions (Herrick and Sarewitz 2000; Rayner 2000; Lempert and Popper 2005; Weaver et al. 2013).

A reliance on prediction can also skew the framing of a decision challenge. President Eisenhower (reportedly) advised "if a problem cannot be solved, enlarge it." But science often reduces uncertainty by narrowing its focus, prioritizing questions that can be resolved by prediction, not necessarily on the most decision-relevant inquiries. So-called wicked problems (Rittel and Webber 1973) present this contrast most starkly. In addition to their irreducible uncertainty and nonlinear dynamics, wicked problems are not well-bounded, are framed differently by various stakeholders, and are not well-understood until after formulation of a solution. Using predictions to adjudicate such problems skews attention toward the proverbial lamp post, not the true location of the keys to a policy solution.

This chapter describes Robust Decision Making (RDM), a set of concepts, processes, and enabling tools designed to re-imagine the role of quantitative models and data in informing decisions when prediction is perilous (Lempert et al. 2003, 2006). Rather than regarding models as tools for prediction and the subsequent prescriptive ranking of decision options, models and data become vehicles for systematically exploring the consequences of assumptions; expanding the range of futures considered; crafting promising new responses to dangers and opportunities; and sifting through a multiplicity of scenarios, options, objectives, and problem framings to identify the most important tradeoffs confronting decisionmakers. That is, rather than making better predictions, quantitative models and data can be used to inform better decisions (Popper et al. 2005).

RDM rests on a simple concept (Lempert et al. 2013a). Rather than using computer models and data as predictive tools, the approach runs models myriad times to stress test proposed decisions against a wide range of plausible futures. Analysts then use visualization and statistical analysis of the resulting large database of model runs to help decisionmakers identify the key features that distinguish those futures in which their plans meet and miss their goals. This information helps decisionmakers identify, frame, evaluate, modify, and choose robust strategies—ones that meet multiple objectives over many scenarios.

RDM provides decision support under conditions of deep uncertainty. As described in Sect. 2.2, RDM builds on strong foundations of relevant theory and practice to provide an operational and newly capable synthesis through the use of today's burgeoning information technology. The most commonly used analytic methods for predictive decision and risk analysis have their roots in the 1950s and 1960s, when relative computational poverty made a virtue of analytics recommending a single best answer based on a single best-estimate prediction. Today's ubiquitous and inexpensive computation enables analytics better suited to more complex problems, many of them "wicked" and thus poorly served by the approximation that there exists such an optimal solution.

The British won the Falklands/Malvinas war, but lost a vessel to an unexpectedly effective anti-ship missile attack and sank an Argentine ship under disputed conditions. From an RDM perspective, the model skeptics at RAND asked the wrong question. Rather than judge the combat simulations by how well they could generate an accurate probability density function, the question should have been—can these quantitative models, and the processes for using them, help the British make better

decisions regarding what forces were appropriate to bring and how they ought to be employed. While useful, even models that produced a perfect probability density function for precisely selected outcomes would not prove sufficient to answer such questions. Nor are they necessary.[1]

2.2 RDM Foundations

How can quantitative, evidence-based analysis best inform our choices in today's fast-paced and turbulent times? RDM and the other DMDU methods in this book aim to answer this question. In particular, RDM does so by providing a new synthesis of four key concepts: Decision Analysis, Assumption-Based Planning, scenarios, and Exploratory Modeling.

Decision Analysis

A large body of empirical research makes clear that people, acting as individuals or in groups, often make better decisions when using well-structured decision aids. The discipline of Decision Analysis (DA) comprises the theory, methodology, and practice that inform the design and use of such aids. RDM represents one type of quantitative DA method, drawing, for instance, on the field's decision structuring frameworks, a consequentialist orientation in which alternative actions are evaluated in each of several alternative future states of the world, a focus on identifying trade-offs among alternative decision options, and tools for comparing decision outcomes addressing multiple objectives.

As one key contribution, DA and related fields help answer the crucial question: What constitutes a good decision? No universal criterion exists. Seemingly reasonable decisions can turn out badly, but seemingly unreasonable decisions can turn out well. Good decisions tend to emerge from a process in which people are explicit about their goals, use the best available evidence to understand the potential consequences of their actions, carefully consider the tradeoffs, contemplate the decision

---

[1]As suggested by the Falklands/Malvinas story, RDM had its origins in debates within the RAND Corporation on how best to use models, typically of social interaction or military combat, that could not be validated in the same way as models of less complex decision spaces, or for which probability distribution functions could not be applied with confidence (Hodges 1991; Bankes 1993). The first published examples of the approach that became RDM include: Lempert et al. (1996); Rydell et al. (1997); Brooks et al. (1999). The Third Assessment report of the Intergovernmental Panel on Climate Change (IPCC 2001, Sect. 10.1.4.4) describes such work under the label "computational, multi-scenario approaches." In the 1990s, Evolving Logic (http://www.evolvinglogic.com) developed the first software dedicated to a DMDU methodology—the Computer-Assisted Reasoning® system (CARs™) (Bankes et al. 2001), to support a methodology called Robust Adaptive Planning (RAP). Parallel work at RAND adopted the term Robust Decision Making (RDM) (Lempert et al. 2003). See Light (2005) for a history of approaches to robust and adaptive planning at RAND.

from a wide range of views and vantages, and follow agreed-upon rules and norms that enhance the legitimacy of the process for all those concerned[2] (Jones et al. 2014).

While broad in principle, in practice the DA community often seeks to inform good decisions using an expected utility framework for characterizing uncertainty and comparing decision options (Morgan and Henrion 1990). This expected utility framework characterizes uncertainty with a single joint probability distribution over future states of the world. Such distributions often reflect Bayesian (i.e., subjective) rather than frequentist probability judgments. The framework then uses optimality criteria to rank alternative options. RDM, in contrast, regards uncertainties as deep and thus either eschews probabilities or uses sets of alternative distributions drawing on the concepts of imprecise probabilities (Walley 1991). RDM uses decision criteria based on robustness rather than optimality.

DA based on expected utility can be usefully termed "agree-on-assumptions" (Kalra et al. 2014) or "predict-then-act" (Lempert et al. 2004) approaches, because they begin by seeking agreement regarding the likelihood of future states of the world and then use this agreement to provide a prescriptive ranking of policy alternatives. In contrast, RDM and many of the approaches described in this book follow an "agree-on-decisions" approach, which inverts these steps.[3] They begin with one or more strategies under consideration, use models and data to stress test the strategies over a wide range of plausible paths into the future, and then use the information in the resulting database of runs to characterize vulnerabilities of the proposed strategies and to identify and evaluate potential response to those vulnerabilities. Such approaches seek to expand the range of futures and alternatives considered and, rather than provide a prescriptive ranking of options, often seek to illuminate trade-offs among not-unreasonable choices. As summarized by Helgeson (2018), agree-on-assumptions approaches generally focus on identifying a normative best choice among a fixed menu of decision alternatives, while agree-on-decision approaches focus on supporting the search for an appropriate framing of complex decisions.

Assumption-Based Planning

As part of an "agree-on-decisions" approach, RDM draws on the related concepts of stress testing and red teaming. The former, which derives from engineering and finance (Borio et al. 2014), subjects a system to deliberately intense testing to determine its breaking points. The latter, often associated with best practice in US and other militaries' planning (Zenko 2015), involves forming an independent group to identify means to defeat an organization's plans. Both stress testing and red teaming aim to reduce the deleterious effects of over-confidence in existing systems and plans by improving understanding of how and why they may fail (Lempert 2007).

---

[2]These attributes follow from a broadly consequentialist, as opposed to rule-based (deontological) view of decisionmaking (March 1994).

[3]The DMDU literature often uses different names to describe this inverted analytic process, including "backwards analysis" (Lempert et al. 2013a), "bottom up" (Ghile et al. 2014), "context first" (Ranger et al. 2010), and "assess risk of policy" (Lempert et al. 2004; Carter et al. 2007; Dessai and Hulme 2007).

In particular, RDM draws on a specific form of this concept, a methodology called Assumption-Based Planning (ABP) (Dewar et al. 1993, 2002). Originally developed to help the US Army adjust its plans in the aftermath of the Cold War, ABP begins with a written version of an organization's plans and then identifies load-bearing assumptions—that is, the explicit and implicit assumptions made while developing that plan that, if wrong, would cause the plan to fail. Planners can then judge which of these load-bearing assumptions are also vulnerable—that is, could potentially fail during the lifetime of the plan.

ABP links the identification of vulnerable, load-bearing assumptions to a simple framework for adaptive planning that is often used in RDM analyses. Essential components of an adaptive strategy include a planned sequence of actions, the potential to gain new information that might signal a need to change this planned sequence, and actions to be taken in response to this new information, i.e., contingent actions (Walker et al. 2001). After identifying the vulnerable, load-bearing assumptions, ABP considers *shaping actions* (those designed to make the assumptions less likely to fail), *hedging actions* (those that can be taken if assumptions begin to fail), and *signposts* (trends and events to monitor in order to detect whether any assumptions are failing).

Scenarios

RDM draws from scenario analysis the concept of a multiplicity of plausible futures as a means to characterize and communicate deep uncertainty (Lempert et al. 2003). Scenarios represent internally consistent descriptions of future events that often come in sets of two or more. Most simply, scenarios are projected futures that claim less confidence than probabilistic forecasts. More generally, a set of scenarios often seeks to represent different ways of looking at the world without an explicit ranking of relative likelihood (Wack 1985).

Scenarios are often developed and used in deliberative processes with stakeholders. Deemphasizing probabilistic ranking—focusing on a sense of possibility, rather than probability—helps stakeholders expand the range of futures they consider, allowing them to contemplate their choices from a wider range of views and vantages, thus helping participants consider uncomfortable or unexpected futures (Schoemaker 1993; Gong et al. 2017). The sense of possibility rather than probability can also help scenarios to communicate a wide range of futures to audiences not necessarily eager to have their vantage expanded. By representing different visions of the future without privileging among them, scenarios can offer a comfortable entry into an analysis. Each person can find an initially resonant scenario before contemplating ones that they find more dissonant.

RDM draws from scenario analysis the concept of organizing information about the future into a small number of distinct cases that help people engage with, explore, and communicate deep uncertainty. In particular, the Intuitive Logics school of scenario analysis (Schwartz 1996) uses qualitative methods to craft a small number of scenarios, distinguished by a small number of key uncertain determinants that differentiate alternative decision-relevant paths into the future. As described below, RDM uses quantitative "Scenario Discovery" algorithms to pursue the same ends.

The resulting scenarios summarize the results of the ABP-style stress tests and can link to the development of adaptive strategies (Groves and Lempert 2007; Groves et al. 2014). Note that while the scenario literature traditionally distinguishes between probabilistic and non-probabilistic treatments, RDM often employs a third alternative—entertaining multiple views about the likelihood of the scenarios.

The scenario literature also describes a process for seeking robust strategies that include choosing a set of scenarios that include the most important uncertainties facing the users and then identifying strategies that perform well across all of them (van der Heijden 1996). This process provides an animating idea for RDM.

Exploratory Modeling

RDM integrates these concepts—DA, ABP, and scenarios—through Exploratory Modeling (EM). Bankes (1993) encapsulated the 1980s RAND debates on useful and predictive models by dividing computer simulations into two types: (1) consolidative models, which gather all known facts together into a single package that, once validated, can serve as a surrogate for the real world, and (2) exploratory models, which map a wide range of assumptions onto their consequences without privileging one set of assumptions over another. Exploratory models are useful when no single model can be validated because of missing data, inadequate or competing theories, or an irreducibly uncertain future.

Running a model many times is not profound. But as perhaps its key insight, EM notes that when used with an appropriate experimental design—that is, appropriate questions and a well-chosen set of cases designed to address those questions—the large database of results generated from non-predictive, exploratory models can prove surprisingly useful toward informing policy choices. Bankes describes several types of questions one may address with EM, including hypothesis generation, reasoning from special cases, and assessing properties of the entire ensemble (Weaver et al. 2013 and Box 2.1). RDM uses them all, but focuses in particular on robustness as a property of the entire ensemble. That is, identifying and evaluating robust strategies become key questions one can inform with EM.

EM provides RDM with a quantitative framework for stress testing and scenario analysis. While consolidative models most usefully support deductive reasoning, exploratory models serve best to support inductive reasoning—an iterative cycle of question and response. As described in the process section below (Sect. 2.3), RDM also aims to support a decision-analytic human-machine collaboration that draws upon what each partner does best.

RDM also exploits another EM advantage: the focus on the simple computational task of running models numerous times in the forward direction. This facilitates exploration of futures and strategies by reducing the requirements for analytic tractability on the models used in the analysis, relative to approaches that rely on optimization or Dynamic Programming. In addition, EM enables truly global sensitivity explorations, since it privileges no base case or single future as an anchor point.

Box 2.1: Key Elements of RDM

RDM meets its goals by proceeding in multiple iterations, as humans and computers alternatively test each other's conclusions about futures and strategies. Four key elements govern these interactions (Lempert et al. 2003):

- Consider a multiplicity of plausible futures. The ensemble of futures should be as diverse as possible to adequately stress test proposed policies. The ensemble can also facilitate group processes by including futures that correspond to different groups' worldviews.
- Seek robust, rather than optimal strategies. Robust strategies perform well, compared to the alternatives, over a wide range of plausible futures.
- Employ adaptive strategies to achieve robustness. Adaptive strategies are designed to evolve over time in response to new information. Generally, such strategies reflect decisionmaking rules and in practice are often organized around near-term actions, signposts to monitor, and contingency actions to take in response to those signposts.
- Use the computer to facilitate human deliberation over explorations, options, and tradeoffs, not as a device for recommending a particular ordering of strategies.

2.3 RDM Process

RDM explicitly follows a learning process called "deliberation with analysis" in which parties to a decision deliberate on their objectives and options; analysts generate decision-relevant information using system models; and the parties to the decision revisit their objectives, options, and problem framing influenced by this quantitative information (NRC 2009). Among learning processes, deliberation with analysis proves most appropriate for situations with diverse decisionmakers who face a changing decision environment and whose goals can evolve as they collaborate with others. Deliberation with analysis also supports continuous learning based on indicators and monitoring (NRC 2009, p. 74), a process important to the literature on adaptive policymaking (Swanson and Bhadwal 2009; Walker et al. 2010).

Step 1 As shown in Fig. 2.1, the RDM process starts with a decision framing exercise in which stakeholders define the key factors in the analysis: the decisionmakers' objectives and criteria; the alternative actions they can take to pursue those objectives; the uncertainties that may affect the connection between actions and consequences; and the relationships, often instantiated in computer simulation models, between actions, uncertainties, and objectives. This information is often organized

Fig. 2.1 Steps in an RDM analysis. *Source* Lempert et al. (2013a), with earlier versions in Lempert and Groves (2010); Lempert and Kalra (2011)
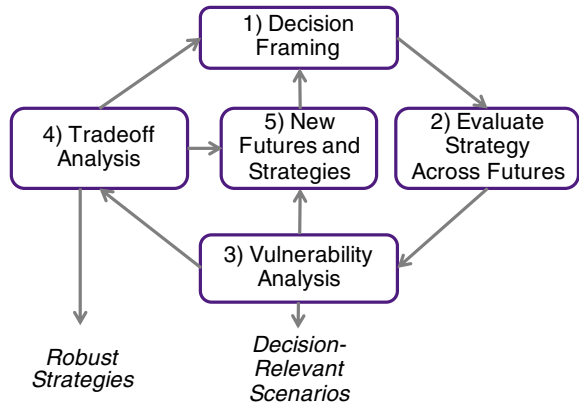
in a 2×2 matrix called "XLRM" (Lempert et al. 2003), for exogenous uncertainties (X), policy levers (L), relationships (R), and measures of performance (M).[4]

Step 2 As an "agree-on-decision" approach, RDM next uses simulation models to evaluate proposed strategies in each of many plausible paths into the future, which generate a large database of simulation model results. The proposed strategies can derive from a variety of sources. In some cases, an RDM analysis might start with one or more specific strategies drawn from the relevant public debate. For instance, an RDM analysis for a water agency might begin with that agency's proposed plan for meeting its supply requirements (Groves et al. 2014) or water quality requirements (Fischbach et al. 2015). In other cases, optimization routines for one or more expected futures or decision criteria might yield the initial proposed strategies (Hall et al. 2012). Additionally, an analysis might begin with a wide span of simple strategies covering the logical spectrum and then refine, select, and modify to yield a small group of more sophisticated alternatives (Popper et al. 2009). Often, such as in the example in Sect. 2.5, an application uses a combination of these approaches.

Step 3 Analysts and decisionmakers next use visualization and data analytics on these databases to explore for and characterize vulnerabilities. Commonly, RDM analyses use statistical Scenario Discovery (SD) algorithms (see below) to identify and display for users the key factors that best distinguish futures in which proposed strategies meet or miss their goals. These clusters of futures are usefully considered policy-relevant scenarios that illuminate the vulnerabilities of the proposed policies. Because these scenarios are clearly, reproducibly, and unambiguously linked to a policy stress test, they can avoid the problems of bias and arbitrariness that sometimes afflict more qualitative scenario exercises (Lempert 2013; Parker et al. 2015).

Step 4 Analysts and decisionmakers may use these scenarios to display and evaluate the tradeoffs among strategies. For instance, one can plot the performance of one or more strategies as a function of the likelihood of the policy-relevant scenarios

---

[4]Chapter 1 (see Fig. 1.2) uses the acronym XPROW for external forces, policies, relationships, outcomes, and weights.
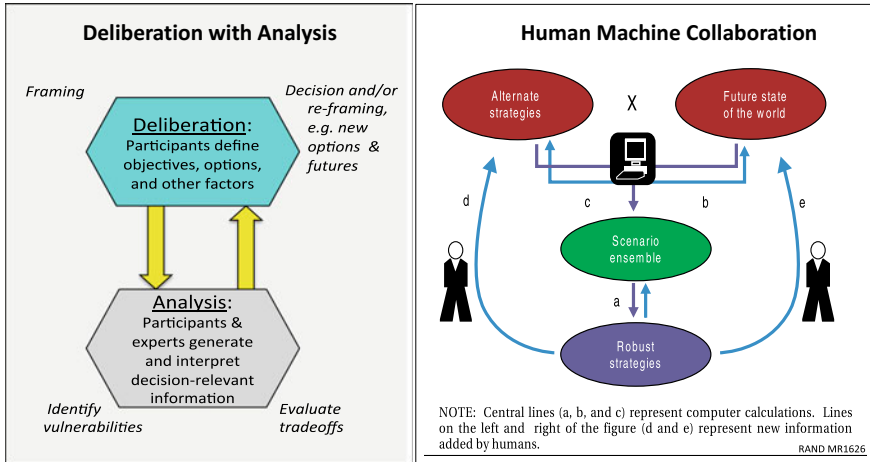
Fig. 2.2 Two views of RDM as a deliberative process. *Source* for right panel, Lempert et al. (2003)

(e.g., see Fig. 2.2) to suggest the judgments about the future implied by choosing one strategy over another. Other analyses plot multi-objective tradeoff curves—for instance comparing reliability and cost (Groves et al. 2012)—for each of the policy-relevant scenarios to help decisionmakers decide how to best balance among their competing objectives.

Step 5 Analysts and decisionmakers could then use the scenarios and tradeoff analyses to identify and evaluate potentially more robust strategies—ones that provide better tradeoffs than the existing alternatives. These new alternatives generally incorporate additional policy levers, often the components of adaptive decision strategies: short-term actions, signposts, and contingent actions to be taken if the pre-designated signpost signals are observed. In some analyses, such adaptive strategies are crafted using expert judgment (e.g., see Lempert et al. 1996, 2000, 2003; Popper et al. 2009; Lempert and Groves 2010). In other analyses, optimization algorithms may help suggest the best combination of near-term actions, contingent actions, and signposts for the new adaptive strategies (Lempert et al. 2006; Lempert and Collins 2007; Kasprzyk et al. 2013; Herman et al. 2014). (This RDM step provides a natural point of connection with the Dynamic Adaptive Policy Pathways (DAPP) methods in Chaps. 4 and 9.)

RDM uses both absolute and relative performance measures to compare strategies in the vulnerability and tradeoff analyses. Absolute performance measures are useful when decisionmakers are focused on one or more outcomes, such as profit, energy produced, or lives saved. Absolute performance measures are also useful when decisionmakers are focused on some invariant standard—for instance, a regulatory requirement on reliability or environmental quality, a required threshold for an economic rate of return, or a requirement that benefits exceed costs. Relative performance measures are often useful when uncertainties create a wide range of outcomes,

so decisionmakers seek strategies that perform well compared to alternatives over a wide range of futures. RDM often uses regret to represent relative performance.

At each of the RDM steps, information produced may suggest a re-framing of the decision challenge. The process produces key deliverables including (1) the scenarios that illuminate the vulnerabilities of the strategies, and (2) potential robust strategies and the tradeoffs among them.

The left panel of Fig. 2.2 shows the RDM steps in support of a process of deliberation with analysis. Stakeholders begin by deliberating over the initial decision framing. In the vulnerability and tradeoff analysis steps, stakeholders and analysts produce decision-relevant information products. Using these products, stakeholders deliberate over the choice of a robust strategy or return to problem framing, for instance seeking new alternatives or stress testing a proposed strategy over a wider range of futures. In practice, the process often moves back and forth between problem framing, generating the scenarios that illuminate vulnerabilities, identifying new alternatives based on those scenarios, and conducting a tradeoff analysis among the alternatives.

People teaming with computers—each doing what they do best—are more capable than computers or people alone.[5] RDM uses EM to support deliberation with analysis in a process of human/machine collaboration (Lempert et al. 2003). As shown in the right panel of Fig. 2.2, people use their creativity and understanding to pose questions or suggest solutions—for instance candidate robust strategies. Computers consider numerous combinations of strategies and futures to help users address their questions, search for initially unwelcome counter-examples to proposed solutions, and help people find new candidate robust strategies to propose (Lempert et al. 2002; Lempert and Popper 2005).

The RDM steps and deliberative processes are consistent with others in the DMDU literature. For instance, Multi-Objective RDM (MORDM) offers a similar iterative process but with the major advance of more articulation of the step of generating alternative strategies (Kasprzyk et al. 2013). The DAPP process (see Chaps. 4 and 9), also similarly, emphasizes monitoring and policy adjustment (Haasnoot et al. 2013). Among related literatures, many-objective visual analytics uses interactive visualizations to support problem framing and re-framing (Kollat and Reed 2007; Woodruff and Reed 2013), often with a posteriori elicitation of preferences (Maass et al. 1962; Cohon and Marks 1975).

Overall, the RDM process aims to provide quantitative decision support that helps meet the criteria for good decisions even in the presence of deep uncertainty and the other attributes of wicked problems. The process encourages participants to be explicit about their goals and consider the most important tradeoffs. The process uses scenario concepts linked to the idea of policy stress tests, along with computer-assisted exploration, to encourage and facilitate consideration of the decision from a wide range of views. It helps recognize the legitimacy of different interests, values, and expectations about the future by using models as exploratory, rather than

---

[5]For example, Thompson (2013) notes that competent chess players teamed with computers can defeat both grand masters without computers and computers without human assistants.

predictive, tools within an evidence-based, multi-scenario, multi-objective decision-making process.

Box 2.2: What is a Robust Strategy?

A robust strategy is one that performs well, compared to the alternatives, over a wide range of plausible futures (Rosenhead et al. 1972; Lempert et al. 2003). Other definitions exist, including trading some optimal performance for less sensitivity to assumptions (Lempert and Collins 2007) and keeping options open (Rosenhead 1990). Choosing a specific quantitative criterion to judge robustness can, however, prove complicated because many robustness criteria exist, and in some cases they can yield a different ordering of strategies (Giuliani and Castelletti 2016). No robustness criterion is best in all circumstances and, as befits a decision support methodology designed to facilitate problem framing, RDM often includes the choice of a robustness criterion as part of its problem-framing step.

RDM and other deliberative "agree-on-decision" DMDU methods generally draw robustness criteria from the normative, "agree-on-assumptions" decision-analytic literature. This literature identifies four traditional criteria —called Wald, Hurwicz, Savage, and Laplace—for ranking choices without well-defined probability distributions over future states of the world (Luce and Raiffa 1957; Schneller and Sphicas 1983). These criteria envision a set of future states of the world f_j and a set of strategies s_i, each with a known utility, u_{ij}, in each state of the world. If the probability of each future, p_j, were known, the best strategy would be that which yielded the maximum expected utility, $\text{Max}_i \sum_j p_j u_{ij}$. Lacking such probabilities, Wald selects the strategy that gives the best worse case, $\text{Max}_i \text{Min}_j (u_{ij})$, and Savage's mini-max regret selects the strategy with the least regret, that is, which deviates least from the best one could choose with perfect information, $\text{Min}_i \text{Max}_j [\text{Max}_i (u_{ij}) - u_{ij}]$. Both Wald and Savage are conservative in that they attempt to avoid worst cases. In contrast, Hurwicz interpolates between selecting the strategy with the best case and the best worst case, $\text{Max}_i [\alpha \text{Max}_j (u_{ij}) + (1 - \alpha) \text{Min}_j (u_{ij})]$, where α represents a level of confidence. Laplace's criterion of insufficient reason assumes equal weighting over all the futures and then selects the strategy that maximizes expected utility. Starr (1962) subsequently proposed the domain criterion, which selects the strategy that has highest utility in the most futures.

The domain criterion comes closest to the idea of performing well over a wide range of futures and can also facilitate tradeoff analyses using scenario-themed visualizations. RDM analyses commonly use some form of the domain criterion, although RDM (see Table 7.2 in Cervigni et al. 2015) and other robustness approaches (e.g., Hansen and Sargent 2008) have also employed Savage's mini-max regret. Traditionally, the domain criterion uses an absolute

performance measure (e.g., the utility u_{ij}, or a benefit–cost ratio (Lempert 2014)), as its argument. RDM often extends the criterion by also using as arguments relative performance measures, e.g., the regret $\text{Max}_i(u_{ij}) - u_{ij}$, or satisficing criteria for both absolute and relative performance measures. For instance, employed with a regret measure and satisficing criterion, the domain criterion might select the strategy that has a regret below some threshold value over the widest range of futures (Lempert and Collins 2007). In contrast to a strict application of Savage's mini-max regret, such a criterion would retain the comparative quality of the regret measure without undue attention to worst cases.

RDM also uses the domain criterion over the space of probabilities rather than the space of futures; that is, the analysis considers a set of plausible probability distributions over the futures (Walley 1991), and the domain criterion selects the strategy with the highest expected utility or expected regret over the largest number of distributions in the set (see, e.g., Fig. 2.3 in this chapter). RDM often uses the domain criterion, or other robustness decision criteria, to focus on tradeoffs rather than on a normative ranking of strategies. For instance, Lempert and Collins (2007) employ a definition of robustness as trading some optimal performance for less sensitivity to assumptions, which it implements with a domain criterion using expected regret as its argument. The study displays the choice of a robust strategy as a function of the confidence one places in the best-estimate distribution. The *Combined* Strategy discussed in Sect. 2.5.3 is robust by both a domain criterion and by mini-max regret.

2.4 Tools

The concepts underlying RDM—scenario thinking, robustness decision criteria, stress testing proposed plans, and the use of exploratory models—have long pedigrees. But over the last decade, new computer and analytic capabilities have made it possible to combine them in practical decision analyses. In particular, RDM often relies on Scenario Discovery and visualization, robust multi-objective optimization, integrated packages for EM, and high-performance computing.

SD algorithms often implement the RDM vulnerability analysis step (Step 3) in Fig. 2.1. SD begins with a large database of runs in which each model run represents the performance of a strategy in one future. The SD cluster-finding algorithms then offer concise descriptions of those combinations of future conditions that best distinguish the cases in which the implementation plan does or does not meet its goals. The requisite classification algorithms—often Patient Rule Induction Method (PRIM) (Friedman and Fisher 1999) or Classification and Regression Tree (CART) (Breiman et al. 1984), combined with a principal component analysis (Dalal et al.

2013)—seek to balance between the competing goals of simplicity and accuracy in order to describe sets of strategy-stressing futures as concise, understandable, and decision-relevant scenarios (Groves and Lempert 2007; Bryant and Lempert 2010). Software to implement the PRIM algorithm is available, both in stand-alone routines[6] and embedded in the EM software packages described below. Overall, SD replicates analytically the ideas of qualitative Intuitive Logics scenario analysis and provides information products that can prove compelling in stakeholder deliberations (Lempert 2013). RDM analyses also use computer visualization of the database of runs to support the vulnerability and tradeoff analyses. Tableau, a commercially available platform, has proven particularly useful in much RDM work (Groves et al. 2013; Cervigni et al. 2015).

The "new futures and strategies" step (Step 5) can employ a variety of methods. While some RDM analyses use only expert judgment to craft responses to potential vulnerabilities (e.g., Lempert and Groves 2010), many applications use some multi-objective robust optimization tool. Some such applications have used constrained optimization to trace out a range of potentially robust solutions, both in single objective (Lempert and Collins 2007) and multi-objective (Groves et al. 2012) cases. The latter instance involved a planning tool that allowed analysts to trace the Pareto optimal tradeoff curves in each scenario for any two objectives using constrained optimizations over the other objectives (Groves et al. 2012), and has been widely used. Other applications, such as the Colorado Basin Supply and Demand Study (Groves et al. 2013; also see Chap. 7), run large portfolio optimizations for many futures; identify the individual actions that occur in the optimal set for most, some, and few of the futures; and use this information to craft adaptive strategies that begin with the actions that occur in most of the futures' optimal sets, and implement the others depending on which future comes to pass (Groves et al. 2013; Bloom 2015). MORDM tools provide a more general solution, using evolutionary algorithms to identify the regions of a Pareto surface over many objectives that are most robust to uncertainty (Kasprzyk et al. 2013). MORDM has been used to identify adaptive strategies modeled as controllers in the control theory sense (Quinn et al. 2017), through a process called direct policy search.

While many RDM analyses use stand-alone software to generate and analyze large numbers of model runs, several integrated EM packages exist that can greatly facilitate such analysis. These packages generally consider a simulation model as a mapping of inputs to outputs and then provide connectivity between the model and various analysis and visualization tools. For instance, EM packages allow the user to specify an experimental design over the inputs; run the resulting cases on one or more processors; store the results in a database; use a variety of visualization routines to examine the data; run SD and other statistical algorithms on the data; and use robust optimization routines to identify new, potentially more robust strategies. While they require study and practice to use, EM software can make it much easier to move seamlessly through the steps of an RDM analysis. Computer-Assisted Reasoning system (CARs) was the earliest such software package (Bankes et al.

---

[6]https://cran.r-project.org/web/packages/sdtoolkit/index.html.

2001). More recently, new open-source platforms have become available, such as the Exploratory Modeling and Analysis Workbench,[7] Open MORDM (Hadka et al. 2015), and Rhodium.[8]

While many RDM analyses may be conveniently run on a laptop or desktop computer (e.g., Lempert et al. 2013b), recent studies have also used high-performance computation—either large-scale cluster computing (Zeff et al. 2014; Groves et al. 2016) or cloud-based computer services (Isley 2014; Cervigni et al. 2015)—to quickly and inexpensively conduct a very large number of runs.

2.5 Example: Carrots and Sticks for New Technology

An early RDM application provides an ideal example of the approach, the types of tools used, and how EM can draw together concepts from decision analysis, scenarios, and ABP. The example, called "carrots and sticks for new technology," focused on determining the most robust combination of two policy instruments—carbon prices and technology subsidies—to reduce climate-altering greenhouse gas emissions (Robalino and Lempert 2000; Lempert 2002). We describe this study using the generic DMDU framework presented in Sect. 1.5 (the major subheadings below) linked to the relevant RDM steps shown in Fig. 2.1.

2.5.1 Frame the Analysis

Formulate Question (*RDM Step 1*)

It is well-understood that an economically ideal greenhouse gas emission reduction policy should include an economy-wide carbon price implemented through mechanisms such as a carbon tax or a cap-and-trade system. This early RDM study addressed the question of whether and under what conditions technology incentives, such as technology-specific tax credits or subsidies, also prove necessary and important as part of a greenhouse gas reduction strategy. Many national and regional jurisdictions worldwide employ such incentives in their climate policies because they prove politically popular and have a compelling logic, if for no other reason than significant technology innovation will prove crucial to limiting climate change. But technology incentives have a mixed record of success (Cohen and Noll 2002), and sometimes link to larger debates about the appropriate role of government (Wolf 1993). Standard economic analysis proves a poor platform to adjudicate such questions, because the extent to which technology incentives prove economically

---

[7]http://simulation.tbm.tudelft.nl/ema-workbench/contents.html.

[8]https://github.com/Project-Platypus/Rhodium/wiki/Philosophy.

important may depend on coordination failures that occur in the presence of increasing returns to scale, imperfect information, and heterogeneous preferences—deeply uncertain factors not well-represented in standard economic models. Addressing the technology incentive question with RDM thus proved useful due to this deep uncertainty and because people's views on technology incentives can be strongly affected by their worldviews.

Identify Alternatives (*RDM Step 1*)

The study was organized as a head-to-head comparison between strategies incorporating two types of "policy instruments" or "actions": an economy-wide carbon tax and a technology price subsidy for low-carbon-emitting technologies. We considered four combinations: (1) neither taxes nor subsidies *(No Action)*, (2) tax only *(Taxes Only)*, (3) subsidies only *(Subsidies Only)*, and (4) a combination of both taxes and subsidies *(Combined)*. As described below, both the tax and subsidy were configured as adaptive strategies designed to evolve over time in response to new information.

Specify Objectives (*RDM Step 1*)

This study compared the strategies using two output measures: the present value of global economic output (pvGDP) and the mid-twenty-first-century level of greenhouse gas emissions. We focused on pvGDP to facilitate comparison of this work with other analyses in the climate change policy literature using more standard economic formulations. As described below, we calculated the regret for each strategy in each future and used a domain criterion for robustness (see Box 2.2), thus looking for strategies with low regret over a wide range of plausible futures.

Specify System Structure (*RDM Step 1*)

To compare these adaptive strategies, we employed an agent-based model of technology diffusion, linked to a simple macromodel of economic growth that focused on the social and economic factors that influence how economic actors choose to adopt, or not to adopt, new emission-reducing technologies. The agent-based representation proved useful, because it conveniently represents key factors potentially important to technology diffusion, such as the heterogeneity of technology preferences among economic actors and the flows of imperfect information that influence their decisions. Considering this tool as an exploratory, rather than predictive, model proved useful, because it allowed the study to make concrete and specific comparisons of price- and subsidy-based strategies even though, as described below, available theory and data allowed the model's key outputs to vary by over an order of magnitude.

   The model was rooted in the microeconomic understanding of the process of technology diffusion (Davies 1979). Each agent in our model represents a producer of a composite good, aggregated as total GDP, using energy as one key input. Each period the agents may switch their choice of energy-generation technology, choosing among high-, medium-, or low-emitting options. Agents choose technology to maximize their economic utility. The agents estimate utility based on their expectations regarding each technology's cost and performance. Costs may or may not decline

significantly due to increasing returns to scale as more agents choose to adopt. The agents have imperfect information and current and future technology cost and performance, but can gain information based on their own experience, if any, and by querying other agents who have used the technology. Thus, the model's technology diffusion rates depend reflexively on themselves, since by adopting a technology each agent generates new information that may influence the adoption decisions of other potential users. The model also used simple, standard, but deeply uncertain relationships from the literature on the connections between greenhouse gas emissions and the economic impacts due to climate change (Nordhaus 1994).

2.5.2 Perform Exploratory Uncertainty Analysis

Specify Uncertainties or Disagreements (*RDM Step 2*)

The agent-based model had thirty input parameters representing the deeply uncertain factors, including the macroeconomic effects of potentially distortionary taxes and subsidies on economic growth, the microeconomic preferences that agents use to make technology adoption decisions, the future cost and performance of high- (e.g., coal), low- (e.g., gas), and non-emitting (e.g., solar) technologies, the way information about new technologies flows through agent networks, and the impacts of climate change.

We employed three sources of information to constrain our EM. First, the agent-based model embodied the theoretical economic understanding of technology diffusion. Second, we drew plausible ranges for each individual parameter using estimates from the microeconomics literature. Third, we required the model to reproduce macroeconomic data regarding the last fifty years of economic growth and market shares for different types of energy technology. We also constrained future technology diffusion rates in the model to be no faster than the fastest such rates observed in the past.

Consistent with these constraints, the model nonetheless was able to generate a vast range of plausible futures. To choose a representative sample of futures from this vast set, the study launched a genetic algorithm over the model inputs, searching for the most diverse set of model inputs consistent with the theoretical, macroeconomic, and microeconomic constraints (Miller 1998). This process yielded an ensemble of 1,611 plausible futures, with each future characterized by a specific set of values for each of the thirty uncertain model input parameters. Each member of the ensemble reproduced the observed history from 1950 to 2000, but differed by up to an order of magnitude in projected mid-twenty-first-century emissions.

Evaluate Strategies Against Futures (*RDM Step 2*)

The study represented its carbon tax and technology subsidies as adaptive strategies using a single set of parameters for each to describe their initial conditions, and how they would evolve over time in response to new information.

As with many carbon price proposals, the study's carbon tax would start with an initial price per ton of CO_2 at time zero and rise at a fixed annual rate, subject to two conditions meant to reflect political constraints. First, we assumed that the government would not let the tax rise faster than the increase in the observed social cost of carbon. Second, we assumed that if the economy dropped into recession (defined as a global growth rate below some threshold), the government would drop the carbon tax back to its original level. Three parameters defined the tax strategy: the initial tax rate, the annual rate of increase, and the minimum threshold for economic growth required not to repeal the tax.

The study's technology subsidy reduces the cost to users of a technology to a fixed percentage of its unsubsidized cost. The subsidy stays in effect until the policy either succeeds in launching the technology (defined as its market share rising above some threshold value) or the policy fails (defined as the market share failing to reach a minimum threshold after a certain number of years). Meeting either of these conditions permanently terminates the subsidy. Four parameters defined the subsidy strategy: the subsidy level, the market share defining success, the market share defining failure, and the number of years before the subsidy can be judged a failure.

The study chose a single set of parameter values to define the tax and subsidy policies, each set chosen to optimize the pvGDP for the future represented by the average value for each of the thirty model parameters. The results of the study were relatively insensitive to this simplification.

We used the agent-based model to calculate the pvGDP and mid-century greenhouse gas emissions for each strategy in each of the 1,611 plausible futures. This ensemble of runs made it immediately clear that the *Taxes Only* and *Combined* strategies consistently perform better than the *No Action* and *Subsidies Only* strategies. The remainder of the analysis thus focused on the first two.

2.5.3   Choose Initial Actions and Contingent Actions

Illustrate Tradeoffs (*RDM Steps 3 and 4*)

Lacking (not-yet-developed) SD algorithms and faced with too many dimensions of uncertainty for an exhaustive search, the study used importance sampling to find the five uncertain input parameters most strongly correlated with mid-century GHG emissions. Four of these key uncertainties related to the potential for coordination failures to slow technology diffusion—the rate of cost reductions for non-emitting technologies caused by increasing returns to scale, the rate at which agents learn from one another about the performance of new technologies, the agents' risk aversion, and the agents' price-performance preferences for new technologies—and one related to the damages from climate change.

We then examined the regret in pvGDP across the cases for the *Taxes Only* and *Combined* strategies as a function of all ten two-dimensional combinations of the

five key uncertainties.[9] Each visualization told a similar story—that the *Combined* strategy cases had lower mean and variability in pvGDP regret than those for *Taxes Only* except in the corner of the uncertainty space with low potential for coordination failures and/or low impacts from climate change.

This analysis provides the study's basic comparison among the strategies: *Taxes Only* performs best when the potential for coordination failures and the impacts from climate change are small, and *Combined* performs best otherwise. Both strategies perform better than *No Action* or *Subsidies Only*.

*Select and Plan for Adaptation (RDM Step 4)*

To help decisionmakers understand the conditions under which the *Taxes Only* and *Combined* strategies would each be favored over the other, we first collapsed the ten visualizations into a single two-dimensional graph by combining all four key uncertain parameters relating to potential coordination failures into a single variable. Figure 2.3 shows the sets of expectations in which pvGDP for the *Taxes Only* strategy exceeds that of *Combined*, and vice versa, as a function of the probability assigned to high rather than low values of the four uncertain parameters related to potential coordination failures, which we labeled "probability of a non-classical world," and "probability of high damages" due to climate change. The region dominated by *Combined* is larger than that of *Taxes Only* for two reasons. First, *Combined* dominates *Taxes Only* over larger regions of the state space. Second, in the regions of the uncertainty space where *Taxes Only* is the better strategy, *Combined*'s regret is relatively small, while in the regions where *Combined* is better, *Taxes Only*'s regret is relatively large. Existing scientific understanding proves insufficient to define with certainty where the future lies in Fig. 2.3. Different parties to the decision may have different views. But the boundary between the two regions is consistent with increasing returns to scale much smaller than those observed for some energy technologies—such as natural gas turbines, wind, and solar—and the middle of the figure is consistent with relatively small levels of risk aversion, learning rates, and heterogeneity of preferences compared to those seen in various literatures. These results suggest that a combination of price instruments and technology subsidies may prove the most robust strategy over a wide range of plausible futures.

Implementation, Monitoring, and Communication

This study addressed a high-level question of policy architecture—the best mix of policy instruments for decarbonization. While the study did not provide detailed implementation plans, it does suggest how a national or state/provincial government might pursue the study's recommendations. The study envisions policymakers choosing a strategy which includes the rules by which the initial actions will be adapted over time (Swanson et al. 2007). The carbon price, presumably set by the legislature, would follow the social cost of carbon as periodically updated by executive agencies (National Academies of Sciences 2016) whenever the economy was not in recession.

---

[9]Ten combinations because $\binom{5}{2} = 10$.

Fig. 2.3 Scenario map comparing "Taxes Only" and "Combined" strategies
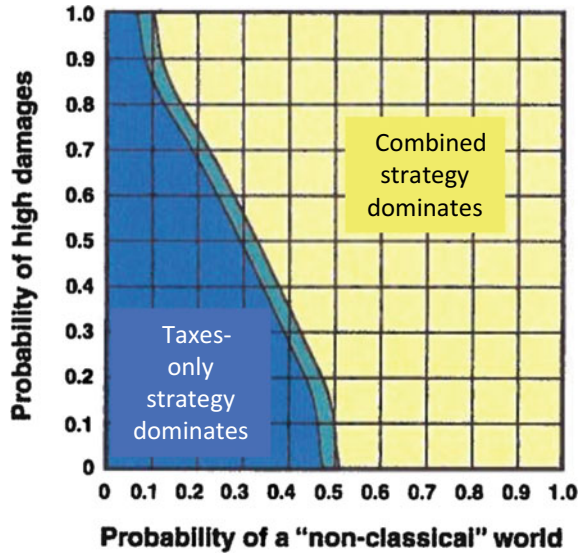
The legislature would also set the technology subsidy and terminate it when the subsidized technologies either succeeded or failed based on market share data gathered by executive agencies. The study did not examine pre-commitment issues—that is, how the current legislature could ensure that future legislatures would in fact follow the adaptive strategy. These interesting issues of political economy have, however, been recently explored using RDM methods (Isley et al. 2015).

2.5.4 Iterate and Re-Examine (RDM Steps 2, 3, and 5)

The study's results are based on an examination of only six of the 30 dimensions of uncertainty in the model, representing a small subset of the full range of plausible futures. As a key final step, we tested the policy recommendations by launching a genetic search algorithm across the previously unexamined dimensions looking for additional futures that would provide counter-examples to our conclusions. This process represents the computer feedback loop in the right panel of Fig. 2.2. The genetic algorithm ran for most of the time the authors spent writing their manuscript and found no plausible counter-examples.

Overall, this study suggests that if decisionmakers hold even modest expectations that market imperfections are likely to inhibit the diffusion of new, emissions-reducing technologies or that the impacts of climate change will turn out to be serious, then strategies combining both carbon taxes and technology incentives may be a promising component of a robust strategy for reducing greenhouse gas emissions.

2.6 Recent Advances and Future Challenges

The "carrots and sticks for new technology" example of Sect. 2.5 includes all the steps of an RDM analysis shown in Fig. 2.1. It used optimization algorithms to define its alternative, adaptive strategies, and generated its futures using genetic algorithms to perform what has more recently been called scenario diversity analysis (Carlsen et al. 2016a, b). The study employed the process of human-machine collaboration shown in Fig. 2.2, in particular in the computer search for counter-examples to the human-derived patterns that constitute its policy conclusions.

Since this early example, the methods and tools for MORDM analyses have approached maturity, now reaching the point at which one can describe with some specificity how to conduct multi-scenario, MORDM for many wicked problems. For instance, a recent study, described in detail in Chap. 7, used RDM on a topic similar to the "carrots and sticks" example—examining how international finance institutions such as the Green Climate Fund (GCF) can best craft long-term investment strategies to speed decarbonization in the face of deep technological and climate uncertainty (Molina Perez 2016). This more recent study was made possible by powerful new SD algorithms and visualization tools.

Recent work for four North Carolina cities illustrates the power of MORDM, a combination of RDM with new evolutionary algorithms for multi-objective robust optimization (Herman et al. 2014, 2016; Zeff et al. 2014, 2016). The study helped the neighboring cities of Raleigh, Durham, Chapel Hill, and Cary link their short-term operational and long-term investment water plans by shifting the former from rule-based procedures to new dynamic risk-of-failure triggers, and the latter from static to adaptive policy pathways. The study also helped the four independent cities coordinate their plans in the presence of different objectives and deep uncertainty. Linking such MORDM with the Dynamic Adaptive Policy Pathways (DAPP) approach of Chaps. 4 and 9 represents an exciting direction for the DMDU community.

Recent work for the US Bureau of Reclamation and the parties to the Colorado Compact, also described in Chap. 7, showcases RDM's ability to facilitate deliberation with analysis, helping contesting parties to agree on the vulnerabilities they face and adaptive strategies for addressing them (Groves et al. 2013; Bloom 2015).

Some important technical hurdles still remain before these capabilities fully mature. First, an approach is needed that provides full Pareto satisficing surfaces. Current MORDM analyses identify Pareto *optimal* surfaces for best-estimate cases and measure the robustness of alternative strategies, represented by different regions on the Pareto surface, to the deep uncertainties (Kasprzyk et al. 2013). In the future, MORDM could produce sets of strategies chosen specifically because their performance across multiple objectives was largely insensitive to the deep uncertainties. Recent work has taken steps toward providing such Pareto satisficing surfaces (Watson and Kasprzyk 2017), but more needs to be done. Furthermore, despite the availability of ubiquitous computation on the cloud, and through high-performance computation facilities, it still remains difficult in many cases to conduct a full MORDM analysis using realistic system models, which would require running many thou-

sands of cases to perform the multi-objective robust optimization over each of many thousands of scenarios. Research is needed on what we might call adaptive sampling approaches to help navigate more efficiently through the set of needed runs. In addition, research could usefully provide guidance on when to use alternative robustness criteria, as well as the conditions under which RDM's iterative analytic process is guaranteed to converge independent of the initial problem framing, or when path dependence may lead analyses to different answers (Kwakkel et al. 2016).

Finally, the cost of developing the needed system models often puts RDM analyses out of reach for many decisionmakers. Research on "RDM-lite"—means to quickly develop such models through approaches such as expert elicitation and participatory modeling—could greatly increase the use of these methods (see, e.g., O'Mahony et. al. 2018).

Evaluation plays a crucial role in the design and use of any successful decision support system (NRC 2009; Pidgeon and Fischhoff 2011; Wong-Parodi et al. 2016). Some evaluations of RDM tools, visualizations, and processes exist, both in the laboratory (Budescu et al. 2013; Parker et al. 2015; Gong et al. 2017) and through field experiments (Groves et al. 2008). Recent work has proposed frameworks for evaluating the impacts of RDM-based decision support in urban environments (Knopman and Lempert 2016). But much more such evaluation work is required to improve the practical application of RDM decision support (Bartels et. al., forthcoming).

More broadly, as DMDU methods reach technical maturity, they offer the opportunity to reshape the relationship between quantitative decision analytics and the way in which organizations use this information with their internal and external audiences and processes. The potential for such reshaping presents a rich menu of research needs to understand the organizational, anthropological, political, and ethical implications.

As one example, the concept of *risk governance* embeds risk management, which often has a narrow, more technocratic perspective, in a broader context that considers institutions, rules conventions, processes, and mechanisms through which humans acting as individuals and groups make choices affecting risk (Renn 2008). Recent work has explored how to embed RDM methods and tools in a risk governance framework (Knopman and Lempert 2016). For instance, RDM can help decisionmakers working within a multi-agent and multi-jurisdictional system organize their strategies into "tiers of transformation," which derives from the ideas of triple-loop learning. Lower tiers represent actions the decisionmakers can address on their own, while the outer tiers represent large-scale, transformative system changes that only the decisionmakers can help to catalyze. Any understanding of how to implement and use such capabilities, and the extent to which they would prove useful, remains nascent.

Future work can also usefully situate the types of moral reasoning and social choice embodied in alternative approaches to decision support. In his treatise, the *Idea of Justice*, Amartya Sen (2009) describes two classes of moral reasoning—the transcendental and the relational. The former, represented by Sen's teacher John Rawls (1971), seeks to inform ethical societal choices by first envisioning a common vision of a perfectly just world. People can then use that vision to inform their near-term choices. The latter, Sen's preferred alternative, rests on the assumption that

irreducible uncertainty about the consequences of our actions, and a diversity of priorities, goals, and values, are fundamental attributes of our world. Thus, no such transcendental vision of the type envisioned by Rawls is possible because the level of agreement and commonality of values it presupposes does not, and should not, exist in a diverse society in which people are free to pursue their lives according to their own, often very different, visions of what is good. In addition, Sen argues, even if a common transcendental vision were possible, it would prove insufficient to inform near-term choices, because human knowledge is too fallible and the uncertainties too deep to chart an unambiguous path to the ideal. But humans can obtain sufficient knowledge to craft near-term options and differentiate the better from the worse. Relational reasoning thus involves an iterative process of debating, choosing, learning, and revisiting choices, always trying to move in the direction of more justice in the face of imperfect knowledge and conflicting goals.

"Agree-on-assumptions" approaches to decision support reflect transcendental reasoning, while "agree-on-decisions" approaches reflect relational reasoning (Lempert et al. 2013c). Sen emphasizes the importance of deliberation in a relational process of social choice. His framework provides attributes for judging what constitutes an ethical process of deliberation with analysis. In particular, such deliberations work best when they recognize the inescapable plurality of competing views; facilitate re-examination and iterative assessments; demand clear explication of reasoning and logic; and recognize 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.

Recent work has pioneered methods for conducting ethical-epistemological analysis on the extent to which decision support products, methods, and systems meet such ethical criteria (Tuana 2013; Bessette et al. 2017; Mayer et al. 2017), but much more remains to be done. Today's world presents numerous, complex decision challenges—from sustainability to national security—that require quantitative decision support to successfully address. But "agree-on-assumptions" methods often lure decisionmakers toward over-confidence and can make it difficult to engage and promote consensus among participants with diverse expectations and interests. Such methods—built on the assumption that the decision analytics aim to provide a normative ranking of decision options—have their foundations in a time of computational poverty and rest on a narrow understanding of how quantitative information can best inform decisions. Recent years have seen an explosion of computational capabilities and a much richer understanding of effective decision support products and processes. RDM—a multi-objective, multi-scenario "agree-on-decision" approach—exploits these new capabilities and understanding to facilitate deliberative processes in which decisionmakers explore, frame, and reach consensus on the "wicked" problems that today's decisionmakers increasingly face.

References

Bankes, S. C. (1993). Exploratory modeling for policy analysis. *Operations Research, 41*(3), 435–449.

Bankes, S. C., Lempert, R. J., & Popper, S. W. (2001). Computer-assisted reasoning. *Computing in Science & Engineering, 3*(2), 71–77.

Bartels, E., Mikolic-Torreira, I, Popper, S. W., & Predd, J. (forthcoming). *What is the value proposition of analysis for decisionmaking*? Santa Monica, CA: RAND Corporation, PR-3485-RC.

Bessette, D. L., Mayer, L. A., Cwik, B., Vezer, M., Keller, K., Lempert, R., et al. (2017). Building a values-informed mental model for New Orleans climate risk management. *Risk Analysis, 37*(10), 1993–2004.

Bloom, E. (2015). *Changing midstream: Providing decision support for adaptive strategies using robust decision making.* Santa Monica, CA: RAND Corporation, RGSD-348.

Borio, C., Drehmann, M., & Tsatsaronis, K. (2014). Stress-testing macro stress testing: Does it live up to expectations? *Journal of Financial Stability, 12*(20), 3–15.

Breiman, L., Friedman, J. H., Olshen, R. A., & Stone, C. J. (1984). *Classification and regression trees*. Wadsworth Statistics/Probability Series. Monterey, CA: Wadsworth.

Brooks, A., Bennett, B., & Bankes, S. C. (1999). An application of exploratory analysis: The weapon mix problem. *Military Operations Research, 4*(1), 67–80.

Bryant, B. P., & Lempert, R. J. (2010). Thinking inside the box: A participatory, computer-assisted approach to scenario discovery. *Technological Forecasting and Social Change, 77,* 34–49.

Budescu, D. V., Lempert, R. J., Broomell, S., & Keller, K. (2013). Aided and unaided decisions with imprecise probabilities. *European Journal of Operational Research, 2*(1–2), 31–62.

Carlsen, H., Eriksson, E. A., Dreborg, K. H., Johansson, B., & Bodin, Ö. (2016a). Systematic exploration of scenario spaces. *Foresight, 18*(1), 59–75.

Carlsen, H., Lempert, R. J., Wikman-Svahn, P., & Schweizer, V. (2016b). Choosing small sets of policy-relevant scenarios by combining vulnerability and diversity approaches. *Environmental Modelling and Software, 84,* 155–164.

Carter, T. R., Jones, R. N., Lu, S. B. X., Conde, C., Mearns, L. O., O'Neill, B. C., et al. (2007). New assessment methods and the characterisation of future conditions. In M. L. Parry, O. F. Canziani, J. P. Palutikof, P. J. V. D. Linden, & C. E. Hanson, (Eds.), *Climate change 2007: Impacts, adaptation and vulnerability. Contribution of working group II to the fourth assessment report of the intergovernmental panel on climate change* (Vol. 1, pp. 33–171). Cambridge, UK: Cambridge University Press.

Cervigni, R., Liden, R., Neumann, J. E., & Strzepek, K. M. (Eds.). (2015). *Enhancing the climate resilience of Africa's infrastructure: The water and power sectors*. Africa Development Forum Series. Washington, DC: World Bank.

Cohen, L. R., & Noll, R. G. (2002). *Technology pork barrel*. Washington, DC: Brookings Institution Press.

Cohon, J., & Marks, D. (1975). A review and evaluation of multiobjective programing techniques. *Water Resources Research*, 11(2).

Dalal, S., Han, B., Lempert, R. J., Jaycocks, A., & Hackbarth, A. (2013). Improving scenario discovery using orthogonol rotations. *Environmental Modeling and Software, 48,* 1–16.

Davies, S. (1979). *The diffusion of process innovations*. Cambridge, MA: Cambridge University Press.

Dessai, S., & Hulme, M. (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.

Dewar, J. A., Builder, C. H., Hix, W. M., & Levin, M. H. (1993). Assumption-based planning—A planning tool for very uncertain times. Santa Monica, CA, RAND Corporation. https://www.rand.org/pubs/monograph_reports/MR114.html. Retrieved July 20, 2018.

Dewar, J. A. (2002). *Assumption-based planning—A tool for reducing avoidable surprises*. Cambridge: Cambridge University Press.

Fischbach, J. R., Lempert, R. J., Molina-Perez, E., Tariq, A., Finucane, M. L., & Hoss, F. (2015). *Managing water quality in the face of uncertainty: A robust decision-making demonstration for EPA's National Water Program*. Santa Monica, CA: RAND Corporation, PR-1148-EPA.

Friedman, J. H., & Fisher, N. I. (1999). Bump hunting in high-dimensional data. *Statistics and Computing, 9,* 123–143.

Ghile, Y. B., Taner, M. Ü., Brown, C., Grijsen, J. G., & Talbi, A. (2014). Bottom-up climate risk assessment of infrastructure investment in the Niger River Basin. *Climatic Change, 122,* 97–110.

Giuliani, M., & Castelletti, A. (2016). Is robustness really robust? How different definitions of robustness impact decision-making under climate change. *Climatic Change, 135*(3–4), 409–424.

Gong, M., Lempert, R. J., Parker, A. M., Mayer, L. A., Fischbach, J., Sisco, M., et al. (2017). Testing the scenario hypothesis: An experimental comparison of scenarios and forecasts for decision support in a complex decision environment. *Environmental Modeling and Software, 91,* 135–155.

Groves, D. G., Bloom, E. W., Lempert, R. J., Fischbach, J. R., Nevills, J., & Goshi, B. (2014). Developing key indicators for adaptive water planning. *Journal of Water Resources Planning Management, 141*(7).

Groves, D. G., Fischbach, J. R., Bloom, E., Knopman, D., & Keefe, R. (2013). *Adapting to a changing Colorado river: Making future water deliveries more reliable through robust management strategies*. Santa Monica, CA: RAND Corporation, RR-242-BOR.

Groves, D. G., Knopman, D., Lempert, R. J., Berry, S., & Wainfan, L. (2008). *Presenting uncertainty about climate change to water resource managers—Summary of workshops with the Inland Empire Utilities Agency*. Santa Monica, CA: RAND Corporation, TR-505-NSF.

Groves, D. G., & Lempert, R. J. (2007). A new analytic method for finding policy-relevant scenarios. *Global Environmental Change, 17,* 73–85.

Groves, D. G., Lempert, R. J., May, D. W., Leek, J. R., & Syme, J. (2016). *Using high-performance computing to support water resource planning, a workshop demonstration of real-time analytic facilitation for the Colorado River Basin*. Santa Monica, CA: RAND Corporation and LLNL, CF-339-RC.

Groves, D. G., Sharon, C., & Knopman, D. (2012). *Planning tool to support Louisiana's decision-making on coastal protection and restoration*. Santa Monica, CA: RAND Corporation, TR-1266-CPRA.

Haasnoot, M., Kwakkel, J. H., Walker, W. E., & ter Maat, J. (2013). Dynamic adaptive policy pathways: A new method for crafting robust decisions for a deeply uncertain world. *Global Environmental Change, 23*(2), 485–498.

Hadka, D., Herman, J., Reed, P., & Keller, K. (2015). An open source framework for many-objective robust decision making. *Environmental Modelling and Software, 74,* 129–144.

Hall, J. M., Lempert, R. J., Keller, K., Hackbarth, A., Mijere, C., & McInerney, D. (2012). Robust Climate Policies under uncertainty: A comparison of Info-Gap and RDM methods. *Risk Analysis, 32*(10), 1657–1672.

Hansen, L. P., & Sargent, T. J. (2008). *Robustness*. Princeton, NJ: Princeton University Press.

Helgeson, C. (2018). Structuring decisions under deep uncertainty. *Topoi*, pp. 1–13.

Herman, J., Zeff, H., Lamontagne, J., Reed, P., & Characklis, G. (2016). Synthetic drought scenario generation to support bottom-up water supply vulnerability assessments. *Journal of Water Resources Planning and Management,* 142(11).

Herman, J., Zeff, H., Reed, P., & Characklis, G. (2014). Beyond optimality: Multistakeholder robustness tradeoffs for regional water portfolio planning under deep uncertainty. *Water Resources Research, 50*(10), 7692–7713.

Herrick, C., & Sarewitz, D. (2000). Ex post evaluation: A more effective role for scientific assessments in evnironmental policy. *Science, Technology and Human Values, 25*(3), 309–331.

Hodges, J. (1991). Six (or so) things you can do with a bad model. *Operations Research, 39*(3), 355–365.

IPCC. (2001). *Climate change 2001: Mitigation, intergovernmental panel on climate change*.

Isley, S. (2014). *The political sustainability of carbon control policies in an evolutionary economics setting.* Santa Monica, CA: RAND Corporation, RGSD-331.

Isley, S. C., Lempert, R. J., Popper, S. W., & Vardavas, R. (2015). The effect of near-term policy choices on long-term greenhouse gas transformation pathways. *Global Environmental Change, 34,* 147–158.

Jones, R. N., Patwardhan, A., Cohen, S., Dessai, S., Lammel, A., Lempert, R. J., et al. (2014). Foundations for decision making. In *Climate change 2014: Impacts, adaptation, and vulnerability.* Intergovernmental Panel on Climate Change (IPCC).

Kalra, N., Hallegatte, S., Lempert, R. J., Brown, C., Fozzard, A., Gill, S., et al. (2014). Agreeing on robust decisions: A new process of decision making under deep uncertainty. *Policy Research Working Paper.* World Bank, WPS-6906.

Kasprzyk, J. R., Nataraj, S., Reed, P. M., & Lempert, R. J. (2013). Many-objective robust decision making for complex environmental systems undergoing change. *Environmental Modeling and Software, 42,* 55–71.

Knopman, D., & Lempert R. J. (2016). *Urban responses to climate change: Framework for decision-making and supporting indicators* (156). Santa Monica, CA: RAND Corporation, RR-1144-MCF.

Kollat, J., & Reed, P. (2007). A framework for visually interactive decision-making and design using evolutionary multi-objective optimization (VIDEO). *Environmental Modeling and Software, 22*(12), 1691–1704.

Kwakkel, J. H., Haasnoot, M., & Walker, W. E. (2016). Comparing robust decision-making and dynamic adaptive policy pathways for model-based decision support under deep uncertainty. *Environmental Modelling and Software, 86,* 168–183.

Lempert, R. J. (2013). Scenarios that illuminate vulnerabilities and robust responses. *Climatic Change, 117,* 627–646.

Lempert, R. J., & Groves, D. G. (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. J., Popper, S. W., Groves, D. G., Kalra, N., Fischbach, J. R., Bankes, S. C. et al. (2013a). *Making Good Decisions Without Predictions: Robust Decision Making for Planning Under Deep Uncertainty.* Santa Monica, CA: RAND, RB-9701.

Lempert, R. J., Kalra, N., Peyraud, S., Mao, Z., Tan, S. B., Cira, D., & Lotsch, A. (2013b). Ensuring Robust Flood Risk Management in Ho Chi Minh City: A robust decision-making demonstration. World Bank, WPS-6465.

Lempert, R. J., Groves, D. G., & Fischbach, J. (2013c). *Is it ethical to use a single probability density function*? Santa Monica, CA: RAND Corporation, WR-992.

Lempert, R. J., & Kalra, N. (2011). *Managing climate risks in developing countries with robust decision making.* Washington, DC: World Resources Institute. https://www.wri.org/our-work/project/world-resources-report/managing-uncertainty. Retrieved October 12, 2018.

Lempert, R. J., Nakicenovic, N., Sarewitz, D., & Schlesinger, M. (2004). Characterizing climate-change uncertainties for decision-makers—An editorial essay. *Climatic Change, 65*(1–2), 1–9.

Lempert, R. J. (2002). A new decision sciences for complex systems. *Proceedings of the National Academy of Sciences, 99*(3), 7309–7313.

Lempert, R. J. (2007). Can scenarios help policymakers be both bold and careful? In F. Fukuyama (Ed.), *Blindside: How to anticipate forcing events and wild cards in global politics.* Washington, DC: Brookings Institute Press.

Lempert, R. J. (2014). Embedding (some) benefit-cost concepts into decision support processes with deep uncertainty. *Journal of Benefit Cost Analysis, 5*(3), 487–514.

Lempert, R. J., & Collins, M. (2007). Managing the risk of uncertain threshold responses: Comparison of robust, optimum, and precautionary approaches. *Risk Analysis, 27*(4), 1009–1026.

Lempert, R. J., Groves, D. G., Popper, S. W., & Bankes, S. C. (2006). A general, analytic method for generating robust strategies and narrative scenarios. *Management Science, 52*(4), 514–528.

Lempert, R. J., & Popper, S. W. (2005). High-Performance Government in an Uncertain World. In R. Klitgaard & P. Light (Eds.), *High Performance Government: Structure, Leadership, and Incentives*. Santa Monica, CA: RAND Corporation.

Lempert, R. J., Popper, S. W., & Bankes, S. C. (2002). Confronting Surprise. *Social Science Computer Review, 20*(4), 420–440.

Lempert, R. J., Popper, S. W., & Bankes, S. C. (2003). *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-term Policy Analysis*. Santa Monica, CA, RAND Corporation, MR-1626-RPC.

Lempert, R. J., Schlesinger, M. E., & Bankes, S. C. (1996). When we don't know the costs or the benefits: Adaptive strategies for abating climate change. *Climatic Change, 33*(2), 235–274.

Lempert, R. J., Schlesinger, M. E., Bankes, S. C., & Andronova, N. G. (2000). The impact of variability on near-term climate-change policy choices. *Climatic Change, 45*(1), 129–161.

Light, P. C. (2005). *The four pillars of high performance: how robust organizations achieve extraordinary results*. New York, NY: McGraw-Hill.

Luce, R. D., & Raiffa, H. (1957). *Games and decisions*. New York: Wiley.

Maass, A., Hufschmidt, M., Dorfman, R., Thomas, H., Marglin, S. Fair, G. et al. (1962). *Design of water resources systems; new techniques for relating economic objectives, engineering analysis, and governmental planning*. Cambridge, Harvard University Press.

March, J. (1994). *A primer on decision making: how decisions happen*. Toronto, Canada: The Free Press.

Mayer, L. A., Loa, K., Cwik, B., Tuana, N., Keller, K., Gonnerman, C., et al. (2017). Understanding scientists' computational modeling decisions about climate risk management strategies using values-informed mental models. *Global Environmental Change, 42,* 107–116.

Miller, J. H. (1998). Active Nonlinear Tests (ANTs) of complex simulations models. *Management Science, 44*(6), 820–830.

Molina Perez, E. (2016). *Directed international technological change and climate policy: New methods for identifying robust policies under conditions of deep uncertainty*, Pardee RAND Graduate School, RGSD-369.

Morgan, M. G., & Henrion, M. (1990). *Uncertainty: A guide to dealing with uncertainty in quantitative risk and policy analysis*. Cambridge, UK: Cambridge University Press.

National Academies of Sciences, E., and Medicine. 2016 (NAS) (2016). *Assessment of Approaches to Updating the Social Cost of Carbon: Phase 1 Report on a Near-Term Update*. Washington, DC, The National Academies Press.

National Research Council (NRC) (2009). *Informing decisions in a changing climate*. National Academies Press.

Nordhaus, W. D. (1994). *Managing the global commons: the economics of climate change*. Cambridge, MA: MIT Press.

O'Mahony, A., Blum, I., Armenta, G., Burger, N., Mendelsohn, J., McNerney, M. et al. (2018). *Assessing, Monitoring and Evaluating Army Security Cooperation: A Framework for Implementation*. Santa Monica, CA: RAND Corporation, RR-2165-A.

Parker, A. M., Srinivasan, S., Lempert, R. J., & Berry, C. (2015). Evaluating simulation-derived scenarios for effective decision support. *Technological Forecasting and Social Change, 91,* 64–77.

Pidgeon, N., & Fischhoff, B. (2011). The role of social and decision sciences in communicating uncertain climate risks. *Nature Climate Change, 1,* 35–41.

Popper, S. W., Berrebi, C., Griffin, J., Light, T., Min, E. Y., & Crane, K. (2009). *Natural gas and Israel's energy future: Near-term decisions from a strategic perspective*. Santa Monica, CA, RAND Corporation, MG-927.

Popper, S. W., Lempert, R. J., & Bankes, S. C. (2005). Shaping the future. *Scientific American, 292*(4), 66–71.

Quinn, J. D., Reed, P. M., & Keller, K. (2017). Direct policy search for robust multi-objective management of deeply uncertain socio-ecological tipping points. *Environmental Modelling and Software, 92,* 125–141.

Ranger, N., Millner, A., Dietz, S., Fankhauser, S., Lopez, A., & Ruta, G. (2010). *Adaptation in the UK: A decision-making process*. London: Granthan/CCEP Policy Brief.

Rawls, J. (1971). *A theory of justice*. Harward University Press

Rayner, S. (2000). Prediction and other approaches to climate change policy. In D. Sarewitz (Ed.), *Prediction: Science, decision making, and the future of nature* (pp. 269–296). Washington, DC: Island Press.

Renn, O. (2008). *Risk governance: Coping with uncertainty in a complex world*. London Earthscan.

Rittel, H., & Webber, M. (1973). Dilemmas in a general theory of planning. *Policy Sciences, 4,* 155–169.

Robalino, D. A., & Lempert, R. J. (2000). Carrots and sticks for new technology: Abating greenhouse gas emissions in a heterogeneous and uncertain world. *Integrated Assessment, 1*(1), 1–19.

Rosenhead, J. (1990). Rational analysis: Keeping your options open. In: J. Rosenhead & J. Mingers (Eds.), *Rational analysis for a problematic world: Problem structuring methods for complexity, uncertainty and conflict*. Chichester, England, Wiley.

Rosenhead, M. J., Elton, M., & Gupta, S. K. (1972). Robustness and optimality as criteria for strategic decisions. *Operational Research Quarterly, 23*(4), 413–430.

Rydell, C., Peter, J. P. Caulkins, & S. S. Everingham. (1997). *Enforcement or treatment? Modeling the relative efficacy of alternatives for controlling cocaine*. Santa Monica, CA, RAND Corporation, RP-614.

Sarewitz, D., & Pielke, R. A. (2000). *Science, prediction: Decisionmaking, and the future of nature*. Washington, DC: Island Press.

Schneller, G. O., & Sphicas, G. P. (1983). Decision making under uncertainty: Starr's Domain Criterion. *Theory and Decision, 15,* 321–336.

Schoemaker, P. J. H. (1993). Multiple scenario development: Its conceptual and behavioral foundation. *Strategic Management Journal, 14*(3), 193–213.

Schwartz, P. (1996). *The art of the long view—planning for the future in an uncertain world*. New York, NY: Currency-Doubleday.

Sen, A. (2009). *The idea of justice*. Cambridge, Massachusetts: Belknap Press.

Starr, M. K. (1962). *Product design and decision theory*. Englewood Cliffs, NJ: Prentice-Hall.

Swanson, D., & Bhadwal, S. (2009). *Creating adaptive policies: A guide for policy-making in an uncertain world*. Sage Publications.

Swanson, D., Venema, H., Barg, S., Tyler, S., Drexage, J., Bhandari P., & Kelkar, U. (2007). *Initial conceptual framework and literature review for understanding adaptive policies*.

Tetlock, P. E., & Gardner, D. (2016). *Superforecasting: The art and science of prediction*. Random House.

Thompson, C. (2013). *Smarter than you think: How technology is changing our minds for the better*. New York: Penguin.

Tuana, N. (2013). Embedding philosophers in the practices of science: Bringing humanities to the sciences. *Synthese, 190*(11), 1955–1973.

Van der Heijden, K. (1996). *Scenarios: The art of strategic conversation*. New York: Wiley.

Wack, P. (1985). The gentle art of reperceiving—scenarios: Uncharted waters ahead (part 1 of a two-part article). *Harvard Business Review (September–October)*: 73–89.

Walker, W., Marchau, V., & Swanson, D. (2010). Addressing deep uncertainty using adaptive policies. *Technology Forecasting and Social Change, 77,* 917–923.

Walker, W. E., Rahman, S. A., & Cave, J. (2001). Adaptive policies, policy analysis, and policy-making. *European Journal of Operational Research, 128,* 282–289.

Walley, P. (1991). *Statistical reasoning with imprecise probabilities*. London: Chapman and Hall.

Watson, A. A., & Kasprzyk, J. R. (2017). Incorporating deeply uncertain factors into the many objective search process. *Environmental Modeling and Software, 89,* 159–171.

Weaver, C. P., Lempert, R. J., Brown, C., Hall, J. A., Revell, D., & Sarewitz, D. (2013). Improving the contribution of climate model information to decision making: The value and demands of robust decision frameworks. *WIREs Climate Change, 4,* 39–60.

Wolf, C. (1993). *Markets or governments: Choosing between imperfect alternatives*. MIT Press.

Wong-Parodi, G., Krishnamurti, T., Davis, A., Schwartz, D., & Fischhoff, B. (2016). A decision science approach for integrating social science in climate and energy solutions. *Nature Climate Change, 6,* 563–569.

Woodruff, M., & Reed, P. (2013). Many objective visual analytics: Rethinking the design of complex engineered systems. *Structural and Multidiciplinary Optimization, 48*(1), 201–219.

Zeff, H., Herman, J., Reed, P., & Characklis, G. (2016). Cooperative drought adaptation: integrating infrastructure development, conservation, and water transfers into adaptive policy pathways. *Water Resources Research, 52*(9), 7327–7346.

Zeff, H. B., Kasprzyk, J. R., Herman, J. D., Reed, P. M., & Characklis, G. W. (2014). Navigating financial and supply reliability tradeoffs in regional drought management portfolios. *Water Resources Research, 50*(6), 4906–4923.

Zenko, M. (2015). *Red Team: How to succeed by thinking like the enemy*. Basic Books.

Robert Lempert is a principal researcher at the RAND Corporation and Director of the Frederick S. Pardee Center for Longer Range Global Policy and the Future Human Condition. His research focuses on risk management and decision-making under conditions of deep uncertainty. Dr. Lempert is a Fellow of the American Physical Society, a member of the Council on Foreign Relations, a chapter lead for the Fourth US National Climate Assessment, and a convening lead author for Working Group II of the United Nation's Intergovernmental Panel on Climate Change (IPCC) Sixth Assessment Report. Dr. Lempert was the Inaugural EADS Distinguished Visitor in Energy and Environment at the American Academy in Berlin and serves as the inaugural president of the Society for Decision Making Under Deep Uncertainty. A Professor of Policy Analysis in the Pardee RAND Graduate School, Dr. Lempert is an author of the book *Shaping the Next One Hundred Years: New Methods for Quantitative, Longer-Term Policy Analysis*.