

Structuring Decisions Under Deep Uncertainty*

Casey Helgeson[†]

ABSTRACT. Innovative research on decision making under 'deep uncertainty' is underway in applied fields such as engineering and operational research, largely outside the view of normative theorists grounded in decision theory. Applied methods and tools for decision support under deep uncertainty go beyond standard decision theory in the attention that they give to the *structuring* (also called framing) of decisions. Decision structuring is an important part of a broader philosophy of managing uncertainty in decision making, and normative decision theorists can both learn from, and contribute to, the growing deep uncertainty decision support literature.

1 Introduction

The past decade has seen a rapid expansion of research in the decision sciences aimed at decision-making contexts characterised by 'deep uncertainty'. Also called 'severe', 'extreme', or 'great' uncertainty, *deep uncertainty* refers loosely to contexts in which decision makers lack complete information about (or cannot agree on) the probabilities for key contingencies, the availability of present and future actions, the outcomes to which available actions lead, or the value of these outcomes (see, e.g., Lempert et al. 2006; Walker et al. 2013). Other formulations emphasise the complexity of the system on which a policy choice will intervene, or the richness of the universe of policy options to choose from (Popper 2016). In any case, the motivating idea is that there are contexts for which orthodox decision theory based on maximising expected utility—and more broadly, methods that aim for *optimal* solutions—are poorly suited because the uncertainties (broadly construed) are, in some sense, *too severe*. Here, new approaches to rational decision making are required. Climate change mitigation and adaptation decisions are common examples, but there are many others including contexts in finance, defence, resource management, infrastructure planning, and regulation of new technologies.

A portion of this new research appears in economics or philosophy journals such as *Theory and Decision*, *Games and Economic Behavior*, *Econometrica*, *Journal of Mathematical*

*This work was supported by the Arts and Humanities Research Council through the Managing Severe Uncertainty project (AH/J006033/1), the Agence Nationale de la Recherche through Decision-Making & Belief Change Under Severe Uncertainty: A Confidence-Based Approach (DUSUCA) (ANR-14-CE29-0003-01), and the National Science Foundation through the Network for Sustainable Climate Risk Management (SCRiM) (GEO-1240507).

[†]Earth and Environmental Systems Institute, Pennsylvania State University, casey.helgeson@gmail.com

Economics, Economics and Philosophy, Philosophy of Science, and Synthese, and plants itself firmly in the tradition that philosophers and social scientists know as *decision theory*.¹ A second and largely separate slice of deep uncertainty decision research comes from more applied fields like operational research, policy analysis, management science, and engineering. This work is more likely to be labeled ‘decision analysis’, or ‘decision support’, and appears in journals such as *Risk Analysis, Water Resources Research, Sustainability, Global Environmental Change, Ecological Applications, EURO Journal on Decision Processes*, and *Environmental Modelling and Software*.

This more applied branch of research (hereafter ‘deep uncertainty decision support’) is often tied to the development of best practices in particular areas of application such as climate change adaptation (Wilby and Dessai 2010; Brown et al. 2012; Ranger et al. 2013; Weaver et al. 2013) and water resources management (Haasnoot et al. 2011, 2012; Kasprzyk et al. 2012; Kwakkel et al. 2015; Herman et al. 2016). Frequently discussed methods include Robust Decision Making, Decision Scaling, Info-gap Decision Theory, Adaptive Policymaking, Adaptive Pathways, Robust Optimisation, and Direct Policy Search. Core contributions to this literature have come from non-academic (or not purely academic) research centers such as the RAND Corporation in the United States and Deltares in the Netherlands.

Philosophy has so far paid scant attention to deep uncertainty decision support,² but there are several reasons why this ought to change. Deep uncertainty decision support is of great practical importance to society. Practical methods and tools for supporting decision making under deep uncertainty are developing quickly and will be used ever more broadly given—among many other applications—a looming host of climate change adaptation decisions (Hewitt et al. 2012; American Meteorological Society 2015). These tools are often technically complex and computationally intensive, making it particularly important that outside observers equipped to do so ‘look under the hood’ and critique what they see. Philosophers working in a variety of subfields are well placed to both critique such decision support tools, and contribute to their further development.

Philosophy may also have much to gain from engaging with deep uncertainty decision support. Despite residing at the more applied end of the spectrum, research on decision support is continuous with the (often more theoretical) concerns of philosophers working in rational choice and adjacent areas of formal epistemology and philosophy of science. Indeed, some of this applied research appears to be addressing philosophically important questions about managing uncertainty in decision making that are at present neglected by the philosophical literature (more on this below). Moreover, in many areas of application, deep uncertainty decision support is changing the landscape of science-policy interaction over which philosophical debates such as the role of values in science advising (Douglas 2009, 2016; Steele 2012; Betz 2013) will play out in the future.

The purpose of this paper is to draw attention to this applied literature and encourage

¹While economists develop decision models mainly for descriptive purposes, philosophers, policy analysts, and others evaluate the merits of those models as normative guides or apply them as such. It is only this normative reading that I address here.

²Mitchell (2009, 85–104) and Sprenger (2012) are two exceptions.

philosophical engagement with it by lowering some of the disciplinary hurdles that stand in the way. (And though this paper should not be read as a survey of deep uncertainty decision support research, I have endeavoured to include enough references for it to serve as a useful gateway into that literature.)

Anecdotally, factors contributing to deep uncertainty decision support remaining largely off normative decision theorists' radar include a number of significant disciplinary barriers between the two branches of research outlined above, including differences in language, assumptions, notation, methods, and style. Axiomatisations and representation theorems, for example, are entirely absent from the applied literature. Even more disorienting, applied frameworks often lack a discernible *decision rule*—the method of comparing acts and determining which is best. There is little talk of *rationality*, or of what constitutes rational behaviour. Consequently, researchers versed in decision theory can find themselves puzzled about what the applied research really contributes to the study of making good decisions. Here I argue for what I think is (at least part of) the answer.

Real-world decision making is a process comprising several stages. A key division widely recognised in the decision sciences separates the *framing* or *structuring* of the decision from the subsequent *choice task*. (The distinction needn't be sharp, and a process may be iterative, returning to the framing task after a provisional choice.) The resulting two-stage picture of decision processes provides a useful schema for locating the applied field of deep uncertainty decision support with respect to the parallel exploration of deep uncertainty that is based strictly within the tradition of decision theory. A key difference between the two branches of research lies in the attention they give to the *structuring of decisions*. The principal focus of many deep uncertainty decision support practices is to bring about better decisions through sound decision structuring. While the importance of structuring is recognised within the theoretical literature, normative guidance offered on the topic is limited.

In what follows, I briefly review approaches to deep uncertainty decision making in the decision theory literature, using the formation of the state-consequence matrix as a convenient marker of the transition between structuring and choosing, and highlighting that the innovations often thought to adapt decision theory models to contexts of deeper uncertainty primarily target the choosing phase. I then introduce a selection of five techniques from the deep uncertainty decision support literature, and explain how the advice contained in those methods contributes to the structuring of decisions. *Bottom-up* approaches to decision making (§3.1), and a specific bottom-up approach called *scenario discovery* (§3.2), offer advice on partitioning state spaces. *Iterative restructuring* (§3.3) of a decision problem is a common feature of deep uncertainty decision support through which state space partitions and menus of acts coevolve. *Adaptive policymaking* (§3.4) and *multi-objective robust optimisation* (§3.5) contribute to forming the menu of acts for further consideration.

2 Deep uncertainty in decision theory

Formal decision models of the kind widely used and discussed in the social sciences and in

philosophy can be applied to a decision only once the problem has been formulated as a state-consequence matrix (Figure 1), which lists the actions under consideration (acts 1– n) and exhaustively divides up the different ways that relevant contingencies could play out (states 1– m). For example, the acts could be bets placed prior to a roll of the roulette wheel, and the states the slots into which the ball can come to rest. The state-consequence matrix also associates each act-state pair with a consequence of interest to the decision maker: if you bet on thirteen and the ball lands there, you win!

	state 1	state 2	...	state m
act 1	$c_{1,1}$	$c_{1,2}$...	$c_{1,m}$
act 2	$c_{2,1}$	$c_{2,2}$...	$c_{2,m}$
⋮				
act n	$c_{n,1}$	$c_{n,2}$...	$c_{n,m}$

Table 1: A state-consequence matrix. States 1– m are alternative resolutions to uncertain contingencies. Consequence $c_{i,j}$ results from choosing act i in state j .

Presupposing such a matrix means starting from the assumption that the decision maker has already *structured* the decision in the sense that they have conceived of and articulated all of the actions worthy of consideration, that they have identified every contingency that matters to how their action will turn out, and that they have carved up this space of possibilities in a way that allows them to perceive the possible outcomes of each action.

Building on the foundation of the state-consequence matrix, orthodox decision theory further supposes beliefs and preferences that are sufficiently rich as to determine a unique probability distribution over the states and a cardinal utility function over the consequences. In this case, the rational choice is whichever action produces the most utility on average, weighted by the probabilities of the states. Agents should, in other words, *maximise expected utility*.

Other decision rules aim to order the acts rationality-wise for agents with beliefs and preferences that are sparser, fuzzier, or less tractable than precise probabilities and cardinal utilities. These approaches to evaluating one's options can thus be seen as addressing contexts of deep uncertainty—or at least *deeper* uncertainty. Appeals to *dominance*, for example, can (sometimes) rank options in the absence of any probability-like information. The same goes for judging acts by their worst-case consequence (minimax). A variety of approaches represent the agent's beliefs as a set of probability functions rather than a single one (Gilboa and Schmeidler 1989; Ghirardato et al. 2004; Binmore 2009; Gilboa and Marinacci 2013), or a set of probability functions with second-order weights (Gärdenfors and Sahlin 1982; Klibanoff et al. 2005; Maccheroni et al. 2006; Chateauneuf and Faro 2009) or other super-structure (Hill 2013, 2016). Analogous moves have also been explored for the utilities (e.g., Fishburn 1964; Galaabaatar and Karni 2012, 2013; Bradley 2017).

But while these approaches relax traditionally demanding requirements on the richness and specificity of the decision maker's attitudes, they continue to rely on the foundation of a pre-formed state-consequence matrix that defines the objects of those attitudes. That is, they

assume a decision maker who has already structured the decision before them.³

This is not to say that decision theory has no advice to give about structuring. Formal decision models make a variety of assumptions about the contents of the state-consequence matrix—for example, about dependencies between acts, states, and consequences (see, e.g., Binmore, 2009, §1.4; Bradley, 2017, §1.2–1.3). When the models are read normatively, these assumptions can be construed as advice about how one ought to structure one’s decisions. (For example, structure your decision in such a way that states and acts are probabilistically and causally independent.) Such advice offers constraints on how one should encode the decision problem into the form of a matrix so that the mathematics of the subsequently applied decision rule or value function will work as designed.

Beyond these minimal constraints hard-wired into particular decision-theoretic frameworks, decision structuring is often considered more art than science. Still, some theoreticians have offered supplemental qualitative structuring advice that goes beyond the minimal constraints (some of which will be discussed below). For the purpose of aligning the theoretical and applied literatures to orient theoreticians within the applied work, structuring advice from deep uncertainty decision support can be set next to the combined collection of both the model-based constraints and the supplemental soft structuring advice from the decision theory side. In order to facilitate deeper comparisons across the literatures, I will gesture at specific points of contact throughout the following section.

3 Deep uncertainty decision support

Moving on to the applied research in deep uncertainty decision support, I now pick out five approaches from this literature, in each case highlighting how the approach contributes to the task of structuring decisions.

3.1 Bottom-up approaches

What has been called the *top-down* (also *scenario-based*) approach to climate change adaptation planning begins by predicting—as best we can—the future climate of a given region, at the greatest feasible level of detail. This may, for example, take the form of probabilistic forecasts for a large array of climate variables that might matter to regional or municipal adaptation planning (like managing a city’s water supply). In light of the forecast, decision makers devise a menu of planning options and then choose among them. The general idea is to first understand how the future is likely to unfold, then formulate options around that understanding.

³Though the literature does also include proposals about how to redistribute probability after minor restructuring of the state space partition (Walker and Dietz 2011; Binmore 2015) and incorporate awareness that a state space is incomplete (Karni and Vierø 2013, 2014).

In more explicit decision theoretic terms, the top-down approach begins by delimiting and partitioning the *state space* (states 1– m , Table 1), and modelling uncertainty over those states (for example, by assigning a probability distribution). This much is done on the basis of scientific understanding and general-purpose modelling of regional climate change, with no special attention to the particulars of the decision context. The list of *acts* to be evaluated (acts 1– n , Table 1) is developed in response to this context-independent structuring of the state space, with the potential *consequences* of each act (the $c_{i,j}$ of Table 1) spelled out at a level of detail determined by the context-independent state space partition.

One difficulty with this approach is that generating the required climate change projections can be costly and time consuming. Moreover, methodological decisions within the modelling exercises to generate projections and characterise uncertainty inevitably prioritise accuracy in some variables over others (Parker 2014; Parker and Winsberg 2018), and within a variable, some value ranges over others (Garner and Keller 2018). *A priori* modelling choices may not line up well with what turns out to matter most to the particular decision at hand.

Another issue is that the uncertainties associated with fine-grained, regional climate change projections are generally large and difficult to quantify. If decision makers focus on devising acts to perform well in the most likely futures, and if uncertainties about those futures are underrepresented (the projections turn out to be overconfident), then decision makers will end up evaluating an inappropriately narrow set of acts.

Within the climate change adaptation literature, a shift towards *bottom-up* approaches (Lempert et al. 2004, 2006; Lempert and Collins 2007; Dessai and Sluijs 2007; Dessai et al. 2009; Wilby and Dessai 2010; Brown and Wilby 2012) has begun to complement the default *top-down* approach. A bottom-up approach (also called *vulnerability-first*, *policy-first*, or *assess-risk-of-policy*) structures the problem starting from an understanding of the specific decision context rather than the broader climatic conditions under which the chosen action will ultimately play out.

Specifically, it begins with a well defined *objective* and a concrete *policy option* already on the table. The objective is defined in terms of a measure of policy performance and a critical threshold that draws a line between adequate and inadequate performance on that measure. Starting from this objective, one works backwards to work out the climatic (or other external) conditions under which the stated performance threshold would be reached, were the policy in question implemented.

Brown et al. (2012) illustrate the approach with a stylised example based on the reservoir system supplying water to the metropolitan Boston area. The decision concerns whether and how to expand this reservoir system in light of the changing climate, and the initial policy option is simply retaining the existing water supply infrastructure *as is* (i.e., no action). In this illustration, the performance metric is the system's *reliability*, understood as the chance of fully satisfying water demand in any given year, and the critical threshold is fixed at 95% reliability (see Brown et al. (2012) for details).

Brown et al. (2012) develop a model of system reliability as a function of annual mean precipitation and temperature (higher temperatures mean greater loss of water into the

atmosphere through evaporation and plant transpiration). Working backwards from the desired performance threshold, they then uncover the climatic conditions that lead to inadequate policy results. The shaded area in Figure 1 shows the conditions under which system reliability falls below 95%. (In the non-shaded region, increased precipitation more than compensates for water loss resulting from higher temperatures.)

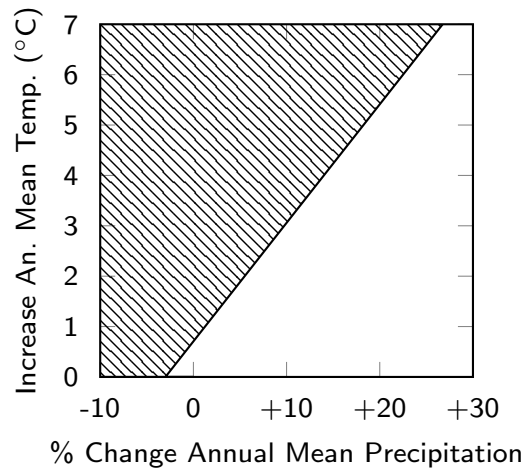


Figure 1: A continuous 2D state space for a municipal water supply management decision. The diagonal line partitions the space, with the shaded region indicating climatic conditions under which the existing water supply system would fail to meet a performance goal of 95% reliability. (Based on Brown et al. 2012, 9)

In decision-theoretic terms, Figure 1 illustrates a *partition of the state space*. It is a very simple, coarse-grained partition (only two states), but it is useful for the decision problem at hand, and tailor-made based on decision makers' objectives (the performance threshold). Labeling the two states (shaded and non-shaded) 'Too Dry' and 'Wet Enough', Table 2 illustrates the role of the partition in structuring the decision problem, as seen through the decision matrix.

		Too Dry	Wet Enough
act 1		-	+
act 2			
⋮			

Table 2: An illustrative decision matrix featuring the state space partition resulting from a bottom-up vulnerability analysis in Brown et al. (2012).

On this bottom-up approach, it is only this coarse-grained partition over which climate

services providers, or other science advisors, must venture to assign probabilities (or some fuzzier assessment of chances). The approach thus reduces the decision process' demands for sharp and specific probability information—which may be of suspect quality or trustworthiness in contexts such as local climate change adaptation. As Brown et al. (2012, 52) explain, 'The strategy rests on using the insights from a vulnerability analysis to inform the selection and processing of the climate model information—to tailor the choice and use of model outputs to maximize their credibility and utility in the assessment.'

Read as decision-structuring advice, we might see bottom-up vulnerability analyses as offering a strategy for managing and balancing various uncertainties through smart partitioning. The credibility of probabilistic climate information overlaid on the state space will generally be improved by tailoring the scientific assessment to best distinguish between specific decision-relevant states—and the simpler the partition, the better. But simple partitions also have a downside: they lead to *heterogeneity within consequences* (the $c_{i,j}$ of Table 1). No single reliability number can be assigned to the status-quo policy under the state of the world 'Too Dry' because that state encompasses a wide variety of conditions, which lead to *different* reliabilities (though all are below 95%). More credible climate inputs thus come at the cost of greater uncertainty in how to evaluate (e.g., in utility or monetary terms) the miscellaneous baskets of outcomes now associated with each act under each state of the world.

Bottom-up approaches ride this trade-off all the way down, buying all the climate credibility they can get, and paying the price in lost capacity to discriminate between outcomes.⁴ But they minimise the pain by asking decision makers to *prioritise* where, along the scale of the policy performance metric, they wish to retain the capacity to discriminate. When decision makers declare the performance threshold around which the partitioning exercise is based, they are saying 'What matters most to me is whether the policy will do at least *this* well.' The procedure then delivers the simplest partition under which consequences can still be categorised as meeting or failing to meet that goal (but note that this applies only for the particular policy option from which the analysis began—the partition may be less appropriate for assessing alternative policies).

While the main purpose of this and subsequent subsections is simply to illustrate the attention given to the structuring task within the decision support literature, the larger aim of facilitating engagement with that literature will be furthered by a few (selective, and far too brief) comments noting elementary points of contact between each method and ideas about decision structuring familiar within the philosophical literature.

By deliberately inviting decision-relevant heterogeneity into the decision consequences, bottom-up partitioning as described above disregards Savage's notion of 'small world' decision making (Savage, 1954, 82–91; Joyce, 1999, §2.6), in which the consequence of an act under a state should be fully resolved, leaving nothing that is relevant to the valuation of that outcome unspecified. Coarsening states to the point of introducing such heterogeneity can be

⁴Bottom-up approaches to partitioning, and the robustness-based decision support frameworks in which they are typically embedded, also have wider consequences for modelling strategies and distribution of resources within climate change modelling intended to inform decision-making (see Dessai and Hulme 2004; Dessai and Sluijs 2007; Weaver et al. 2013).

thought of as relocating uncertainty from the states to the consequences; Bradley and Drechsler (2014) discuss this and other such accounting manoeuvres, with associated trade-offs, in more careful decision-theoretic terms.

Setting aside the extreme simplicity of the partition seen above (just two states), the process for drawing the line between them also defies decision theory orthodoxy. Bottom-up framing offers a quantitative partitioning tool driven directly by the decision maker's aims (the performance threshold). The mechanism that translates the performance threshold into a state space partition is the initial policy option, and the resulting partition is, in this way, policy-relative. In normative discourse in and around decision theory, it generally goes without saying that no *rational* decision process would prejudge a matter by framing the decision in a way that privileges one act (e.g., the status quo) over others. Bottom-up framing displays a willingness to compromise on this desideratum. (A bit more on this in §3.3.)

While the Brown et al. (2012) example concerns partitioning a space of fixed dimensions rather than adding and subtracting new dimensions to that space, the rationale for bottom-up, threshold-based partitioning explored above shares some similarities with Bradley's (2017, §1.2) discussion of choosing which and how many factors (dimensions) to include through balancing the 'quality' gained by including more factors against the 'efficiency' lost through the time and effort taken to make associated additional—and perhaps more difficult—judgements.

The method discussed next simultaneously addresses *both* the inclusion or exclusion of factors in a state space *and* the number and placement of the partitioning lines drawn within that space.

3.2 Scenario discovery

A technique called *scenario discovery* (Lempert et al. 2006; Groves and Lempert 2007; Bryant and Lempert 2010; Hall et al. 2012; Lempert 2013) builds on the bottom-up approach, giving additional advice on partitioning the state space in contexts where a large number of external factors interact, sometimes in complex ways, to influence how a policy choice will play out. Such contexts present two complications for bottom-up analyses. First, complex interactions make it unlikely that any simple border (like the straight line in Figure 1) will separate the combinations of external factors that lead to adequate policy outcomes from those that don't. Instead, pockets of good and bad policy performance may be scattered across the many dimensions of the state space. Second, applicable computational models are often *simulations* that cannot be solved analytically, meaning the partition line(s) can be found only by running the model at a finite number of points, observing the results, and interpolating where the line might be.

Scenario discovery addresses the computational issue by sampling points evenly across the space of external factors (as densely as computational constraints will allow), then calculating (simulated) policy outcomes at these points—providing an imperfect but useful view of policy

performance across that space.⁵ To address complexity, a data-mining algorithm is applied that seeks to find a small number of relatively simple borders that separate, for the most part, the combinations of external factors leading to inadequate performance from those leading to adequate performance. The resulting regions are the ‘scenarios’, and each scenario becomes a state in the partition.⁶

One software instantiation of scenario discovery (Bryant and Lempert 2010) seeks box-shaped regions that can be described using only a small number of the original dimensions, but which contain a large proportion of the bad policy outcomes and few of the good ones. To achieve this, the algorithm jointly maximises what its authors call *interpretability*, *density*, and *coverage*. Simply-defined boxes (and few of them) are easier for human decision makers to get their heads around (*interpretability*). And the more homogeneous the boxes (*density*), the better the simple partition will reflect the more complex underlying picture of policy performance across the space of external factors. Added together, the boxes should cover a large proportion of the total number of bad outcomes (*coverage*) so that few of these are left floating in the space between boxes. The three desiderata typically compete, so a choice must be made by the user regarding how much of one at the expense of the others.

Groves and Lempert (2007) illustrate scenario discovery with an application to water supply management in the South Coast hydrological region of California. For simplicity, the authors focus on uncertainties in *urban sector water demand*, holding fixed other considerations (including climate). They use a model of urban sector water demand that includes 16 uncertain inputs, including demographic, behavioural, and economic variables affecting water usage. These inputs form a 16-dimensional state space to be partitioned through scenario discovery.

The authors consider a fixed menu of 24 management strategies (acts) consisting of different combinations of new supply projects and improvements to efficiency. One of these acts—the best option based on California Department of Water Resources demand projections—is privileged as the *base case strategy*. Policies were evaluated using an aggregate cost measure on simulations run from the year 2000 to 2030. In contrast to the previous example, the performance threshold used in this application was comparative rather than absolute: performance of the base case strategy at each sampled point in the state space was judged a success if no other strategy outperformed it (on the cost measure) by more than a chosen amount. (In other words, a *regret* threshold was used.)

The data-mining algorithm then searches for boxes encompassing clusters of such policy failure. From among the options presented by the algorithm (which differ in the weight given to interpretability, density, and coverage) the authors chose the two boxes/scenarios shown in

⁵Bryant and Lempert (2010) use Latin Hypercube sampling within limits given by expert judgement of upper and lower bounds for each dimension in the space.

⁶Here, a *scenario* is a precisely defined *region* within the space of external conditions. Other uses of ‘scenario’ in the decision support literature understand scenarios as *points* within that space (e.g., Schwartz, 1996; Rounsevell and Metzger, 2010; Carlsen et al., 2013, 2016a, 2016b). Point-scenarios do not partition a space, though they can be understood to structure the decision in a looser sense that is not captured by the notion of structuring used in this paper (establishing the state-consequence matrix).

Figure 2 as the best overall representation of the conditions under which the base case policy fails to perform adequately. (In the scenario labeled ‘Soft Landing’, the policy over-invests, resulting in excess supply; the opposite happens under ‘Rapid Growth’.) The two boxes constrain only 3 of the original 16 dimensions (interpretability) and together encompass 75% of the simulations resulting in policy failure (coverage); their densities are 63% and 61% respectively.

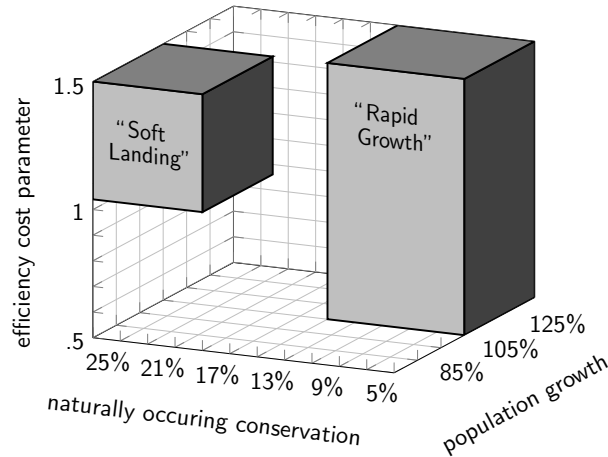


Figure 2: Scenarios generated through application of the scenario discovery method to a stylised water supply management problem (based on Groves and Lempert 2007, 80). Boxes indicate combinations of external factors that lead to policy failure.

These two boxes then become states in a partition, along with a third state for the leftover catch-all region in which the base case policy is largely successful (Table 3). Like the simpler case of bottom-up partitioning describe above (§3.1), it is only through this partition that probabilistic information is then applied to the decision problem (see Figure 5 in Groves and Lempert, 2007; also Figure 8 in Hall et al., 2012).

		Soft Landing	Rapid Growth	all else
base case	-	-	+	
strategy 2				
strategy 3				
⋮				

Table 3: A decision matrix with state space partition produced through scenario discovery (based on Groves and Lempert 2007, 82). Symbols ‘-’ and ‘+’ indicate states in which the base case strategy largely fails/meets the stated performance threshold.

As a variation on the general strategy of bottom-up decision structuring, much of the closing discussion from §3.1 applies here to scenario discovery as well. But the complications presented by complex interactions and computational constraints also bring new considerations to the discussion of rationale.

We saw in §3.1 how simple partitions create broad-brush consequences that obscure decision-relevant heterogeneity (such as the difference between 90% and 30% reliability). Consequences under a scenario-discovered partition are even harder to get a grip on, as they encompass variation not only *within* one of the two value categories supplied by decision makers (either above or below the threshold), but even *across* that threshold. For example, while the ‘Rapid Growth’ state in Figure 2 is to be interpreted as a state of the world under which the base case policy *fails* to meet decision makers’ performance goal, its density is only 61%, meaning that 39% of the simulations run at points within that state *satisfied* that policy goal. And the combined coverage of the two scenarios is only 75%, meaning that 25% of the simulated policy failures happened within the leftover catch-all state that is understood to characterise conditions under which the base case policy *succeeds*.

Why accept such a representation of the problem? These compromises result from projecting the state space into just 3 of its original 16 dimensions, and drawing only a few simple borders within that reduced space; together these choices make for a more *interpretable* partition. The rationale for valuing interpretability appeals to a number of *cognitive benefits* for decision makers (Bryant and Lempert 2010, 35, 44).

The density and coverage numbers are, moreover, only *estimates* of the true (according to the model) coverage and densities, based on what may be a relatively sparse sampling of the state space (in the CA water demand example, 500 points are spread over 16 dimensions). This makes for additional uncertainty about exactly what outcomes are lumped under a given consequence within the decision matrix. And while this particular uncertainty can be reduced by applying more computing power, the computational demands can be significant and limiting (Bhave et al. 2016). Just how great those demands are (or on a fixed computing budget, how densely the state space can be sampled) of course depends on the choice of system model to be used in the analysis. So from a bigger-picture perspective, the compromises struck within an application of scenario discovery are also linked—through the mediating currency of computational cost—to an array of pros and cons associated with simpler versus more complex models, and their influence on interactions between scientists, analysts, decision makers, and stake-holders (Vezér et al. 2018).

Turning back to flagging points of contact with the philosophical literature, scenario discovery—and bottom-up approaches in general—can be contrasted with the idea, expressed, e.g., by Joyce (1999, 70–73), of the deliberation involved in structuring decisions as a process of *refinement* wherein each step of deliberation further subdivides the state space (and thus the consequences). Scenario discovery proceeds in the opposite direction, moving from a continuous, high-dimensional proto-state space to a simple and discrete, low-dimensional partition, deliberately disregarding detail that has been explicitly considered.

Discrete partitions can differ in level of detail either through finer/coarser division of a fixed

number of dimensions, or by including/excluding whole factors (dimensions). The §3.1 example of bottom-up partitioning included only the former, while scenario discovery also addresses the latter. The advice that scenario discovery gives can be contrasted with the suggestion of Bradley and Steele (2015) and Bradley (2017, §1.2) that severity of uncertainty should be a consideration in choosing which factors to explicitly represent within the state space (they say *the more uncertainty* about the true value of a factor, *the lower* its priority for representation in the state space). Narrowly construed, scenario discovery prioritises dimensions from its proto-state space with no regard for severity of uncertainty; all that matters is whether a given dimension is used within the best overall representation of conditions separating failure from success for the base case policy.

But taking into account the broader context within which scenario discovery operates, the bigger-picture procedure does pay heed to something like severity of uncertainty. Bounds on each dimension are set by expert judgement, where more uncertainty about the true value may translate to wider bounds in that dimension. And wider bounds will tend to make the variation explored along that dimension by the sampling procedure more important to explaining variance in policy success. Factors that are more uncertain are, therefore, somewhat *more likely* to end up represented explicitly in the recommended state space partition—the opposite of Bradley and Steele’s (2015) advice.

Whether divergent structuring recommendations lead to contradictory advice about the ultimate decision depends on what goes on in the other parts in the decision process, including, for example, how heterogeneity within a decision consequence is treated in the valuation of that consequence. And if structuring can in some sense be optimised for a particular decision rule (or vice versa), then norms for deciding and structuring are not independent, and one reason for clashing advice on the structuring phase might be contrasting ideas about the *decision rule* to be applied once the problem has been structured.⁷ Another reason for divergent advice on setting up the state space may be that facilitating the immediate evaluation and comparison of acts is not the only purpose of a state-space partition—which brings us to the next topic.

3.3 Iterative restructuring

In the §3.2 illustration of scenario discovery, the menu of acts was fixed from the start. But this needn’t be the case. Indeed, helping to develop new policy options to be included on the menu is an important secondary aim of bottom-up structuring approaches. This is particularly explicit within the broader framework of Robust Decision Making (Lempert et al. 2006; Lempert and Collins 2007; Dessai and Sluijs 2007), where scenario discovery reveals the combinations of conditions under which a given policy fails to meet the decision makers’ performance goal, and consideration of these conditions can then help decision makers to devise improvements or variants on that policy, or lead them to previously unconceived

⁷Bradley and Steele (2015) have in mind a decision model from Walker and Dietz (2011), while §4.1.2 of Hall et al. (2012) gives the flavour of the deciding norms promoted within *Robust Decision Making*—the broader framework within which scenario discovery has been developed and deployed.

alternative actions. So deep uncertainty decision support also includes advice about *constructing policy options* (acts 1–n, Table 1) to be considered and evaluated, marking another contribution to the decision-structuring task.

Editing the menu of acts may in turn motivate re-partitioning of the state space, either because a new policy is highlighted as the base case policy, or because decision makers employ a regret-based performance threshold (as in the §3.2 example), in which case changes to the lineup of alternative policies may redefine the conditions under which the existing base case policy is sufficiently outperformed by some alternative in order to count as a policy failure. Such iterative reformulating and restructuring of decision problems is typical of applied decision analysis in general (Phillips 1984; Clemen and Reilly 2013), though may be even more central to decision support aimed specifically at contexts of deeper uncertainty (compare, e.g., the process flowchart of Weaver et al. (2013, 46) with that of Clemen and Reilly (2013, 9)).

Steps within such iterative restructuring routines fall under what (Peterson 2009, 33–35) calls ‘transformative decision rules’, rules for transforming one’s representation of the decision problem rather than for solving the problem by picking an action.

As noted above, the partitions illustrated in both §3.1 and §3.2 privilege one act (the ‘base case’, or default policy) above the others, carving up the state space to reveal what matters most to the success of this privileged act, and potentially hampering the ability of a decision rule—or a less formal human judgement process—to even-handedly assess the value of the alternatives vis-à-vis the base case. One reply to the charge of bias is that through iterative restructuring of the problem, decision makers are exposed to a number of different representations of the decision problem (with different acts granted this privileged status, and correspondingly different state space partitions), each showing a different and potentially valuable perspective. This multiple-framings reply sits uncomfortably with the strong presumption within decision-theory commentary (e.g., Joyce 1999, 70) that any deliberation over problem structure will eventually settle on a single, definitive representation of the problem before proceeding to assessment.

3.4 Adaptive policymaking

An approach called *Adaptive Policymaking* (Walker et al. 2001; Kwakkel et al. 2010) goes even further to facilitate the policy-generation side of decision structuring. Adaptive Policymaking is a qualitative, step-by-step framework for developing and implementing a particular type of policy option, namely *adaptive policies*, or policies that can evolve and respond to changing conditions as uncertainties begin to resolve over time. Adaptive Policymaking is a form of *iterative risk management* (Morgan et al. 2009) and builds on the idea that some decisions are better served by a strategy of *monitor-then-react* rather than *predict-then-act* (Walker et al. 2013).

The policy-generation phase of Adaptive Policymaking prescribes a series of concrete steps for further developing, and building alternatives to, a promising *basic policy* option (Walker et al.

2001). Analysts first identify the conditions under which the basic policy will succeed, as well as its potential adverse consequences—the latter called *vulnerabilities* (a somewhat different usage from §3.1). For vulnerabilities that are certain to occur, *mitigating actions* can be added to the basic policy to be put in place immediately upon implementation. For those that may or may not materialise, *hedging actions* can be added to reduce exposure or cushion potential impacts.

Kwakkel et al. (2010) illustrate these ideas with an application to strategic planning for Amsterdam's Schiphol Airport, a major aviation hub in the European Union. In the face of uncertain developments in the aviation sector and deteriorating wind conditions due to climate change, airport managers aim to maintain or expand their share of European air traffic while minimizing negative local impacts such as noise and air pollution. Here, a *basic policy* option might be to tweak operations of the existing runway system to increase capacity and reduce noise impacts, while also starting work on an additional runway and terminal.

Regarding vulnerabilities, the new runway will certainly increase noise, so managers should build into the plan *mitigating actions* such as investing in noise insulation or offering financial compensation to affected residents. Uncertain vulnerabilities include the possibility of the national carrier (KLM) shifting its main hub from Schiphol to Paris Charles de Gaulle; managers should *hedge* against this by diversifying carriers.

Policy generation under the Adaptive Policymaking framework continues with the development of a *monitoring scheme* and designation of critical levels—called *triggers*—at which pre-specified contingency plans will be activated. For example, emissions from the fleet serving Schiphol might be monitored, with an emissions increase of more than 10% triggering a rise in landing fees for environmentally unfriendly planes.

By building in the monitoring scheme and contingency plans, as well as the hedging and mitigating actions, a basic policy option is remodeled into a new (adaptive) policy option not previously considered by decision makers, thereby contributing to the decision-structuring task by generating new (and better) policy options (acts 1–n, Table 1) to be considered and evaluated.⁸

The question of where acts come from is largely neglected both in decision theory itself and in commentary around the theory. Consequently, there is little within the more theoretical literature against which to compare or contrast the act-development advice offered by Adaptive Policymaking. The mere fact of treating acts as concrete *courses of action* that a person might plan and enact marks a departure from the decision-theoretic approach to acts as featureless abstractions mapping states of the world to consequences.⁹ Indeed, positing the

⁸See Walker et al. (2001) and Kwakkel et al. (2010) for more details and additional steps in the Adaptive Policymaking framework. Further developments of the basic ideas of Adaptive Policymaking include a computer-assisted approach to adaptive policy design (Kwakkel et al. 2012, 2015) and an expanded, hybrid approach called *Dynamic Adaptive Policy Pathways* (Haasnoot et al. 2013; Kwakkel et al. 2015) that is the result of incorporating elements of Adaptive Policymaking into the *Adaptation Pathways* (Haasnoot et al. 2011, 2012; Ranger et al. 2013) decision support framework.

⁹Case-based decision theory (Gilboa and Schmeidler 2001) is an exception to this and many other generalisations about decision theory.

availability of all logically possible mappings (Savage 1954) makes the formulation of a list of acts an automatic side-effect of other structuring choices. For some theoreticians, this may tend to obscure the reality that devising a menu of acts can be an important part of structuring real decisions. In practice, after all, ‘I cannot choose it if I do not think of it’ (Baron 2008, 62).

One tenuous point of contact in the philosophical literature is Hansson’s (1996; 2005) discussion of what he calls *demarcation uncertainty*, one aspect of which is the *unfinished list* of acts. But Hansson’s suggestions address a higher-level decision about how to proceed in the face of a stubbornly unfinished list of options rather than how to go about improving the list; they include going with the options you have, and postponing the decision while looking for more—perhaps implementing a temporary fix in the meantime.

3.5 Multi-objective robust optimisation

While Adaptive Policymaking begins from a single policy and builds out constructively, another approach is to cast a wide net over a broad class of policy options and then narrow down from there. A common approach to this second sort of problem—widely practised in engineering and operational research—is *optimisation*. To narrow down through optimisation, a *solution space* is described precisely in mathematical terms, and an *objective function* is defined that evaluates performance at each point in the solution space. Optimisation then consists in finding the point(s) in solution space that maximise the objective function. To map this language onto a decision problem, let the solution space be a space of possible acts and the objective function a formula by which those acts will be evaluated. For example, the solution space might be all possible allocations of a fixed sum of money across a number of investments and the objective function might be the expected return on investment.

While the term ‘optimisation’ suggests a unique solution, this need not be the case, in particular where there are multiple objectives. Within decision theory, a standard approach to multiple objectives (*multi-criteria decision theory*) expresses each objective in a common currency (such as utility)—or otherwise weights the objectives relative to one another—thereby combining them into a single all-inclusive score. But if multiple objectives are instead kept separate and incommensurable (an approach sometimes called *a posteriori decision support*, or *generate-first-choose-later*; see Herman et al. 2015, and references therein), then solutions can be only partially ranked. In this case, the goal of optimisation is to identify the *Pareto optimal* set of solutions: those that are not dominated (worse on every objective) by any other solution (Figure 3). Finding the Pareto optimal set within a solution space produces a set of acts for further consideration and evaluation—one of the tasks required to structure a decision.

Where individual objectives quantify concrete consequences of a decision, identifying the Pareto set requires knowing the true state of the world. Where objectives are expectations, *probabilities* must be assigned to states. *Robust optimisation* instead defines the objectives as performance measures across a specified range of possible futures (Gabrel et al. 2014), often




Figure 3: An illustrative Pareto optimal set of acts in a multi-objective decision problem. No act is beaten by another on every objective.

with no appeal to probabilities.

For example, Quinn et al. (2017) seek to optimise the operating procedures for a group of four interconnected reservoirs on the Red River basin in Vietnam. Here, an operating procedure is a mathematical formula describing how much water is released by each reservoir on any given day, based on the current water levels in each reservoir, the level of water flowing into the system, and the day of the year. The combinatorics of four reservoirs, six inputs, and a dozen candidate functional forms means that specifying a single operating procedure requires fixing each of 176 decision variables.

Within this vast space of possible procedures (the acts between which managers must choose), Quinn et al. (2017) identify the Pareto set with respect to three competing objectives: hydropower production, agricultural water supply, and flood protection for the city of Hanoi—which is downstream from the reservoirs. Each of these three objectives is quantified as a performance measure across a large ensemble of simulated water inflows designed to represent a range of conditions that might be experienced. On one formulation, each metric was evaluated as the first percentile performance across 1000 simulations; in other words, each operating procedure was judged only as good as its tenth-worst performance (on hydropower production, agricultural water supply, and flood protection) in 1000 simulated trials, none of which is assumed more probable than any other.

The Pareto set to emerge of course depends on methodological choices made in the process of completing such an exercise, including delineation and structuring of the space of potential acts, and the choice of objectives to be maximised. Indeed, a central point of the Quinn et al. (2017) study is to demonstrate that even subtly different formulations of the optimisation problem can yield significantly different Pareto sets.

In one form or another, multi-objective robust optimisation is used to identify Pareto optimal sets of policy options—and thus to structure subsequent stages of decision making—within Robust Decision Making (Lempert et al. 2006), Dynamic Adaptive Policy Pathways (Kwakkel et al. 2015), Many-Objective Robust Decision Making (MORDM, Kasprzyk et al. 2013; Hadka et al. 2015), Direct Policy Search (Quinn et al. 2017), and other decision support frameworks (Herman et al. 2015).

Dominance is, of course, a familiar idea within decision theory, as is the advice to narrow

down one's options by removing dominated acts. Note, however, that the implementation of dominance reasoning in the Quinn et al. (2017) example differs from state dominance (equal or better utility in every state of the world), stochastic dominance (equal or better probability for every utility minimum), and multi-objective dominance *under certainty* (equal or better on each of several ways to value a sure outcome). Each of Quinn et al.'s objectives is akin to a separate decision rule that integrates and weighs up both belief-like information, in the form of the ensemble of simulations, and value-relevant information, e.g., hydropower performance, in a way that dominance reasoning generally avoids.

From a decision-theoretic perspective, decision structuring is merely a precursor to the main event of imposing rationality constraints on the choice problem. But while multi-objective robust optimisation performs a task central to decision structuring (identification of the acts), this is not to say that the resulting menu of acts is typically fed to a decision-theoretic model for further evaluation. Subsequent evaluation of acts may consist of informal deliberation on what constitutes the best trade-off between the objectives used to define the Pareto set.

On the other hand, multi-objective robust optimisation can be viewed from another perspective by labelling different parts of the procedure as the structuring and the choosing. The initial parameterization of the full space of potential acts, as well as the generation of the ensemble of simulated external conditions, might already be viewed as structuring, while the subsequent dominance reasoning with respect to the three objective functions constitutes choosing. Yet a decision maker who gets only this far on the problem may still be very far from having made a choice.

More broadly, structuring that is not intended to prepare the ground for a formal decision model raises questions about the compatibility of methods for structuring and for choosing that are drawn from different literatures or research traditions. Different methods may draw the line between the two tasks in different places, or disagree about which bits of the overall decision process should be targeted with tools and theories, and which are best left to less structured deliberation. Deeper incompatibilities might result from fundamentally different underlying visions of the decision making process (see, e.g., Zeleny, 1989; Hansson, 2005, §3) or of the proper role of analysts and models within that process (see, e.g., Tsoukiàs 2008).

4 Conclusion

There is broad recognition, across many areas of science and policy, of a need for better normative guidance for decision making in situations of severe uncertainty. As the formalisms, suppositions, and disciplinary practices of decision theory are stretched to acknowledge and address this need, there is, I suggest, an important point deserving of more attention: The deeper the uncertainty, the more of what is most difficult about the decision problem is in the structuring. At any rate, this is a conclusion that might be drawn from observing where the greatest efforts are being made within the more applied branch of deep uncertainty decision research referred to here as 'deep uncertainty decision support'.

Decision structuring can be done well or done poorly, and it is possible to give sound advice about how to do it better. Greater appreciation that this is what (among other things) the applied methods offer may help clarify the contributions that these methods make to the broader study of managing uncertainty and making good decisions. Normative theorising about decision making that aims for relevance in contexts of deep uncertainty might look to the decision support literature for ideas and salutary challenges on the topic of decision structuring, and as a counterbalancing influence to decision theory's comparative focus on the choice task.

The rationales given or implied in the design and evaluation of deep uncertainty decision support tools draw on a fluid and sometimes un-differentiated mix of considerations, including intuitive standards of rationality or reasonableness, but also tractability and computational cost of the required modelling and analysis, domain-specific considerations of the relevant scientific modelling and prediction, the cost-effectiveness of a decision analyst's time, and the comfort, affect, and cognitive limitations of decision makers. There is room for more rigorous interrogation of these rationales and of the normative foundation behind them, including the navigation of trade-offs made among competing considerations, in both the structuring and the choosing tasks, and in the integration of the two.

References

- American Meteorological Society (2015). A policy statement of the American Meteorological Society: Climate services. Technical report, The AMS Council, <https://www2.ametsoc.org/ams/index.cfm/about-ams/ams-statements/statementsof-the-ams-in-force/climate-services1/>.
- Baron, J. (2008). *Thinking and deciding* (4th ed.). Cambridge University Press.
- Betz, G. (2013). In defence of the value free ideal. *European Journal for Philosophy of Science* 3(2), 207–220.
- Bhave, A. G., D. Conway, S. Dessai, and D. A. Stainforth (2016). Barriers and opportunities for robust decision making approaches to support climate change adaptation in the developing world. *Climate Risk Management* 14, 1–10.
- Binmore, K. (2009). *Rational Decisions*. Princeton University Press.
- Binmore, K. (2015). A minimal extension of bayesian decision theory. *Theory and Decision*, 1–22.
- Bradley, R. (2017). *Decision Theory With a Human Face*. Cambridge University Press.
- Bradley, R. and M. Drechsler (2014). Types of uncertainty. *Erkenntnis* 79(6), 1225–1248.
- Bradley, R. and K. Steele (2015). Making climate decisions. *Philosophy Compass* 10(11), 799–810.
- Brown, C., Y. Ghile, M. Laverty, and K. Li (2012). Decision scaling: Linking bottom-up vulnerability analysis with climate projections in the water sector. *Water Resources Research* 48(9).
- Brown, C. and R. L. Wilby (2012). An alternate approach to assessing climate risks. *Eos, Transactions American Geophysical Union* 93(41), 401–402.

- 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(1), 34–49.
- Carlsen, H., K. H. Dreborg, and P. Wikman-Svahn (2013). Tailor-made scenario planning for local adaptation to climate change. *Mitigation and Adaptation Strategies for Global Change* 18(8), 1239–1255.
- Carlsen, H., E. A. Eriksson, K. H. Dreborg, B. Johansson, and Ö. Bodin (2016). Systematic exploration of scenario spaces. *Foresight* 18(1), 59–75.
- Carlsen, H., R. Lempert, P. Wikman-Svahn, and V. Schweizer (2016). Choosing small sets of policy-relevant scenarios by combining vulnerability and diversity approaches. *Environmental Modelling & Software* 84, 155–164.
- Chateauneuf, A. and J. H. Faro (2009). Ambiguity through confidence functions. *Journal of Mathematical Economics* 45(9), 535–558.
- Clemen, R. T. and T. Reilly (2013). *Making hard decisions with DecisionTools*. Cengage Learning.
- Dessai, S. and M. Hulme (2004). Does climate adaptation policy need probabilities? *Climate policy* 4(2), 107–128.
- Dessai, S., M. Hulme, R. Lempert, and R. Pielke (2009). Do we need better predictions to adapt to a changing climate? *EOS, Transactions American Geophysical Union* 90(13), 111–112.
- Dessai, S. and J. P. Sluijs (2007). *Uncertainty and climate change adaptation: A scoping study*. Copernicus Institute for Sustainable Development and Innovation, Department of Science Technology and Society.
- Douglas, H. (2009). *Science, policy, and the value-free ideal*. University of Pittsburgh Press.
- Douglas, H. (2016). Values in science. In P. Humphreys (Ed.), *The Oxford Handbook of Philosophy of Science*, pp. 609–630. Oxford University Press.
- Fishburn, P. C. (1964). *Decision and value theory*. Wiley.
- Gabrel, V., C. Murat, and A. Thiele (2014). Recent advances in robust optimization: An overview. *European journal of operational research* 235(3), 471–483.
- Galaabaatar, T. and E. Karni (2012). Expected multi-utility representations. *Mathematical Social Sciences* 64(3), 242–246.
- Galaabaatar, T. and E. Karni (2013). Subjective expected utility with incomplete preferences. *Econometrica* 81(1), 255–284.
- Gärdenfors, P. and N.-E. Sahlin (1982). Unreliable probabilities, risk taking, and decision making. *Synthese* 53(3), 361–386.
- Garner, G. and K. Keller (2018). When tails wag the decision: The role of distributional tails on climate impacts on decision-relevant time-scales. *manuscript*.
- Ghirardato, P., F. Maccheroni, and M. Marinacci (2004). Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory* 118(2), 133–173.
- Gilboa, I. and M. Marinacci (2013). Ambiguity and the Bayesian paradigm. In *Advances in Economics and Econometrics, Tenth World Congress*, Volume 1.

- Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of mathematical economics* 18(2), 141–153.
- Gilboa, I. and D. Schmeidler (2001). *A theory of case-based decisions*. Cambridge University Press.
- Groves, D. G. and R. J. Lempert (2007). A new analytic method for finding policy-relevant scenarios. *Global Environmental Change* 17(1), 73–85.
- Haasnoot, M., J. H. Kwakkel, W. E. Walker, and J. ter Maat (2013). Dynamic adaptive policy pathways: A method for crafting robust decisions for a deeply uncertain world. *Global environmental change* 23(2), 485–498.
- Haasnoot, M., H. Middelkoop, A. Offermans, E. Van Beek, and W. P. Van Deursen (2012). Exploring pathways for sustainable water management in river deltas in a changing environment. *Climatic Change* 115(3-4), 795–819.
- Haasnoot, M., H. Middelkoop, E. Van Beek, and W. Van Deursen (2011). A method to develop sustainable water management strategies for an uncertain future. *Sustainable Development* 19(6), 369–381.
- Hadka, D., J. Herman, P. Reed, and K. Keller (2015). An open source framework for many-objective robust decision making. *Environmental Modelling & Software* 74, 114–129.
- Hall, J. W., R. J. Lempert, K. Keller, A. Hackbarth, C. Mijere, and D. J. McInerney (2012). Robust climate policies under uncertainty: A comparison of robust decision making and info-gap methods. *Risk Analysis* 32(10), 1657–1672.
- Hansson, S. (2005). *Decision theory: A brief introduction*.
- Hansson, S. O. (1996). Decision making under great uncertainty. *Philosophy of the social sciences* 26(3), 369–386.
- Herman, J. D., P. M. Reed, H. B. Zeff, and G. W. Characklis (2015). How should robustness be defined for water systems planning under change? *Journal of Water Resources Planning and Management* 141(10).
- Herman, J. D., H. B. Zeff, J. R. Lamontagne, P. M. Reed, and G. W. Characklis (2016). Synthetic drought scenario generation to support bottom-up water supply vulnerability assessments. *Journal of Water Resources Planning and Management* 142(11).
- Hewitt, C., S. Mason, and D. Walland (2012). The global framework for climate services. *Nature Climate Change* 2(12), 831.
- Hill, B. (2013). Confidence and decision. *Games and Economic Behavior* 82, 675–692.
- Hill, B. (2016). Incomplete preferences and confidence. *Journal of Mathematical Economics* 65, 83–103.
- Joyce, J. M. (1999). *The foundations of causal decision theory*. Cambridge University Press.
- Karni, E. and M.-L. Vierø (2013). “Reverse Bayesianism”: A choice-based theory of growing awareness. *The American Economic Review* 103(7), 2790–2810.
- Karni, E. and M.-L. Vierø (2014). Awareness of unawareness: A theory of decision making in the face of ignorance. Technical report, Queen’s Economics Department Working Paper.

- Kasprzyk, J. R., S. Nataraj, P. M. Reed, and R. J. Lempert (2013). Many objective robust decision making for complex environmental systems undergoing change. *Environmental Modelling & Software* 42, 55–71.
- Kasprzyk, J. R., P. M. Reed, G. W. Characklis, and B. R. Kirsch (2012). Many-objective de novo water supply portfolio planning under deep uncertainty. *Environmental Modelling & Software* 34, 87–104.
- Klibanoff, P., M. Marinacci, and S. Mukerji (2005). A smooth model of decision making under ambiguity. *Econometrica* 73(6), 1849–1892.
- Kwakkel, J., M. Haasnoot, and W. Walker (2012). *Computer assisted dynamic adaptive policy design for sustainable water management in river deltas in a changing environment*. International Environmental Modelling and Software Society.
- Kwakkel, J. H., M. Haasnoot, and W. E. Walker (2015). Developing dynamic adaptive policy pathways: a computer-assisted approach for developing adaptive strategies for a deeply uncertain world. *Climatic Change* 132(3), 373–386.
- Kwakkel, J. H., W. E. Walker, and V. Marchau (2010). Adaptive airport strategic planning. *European Journal of Transport and Infrastructure Research (EJTIR)*, 10 (3), 2010.
- Lempert, R. (2013). Scenarios that illuminate vulnerabilities and robust responses. *Climatic Change* 117(4), 627–646.
- Lempert, R., N. Nakicenovic, D. Sarewitz, and M. Schlesinger (2004). Characterizing climate-change uncertainties for decision-makers. *Climatic Change* 65(1), 1–9.
- Lempert, R. J. and M. T. Collins (2007). Managing the risk of uncertain threshold responses: comparison of robust, optimum, and precautionary approaches. *Risk analysis* 27(4), 1009–1026.
- Lempert, R. J., D. G. Groves, S. W. Popper, and S. C. Bankes (2006). A general, analytic method for generating robust strategies and narrative scenarios. *Management science* 52(4), 514–528.
- Maccheroni, F., M. Marinacci, and A. Rustichini (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica* 74(6), 1447–1498.
- Mitchell, S. D. (2009). *Unsimple truths: Science, complexity, and policy*. University of Chicago Press.
- Morgan, M. G., H. Dowlatabadi, M. Henrion, D. Keith, R. Lempert, S. McBride, M. Small, and T. Wilbanks (contributing authors) (2009). Best practice approaches for characterizing, communicating and incorporating scientific uncertainty in climate decision making. A report by the climate change science program and the subcommittee on global change research, National Oceanic and Atmospheric Administration.
- Parker, W. (2014). Values and uncertainties in climate prediction, revisited. *Studies in History and Philosophy of Science Part A* 46, 24–30.
- Parker, W. S. and E. Winsberg (2018). Values and evidence: How models make a difference. *European Journal for Philosophy of Science* (8), 135–142.
- Peterson, M. (2009). *An introduction to decision theory*. Cambridge University Press.
- Phillips, L. D. (1984). A theory of requisite decision models. *Acta psychologica* 56(1), 29–48.
- Popper, S. W. (2016). What is decision making under deep uncertainty and how does it work? Conference presentation, Annual Meeting of the Society for Decision Making under Deep Uncertainty.

- Quinn, J., P. Reed, M. Giuliani, and A. Castelletti (2017). Rival framings: A framework for discovering how problem formulation uncertainties shape risk management trade-offs in water resources systems. *Water Resources Research* 53(8), 7208–7233.
- Quinn, J. D., P. M. Reed, and K. Keller (2017). Direct policy search for robust multi-objective management of deeply uncertain socio-ecological tipping points. *Environmental Modelling & Software* 92(Supplement C), 125–141.
- Ranger, N., T. Reeder, and J. Lowe (2013). Addressing ‘deep’ uncertainty over long-term climate in major infrastructure projects: Four innovations of the Thames Estuary 2100 project. *EURO Journal on Decision Processes* 1(3-4), 233–262.
- Rounsevell, M. D. and M. J. Metzger (2010). Developing qualitative scenario storylines for environmental change assessment. *Wiley Interdisciplinary Reviews: Climate Change* 1(4), 606–619.
- Savage, L. J. (1954). *The Foundations of Statistics*. New York: John Wiley & Sons.
- Schwartz, P. (1996). *The Art of the Long View: Planning in an uncertain world*. Currency-Doubleday, New York.
- Sprenger, J. (2012). Environmental risk analysis: Robustness is essential for precaution. *Philosophy of Science* 79(5), 881–892.
- Steele, K. (2012). The scientist qua policy advisor makes value judgments. *Philosophy of Science* 79(5), 893–904.
- Tsoukiàs, A. (2008). From decision theory to decision aiding methodology. *European Journal of Operational Research* 187(1), 138–161.
- Vezér, M., A. Bakker, K. Keller, and N. Tuana (2018). Epistemic and ethical trade-offs in decision analytical modelling. *Climatic Change* 147(1), 1–10.
- Walker, O. and S. Dietz (2011). A representation result for choice under conscious unawareness. Technical Report 59, Grantham Research Institute on Climate Change and the Environment Working Paper.
- Walker, W. E., M. Haasnoot, and J. H. Kwakkel (2013). Adapt or perish: A review of planning approaches for adaptation under deep uncertainty. *Sustainability* 5(3), 955–979.
- Walker, W. E., R. J. Lempert, and J. H. Kwakkel (2013). Deep uncertainty. In *Encyclopedia of operations research and management science*, pp. 395–402. Springer.
- Walker, W. E., S. A. Rahman, and J. Cave (2001). Adaptive policies, policy analysis, and policy-making. *European journal of operational Research* 128(2), 282–289.
- Weaver, C. P., R. J. Lempert, C. Brown, J. A. Hall, D. Revell, and D. Sarewitz (2013). Improving the contribution of climate model information to decision making: The value and demands of robust decision frameworks. *Wiley Interdisciplinary Reviews: Climate Change* 4(1), 39–60.
- Wilby, R. L. and S. Dessai (2010). Robust adaptation to climate change. *Weather* 65(7), 180–185.
- Zeleny, M. (1989). Cognitive equilibrium: A new paradigm of decision making? *Human Systems Management* 8(3), 185–188.