Delft University of Technology

Supporting DMDU

A Taxonomy of Approaches and Tools

Kwakkel, Jan H.; Haasnoot, Marjolijn

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Chapter 15
Supporting DMDU: A Taxonomy of Approaches and Tools

Jan H. Kwakkel and Marjolijn Haasnoot

Abstract

- A wide variety of tools and approaches for supporting the making of decisions under deep uncertainty have been put forward, but we lack a comparative overview.
- This chapter presents a taxonomy of approaches and tools for supporting decisionmaking under deep uncertainty.
- The taxonomy is based on a decomposition of the tools and approaches into a set of common building blocks.
- Analysts can use the taxonomy for designing context-specific approaches to support DMDU.

15.1 Introduction

Over the last decade, various researchers have put forward approaches for supporting decisionmaking under deep uncertainty (DMDU). For example, Lempert et al. (2006) put forward Robust Decision Making. This was later expanded by Kasprzyk et al. (2013) into Many-objective Robust Decision Making. Other researchers put forward adaptive policymaking (Walker et al. 2001; Kwakkel et al. 2010a) and adaptation pathways (Haasnoot et al. 2012), which were subsequently combined into Dynamic Adaptive Policy Pathways (Haasnoot et al. 2013). At present, a variety of approaches and tools are available, but there is little insight into how the approaches and tools are similar or different, where they overlap, and how they might be meaningfully combined in offering decision support in a specific context. The aim of this chapter is to offer some thoughts on these questions.

J. H. Kwakkel (✉)
Delft University of Technology, Delft, The Netherlands
e-mail: J.H.Kwakkel@tudelft.nl

M. Haasnoot
Water Resources and Delta Management, Deltares, Delft, The Netherlands

Faculty of Physical Geography, Utrecht University, Utrecht, The Netherlands

To provide a tentative, more synthetic view of the DMDU field, this chapter will first discuss the key ideas that underpin the field, followed by a proposed taxonomy. These key ideas are an attempt at articulating the presuppositions that emanate from much of the DMDU literature. They have emerged out of discussions over the last few years with various people. The taxonomy builds on an earlier taxonomy put forward by Herman et al. (2015). It decomposes each approach into its key building blocks. This exercise is particularly useful, for it enables one to move beyond a single approach. By focusing on the building blocks that are available, analysts can start to compose context-specific approaches for supporting DMDU. Moreover, it allows academics to move away from very general statements at the level of comprehensive approaches and instead focus on specific differences and similarities at the level of building blocks.

A challenge one faces when trying to put forward a taxonomy of approaches is that they are subject to interpretation, offer quite some leeway to the user in how they are being used, and are changing over time. A taxonomy pigeonholes approaches and thus runs the risk of failing to do sufficient justice to the intrinsic flexibility that currently exists in practice with respect to DMDU approaches. This risk is somewhat alleviated in the present case, since the focus is on the building blocks that make up the various approaches. By understanding an approach as being composed of building blocks, changes to approaches may be construed as swapping one building block with another one.

The remainder of this chapter is organized as follows. In Sect. 15.2, a broad overview is given of the key ideas that underpin the literature on supporting DMDU. In Sect. 15.3, we present a taxonomy for the various approaches and tools. In Sect. 15.4, we discuss how each of the approaches and tools fit into this taxonomy. Section 15.5 presents our concluding remarks.

15.2 Key Ideas

Decisionmaking with respect to complex systems requires coming to grips with irreducible uncertainty. This uncertainty arises out of intrinsic limits to predictability that occur when dealing with a complex system. Another source of uncertainty is that decisionmaking on complex systems generally involves a variety of stakeholders with different perspectives on what the system is, and what problem one is trying to solve. A third source of uncertainty is that complex systems are subject to dynamic change and are never completely understood.

The intrinsic limits to predictability, the existence of legitimate alternative interpretations of the same data, and the limits to knowability of a system have important implications for decisionmaking. Under the label of 'decisionmaking under deep uncertainty', these are now being explored. Deep uncertainty means that the various parties to a decision do not know or cannot agree on how the system works, how likely various possible future states of the world are, and how important the various outcomes of interest are (Lempert et al. 2003). This suggests that under deep

uncertainty, it is possible to enumerate possible representations of the system, to list plausible futures, and to enumerate relevant outcomes of interest without being able to rank order these representations of the system, the list of plausible futures, or the enumeration of outcomes of interest in terms of likelihood or importance (Kwakkel et al. 2010b).

In the literature, there is an emerging consensus that any decision regarding a complex system should be robust with respect to the various uncertainties. Intuitively, a decision is robust if its expected performance is only weakly affected by the actual future states that emerge as a function of the values actually observed among the various deeply uncertain factors. Various operationalizations of this intuition may be found in the literature (Herman et al. 2015; Giuliani and Castelletti 2016; Kwakkel et al. 2016a; McPhail et al. 2018). On the one hand, there are robustness metrics that focus on the performance of individual policy options and assess their performance over a set of plausible futures. Well-known examples include minimax and the domain criterion (Starr 1963; Schneller and Sphicas 1983). On the other hand, there are metrics of the performance of policy options relative to a reference point. The best-known example of this type is Savage's minimax regret (Savage 1951), which uses the best possible option for a given future as the reference point against which all other options are to be evaluated.

Over the last decade, a new strategic planning paradigm, known as 'decisionmaking under deep uncertainty,' has emerged that aims to support the development of robust plans. This paradigm rests on three key ideas: (i) Exploratory Modeling; (ii) adaptive planning; and (iii) decision support.

15.2.1 Exploratory Modeling

The first key idea is Exploratory Modeling (EM). In the face of deep uncertainty, one should explore the consequences of the various presently irreducible uncertainties for decisionmaking. Typically, in the case of complex systems, this involves the use of computational scenario approaches.

The idea to systematically explore the consequences of the various uncertainties that are present is rooted in the idea of what-if scenario thinking. Scenarios are (plausible) descriptions of what the future might look like. Scenario thinking is a means for thinking about possible threats and opportunities that the future might hold and their impacts on an organization, business, or system. Scenario thinking gained prominence in part due to pioneering work by Shell in the late 1960s. One of the scenarios that emerged described a very rapid rise in oil price, forcing Shell to consider futures quite different from business as usual. This was believed to have given Shell a competitive advantage during the oil crisis of 1973. Thinking with scenarios when making decisions may help in choosing options that perform reasonably well under a wide range of conditions.

Why is there such a strong insistence on the use of models? There is ample evidence that human reasoning with respect to complex uncertain systems is

intrinsically insufficient. Often, mental models are event based, have an open-loop view of causality, ignore feedback, fail to account for time delays, and are insensitive to nonlinearity (Sterman 1994). In complex systems, the overall dynamics, however, are due to accumulations, feedbacks, and time delays, with nonlinear interactions among them. Thus, mental simulations of complex systems are challenging to the point of infeasibility. This is confirmed empirically in various studies (Sterman 1989; Brehmer 1992; Kleinmuntz 1992; Diehl and Sterman 1995; Atkins et al. 2002; Sastry and Boyd 1998). This strongly suggests that it is worthwhile to support human reasoning on uncertain complex systems with simulation models that are much better at adequately deriving the consequences from sets of hypotheses pertaining to the functioning of these systems (Sterman 2002).

EM is a research method that uses computational experimentation for analyzing complex and uncertain systems (Bankes 1993; Bankes et al. 2013). In the presence of deep uncertainty, the available information enables the development of a set of models, but the uncertainty precludes the possibility of narrowing down this set to a single true representation of the system of interest. A set of models that is plausible or interesting in a given context is generated by the uncertainties associated with the problem of interest and is constrained by available data and knowledge. A single model drawn from the set is not a prediction. Rather, it is a computational what-if experiment that reveals how the real-world system would behave if the specific assumptions about the various uncertainties encapsulated in this model were correct. A single what-if experiment is typically not that informative other than to suggest the plausibility of its outcomes, which in turn may contribute to the substantiation of the necessity to intervene. Instead, EM aims to support reasoning and decisionmaking on the basis of a comprehensive set of such models for the system of interest. In contrast to more traditional scenario planning approaches, EM allows reasoning over a much larger set of cases than a scenario process can generate, while maintaining consistency across the set of cases. The analysis of this set of cases allows humans to infer systematic regularities among subsets of the full ensemble of cases. Thus, EM involves searching through the set of models using (many-objective) optimization algorithms, and sampling over the set of models using computational design of experiments and global sensitivity analysis techniques.

15.2.2   Adaptive Planning

The second key idea underpinning many deep uncertainty approaches is the idea of adaptive planning. Adaptive planning means that plans are designed from the outset to be adapted over time in response to how the future may actually unfold. The way a plan is designed to adapt in the face of potential changes in conditions is announced simultaneously with the plan itself rather than taking place in an ad hoc manner *post facto*. The flexibility of adaptive plans is a key means of achieving decision robustness. While the future is unfolding, many deep uncertainties are being resolved. Having an adaptive plan allows decisionmakers to adapt the implementation

of the plan in response to this. This means that a wide variety of futures have to be explored during plan design. Insight is needed into which actions are best suited to which futures, as well as what signals from the unfolding future should be monitored to ensure the timely implementation of the appropriate actions. The timing of plan adaptation is not known a priori; it depends on how the future unfolds. In this sense, adaptive planning differs from planned adaptation (see Chapter 13), where changes generally occur at predetermined moments (e.g., every 5 years) and which entails a review of conditions that results in an adaptation of the original plan. Adaptive planning involves a paradigm shift from planning in time to planning conditional on observed developments.

15.2.3   Decision Support

The third key idea is decision support. Decisionmaking on complex and uncertain systems generally involves multiple actors coming to agreement. In such a situation, decisionmaking requires an iterative approach that facilitates learning across alternative framings of the problem, and learning about stakeholder preferences and trade-offs, in a collaborative process of discovering what is possible (Herman et al. 2015). In this iterative approach, the various approaches for decisionmaking under deep uncertainty often put candidate policy decisions into the analysis by stress testing them over a wide range of uncertainties. The uncertainties are then characterized by their effect on the decision. The challenges inherent in such processes are reviewed in depth by Tsoukiàs (2008). He envisions that various decision support approaches and tools are used to enable a constructive learning process among the stakeholders and analysts. Decision support in this conceptualization moves from an a priori agreement on (or imposition of) assumptions about the probability of alternative states of the world and the way in which competing objectives are to be aggregated with the aim of producing a preference ranking of decision alternatives, to an a posteriori exploration of trade-offs among objectives and the robustness of this performance across possible futures. That is, under deep uncertainty decision support should move away from trying to define what is the right choice and instead aim at enabling deliberation and joint sense making among the various parties to a decision.

15.3   A Taxonomy of Approaches and Tools for Supporting Decision Making Under Deep Uncertainty

The availability of a variety of approaches for supporting the making of decisions under deep uncertainty raises a new set of questions. How are the various approaches different? Where do they overlap? Where are they complementary? What computa-

tional support tools do they use? Answering these questions may help pave the way for the future harmonization and potential integration of the various approaches. It might also help in assessing if certain approaches are more applicable in certain decisionmaking contexts than others. Both Hall et al. (2012) and Matrosov et al. (2013b) compare Info-Gap Decision Theory and Robust Decision Making. They conclude that, along quite different analytical paths, both approaches arrive at fairly similar but not identical results. Matrosov et al. (2013a) compare Robust Decision Making with an economic optimization approach. In this case, the results from applying both techniques yield different analytical results, suggesting value to efforts seeking to combine both approaches. Roach et al. (2015, 2016) compare Info-Gap Decision Theory and robust optimization. They conclude that there are substantial differences between the plans resulting from these two approaches and argue in favor of mixed methodologies. Gersonius et al. (2015) compare a Real Options Analysis (in detail reported in Gersonius et al. 2013) with an Adaptation Tipping Point analysis (Kwadijk et al. 2010). They highlight the substantial differences in starting points and suggest that both approaches could be applied together. Essentially, the same is argued by Buurman and Babovic (2016), who compare DAPP and Real Options Analysis. Kwakkel et al. (2016b) compare DAPP and RDM. They argue in favor of combining both approaches. DAPP primarily provides a systematic structure for adaptive plans, while RDM provides a clear iterative model-based process for designing adaptive plans.

To move beyond a discussion of the similarities and differences among the various DMDU approaches, a taxonomy of the components that make up the approaches is useful. A first such taxonomy was put forward by Herman et al. (2015). This taxonomy focused on model-based robustness frameworks for supporting decisionmaking under deep uncertainty. In light of the key ideas introduced in the previous section, one could say that the Herman taxonomy focused exclusively on EM. The purpose of this chapter, however, is to cover the broader field and not restrict ourselves to approaches based exclusively on EM. Practically, this means that there is a need to supplement the Herman taxonomy with an additional category related to adaptive planning (the second key idea).

Figure 15.1 shows the proposed taxonomy of the components that make up the various DMDU approaches. It covers five broad categories:

1. *Policy architecture*, which covers the various ways in which adaptive policies may be structured.
2. *Generation of policy alternatives*, which covers how policy alternatives or components thereof are identified, given a specification of the available policy levers.
3. *Generation of scenarios*, which covers how context scenarios are identified given a variety of uncertainties.
4. *Robustness metrics*, which cover the various ways in which policy robustness is operationalized.
5. *Vulnerability analysis*, which covers the various analysis techniques that are used to understand how policy robustness is influenced by both uncertainties and policy levers.
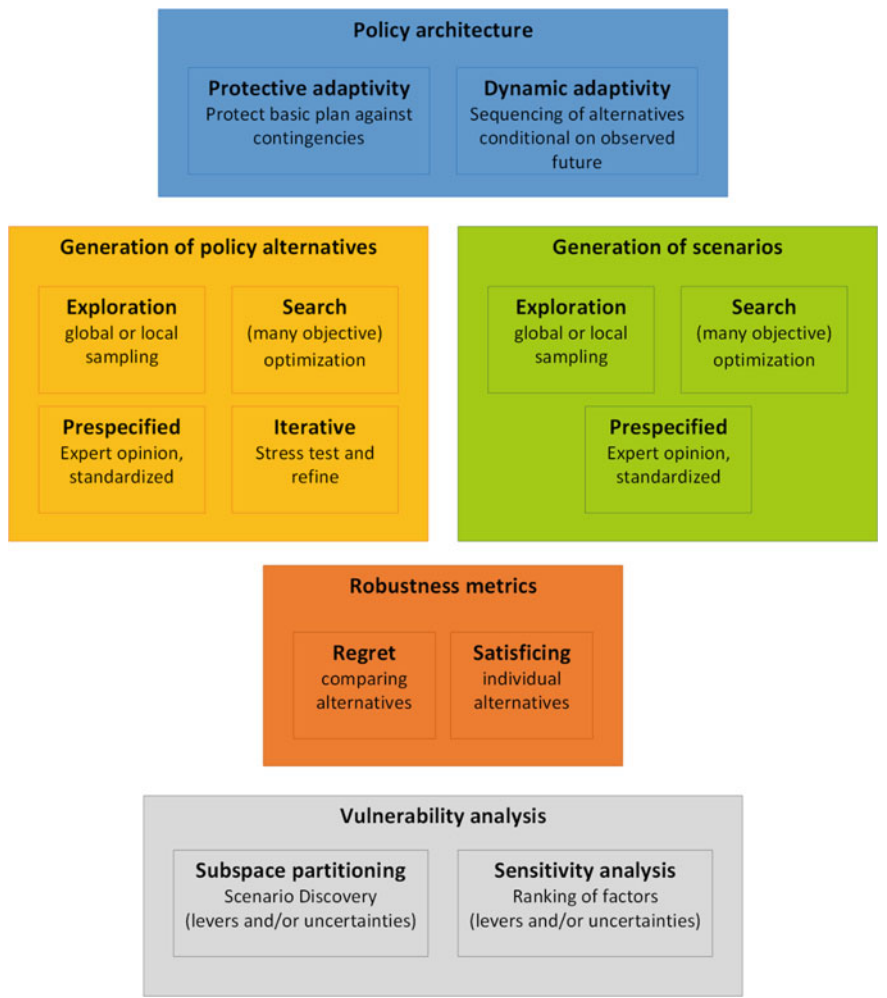
Fig. 15.1 A taxonomy of components that make up the DMDU approaches

In the following subsections, we discuss each of these categories in more detail.

15.3.1   Policy Architecture

In the deep uncertainty literature and before, various ideas on adaptive plans or policies and the importance of flexibility in planning and decisionmaking can be found. Under deep uncertainty, static plans are likely to fail, become overly costly to protect against failure, or are incapable of seizing opportunities. An alternative

is to design flexible plans that can be adapted over time. In this way, a policy may be yoked to an evolving knowledge base (McCray et al. 2010). Adaptive actions are implemented only if, in monitoring how the various uncertainties are resolving, there is a clear signal that these actions are needed. Signals can come both from monitoring data and from computer simulations of future developments. Many examples of authors arguing in favor of this paradigm may be found (Albrechts 2004; Schwartz and Trigeorgis 2004; Swanson et al. 2010; Eriksson and Weber 2008; Walker et al. 2001; Kwakkel et al. 2010a; Lempert et al. 2003; de Neufville and Odoni 2003).

The initial seeds for the adaptive planning paradigm were sown almost a century ago. Dewey (1927) put forth an argument proposing that policies be treated as experiments with the aim of promoting continual learning, and adapted in response to experience over time. Early applications of this idea can be found in the field of environmental management (Holling 1978; McLain and Lee 1996) where, because of the uncertainty about system functioning, policies are designed from the outset to test clearly formulated hypotheses about the behavior of an ecosystem being changed by human use (Lee 1993). A similar attitude is also advocated by Collingridge (1980) with respect to the development of new technologies. Given ignorance about the possible side effects of technologies under development, he argues that one should strive for correctability of decisions, extensive monitoring of effects, and flexibility. Policy learning is also major issue in evolutionary economics of innovation (De La Mothe 2006; Faber and Frenken 2009; Mytelka and Smith 2002).

More recently, Walker et al. (2001) developed a structured, stepwise approach for dynamic adaptation. They advocate that policies should be adaptive: One should take only those actions that are non-regret and time-urgent, and postpone other actions to a later stage. They suggest that a monitoring system and a pre-specification of responses when specific trigger values are reached should complement a basic policy. The resulting policy is flexible and adaptive to the future as it unfolds. The idea of adaptive policies was extended further by Haasnoot et al. (2012), who conceptualized a plan as a sequence of actions to be realized over time. In later work, they called this adaptive policy pathways (Haasnoot et al. 2013).

A policy architecture is the overarching structure that is used to design a plan. A range of adaptive policy architectures is possible. At one extreme (left box under 'Policy architecture' in Fig. 15.1), we have a basic plan to be implemented immediately, complemented by a set of contingency actions that are to be implemented if and when necessary (see Fig. 15.2).

This style of policy architecture can be categorized as 'protective adaptivity.' It is implicitly advocated in Assumption-Based Planning (Dewar 2002; Dewar et al. 1993). Dynamic Adaptive Planning (Walker et al. 2001; Kwakkel et al. 2010a; Chap. 3) also sits in this box. The other extreme no longer views a basic policy as a basic plan plus contingency actions but as a series of actions, the implementation of which coevolve with how the future unfolds. This style of policy architecture can be categorized as 'dynamic adaptivity.' Dynamic Adaptive Policy Pathways (Haasnoot et al. 2013; Chap. 4) exemplifies this type of policy architecture (see Fig. 15.3).
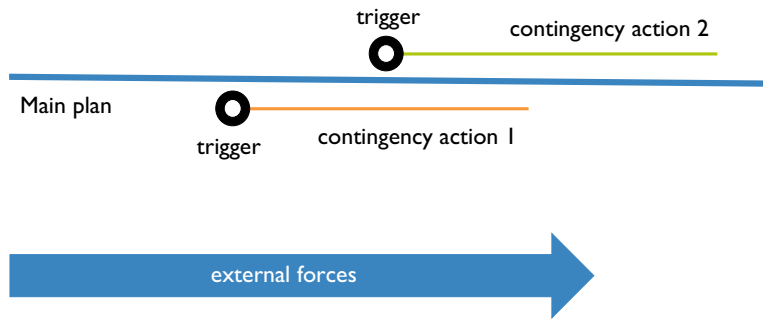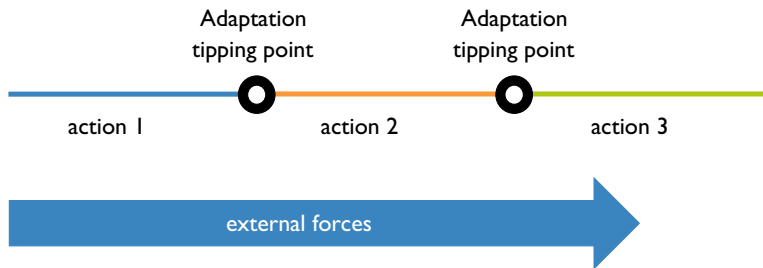
Fig. 15.2 Adaptive policy

Fig. 15.3 Adaptation pathway

15.3.2 *Generation of Policy Alternatives and Generation of Scenarios*

Given a wide variety of deeply uncertain factors and a set of policy levers that may be used to steer the system toward more desirable functioning, the analyst has to choose how to investigate the influence of both uncertainties and levers on outcomes. Broadly speaking, a distinction can be drawn between two different strategies: 'exploration' and 'search.' Exploration strategies investigate the properties of the uncertainty space and the policy lever space by systematically sampling points in these spaces and evaluating their consequences. Exploration relies on the careful design of experiments and can use techniques such as Monte Carlo sampling, Latin Hypercube sampling, or factorial methods. Exploration may be used to answer questions such as 'under what circumstances would this policy do well?', 'under what circumstances would it likely fail?', and 'what dynamics could this system exhibit?'. Exploration provides insight into the global properties of the uncertainty space and the policy lever space.

In contrast, search strategies hunt through these spaces in a more directed manner, searching for points with particular properties. Search may be used to answer questions such as 'What is the worst that could happen?', 'What is the best that could happen?', 'How big is the difference in performance between rival policies?', 'What would a good strategy be given one or more scenarios?' Search provides

detailed insights into particular points in the uncertainty space (called *scenarios*) or in the policy lever space (called *policy alternatives*). Search relies on the use of (many-objective) optimization techniques.

A third strategy is to have pre-specified scenarios or policies, instead of requiring systematic investigation of the entire uncertainty space and policy lever space. For example, in its original inception, RDM assumed that a set of candidate policies is pre-specified. The performance of these policies is then investigated over a wide range of scenarios. Here, the focus of the exploration is on the impact of the uncertainties on the performance of the pre-specified policies.

In practice, the different strategies for investigating the impacts of uncertainties and policy levers can be combined, as well as executed in an iterative manner. For example, if exploration reveals that there are distinct regions in which a policy fails, search may be employed to identify more precisely where the boundary is located between these distinct regions. (This is the essence of Scenario Discovery.) Another example is to use exploration to identify the conditions under which a policy fails and use this insight to modify the policy. By iterating, a policy can be designed that performs acceptably under a wide range of future conditions. RDM, for example, strongly emphasizes the iterative refinement of candidate policies.

15.3.3 Robustness Metrics

A well-established distinction regarding robustness metrics is the distinction between 'regret' metrics and 'satisficing' metrics (Lempert and Collins 2007). *Regret metrics* are comparative and originate from Savage (1954, p. 21), who defines the regret of a policy option as the difference between the performance of the option in a specific state of the world and the performance of the best possible option in that state of the world. Using Savage's definition, a robust policy is one that minimizes the maximum regret across the states of the world (SOWs). Other regret metrics use some type of baseline performance for a given state of the world instead of the performance of the best policy alternative (Popper et al. 2009; Kasprzyk et al. 2013; Lempert and Collins 2007).

Satisficing metrics build on the work of Simon (1996), who pointed out that decisionmakers often look for a decision that meets one or more requirements, but may not achieve the optimal possible outcomes. Satisficing metrics aim at maximizing the number of states of the world in which the policy alternative under consideration meets minimum performance thresholds. A well-known example of this is the 'domain criterion' (Starr 1963; Schneller and Sphicas 1983), which focuses on the fraction of the space in which a given performance threshold is met; the larger this space, the more robust the policy. Often, this is simplified to looking at the fraction of states of the world, rather than the volume of the space.

Synthesizing the foregoing, in defining robustness an analyst must make two choices. The first choice is about the policy alternatives. Is the analyst interested in the performance of individual policies or in comparing the performance of the

Table 15.1 Conceptual representation of various robustness metrics

| | | Characterizing performance of policy alternatives | |
		Comparing alternatives	Performance of individual policies
Characterizing performance over SOWs	Threshold	Satisficing regret	• Domain criterion (Starr 1963) • Radius of stability, Info-Gap (Ben Haim 2001)
	Descriptive statistics	• Minimax regret (Savage 1954) • 90th percentile baseline regret (Kasprzyk et al. 2013) • 90th percentile best option regret (Herman et al. 2015)	Moments of the distribution (e.g., mean, variance), minimum, maximum, and functions thereof (Hurwicz, signal-to-noise, coefficient of variation)

alternative policies? Out of this choice comes a distribution of performance of each policy alternative over the set of states of the world. The second choice is how to succinctly describe this distribution. Here, the analyst also has two options: imposing some user-specified performance threshold, or characterizing the distribution using descriptive statistics. Table 15.1 presents a simple typology of robustness metrics building on these two choices.

Many of the classic decision analytic robustness metrics belong to the lower right hand corner of Table 15.1; they focus on the performance of individual policies and try to describe the performance of a policy over the set of states of the world using descriptive statistics. Maximin and minimax focus on the best and worst performance over the set of scenarios, while Hurwicz is a function of both. Similarly, Laplace's principle of insufficient reason assigns an equal weight to each of the states of the world and then suggests that the best policy is the one with the best mean performance. More recently, there has been interest in using higher-order moments as well. Hamarat et al. (2014) use a signal-to-noise ratio that considers both the average and the variance. A problem here is that combinations of the mean and variance are not always monotonically increasing (Ray et al. 2013). Moreover, focusing on the variance or standard deviation means that positive and negative deviations from the mean are treated equally (Takriti and Ahmed 2004). This explains why higher-order moments, skewness, and kurtosis have attracted attention (Kwakkel et al. 2016a).

15.3.4 Vulnerability Analysis

Vulnerability analysis is often used in combination with an exploration strategy for the generation of scenarios. That is, exploration is used to generate an ensemble of scenarios. Then, one or more alternative policies are evaluated over those scenarios, with vulnerability analysis being used to discover the influence the various uncertainties have on the success or failure of these policies. However, the vulnerability analysis techniques need not be restricted to understanding the role of the uncertainties; they can be used equally well to investigate the role of the policy levers.

Broadly speaking, two distinct styles of vulnerability analysis are being used. On the one hand, we can use global 'sensitivity analysis' to identify the relative importance of the various uncertainties or policy levers on the outcomes of interest. As Herman et al. (2015) point out, this is underutilized in the deep uncertainty literature. Global sensitivity analysis can serve various functions in the context of a deep uncertainty study. It may be used for factor prioritization—that is, to identify the relative influence of the various uncertainties or policy levers on the outcomes of interest. This helps reduce the dimensionality of the problem by focusing subsequent analyses on the key sources of uncertainties or to search over the most influential policy levers. The results of a sensitivity analysis may also enhance understanding of which uncertainties really matter. This is valuable information for designing policies—for example, by designing policies that are capable of adaptation conditional on how these uncertainties resolve over time, or by designing policies that have a reduced sensitivity to these factors. See Herman et al. (2014) for some examples of the usefulness of global sensitivity analyses for supporting the making of decisions under deep uncertainty.

On the other hand, we can use 'subspace partitioning' to find particular subspaces in either the uncertainty space or the policy lever space that result in a particular class of model outcomes. A well-established tool for this is Scenario Discovery (SD), in which one tries to find subspaces of the uncertainty space in which a candidate policy fails. The Adaptation Tipping Point (ATP) concept, which is central to DAPP, is essentially the same: An ATP specifies the uncertain conditions under which the existing actions on the pathway fail to achieve the stated objectives. Decision Scaling (DS) also employs a similar idea. In all three cases, the analyst tries to partition the uncertainty space into distinct regions based on the success or failure of a candidate policy.

15.4 Application of the Taxonomy

Table 15.2 contains a summary overview of the various approaches and computational support tools for DMDU using the taxonomy presented in Fig. 15.1. RDM is not explicitly based on a policy architecture. In most of its applications, however, it uses a form of protective adaptivity, drawing on ideas from Assumption-Based

Planning (ABP) and adaptive policymaking. For the generation of scenarios, it uses sampling. Policies typically are pre-specified and iteratively refined. RDM uses a domain criterion in the vulnerability phase, while it closes with a regret-based analysis of the leading decision alternatives. For the vulnerability analysis, RDM uses SD. Many-Objective Robust Decision Making (MORDM) uses the RDM structure, but replaces pre-specified policies with a many-objective search for a reference scenario. The advantage of this is that the set of initial decision policy alternatives resulting from this process have good performance and represent a careful exploration of the design space.

Three approaches that clearly fit into the policy architecture are Assumption-Based Planning (ABP), Dynamic Adaptive Planning (DAP), and Dynamic Adaptive Policy Pathways (DAPP). They represent opposing ends of the spectrum of adaptive policy architectures. ABP and DAP focus on the use of adaptivity to protect a basic plan from failing. ABP relies primarily on qualitative judgments. DAP says little with respect to how to use models for designing adaptive polices. Hamarat et al. (2013) use RDM for this. In contrast, DAPP can only be assembled from a given set of actions. These need to be found first. In most applications of DAPP, the actions are assumed to be given. For vulnerability analysis, DAPP uses ATPs, which can be seen as a one-dimensional version of SD.

SD, ATP, DS, and Info-Gap Decision Theory (IG) are all essentially vulnerability analysis tools that may be used to design adaptive plans. Of these, SD is the most general form. ATP analysis is, in essence, a one-dimensional SD. Similarly, DS is a one- or at most two-dimensional form of SD that is exclusively focused on climate information. All three focus on finding subspace(s) of the uncertainty space in which a policy fails. For DS and ATP, this subspace has generally been characterized by climate change information, although the latter can easily be used more generally. SD is agnostic with respect to the uncertainties to be considered. IG is a bit different in this respect, since it requires a reference scenario from which the 'radius of stability' is calculated. As such, it does not produce insights into the subspace in which a policy fails, but instead into how far the future is allowed to deviate from the reference scenario before a policy starts to fail. This means that the choice of the reference scenario and the distance metric used become analytical issues.

Engineering Options Analysis (EOA) is a bit of an outlier in this overview. Its primary focus is on assigning economic value to flexibility. As such, it might be used as part of an RDM study. In this case, SD might be used to identify when the net present value of a given option is lower than its costs. Note also that, in contrast to Real Options Analysis, EOA is easier to combine with both types of policy architectures. Again, EOA might be used to analyze in more detail the value of actions on a given pathway, or across pathways.

Many-Objective Robust Optimization (MORO) is a generic computational tool for designing adaptive plans. It has been used in combination with both a DAP approach (Hamarat et al. 2014) and DAPP (Kwakkel et al. 2015). It extends the argument found in the MORDM literature on the relevance of finding promising designs prior to performing in-depth analyses by bringing robustness considerations into the search itself. Because of this, it does not include a vulnerability analysis.

Table 15.2 Application of the taxonomy to various approaches and tools for supporting the making of decisions under deep uncertainty

	Key references	Policy architecture	Generation of scenarios	Generation of policy alternatives	Robustness metrics	Vulnerability analysis
Assumption-Based Planning (ABP)	Dewar (2002), Dewar et al. (1993)	Protective adaptivity	Not considered; focus is on assumptions that might fail, not on scenarios where they might fail	Pre-specified	Typically focused on satisficing measures	Qualitative judgment
Robust Decision Making (RDM)	Lempert et al. (2006)	Not explicitly considered, but often combined with protective adaptivity	Sampling	Generally pre-specified, and iteratively refined	The SD phase uses a satisficing metric (domain criterion); the trade-off analysis phase relies on regret	Scenario Discovery (Bryant and Lempert 2010; Kwakkel and Jaxa-Rozen 2016)
Many-Objective Robust Decision Making (MORDM)	Kasprzyk et al. (2013), Watson and Kasprzyk (2017)	Not explicitly considered	Sampling	Many-objective search using one or a few scenarios separately	The SD phase uses a satisficing metric (domain criterion); the trade-off analysis phase relies on regret	Scenario Discovery
Dynamic Adaptive Planning (DAP)	Kwakkel et al. (2010a), Walker et al. (2001)	Protective adaptivity	Not explicitly considered	Not explicitly considered	Typically focused on satisficing measures	Not explicitly considered
Dynamic Adaptive Policy Pathways (DAPP)	Haasnoot et al. (2012, 2013)	Dynamic adaptivity	Not explicitly considered, but strong emphasis on the need for transient scenarios	Typically pre-specified	Typically focused on satisficing measures	ATP (Kwadijk et al. 2010)
Info-Gap Decision Theory (IG)	Ben Haim (2001, 2004)	Not considered	Sampling outward from a reference scenario	Pre-specified	Satisficing (Radius of stability)	Not considered

Table 15.2 (continued)

	Key references	Policy architecture	Generation of scenarios	Generation of policy alternatives	Robustness metrics	Vulnerability analysis
Engineering Options Analysis (EOA)	de Neufville and Scholtes (2011)		Monte Carlo simulation	Pre-specified	Values at risk	Values at risk
Decision Scaling (DS)	Brown et al. (2012), LeRoy Poff et al. (2015)	Not considered	Pre-specified, or sampling constrained by climate information	Pre-specified	Satisficing (domain criterion)	Climate response function (visual), ANOVA ranking
Scenario Discovery (SD)	Bryant and Lempert (2010), Groves and Lempert (2007), Kwakkel and Jaxa-Rozen (2016)	Not considered	Sampling	Typically pre-specified	Domain criterion (implicit)	Dedicated vulnerability technique in itself
Adaptation Tipping Points (ATP)	Kwadijk et al. (2010)	Not considered	Pre-specified, or sampling	Typically pre-specified	Satisficing (domain criterion)	Dedicated vulnerability technique in itself
Many-Objective Robust Optimization (MORO)	Hamarat et al. (2014), Kwakkel et al. (2015)	Not considered, can be used for both protective and dynamic adaptivity	Pre-specified, or sampling	Many-objective search	Compatible with all satisficing metrics, and reference scenario regret	Not considered

15.5 Concluding Remarks

This chapter has presented and described a taxonomy of the components that make up the various approaches and tools for supporting the making of decisions under deep uncertainty. In short, there are five categories of components: policy architecture, generation of scenarios, generation of alternatives, definition of robustness, and vulnerability analysis. Any given DMDU approach makes choices with respect to these five categories. For some, these choices are primarily or almost exclusively in one category while remaining silent on the others. For others, implicit or explicit choices are made with respect to each category. Table 15.2 summarizes this.

Going forward, it might be useful to use the taxonomy to articulate which choices are made in a given case, rather than having to argue over the exact differences among approaches and the merits of these differences. That is, rather than arguing over whether to apply RDM or DAPP, the discussion should be which combination of tools are appropriate to use given the nature of the problem situation. Different situations warrant different combinations. For example, when designing a new piece of infrastructure (say a large reservoir), it makes sense to use a many-objective search to find promising design alternatives prior to performing a vulnerability analysis. In other cases, a set of policy alternatives might already be available. In such a case, doing additional search might be less relevant.

References

Albrechts, L. (2004). Strategic (spatial) planning reexamined. *Environment and Planning B: Planning and Design, 31,* 743–758.

Atkins, P. W., Wood, R. E., & Rutgers, P. J. (2002). The effects of feedback format on dynamic decision making. *Organizational Behavior and Human Decision Processes, 88*(2), 587–604.

Bankes, S. C. (1993). Exploratory modeling for policy analysis. *Operations Research, 4*(3), 435–449.

Bankes, S. C., Walker, W. E., & Kwakkel, J. H. (2013). Exploratory Modeling and Analysis. In S. Gass & M. C. Fu (Eds.), *Encyclopedia of operations research and management science* (3rd ed.) Berlin, Germany: Springer.

Ben Haim, Y. (2001). *Information-gap decision theory: Decision under severe uncertainty*. London, UK: Academic Press.

Ben Haim, Y. (2004). Uncertainty, probability and information-gaps. *Reliability Engineering and System Safety, 85*(1–3), 249–266.

Brehmer, B. (1992). Dynamic decision making: Human control of complex systems. *Acta Psychologica, 81*(3), 211–241.

Brown, C., Ghile, Y., Laverty, M., & Li, K. (2012). Decision scaling: Linking bottom-up vulnerability analysis with climate projections in the water sector. *Water Resources Research, 48*(9), 1–12. https://doi.org/10.1029/2011WR011212.

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*(1), 34–49. https://doi.org/10.1016/j.techfore.2009.08.002.

Buurman, J., & Babovic, V. (2016). Adaptation pathways and real options analysis—An approach to deep uncertainty in climate change adaptation policies. *Policy and Society*. http://www.dx.org/10.1016/ j.polsoc.2016.05.002.

Collingridge, D. (1980). *The social control of technology*. London, UK: Frances Pinter Publisher.

De La Mothe, J. (2006). *Innovation strategies in interdependent states*. Gloucestershire: Edward Elgar Publishing Ltd.

de Neufville, R., & Odoni, A. (2003). *Airport systems: Planning, design, and management*. New York: McGraw-Hill.

de Neufville, R., & Scholtes, S. (2011). *Flexibility in engineering design*. Cambridge, Massachusetts: The MIT Press.

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

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 Report MR-114-A. Available at: http://www.rand.org/pubs/monograph_reports/2005/MR114.pdf.

Dewey, J. (1927). *The public and its problems*. New York: Holt and Company.

Diehl, E., & Sterman, J. D. (1995). Effects of feedback complexity on dynamic decision making. *Organizational Behavior and Human Decision Processes, 62*(2), 198–215.

Eriksson, E. A., & Weber, K. M. (2008). Adaptive foresight: Navigating the complex landscape of policy strategies. *Technological Forecasting and Social Change, 75*(4), 462–482.

Faber, A., & Frenken, K. (2009). Models in evolutionary economics and environmental policy: Towards an evolutionary environmental economics. *Technological Forecasting and Social Change, 76*(4), 462–470. https://doi.org/10.1016/j.techfore.2008.04.009.

Gersonius, B., Ashley, R., Jeuken, A., Pathinara, A., & Zevenbergen, C. (2015). Accounting for uncertainty and flexibility in flood risk management: Comparing real-in-options optimisation and adaptation tipping points. *Journal of Flood Risk Managmeent, 8,* 135–144. https://doi.org/10.1111/jfr3.12083.

Gersonius, B., Ashley, R., Pathirana, A., & Zevenbergen, C. (2013). Climate change uncertainty: building flexibility into water and flood risk infrastructure. *Climatic Change, 116,* 411. https://doi.org/10.1007/s10584-012-0494-5.

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. https://doi.org/10.1007/s10584-015-1586-9.

Groves, D. G., & Lempert, R. J. (2007). A new analytic method for finding policy-relevant scenarios. *Global Environmental Change, 17*(1), 73–85.

Haasnoot, M., Kwakkel, J. H., Walker, W. E., & ter Maat, J. (2013). Dynamic adaptive policy pathways: A method for crafting robust decisions for a deeply uncertain world. *Global Environmental Change, 23*(2), 485–498. https://doi.org/10.1016/j.gloenvcha.2012.12.006.

Haasnoot, M., Middelkoop, H., Offermans, A., van Beek, E., & van Deursen, W. P. A. (2012). Exploring pathways for sustainable water management in river deltas in a changing environment. *Climatic Change, 115*(3–4), 795–819. https://doi.org/10.1007/s10584-012-0444-2.

Hall, J. W., Lempert, R. J., Keller, A., Hackbarth, A., Mijere, C., & McInerney, D. (2012). Robust climate policies under uncertainty: A comparison of robust decision making and info-gap methods. *Risk Analysis, 32*(10), 1527–1672. https://doi.org/10.1111/j.1539-6924.2012.01802.x.

Hamarat, C., Kwakkel, J. H., & Pruyt, E. (2013). Adaptive robust design under deep uncertainty. *Technological Forecasting and Social Change, 80*(3), 408–418. https://doi.org/10.1016/j.techfore.2012.10.004.

Hamarat, C., Kwakkel, J. H., Pruyt, E., & Loonen, E. (2014). An exploratory approach for adaptive policymaking by using multi-objective robust optimization. *Simulation Modelling Practice and Theory, 46,* 25–39. https://doi.org/10.1016/j.simpat.2014.02.008.

Herman, J. D., Reed, P. M., Zeff, H. B., & Characklis, G. W. (2015). How should robustness be defined for water systems planning under change. *Journal of Water Resources Planning and Management*, 141(10). https://doi.org/10.1061/(asce)wr.1943-5452.0000509.

Herman, J. D., Zeff, H. B., Reed, P. M., & Characklis, G. (2014). Beyond optimality: Multistake-holder robustness tradeoffs for regional water portfolio planning under deep uncertainty. *Water Resources Research, 50*(10), 7692–7713. https://doi.org/10.1002/201-4WR015338.

Holling, C. S. (1978). *Adaptive environmental assessment and management*. New York: Wiley.

Kasprzyk, J. R., Nataraj, S., Reed, P. M., & Lempert, R. J. (2013). Many objective robust decision making for complex environmental systems undergoing change. *Environmental Modelling and Software, 42,* 55–71. https://doi.org/10.1016/j.envsoft.2012.007.

Kleinmuntz, B. (1992). Computers as clinicians: An update. *Computers in Biology and Medicine, 22*(4), 227–237.

Kwadijk, J. C. J., Haasnoot, M., Mulder, J. P. M., Hoogvliet, M. M. C., Jeuken, A. B. M., van der Krogt, R. A. A., et al. (2010). Using adaptation tipping points to prepare for climate change and sea level rise: A case study in the Netherlands. *Wiley Interdisciplinary Reviews: Climate Change, 1*(5), 729–740. https://doi.org/10.1002/wcc.64.

Kwakkel, J. H., Walker, W. E., & Marchau, V. A. W. J. (2010a). Adaptive airport strategic planning. *European Journal of Transportation and Infrastructure Research, 10*(3), 227–250.

Kwakkel, J. H., Walker, W. E., & Marchau, V. A. W. J. (2010b). *From predictive modeling to exploratory modeling: How to use non-predictive models for decisionmaking under deep uncertainty*. Paper presented at the Proceedings of the 25th Mini-EURO Conference—Uncertainty and Robustness in Planning and Decision Making, Coimbra, Portugal.

Kwakkel, J. H., Haasnoot, M., & Walker, W. E. (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. https://doi.org/10.1007/s10584-014-1210-4.

Kwakkel, J. H., & Jaxa-Rozen, M. (2016). Improving scenario discovery for handling heterogeneous uncertainties and multinomial classified outcomes. *Environmental Modelling and Software, 79,* 311–321. https://doi.org/10.1016/.envsoft.2015.11.020.

Kwakkel, J. H., Eker, S., & Pruyt, E. (2016a). How robust is a robust policy? Comparing alternative robustness metrics for robust decision-making. In M. Doumpos, C. Zopounidis, & E. Grigoroudis (Eds.), *Robustness analysis in decision aiding, optimization, and analytics*. Springer.

Kwakkel, J. H., Haasnoot, M., & Walker, W. E. (2016b). Comparing Robust decision-making and dynamic adaptive policy pathways for model-based decision support under deep uncertainty. *Environmental Modelling and Software, 86,* 168–183. https://doi.org/10.1016/j.envsoft.2016.09.017.

Lee, K. (1993). *Compass and gyroscope: Integrating science and politics for the environment*. Washington, DC, USA: Island Press.

Lempert, R. J., & Collins, M. (2007). Managing the risk of uncertain threshold response: Comparison of robust, optimum, and precautionary approaches. *Risk Analysis, 24*(4), 1009–1026. https://doi.org/10.1111/j.1539-6924.2007.00940.x.

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. https://doi.org/10.1287/mnsc.1050.0472.

Lempert, R. J., Popper, S., & Bankes, S. (2003) *Shaping the next one hundred years: New methods for quantitative, long term policy analysis*. Santa Monica, CA: RAND Report MR-1626-RPC.

LeRoy Poff, N., Brown, C., Grantham, T. E., Matthews, J. H., Palmer, M. A., Spence, C. M., et al. (2015). Sustainable water management under future uncertainty with eco-engineering decision scaling. *Nature Climate Change, 6*(1), 25–34. https://doi.org/10.1038/NCLIMATE2765.

Matrosov, E. S., Padula, S., & Harou, J. J. (2013a). Selecting portfolios of water supply and demand management strategies under uncertainty—Contrasting economic optimisation and 'robust decision making' approaches. *Water Resource Management, 27*(4), 1123–1148. https://doi.org/10.1007/s11269-012-0118-x.

Matrosov, E. S., Woords, A. M., & Harou, J. J. (2013b). Robust decision making and info-gap decision theory for water resource system planning. *Journal of Hydrology*, *494*(28 June 2013), 43–58.

McCray, L. E., Oye, K. A., & Petersen, A. C. (2010). Planned adaptation in risk regulation: An initial survey of US environmental, health, and safety regulation. *Technological Forecasting and Social Change, 77*(6), 951–959.

McLain, R. J., & Lee, R. G. (1996). Adaptive management: Promises and pitfalls. *Environmental Management, 20,* 437–448.

McPhail, C., Maier, H. R., Kwakkel, J. H., Giuliani, E., Castelletti, A., & Westra, S. (2018). Robustness metrics: How are they calculated, when should they be used and why do they give different results? *Earth's Future*. https://doi.org/10.1002/2017ef000649.

Mytelka, L. K., & Smith, K. (2002). Policy learning and innovation theory: An interactive and co-evolving process. *Research Policy, 31*(8–9), 1467–1479.

Popper, S., Griffin, J., Berrebi, C., Light, T., & Min, E. Y. (2009). *Natural gas and Israel's energy future: A strategic analysis under conditions of deep uncertainty*. Santa Monica, California: RAND.

Ray, P. A., Watkins, D. W., Vogel, R. M., & Kirshen, P. H. (2013). Performance-based evaluation of an improved robust optimization formulation. *Journal of Water Resources Planning and Management*, *140*(6). https://doi.org/10.1061/(asce)wr.1943-5452.0000389.

Roach, T., Kapelan, Z., & Ledbetter, R. (2015). Comparison of info-gap and robust optimisation methods for integrated water resource management under severe uncertainty. *Procedia Engineering, 119,* 874–883. https://doi.org/10.1016/j.proeng.2015.08.955.

Roach, T., Kapelan, Z., Ledbetter, R., & ledbetter, M. (2016). Comparison of robust optimization and info-gap methods for water resource management under deep uncertainty. *Journal of Water Resources Planning and Management*, 142(9). https://doi.org/10.1061/(asce)wr.1943-5452.0000660.

Sastry, L., & Boyd, D. R. (1998). Virtual environments for engineering applications. *Virtual Reality, 3*(4), 235–244.

Savage, L. T. (1951). The theory of statistical decisions. *Journal of the American Statistical Association, 46*(253), 55–67.

Savage, L. T. (1954). *The foundations of statistics*. New York: Wiley.

Schneller, G. O. I., & Sphicas, G. P. (1983). Decision making under uncertainty: Starr's domain criterion. *Theory and Decision, 15*(4), 321–336. https://doi.org/10.1007/BF00162111.

Schwartz, E. S., & Trigeorgis, L. (2004). *Real options and investment under uncertainty: Classical readings and recent contributions*. The MIT Press.

Simon, H. A. (1996). *The sciences of the artificial*. Cambridge, Massachusetts: The MIT Press.

Starr, M. K. (1963). *Product design and decision theory*. Englewood Cliffs, NJ: Prentice-Hall.

Sterman, J. D. (1989). Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. *Management Science, 35*(3), 321–339.

Sterman, J. D. (1994). Learning in and about complex systems. *System Dynamics Review, 10*(2–3), 291–330.

Sterman, J. D. (2002). All models are wrong: reflections on becoming a systems scientist. *System Dynamics Review, 18*(4), 501–531.

Swanson, D., Barg, S., Tyler, S., Venema, H., Tomar, S., Bhadwal, S., et al. (2010). Seven tools for creating adaptive policies. *Technological Forecasting and Social Change, 77*(6), 924–939. https://doi.org/10.1016/j.techfore.2010.04.005.

Takriti, S., & Ahmed, S. (2004). On robust optimization of two-stage systems. *Mathematical Programming, 99*(1), 109–126. https://doi.org/10.1007/s10107-003-0373-y.

Tsoukiàs, A. (2008). From decision theory to decision aiding methodology. *European Journal of Operational Research, 187,* 138–161. https://doi.org/10.1016/j.ejor.2007.02.039.

Walker, W. E., Rahman, S. A., & Cave, J. (2001). Adaptive policies, policy analysis, and policymaking. *European Journal of Operational Research, 128*(2), 282–289.

Watson, A. A., & Kasprzyk, J. R. (2017). Incorporating deeply uncertain factors into the many objective search process. *Environmental Modelling and Software, 89,* 159–171. https://doi.org/10.1016/j.envsoft.2016.12.001.

Dr. Jan H. Kwakkel is Associate Professor at the Faculty of Technology, Policy and Management (TPM) of Delft University of Technology. He is also the Vice President of the Society for Decision Making under Deep Uncertainty, and a member of the editorial boards of *Environmental Modeling and Software*, *Futures, and Foresight*. His research focuses on model-based support for decision making under uncertainty. He is involved in research projects on smart energy systems, high-tech supply chains, transport, and climate adaptation. Jan is the developer of the Exploratory Modeling Workbench, an open-source software implementing a wide variety of state-of-the-art techniques for decisionmaking under uncertainty.

Dr. Marjolijn Haasnoot is a Senior Researcher and Advisor at Deltares and Associate Professor in adaptive delta management at Utrecht University, the Netherlands. Her research focus is on water management, integrated assessment modeling, and decisionmaking under deep uncertainty. She is working and consulting internationally on projects assessing impacts of climate change and socio-economic developments, and alternative management options to develop robust and adaptive plans. Marjolijn developed the Dynamic Adaptive Policy Pathways planning approach. She has a Ph.D. in engineering and a Masters in environmental science.