



Decision Analysis

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Decision Analysis for Practitioners

Lawrence D. Phillips

To cite this article:

Lawrence D. Phillips (2025) Decision Analysis for Practitioners. *Decision Analysis* 22(4):255-272. <https://doi.org/10.1287/deca.2025.0356>

This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “*Decision Analysis*. Copyright © 2025 The Author(s). <https://doi.org/10.1287/deca.2025.0356>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Copyright © 2025 The Author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Decision Analysis for Practitioners

Lawrence D. Phillips^a

^aDepartment of Management, London School of Economics and Political Science, London WC2A 2AE, United Kingdom

Contact: larry_phillips@msn.com,  <https://orcid.org/0000-0002-4267-4973> (LDP)

Received: March 5, 2025

Revised: April 30, 2025


Accepted: May 23, 2025

Published Online in Articles in Advance:
June 26, 2025

<https://doi.org/10.1287/deca.2025.0356>

Copyright: © 2025 The Author(s)

Abstract. Many practitioners consider decision analysis as a sociotechnical discipline, with probability, utility, and trade-offs as the core components of a model that enables an accountable decision maker experiencing a sense of unease about the present to explore different assumptions about the future and develop a plan about the way forward for the organization. The decision analyst acts as a process consultant, working with the decision maker and key players as a problem solver and applying any of five structural and five content ingredients of decision analysis in building a requisite model that is sufficient in form and content to resolve the problem while acting as a transitional object, which holds and contains the decision maker's unease and anxiety about the future. Ten social skills that enable the decision analyst to serve as a process consultant are explained. Six case studies representing problem types for evaluating options, allocating resources, bargaining and negotiating, choosing and deciding, managing risk, and revising opinion demonstrate the many ways that their acting as a transitional object enables exploring the future.

 **Open Access Statement:** This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “*Decision Analysis*. Copyright © 2025 The Author(s). <https://doi.org/10.1287/deca.2025.0356>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Funding: This work was supported by Facilitations Ltd.

Keywords: decision analysis models • requisite models • transitional objects • decision conferences • process consultancy skills • preferences • judgements

Introduction

Readers of this journal are well aware that the preferences of an informed, rational decision maker can be expressed as axioms, which lead to proofs of theorems establishing the existence of probabilities (von Neumann and Morgenstern 1947), probabilities and utilities (Savage 1954), and preferences and value trade-offs (Keeney and Raiffa 1976). But, real-world decision makers are not necessarily rational (Tversky and Kahneman 1981), so what is the point of decision analysis? This paper explores a new answer to that question, which begins by recognizing that although the axioms of preference logically and uniquely imply the existence theorems, the logic can be considered in reverse; the theorems logically imply the axioms, though not uniquely (Fishburn 1972).

This suggests that the decision analyst might work with an accountable decision maker along with a group of key players applying the technical ingredients of good

decisions to enable the creation of a base-case decision model that includes all the elements needed to represent everyone's concerns even if they do not agree with the base-case results. Subsequent sensitivity analyses exploring these judgement differences will result in persuasive arguments and shifts of risk attitudes (Burnstein and Vinokur 1975), enabling the group to agree on the way forward, even though consensus is lacking about many of the inputs. In short, I am suggesting that the theorems provide the building blocks for a decision model that is good enough to enable agreement about what to do next. What follows is a summary of my experience over 50+ years as an academic/practitioner working with accountable decision makers in private, public, and voluntary organizations of all sizes from start-ups to large corporations.

My initial understanding, as a postgraduate student in the 1960s, led me to think decision theory was mainly Bayesian statistics as Raiffa and Schlaifer (1961) had

presented it, but the textbook of Schlaifer (1969) changed my mind, reinforced when I later read the classic *Decision Analysis* (Raiffa 1968), with both books focusing on modeling uncertainty. Today, a decision tree starting with mutually exclusive courses of action followed by uncertain events with their outcomes and then, more action-event sequences and finally, with financial consequences at the end of all branches of the tree is the structure deployed today by many decision analysts. It gives an answer of what to do.

Requisite Decision Models

That was how I saw decision analysis in my first published case study (Phillips 1982). The managing director (MD) of a medium-sized UK manufacturing firm engaged me to determine what could be done to speed up a decision by his organization's board about changing the design of their best-selling product to improve its safety. Successive presentations to the organization's board of directors over 11 months led them to reject any change, even though the product had nearly killed one user, and the MD feared that the government might ban the product. Furthermore, he knew of developments in new technology that might even make their product obsolete, so he had established an internal group to develop better technology as the current product accounted for 70% of the company's revenues.

Behaving as a traditional 1980s management consultant, I quickly arranged separate meetings with the key internal owners of the problem, the sales, finance, and production managers, and some of their staff, making notes as I learned more about the problem and identifying which elements of good decisions seemed relevant. A decision tree emerged from my thinking with two mutually exclusive decisions (change the design or retain the original design) and uncertainty represented as event nodes (when the new technology threat might become evident, when the company's new technology would be ready, and what the sales volume of the new product would be, all conditional on the initial decision).

After many interviews, I showed the MD a plausible decision tree, which prompted him to remember that the previous introduction of a new product had been troubled; so, that added another event node but just for the new design decision. The final decision tree emerged with 20 new product paths and 30 old product paths, not counting additional branches for the profit calculations. And, the MD guided his team to agree about the

probabilities. At the next board meeting, the MD explained the decision tree, which showed that the new product decision was better than remaining with the old product. Several disagreements by various board members about probabilities and financial values were subjected to sensitivity analysis on the IBM 5110 luggable¹ computer, which had been programmed to facilitate creation of a decision tree and make all the calculations necessary so that the tree could be rolled back to display the expected value of each option. Only if three changes were made cumulatively would the result support the old product, and no member believed all three. So, the board approved the new product at the meeting. (More is said about this case study below in the section Types of Decision Analysis Models.)

I had moved from discussing the problem with the MD to discussing the problem with each of the major problem owners and found that sufficient modeling had been carried out to satisfy the board. And, they had discovered that consensus about every figure was not needed because the model was very robust. In short, it was *requisite*, sufficient in form, and content to resolve the problem (Phillips 1984). In this case, the model represented an existing reality, a product that needed to be changed to a better one, retaining its functionality but safer.

Decision Conferences and Workshops

About this same time, I became aware that Cameron Peterson, the Technical Director of Decisions and Designs, Inc. (DDI), a Washington, DC consulting firm that employed over two dozen PhD decision analysts in the 1970s, had "discovered" what we now know as a decision conference (Phillips 2007). He had been expecting the MD of Westinghouse's Elevator Division to arrive with a small top team to attend a workshop at DDI's offices. Cam was surprised when the MD brought his entire team of about 20 managers, so he quickly adapted to working with them over a three-day period and engaged DDI's report specialists in the evenings to prepare a report of the meeting, copies of which were given to every attendee at the end of the final day after all members had agreed with the results.

How could this happen so quickly? Largely because sensitivity analyses established that disagreements among the team members and imprecision about the future could be explored by displaying the output, helping the team to agree about results, even though they did

not always agree about the inputs. This way of working with all the key players led me to formulate the working hypothesis that we did not need a perfect model, only one that was requisite, which could be recognized at the point when no new intuitions arose among participants during the process of creating the model. It was not an optimal model or a normative model, not even a satisficing model (Simon 1955). And, it was a model about a potential future reality, a new type of elevator, generated by a group of experts with different views about what could be, capitalizing on our human ability to imagine, anticipate, and plan for the future. This seemed to me a new use of decision analysis.

A little later, a further interaction with a client suggested another criterion of requisiteness: when the accountable decision maker's initial sense of unease about what to do is resolved. I had facilitated a decision conference for Robb Wilmot, then the MD of Britain's largest computer company, International Computers Ltd. He was very keen on decision support systems, which predicted that computers would soon take over the drudgery of paper-based systems (Keen and Morton 1978), and Robb foresaw that one day, computers would become commodities, a rare belief in the early 1980s when mainframe computers dominated and social media did not exist. Robb's problem was that he needed clarity about how best to allocate £200 million to develop alternatives to mainframes. A few days after the decision conference, Robb came to the London School of Economics to further explore the model.

He wanted to try out several judgements and assumptions about the future that the group had agreed on. So, he started by questioning one part of the model, and I typed in the changes. But, when he saw the subsequent result, he said he did not like that outcome. So, back to the original model as he tested other judgements, some looking better and some looking worse. Suddenly, he said, "OK, you can close the computer. Now I know what to do." He turned to me and exclaimed, "This is great. I get to try out the future before I have to live it!"

By then, Keeney and Raiffa (1976) had extended decision theory to include multiple objectives and trade-offs, and later, in the 1993 Cambridge University Press Edition that incorporated new work, to say:

When a decision involves multiple objectives—and this is almost always the case with important problems—multiattribute utility theory forms the basic foundation

for applying decision analysis. (Keeney and Raiffa 1993, p. 3)

And, it was *Value-Focused Thinking: A Path to Creative Decisionmaking* by Keeney (1992) that suggested modeling should start by considering objectives first and then, work backward to courses of action that might realize those objectives. The focus on objectives substantially shifted problem-solving and opened a new era in decision analysis by suggesting decisions that were not necessarily even modeled. Sometimes, a model simply compares options, systems, strategies, policies, or anything at all without considering them as alternatives. For example, comparing 20 recreational drugs on 16 criteria of harms to users and to others with a multicriteria decision analysis model (MCDA) challenged the United Kingdom's 50-year-old drug classification system, which was supposedly based on harm (Nutt et al. 2010). Decisions do not appear in the model, which was intended to provide information from experts so that others can make more informed decisions about drugs.

Decision Models as Transitional Objects

In short, decisions themselves are often not the real problem of the client. To this day, I do not know what Robb Wilmot decided. But, that does not matter because what I saw was a client using the decision model to explore what might happen in the future and formulating what he could do now to best move forward. There was clearly no single "best" option for him. I think he probably formulated a plan in his head because not all the £200 million had to be allocated at once. He could position the early research to inform later decisions.

Instead, it is the gap felt by the decision maker between where their organization is now and where they would rather be that provides the motivation for sociotechnical decision analysis. It is our unique human ability to imagine alternative futures, make judgments, and form preferences, guided by our values and goals, that can see us through uncertain times. But, that very ability may also leave us confused, with a sense of unease about the present and unsure how to proceed. Close inspection of that unease reveals that the decision maker's preferences are ill formed or absent, sometimes because there are too many alternatives, conflicting objectives, or uncertain future events, with no guidance on what to trade-off against what. Or, they may feel that

making a judgement is premature, that additional data are required, or that deeper exploration of the problem is required. Furthermore, culture influences preferences in the family, group, organization, or country (Falk et al. 2018), making agreement about what to do next even more problematical.

Fortunately, decision analysis gives us the vocabulary and its grammar to clarify the nature of these types of problems. As we will see below, building a decision-analytic model might start with any of these three features to gain clarity about what the future might hold and how best to move forward: no optimizing, no providing the “right” answer, and no instructing the decision maker about what to do. Rather, apply decision analysis as a process for the decision maker to analyze the problem, with the decision model serving as a *transitional object*, containing many perspectives and judgements, allowing time for the decision maker to “try out the future,” and holding the anxiety experienced by the decision maker and others while exploring the future. I have borrowed the concept of transitional change from Bridger (2001), whose two requirements for the transitional process of change are (1) time and space for the meeting and (2) a form of “cover” for safely exploring a new relationship with the external world. Both are met by the decision-conferencing process; time and space are carefully planned with the client, whereas the facilitator and the computer model provide safe cover.

A requisite decision model serves as a transitional object, which holds the decision maker’s anxiety about what to do next by providing a safe zone for participants to explore the model, simulating and trying out possible futures without commitment. The transitional object suspends time so that the decision maker can think about the future. As new intuitions arise about the problem, this enables decision makers to “try out the future before having to live it” and to plan and move forward confidently. The transitional object may then be discarded or perhaps, continued to monitor the process of moving in a new direction.

You can now understand how I arrived at this new view of decision analysis. In summary, a requisite decision model serves as a transitional object, enabling the accountable decision maker who is experiencing a sense of unease about the present to safely explore possible futures and construct a plan that will resolve the sense of unease as the organization moves ahead. When working

with a group, the model helps participants to develop a shared understanding of the key issues, generate a sense of common purpose, and gain a degree of commitment to the way forward. So, social purposes are achieved by the technical model, which justifies labeling this form of decision analysis as sociotechnical.

The next section provides an overview of what is meant by decision analysis as technology: in particular, the ingredients of structure, content, and process. It is followed by a second section considering the social skills of the decision analyst that are required to facilitate the process of building and exploring the model. The third section discusses the types of models that apply the ingredients of decision analysis in different ways.

Decision Analysis Technology

Subsequent developments of the theory to accommodate multiple objectives and trade-offs (Keeney and Raiffa 1976, Dyer and Sarin 1979) were axiomatic treatments for defining utility as representing an individual’s strength of preference, providing an interval scale of measurement between alternatives or consequences. These developments simplified decision analysis by replacing “utility” with “preference value,” thereby clarifying the distinction in decision theory between risky and riskless choice, with anything considered as risky defined by the probabilities of uncertain events, by separate risk criteria, or by both. This makes risk more explicit and understandable by the client.

It is useful to define the elements of decision theory as the ingredients of the theory’s applied form: decision analysis. Table 1 shows my view, which reflects the order in which a decision-analytic model is usually built: first, the *structure* of a decision model and then, its *content*, both requiring judgements from the client about whether to use them.

As in a cookbook, ingredients can be combined in different ways to suit the intended purpose. A five-step generic process that helps to guide the decision analyst is shown in Table 2. Of course, these five steps will be different depending on the type of problem, and there is a different recipe of ingredients and how to combine them for each of six main problem types: evaluate options, allocate resources, bargain and negotiate, choose and decide, manage risk, and revise opinion. I will discuss these later. In practice, iteration between identifying and

Table 1. The 10 Ingredients of Good Decisions

Ingredient	Definition
Structure	
Objectives	The aims or purposes to be achieved
Criteria	Standards against which achievement of the objectives is assessed
Options	Alternatives, decisions, choices, or courses of action for achieving the objectives
Events	Happenings that can influence achievement of the objectives
Outcomes	Ways by which the happenings influence achievement of the objectives
Content	
Consequences	The results or effects of the event’s outcomes
Preference values	Extent to which the consequences are judged to achieve the objectives
Trade-offs	The extent to which more value on one criterion can be balanced by less on another
Probabilities	Degrees of belief about the occurrence of the outcomes
Risk attitude	Extent to which the possibility of harm is judged to be tolerable

defining the two sets of ingredients and following the process typically accompany the model’s construction.

Overall, these 15 ingredients of decision technology, 5 each for structure, content, and process, show that the 3 key ingredients of decision theory (preference values, trade-offs, and probabilities) require at least some of the other ingredients to be incorporated in any realistic, useful decision model. But, which ones are required for what problems (that is, for the decision analyst to determine in the initial stages of interacting with the client)?

The next section considers the social skills required of the decision analyst so that he or she can guide the construction of a decision-analytic model that serves to help the client explore the future.

Social Skills

The above case studies show that working with groups of key players and developing requisite models are effective ways of enabling the accountable decision maker to

resolve his or her unease and ensure that the problem owners are now moving forward in the same direction. But, how is the sense of unease resolved? Much depends on the social skills of the facilitator, which are different from the role of the group’s leader.

The role of the facilitator is to guide the process of constructing a decision model, whereas the leader’s role is to monitor and contribute to the content. But, this requires a particular set of social skills on the part of the facilitator, and after researching the skills of professional facilitators, I discovered the work of Schein (1999). Schein (1999) explains that clients approach management consultants to act in one of three relationships. The first is the expert; the client wants information that the consultant can pass on to them. The second is a doctor-patient relationship; the client wants the consultant to diagnose an organizational problem and prescribe the solution. The third is a process consultancy relationship; the consultant uses his or her skills to help the client solve the

Table 2. The Process for Creating Any Decision Model

Process
1. Consider context. What has given rise to the sense of unease? What is the problem? What aspects of the physical and social environments are relevant? Understanding the context will influence assessment of preference values, trade-off weights, and probabilities.
2. Frame the problem. Can the problem be represented by a decision? Which of the 10 ingredients should be represented in the model? Is the problem more one of resolving uncertainty, one of conflicting values, or a combination of those two main features?
3. Provide content. What information is relevant to judgments of preference values, trade-offs, and probabilities? What expertise is needed for making those judgments, and what potentially different expertise is needed for judging value trade-offs?
4. Explore results. What results might change when considering different assumptions or judgements about the future? Conduct sensitivity analyses showing the extent to which results are robust to differences in experts’ judgements and imprecision in the data.
5. Agree on the way forward. What can be agreed by those engaged in the modeling that would be helpful to the decision maker? Can a narrative be constructed based on the modeling that points to possible ways forward? If not, appraise the accountable person who might do so.

organization's problems. This latter relationship best serves the construction of a requisite decision model.

In the 1980s, I was invited by a large UK organization to conduct a decision conference for them about whether to build a new factory to supplement their existing ones. They had previously spent a substantial amount of money to hire a decision analysis consultancy operating in doctor-patient mode: interviewing key people in the client organization, then leaving to build a decision model, and returning to the client with a PowerPoint presentation that explained the model and showed them the best course of action. The company subsequently told me that the consultants had not adequately tackled the underlying problem, and they hoped I could work *with* them in modeling the real problem.

Fortunately, information gathered by the consultancy group remained alive in the heads of those attending the decision conference, and together, we created a multicriteria model whose weights were conditional on the future direction of the market for their main product, with options of building a new factory, refurbishing an existing one, or making no changes. Sensitivity analyses on the market weights led the group to prefer the refurbishment option, and that was soon implemented. The sense of unease increased with the doctor-patient treatment but dissipated with the process consultancy approach.

Process consultancy is not easy, largely because working *with* a client is a very different type of relationship than working *for* them. In the following paragraphs, I show what Schein's 10 principles mean for decision analysts, though in a slightly different order from shown in his book (Schein 1999).

Always Try to Be Helpful

For decision analysts, this is the most important principle. Whatever you personally feel about the project or any of the participants, you carry on helping them to achieve their objective or primary task. When senior management halved the budget of a major project I was facilitating, I quit in protest. Bad mistake. I should have engaged with the management, saying that our original objective was now compromised; so, could I help them reformulate a more realistic purpose? Big projects rarely move smoothly, I have come to realize, yet it is essential to maintain the helping principle.

It Is the Client Who Owns the Problem and the Solution

Never, ever tell the client what to do, no matter how much they would like you to take on the doctor-patient role. You, the facilitator, are not the accountable decision maker, so if you are operating as a process consultant, you cannot take the monkey off the client's back. Anyway, you do not know enough about the background of the project, the culture of the company, or many other features that are relevant to a solution to take the decision maker's unease onto yourself. You must do your best to help your clients shoulder their responsibilities, but that is it.

Always Stay in Touch with the Current Reality

As empathic human beings, we are easily influenced by what others are feeling, and this is particularly true in a group, especially at the start of a decision conference when decisions, judgements, and assumptions are thrown into the discussion. The facilitator may feel the decision maker's sense of unease and begin to doubt whether any help at all can be made available to the client. Your feeling of helplessness may be telling you what the social reality is in the group. The next principle will help you to sort out what is *your* reality as distinct from the group's reality.

Access Your Ignorance

Access what you know, what you assume you know, and what you do not know. Of course, you cannot possibly know more about your client's problem, especially about its context, but you do know decision analysis. Pushing your feelings away because your client seems to know far more than you about his or her problem will make it more difficult to deploy what you know about decision analysis that might help.

The next two principles are particularly important at the start of a workshop or decision conference and will be easier to implement if a calling note has been sent to all participants well ahead of the actual meeting. The calling note is a letter or email from the client to each participant that explains the administrative arrangements for the meeting (date, time, location, etc.), the purpose of the meeting (its main objective or primary task), a brief description of the anticipated stages in the meeting (not a fixed agenda but stages, such as those shown in Table 2), and how to prepare for it (reflect on what information,

experience, and data could be helpful to achieving the purpose). This ensures that each participant knows why they have been invited and therefore, arrives at the meeting with an expectation that they can make a useful contribution.

Go with the Flow

To understand your client's reality and the culture of the group, ask questions, and actively attend to the answers, responding by following through by asking for elaboration. You start with pure enquiry, waiting, prompting, and listening, just focusing on what the problem is. A little later, ask questions that explore your client's sense of unease by asking what they feel about attempts to deal with the problem and what they are trying to achieve. Once you have begun to see the situation more clearly as an outsider and begun to consider which of the 10 elements of good decisions are relevant, start to ask questions that are sharing your ideas, usually beginning the question with "I wonder if [your idea]," which allows the client to agree or disagree.

Timing Is Crucial

As you become more confident that decision modeling might help, especially if you have worked for organizations in the same field and have accumulated knowledge that could be relevant to this new client, consider asking, "Have you thought about ... ?" But, take care about doing this too early in your interaction because you are then potentially contributing to content, and you do not want to encourage the client to become dependent on you, moving to either the expert or the doctor-patient relationship. By listening carefully, you will be able to judge whether the client might be receptive to a suggestion and take it in.

Everything You Do Is an Intervention

Anything you do or say, including saying nothing, should be in the service of helping the group to move forward without becoming dependent on you or expecting you to do their work for them. Every interaction with the client, even when you are only observing, has consequences for both you and the client, so you must ensure that the consequences contribute to building trust and the helping relationship. Your main interventions include the following (Phillips and Phillips 1993).

- Pacing the task. Set start and stop times, and ensure that all five process steps are completed when the meeting is scheduled to end. As the meeting progresses, the facilitator continuously revises the anticipated time for each stage, informing the group about progress (neither praising nor condemning) with a simple "We're on track" or "We need to speed up" if time is pressing.

- Directing. Participants may allay anxiety by diverting the discussion to something safer. Rather than ignoring or dismissing the comment, the facilitator might say, "Sorry, I'm unclear how this is relevant to our current task, which is to [restate it]." Depending on the response, you may need to bring the discussion back to the difficult issue.

- Questioning. Ask questions, such as "Why do you say that?," "Say more," "What is the evidence?," "What are the implications?," "What might the consequences be?," "Are there any other perspectives?," and so forth. Be sure that you and your client agree about the meaning of the words and jargon. Common terms, like risk, benefit, impact, safety, harm, strategy, vision, and mission, can take on different meanings from one discipline to the next and even within a discipline. For example, we decision analysts usually think of risk as the probability of something happening. But, in healthcare, risk is often seen as the probability only of harm, whereas other disciplines also include the magnitude of harm in the definition (Slovic 1987) and even the number of people who are exposed to the risk.

- Summarizing. Frequent summaries by the facilitator serve to ensure that everyone shares the same understanding of the issues. They also create milestones about progress and if saved on a computer, will help the facilitator in writing the report after the meeting.

- Reflecting. As the facilitator's knowledge of the primary task deepens, it is worth checking that your understanding is accurate. Ask a participant who has stated their position on an issue, "You seem to be saying [restate what you have heard]. Am I hearing you correctly?"

- Handing back in changed form. In effect, that happens as the model is being constructed, and it is possible to show expected or weighted values when only a part of the model has been completed. It was the participants' scores, weights, or probabilities that were

input to the model, which now shows the results but in a changed form as expected values. Sometimes, the facilitator can question assertions like “We’ve already discussed and dismissed that issue” by saying “So are we no longer allowed to discuss it now that it has been raised again?” You may agree with the assertion but add a possible implication that “hands back” the assertion with a new thought that might change the mind of the questioner. This is rather like an architect handing back to a client an agreed-on plan that can now be explored in three dimensions as if walking through the proposed building.

The one intervention that must not be made is to interpret the behavior of individual participants or the group. Decision conferences and model-building workshops are not therapy groups, so any attempt to interpret will be met with anger and even hostility as the facilitator has overstepped his or her limited authority. For example, when training facilitators, I often engage participants in role-playing; I recall one session in which the trainee was having difficulty facilitating the group, and the session ended on a low note at the end of the day. At the start of the next session, I asked the trainee how he had felt about his facilitation the previous day; he severely criticized the group for being so uncooperative, but he had thought about it overnight and now believed he knew what to do next.

He started to redirect the group, but the resistance only increased his anger. I suggested he stop and instead, ask the group what they thought about how we could move forward. Within just a few minutes, he calmed down and worked with the group in a productive way. In reflecting later on the rapid shift to constructive work, he admitted that he had not appreciated how important it was to maintain his impartiality whatever he was feeling.

Be Constructively Opportunistic with Confrontative Interventions

Watch for the moments when the door is partly open and you feel it might be possible to confront the group with something new that could help bring about needed change. For example, in a decision conference attended by a company’s marketers and accountants, the marketers extolled the virtues of a new product, whereas the accountants derided it as being too expensive to develop. The facilitator challenged the group to find a way for the

company to gain the advantages of the new product while also reducing the development cost. After many iterations of the multiattribute model, the group was able to agree on a new option that would share the development cost with an outside organization and split the profits.

I have occasionally used a version of the “double-task” method proposed by Bridger (1990) as a confrontative intervention, which is particularly useful if a verbal fight develops in a group. For example, leaders of teams developing new drugs were gathered for a decision conference to agree how best to fund individual teams within a limited overall budget. Unsurprisingly, each leader defended their own patch as the verbal temperature in the room became increasingly heated, with many conversations going at once. When two people on opposite sides of a large unmovable table stood shaking their fingers at each other, I stood and asked them to sit and for the group to stop talking. “We have gathered today to develop a portfolio of drugs,” I reminded the group. “Let’s call that Task A. I would like you now to stop that task. Instead, I want you to work on Task B, which is to consider how well we are doing on Task A.”

After a brief silence, someone said, rather quietly, “Not very well.” I waited, and gradually, others spoke; a thoughtful discussion ensued. After about 10–15 minutes, I felt that the group had confronted several reasons for their unhelpful behavior, so without commenting on their discussion, I suggested we stop Task B and resume Task A. Immediately, the group took ownership of the task and worked constructively on it.

Now, let us turn to the final two process consultancy principles. They are both about you, the facilitator.

Everything Is Data; Errors Will Always Occur and Are Your Prime Source of Learning

The facilitator can never know enough about the context of a problem to avoid errors in at least some interventions. For example, I helped a private healthcare company wishing to compare the benefit-risk balance of the main imaging agents that enhance the image of a magnetic resonance imaging scan. The experts in a decision conference agreed about the available data on two favorable and six unfavorable events, so scoring was not a problem. But, when they attempted to judge the weights on the eight criteria, they could not agree, arguing that the weights depended on the patient’s age, clinical

characteristics, and many other factors. I asked them to consider a typical patient, so they invented one, enabling them to determine weights for that patient.

I was ready to leave it there, but the experts then questioned the usefulness of the model as it would not apply to all or even some of their patients. Eventually, it dawned on me that I had imported the “average patient” from my previous experience modeling the benefit-risk balance of new drugs for government regulators that approve new drugs. They use statistical data, obtained by pharmaceutical companies, from many patients, mainly means and percentages, representing the average patient as there was substantial variability from one patient to the next, but this was an irrelevant concept for the imaging company. Eventually, the experts developed eight different vignettes of patient types whose characteristics enabled the group to agree on different sets of weights for each vignette. My error was to import my understanding of drug approvals, which applies to populations of patients, into a situation where radiologists make decisions about individual patients.

When in Doubt, Share the Problem

Inevitably, there will be times when you are uncertain about what to do next, and this is true even for the most experienced consultants. Sometimes admitting that to the group will engage their helpfulness, especially if you can start the discussion by suggesting some options. At other times, it may be better to call a five-minute break and discuss your doubt in an aside with the leader.

Especially when I have been working with an organization on several projects, I might eventually feel there is something I am not quite getting right in my facilitation, so I have hired an outside process consultant to explore my own sense of unease. They may well pick up a social reality that defines the culture or potential intelligence of the whole organization (Armstrong 2005), which might be affecting my facilitation without being aware of it.

Guided by these 10 process consultancy principles, the facilitator will build a trusting relationship with the client, which is crucial to ensuring that the construction of a model will help to resolve the client’s sense of unease. As the Irish philosopher O’Neill (2013) has argued, trust is a gift to be earned, which requires the trusted person to be competent, honest, and consistent.

Types of Decision Analysis Models

Six types of decision model that deploy at least one of probabilities, preference values, or trade-offs share a common feature that explains why they can effectively resolve the decision maker’s sense of unease, support change and innovation, and enable strategic planning for the future; they all serve as transitional objects.

Let us now see how that can work differently for each of the six. All these examples applied some version of decision conferencing or a combination with workshops to create the model, with groups of experts and key players participating in the modeling process; each group’s participants were engaged in collaborative discourse, sharing experiences, debating pros and cons, and constructing their preferences (Lichtenstein and Slovic 2006) to the point of substantial agreement for a base case, with disagreements tested in sensitivity analyses. Part 3 of *Decision Analysis for Creating the Future* (Phillips 2025) describes the following case studies in more detail. The next two case studies show the importance of the decision model as a transactional object as it helped two UK governmental agencies make decisions about multimillion-pound major projects.

Evaluate Options

In November 2003, UK Government ministers formed the Committee on Radioactive Waste Management (CoRWM), whose purpose was to recommend a long-term solution for managing the country’s high and intermediate levels of radioactive waste, which had been accumulating for many years and was mainly stored above ground or at shallow depth where it was created. The review of options was to be open, transparent, and inclusive and to “inspire public confidence.”

In 2004, CoRWM commissioned a countrywide survey to obtain people’s opinions and concerns about radioactive waste, but the diversity of views was so wide that committee members felt uneasy about what to do with the findings. In the spring of 2005, CoRWM engaged Catalyze Ltd, the United Kingdom’s decision analysis consulting firm, which was formally associated with the London School of Economics, to explore how to use the survey information.

Many decision conferences with the committee throughout 2005 transformed the survey data into a value tree of nine nonfinancial impacts as objectives and 26 criteria in clusters under the objectives.² In facilitated

workshops, outside experts about the impacts developed fixed value functions as nine-point interval scales, with nine representing the maximum that was technologically feasible today and one as minimally acceptable. The experts also scored the six storage options and three deep disposal options on each of the criteria. Figure 1 shows two examples. When agreement about scores could not be agreed on, a median score was entered for the base case, and high and low scores were recorded for later sensitivity analyses.

Although three subcriteria (monitoring, adaptability, and retrievability) were invoked in the definition of flexibility, after the experts defined points 9 and 1, they then defined point 5 as half the value of the points 1–9 difference in value. (For some criteria, definitions for points 3 and 7 also halved their five-point differences.) Thus, all criteria were constructed to be equal interval value scales. To assess weights, Catalyze devised a simple paired-comparisons approach (Thurstone 1959) comparing cards showing the one to nine swings in value (indicated by the up arrows in Figure 1) on the criteria so that weights could be assessed to establish a common unit across all the criteria.

This process was applied by an outside organization to various stakeholders and members of the public, summarized in a report to CoRWM members. Stakeholders included older and younger citizens, local governments and communities, nongovernmental organizations,

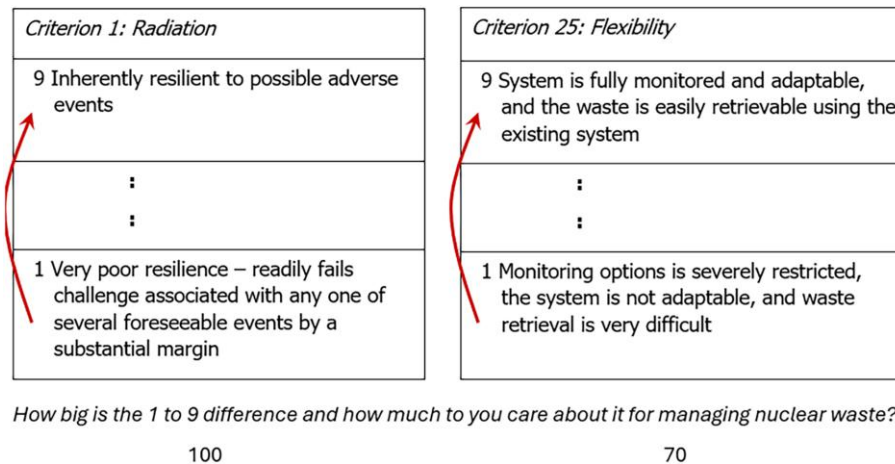
nondevelopmental governmental bodies, learned societies, and the Environment Agency.

A final three-day decision conference facilitated by two members of the CoRWM team and an external member from Quintessa Ltd brought together the experts’ scores along with the average weights from each of the different stakeholder groups to establish a base-case MCDA model. During the consolidation, any disagreements among the CoRWM participants were noted and tested on the third day in sensitivity analyses.

Results consistently favored deep disposal, even for weights from the different stakeholders. When experts’ optimistic scores for burial and pessimistic ones for disposal replaced all the base-case scores, the difference between burial and disposal reduced a modest amount but still favored deep disposal. This was the first time that all nine members of the CoRWM committee were in agreement. One member, Pete Wilkinson, a former chair of UK Greenpeace, had consistently said he would never allow deep disposal because all rock is fissured, so radiation would inevitably reach the biosphere. When I later asked him why he had changed his mind, he replied, “Because all the other options are worse!”

A final meeting of CoRWM, without Catalyze, to conduct a “holistic analysis” and construct a narrative from the decision analysis model led to 15 recommendations sent to the Environment Secretary, David Miliband, in July 2006. On October 25, 2006, he announced that

Figure 1. Participants Compared Two Cards at a Time to Find Which of the 26 Cards Was the Largest Difference That Mattered



Notes. Then, all of the other 25 cards were compared with the 100 card. This is a typical example.

nuclear waste would be buried deep underground and that local councils will be invited to volunteer sites for this multimillion-dollar investment, appointing the Nuclear Decommissioning Authority (NDA) responsible for the process of site selection. Now 19 years later, no site has yet been chosen, and research continues at each of the three locations on the NDA's current short list. An overview of how CoRWM managed to translate a troublesome start into a solid final success can be found in Morton et al. (2009).

The decision model developed by CoRWM played several roles as a transitional object for well over a year. First, it served to summarize progress in constructing an MCDA model in a way that could be understood by an ordinary person at each of three stages: (1) structuring the original public enquiry into well-defined options and criteria, (2) scoring the options by experts chosen for their lived experience and knowledge about the criteria, and (3) weighting the criteria by different stakeholders.

Second, the various workshops and decision conferences were facilitated impartially by group-process-trained specialists, with the developing MCDA model projected to encourage deliberative discourse analysis (Renn 1999). They ensured that the transparency of the MCDA model was maintained throughout its construction, thereby holding committee members' uncertainty as they could see progress and anticipate the next steps.

Third, CoRWM received written feedback from the decision conferences from any person who had observed these open meetings, including from two observers who formally represented the United Kingdom's Department for Environment, Food & Rural Affairs and the National Nuclear Corporation, a nuclear power plant design and construction contractor.

Fourth, the MCDA model held far more information than any one brain can manage without simplifying the problem down to the two or three most important objectives. I have no idea how a "holistic analysis" could have been carried out. That holding function along with the decision conferencing process helped CoRWM members to develop a shared understanding of the issues, create a sense of common purpose, and build commitment to the way forward.

Allocate Resources

Imagine an organization that consists of many departments, each with its own budget. How does the

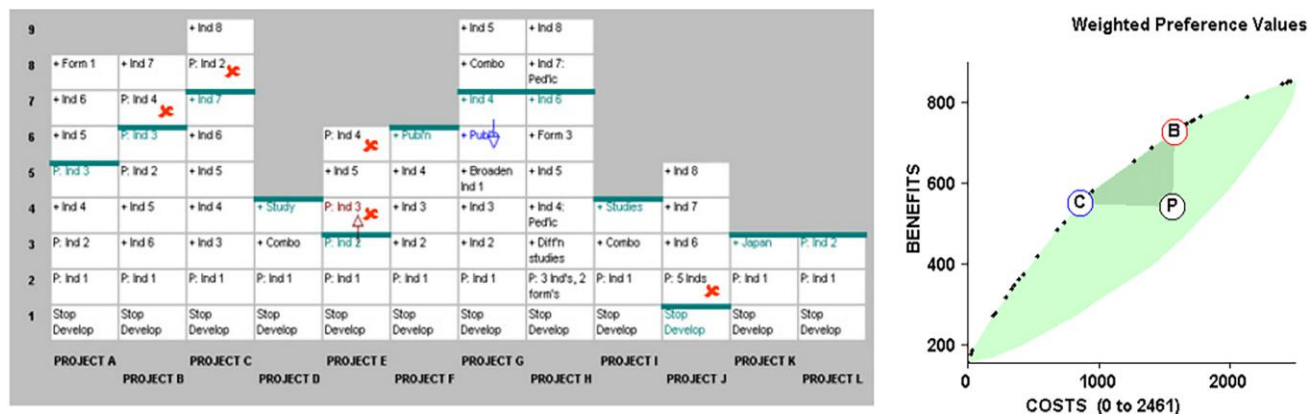
president of the organization, the managing director, allocate resources to the heads of the department, the vice presidents (VPs)? One possibility is at a collegiate meeting, the MD and VPs collectively examining past performance and discussing what could be done better, thereby increasing budgets in areas of greater opportunities along with reductions for areas of lesser opportunities. Reducing is always difficult and can lead to behind-the-scenes "you scratch my back, and I'll scratch yours" or at the very least, a kind of salami slicing, cutting off small funds for some projects and adding other small amounts to other projects over time.

Another approach is portfolio decision analysis, in which each budget category is subjected to an MCDA decision analysis followed by a "merge meeting" where the MCDA models are brought together and a common unit of benefit is created so that the best combinations of options across the different budget categories, a portfolio, can be determined. Money is the resource being allocated, where "best" is defined in corporate finance textbooks as the priority index (PI):

$$PI = \frac{\text{Risk adjusted Benefit}}{\text{Forward cost}}$$

For example, each set of vertical white blocks in the left panel of Figure 2 shows possible options for 12 projects, A through L, for a pharmaceutical company, each a medicinal product. The letter P identified the current plan for each project, numbered as they were originally ordered while the model was created.³ A plus sign represents a new option that would require additional resources. The software calculates a PI for each option in every project and reorders each project's options from level 2 upward in decreasing order of PI as shown by the "cityscape" display of Figure 2. Benefits were the weighted averages of three criteria (meeting unmet medical need, NPV, and future potential) multiplied by the probability that the regulator would approve the product.

The options below the bold horizontal bars in the left panel of Figure 2 collectively maximize the PIs for the same total cost as the current plan. The lightly shaded football shape in the right panel of Figure 2 is an area that contains all possible combinations of options, with the best options indicated by the dots along the efficient frontier at the top. Also shown is the better portfolio for the same cost as the current plan as well as a cheaper

Figure 2. (Left Panel) Options for Each of the 12 Projects

Notes. Options below the bold horizontal line show the better plan. The X's are the plan options that are excluded from the better portfolio. B, better; C, cheaper; Combo, combination of drugs; Diff'n, differentiation; Form, formulation; Ind, indication; P, current plan; Ped'ic, pediatric drug; Pub'n, publication.

portfolio for the same cost as the current plan. Note that five current options, the X's in the left panel of Figure 2, fall outside the better portfolio. Any portfolio within the darker-shaded triangle in the right panel of Figure 2 would be preferred to the current portfolio.

The group of project managers was surprised by this result and wondered what would emerge if the portfolio was driven only by NPV, so weights on the other two criteria were set to zero. The group agreed that the resulting portfolio missed too many drugs developed for meeting unmet medical need. Another run with only the medical need criterion missed too many opportunities for making money, and driving only by future potential was also found to be too future oriented so that might endanger the company's current share price. So, the more balanced portfolio was then somewhat modified to reflect realistic constraints. The director of research and development commented that the cityscape display was the first time he had seen all projects' options in one display, which he found very helpful.

In general, resource allocation models are first created individually by each project team, starting with defining their project followed by a strengths, weaknesses, opportunities, threats (SWOT) analysis to develop new options assuming that more resources are available, scoring the options on the three criteria, and assessing the probability of the drug's approval. Weighting in this type of model is of two types, both within and across criteria,

which require knowledge across the projects, usually assessed by the organization's senior management, as these require trade-offs between the projects and among the three criteria.

Thus, the holding property of the individual project models and bringing them together is the primary function of the final model as a transitional object. For this organization, it kept building over a period of three months. For fewer projects, it can take days or weeks. All the features of MCDA models as transitional objects described in the previous subsection are evident for resource allocation MCDA models plus the advantage of all project team leaders seeing that their projects are fairly considered by senior management and given in-depth consideration.

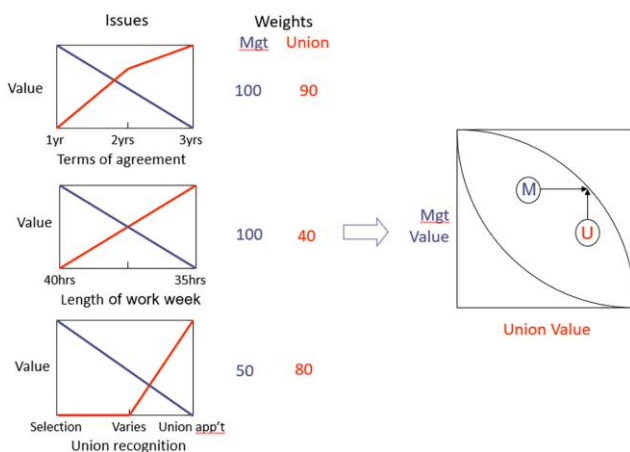
Accepting that it takes time to develop this type of model might tempt the organization to speed up the process by the more usual salami-slicing process. But, I should point out that this project revealed a large discrepancy between the current plan and the better one: a probability-weighted Net Present Value (NPV) of \$1.37 billion. That is easily the equivalent of a blockbuster drug hidden by the previously inefficient method of allocating resources. As an organization consistently applies the MCDA approach to resource allocation, the green area in the right panel of Figure 2 becomes thinner as low-hanging fruits are either accepted or discarded, and the current plan moves close to the efficient frontier (Phillips and Bana e Costa 2007).

Bargain and Negotiate

Negotiating among different parties to arrive at an agreed solution is an obvious candidate for help from decision-analytic modeling. Figure 3 shows an application between management and unions, in which a negotiator hired by a company was helped by decision analysts to build a model of 28 issues, with the two sides taking different stands (options) on the issues. Illustrated here are three of the issues: the terms of the agreement (one, two, or three years), length of the working week (35 or 40 hours), and union recognition (only by National Labor Relations Board election and various other methods or based on card check).

The three issues clearly show many opposing values but also, asymmetries of interest in the top and bottom curves of Figure 3, not just zero-sum relationships (what one side gains the other side loses) as in the length of the work week. Furthermore, each side’s weights reveal that the two sides differ on the relative importance of the issues. Swing weighting the values and summing the products over all combinations of the 28 issues enable all possible combinations of issues to be represented within the football-shaped contract space on the right side of Figure 3.

Figure 3. Three of 28 Issues and the Preference Value Curves for Each Side Showing Their Different Stands on the Issues



Notes. Adding the weighted preference values for each side provides a possible contract, which falls within the area bounded by the concave and convex curves. Negotiations may begin with each side proposing a contract that favors them as indicated by M and U. The model may show a contract on the concave efficient curve that is better than both starting contracts. Mgt, management.

Despite Howard Raiffa’s attempts to explain how this can work for many kinds of negotiation situations, most extensively in his last book with two colleagues (Raiffa et al. 2002), which contains the theory, practice, and case studies, this methodology is underused compared with all the other uses of decision analysis. DDI applied it extensively in the 1960s and 1970s for the Panama Canal treaty negotiations, for base rights and arms negotiation for the Pentagon and the State Department, and in various cold war negotiations. A particularly notable success was in establishing worldwide standards for oil tankers to reduce pollution of land and seas, which involved the U.S. Coast Guard in cooperation with the State Department and partners from many other countries supported by President Carter in 1977 and 1978 (Ulvila and Snider 1980).

Negotiation models are dynamic in the sense that they change during the bargaining process as the sides become clearer about each other’s values. As the models change and identify asymmetry of interests and as participants revise the issues and the stands on the issues, the model holds the substantial complexity, enabling trade-offs to be examined and iterating through possible treaties, all the while containing the current understanding of the problem. A final agreement was developed that was acceptable to all the participating countries. More than 750 tanker oil spills worldwide in the 1970s were reduced to just 17 in the 2020s.

Overall, Cam Peterson reported that DDI’s experience with bargaining and negotiation models helped negotiators prepare in advance, facilitate coordination between negotiators and their organization, improve implementation, anticipate positions of other delegates, and improve communications within the negotiation team. Thus, as transitional objects, they make it possible for all parties to continue engaging in a process of change that is perceived to be fair and unbiased, allowing negotiations to be flexible and to widen the scope of issues, and that encourages the development of new alternatives.

Choose and Decide

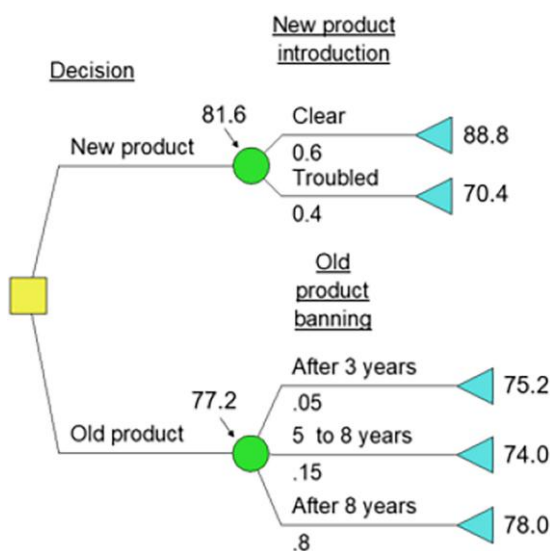
Return now to the “new product” model mentioned at the start of this paper. Recall that the MD of a manufacturing company was concerned about introducing a new product to replace its highest-selling product because a recent safety issue might cause the government to ban it. He had presented versions of the new

product to his board for their approval, but they had been rejected 11 times in as many months. I helped him develop a decision tree; Figure 4 shows the first part of the tree, with probabilities of outcomes, whereas all other figures are expected values. Additional uncertain events are attached to each of the triangles in Figure 4.⁴

When the decision tree was rolled back, he could see that the expected value of the new product was £81.6 and that the expected value of the old product was £77.2. This is not very different, so no wonder he and the board were unable to make a decision. However, he pointed to that key uncertainty, introduction of the new product clear or troubled, and said that the expected value at clear was £88.8 million but with a probability of only 0.6. He asked if he would then be justified in spending up to the difference, £11.6 million (88.8 minus 77.2), to ensure a clear introduction. “Yes,” I replied, “but can you be so sure?” “No, but I can get close even by spending less than half that amount.”

When board members were shown the original decision tree in a decision conference, each of three members disagreed with several probabilities, so their judgements were subjected to sensitivity analyses that collectively changed the expected values to favor the old product.

Figure 4. The Beginning of the New Product Decision Tree



Notes. The values on the far right are the expected values from rolling the tree back from the preference values at their end positions. The final rollback provides the weighted preference values shown at the circles.

However, any two combinations favored the new product; because none of the board member believed all three, they approved the new product, and the MD hired a design consultant to redesign the product (Phillips 1982).

The point here is that the explicit structure of the problem had enabled the MD and board to hold the results of the 50 paths in the decision tree while they debated the probabilities and values. In particular, the MD could then see how to figure out the value of better information, which would allow him to take control of the risk inherent in the new product. He could now tolerate the remaining risk of the new project. This transitional object had revealed something that nobody had realized, and sensitivity analyses showed that even if some changes were made, the new product alternative remained the best; consensus about the inputs was not required. A requisite model was realized. What had originally been implicit and subjective was made explicit, quantitative, and less scary.


Manage Risk

In June 2009, the World Health Organization declared that the H1N1 virus was a pandemic, causing influenza. It was not clear how serious the pandemic would be, but the numbers of people with the disease were rising. This presented a challenge to the European Medicines Agency (EMA) about how soon to approve vaccines whose full trials were incomplete. An early approval without the full data about safety and effectiveness could be criticized, but so could a later trial after the vaccine data were complete because deaths could have been avoided. The new benefit-risk team at the EMA recognized that this presented an opportunity for decision-analytic modeling. A small team of experts from the EMA staff met for a decision conference to see if an “explicit modeling-based approach to regulatory decision making could improve understanding, transparency, and communication” of the results (Phillips et al. 2013).

The group quickly constructed the decision tree shown in Figure 5.⁵ The approval alternatives of the end of September or October were followed by outcomes of the key uncertainties. Fully attaching those events led to 24 future scenarios, 12 for each alternative, describing the deaths and serious diseases (DSDs) if 500,000 people were vaccinated in Europe. Thus, serious risks attended both options.

Deep discussion among participants about similar past worldwide pandemics suggested what might happen

Figure 5. Decision Tree About When to Approve the New H1N1 Vaccines, Taking Account of the Uncertainty in August 2009 of the Subsequent Disease Seriousness and the Efficacy and Safety of the Vaccines



Note. DSD, death and serious disease.

and led them to formulate 11 assumptions about the future as well as probabilities that were often conditional on other events or independent. Subsequently, the model was slightly refined by a subset of participants, and a third one-hour telephone meeting engaged an expert from the European Centre for Disease Prevention and Control to help refine some of the inputs.

When the completed tree was rolled back, it showed expected DSDs for September to be 216,500 and for October to be 291,547. Thus, an expected (probability-weighted) 75,047 DSDs would be prevented by approving the vaccines at the earlier date. Many sensitivity analyses showed that probabilities about disease efficacy and safety only affected the number of lives saved for extreme probabilities, which nobody believed. However, if the probability of the disease being moderate, given delay to the end of October, were to exceed 0.84, then waiting until October for authorization could be justified, but nobody at the first meeting believed it would be that high. The vaccines were approved at the end of September.

And, what happened? The pandemic turned out to be milder compared with the 1918 virus infection. For this case, the transitional object only helped the EMA insofar as it could now justify an early approval of the vaccines. But, I have included it here because it shows that for any subsequent pandemic, decisions to be made by businesses, governments, and policymakers could extend beyond the issues of concern to epidemiologists and statisticians and could include future consequences for the many socioeconomic and health impacts and their associated decisions (Dillon et al. 2023) in the country affected.

It is transitional objects that can hold the uncertainties and anxieties that decision makers feel when faced with life-threatening events and when they must act despite a paucity of data. Models of possible future consequences

that engage informed, expert human judgement are capable of supporting decision and policymakers about what we should do now to minimize future risks.

Revise Opinion

At an early meeting of Drug Science, I was asked to present an overview of Bayesian statistics. When I showed graphically how probability distributions can be revised when new information is made available, one member observed that this is what we doctors do whenever we see a patient. We revise our opinions as the patient reports activities relevant to their medical condition and their symptoms, and we continue that process as we examine the patient and gather more evidence. Medical training in the United Kingdom apparently does not mention Bayesian ideas or the research showing that more typically than not, most people do not revise probability beliefs very much when new information is presented to them compared with revision with Bayes' rule (Phillips et al. 1966). But, they are all informed about classical statistics as they must be able to interpret random controlled trials (RCTs). However, practicing physicians mainly now recall only that a p -level less than 0.05 is required for declaring a statistically significant result and that $p < 0.01$ is even better. And, sample sizes must be big, preferably larger than 100 (or variations on these rules of thumb).

I was reminded of this when one of Professor David Nutt's PhD students, Rayyan Zafar, started examining the effect of medical cannabis on children with chronic epilepsy. He had observed that 10 children's epileptic seizures had substantially declined in number within the month after taking medical cannabis and had presented before-and-after bar graphs, one pair for each child. This seemed a perfect case to illustrate and report the data using Bayesian revision of the proportion of children whose seizures become fewer. At this early point, $N = 10$,

and seizures had substantially reduced or been eliminated in all 10 children.

A plausible Bayesian analysis might start with a non-informative prior about the proportion, θ , of severely epileptic children who benefit from treatment with medical cannabis. A Beta distribution is conjugate to a Bernoulli process (with successes and failures), which means that both prior and posterior are Beta distributions with parameters a and b . With both a and b set to one, the Beta is simply a constant 1 over values of θ from 0 to 1.0 as shown in Figure 6.⁶

The mean of any Beta distribution is $a/(a + b)$, and applying Bayes' rule simply adds successes (S) and failures (F) to the prior parameters. Thus, the mean of the posterior is $(S + 1)/(S + F + 2) = 11/12 = 0.92$. But, this is not a random trial with a small number of patients, so no reputable medical journal would accept the finding. Rayyan continued to treat children whose parents he had been contacting through two child epileptic charities. Again, all 10 improved. That brought the mean up to $21/22 = 0.95$, and coincidentally, the posterior distribution gives a 95% credible interval that the true proportion is between 0.86 and 1.00. The paper was finally published on the condition that the first 10 be considered a field study and that the second 10 be considered a confirmative study (Zafar et al. 2021).

This was, of course, a humane decision as seizures per month reduced by an average of 82.4%, which is a massive improvement for epilepsy, so withholding the treatment from a control group in an RCT would be seen as

inhumane. Furthermore, for a prescribing physician, to be 95% sure of an improvement or at least more than 86% sure provides information that is directly relevant to a decision, which is not true of significance levels and confidence intervals, which are both statements about the data, not about the medical condition.

For Rayyan, the Bayesian model served as a transitional object by encouraging him to continue with his PhD research gathering more patients, knowing that medical cannabis was almost certain to help severely epileptic children, despite critics saying that observational trials would require many more patients to be certain (Rawlins 2008).

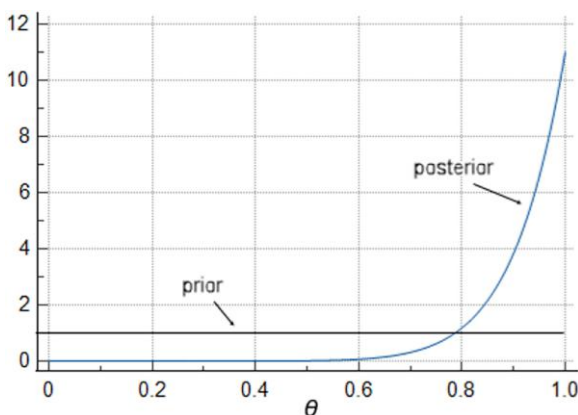
Discussion

This paper has traced Savage's 1952 small world of individual decision theory through Raiffa and Schlaifer's decision analysis in the 1960s as a way of modeling uncertainty about the future to the recognition in the 1980s by Keeney and Raiffa that multiple objectives and trade-offs must also be considered by the decision maker. All decision analysis models now include multiple criteria, many uncertain events, or both, with acceptance today that this technical discipline is now well established as a sociotechnical system because technical reality and social reality interact, affecting the accountable decision maker as well as organizational and societal life.

This came about as applications of decision analysis to real-world problems and issues showed that models only needed to be "good enough" or requisite: that is, sufficient in form and content to resolve the issues at hand and later recognize that the model resolved the decision maker's sense of unease about how to create a better future. We now know that requisite models serve as transitional objects, containing the important elements of good decisions and enabling the client and key players to hold their sense of unease and the anxiety it creates about what to do next while using the model to explore different scenarios about the future and better understand how imprecision in the data and disagreements among the experts can be managed.

The decision analyst acts as a process consultant, exploring the client's sense of unease, and builds with the client a model that will enable exploring possible futures. In a sense, providing technology that is a simulator and handing back to the client in changed form (weighted and expected preference values) coherent

Figure 6. A Uniform Prior with Parameters (1, 1) Is Revised by 10 Successes and No Failures to Give This Posterior Probability Density Function (11, 12)



results about how decisions (alternatives and actions), judgments (assumptions and opinions), and preferences (likes and dislikes) can impact the future.

We have seen that decision analysis provides 10 ingredients, 5 representing the structure of a model, objectives, criteria, options, events, and outcomes and 5 providing content, consequences, preference values, trade-offs, probabilities, and risk attitude. The process for selecting the right ingredients for any problem and combining them can be accomplished in five steps that consider context, frame the problem, provide content, explore results, and agree on the way forward.

Those 10 ingredients and the five-step process of assembling them for use are similar to Problem, Objectives, Alternatives, Consequences, Tradeoffs, Uncertainty, Risk Tolerance, Linked Decisions, with PrOACT as characteristic of all models and URL as characteristic of uncertainty and linked decisions (Hammond et al. 1999). However, my split between structure and content ingredients is the order in which they are created, with structure first and then, content. Thus, replicating a model for a different context often maintains the structure, whereas content will be different. An example is the 2010 model about the harm of psychoactive drugs in the United Kingdom (Nutt et al. 2010), which has been replicated in the European Union (van Amsterdam et al. 2015), Australia (Bonomo et al. 2019), New Zealand (Crossin et al. 2023), and the United States and Canada (whose reports have not yet been submitted for publication). Most of the 16 criteria have been applied in all these countries, with some modifications for the country, whereas about half or more of the 20 drugs are identical. The main differences are in the scoring of the drugs and the weighting of the harm criteria.

Creating a requisite model is an art that depends on the social skills of the decision analyst, which are rarely discussed in textbooks about decision analysis. Applying Edgar Schein's 10 principles of process consultancy when working with a client goes a long way toward providing the skills that will develop the decision analyst's trustworthiness as the helping relationship grows. Along the way, the decision analyst might need to remind the group that nothing is yet set in concrete; the result is merely a base-case model, which will be subjected to many sensitivity analyses.

The final section of the paper describes the six main types of decision models that have been developed and applied: evaluate options, allocate resources, bargain

and negotiate, choose and decide, manage risk, and revise opinion. The initial structuring of any of the six models relies on the help of the decision analyst working with the client. Mainly for that reason, decision analysis models are unique to the situation at hand, largely because they are contingent on the context.

Conclusion

Whereas many decision analysts use their expertise to diagnose a client's problem and prescribe an answer, this paper argues for a process consultancy approach: that is, a problem-solving approach to better understand the client's sense of unease and work with the client to develop a model that can serve as a transitional object enabling the client to safely explore alternative futures. This process will impact the client's risk attitude, making it possible to move forward confidently, despite there being no single correct answer. At no point does the decision analyst suggest to the client what they should do.

Acknowledgments

The author thanks the associate editor and two referees, whose suggested revisions have greatly improved the paper.

Endnotes

¹ This was a 55-pound computer with a keyboard, 5-inch cathode ray tube display, and tape drive that was programmable in either Basic or A Programming Language with 64 K (not a mistype) of memory.

² We used Hiview3, which is available at <http://www.catalyzeconsulting.com/downloads/executables/Hiview-3.2.0.7.exe>. The program is now free and can be unlocked with the security code 10741010EHDD.

³ We used Equity3, which can be obtained at <https://www.catalyzeconsulting.com/downloads/executables/equity-3.4.0.6.exe>. The program is now free and can be unlocked with the security code 10061602KCDZ.

⁴ They are displayed here using DPL9, which is available at www.syncopation.com.

⁵ This is shown here with DPL, but originally, it was shown with Excel Add-in TreePlan software, which is available at <https://treeplan.com/download/>.

⁶ See <https://www.medcalc.org/manual/beta-distribution-functions.php>.

References

- Armstrong D (2005) *Organization in the Mind: Psychoanalysis, Group Relations, and Organizational Consultancy* (Routledge, London).
- Bonomo Y, Norman A, Biondo S, Bruno R, Daghli M, Dawe S, Egerton-Warburton D, et al. (2019) The Australian drug harms ranking study. *J. Psychopharmacology* 33(7):759–768.
- Bridger H (1990) Courses and working conferences as transitional learning institutions. Trist E, Murray H, eds. *The Social Engagement of Social Science, Volume I: A Tavistock Anthology: The Socio-*

- Psychological Perspective* (University of Pennsylvania Press, Philadelphia), 221–245.
- Bridger H (2001) Foreword. Amado G, Ambrose A, eds. *The Transitional Approach to Change* (Routledge, London), xi–xiv.
- Burnstein E, Vinokur A (1975) What a person thinks upon learning he has chosen differently from others: Nice evidence for the persuasive-arguments explanation of choice shifts. *J. Experiment. Soc. Psych.* 11(5):412–426.
- Crossin R, Cleland L, Wilkins C, Rychert M, Adamson S, Potiki T, Pomerleau AC, et al. (2023) The New Zealand drug harms ranking study: A multi-criteria decision analysis. *J. Psychopharmacology* 37(9):891–903.
- Dillon RL, Bier VM, John RS, Althenayyan A (2023) Closing the gap between decision analysis and policy analysts before the next pandemic. *Decision Anal.* 20(2):109–132.
- Dyer JS, Sarin RK (1979) Measurable multiattribute value functions. *Oper. Res.* 27(4):810–822.
- Falk A, Becker A, Dohmen T, Enke B, Huffman D, Sunde U (2018) Global evidence on economic preferences. *Quart. J. Econom.* 133(4):1645–1692.
- Fishburn PC (1972) *Mathematics of Decision Theory* (De Gruyter Mouton, Berlin).
- Hammond JS, Keeney RL, Raiffa H (1999) *Smart Choices: A Practical Guide to Making Better Decisions* (Harvard University Press, Cambridge, MA).
- Keen PGW, Morton MSS (1978) *Decision Support Systems: An Organizational Perspective* (Addison-Wesley Publishing, Reading, MA).
- Keeney RL (1992) *Value-Focused Thinking: A Path to Creative Decision-making* (Harvard University Press, Cambridge, MA).
- Keeney RL, Raiffa H (1976) *Decisions with Multiple Objectives: Preferences and Value Tradeoffs* (John Wiley, New York).
- Keeney RL, Raiffa H (1993) *Decisions with Multiple Objectives: Preferences and Value Tradeoffs* (Cambridge University Press, Cambridge, UK).
- Lichtenstein S, Slovic P, eds. (2006) *The Construction of Preference* (Cambridge University Press, Cambridge, UK).
- Morton A, Airoidi M, Phillips L (2009) Nuclear risk management on stage: The UK's Committee on Radioactive Waste Management. *Risk Anal.* 29(5):764–779.
- Nutt DJ, King LA, Phillips LD (2010) Drug harms in the UK: A multicriteria decision analysis. *Lancet* 376(9752):1558–1565.
- O'Neill O (2013) What we don't understand about trust. Accessed June 5, 2025, https://www.ted.com/speakers/onora_o_neill.
- Phillips LD (1982) Requisite decision modelling: A case study. *J. Oper. Res. Soc.* 33(4):303–311.
- Phillips LD (1984) A theory of requisite decision models. *Acta Psychologica* 56(1–3):29–48.
- Phillips LD (2007) Decision conferencing. Edwards W, Miles RF, von Winterfeldt D, eds. *Advances in Decision Analysis: From Foundations to Applications* (Cambridge University Press, Cambridge, UK), 375–399.
- Phillips LD (2025) *Decision Analysis for Creating the Future* (Cambridge University Press, Cambridge, UK).
- Phillips LD, Bana e Costa CA (2007) Transparent prioritisation, budgeting and resource allocation with multi-criteria decision analysis and decision conferencing. *Ann. Oper. Res.* 154(1):51–68.
- Phillips LD, Phillips MC (1993) Facilitated work groups: Theory and practice. *J. Oper. Res. Soc.* 44(6):533–549.
- Phillips LD, Hays WL, Edwards W (1966) Conservatism in complex probabilistic inference. *IEEE Trans. Human Factors Electronics* HFE-7(1):7–18.
- Phillips LD, Fasolo B, Zafiropoulos N, Eichler H-G, Ehmann F, Jekerle V, Kramarz P, Nicholl A, Lönnngren T (2013) Modelling the risk-benefit impact of H1N1 influenza vaccines. *Eur. J. Public Health* 23(4):674–678.
- Raiffa H (1968) *Decision Analysis* (Addison-Wesley, Reading, MA).
- Raiffa H, Schlaifer R (1961) *Applied Statistical Decision Theory* (Harvard University Press, Cambridge, MA).
- Raiffa H, Richardson J, Metcalfe D (2002) *Negotiation Analysis: The Science and Art of Collaborative Decision Making* (The Belknap Press of Harvard University Press, Cambridge, MA, and London).
- Rawlins M (2008) De testimonio: On the evidence for decisions about the use of therapeutic interventions. *Lancet* 372(9656):2152–2161.
- Renn O (1999) A model for an analytic-deliberative process in risk management. *Environ. Sci. Tech.* 33(18):3049–3055.
- Savage LJ (1954) *The Foundations of Statistics* (Wiley, New York).
- Schein EH (1999) *Process Consultation Revisited: Building the Helping Relationship* (Addison-Wesley, Reading, MA).
- Schlaifer R (1969) *Analysis of Decisions Under Uncertainty* (McGraw-Hill, New York).
- Simon HA (1955) A behavioral model of rational choice. *Quart. J. Econom.* 69(1):99–118.
- Slovic P (1987) Perception of risk. *Science* 236(4799):280–285.
- Thurstone LL (1959) *The Measurement of Values* (University of Chicago Press, Chicago).
- Tversky A, Kahneman D (1981) The framing of decisions and the psychology of choice. *Science* 211(30):453–458.
- Ulvila JW, Snider WD (1980) Negotiation of international oil tanker standards: An application of multiattribute value theory. *Oper. Res.* 28(1):81–96.
- van Amsterdam J, Nutt D, Phillips L, van den Brink W (2015) European rating of drug harms. *J. Psychopharmacology* 29(6):655–660.
- von Neumann J, Morgenstern O (1947) *Theory of Games and Economic Behavior*, 2nd ed. (Princeton University Press, Princeton, NJ).
- Zafar R, Schlag A, Phillips L, Nutt DJ (2021) Medical cannabis for severe treatment resistant epilepsy in children: A case-series of 10 patients. *BMJ Paediatrics Open* 5(1):e001234.

Lawrence D. Phillips is an emeritus professor of decision sciences at the London School of Economics and Political Science and a Director of Facilitations Ltd. He applies decision analysis as a process consultant, assisting individuals and teams dealing with complex issues to arrive at an agreed way forward. In November 2005, the Decision Analysis Society of INFORMS awarded him the Frank P. Ramsey medal for distinguished contributions to decision analysis.