CHAPTER 1

Introduction: Five Breakthroughs in Decision and Risk Analysis

Louis Anthony (Tony) Cox, Jr.
Cox Associates, NextHealth Technologies, University of Colorado-Denver, Denver, CO, USA

This book is about breakthroughs in decision and risk analysis—new ideas, methods, and computational techniques that enable people and groups to choose more successfully when the consequences of different choices matter, yet are uncertain. The twentieth century produced several such breakthroughs. Development of subjective expected utility (SEU) theory combined with Bayesian statistical inference as a model of ideal, rational decision-making was among the most prominent of these. Chapter 2 introduces SEU theory as a point of departure for the rest of the book. It also discusses more recent developments—including prospect theory and behavioral decision theory—that seek to bridge the gap between the demanding requirements of SEU theory and the capabilities of real people to improve their decision-making. Chapters 5 and 8 address practical techniques for improving risky decisions when there are multiple objectives and when SEU cannot easily be applied, either because of

Breakthroughs in Decision Science and Risk Analysis, First Edition.
Edited by Louis Anthony Cox, Jr.

uncertainty about relevant values, causal models, probabilities, and consequences; or because of the large number and complexity of available choices.

HISTORICAL DEVELOPMENT OF DECISION ANALYSIS AND RISK ANALYSIS

Perhaps the most audacious breakthrough in twentieth-century decision analysis was the very idea that a single normative theory of decision-making could be applied to all of the varied risky decisions encountered in life. It may not be obvious what the following problems, discussed in ensuing chapters, have in common:

- *Investment decisions:* How should investors allocate funds across investment opportunities in a financial portfolio? (Chapter 3)
- *Operations management decisions:* How should a hospital emergency room be configured to make the flow of patients as easy and efficient as possible? How should an insurance company staff its claims-handling operations? (Chapter 3)
- *Inventory management and retail decisions:* How much of an expensive, perishable product should a business buy if demand for the product is uncertain? (Chapter 4)
- *Trial evaluation and selection decisions:* How much trial, testing, and comparative evaluation should be done before selecting one of a small number of costly alternatives with uncertain consequences, for example, in choosing among alternative new public policies, consumer or financial products, health care insurance plans, research and development (R&D) projects, job applicants, supply contracts, locations in which to drill for oil, or alternative drugs or treatments in a clinical trial? (Chapter 4)
- *Adversarial risk management decisions:* How should we model the preferences and likely actions of others, in order to make effective decisions ourselves in situations where both their choices and ours affect the outcomes? (Chapters 2, 5, 9, and 10)
- *Regulatory decisions:* When experimentation is unethical or impractical, how can historical data be used to estimate and compare the probable consequences that would be caused by alternative choices,

such as revising versus maintaining currently permitted levels of air pollutants? (Chapter 6)

- *Learning how to decide in uncertain environments:* Suppose that not enough is known about a system or process to simulate its behavior. How can one use well-designed trial-and-error learning to quickly develop high-performance decision rules for deciding what to do in response to observations? (Chapters 4 and 7)
- *Medical decision-making:* How should one trade off the ordinary pleasures of life, such as consumption of sugar-sweetened drinks, against the health risks that they might create (e.g., risk of adult-onset diabetes)? More generally, how can and should individuals make decisions that affect their probable future health states in ways that may be difficult to clearly imagine, evaluate, or compare? (Chapter 8)

That the same basic ideas and techniques might be useful for decision-making in such very different domains is a profound insight that might once have excited incredulity among experts in these fields. It is now part of the canon of management science, widely taught in business schools and in many economics, statistics, and engineering programs.

Decision analysis views the "success" of a decision process in terms of the successes of the particular decisions that it leads to, given the information (usually incomplete and possibly incorrect or inconsistent) that is available when decisions must be made. The "success" of a single choice, in turn, can be assessed by several criteria. Does it minimize expected post-decision regret? Is it logically consistent with (or implied by) one's preferences for and beliefs about probable consequences? In hindsight, would one want to make the same choice again in the same situation, if given the same information? The giants of twentieth-century decision theory, including Frank Ramsey in the 1920s, John von Neumann in the 1940s, and Jimmy Savage in the 1950s, proved that, for perfectly rational people (*homo economicus*) satisfying certain mathematical axioms of coherence and consistency (i.e., complete and transitive preference orderings for outcomes and for probability distributions over outcomes), all of these criteria prescribe the same choices. All imply that a decision-maker should choose among risky prospects (including alternative acts, policies, or decision rules with uncertain consequences) as if she were maximizing subjective expected utility (SEU). Chapters 2 and 7 introduce SEU theory and some more recent alternatives. Decision-making processes and environments that encourage high-quality decisions as judged by one of these criteria will also promote the rest.

However, real people are not perfectly rational. As discussed in Chapter 2, *homo economicus* is a fiction. The prescriptions of decision theory are not necessarily easy to follow. Knowing that SEU theory, the long-reigning gold standard for rational decision-making, logically implies that one *should* act as if one had coherent (e.g., transitive) preferences, and clear subjective probabilities are cold comfort to people who find that they have neither. These principles and limitations of decision theory were well understood by 1957, when Duncan Luce and Howard Raiffa's masterful survey *Games and Decisions* explained and appraised much of what had been learned by decision theorists, and by game theorists for situations with multiple interacting decision-makers. Chapter 2 introduces both decision theory and game theory and discusses how they have been modified recently in light of insights from decision psychology and behavioral economics.

During the half-century after publication of *Games and Decisions*, a host of technical innovations followed in both decision analysis and game theory. Decision tree analysis (discussed in Chapters 8 and 10) was extended to include Monte Carlo simulation of uncertainties (see Chapter 3). Influence diagrams were introduced that could represent large decision problems far more compactly than decision trees, and sophisticated computer science algorithms were created to store and solve them efficiently. Methods of causal analysis and modeling were developed to help use data to create risk models that accurately predict the probable consequences of alternative actions (see Chapter 6). Markov decision processes for dynamic and adaptive decision-making were formulated, and algorithms were developed to adaptively and robustly optimize decision rules under uncertainty (see Chapter 7). SEU theory was generalized, e.g., to allow for robust optimization with ambiguity aversion when probabilities are not well known. Practical constructive approaches were created for structuring and eliciting probabilities and utilities, as discussed and illustrated in Chapters 5, 8, and 10.

These technical developments supported a firmly founded discipline of applied decision analysis, decision aids, and decision support consulting. The relatively new discipline of applied decision analysis, developed largely from the 1960s on, emphasized structuring of decision problems (especially, identifying and solving the right problem(s)); clearly separating beliefs about facts from values and preferences for outcomes; eliciting or constructing well-calibrated probabilities and coherent utilities; presenting decision recommendations, together with sensitivity and uncertainty analyses, in understandable ways that decision-makers find useful; assessing value of information and optimal timing of actions; and deliberate, careful learning from results, for both individuals and organizations. The 1976 publication of the landmark *Decisions with Multiple Objectives: Preferences and Value Tradeoffs* by Ralph Keeney and Howard Raiffa summarized much of the state of the art at the time, with emphasis on recently developed multiattribute value and utility theory and methods. These were designed to allow clearer thinking about decisions with multiple important consequence dimensions, such as costs, safety, profitability, and sustainability. Chapters 5, 8, and 10 review and illustrate developments in elicitation methods and multiattribute methods up to the present.

While decision analysis was being developed as a prescriptive discipline based on normative theory (primarily SEU theory), an increasingly realistic appreciation of systematic "heuristics and biases" and of predictable anomalies in both laboratory and real-world decision-making was being developed by psychologists such as Amos Tversky, Daniel Kahneman, Paul Slovic, and Baruch Fischhoff, and by many other talented and ingenious researchers in what became the new field of behavioral economics. Chapter 2 introduces these developments. Striking differences between decision-making by idealized, rational thinkers (*homo economicus*) and by real people were solidly documented and successfully replicated by different teams of investigators. For example, whether cancer patients and their physicians preferred one risky treatment procedure to another might be changed by presenting risk information as the probability of survival for at least 5 years instead of as the probability of death within 5 years—two logically equivalent descriptions (gain faming vs. loss framing) with quite different emotional impacts and effects on decisions. Many of these developments were reflected in the 1982 collection *Judgment under Uncertainty: Heuristics and Biases*, edited by Kahneman, Slovic, and Tversky. Chapter 10 summarizes key insights from the heuristic-and-biases literature in the context of eliciting expert judgments about probabilities of adversarial actions.

Twenty-five years later, the 2007 collection *Advances in Decision Analysis: From Foundations to Applications*, edited by Ward Edwards, Ralph Miles, and Detlof von Winterfeldt, took stock of the thriving and increasingly well-developed field of decision analysis, which now integrated both normative (prescriptive) theory and more descriptively realistic considerations, e.g., using Kahneman and Tversky's prospect theory. This collection looked back on decades of successful developments in decision analysis, including the field's history (as recalled by founding luminaries, including Ron Howard of Stanford and Howard Raiffa of the Harvard Business School), surveys of modern progress (including influence diagrams, Bayesian network models, and causal networks), and important practical applications, such as to engineering and health and safety risk analysis, military acquisitions, and nuclear supply chain and plutonium disposal decisions.

OVERCOMING CHALLENGES FOR APPLYING DECISION AND RISK ANALYSIS TO IMPORTANT, DIFFICULT, REAL-WORLD PROBLEMS

Despite over five decades of exciting intellectual and practical progress, and widespread acceptance and incorporation into business school curricula (and into some engineering, statistics, mathematics, and economics programs), decision analysis has limited impact on most important real-world decisions today. Possible reasons include the following:

- *Many real-world problems still resist easy and convincing decision-analytic formulations.* For example, a dynamic system with random events (i.e., patient arrivals, departures, and changes in condition in a hospital ward) with ongoing opportunities to intervene (e.g., by relocating or augmenting staff to meet the most pressing needs) cannot necessarily be represented by a manageably small decision tree, influence diagram, Markov decision process, or other tractable model—especially if the required transition rates or other model parameters are not known and data from which to estimate them are not already available. Chapters 3, 4, 6, and 7 present breakthroughs for extending decision and risk analysis principles to such realistically complex settings. These include increasingly well-developed *simulation–optimization* methods for relatively well-characterized systems (Chapter 3) and *adaptive learning*, statistical methods for estimating causal relations from data (Chapter 6), model ensemble, and robust optimization methods for settings where not enough is known to create a trustworthy simulation model (Chapters 4 and 7).

- *It has often not been clear that individuals and organizations using formal decision analytic models and methods outperform and outcompete those who do not.* Chapters 4, 6, and 7 emphasize methods for causal analysis, adaptive learning, and empirical evaluation and comparison of alternative choices. These methods can help decision-makers make choices that demonstrably outperform (with high statistical confidence) other available choices in a wide variety of practical applications.
- *While decision analysts excel at distinguishing clearly between matters of fact and matters of preference, real-world decision-makers often prefer to fall back on judgments that conflate the two*, perhaps feeling that no matter what academic theory may say, such holistic judgments give more satisfactory and trustworthy recommendations than calculations using hypothetical (and not necessarily clearly perceived or firmly believed) subjective utilities and probabilities. (This tendency may perhaps explain some of the popularity of simplistic decision aids that use ill-defined concepts, such as "relative importance" of goals or attributes, without clear definition of what "relative importance" means and of how it should reflect interdependencies.) Too often, there is simply no satisfactory way to develop or elicit credible, defensible, widely shared probabilities and utilities for situations or outcomes that are novel, hard to imagine, or controversial. Chapters 5 and 8–10 discuss innovations for alleviating this problem with new methods for eliciting and structuring utilities, value trade-offs, and probabilistic expert beliefs.
- *Most real-world decisions involve multiple stakeholders, influencers, and decision-makers,* but SEU is preeminently a theory for single-person decisions. (Extensions of SEU to "social utility" for groups, usually represented as a sum of individual utilities, can certainly be made, but an impressive list of impossibility theorems from collective choice theory establish that *homo economicus* will not necessarily provide the private information needed for collective choice mechanisms to produce desirable, e.g., Pareto-efficient, outcomes.) Less theoretically, the notorious Prisoner's Dilemma, discussed in Chapter 2, illustrates the tension between individual and group rationality principles. In the Prisoner's Dilemma, and in many other situations with externalities, individuals who make undominated or otherwise

individually "rational" choices will thereby collectively achieve Pareto-dominated outcomes (a hallmark of collectively suboptimal choice), meaning that everyone would have been better off if all had made different (not "rational") choices. Chapter 2 discusses both classical game theory and its behavioral modifications to better apply to real people, who often cooperate far better than theories for merely "rational" individuals would predict. Chapters 9 and 10 consider applications of game theory and alternatives for defending electrical grids (Chapter 9) and other targets (Chapter 10) against terrorists or other adversaries, including natural disasters in Chapter 9.

In addition to these major conceptual challenges, there are also purely technical challenges for making decision-analytic principles more widely applicable. For example, decision trees (Chapter 5) are well suited to model (and if necessary simulate) alternative possible sequences of events and decisions when there are only a few of each. However, they are far from ideal when the number of choices is large or continuous.

Example: Searching for a Hidden Prize

Suppose that a prize is hidden in one of 100 boxes, and that the cost of opening each box to see whether the prize is in it, as well as the prior probability that it is, are known (say, $c(j)$ to open box j, which has prior probability $p(j)$ of containing the prize). Then in what order should the boxes be opened to minimize the expected cost of finding the prize? (This is a very simple model for sequential search and R&D problems.) It would clearly be impracticable to create a decision tree describing the 100! possible orders in which the boxes might be opened. Yet, it is easy to determine the optimal decision. Simple optimization reasoning establishes that the boxes should be opened in order of descending probability-to-cost ratio (since interchanging the order of any two boxes that violate this rule can readily be seen to reduce expected cost).

The remainder of this book explains and illustrates breakthrough methods to help real people make real decisions better. It presents ideas and methods that the authors and editors believe are mature enough to be highly valuable in practice and that deserve to be more widely known and applied. The starting point, developed in Chapter 2, is a candid acknowledgment that:

- SEU and classical Bayesian decision analysis together provide a logically compelling model for how individual decisions ideally should be made; but
- Real people have not evolved to be always capable of producing (or agreeing on) the crisply specified, neatly decoupled subjective probabilities and utilities that are required (and implied) by SEU.

Instead, decision-makers in real-world organizations and institutions typically have access to some imperfect data and knowledge bearing on the causal relations between alternative choices and their probable consequences. From this partial information, and through processes of deliberation and analysis, they may construct preferences for alternative actions, with supporting rationales that are more or less convincing to themselves and to others.

The following five breakthroughs, explained and illustrated in the chapters that follow, can help to understand and improve these decision processes.

Breakthrough 1: Behavioral Decision Theory and Game Theory

It is now known that different neural pathways and parts of the brain are activated by different aspects of risky prospects, such as probabilities versus amounts of potential perceived gains or losses; immediate versus delayed consequences; positive versus negative emotional affects of cues used in describing them; trust versus suspicion of others involved in joint decisions; and moral reactions to risks and to imposing risks on others. To a useful first approximation, heart and head (or, more formally, "System 1" and "System 2," referring to the quick, intuitive and slower, more cognitive aspects of decision-making, respectively) may disagree about what is best to do, using different parts of the brain to evaluate alternatives and to arrive at these conclusions. Chapter 2 further describes these aspects of the divided decision-making self, situating the problem of wise decision-making (as well as moral and social judgment-making) in the context of competing decision pathways and emphasizing the need to use both emotional and intuitive judgments and rational calculations in making effective decisions. Preferences, judgments, and beliefs are often transient and contingent on context and cues (some of which may be logically irrelevant) when they are elicited.

Assessing a single coherent utility function for a person whose preferences arise from multiple cues and competing pathways may not yield a reliable basis for prescribing what to do if the goal is to minimize post-decision regret.

Recognizing these realities of human nature and behavior motivates *behavioral decision theory* and *behavioral game theory*. These address how to use well-placed "nudges," design of occupational and consumer environments, and other practical methods to help real people make better decisions while taking into account their heuristics, biases, inconsistencies, limited attention span and will-power, irrational altruism, moral aspirations, perceptions of fairness, desire for social approval, and other realities of motivation and behavior. "Better" decisions can no longer necessarily be defined as those that are consistent with SEU axioms for preferences and beliefs, as assessed at a given moment and in a given context. What constitutes desirable decision-making must be defined afresh when preferences and beliefs are seen as being constructed on the fly and subject to motivated and self-serving reasoning, wishful thinking, salience of cues and their emotional affects, priming by context, and other biases. For example, "good" choices might be defined as those that are consistent with the guidance or principles (more formally, with the if-then "decision rules" mapping available information to available actions) that one would ideally want one's self to follow, if one were given time, resources, and ability to develop such principles outside the context of short-run distractions, passions, and temptations. Such reflective and reflexive considerations, which have a long tradition in deontological, utilitarian, and virtue ethics, are gaining new currency and an applied focus through behavioral decision theory and game theory. The core breakthrough in Chapter 2 is the insight that advice on how to make decisions, to be most useful, should be rooted to an understanding of real human behavior and realistic possibilities for changing it.

Breakthrough 2: Simulation–Optimization of Risk Models

For decades, one vision of applied decision analysis has been that knowledge and information about the system or situation that a decision-maker seeks to influence via her decisions should be represented in an explicit *risk model* relating decisions (controllable inputs) and uncertainties (e.g., modeled as random inputs from the environment, not controlled by the decision-maker) to resulting probabilities of different outputs (consequences). If expected utility, or any other "objective function" whose expected value is to be maximized, is used to evaluate the probabilities of consequences induced by alternative choices of controllable inputs, then the decision problem of selecting those inputs can be decomposed into the following two technical tasks:

1. *Simulate* output (consequence) probability distributions, for any choice of inputs; and
2. *Optimize* inputs, that is, identify a combination of input values to produce the most desirable probability distribution of outputs, as evaluated by the objective function (e.g., expected utility).

If the risk model and simulation–optimization process can be trusted to model adequately the real system or situation of interest and to automatically find the best combination of controllable inputs to choose for that system, then the decision-maker is freed to focus on specifying the controllable inputs and the objective function to be used in evaluating results. Appropriate subject matter experts and modelers can focus on developing and validating the risk model describing the probabilities of consequences caused by different choices of controllable inputs (together with the uncontrollable ones chosen by "nature" or by others). The simulation–optimization engine can handle the details of solving for the best choice of inputs, much as software products such as the Excel Solver or Wolfram Alpha can solve simper problems, freeing users to focus on model development and input specification.

Example: Optimal Level of R&D

Suppose that a pharmaceutical company can invest in investigating any of a large number of new leads (molecules and molecular signaling pathways) in parallel for developing a drug to treat a disease in patients with a certain genotype. Each lead costs \$5M to investigate, and each has a probability 0.1 of proving successful within 5 years. The value of a success within 5 years is \$100M. The company must decide how many leads to investigate in parallel. What number of leads should be investigated to maximize the expected profit? This objective function, in units of millions of dollars, is given by the formula:

$$\begin{aligned}\text{expected profit} &= \text{probability of success} \times \$100\text{M} \\ &\quad - (\text{number of parallel investigations}) \times \$5\text{M} \\ &= \left(1-(1-p)^N\right) \times 100 - 5N,\end{aligned}$$

where N is the number of leads investigated—the decision variable in this problem—and p is the probability of success for each investigated lead. (The probability that all N investigated leads are unsuccessful is $(1-p)^N$, and hence the probability of success for the whole effort, that is, the probability that not all fail, is $(1-(1-p)^N)$.) The expected profit-maximizing value of N can readily be found in this simple example by searching over a range of values of N. For example, for those familiar with R, the following code generates Fig. 1.1: N = c(1:20); EMV = 100 × (1 − 0.9^N) − 5 × N; plot (N, EMV). The number of leads that maximizes the objective function is $N = 7$.

Now, suppose that the problem were more complicated, with unequal success probabilities and different costs for the different leads, and with annual budgets and other resource constraints limiting the number of projects (investigations) that could be undertaken simultaneously. Then instead of searching for the best solution over a range of values for N, it would be necessary to search a more complicated space of possible solutions, consisting of all subsets of projects (lead investigations) that can be investigated simultaneously (i.e., that satisfy the budget and resource constraints). If, in addition, the objective function could not easily be described via a formula, but instead had to be estimated by simulating many realizations of the uncertain quantities in the model for each choice of inputs, then efficient search and evaluation of different input combinations might become important, or even essential, for finding a good solution to the decision problem. Simulation-optimization provides technical methods for efficiently searching complex sets of feasible decisions, performing multiple simulation-based evaluations of alternative combinations of controllable inputs to identify those that (approximately) optimize the user-specified objective function.

The vision of decision-making as optimization of an appropriate objective function subject to constraints, corresponding to a model of how choices affect consequence probabilities, is fundamental in economics, operations research, and optimal control engineering (including stochastic, robust, and adaptive control variations). However, to make it practical, both the simulation and the optimization components must be

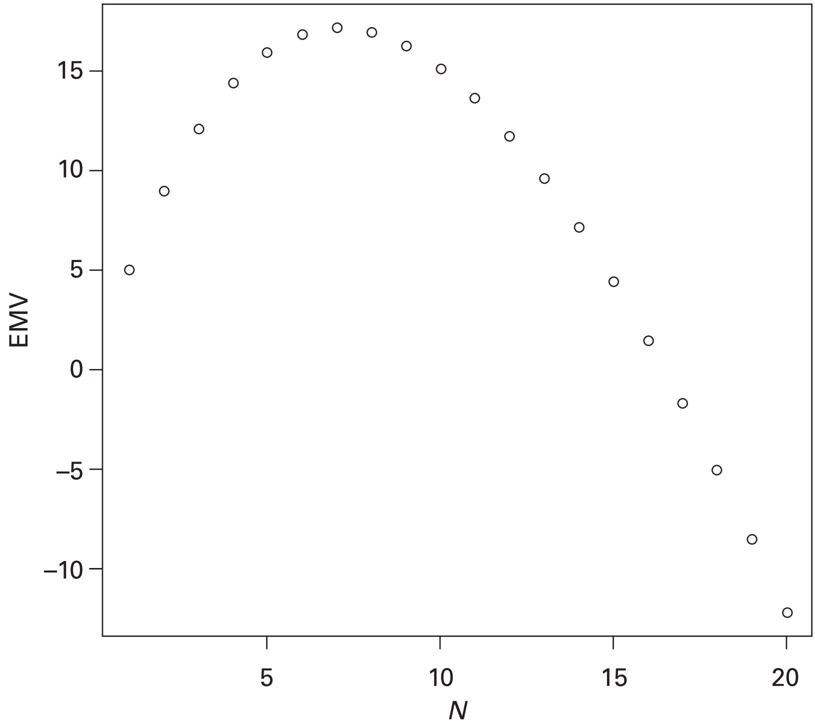

Figure 1.1 Plot of expected monetary value (EMV) for profit vs. N in the R&D example. Choosing $N=7$ parallel projects (leads to investigate) maximizes expected net return.

well enough developed to apply to realistically complex systems. This is an area in which huge strides have been made in the past two decades, amounting to a breakthrough in decision problem-solving technology. Sophisticated stochastic simulation techniques (e.g., Gibbs sampling, more general Markov chain Monte Carlo (MCMC) techniques, Latin Hypercube sampling, importance sampling, and discrete-event simulation methods) and sophisticated optimization methods that work with them (e.g., evolutionary optimization, simulated annealing, particle filtering, tabu search, scatter search, and other optimization meta-heuristics) are now mature. They have been encapsulated in user-friendly software that presents simple interfaces and insightful reports to users, who do not need to understand the details of the underlying algorithms. For example, the commercially available OptQuest simulation-optimization engine discussed in Chapter 3 is now embedded in software products such as Oracle, Excel, and Crystal Ball. Its state-of-the-art optimization meta-heuristics make it practical to easily formulate and solve decision problems that once would have been formidable or impossible.

The basic breakthrough discussed in Chapter 3 is to extend to realistically complex decisions the key decision and risk analysis principles of (i) predicting probable consequences of alternative actions (using stochastic simulation-based risk models of the relation between actions and their probable consequences, which may include complex, nonlinear interactions and random transitions) and (ii) finding the "best" feasible actions, defined as those that create preferred probability distributions for consequences. This is accomplished via *simulation-optimization models*, in which computer simulation models are used to represent the probable behavior of the system or situation of interest in response to different choices of controllable inputs. Powerful heuristic optimization methods, extensively developed over the past two decades, then search for the combination of controllable inputs that produces the most desirable (simulated) distribution of outcomes over time. Modern simulation-optimization technology substantially extends the practical reach of decision analysis concepts by making them applicable to realistically complex problems, provided that there is enough knowledge to develop a useful simulation model.

Breakthrough 3: Decision-Making with Unknown Risk Models

Simulation-optimization technology provides a breakthrough for deciding what to do if the causal relation between alternative feasible actions and their probable consequences—however complex, nonlinear, dynamic, and probabilistic it may be—is understood well enough to be described by a risk model that can be simulated on a computer. But suppose that the relation between controllable inputs and valued outputs is unknown, or is so uncertain that it cannot usefully be simulated. Then different methods are needed. One possibility, discussed in Chapters 4 and 7, is to *learn from experience, by intelligent trial and error*. In many applications, one might dispense with models altogether, and experiment and adaptively optimize interactions with the real system of interest, in effect replacing a simulation model with reality. (This is most practical when the costs of trial and error are relatively small.) For example, a marketing campaign manager might try different combinations of messages and media, study the results, and attempt to learn what combinations are most effective for which customers. Or, in the domain of healthcare, a hospital might try different drugs, treatment options, or procedures (none of which is known to be worse than the others) with patients, collect data on the results, and learn what works best for whom—the basic idea behind clinical trials. Chapter 4 considers optimal learning and anticipatory decision-making, in which part of the value of a possible current decision is driven by the information that it may reveal, and the capacity of such information to improve future decisions (the "value-of-information" (VOI) concept from traditional decision analysis). Recognizing that any current "best" estimated model may be wrong, and that future information may lead to changes in current beliefs about the best models, and hence to changes in future decision rules, can help to improve current decisions. Chapter 7 also discusses "low-regret" decision-making, in which probabilities of selecting different actions, models to act on, or decision rules are adaptively adjusted based on their empirical performance. Such adaptive learning leads in many settings to quick convergence to approximately optimal decision rules.

A second approach to decision-making with initially unknown risk models is possible if plentiful historical data on inputs and outputs are available. This is to estimate a relevant causal risk model from the available historical data. The estimated model can then be used to optimize decisions (e.g., selection of controllable inputs, or design of policies or decision rules for selecting future inputs dynamically, based on future observations as they become available). Chapter 6 briefly surveys methods of causal analysis and modeling useful for constructing risk models from historical (observational) data; for testing causal hypotheses about the extent to which controllable inputs (e.g., exposures) actually cause valued outputs (e.g., changes in health risks); and for estimating the causal impacts of historical interventions on outcomes. Chapter 7 discusses what to do when more than one possible model fits available data approximately equally well, making it impossible to confidently identify a unique risk model from the data. In this case, model ensemble methods, which combine results from multiple plausible models, can give better average performance than any single model.

Finally, a third possible approach to decision-making with unknown or highly uncertain models, also discussed in Chapter 7, is to seek "robust" decisions—that is, decisions that will produce desirable consequences no matter how model uncertainties are eventually resolved.

Of course, such a robust decision may not always exist. But a rich theory of robust optimization (and of its relations to related innovative concepts, including coherent risk measures and to older techniques such as stochastic programming) has been developed relatively recently, and this theory shows that many risk management decision problems of practical interest can be formulated and solved using robust optimization techniques. Chapter 7 and its references discuss these recent developments further.

Taken together, these relatively recent techniques for dealing with model uncertainty in decision-making constitute a distinct improvement over earlier methods that required decision-makers to, in effect, specify a single best-estimate model (typically based on subjective probabilities). Allowing for model uncertainty leads to new and useful principles for adaptive learning from data and for low-regret and robust optimization. These hold great promise for a wide variety of practical risk management decision problems, as illustrated by examples in Chapters 4 and 7.

Breakthrough 4: Practical Elicitation and Structuring of Probabilities and Multiattribute Utilities

The first three breakthroughs—behavioral decision and game theory, simulation-optimization, and methods for learning risk models from data and, in the interim, for making decisions with unknown or highly uncertain risk models—represent substantial enhancements to or departures from traditional SEU-based decision analysis. Breakthrough 4 consists of methods for making SEU theory more applicable to complex and difficult real-world problems by eliciting preferences and beliefs via techniques that impose less cognitive load on users and/or that achieve greater consistency and reliability of results than older methods. Chapter 5 considers state-of-the-art methods for developing multiattribute utility functions. This is a potentially painstaking task that once involved estimating multiple trade-off parameters and verifying quite abstract independence conditions for effects of changes in attribute levels on preferences. Chapter 5 presents much simpler and more robust methods, developed largely in marketing science, to enable relatively quick and easy development of multiattribute utility functions from simple preference orderings.

Breakthrough 5: Important Real-World Applications

The final category of breakthroughs consists of applications of decision and risk analysis principles to important and difficult fields that have historically relied on other methods. Chapters 5, 9, and 10 discuss applications of expert elicitation, game theory, decision tree analysis, Bayesian networks (for text mining of natural language), and machine learning techniques to the challenges of modeling and defending against adversarial actions. Chapter 8 illustrates the application of multiattribute utility theory to medical decision-making problems, at both the individual and the societal levels, by assessing utility functions for making trade-offs between consumption of sugar-sweetened beverages and risks of morbidity (type 2 diabetes) and early mortality. Chapter 9 discusses vulnerability, resilience, and defense of complex systems—specifically, electric power networks—and compares the insights gleaned from game-theoretic models and considerations to those from less sophisticated methods, concluding that the more sophisticated methods are often very worthwhile. Chapter 6 suggests that many public health decisions that are based on attempts to interpret associations causally would be much better served by applying more objective (and now readily available) methods of causal analysis and risk modeling.

In each of these application areas, and countless others, decision support methods have long been used that do not incorporate the precepts of SEU theory, modern decision analysis, simulation, optimization, optimal learning, or analysis and deliberation using causal risk models of causal relations. In each of these areas, adopting improved methods can potentially achieve dramatic objective improvements in average outcomes (and in entire probability distributions of outcomes). This point is made and illustrated by dozens of examples and case studies in the chapters that follow. Adopting the methods discussed in this book, implementing them carefully, and monitoring and learning from the results can yield breakthrough improvements in areas including marketing, regulation, public health, healthcare and disease risk management, infrastructure resilience improvement, network engineering, and homeland security. The conceptual and methodological breakthroughs presented in the following chapters were selected because they are ready for practical use and because they have been found to create great benefits in practice. Opportunities to apply them more widely are many, and the likely rewards for doing so are great.