TAMING UNCERTAINTY

Ralph Hertwig,
Timothy J. Pleskac,
Thorsten Pachur, and
the Center for
Adaptive Rationality

Taming Uncertainty

Ralph Hertwig, Timothy J. Pleskac, Thorsten Pachur, and The Center for Adaptive Rationality

The MIT Press
Cambridge, Massachusetts
London, England

This book was set in Stone Serif by Westchester Publishing Services. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Names: Hertwig, Ralph, author. | Pleskac, Timothy J., author. | Pachur, Thorsten, author.
Title: Taming uncertainty / Ralph Hertwig, Timothy J. Pleskac, and Thorsten Pachur.
Description: Cambridge, MA : MIT Press, [2019] | Includes bibliographical references and index.
Identifiers: LCCN 2018041802 | ISBN 9780262039871 (hardcover : alk. paper)
Subjects: LCSH: Uncertainty. | Decision making--Psychological aspects.
Classification: LCC BF463.U5 H47 2019 | DDC 153.8/3--dc23
LC record available at https://lccn.loc.gov/2018041802

10 9 8 7 6 5 4 3 2 1

Contents

Interactive Elements and Supplementary Material:
https://taming-uncertainty.mpib-berlin.mpg.de/

Preface

Almost every decision you make represents a leap into the unknown. You do not know with certainty what the consequences of your actions will be, nor do you know how likely the consequences are or when they will materialize—not to mention whether you and others will like the repercussions of your decisions. Faced with these myriad uncertainties, you may be tempted to throw in the towel. But that is not what people do. In fact, most of the time people are pretty good at navigating the unknown. The mind seems to be equipped with cognitive tools that empower people not only to reduce uncertainties where possible, but also to proceed in light of uncertainties that defy reduction.

In *Taming Uncertainty*, we aim to shed light on the cognitive tools in the mind's adaptive toolbox that help people make the leap into the unknown. For many decades, scholarly work in psychology and economics has understood the human response to uncertainty in terms of an attempt to uncover objective probabilities or, when this proves impossible, to conjure up subjective probabilities. Working with the currency of probabilities, this research has explored how extravagant Bayesian brains might update their estimates. In this book, by contrast, we take a different perspective—one that is rooted in what is known about what real, humble minds can do and resists whittling the human response to uncertainty down to an act of juggling probability quantities. In our view, adaptive intelligence in an uncertain world arises from a variety of simple tools. We focus on three types: first, decision strategies that efficiently permit people to infer and decide based on limited information by cleverly making use of key regularities in the environment; second, flexible search processes that guide where to look for further information and, importantly, when to stop searching and to act; and, third, cognitive tools that help people respond to the opportunities and challenges presented by others. We also trace changes in these cognitive tools brought about by the development of the human mind, both within and across generations. We argue that these tools empower the human mind with what the poet John Keats (1891) called "negative capability"—the ability to survive and thrive in uncertainty.

Ideas and arguments expressed on paper cannot fully replace experience. We have therefore created interactive elements to accompany many chapters. These online companions provide dynamic, hands-on encounters with some of the experimental paradigms, formal theories, and data featured throughout the book. We invite you to explore these elements—which are also referenced individually in the respective chapters—at https://taming-uncertainty.mpib-berlin.mpg.de/.

Taming Uncertainty is not an edited book; it is a joint product of the members of the Center for Adaptive Rationality (ARC), a multidisciplinary team of psychologists, economists, biologists, philosophers, computer scientists, and physicists, that was founded in 2012 at the Max Planck Institute for Human Development in Berlin. We worked on this book as a group in order to reflect and tap into our complementary disciplinary interests, skills, and knowledge; it summarizes our progress so far on our journey to unravel the nature of adaptive rationality in an uncertain world. *Taming Uncertainty* would have been impossible without the generous funding of the Max Planck Society, which has helped create this interdisciplinary group of researchers who are profoundly curious, challenging each other's viewpoints and exploring the mind's toolbox together.

We are also grateful to the colleagues who read drafts of individual chapters and provided their insightful feedback: Joshua Abbott, Ruben Arslan, Florian Artinger, Judith Avrahami, Simon Ciranka, Nadine Fleischhut, Thomas Hills, Ulrich Hoffrage, Sebastian Horn, Anika Josef, Juliane Kämmer, Yaakov Kareev, Robert Lorenz, Lucas Molleman, Shabnam Mousavi, Paul Pedersen, Amnon Rapoport, Samuli Reijula, Job Schepens, Oliver Schürmann, Warren Thorngate, Claus Vögele, Michael Waldmann, Charley Wu, Shuli Yu, and Veronika Zilker. We owe a further debt of gratitude to the scholars whose formidable brains we were able to pick during the production of the book: Bahador Bahrami, Gordon Brown, Jerome R. Busemeyer, Mike DeKay, Adele Diederich, Ido Erev, Craig Fox, Gerd Gigerenzer, Ulrike Hahn, Robin Hogarth, Joachim Krueger, Steve Lewandowsky, Taosheng Liu, John McNamara, Ben R. Newell,

Roger Ratcliff, Lael Schooler, Neil Stewart, Chris Summerfield, Kinneret Teodorescu, Peter M. Todd, Peter Wakker, Tom Wallsten, and Greg Wheeler.

Our most heartfelt thanks go to the many people in the Max Planck Institute without whose support this book would not have flourished—to Susannah Goss and Deb Ain for putting their careful eyes and unending patience to work editing the manuscript; to Diana Schieck for managing the references; to Antonio Amaddio and Philip Jakob for leading the development of the interactive elements; to Marianne Hauser for her meticulous work creating the figures; to Katja Münz and Kate Pleskac for keeping the authors and the book organized and running efficiently; and to Valerie Chase for helping us shape chapter 1 and getting us off to a good start. And, of course, our special thanks go to our families for their support and patience when we cruised off into the unknown once again.

This book marks the first step in our exploration of the murky waters of uncertainty. We are looking forward to our future discoveries of the mind's tools and feats. If you want to check on our progress, please visit the ARC website at www.mpib-berlin.mpg.de/en/research/adaptive-rationality.

Ralph Hertwig
Timothy J. Pleskac
Thorsten Pachur

Contributors

Ralph Hertwig Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany (hertwig@mpib-berlin.mpg.de)

Timothy J. Pleskac Department of Psychology, University of Kansas, USA (pleskac@ku.edu)

Thorsten Pachur Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany (pachur@mpib-berlin.mpg.de)

Eduard Brandstätter Department of Economic Psychology, Johannes Kepler University, Linz, Austria (eduard.brandstaetter@jku.at)

Larissa Conradt Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany (animcolldec@gmail.com)

Junyi Dai Department of Psychology and Behavioral Sciences, Zhejiang University, Hangzhou, China (junyidai@zju.edu.cn)

Mattea Dallacker Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany (dallacker@mpib-berlin.mpg.de)

Renato Frey Center for Cognitive and Decision Sciences, University of Basel, Switzerland (renato.frey@unibas.ch)

Philipp Gerlach Hochschule Fresenius, University of Applied Sciences, Hamburg, Germany (philipp.gerlach@hs-fresenius.de)

Stefan M. Herzog Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany (herzog@mpib-berlin.mpg.de)

Arend Hintze Evolving Artificial Intelligence Laboratory, Michigan State University, USA (hintze@msu.edu)

Anastasia Kozyreva Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany (kozyreva@mpib-berlin.mpg.de)

Ralf H. J. M. Kurvers Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany (kurvers@mpib-berlin.mpg.de)

Peter D. Kvam Department of Psychology, University of Florida, Gainesville, USA (kvam.peter@gmail.com)

Corinna Laube Center for Lifespan Psychology, Max Planck Institute for Human Development, Berlin, Germany (laube@mpib-berlin.mpg.de)

Tomás Lejarraga Departament d'Economia de l'Empresa, Universitat de les Illes Balears, Spain (tomas.lejarraga@uib.eu)

Christina Leuker Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany (leuker@mpib-berlin.mpg.de)

Aleksandra Litvinova Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany (litvinova@mpib-berlin.mpg.de)

Doug Markant Department of Psychological Science, University of North Carolina at Charlotte, USA (dmarkant@uncc.edu)

Jutta Mata Health Psychology Lab, University of Mannheim, Germany (mata@uni-mannheim.de)

Rui Mata Center for Cognitive and Decision Sciences, University of Basel, Switzerland (rui.mata@unibas.ch)

Mehdi Moussaïd Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany (moussaid@mpib-berlin.mpg.de)

David Pietraszewski Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany (pietraszewski@mpib-berlin.mpg.de)

Christin Schulze Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany (cschulze@mpib-berlin.mpg.de)

Leonidas Spiliopoulos Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany (spiliopoulos@mpib-berlin.mpg.de)

Dries Trippas Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany (dries.trippas@gmail.com)

Alan N. Tump Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany (tump@mpib-berlin.mpg.de)

Wouter van den Bos Faculty of Social and Behavioural Sciences, University of Amsterdam, Netherlands (W.vandenBos@uva.nl)

Jan K. Woike Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany (woike@mpib-berlin.mpg.de)

Dirk U. Wulff Center for Cognitive and Decision Sciences, University of Basel, Switzerland (dirk.wulff@unibas.ch)

Kyanoush S. Yahosseini Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany (yahosseini@mpib-berlin.mpg.de)

I The Research Agenda

1 Reckoning with Uncertainty: Our Program of Research

Ralph Hertwig, Timothy J. Pleskac, and Thorsten Pachur

1.1 Uncertainty as *the* Human Condition

At a press conference held in East Berlin on the evening of November 9, 1989, Günter Schabowski, a spokesman for the socialist regime of the German Democratic Republic (GDR), read out what seemed like a secondary item on a list of mundane announcements:

> A decision was made today, as far as I know…. A recommendation from the Politburo was taken up that we take a passage from [a draft of] the travel regulation and put it into effect, that, (um)—as it is called, for better or worse—that regulates permanent exit, leaving the Republic. Since we find it (um) unacceptable that this movement is taking place (um) across the territory of an allied state, (um) which is not an easy burden for that country to bear. Therefore (um), we have decided today (um) to implement a regulation that allows every citizen of the German Democratic Republic (um) to (um) leave the GDR through any of the border crossings. (Hertle, 2001, p. 157)

Schabowski continued on, but was interrupted by a reporter with a question that would change the world: "When does it come into effect?" Visibly uncertain, mumbling to his aides and thumbing through his papers, he finally said, "to my knowledge … immediately, without delay" (Hertle, 2001, p. 158). It is unlikely that he could have imagined the chain of events that these words would set off. Suddenly, the people of the GDR were thrust into their own state of uncertainty: Would the country's armed border guards obey the new travel regulations, granting complete freedom to travel, or would they use force to prevent any border crossing? Would the country's leaders revoke the announcement once its effects became clear, choosing to shed the blood of their people rather than risk a hemorrhaging of their population? Within hours, as joy vied with fear, tens of thousands of East Germans converged on the crossing points to West Berlin chanting "*Tor* *auf!*" ("Open the gate!"). Stunned, outnumbered, and lacking any information on the new policy or orders from the military leadership, the border guards ad-libbed—and opened the crossings. This division at the heart of Europe, which was literally set in stone, vanished overnight—and with it the certitudes, plans, and projections of millions of people.

Predictions about repercussions varied wildly. The GDR leadership had hoped the travel law would strengthen its regime (Hertle, 2001; Meyer, 2015). French President François Mitterrand, in contrast, believed that the Soviet leadership would never accept this development and that the Germans were unwittingly risking a world war. The very next day, Soviet leader Mikhail Gorbachev warned the leaders of France, the United Kingdom, and the United States of a possible "destabilization of the situation not only in the center of Europe but also beyond" (quoted in Hertle, 2001, p. 138). None of these things came to pass. As the saying often attributed to physicist Niels Bohr goes: Prediction is very difficult, especially about the future.

The quip highlights three more general properties of the human struggle with uncertainty. First, much of what people do—not only in politics but also in ordinary life—is predicated on forecasts of the future. Whether choosing a job, an apartment, or a spouse; whether deciding when to travel the world, have children, or start saving for old age; whether voting in an election or deciding between medical treatments, people base their decisions on predictions about what the future holds. Second, prediction can be difficult because most future events are shrouded in uncertainty—indeed, virtually all the political players in the fall of the Berlin Wall got it wrong (Hertle, 2001). Third, uncertainty and lack of knowledge bedevil not only people's predictions of the future but also their mental constructions of the present and reconstructions of the past (Loftus, 1993; Schacter, 1999; S. M. Stigler, 1980). "Ask any American who brought down the Berlin Wall, and nine of 10 will say Ronald Reagan," said former Secretary of State James Baker, but "we had hardly anything to do with it" (Meyer, 2015, para. 14).

1.2 Our Program of Research: Understanding the Adaptive Toolbox for Taming Uncertainty

If people inhabit "isolated islands of certainty in an ocean of uncertainty" (Arendt, 1958/2013, p. 244), this book is about how people navigate the high seas. Even under daunting conditions—where knowledge is imperfect, complexity is high, and time is short—people make predictions, inferences, and decisions so close to effortlessly that they have been called "masters of prediction" (Clark, 2016), and the accuracy of many of their predictions excels that of the political players in the fall of the Berlin Wall. What are the foundations of this mastery? In this book, we argue that they are a set of tools that the human mind—as an evolved and continuously learning cognitive system—has developed to grapple with uncertainty. And just as a good mechanic has multiple tools, each designed for a specific purpose, the human cognitive system has specific tools for dealing with the different forms of uncertainty it encounters (see chapter 18; Lo & Mueller, 2010).

With this notion of cognition, we challenge the idea that people manage uncertainty as if they were reducing it to risks, that is, as if they could summon at least subjective numerical probabilities for the outcomes of every decision they make (originally proposed by Savage, 1954). If they could, probability would constitute the very currency of human thought. Instead, we argue that people master the myriad types of uncertainty they face, from the inherent unpredictability of their environment (environmental uncertainty) to their limited knowledge and understanding of other people's actions and intentions (strategic uncertainty), by deploying a wide range of cognitive tools, many of which have no or only a rudimentary need to know the probabilities, let alone the utilities, of outcomes. We refer to this repertoire of tools for making predictions, inferences, and decisions as the mind's *adaptive toolbox*.

Guided by this concept, we build on our own and others' work on cognitive tools known as simple heuristics (Gigerenzer, Hertwig, & Pachur, 2011; Gigerenzer, Todd, & the ABC Research Group, 1999; Hertwig & Hoffrage, 2013; J. W. Payne, Bettman, & Johnson, 1993; Todd & Gigerenzer, 2012). Inspired by Herbert Simon and his vision of bounded rationality (1955, 1982, 1990), this work has proposed that much of human (and animal) reasoning, decision making, and behavior can be modeled in terms of such heuristics, which rest on simple principles of information processing and often consider only a subset of the information available. Heuristics are a realistic alternative to more classical approaches to decision making under uncertainty, such as expected utility theory and subjective expected utility theory, which often rest on extravagant implicit requirements and assumptions about the information, processing capacity, and time available (see Simon, 1955, pp. 103–104; see also chapter 2).

Simple heuristics constitute one important class of cognitive tools for reckoning with uncertainty. But a good mechanic's toolbox holds many handy implements in its drawers and trays: calipers for taking measurements; a flashlight to search for information and diagnose the problem so as to select the right tool; wrenches, ratchets, and pliers to fix it. And just as a mechanic can consult with colleagues and friends on tricky problems, people facing uncertainty can deploy social tools. Acknowledging the many facets of the adaptive toolbox, our research program examines two other important sets of cognitive tools, both of which serve to reduce reducible uncertainty or deal with irreducible uncertainty. One set supports search, permitting people to glean information about the future by learning from the environments they encounter. The second set has people teaming up with others and tapping their collective intelligence, that is, a group's ability to sometimes outperform individual decision makers when solving cognitive problems.

Our first goal in this book is to train the spotlight on these three sets of cognitive tools—heuristics, search strategies, and crowd aggregation rules—which we consider indispensable for reckoning with uncertainty. A second goal is to reveal the dynamic nature of the adaptive toolbox. Cognitive tools develop in response to changes inside the mind—for example, when cognitive resources like memory or knowledge grow or decline, or preferences concerning risk change. They also respond to changes outside the mind, such as when environmental demands shift (e.g., from exploring the world and forming alliances to finding a partner and raising children). Such developmental changes influence how the mind reckons with uncertainty and what kind of uncertainty it faces. How cognitive tools such as heuristics or search strategies are selected generally depends on factors such as the person's amount of accumulated knowledge (e.g., what they do and do not recognize), working memory capacity, value processing, and cognitive control, each of which has a distinct developmental trajectory. To understand how the mind handles uncertainty, it is therefore also imperative to understand how the adaptive toolbox is impacted by and develops along with the mind using it, as this book's handful of examples of the toolbox's mutability over the course of developmental change within individuals (ontogenetic change) and across generations of individuals (phylogenetic change) make clear.

Our third goal pertains to the kind of rationality that can arise from the relationship between the adaptive toolbox and the environment. Any cognitive tool will work in some contexts but not in others: its rationality is domain-specific rather than general. Like a lock and key (Barrett, 2005; see also Barrett & Kurzban, 2006), both the architecture of the cognitive tool and the respective environment must be investigated to determine how well they fit together. In this book, we extend this ecological rationality approach beyond the study of heuristics (e.g., Arkes, Gigerenzer, & Hertwig, 2016; Todd, Gigerenzer, & the ABC Research Group, 2012) to search strategies and crowd aggregation rules as well. Like heuristics, search strategies will succeed in some environments but fail in others; the same applies to strategies that harness collective intelligence. The ecological rationality of all these cognitive tools means that there is no master key. The lack of a master key, in turn, implies there are notable costs to using an adaptive toolbox, since the tools that empower the mind to deal with uncertainty are themselves a source of uncertainty. When should a specific tool be deployed? In this book, we seek to identify where tools and environments fit together, and where they do not; in so doing, we advocate a systemic view of uncertainty. This approach locates uncertainty neither solely in the mind (epistemic uncertainty) nor solely in the environment (aleatory uncertainty) but highlights the interactive dynamic of the two (see also chapter 18).

In summary, this book aims to advance the understanding of how the mind deals with uncertainty on three fronts. First, we integrate three important dimensions of human decision making into the concept of the adaptive toolbox: (a) boundedly and ecologically rational heuristics, (b) cognition as a search process, and (c) the tools the mind uses to tap into collective intelligence. Each dimension represents influential visions of human (and animal) cognition that have previously existed apart. Second, we advocate the view that the mind's repertoire of cognitive tools is anything but static—not only the toolbox but also its cognitive foundation and the environment are in constant flux and subject to developmental change. Finally, we demonstrate that each cognitive tool can be analyzed by enlisting the concept of ecological rationality, that is, the fit between specific tools and specific environments. All three goals are informed by the desire to further the understanding of what Arrow (1951) called "realistic" (p. 404) theories of how people make decisions without complete knowledge and, it should be added, where time and computational capacities are limited.

In seeking to capture how real people, alone and in tandem with others, make decisions under uncertainty, we take a decidedly different path than the ones trod by others. Many scholars of the mind see cognition in an uncertain world in terms of a single universal, computationally powerful and optimizing prediction machine—whether Bayesian in nature (Clark, 2016; Griffiths, Chater, Kemp, Perfors, & Tenenbaum, 2010; but see Jones & Love, 2011), resting on neural networks (McClelland et al., 2010; Rumelhart, McClelland, & the PDP Research Group, 1986), or involving a combination of both (e.g., Bayesian deep learning). Others leave the well-worn path of assuming classical information processing properties in favor of alternatives such as quantum information processing (Busemeyer & Bruza, 2012). Each approach has merits and has yielded new insights. For instance, it is useful to understand how the mind might go about implementing an optimal solution (Gershman, Horvitz, & Tenenbaum, 2015), or to what extent the principles of quantum information processing shed light on certain psychological regularities (Kvam, Pleskac, Yu, & Busemeyer, 2015). The adaptive toolbox approach we offer here seeks to understand the mind's amazing machinery for prediction, inference, and decision making as a repertoire of psychologically realistic strategies that fit specific environments. This toolbox promises flexibility and efficiency, thus equipping cognition with the ability to unlock the information that ever-changing environments carry about the uncertain future. We now turn to the conceptual issues that guide our investigations in the four major parts of this book: the heuristic mind, the exploring mind, the social mind, and the unfinished mind.

1.3 The Heuristic Mind

Heuristics are one of the most important types of realistic decision-making tools for coping with uncertainty. Each heuristic's policy represents a wager on the structure of the environment in question; it bets that ignoring some of the (often noisy) available information will enable faster, and potentially even more accurate, decisions (Gigerenzer & Gaissmaier, 2011). Indeed, one of the major discoveries of research on simple heuristics is that they are sometimes more accurate than complex procedures (e.g., Gigerenzer & Goldstein, 1996; Hertwig & Todd, 2003; Pachur, 2010; Pleskac, 2007). This discovery challenges the standard explanation of people's reliance on heuristics, which frames it as a kind of compromise between minimizing cognitive effort and maximizing accuracy (Shah & Oppenheimer, 2008). From this perspective, people rely on heuristics because searching for and processing information is taxing. Heuristics offer relief by trading accuracy for faster and more frugal cognition. This accuracy–effort trade-off—which is sometimes conceptualized as rational (J. W. Payne et al., 1993) and at other times as seriously flawed (leading to "cognitive illusions"; see Gilovich, Griffin, & Kahneman, 2002; Kahneman, Slovic, & Tversky, 1982)—seems to have been accepted as a potentially universal law of cognition.

The world's complexity is unquestionably beyond the grasp of the individual mind, but this David-versus-Goliath imbalance may not be why people rely on heuristics. The fact that simple heuristics can outperform more complex strategies raises a different possibility: the use of heuristics may reflect the power of heuristics rather than the weakness of the mind. Heuristics can be versatile and competitive under at least two conditions: when they are deployed in appropriate environments—that is, when the degree of tool-to-environment fit is high—and equally important, when information about the decision environment is scarce (e.g., Gigerenzer & Brighton, 2009; Katsikopoulos, Schooler, & Hertwig, 2010). This second condition is likely to be the rule rather than the exception in the complex environments that humans inhabit.

Brunswik (1957/2001) likened cognition and the environment to a married couple who have to come to terms with each other through mutual adaptation. The concept of ecological rationality (Todd, Gigerenzer, et al., 2012) highlights the fit between a heuristic—or, as we propose in this book, of any decision-making tool—and an environment. It also raises the following questions:

What structures does the environment offer that a decision-making tool could exploit?

In which environments does a particular decision-making tool succeed?

Which decision-making tools succeed in a particular environment?

How do decision-making tools and environments coevolve?

Cognitive science, psychology, and behavioral economics have often been content to study just one half of the couple: the mind's "software" and capacities. But as any marriage counselor knows, listening to just one side of the story will probably not illuminate why and when a marriage does or does not work. The same holds for cognitive strategies. In research on ecologically rational heuristics, considerable progress has been made toward understanding the intersection between the environment and the mind (Fawcett et al., 2014; Hogarth & Karelaia, 2006; Martignon & Hoffrage, 1999, 2002; Pachur, Hertwig, & Rieskamp, 2013a; Pleskac & Hertwig, 2014; Şimşek, 2001; Şimşek & Buckmann, 2015; Todd, Gigerenzer, et al., 2012). Yet the study of the interplay between the mind and the environment is still in its infancy; important challenges remain. In the second part of this book, titled "The Heuristic Mind," we take on three of these challenges.

1.3.1 How Do Preferential Choice Heuristics Handle Uncertainty?

From the outset, research on ecologically rational heuristics and "Homo heuristicus" (Gigerenzer et al., 2011) has focused on the domain of inference (e.g., which of these paintings is more valuable?) rather than preference (e.g., which of these two paintings would you like to have?). The key reason is that the domain of inference involves external criteria (e.g., the monetary value of a painting) and, therefore, commonly accepted benchmarks for evaluating the performance of inferential heuristics and people's use of them. Since benchmarks for preference are less clear, "accuracy" of choice in this context has been defined in many ways, ranging from adherence to coherence criteria (e.g., transitivity; see Pleskac, Diederich, & Wallsten, 2015) to "gold standards" such as expected value or utility maximization. Examination of choice heuristics in the domain of preference has been limited to the world of risk, where, according to Luce and Raiffa (1957), the outcomes of actions and the probabilities of those outcomes are known (e.g., Brandstätter, Gigerenzer, & Hertwig, 2006; J. W. Payne et al., 1993). In addition, choice heuristics have typically been invoked to explain systematic violations of, for instance, transitivity and axioms such as independence (Katsikopoulos & Gigerenzer, 2008; Tversky, 1969, 1972). But what happens when choice heuristics and strategies of rational choice (e.g., expected value theory) engage in *decision making under uncertainty*, where each action has a set of possible outcomes whose probabilities are not known (Luce & Raiffa, 1957)? Will some heuristics be on a par with, or even more accurate than, computationally more complex strategies in the domain of preference, where the "Olympian" models of rationality (Simon, 1983, p. 19) were originally proposed (e.g., subjective expected utility theory; Savage, 1954)? If so, in which environments can different heuristics be expected to succeed or fail? Chapter 2 addresses these questions.

1.3.2 Which Environmental Structures Are There to Be Exploited?

The second challenge we take on focuses on the environmental half of Brunswik's (1957/2001) married couple. What structures does the environment offer that might enable ecologically rational decisions? In the domain of inference, for instance, recognition has been identified as an important predictive indicator. It builds on the environmental regularity that objects scoring high on a criterion (such as large cities, wealthy people, or successful athletes) are seen and talked about more frequently than objects that score low. This ecological structure is exploited by the recognition heuristic (D. G. Goldstein & Gigerenzer, 2002), which interprets the failure to recognize one object and the recognition of another as a sign that the latter scores higher on a given criterion. But there is also an environmental regularity that binds extreme values and frequency differently: the relationship between risk and reward. Chapter 3 focuses on this structure, which exists in many environments. For instance, the bigger the jackpot one can win in a lottery or casino, the smaller the chances of actually winning it. Being cognizant of this regularity is an ecologically smart way to estimate unknown probabilities, thus reducing uncertainty. Chapter 4 demonstrates that structures in the social world (see also chapter 14), such as spatial clustering of social phenomena and people's hierarchical social network structure, can also be exploited to make accurate and frugal inferences about social statistics such as the frequencies of health hazards in the population.

Like Brunswik (1957/2001), Simon (1990) emphasized the collaboration between cognition and environment and insisted that "to describe, predict and explain the behavior of a system of bounded rationality, we must both construct a theory of the system's processes and *describe the environments* to which it is adapting" (pp. 6–7; emphasis added). The science of bounded rationality has proceeded along these lines by surveying and cataloging choice environments and asking which strategies and solutions would be most effective in each (Marewski & Schooler, 2011; Şimşek & Buckmann, 2015; Todd, Gigerenzer, et al., 2012). One downside to this approach is that one may end up with a different heuristic or set of heuristics for each discernable environment or environmental structure, resulting in a multitude of descriptions of environment–heuristic associations. In order to avoid such "description inflation," let us reframe Simon's goal as follows: In order to explain the behavior of a system of bounded rationality, we must eventually also construct theories of the system's processes as well as *theories of the mechanisms* underlying the emergence of classes of environmental structures. Admittedly, this goal is extremely ambitious, but we have already made modest progress toward it. In chapter 3, we outline a theory that explains and predicts when and why risk–reward structures emerge and, by extension, where a heuristic exploiting this structure can be expected to work well or falter (see Pleskac, Conradt, Leuker, & Hertwig, 2018).

1.3.3 Can Heuristics Succeed under Strategic Complexity?

A third challenge we address is whether the success of heuristics is restricted to static environments that do not require sophisticated strategic responses. There is a firm belief that simple heuristics are destined to fail when employed in "interactions with other intelligent agents, especially competitive agents" (Sterelny, 2003, p. 53). The rationale behind this belief is that social environments populated with other, competitive, agents are much more complex than physical environments (see Hertwig, Hoffrage, & the ABC Research Group, 2013). In competitive environments, strategies face counterstrategies, ostensibly requiring individuals to proactively interpret and forecast the behavior of others. In this view, individuals need to be aware that others will try to get the better of them, whereas nature, in its dispassionate amorality, will not. Is there any way that heuristics can prevail in these competitive interactions, which seem to require that an individual generate a model of the opponents' behavior, as well as a model of the opponents' model of the individual's behavior, and so on? Chapter 5 gives an answer: Heuristics also hold up in worlds invoking strategic interactions (see also Hertwig & Herzog, 2009).

To conclude, perhaps the most important discovery in research on simple heuristics is that they can be as accurate as, and sometimes even more accurate than, strategies that make the greatest possible use of information and computation. We address two important challenges to the generality of this finding: In preferential choice, it appears as if heuristics cannot escape an accuracy–effort trade-off; in strategic interactions with competitive agents, it is commonly assumed that heuristics will crash and burn. We also show that heuristics are not just vehicles for describing how people reckon with uncertainty; they also can be explicitly designed and enlisted to help people make better decisions. Chapter 6 analyzes the many uncertainties of the modern food environment, which cannot easily be reduced to the stock dimensions of outcomes and probabilities, and illustrates how people can use heuristics to safely navigate a world of carefully crafted temptations (see also chapter 14). But heuristics are just one set of indispensable tools for coping with uncertainty; there is another class of cognitive strategy that is equally important. Adaptive cognition is first and foremost about the smart search for information—the focus of the third part of this book, "The Exploring Mind."

1.4 The Exploring Mind

> *September 30, 1659.* I, poor, miserable Robinson Crusoe, being shipwrecked, during a dreadful storm in the offing, came on shore on this dismal unfortunate island, which I called 'the Island of Despair,' all of the ship's company being drowned, and myself dead.
>
> —Daniel Defoe, *Robinson Crusoe*

Can there be a lonelier and more uncertain world than the one described by Robinson Crusoe? What should he expect? Is the island inhabited? Are the locals friendly or hostile? Should he be worried about wild animals? What food sources are there? Is there drinking water? Where can he find shelter? Will anybody come to his rescue? To cope with these and other existential uncertainties, Crusoe recruited what is perhaps the quintessential tool for making decisions: he explored the island to ascertain the lay of the land, his options, and their possible consequences. In fact, search for information—either within the bounds of one's mind or in the external world—is what much of cognition is about.

At least two variants of search strategies can be used to reduce a knowledge gap. One is to simulate and forecast the consequences of one's actions on the basis of past data stored in memory—one's own or others' experiences in similar circumstances (e.g., Dudai & Carruthers, 2005). Of course, relying on past experience to anticipate the future only works if relevant past data are available. Sometimes there are none, as in the case of Crusoe, and sometimes, for whatever reason, the past is a poor predictor of the present or future. In such cases, another search strategy is required, one that permits the individual to explore the world, thus acquiring novel data before deciding. Because the process of sampling information is limited by the decision maker's time and capacity, unbounded search is not an option. But limited search does not necessarily imply poor decision making. On the contrary, mirroring the surprising accuracy of simple heuristics, limited search can yield surprisingly good results—whether it targets the information needed for heuristic inference (e.g., Katsikopoulos et al., 2010; Pachur et al., 2013a) or the properties of options in the context of preferential choice (e.g., Hertwig & Pleskac, 2010; Vul, Goodman, Griffiths, & Tenenbaum, 2014).

Many normative and descriptive theories of choice, including expected utility theory, subjective expected utility theory, and prospect theory (Kahneman & Tversky, 1979; Wakker, 2010), are mute on how people search for and learn from information. This reticence might be taken to suggest that how people search contributes little to comprehending how they handle uncertainty. But nothing could be further from the truth. Like Robinson Crusoe, people survive and even thrive in the ocean of uncertainty by enlisting search processes in external and internal environments: visually searching for targets of interest, looking up information on the Internet, or searching their semantic memory (Hills, Jones, & Todd, 2012). Unless decision scientists comprehend cognition as a search process, they will fail to understand important aspects of human behavior. Consider, by way of illustration, research on risky choice. For at least five decades (E. U. Weber, Shafir, & Blais, 2004), the field has predominantly studied how people make risky choices by asking them to choose between monetary gambles such as the following:

(A) An 80% chance of winning $4, otherwise nothing, and

(B) $3 for sure.

Although the expected value of option A is higher, most people prefer option B. Using choice problems like this, researchers have proposed elegant theories of risky choice that postulate, for instance, how people subjectively represent the objective information given (e.g., subjective functions of probability and outcome information). Yet the process of search is completely missing from these subjective representations and choices; all the necessary information is handed to decision makers on a silver platter. They are thus making *decisions from description*, in what Edwards (1962a) described as "static" decision tasks where, as Busemeyer (1982) highlighted, there is no need to learn from the past: "When a static decision task is used, the decision maker does not have to learn from past experience with the outcomes of previous decisions.... This feature of the static decision task becomes a problem when generalizing results to the many day-to-day decisions that repeatedly confront individuals, since explicit information concerning outcome probabilities is frequently not available and must be learned from previous experience" (p. 176).

Indeed, in everyday life, "it is hard to think of an important natural decision for which probabilities *are* objectively known" (Camerer & Weber, 1992, p. 325). When people decide whether to start a business or ask someone out on a date, there are no actuarial risk tables to consult. Only by considering uncertainty, the opportunity to search, and the process of stopping search in the choice situation can researchers begin to predict and explain how people arrive at *decisions from experience* (Hertwig, Barron, Weber, & Erev, 2004). We and others have studied decisions from experience using the same kind of monetary gambles commonly employed in studies of decisions from description. Systematic comparison of decisions from description with decisions from experience reveals that choices can diverge systematically, an empirical regularity described as the *description–experience gap* (Hertwig & Erev, 2009). Chapters 7 and 8 examine this gap—which could be called a risk–uncertainty gap—and its potential causes. What is increasingly clear is that the major theories developed by decision science to understand decisions from description do not readily apply to decisions from experience. One reason, though not the only one, is that these theories pay no attention to search, the key process through which people pick up on environmental regularities that, in turn, enable them to reduce uncertainty.

1.4.1 How Search and Choice Are Intertwined

Search and learning are at the heart of decisions from experience. Therefore, one path forward is obvious: a computational model of a system that explains decisions from experience must specify (a) the processes that guide sampling, (b) the processes that terminate sampling, and (c) the processes that generate a decision on the basis of the sampled information. Models of decisions from experience have attempted to do this in different ways. Often, however, the models treat search and choice as independent processes—and past and new experience as a record or set of records that can be consulted when making a final choice (e.g., Baron, 2005; Fox & Hadar, 2006; C. Gonzalez & Dutt, 2011; Hau, Pleskac, Kiefer, & Hertwig, 2008; Hertwig, Barron, rife with unforeseeable hazards, diseases, and shortages. Social worlds and our knowledge about them present both opportunities and pitfalls (see also chapter 4). In the fourth part of this book, "The Social Mind," we turn to this duality and the challenges it entails.

1.5.1 Search and the Challenges of Competition

Animals and humans are constantly faced with important adaptive problems: what to eat, where to live, which mate to choose. Not infrequently, answers must be found under the threat of competition. Under these circumstances, the mind faces different types of uncertainty (e.g., environmental, strategic) and must learn how to balance their consequences. Take, for instance, hermit crabs, which live in the abandoned shells of other sea creatures—usually snails. Since they require larger shells as they grow, they are always on the lookout for a new mobile home. Since the quality of the available shells in the environment varies, a solitary crab faced with merely this kind of environmental uncertainty will thoroughly inspect any potential new home before moving in. However, if several crabs encounter an empty shell at the same time, each individual crab also faces strategic uncertainty, because it does not know the intentions and strategies of its competitors. Under these conditions, what would otherwise be a meticulous exploration is dramatically curtailed (Rotjan, Chabot, & Lewis, 2010). The crab nearest to the shell will make a split-second decision on whether or not to take it based on a brief visual inspection alone. Once a shell is taken, a chain reaction of shell upgrades may quickly ensue.

This strategic shell game illustrates the challenges and opportunities of sharing the world with others. Exposed to environmental uncertainty alone, the crabs face what is known as the exploration–exploitation trade-off: whether to continue inspecting a shell (exploration) or to move into it (exploitation). In human choice, the exploration–exploitation trade-off in solitary choice situations has been extensively studied, both theoretically (Brezzi & Lai, 2002; Gittins, 1979; Gittins, Glazebrook, & Weber, 2011) and empirically (Cohen, McClure, & Yu, 2007). The presence of competitors creates an even more complex dilemma. For instance, the more time one takes to explore and thereby reduce environmental uncertainty, the higher the risk that a competitor will act first. There is no perfect escape from this dilemma. Hermit crabs solve it by adjusting their exploration strategies and aspiration levels to the situation at hand, acting as meticulous product testers when alone but taking the leap after just a quick peek when faced with fierce competition. Little is known about how people try to come to terms with competition during search. In chapter 12, we examine the extent to which people's search in the external world adapts to the challenges of a competitive environment. Are their search tools as ecologically rational as those of the hermit crab when environmental and strategic uncertainty conspire to create a difficult trade-off?

1.5.2 Complex Collective Behaviors Often Arise from Simple Rules

The presence of others does not always mean competition. In many domains, people are better off cooperating than competing. A fascinating domain in which to examine how large groups of people cooperate is individuals' movement through a shared physical space. The motions of many individuals can combine to create complex collective patterns. Consider pedestrian behavior and other self-organization phenomena in crowds. In many cases, such as when unidirectional lanes form spontaneously as people move, collective behavior can be beneficial. In other cases, such as stop-and-go waves and crowd turbulence, collective behavior can be highly detrimental and potentially disastrous, as when people rush to building exits in an emergency. One modeling approach, inspired by Newtonian mechanics, uses models from fluid dynamics and social force theory to conceptualize crowd dynamics (see Moussaïd, Helbing, & Theraulaz, 2011). This approach treats pedestrians as molecules and their motion as the result of attractive, repulsive, driving, and fluctuating forces, but does not capture the cognitive processes in each mind. Chapter 14 offers an alternative to physics-based models that aims to describe the underlying actual processes: pedestrian heuristics. The view that begins to emerge is that neither strategically demanding interactions (see also chapter 5) nor complex collective behaviors require complex mechanisms (see also Hertwig, Davis, & Sulloway, 2002). Simple heuristics offer a good starting point for explaining both.

1.5.3 The Ecological Rationality of the Wisdom of Crowds

Decision making in humans and animals often occurs in groups. Grouping organisms such as social insects (e.g., ants) must often make rapid decisions about which direction to move in or what action to take in uncertain and dangerous environments. These decisions are rarely solitary. In fact, in swarming ants, schooling fish, and flocking birds, effective distributed decision making occurs across a range of environmental contexts (see Couzin, 2009). Pooling information, votes, and preferences in order to make group decisions can be a powerful way of outwitting both the physical and the social environment (e.g., Krause, Ruxton, & Krause, 2010; Woolley, Chabris, Pentland, Hashmi, & Malone, 2010). But the wisdom of the crowd is not fail-safe. By all means, tapping the wisdom of others and facing environmental uncertainties together can facilitate solutions that go beyond the capacities of the individual, especially when the problem is difficult to solve alone (e.g., catching larger prey). But collective decisions are not invariably better than individual ones (Sunstein & Hastie, 2015). Groups can fail to reach better decisions for a number of reasons: for instance, people do not usually become members of a group at random but are selected for specific reasons; selection processes can produce biased groups. More generally, like other decision-making tools, tools for harnessing the wisdom of the crowd are not good or bad per se; their success depends on the problem and the environment at hand. Again, this set of tools can be understood in terms of ecological rationality, as chapter 13 demonstrates.

1.6 The Unfinished Mind

The mind's adaptive toolbox is a work in progress. It is never completed. In the fifth part of this book, "The Unfinished Mind," we examine how the life-span trajectory of cognitive development shapes the use of the mind's decision-making tools. Admittedly, little is yet known about the intricate dynamics between cognitive development and the adaptive toolbox, but we can hope to gain glimpses into their interplay (see also chapter 4).

One way we have gained traction on uncovering the dynamics involved is by looking for developmental change that is rooted in the change of core cognitive abilities. Take, for instance, the changes during a person's lifetime in "crystallized" cognitive abilities, such as vocabulary and world knowledge, versus "fluid" abilities, such as reasoning, attention, processing speed, and working memory (Li, Lindenberger, & Sikström, 2001). Whereas crystallized abilities increase throughout young adulthood and middle age and then plateau, fluid abilities increase in childhood and adolescence, peak in young adulthood, and decline from middle adulthood through old age. These developmental changes in the mind's cognitive abilities impose age-specific constraints for the processes that draw on them. The coupling of these abilities and decision-making tools—as well as their basic building blocks, such as the abilities to activate, represent, maintain, and process information—can be expected to be particularly strong in childhood, when crystallized abilities are least developed, and in old age, when fluid abilities are in decline. It follows that, depending on the cognitive abilities that make up specific decision-making tools, the degree to which these tools are used efficiently will be higher or lower during some periods of cognitive development than others.

By way of illustration, let us return to the recognition heuristic (D. G. Goldstein & Gigerenzer, 2002). For choices between two alternatives, the recognition heuristic is stated as follows: If one of two objects is recognized and the other is not, then infer that the recognized object has the higher value with respect to the criterion. This heuristic is useful when there is a strong correlation—in either direction—between recognition and the criterion, such as in competitive sports, where successful athletes are more likely to be mentioned in the media and in general conversation than less successful ones and are therefore more likely to be recognized. In contrast, in environments where people or media outlets talk about all objects in question equally often (or equally rarely), recognition is not correlated with the criterion, and the recognition heuristic should therefore not be used to make inferences. For an ecologically rational use of the recognition heuristic, some knowledge of the environment—particularly of the predictive power of recognition for the target in question—is required. Assuming that world knowledge (crystallized intelligence) grows with age, one might therefore expect that young children use the recognition heuristic less discriminately. This is exactly what Horn, Ruggeri, and Pachur (2016) observed when they investigated the use of the recognition heuristic across individuals in three age groups (9, 12, and 17 years old). First, they found that elementary school children already made systematic use of the recognition heuristic; second, 9- and 12-year-olds did not adjust their strategy use between domains in which the recognition heuristic resulted in accurate versus inaccurate inferences; third, older adolescents adaptively adjusted their use of the heuristic between domains. These findings suggest that, when the adaptive use of a heuristic requires crystallized abilities (knowledge and experience) and those abilities are still in the making, the heuristic may be used less adaptively.

The other developmental period during which one may expect a strong connection between cognitive abilities and cognitive tools is old age. During this period, fluid abilities such as working memory are in decline, which can be expected to impact search and learning tools. But what are the consequences? Will aging decision makers explore the world more as their ability to store, extract, and synthesize signals from the sampled data fades? Or will they explore the world less, because their ability to represent, maintain, and process information has declined? These are just a few of the conceptual questions raised by a lifespan perspective on the adaptive toolbox, most of which have barely been investigated. Chapters 16 and 17 report on the progress that we have made in addressing them so far. Chapter 16 deals with the period of adolescence, demonstrating that adolescents respond to uncertainty differently than adults or children and revealing how adolescence is a developmental period fraught with uncertainty. Chapter 17 addresses the development of risk-taking propensity, a major concept in theories of human choice. Two of its main findings are that risk-taking propensity almost universally declines over the adult lifespan and that it is systematically associated not only with the properties and requirements of the decision task but also with environmental conditions such as harsh living conditions.

We also take the analysis of developmental change one step further to include phylogeny. Evolution by natural selection—the longest-running process that produces a fit between environment and behavior—drives developmental change across generations. Chapter 15 demonstrates how computational evolution can be used to examine how natural selection may have shaped the tools in the adaptive toolbox and their ecological rationality.

1.7 A Systemic View of Uncertainty

The discussion of uncertainty has a long and winding history, in which scholars from different disciplines have proposed classificatory dichotomies such as that between measurable and immeasurable uncertainties (Knight, 1921/2002) or between epistemic and aleatory uncertainties (Hacking, 1975/2006). The former dichotomy pertains to whether the calculus of probability theory can be called upon to contain uncertainty by quantifying it. The latter concerns the source and cause of uncertainty, namely, lack

image
not
available

them, harvesting the knowledge and wisdom of others. But in decision making, as in life, there is no such thing as a free lunch. Having many tools means having to choose between them, and this choice can go wrong. The adaptive toolbox is thus both the solution to and the source of uncertainty. If we had to make a prediction, our bet would be that the advantages of the former greatly outweigh the disadvantages of the latter.

II The Heuristic Mind

2 The Robust Beauty of Heuristics in Choice under Uncertainty

Ralph Hertwig, Jan K. Woike, Thorsten Pachur, and Eduard Brandstätter

2.1 Axiomatizing Rational Choice—within Two Hours

> April 14, 1942. Today at Johnny's: axiomatization of measurable utility together with the numbers. It developed slowly, more and more quickly, and at the end, after two hours (!) it was nearly completely finished. It gave me great satisfaction, and moved me so much that afterwards I could not think about anything else.... (Oskar Morgenstern, cited in Leonard, 1995, p. 753)

The diary writer is the Austrian economist Oskar Morgenstern; Johnny is the Hungarian-born mathematician John von Neumann. The two men first met in the fall of 1938, by which time they had both left Europe for good and were working at Princeton. The culmination of their collaboration, the book *Theory of Games and Economic Behavior* (von Neumann & Morgenstern, 1944/2007), was an intellectual coup that would thoroughly transform a range of fields. One major step on their route to game theory was to formulate—within "two hours (!)"—an axiomatic foundation of Daniel Bernoulli's (1738/1954) path-breaking expected utility theory. According to this theory, a rational decision maker will choose among risky options in such a way as to maximize expected utility. Von Neumann and Morgenstern derived a set of axioms—such as transitivity, completeness, and independence (see Luce & Raiffa, 1957)—that the preferences and choices of a decision maker obeying expected utility theory would have to satisfy. The axiomatized version of utility theory swiftly became a framework for research in areas as diverse as statistical decision theory, management science, operation research, and the theory of the firm. Within a decade, expected utility theory was generalized from "objective" probabilities (or "risk," to use Knight's, 1921/2002, terminology) to "subjective" probabilities (Savage, 1954), giving rise to what is now called Bayesian decision theory. According to the Bayesian approach, a rational person can translate any uncertainty into numbers, that is, subjective probabilities, which must, first and foremost, be consistent—but not necessarily plausible. Diehard Elvis Presley fans who believe he is living among us, now in his early eighties, may estimate this probability to be 99% and assign 1% to him being dead. Largely unconstrained by facts, such beliefs can nevertheless be coherent.

2.2 The Olympian Model and Its Unrealistic Assumptions

Von Neumann and Morgenstern's (1944/2007) axiomatized utility theory also evoked fierce criticism. One challenge was empirical in nature. French economist and later Nobel laureate Maurice Allais (1953) did not mince his words: "Whatever their attraction might be, none of the fundamental postulates leading to the Bernoulli principle as formulated by the American school can withstand analysis. All are based on false evidence" (p. 505). Another challenge was conceptual. Although Herbert Simon respected utility theory's normative appeal (at least for the domains in which its assumptions hold)[1]—he profoundly criticized its unrealistic assumptions. In his article *A Behavioral Model of Rational Choice* (1955), Simon spelled out the "severe demands" of what he later described as an "Olympian model" (Simon, 1983, p. 19)—an ideal that might work for omniscient gods, but was simply out of place in the real world:

> If we examine closely the "classical" concepts of rationality […], we see immediately what severe demands they make upon the choosing organism. The organism must be able to attach definite pay-offs (or at least a definite range of pay-offs) to each possible outcome. This, of course, involves also the ability to specify the exact nature of outcomes—there is no room in the scheme for "unanticipated consequences." The pay-offs must be completely ordered—it must always be possible to specify, in a consistent way, that one outcome is better than, as good as, or worse than any other. And, if the certainty or probabilistic rules are employed, either the outcomes of particular alternatives must be known with certainty, or at least it must be possible to attach definite probabilities to outcomes. (pp. 103–104)

In actual human choice, such demands are rarely met. Beyond what Savage (1954) called "small worlds"—highly simplified environments such as

1. In fact, Simon (1945) wrote a glowing review of *Theory of Games and Economic Behavior.*

monetary gambles, where the consequences (e.g., monetary payoffs) and probabilities of all outcomes are known ("decisions from description"; Hertwig, 2015)—it is impossible for real people to live up to the decision-making ideal of specifying all possible outcomes, assigning them probabilities, and then maximizing the expected payoff. Instead, mere mortals often have access to only some of the information, or are unable to integrate that information in the sophisticated way mandated by expected utility theory, and instead rely on simplifying procedures. What do these constraints mean for the quality of people's choices? Simon (1956) conjectured that real organisms' behavior, although adaptive and satisficing, probably "*falls far short* of the ideal of 'maximizing' as postulated in economic theory" (Simon, 1956, p. 129, emphasis in the original). He further suggested that "the environments to which organisms must adapt possess properties that permit further simplification of its choice mechanisms" (p. 129).

Our goal in this chapter is to examine both of Simon's key theses. First, we investigate the price of simplicity: How far short of the ideal of maximization do simple choice strategies that fail to obey the demands of classical rationality fall? Second, we examine which statistical properties of the environment support or impede the performance of such heuristics. We analyze both questions in the time-honored environment of monetary gambles, that is, the very environment from which the concept of mathematical expectation and the classic notion of rational choice emerged (Hacking, 1975/2006; see also chapter 8). We do not, however, implement the monetary gambles as "small worlds" (Savage, 1954), fully described and with all outcomes and probabilities known; instead, we introduce uncertainty into the simulated environment in the form of imperfect knowledge. Before we describe this paradigm in more detail, let us briefly review two investigations that have inspired our own.

2.3 How Short Do Risky Choice Heuristics Fall of the Ideal of Maximization?

This analysis builds on the foundation of two previous investigations: Thorngate's (1980) strategy tournament and J. W. Payne, Bettman, and Johnson's (1988, 1993) influential research on the adaptive decision maker. Both focused on simple choice strategies and analyzed how short they fall of maximization. They explored this question in the world of risk, in which each choice leads to one of a set of possible specified outcomes, and each outcome occurs with a known probability (Luce & Raiffa, 1957, p. 13)—that is, in a world in which there are no surprises.

2.3.1 Efficient Decision Heuristics and Measures of Success

Since the 1970s, many decision scientists have joined the quest to identify systematic biases originating from people's reliance on heuristics (Kahneman, Slovic, & Tversky, 1982). Ignoring the zeitgeist, Thorngate (1980) instead asked how simple choice strategies can be: How much information can they ignore and still permit successful choices? In one of the early computer simulations in the decision sciences, he orchestrated a computer tournament in which 10 choice heuristics competed against one another. The heuristics were tested in a randomly generated choice environment involving choice problems with two, four, or eight options, with each option offering two, four, or eight outcomes (for details, see Thorngate, 1980). Here, we focus on the performance of the two top-performing heuristics in Thorngate's competition—the equiprobable heuristic and the probable heuristic—and the worst-performing heuristic, the least-likely heuristic. The three heuristics' policies are outlined in box 2.1. Each of the heuristics ignores a different aspect of the available information and thus implements a different variant of cognitive simplification (Gigerenzer & Gaissmaier, 2011). The equiprobable heuristic ignores all probabilities, meaning that it does not multiply (weigh) outcomes by their probabilities. Like Dawes' (1979) "improper linear models," it simply calculates the arithmetic mean of all outcomes per option and chooses the option with the highest mean. It thus acts as if each outcome, no matter how small or large, is as probable as any other outcome. The equiprobable heuristic thus embodies the "principle of indifference," a coinage attributed to Keynes (1921/1973b). It states that whenever there is no evidence favoring one possibility over another, they have the same probability (see also chapter 5). The probable heuristic, in contrast, considers probabilities, but only to classify outcomes into two sets (probable vs. improbable outcomes); it then removes all improbable outcomes from consideration. The least-likely heuristic considers only one kind of outcome per option, namely, the worst possible outcome, and chooses the option with the smallest probability that this worst outcome will occur.

How high is the price that the heuristics pay for straying from the ideal of maximization and ignoring some or much of the available information?

In sum, although Thorngate (1980) applauded the efficiency of the heuristics, they fell noticeably short of the ideal of expected value maximization on his performance measure (which used the "best" option as a benchmark). In contrast, on a performance measure that gauges how costly it is to fail to choose the best alternative when choosing heuristically, we found that simple choice policies fared substantially better. Furthermore, Thorngate's analysis did not address the role of environmental properties in fostering or impeding the heuristics' performance. For instance, how robust will the top performer, the equiprobable heuristic, prove to be if the environment's probability distribution is skewed, causing high variance in probability information and potentially rendering the assumption of equal weights dangerously inaccurate?

2.3.2 The Adaptive Decision Maker

The interaction between environmental structures and the information-processing architecture of heuristics was a focus of J. W. Payne et al.'s (1993) research program. Its premise was that people can select from a multitude of available strategies. Each strategy combines attractive properties (e.g., simplicity, low cognitive effort, accuracy) with unattractive ones (e.g., higher cognitive effort, lack of accuracy). The impact of these properties varies across choice environments and conditions such as time pressure and cognitive abilities. An adaptive decision maker considering candidate strategies for a task will select the one that affords the best trade-off between anticipated accuracy and effort. The assumption here is that there is an inescapable and law-like *accuracy–effort trade-off*, meaning that the less information, computation, or time a strategy requires and the decision maker invests, the less accurate (rational or optimized) the ensuing behavior will be.

J. W. Payne et al. (1988) used computer simulations to analyze this accuracy–effort trade-off in risky and riskless choice in a range of environments and conditions. We focus on the simulation of risky choice and on a single environmental condition: variance in probabilities. For example, one option may have four possible outcomes, with probabilities of .28, .25, .25, and .22, respectively. In this case, variance in probabilities is low. The outcomes of another option may have probabilities of .7, .2, .08, and .02; in this case, variance in probabilities is high. How did heuristics perform under high and low probability variance? Here, we focus on the equiprobable heuristic (or "equal-weight heuristic," to use the terminology of J. W. Payne et al., 1988), the best performer in the Thorngate (1980) tournament, and the lexicographic heuristic (LEX), the best performer in the J. W. Payne et al. tournament. Like the other heuristics, LEX ignores part of the information (see box 2.1). It is a noncompensatory strategy; note that the environmental cue that in this heuristic overrides all other cues and determines the choice depends on the properties of the choice problem at hand. Details of the simulations are provided in J. W. Payne et al. (1988).

Three key results emerged (J. W. Payne et al., 1988). First, heuristics can be highly competitive in one environment but fail in another. The lexicographic heuristic, for instance, was very successful in the high-variance environment, but its performance dropped by more than 20 percentage points in the low-variance environment. No generalist heuristic emerged that was able to perform consistently well in all environments. Second and relatedly, variance in probabilities affected processing policies differently. Whereas high probability variance fostered the performance of the lexicographic heuristic, it undermined the success of the equiprobable heuristic; low probability variance had the opposite effect. Third, when averaged across all environments, the equiprobable heuristic and the lexicographic heuristic chose the best option in 79% and 56% of cases, respectively (the performance criterion in this simulation was equivalent to the expected value criterion employed by Thorngate, 1980). That is, the two heuristics paid a substantial premium for simplicity, similar in magnitude to that observed in Thorngate's analysis: the top choice performance was more than 20 percentage points below that of the expected value benchmark.

Let us summarize the findings so far. The most successful choice heuristics in Thorngate (1980) and J. W. Payne et al. (1988) represent two very different paths to cognitive simplification (see box 2.1). Whereas the equiprobable heuristic considers all outcomes per option but neglects probabilities altogether, the lexicographic heuristic examines one cue at a time, ignoring all others. Furthermore, low variance in outcome probabilities supports a focus on outcomes, whereas high variance in probabilities is more compatible with a noncompensatory choice policy. These insights were obtained in the world of risk, where the probability and outcome space are known. What will be the trade-off between performance and simplicity—or, to use J. W. Payne et al.'s terminology, accuracy and effort—when knowledge is imperfect and surprises can happen? In other words, how do the heuristics perform in a world of incomplete knowledge and uncertainty? The bias–variance framework, originally developed in machine learning (Geman, Bienenstock, & Doursat, 1992; T. Hastie, Tibshirani, & Friedman, 2001) offers a conceptual approach for understanding the impact of uncertainty on prediction models.

2.4 The Bias–Variance Dilemma

Bounds on people's knowledge about the environment can be hugely consequential, to the extent that limited knowledge intensifies the *bias–variance dilemma* (e.g., Gigerenzer & Brighton, 2009; Katsikopoulos, Schooler, & Hertwig, 2010). To introduce this concept, which is relevant for any kind of prediction model, we offer an example. Bias and variance both contribute to the total error committed by any prediction model. Let us imagine that a prediction model is attempting to learn an underlying (true) function from a sample of (potentially noisy) data that was generated by this function. Averaged across all possible data samples of a given size, the bias of the algorithm is defined as the difference between the underlying function and the mean function computed by the algorithm from these data samples. Consequently, if this mean function is the same as the underlying function, bias will be zero. Variance reflects the sensitivity of the prediction model to different samples drawn from the same environment. High variance implies that the predictions of a model may differ greatly depending on the specific properties of the observed samples. This type of variance increases with model flexibility. For example, the more flexible the model, the more likely it will capture not only the true structure (assuming that there is a true structure and that the models are complex enough to capture it) but also unsystematic patterns, such as noise. Bias and variance both depend on the structure to be predicted (e.g., daily temperature across a year in a specific location; daily stock market fluctuation across a year of a specific index), and at least the variance also depends on how many sampled observations are available for this structure. Therefore, it will not always be adaptive to seek low bias in a prediction model by including as many adjustable components as possible to flexibly capture patterns in the sampled data. Model flexibility can itself become a curse when there is a high risk of increasing error through variance. From this, it follows that a model should be complex enough to avoid excessive bias, but simple enough to avoid overfitting idiosyncratic noise in the limited samples on which the estimates are made (e.g., Pitt, Myung, & Zhang, 2002).

The principles underlying the bias–variance dilemma can be applied to decisions about monetary lotteries under uncertainty—that is, when the task is to predict the value of the lotteries. Let us assume that choice strategies do not enjoy perfect knowledge of outcomes and probabilities but instead gauge them from samples drawn from the environment. Consequently, all strategies face two sources of error—and the possible trade-off between them. One source is error through variance. For illustration, expected value theory assumes that outcomes are weighted (multiplicatively) by their exact probabilities and then summed and maximized. Under conditions of imperfect information, other, simpler forms of information integration—for example, additive rather than multiplicative integration (Juslin, Nilsson, & Winman, 2009)—or forgoing integration altogether (e.g., lexicographic heuristic) may be more robust and less likely to suffer from overfitting. Yet simplifications can go too far, causing a substantial bias in the choice strategy and, consequently, prompting performance to deteriorate. How will these two sources of error shape the performance of choice heuristics under uncertainty? We examined this question by conducting a new set of simulations.

2.5 What Is the Price of Cognitive Simplicity in Choice under Uncertainty?

The benchmark used in our simulations is the performance of the expected value model under perfect knowledge. We call this model the omniscient expected value model. In addition, we used the same performance metric as in our reanalysis of the Thorngate (1980) competition. On this metric, 100% represents the sum of the expected values across choice problems in the case that a choice strategy always selects the option with the highest expected value (as the omniscient expected value model does); 0% represents the sum of the expected values across choice problems in the case that a strategy always selects the option with the lowest expected value. All heuristics were tested in environments in which the available information on the options' outcomes and probabilities was incomplete. Information was acquired through repeated sampling of monetary outcomes. We define a sample as the draw of a single monetary outcome from each option in a choice problem (with replacement). Each new draw thus offered information about the gambles' possible outcomes and their relative frequency of occurrence (i.e., probability). Furthermore, after each new sample, each heuristic rendered a choice, allowing us to analyze how the heuristics' performance changed as knowledge increased. In this simulated environment, the uncertainty the decision maker faced concerned both the outcome space (i.e., at any given point in time, the decision maker did not know for certain whether they were aware of the full outcome space) and the probabilities (i.e., the decision maker estimated the probabilities from the sequences of the encountered outcomes, meaning that the probabilities of possible outcomes that had not been encountered were unknown and, for known outcomes, the probabilities could only be estimated on the basis of the experienced sample of draws).

2.5.1 Competitors

We tested the equiprobable, the probable, and the lexicographic heuristics—the three top performers in the simulations by Thorngate (1980) and J. W. Payne et al. (1988)—as well as the least-likely heuristic, the worst-performing heuristic in Thorngate's competition. In addition, we tested the natural-mean heuristic (Hertwig & Pleskac, 2008, 2010), whose policy is described in box 2.1. This heuristic has some interesting characteristics. For one, it predicts the same choice as expected value theory if the latter also bases its choices on samples of experience rather than on perfect knowledge. It thus defines the level of accuracy of a sampling-based expected value theory (without Bayesian priors). However, the heuristic rests on a much simpler processing policy than expected value theory. Instead of multiplying each sampled outcome by its inferred (sample-based) probability and summing up the products, the heuristic simply totals up all experienced outcomes per lottery and then divides this sum by the sample size per option. In other words, it replaces the multiplicative core of expected value theory by simple summing and division, thus requiring no explicit representation of probabilities.

2.5.2 Environments

We implemented 20 choice environments, designed by combining five outcome distributions with four ways of constructing the associated probabilities (see also chapter 3). The outcome distributions consisted of a