

Heuristics

The Foundations of Adaptive Behavior

Edited by

Gerd Gigerenzer

Ralph Hertwig

Thorsten Pachur

OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

Oxford University Press, Inc., publishes works that further
Oxford University's objective of excellence
in research, scholarship, and education.

Oxford New York
Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in
Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam

Copyright © 2011 Oxford University Press

Published by Oxford University Press, Inc.
198 Madison Avenue, New York, New York 10016

Oxford is a registered trademark of Oxford University Press

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
electronic, mechanical, photocopying, recording, or otherwise,
without the prior permission of Oxford University Press.

Library of Congress Cataloging-in-Publication Data

Heuristics : the foundations of adaptive behavior / edited by
Gerd Gigerenzer, Ralph Hertwig, and Thorsten Pachur.
p. cm.

Includes bibliographical references and index.

ISBN 978-0-19-974428-2

I. Adaptability (Psychology) I. Gigerenzer, Gerd. II. Hertwig, Ralph.
III. Pachur, Thorsten.

BF335.H428 2010

155.2'4-dc22

2010013592



20.12 / 1387

153.82 6459

9 8 7 6 5 4 3 2 1

Printed in the United States of America
on acid-free paper.

INTRODUCTION

Gerd Gigerenzer, Ralph Hertwig, and Thorsten Pachur

How do people make decisions when time is limited, information unreliable, and the future uncertain? On the basis of the work of Herbert A. Simon and with the help of colleagues around the world, the Adaptive Behavior and Cognition (ABC) Research Group at the Max Planck Institute for Human Development in Berlin has developed a research program on *simple heuristics*, also known as *fast-and-frugal heuristics*. Providing a fresh look at how the mind works as well as the nature of rational behavior, this program has stimulated a large body of research; led to fascinating applications in fields as diverse as law, medicine, business, and sports; and instigated controversial debates in psychology, philosophy, and economics. In a single volume, we have brought together key articles that have been previously published in journals across many disciplines. These articles present theory, applications, and a sample of the large number of existing experimental studies. We have shortened many of the articles, corrected errors, Americanized spelling, and updated references that were in press at the time.

What kinds of theories can provide insight into how people make decisions? Logic, probability, and heuristics are three answers to this question and, more generally, to the question of how the mind works. For Aristotle, logic was a theory of ideal human reasoning and inference. For Jean Piaget, logic became the guiding metaphor for cognitive development, with the mature mind operating like an intuitive logician. Probability theory emerged only late, in the mid-seventeenth century, renouncing logical cer-

tainty by acknowledging that humble humans live in the twilight of probability. Since the second half of the twentieth century, the mind has been seen as an intuitive statistician and modeled by probabilistic accounts such as signal detection theory and Bayesian theories of cognition.

In contrast to logic and probability, *heuristics are processes that ignore information and enable fast decisions*. They are a comparatively recent means of understanding how the mind works. Albert Einstein used the term *heuristic* in the title of his 1905 Nobel Prize-winning article on quantum physics to indicate that he considered this view as incomplete, even false, but of great transitory use toward a more accurate theory. Max Wertheimer (a close friend of Einstein) and his fellow Gestalt psychologists spoke of heuristic methods in problem solving, such as “looking around.” Herbert Simon replaced the somewhat vague terms of the Gestalt school with computational models for the “art of guided search.” In artificial intelligence, heuristics were introduced for playing chess and problem solving with the goal of making computers “smart,” whereas in psychology, “heuristics-and-biases” were introduced to explain why people are not so smart.

WHY HEURISTICS?

The classic answer to this question is that because of their cognitive limitations, humans are unable to perform rational calculations and instead rely on error-prone heuristics. A variant of this view says that even when people could optimize, that is, compute the best decision, they often rely on

heuristics in order to save effort at the price of sacrificing some accuracy. The first answer assumes the inability to optimize, the second a pragmatic decision that it may not be worth spending the time. Both answers are based on the principle of an *accuracy–effort trade-off*: The less information, computation, or time one uses, the less accurate one’s judgments will be. This trade-off is believed to be one of the few general laws of the mind.

The rationale for heuristics laid out in this book is entirely different from these classical views. Less effort can in fact lead to better or worse accuracy, depending on the environment in which a heuristic is used, that is, its *ecological rationality* (see later). We begin with two concepts fundamental to understanding the role and potential of heuristics: Savage’s small worlds and Simon’s scissors.

Savage’s Small Worlds

Leonard Jimmie Savage is known as the creator of modern Bayesian decision theory, and his *Foundations of Statistics* (1954) counts as one of the most influential books on this topic. Savage carefully limited its applications to what he called “small worlds,” simple and well-defined microcosms in which all relevant alternatives, consequences, and probability distributions are known and no surprises are allowed. A typical example for a small-world decision is the purchase of a lottery ticket. In a small world, it is always possible to “look before you leap,” whereas in a large world, one can only “cross certain bridges once they are reached.” In what we call a “large world” or, alternatively, an uncertain world, not all alternatives, consequences, and probabilities are known or knowable, and surprises can happen. Savage warned against both using his theory outside of small worlds and believing that Bayesian updating would be the general solution to the problem of induction. Yet his caveat is disregarded when theorists build optimization models for small worlds and leap to the conclusion that the results describe or prescribe how humans make decisions in the large world (see Binmore, 2009).

How should we make decisions in the large world—that is, when Bayesian theory or similar optimization methods are out of reach? The

second half of Savage’s (1954) *Foundations of Statistics* turns to this question, analyzing heuristics such as minimax (i.e., choosing the alternative that minimizes the greatest loss). Savage’s distinction anticipates two research programs in the cognitive sciences. The first is to build complex models of behavior for small worlds; the second is to study the simple heuristics that perform well in large worlds. Ideally, these programs should progress in tandem, yet that is rarely the case. The limitation of optimization models to small worlds appears to be largely overlooked, and heuristics are relegated to mere shortcuts to optimal solutions.

In sum, the accuracy–effort trade-off indeed exists in a small world. In a large world, however, where there is no *general* accuracy–effort trade-off, heuristics can be more accurate than methods that use more information and computation, including optimization methods. Thus, one reason why people rely on heuristics is that these allow for decisions in both small and large worlds.

Simon’s Scissors

Behavior has been explained by intentions, preferences, attitudes, personality traits, cognitive styles, risk aversion, altruism, egoism, and other internal causes. The common denominator is that these causes are located inside the mind, and overt action follows unless circumstances prevent it. This “internalistic” perspective still dominates in both cognitive and psychodynamic psychology. Herbert Simon proposed a different explanation, based on an *ecological* view of behavior (1990, p. 7):

Human rational behavior (and the behavior of all physical symbol systems) is shaped by a scissors whose two blades are the structure of task environments and the computational capabilities of the actor.

Simon’s scissors analogy implies that internal explanations are incomplete explanations of behavior because they ignore the influence of the physical or social environment. For instance, one of the foremost contributions of social psychology has been to demonstrate the power of social environments, as in Milgram’s (1974) obedience experiments, in which the

INTRODUCTION

experimenter instructed the participant to administer electric shocks of increasing intensity to a person in another room every time the latter answered a question incorrectly. The degree to which people obeyed strongly depended on properties of the environment, such as whether the experiment was conducted in a nondescript office building rather than within the walls of a prestigiously ornate hall on Yale's old campus. It is impossible to understand why people's behavior in this and other situations is so intimately connected to the context if one looks for internal causes alone. Similarly, it is impossible to understand why a heuristic succeeds or fails by looking simply at the heuristic. The solution lies in the match between heuristic and environment, the two blades of Simon's scissors. Heuristics are not good or bad, rational or irrational per se, but only relative to an environment. The same holds for optimization methods. Thus, Simon's scissors analogy has given rise to a new question: In what environmental structures will a given heuristic fail, and in which will it succeed? We call this the study of ecological rationality. The structure of the environment, not the accuracy-effort trade-off, provides the key to understanding why and when it is rational to rely on heuristics.

RESEARCH PROGRAM

Research on fast-and-frugal heuristics has three goals: the first is descriptive, the second is normative, and the third is one of design or engineering.

(a) *The adaptive toolbox.* The descriptive goal is to analyze the content of the "adaptive toolbox," that is, the heuristics, their building blocks, and the evolved and learned core capacities on which heuristics operate. Examples of building blocks are search rules, stopping rules, and decision rules. Core capacities include, for instance, recognition memory, frequency monitoring, and the ability to imitate the behavior of others. Heuristics are simple because they take advantage of these capacities. The descriptive study of the adaptive toolbox examines its phylogenetic, ontogenetic, and cultural development, as well as the question of how heuristics

are selected in response to a goal. The main methods are observation and experimentation.

(b) *Ecological rationality.* The normative goal is to determine the environmental structures in which a given heuristic succeeds or fails, that is, the match between mind and environment. A heuristic is ecologically rational to the degree that it is adapted to the structure of an environment. Because ecological rationality dispenses with optimization, it can be applied to both small and large worlds. The study of ecological rationality results in statements about how well a heuristic functions (e.g., predicts faster, with less information, or more accurately) compared to competing strategies in a given environment. This analysis extends to the co-evolution of heuristics and environments. The main methods are computer simulation and mathematical analysis.

(c) *Intuitive design.* The engineering goal is to apply the results from (a) and (b) to design heuristics and/or environments for improving decision making in applied fields such as health care, law, and business. We refer to this goal as "intuitive design" because it relies on heuristics that reflect the way that the human mind works rather than on standard statistical software programs, which many professionals such as medical and legal decision makers find obscure.

To achieve these three goals, two requirements are indispensable.

1. *Process models, not only as-if models.* Optimization theories (such as Bayesian, expected utility maximization, cumulative prospect theory) typically imply complex estimations and computations and are thus presented as *as-if models*, that is, as models of the behavioral outcome, but not the cognitive process. The classical example of an *as-if model* is Ptolemy's theory in which planets move around the earth in circles and epicycles. Although few believed that this theory would describe the actual motions of planets, it was quite accurate in predicting their positions—provided that enough epicycles were included in the model. Kepler's theory in which planets move in ellipses around the sun, in contrast, was meant as a process

theory, that is, it described the actual motions of planets. The classical support for as-if models in the social sciences stems from the economist Milton Friedman (1953), who, like the psychologist B. F. Skinner, saw little value in modeling cognitive processes. In contrast, our aim is to understand actual decision processes, not only the outcomes. There is a good reason for this. Without modeling the cognitive blade of Simon's scissors, it is utterly impossible to determine in what environments heuristics succeed, that is, their ecological rationality.

2. *Computational models, not labels.* A crucial distinction is between models of heuristics and labels for them. Computational models include tit-for-tat, elimination-by-aspects, and the recognition heuristic, whereas "availability" and "representativeness" are often treated as vague one-word labels. Unlike labels, computational models enable the performance of heuristics in specific environments to be studied and can lead to novel predictions. None of the discoveries surveyed in this book, such as less-is-more effects, could have been made without computational models of heuristics.

The simple heuristics program is not the only one that aims to render more psychologically realistic theories of rational behavior. Two other influential ones are the heuristics-and-biases program and the optimization-under-constraints program. The heuristics-and-biases program (Kahneman, Slovic, & Tversky, 1982; Gilovich, Griffin, & Kahneman, 2002) shares with the simple heuristics program a focus on cognitive processes rather than as-if models, but differs in three important respects. First, the heuristics it describes are common-sense labels that have typically not been developed into computational models. Second, its definition of rationality—the benchmarks against which human judgment is evaluated—is not based on Simon's scissors; in it, rationality remains logical instead of ecological. Although it is sometimes said that the two research programs on heuristics differ by the amount of rationality they ascribe to humans, they already part ways

in their definitions of what counts as rational. Third, the heuristics-and-biases program assumes a general accuracy-effort trade-off and therefore also assumes that heuristics may be less effortful but can never be more accurate than more complex strategies. "Heuristics and biases" became inseparable twins and, in the words of economist Richard Thaler, biases or cognitive illusions are the "rule rather than the exception" (1991, p. 4). As a consequence, the role of heuristics is solely descriptive, and never normative as in the study of ecological rationality.

By introducing more realism into theories of rationality, the simple heuristics program also takes a drastically different approach from that of optimization under constraints. Optimization under constraints is a widely adopted modeling strategy in the cognitive sciences, neoclassical economics, behavioral economics, and optimal foraging theory. The idea is to make full rationality, or optimization without constraints, more psychologically realistic by adding psychological phenomena as free parameters to the utility framework (as in cumulative prospect theory or the inequality-aversion theory). This approach shares with the simple heuristics program an emphasis on computational models, but differs in three respects. First, it retains optimization—which becomes more demanding mathematically with each constraint added—and thus tends to generate complex as-if theories rather than process models. As Simon (1955) noted early in his career, "there is a complete lack of evidence that, in actual human choice situations of any complexity, these computations can be, or in fact are, performed" (p. 102). Second, like the heuristics-and-biases program, it assumes a general accuracy-effort trade-off, such as when computing the optimal stopping point at which the costs of further search exceed those of increasing accuracy. Finally, optimization with or without constraints requires known alternatives, consequences, and probability distributions, meaning that the domain of these theories is restricted to small worlds. This is why optimization models are of limited practical use for physicians, managers, politicians, and John Q. Public.

INTRODUCTION

THE STUDY OF HEURISTICS ANSWERS BOTH DESCRIPTIVE AND NORMATIVE QUESTIONS

Whereas logic and probability have been interpreted as descriptive (how we actually reason) and normative (how we ought to reason), the role of heuristics in psychology was until now limited to the description of cognitive processes, with no claim to normative merits. This reflects a longstanding tradition in philosophy, where rationality is defined orthogonally to psychology. For instance, the laws of logic and probability typically define rational judgment in the heuristics-and-biases program, whereas psychology is relegated to explaining why people deviate from these laws. Although it was acknowledged that heuristics are often efficient, nearly every experimental illustration was designed to demonstrate the existence of ever-new judgmental errors, as measured by logical norms. The resulting opposition between the rational and the psychological has grown into an unquestioned article of faith. In this context, psychology was limited to the clinical explanation of bad reasoning, thus implying that if everyone were to become rational, psychological research could be abandoned because it is mute about the nature of good reasoning. As a famous economist once said with conviction, "Look, reasoning is either rational or it's psychological" (Gigerenzer, 2000, p. vii).

In the mid-1990s, when we began our research program, most of our scientific community took this schism for granted. A major contribution of our research to date, as shown in this book, has been to demonstrate that simple heuristics using limited search, stopping rules, and aspiration levels can lead to more accurate inferences or predictions than can algorithms based on the principles of logic, probability, or maximization (e.g., see Chapter 1). This puts psychological processes on a par with models that obey the classical maxims of rationality. More important, it brings a new question to the foreground: In which environments is a heuristic better than, say, a logistic regression or a Bayesian model, and in which is it not? Posing this question changes the nature of rationality and the discourse on it from logical to

ecological. Once it is understood that heuristics can be more accurate than more complex strategies, they are normative in the same sense that optimization methods such as multiple regression and Bayes' rule can be normative—in one class of environments, but not in all. Because heuristics ignore information or exploit forgetting, cognitive limitations can be rethought as properties that may enable good decisions, not merely hinder them (e.g., see Chapter 4).

Homo heuristicus is not the cognitive miser that previous research made him out to be. One of our first discoveries concerned a simple inference heuristic called take-the-best. This heuristic ignores dependencies between cues and relies on the first good cue that allows an inference to be made about the world. Whereas experimental evidence for the use of this kind of lexicographic rule has been documented for decades, for instance, in preferential choice (Ford, Schmitt, Schechtman, Hults, & Doherty, 1989), it was widely believed that a lexicographic rule "is naively simple" and "will rarely pass a test of reasonableness," to quote two eminent decision theorists, Ralph Keeney and Howard Raiffa (1993, pp. 77–78). Cognitive processes that ignore information cannot but be inferior, or such is the die-hard belief of many decision theorists.

When we tested the take-the-best heuristic in a large world (where the order of cues was not known but had to be estimated from limited samples), however, the surprising result was that it could make *more* accurate predictions than strategies that use all information and computation, including optimization models (see Chapters 1 and 2). Hence, it is a conceptual error to equate complex models with rationality and to rate psychological heuristics as second-best or even liken them to irrationality. Which model is better is an empirical question and cannot be answered by an authoritative dictum. In this book, answers are provided by experiment, mathematical proof, and computer simulation.

A BRIEF HISTORY

The simple heuristics program has not evolved in a vacuum. Computational models of heuristics were proposed by Luce (1956), Simon

(1957), and Tversky (1972), among others, and the idea of an adaptive decision maker was advocated by Payne, Bettman, and Johnson (1993). In this earlier work, heuristics were still considered subject to an accuracy–effort trade-off. In the initial phase of our research program (featured in Part I), the focus was on showing that this deeply entrenched belief is not generally true. We demonstrated the existence of simple heuristics that can indeed be equally or even more accurate than complex strategies. Incredulous at first, we checked and double-checked the simulations. The results were reliable. The accuracy–effort trade-off is not an inevitable predicament of cognition: Looking beyond the individual into the mind–environment system, one begins to understand when and why the mind can have it both ways, achieving more accuracy with less effort.

The second phase of the research program focused on experiments. It began in an unexpected way. During decades of research on “heuristics and biases,” few had doubted that people rely on heuristics. Yet once our research challenged the link between heuristics and reasoning errors, the argument was raised that there was little evidence that people actually use heuristics. We include two chapters by researchers who were highly skeptical of the use of fast-and-frugal heuristics such as take-the-best (see Chapters 17 and 18). In one of these, Arndt Bröder tells of how the experimental evidence finally convinced even him, one of the fiercest critics. He also points out how important it is to ask the right questions. For instance, the question is not “does everyone always use heuristic X?” but instead “in which environment do people use heuristic X, and is this use ecologically rational?” The second section of Part II features some of these enlightening experimental studies.

The third phase of research addressed the study of ecological rationality, asking: In which environment is a given heuristic more accurate than a strategy that needs more information and computation, and in which is it not? The study of ecological rationality is inspired by Simon’s scissors analogy, with the mind and the environment as the two blades: In order to understand how a pair of scissors cuts, one needs to study

how the two blades match and interact. Looking at just one blade and observing, for instance, that minds ignore dependencies between cues does not suffice for ascertaining whether this leads to rational or irrational behavior. As Egon Brunswik (1957) pointed out, the mind and the environment are like a husband and wife who have come to terms with each other by mutual adaptation. The question is how heuristics and environmental structures function in tandem. The first section of Part II features some answers to this question.

The most recent development in the simple heuristics program is the study of decision making “in the wild.” Part III of this book features a sample of these exciting studies. They pose three questions that are related to the three goals of the program. The first question is descriptive: What strategies do experts and laypeople rely on in real-world decisions, outside the laboratory? Chapter 30 reports that when deciding on a location to break into (or establishing which one is likely to be burgled), experienced burglars and policemen follow the take-the-best heuristic, whereas inexperienced students in the laboratory appear to weight and add cues the “rational” way. The second question is both descriptive and prescriptive: How well does a heuristic perform compared to sophisticated statistical techniques developed for the same problem? Chapter 36 observes that experienced managers rely on only one reason to identify inactive customers and tests the accuracy of this heuristic compared to the accuracy of Pareto/NBD models that use more information and computation. Managers’ heuristics use less information but predict actual customer behavior better. Yet this fascinating study does not tell us *in what world* relying on one good reason is better. The third question asks: How ecologically rational are the involved strategies? Chapter 34 shows that the $1/N$ heuristic—invest equally in N assets—can make more money than optimal investment models, including Markowitz’s Nobel Prize-winning mean-variance portfolio. The point is not that simplicity can make you wealthier, but to understand both when the heuristic outperforms an optimization model and when optimization

INTRODUCTION

would pay. As in other ecological analyses, key properties of the environment include the predictability of the criterion, the sample size, and the number N of alternatives. With $N = 50$ assets and a window of 10 years of stock data (which is what many investment firms use), for example, $1/N$ would outperform the mean-variance portfolio, whereas around 500 years of stock data are needed to outperform $1/N$. The smaller the predictability and sample size and the larger the N , the greater the expected advantage of the heuristic over optimizing models.

TRANSPARENCY AND ROBUSTNESS

Last but not least, let us point out two methodological values that we have come to appreciate: transparency and robustness. Both are consequences of simplicity. Transparency can be reached by simple models that use zero free parameters (such as $1/N$ and the priority heuristic) or only few parameters (such as take-the-best, which needs to estimate the order of cues). It suffers when the number of adjustable parameters increases and when verbal labels are proposed instead of formal models. Of course, every model has parameters; the difference is whether they are free, adjustable within a range, or fixed. Unlike models of cognition that feature half-a-dozen adjustable parameters, those with zero adjustable parameters can be tested by hand. Thus, everybody can establish for themselves when a model does or does not predict the data. Transparency is essential for intuitive design. For instance, physicians tend to reject diagnostic systems based on logistic regression because they do not understand them. Yet a fast-and-frugal tree for coronary care unit allocation, as explained in Chapter 6, is transparent; a dozen years after being introduced, it was still used for predicting heart attacks by the doctors in the clinic under study and has been updated since for specific patient groups.

Robustness, a key feature of the evolved design of humans and other animals, concerns the ability of a mental strategy to operate successfully when the environment and inner

states change. Robustness often follows from simplicity because simple models that have no or few free parameters are less vulnerable to overfitting (i.e., increasing the fit by fitting noise) and tend to be more robust. Robustness is not the same as optimization. An organism that is optimally adapted to its past may fail if the future changes. A robust design, in contrast, is one that is not optimally adapted to its past but has a good chance of performing well when the future holds surprises. As Chapter 1 explains, simplicity leads to bias, and bias often enhances robustness. Thus, in the simple heuristics program, the notion of *heuristics and biases* takes on a new, more favorable meaning. Rather than necessarily leading to biases in the sense of errors, heuristics have biases built in to enable good judgments. Without bias, a mind could not function well in our uncertain world.

How do people make decisions when optimization is out of reach? This is the central question this book addresses. Simon once said that this was the only question that he tried to pursue in his scientific life. Yet it is far from being solved. As Simon wrote in a letter to one of us (GG) late in his life, in 1999:

I guess a major reason for my using somewhat vague terms—like bounded rationality—is that I did not want to give the impression that I thought I had “solved” the problem of creating an empirically grounded theory ... There still lies before us an enormous job of studying the actual decision processes... .

This book presents some of the progress made in answering Simon’s question. It conceives of the mind as a modular system, composed of heuristics, their building blocks, and evolved capacities. We thank all contributors and critics for carving out the details of this vision of adaptive rationality, and thus contributing, in the words of Herbert Simon (1999), “to this revolution in cognitive science, striking a great blow for sanity in the approach to human rationality.” Yet the progress is far from complete, with many surprises generating new questions. In the search for answers to these, we look forward to the active help of our readers.

ACKNOWLEDGMENTS

We would like to thank Rona Unrau and Laura Wiles for helping at all stages of the development of this book; they provided wonderful and constructive support. We are grateful to Jürgen Rossbach and Marianne Hauser, who redesigned the figures and tables, and to Emina Canic, Florence Ettlin, Christel Fraser, Daniela Fromm, Eva Günther, Renate Hoffmann, Carmen Kaiser, Erna Schiwietz, Lea Schweizer, Sarah Turowski,

and Françoise Weber for their help in compiling the reference list. We also thank the authors for taking the time to go through the shortened versions of their articles. Last not least, we express our gratitude to the Max Planck Society and the University of Basel, who provided the unique resources and the splendid intellectual atmosphere in which the science of heuristics could grow and mature.